
Subject: Re: Network namespaces a path to mergable code.

Posted by [Sam Vilain](#) on Wed, 28 Jun 2006 06:19:19 GMT

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

Eric W. Biederman wrote:

> In general it is possible to get file descriptors opened by someone
> else because unix domain sockets allow file descriptor passing. Similarly
> I think there are cases in both unshare and fork that allows you to sockets
> open before you entered a namespace.
>

This is an interesting point; it is known to be possible to do this on a traditional system, because with a Unix Domain socket, the other end is always in the same Unix Domain.

However what we're doing is saying that, well, the other end of the socket might not be in the same Unix Domain. In fact, we've already smashed to pieces this monolithic concept of a Unix Domain, to the point where the other end might be in a different network domain, but is in the same filesystem domain, for instance. Does it get to pass file descriptors through?

We would appear to be stretching the definition of "Unix Domain" somewhat if we allow these sockets to exist between network namespaces. Maybe it doesn't matter; this is just a VFS namespace feature/caveat.

Sam.
