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

Posted by [ebiederm](#) on Wed, 28 Jun 2006 04:33:48 GMT

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

Sam Vilain <sam@vilain.net> writes:

> It sounds then like it would be a good start to have general socket
> namespaces, if it would merge more easily - perhaps then network device
> namespaces would fall into place more easily.

I guess I really see both sockets and devices as the fundamental entities of a network namespace. Sockets need to be tagged because in the general case there is no guarantee that a socket that you are using was created in the network namespace of your current process.

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.

Since you can't create a new socket in a different network namespace I can't see any real problems with allowing them to be used, but they are something to be careful about in container creation code.

Something to examine here is that if both network devices and sockets are tagged does that still allow implicit network namespace passing.

Eric
