
Subject: [RFC] Network namespaces a path to mergable code.

Posted by [ebiederm](#) on Tue, 27 Jun 2006 17:20:40 GMT

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

Thinking about this I am going to suggest a slightly different direction for get a patchset we can merge.

First we concentrate on the fundamentals.

- How we mark a device as belonging to a specific network namespace.
- How we mark a socket as belonging to a specific network namespace.

As part of the fundamentals we add a patch to the generic socket code that by default will disable it for protocol families that do not indicate support for handling network namespaces, on a non-default network namespace.

I think that gives us a path that will allow us to convert the network stack one protocol family at a time instead of in one big lump.

Stubbing off the sysfs and sysctl interfaces in the first round for the non-default namespaces as you have done should be good enough.

The reason for the suggestion is that most of the work for the protocol stacks ipv4 ipv6 af_packet af_unix is largely noise, and simple replacement without real design work happening. Mostly it is just tweaking the code to remove global variables, and doing a couple lookups.

Eric
