
Subject: Re: Network namespaces a path to mergable code.
Posted by [Andrey Savochkin](#) on Tue, 27 Jun 2006 17:58:59 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric,

On Tue, Jun 27, 2006 at 11:20:40AM -0600, Eric W. Biederman wrote:

- >
- > 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.

I agree with the direction of your thoughts.

I was trying to do a similar thing, define clear steps in network namespace merging.

My first patchset covers devices but not sockets.

The only difference from what you're suggesting is ipv4 routing.

For me, it is not less important than devices and sockets. May be even more important, since routing exposes design deficiencies less obvious at socket level.

- >
- > 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.

Fine

Can you summarize you objections against my way of handling devices, please?
And what was the typo you referred to in your letter to Kirill Korotaev?

Regards
Andrey
