
Subject: Re: [PATCH 4/9] network namespaces: socket hashes
Posted by [Daniel Lezcano](#) on Thu, 21 Sep 2006 12:34:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Andrey Savochkin wrote:

- > The main reason is that socket hash tables should be large enough to work
- > efficiently, but it isn't good to waste a lot of memory for each namespace.
- > Namespaces should be cheap enough, to allow to have hundreds of them.
- > This reason of memory efficiency, of course, has a priority unless/until
- > socket hash tables start to resize automatically.
- >
- > Another point is that routing lookup is much more complicated than the
- > socket's one to add another search key.
- > Routing also have additional routines for deleting entries matching some
- > patterns, and so on.
- > In short, routing is much more complicated, and it already quite efficient
- > for various sizes of routing tables.

That makes sense, thx for the explanation.

Cheers.
-- Daniel.
