

---

Subject: Re: Network namespaces a path to mergable code.  
Posted by [Andrey Savochkin](#) on Wed, 28 Jun 2006 11:06:05 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi Eric,

On Tue, Jun 27, 2006 at 10:20:32PM -0600, Eric W. Biederman wrote:

> Andrey Savochkin <saw@swsoft.com> writes:

[snip]

> > 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.

>

> I agree we need to do it. I mostly want a base that allows us to

> not need to convert the whole network stack at once and still be able

> to merge code all the way to the stable kernel.

>

> The routing code is important for understanding design choices. It

> isn't important for merging if that makes sense.

Ok, fine.

Now I'm working on socket code.

We still have a question about implicit vs explicit function parameters.

This question becomes more important for sockets: if we want to allow to use sockets belonging to namespaces other than the current one, we need to do something about it.

One possible option to resolve this question is to show 2 relatively short patches just introducing namespaces for sockets in 2 ways: with explicit function parameters and using implicit current context.

Then people can compare them and vote.

Do you think it's worth the effort?

>

> For everyone looking at routing choices the IPv6 routing table is

> interesting because it does not use a hash table, and seems quite

> possibly to be an equally fast structure that scales better.

>

> There is something to think about there.

Sure

[snip]

> >

> > 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?  
>  
> I have no fundamental objects to the content I have seen so far.  
>  
> Please read the first email Kirill responded too. I quoted a couple  
> of sections of code and described the bugs I saw with the patch.

I found your comments, thank you!

>  
> All minor things. The typo I was referring to was a section where the  
> original iteration was on an ifp variable and you called it dev  
> without changing the rest of the code in that section.  
>  
> The only big issue was that the patch too big, and should be split  
> into a patchset for better review. One patch for the new functions,  
> and the an additional patch for each driver/subsystem hunk describing  
> why that chunk needed to be changed.

I'll split the patch.

> I'm still curious why many of those chunks can't use existing helper  
> functions, to be cleaned up.

What helper functions are you referring to?

Best regards

Andrey

---