
Subject: Re: Network virtualization/isolation
Posted by [ebiederm](#) on Mon, 04 Dec 2006 16:52:10 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dmitry Mishin <dim@openvz.org> writes:

> On Monday 04 December 2006 18:35, Eric W. Biederman wrote:
> [skip]
>> Where and when you look to find the network namespace that applies to
>> a packet is the primary difference between the OpenVZ L2
>> implementation and my L2 implementation.
>>
>> If there is a better and less intrusive while still being obvious
>> method I am all for it. I do not like the OpenVZ thing of doing the
>> lookup once and then stashing the value in current and the special
>> casing the exceptions.
> Why?

I like it when things are obvious and not implied.

The implementations seems to favor fewer lines of code touched over maintainability of the code. Which if you are maintaining out of tree code is fine. At least that was my impression last time I looked at the code.

I know there are a lot of silly things in the existing implementations because they were initially written without the expectation of being able to merge the code into the main kernel. This resulted in some non-general interfaces, and a preference for patches that touch as few lines of code as possible.

Anyway this has been discussed before and we can discuss it seriously in the context of patch review.

Eric
