
Subject: Re: [RFC] network namespaces
Posted by [ebiederm](#) on Sun, 10 Sep 2006 11:48:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

Dmitry Mishin <dim@openvz.org> writes:

> On Sunday 10 September 2006 07:41, Eric W. Biederman wrote:
>> I certainly agree that we are not at a point where a final decision
>> can be made. A major piece of that is that a layer 2 approach has
>> not shown to be without a performance penalty.
> But it is required. Why to limit possible usages?

Wrong perspective.

The point is that we need to dig in and show that there is no measurable penalty for the current cases. Showing that there is little penalty for the advanced configurations is a plus.

The practical question is, do we need to implement the grand unified lookup before we can do this cheaply, or can we implement this without needing that optimization?

To get a perspective, to get a good implementation of the pid namespace I am having to refactor significant parts of the kernel so it uses abstractions that can cleanly express what we are doing. The networking stack is in better shape but there is a lot of it.

>> A practical question. Do the IPs assigned to guests ever get used
>> by anything besides the guest?
> In case of level2 virtualization - no.

Actually that is one of the benefits of a layer 2 implementation you can set up weird things like shared IPs, that various types of fail over scenarios want.

My question was really about the layer 3 bind filtering techniques, and how people are using them.

The basic attraction with layer 3 is that you can do a simple implementation, and it will run very fast, and it doesn't need to conflict with the layer 2 work at all. If you can make that layer 3 implementation clean and generally mergeable as well it is worth pursuing.

Eric

Containers mailing list
Containers@lists.osdl.org

