Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view Posted by Daniel Lezcano on Mon, 26 Jun 2006 14:56:32 GMT

View Forum Message <> Reply to Message

Andrey Savochkin wrote:

> Hi Daniel,

Hi Andrey,

>

- > It's good that you kicked off network namespace discussion.
- > Although I wish you'd Cc'ed someone at OpenVZ so I could notice it earlier:).

devel@openvz.org?

- > When a device presents an skb to the protocol layer, it needs to know to which
- > namespace this skb belongs.
- > Otherwise you would never get rid of problems with bind: what to do if device
- > eth1 is visible in namespace1, namespace2, and root namespace, and each
- > namespace has a socket bound to 0.0.0.0:80?

Exact. But, the idea was to retrieve the namespace from the routes.

IMHO, I think there are roughly 2 network isolation implementation:

- make all network ressources private to the namespace
- keep a "flat" model where network ressources have a new identifier which is the network namespace pointer. The idea is to move only some network informations private to the namespace (eg port range, stats, ...)

Daniel.