Subject: Re: [patch 2/6] [Network namespace] Network device sharing by view Posted by ebjederm on Wed, 28 Jun 2006 14:21:38 GMT

View Forum Message <> Reply to Message

Herbert Poetzl herbert@13thfloor.at writes:

- > last time I pointed to such 'misguided' apps which
- > made assumptions that are not necessarily true
- > inside a virtual environment (e.g. pstree, initpid)
- > the general? position was that those apps should
- > be fixed instead adding a 'workaround'

I agree that if it was solely misguided apps. There would be no justification.

One of the standard applications interfaces we support is renaming a network interface. So supporting those misguided apps is a actually a side effect of supporting one the standard operations on a network interface.

Another way of looking at it is that the names of networking devices like the names of devices in /dev are a user space policy (today). In the configuration of the networking stack historically we had those identifiers hard coded. It isn't until just recently that user space has been able to cope with dynamically added/removed network devices.

As for initpid and friends. In the context where you are simply isolating pids and not doing a full pid namespaces it was felt that changing the few user space applications that care was easier and probably worth doing anyway.

- > note: personally I'm absolutely not against virtualizing
- > the device names so that each guest can have a separate
- > name space for devices, but there should be a way to
- > 'see' _and_ 'identify' the interfaces from outside
- > (i.e. host or spectator context)

Yep. Basically that interface comes when we fix the sysfs support, to support per namespace reporting.

Eric