
Subject: Re: [patch 3/4] Network namespaces: IPv4 FIB/routing in namespaces
Posted by [Ben Greear](#) on Wed, 28 Jun 2006 17:10:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daniel Lezcano wrote:

> Kirill Korotaev wrote:

>

>>>>> Structures related to IPv4 routing (FIB and routing cache)

>>>>> are made per-namespace.

>>>>

>>>>

>>>>

>>>> Hi Andrey,

>>>>

>>>> if the resources are private to the namespace, how do you will

>>>> handle NFS mounted before creating the network namespace ? Do you

>>>> take care of that or simply assume you can't access NFS anymore ?

>>>

>>>

>>>

>>>

>>> This is a question that brings up another level of interaction between

>>> networking and the rest of kernel code.

>>> Solution that I use now makes the NFS communication part always run in

>>> the root namespace. This is discussable, of course, but it's a far more

>>> complicated matter than just device lists or routing :)

>>

>>

>> if we had containers (not namespaces) then it would be also possible

>> to run NFS in context of the appropriate container and thus each user

>> could mount NFS itself with correct networking context.

With a relatively small patch, I was able to make NFS bind to a particular local IP (poor man's namespace with existing code). I also changed it so that multiple mounts to the same destination (and with unique local mount points) are treated as unique mounts. This patch was done so that I could stress test NFS servers, but similar logic might work for namespace isolation as well...

Ben

--

Ben Greear <greearb@candelatech.com>

Candela Technologies Inc <http://www.candelatech.com>
