
Subject: Re: [patch 3/4] Network namespaces: IPv4 FIB/routing in namespaces
Posted by [Daniel Lezcano](#) on Wed, 28 Jun 2006 16:56:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

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.

I was asking the question because in some case, we want a lightweight container for running applications (aka application container) who need to share the filesystem and it will be too bad to have a network namespace which brings isolation and prevents to implement application containers. By the way, I agree from a point of view of a system container, a complete network isolation is perfect.

Regards.

Daniel.
