Subject: Re: [patch 3/4] Network namespaces: IPv4 FIB/routing in namespaces
Posted by dev on Wed, 28 Jun 2006 14:34:33 GMT

>>>>Structures related to IPv4 rounting (FIB and routing cache)
>>>>are made per-namespace.
>>
>>Hi Andrey,
>>
>>if the ressources 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.

it's another thing which ties subsytems and makes namespaces ugly :/

Kirill