
Subject: Re: [patch 3/4] Network namespaces: IPv4 FIB/routing in namespaces
Posted by [ebiederm](#) on Wed, 28 Jun 2006 15:05:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daniel Lezcano <dlezcano@fr.ibm.com> writes:

> Daniel Lezcano wrote:
>> Andrey Savochkin wrote:
>>
>>> Structures related to IPv4 routing (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 ?

Through the filesystem namespace. It is a weird case but it works :)

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

It is actually a noop. Unless I goofed this is basically handled by looking at which socket NFS is using to communicate, and plucking the namespace from there.

As I recall NFS gets the socket at mount time when it still has user space context available.

So regardless if I implemented it correctly you can implement it that way and always get the namespace context from whoever implemented it.

Eric
