Subject: Re: [patch 2/2][NETNS][RFD] use dst_entries to retrieve the network namespace pointer
Posted by Daniel Lezcano on Wed, 12 Dec 2007 09:03:17 GMT
View Forum Message <> Reply to Message

Denis V. Lunev wrote:
> Daniel Lezcano wrote:
>> The objective we have is to remove the fl_net field from the struct flowi.
>>
>> In the previous patch, we make the network namespace to go up to the dst_entry.
>> This patch makes an example in how we can use the dst_entry/rtable combined with
>> a container_of to retrieve the network namespace when we have the flowi parameter.
>>
>> One restriction is we should pass always a flowi parameter coming from a struct
>> dst_entry.
> this is an obfuscation IMHO. You use rtable instead of flowi, so
> - rtable obtain flowi meaning and keeps original meaning
> - stack usage in increased

Yes, you are right. Using container_of is not a good idea finally.

> So, for the case, I think it is better to have an additional parameter
> on IP route output path. (There is a device on the input one).

I thought that at the beginning, but I was not inclined to change
exported functions API, so I looked a way to avoid that but I failed
miserably :)
Obviously, the ip route output path extra parameter is the cleanest way
to introduce the network namespace pointer in the routing code. So why not ?

_____
Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers