Subject: Re: Pid namespaces problems
Posted by Daniel Lezcano on Thu, 08 Nov 2007 13:29:03 GMT

Denis V. Lunev wrote:
> Daniel Lezcano wrote:
>
>> * the first one is the locking of the network namespace list by
>> rtnl_lock, so from the timer callback we can not browse the network
>> namespace list to check the age of the routes. It is a problem I would
>> like to talk with Denis if he has time
>
> From my point of view, the situation is clear. The timer should be
> per/namespace. The situation is completely different as one in IPv4.

We thought to make a timer per namespace for ipv6, but we are a little
afraid for the performances when there will be a lot of containers.
Anyway, we can do a timer per namespace and optimize that later. I will
cook a new patch to take into account that for the next week.

>> * the loopback refcounting is not correctly handled in ipv6. This
>> protocol do not expect to have the loopback to be unregistered, so there
>> is some problem with the addr_ifdown function when exiting the network
>> namespace
>
> I think that default routing targets and similar staff should be
> dynamically allocated as a start and submitted now. No need to wait
> NETNS infrastructure. I have sent similar cleanup for fib rules
> recently. No answer from David yet.

Interesting.

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