

Pavel Emelyanov wrote:

> Daniel Lezcano wrote:

>> Denis V. Lunev wrote:

>> > Daniel Lezcano wrote:

>> >> 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.

>> >

>> > IMHO not a problem. tcp_write_timer is per/socket timer. If this works
>> > efficiently, per/namespace one will work also.

>>

>> That's right, this is a good argument. By the way, the amount of work to
>> be done in the tcp_write_timer is perhaps smaller than the one done in
>> the ipv6 routing age check, no ? Anyway, I'm not against a timer per
>> namespace in this case, I already did a try before rolling back to a
>> for_each_net in the gc timer, that changes a little the API, but nothing
>

> We can easily make the netns list rcu protected to address this issue.
> If you're interested, I can prepare a patch tomorrow.

That would be great if you manage to do it.

This was our initial idea with Daniel, but as I have a limited
knowledge of RCU, I didn't manage to obtain an acceptable patch.
One of the more problematic area is rtnl_unlock().

Benjamin

>

>> we can handle easily.

>>

>>

>

>

--

B e n j a m i n T h e r y - BULL/DT/Open Software R&D

<http://www.bull.com>

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
