
Subject: Re: [patch 1/1][NETNS][IPV6] protect addrconf from loopback registration
Posted by [ebiederm](#) on Tue, 13 Nov 2007 12:59:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

David Miller <davem@davemloft.net> writes:

Well this is a weird way to get to this part of the conversation.

> From: "Denis V. Lunev" <den@sw.ru>
> Date: Mon, 12 Nov 2007 19:49:03 +0300
>
>> Unregister for a loopback in !init_net is a _valid_ operation and should
>> be clean, i.e. without kludges in the path. This is the only way to
>> check the ref-counting.
>
> For ipv6 the stack really wants to pin down the loopback
> device because we need a valid inet6_dev object to reference
> at all times in order to simplify the per-device SNMP
> statistic bumping.
>
> When a non-loopback device goes down, we point any existing
> references to that device's idv to the loopback one instead.
>
> I really consider taking down the loopback device to be
> an invalid operation at least how things are implemented
> currently.

In a secondary network namespace the current implementation registers the loopback device before all other network devices in a network namespace and it unregisters the loopback device after all other network devices.

So we don't endanger the current scheme of pointing any existing reference to the loopback device. In fact the current ipv6 addrconf_cleanup largely does the same thing.

Unregistering the loopback device is definitely a case where we need to tread very carefully.

Bug we absolutely need to do all of our cleanup for a network namespace went it goes away and that includes removing the per network namespace copy of the loopback device.

Eric

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