
Subject: Re: [patch 1/1][NETNS][IPV6] protect addrconf from loopback registration
Posted by [dlunev](#) on Mon, 12 Nov 2007 20:59:21 GMT

[View Forum Message](#) <> [Reply to Message](#)

> Let me try to clarify:

>

> * when the init network namespace is created, the loopback is created
> first, before ipv6, and the notifier call chain for ipv6 is not setup,
> so the protocol does not receive the REGISTER event

>

> * when the init network namespace is destroyed during shutdown, the
> loopback is not unregistered, so there is no UNREGISTER event

>

> * when we create a new network namespace, a new instance of the
> loopback is created and a NETDEV_REGISTER is sent to ipv6 because the
> notifier call chain has been setup by the init netns (while ipv6
> protocol is not yet configured for the namespace which is being created)

>

> * when the network namespace exits, the loopback is unregistered after
> the ipv6 protocol but the NETDEV_UNREGISTER is sent to addrconf_notify
> while the ipv6 protocol has been destroyed.

this should not be a problem :). IPv6 exiting code could remove in_dev6 at the protocol layer. In this case the notifier will be noop. This approach is completely equivalent to the unloading of IPv6 module with persistent loopback.

The registration is still the problem. May be we need to separate registration from initialization and perform it after protocol layer for all loopback devices? I'd like this right now, but this is tight change...

>

>

> The objective of the patch is to discard these events because they were
> never taken into account and they are not expected to be receive by ipv6
> protocol.

>

Containers mailing list

Containers@lists.linux-foundation.org

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