Subject: Re: [patch 1/1][NETNS][IPV6] protect addrconf from loopback registration Posted by ebjederm on Tue, 13 Nov 2007 13:11:37 GMT

View Forum Message <> Reply to Message

- > In article <m16407450i.fsf@ebiederm.dsl.xmission.com> (at Mon, 12 Nov 2007
- > 12:50:53 -0700), ebiederm@xmission.com (Eric W. Biederman) says:

>

- >> My opinion is that both your analysis is slightly off (as to the cause
- >> of your problems) and that your approach to fix your problem is wrong
- >> because you don't untangle the knot you keep it.

>:

- >> I have register_pernet_subsys and register_per_net_device to ensure
- >> that when we create a new network namespace all of the subsystems are
- >> initialized before the network devices are initialize. So ipv6 should
- >> be ready before we initialize the new loopback device comes into
- >> existence.

>

- > User may not load ipv6.ko at boot, and then do "modprobe ipv6".
- > Do you take this into account?

Absolutely.

In the general case the infrastructure has to work for netfilter, ipv6, and other parts of the networking stack that can be made modular.

The only limitation is that if you update struct net to add a new field to help a modular ipv6 the core kernel needs to be recompiled.

When you load ipv6.ko late in the game first we call the init methods which will eventually register the per network namespace registration methods. Then register_netdevice_notifier is called. At which point ipv6 is ready for the registration method.

For additional network namespace (which is the case that was claimed was in trouble) the pernet_susbsys logic initializes all of the subsystems before it initializes any of the network devices. Effectively persevering the initialization order that exists today with just the init methods and register_netdevice_notifier.

_		
⊢	rı	C

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