
Subject: Re: [PATCH 1/3] [IPV6]: Event type in addrconf_ifdown is mis-used.
Posted by [den](#) on Sun, 23 Mar 2008 14:34:59 GMT

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

On Sun, 2008-03-23 at 03:17 -0700, David Miller wrote:

```
> From: "Denis V. Lunev" <den@openvz.org>
> Date: Sun, 23 Mar 2008 11:13:16 +0300
>
> > First, this behaviour is broken for a namespace right now in the 2.6.26
> > tree. inet6_dev pointer will be NULL for a loopback inside the
> > namespace. The case is simple. Just remove all INET6 addresses from a
> > loopback device inside a VE. This will call
> >   inet6_addr_del
> >   addrconf_ifdown(dev, 1);
> >   if (dev == init_net.loopback_dev && how == 1)
> >       how = 0;
> > the condition will be false and how will not be changed here.
>
> That's a bug.
>
> You can't mark any namespace's loopback device's inet6_dev as NULL
> until you know that all routes, devices, and packets referring to such
> devices and routes in that namespace are %100 gone and unreferenced.
>
> It is now obviously apparent that there are several severe errors
> here.
```

You are perfectly correct and the place in addrconf_cleanup is that place when we believe that we should destroy all the stuff.

You see, it is pretty useless to call addrconf_ifdown(dev, 2) after addrconf_dev(dev, 0) for a loopback in the current code! No new cleanups will be performed for 2, pls check :)

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