
Subject: Re: [PATCH 1/3] [IPV6]: Event type in addrconf_ifdown is mis-used.
Posted by [davem](#) on Sun, 23 Mar 2008 00:38:43 GMT

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

From: "Denis V. Lunev" <den@openvz.org>

Date: Tue, 18 Mar 2008 17:35:23 +0300

> addrconf_ifdown is broken in respect to the usage of how parameter. This
> function is called with (event != NETDEV_DOWN) and (2) on the IPv6 stop.
> It the latter case inet6_dev from loopback device should be destroyed.
>
> Signed-off-by: Denis V. Lunev <den@openvz.org>

The code purposefully treats "2" specially because when IPV6 routes are destroyed they are changed to point to the loopback device's inet6_dev object.

This allows statistic bumping code to not have to check if it has a NULL inet6_dev pointer or not, because that's now impossible.

Since ipv6 is not unloadable, addrconf_cleanup(), and thus the "how == 2" case can only occur when ipv6 fails to load properly. The only real consequence of this bug is that if ipv6 fails to load properly, a subsequent successful load of ipv6 will leak the loopback device's inet6_dev object, which isn't that much of a big deal.

I understand that for namespaces you have to deal with multiple loopback devices, but you'll need to solve that problem while still handling the wish of the ipv6 stack for inet6_dev objects of loopback devices to be permanent and guaranteed to always be around for the sake of statistics bumping.

I thus can't apply any of these patches until those issues are resolved.

Containers mailing list

Containers@lists.linux-foundation.org

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