Subject: [PATCH] vlan: fix potential race in vlan_cleanup_module vs vlan_ioctl_handler
Posted by Pavel Emelianov on Tue, 11 Dec 2007 10:25:11 GMT
View Forum Message <> Reply to Message

The vlan module cleanup function starts with

```
 vlan_netlink_fini();
 vlan_ioctl_set(NULL);
```

The first call removes all the vlan devices and
the second one closes the vlan ioctl.

AFAIS there's a tiny race window between these two
calls - after rtnl unregistered all the vlans, but
the ioctl handler isn't set to NULL yet, user can
manage to call this ioctl and create one vlan device,
and that this function will later BUG_ON seeing
non-emply hashes.

I think, that we must first close the vlan ioctl
and only after this remove all the vlans with the
vlan_netlink_fini() call.

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

```
diff --git a/net/8021q/vlan.c b/net/8021q/vlan.c
index 5b18315..4add9bd 100644
--- a/net/8021q/vlan.c
+++ b/net/8021q/vlan.c
@@ -124,8 +124,8 @@ static void __exit vlan_cleanup_module(void)
 {
  int i;

- vlan_netlink_fini();
  vlan_ioctl_set(NULL);
+ vlan_netlink_fini();

  /* Un-register us from receiving netdevice events */
  unregister_netdevice_notifier(&vlan_notifier_block);
```

Subject: Re: [PATCH] vlan: fix potential race in vlan_cleanup_module vs vlan_ioctl_handler
Posted by Patrick McHardy on Tue, 11 Dec 2007 10:38:38 GMT

Pavel Emelyanov wrote:
> The vlan module cleanup function starts with
>
> vlan_netlink_fini();
> vlan_ioctl_set(NULL);
>
> The first call removes all the vlan devices and
> the second one closes the vlan ioctl.
>
> AFAIS there's a tiny race window between these two
> calls - after rtnl unregistered all the vlans, but
> the ioctl handler isn't set to NULL yet, user can
> manage to call this ioctl and create one vlan device,
> and that this function will later BUG_ON seeing
> non-emply hashes.

Indeed, I can't see anything preventing this.

> I think, that we must first close the vlan ioctl
> and only after this remove all the vlans with the
> vlan_netlink_fini() call.

That looks correct, thanks Pavel. Dave, please apply.

Subject: Re: [PATCH] vlan: fix potential race in vlan_cleanup_module vs vlan_ioctl_handler
Posted by davem on Tue, 11 Dec 2007 10:41:38 GMT

From: Patrick McHardy <kaber@trash.net>
Date: Tue, 11 Dec 2007 11:38:38 +0100

> Pavel Emelyanov wrote:
> > AFAIS there's a tiny race window between these two
> > calls - after rtnl unregistered all the vlans, but
> > the ioctl handler isn't set to NULL yet, user can
> > manage to call this ioctl and create one vlan device,
> > and that this function will later BUG_ON seeing
> > non-emply hashes.
>
> Indeed, I can't see anything preventing this.
>
> > I think, that we must first close the vlan ioctl
> > and only after this remove all the vlans with the
> > vlan_netlink_fini() call.

>
> That looks correct, thanks Pavel. Dave, please apply.

Applied to net-2.6, thanks!
