
Subject: [PATCH] Remove notifier block from chain when register_netdevice_notifier fails

Posted by [Pavel Emelianov](#) on Wed, 14 Nov 2007 12:08:03 GMT

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

Commit fcc5a03ac42564e9e255c1134dda47442289e466 makes the register_netdevice_notifier() handle the error from the NETDEV_REGISTER event, sent to the registering block.

The bad news is that in this case the notifier block is not removed from the list, but the error is returned to the caller. In case the caller is in module init function and handles this error this can abort the module loading. The notifier block will be then removed from the kernel, but will be left in the list. Oops :(

I think that the notifier block should be removed from the chain in case of error, regardless whether this error is handled by the caller or not. In the worst case (the error is `_not_` handled) module will not receive the events any longer.

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

```
diff --git a/net/core/dev.c b/net/core/dev.c
index dd40b35..86d6261 100644
--- a/net/core/dev.c
+++ b/net/core/dev.c
@@ -1171,6 +1171,8 @@ rollback:
     nb->notifier_call(nb, NETDEV_UNREGISTER, dev);
 }
 }
+
+ raw_notifier_chain_unregister(&netdev_chain, nb);
+ goto unlock;
 }
```

Subject: Re: [PATCH] Remove notifier block from chain when register_netdevice_notifier fails

Posted by [Herbert Xu](#) on Wed, 14 Nov 2007 13:37:17 GMT

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

On Wed, Nov 14, 2007 at 03:08:03PM +0300, Pavel Emelyanov wrote:
> Commit fcc5a03ac42564e9e255c1134dda47442289e466 makes the
> register_netdevice_notifier() handle the error from the

> NETDEV_REGISTER event, sent to the registering block.
>
> The bad news is that in this case the notifier block is
> not removed from the list, but the error is returned to the
> caller. In case the caller is in module init function and
> handles this error this can abort the module loading. The
> notifier block will be then removed from the kernel, but
> will be left in the list. Oops :(
>
> I think that the notifier block should be removed from the
> chain in case of error, regardless whether this error is
> handled by the caller or not. In the worst case (the error
> is `_not_handled`) module will not receive the events any
> longer.
>
> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Good catch. Thanks Pavel!

Acked-by: Herbert Xu <herbert@gondor.apana.org.au>

--

Visit Openswan at <http://www.openswan.org/>

Email: Herbert Xu ~{PmV>Hl~} <herbert@gondor.apana.org.au>

Home Page: <http://gondor.apana.org.au/~herbert/>

PGP Key: <http://gondor.apana.org.au/~herbert/pubkey.txt>

Subject: Re: [PATCH] Remove notifier block from chain when
register_netdevice_notifier fails
Posted by [Pavel Emelianov](#) on Wed, 14 Nov 2007 13:47:24 GMT
[View Forum Message](#) <> [Reply to Message](#)

Herbert Xu wrote:

> On Wed, Nov 14, 2007 at 03:08:03PM +0300, Pavel Emelyanov wrote:
>> Commit fcc5a03ac42564e9e255c1134dda47442289e466 makes the
>> register_netdevice_notifier() handle the error from the
>> NETDEV_REGISTER event, sent to the registering block.
>>
>> The bad news is that in this case the notifier block is
>> not removed from the list, but the error is returned to the
>> caller. In case the caller is in module init function and
>> handles this error this can abort the module loading. The
>> notifier block will be then removed from the kernel, but
>> will be left in the list. Oops :(
>>
>> I think that the notifier block should be removed from the
>> chain in case of error, regardless whether this error is
>> handled by the caller or not. In the worst case (the error

>> is `_not_ handled`) module will not receive the events any
>> longer.
>>
>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
>
> Good catch. Thanks Pavel!
>
> Acked-by: Herbert Xu <herbert@gondor.apana.org.au>

Thanks :)

BTW, I have one more question I can't find the answer to by myself. Why the `unregister_netdevice_notifier` doesn't produce the `GOING_DOWN-DOWN-UNREGISTER` event sequence for the notifier block which is being unregistered? Was this made deliberately?

Pavel

Subject: Re: [PATCH] Remove notifier block from chain when `register_netdevice_notifier` fails
Posted by [Herbert Xu](#) on Wed, 14 Nov 2007 13:50:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, Nov 14, 2007 at 04:47:24PM +0300, Pavel Emelyanov wrote:

>
> BTW, I have one more question I can't find the answer to by
> myself. Why the `unregister_netdevice_notifier` doesn't produce
> the `GOING_DOWN-DOWN-UNREGISTER` event sequence for the notifier
> block which is being unregistered? Was this made deliberately?

I think it's more of an oversight :)

Although it has no real impact because the only user that matters here is IPv6 and it can't be unloaded.

Cheers,

--

Visit Openswan at <http://www.openswan.org/>
Email: Herbert Xu ~{PmV>Hl~} <herbert@gondor.apana.org.au>
Home Page: <http://gondor.apana.org.au/~herbert/>
PGP Key: <http://gondor.apana.org.au/~herbert/pubkey.txt>

Subject: Re: [PATCH] Remove notifier block from chain when `register_netdevice_notifier` fails
Posted by [davem](#) on Wed, 14 Nov 2007 23:53:43 GMT

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

From: Herbert Xu <herbert@gondor.apana.org.au>

Date: Wed, 14 Nov 2007 21:37:17 +0800

> Acked-by: Herbert Xu <herbert@gondor.apana.org.au>

Indeed, good catch, patch applied.

Thanks!
