batman-adv - Bug #307

Can't re-add interface via rtnetlink in 2016.4

11/02/2016 05:14 PM - Julian Labus

Status:	Closed	**Start date:**	11/02/2016
Priority:	Normal	**Due date:**	
Assignee:	Linus Lüssing	**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:	2016.5		

Description

Since version 2016.4 it is no longer possible to re-add an interface via rtnetlink if it was removed from an batadv interface.

It is really easy to reproduce:
ip link add name bat0 type batadv
ip link set dev eth1 master bat0
ip link set dev eth1 nomaster
ip link set dev eth1 master bat0

The second command to set master of eth1 to bat0 again fails with the message *"RTNETLINK answers: Invalid argument"*

I couldn't reproduce this with 2016.2 and 2016.3. The tested kernels where 4.7.0 (amd64) and 3.10.14 (mipsel)

History

#1 - 11/02/2016 05:54 PM - Sven Eckelmann

- Assignee set to Linus Lüssing

Thanks for the report. It looks like following commit causes this problem. Linus please have a look

```
bac7733d06fac28ce68a79bcdf88b2b265600cf2 is the first bad commit
commit bac7733d06fac28ce68a79bcdf88b2b265600cf2
Author: Linus Lüssing <linus.luessing@c0d3.blue>
Date:   Thu Oct 6 01:43:08 2016 +0200

    batman-adv: fix splat on disabling an interface

    As long as there is still a reference for a hard interface held, there might
    still be a forwarding packet relying on its attributes.

    Therefore avoid setting hard_iface->soft_iface to NULL when disabling a hard
    interface.

    This fixes the following, potential splat:

        batman_adv: bat0: Interface deactivated: eth1
        batman_adv: bat0: Removing interface: eth1
        cgroup: new mount options do not match the existing superblock, will be ignored
        batman_adv: bat0: Interface deactivated: eth3
        batman_adv: bat0: Removing interface: eth3
        ------------[ cut here ]------------
        WARNING: CPU: 3 PID: 1986 at ./net/batman-adv/bat_iv_ogm.c:549 batadv_iv_send_outstanding_bat_ogm_pack
et+0x145/0x643 [batman_adv]
        Modules linked in: batman_adv(O-) <...>
        CPU: 3 PID: 1986 Comm: kworker/u8:2 Tainted: G        W  O    4.6.0-rc6+ #1
        Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS 1.7.5-20140531_083030-gandalf 04/01/2014
        Workqueue: bat_events batadv_iv_send_outstanding_bat_ogm_packet [batman_adv]
         0000000000000000 ffff88001d93bca0 ffffffff8126c26b 0000000000000000
         0000000000000000 ffff88001d93bcf0 ffffffff81051615 ffff88001f19f818
         000002251d93bd68 0000000000000046 ffff88001dc04a00 ffff88001becbe48
        Call Trace:
         [<ffffffff8126c26b>] dump_stack+0x67/0x90
         [<ffffffff81051615>] __warn+0xc7/0xe5
         [<ffffffff8105164b>] warn_slowpath_null+0x18/0x1a
         [<ffffffffa0356f24>] batadv_iv_send_outstanding_bat_ogm_packet+0x145/0x643 [batman_adv]
         [<ffffffff8108b01f>] ? __lock_is_held+0x32/0x54
```

```
         [<ffffffff810689a2>] process_one_work+0x2a8/0x4f5
         [<ffffffff81068856>] ? process_one_work+0x15c/0x4f5
         [<ffffffff81068df2>] worker_thread+0x1d5/0x2c0
         [<ffffffff81068c1d>] ? process_scheduled_works+0x2e/0x2e
         [<ffffffff81068c1d>] ? process_scheduled_works+0x2e/0x2e
         [<ffffffff8106dd90>] kthread+0xc0/0xc8
         [<ffffffff8144de82>] ret_from_fork+0x22/0x40
         [<ffffffff8106dcd0>] ? __init_kthread_worker+0x55/0x55
        ---[ end trace 647f9f325123dc05 ]---
```

```
    What happened here is, that there was still a forw_packet (here: a BATMAN IV
    OGM) in the queue of eth3 with the forw_packet->if_incoming set to eth1 and the
    forw_packet->if_outgoing set to eth3.

    When eth3 is to be deactivated and removed, then this thread waits for the
    forw_packet queued on eth3 to finish. Because eth1 was deactivated and removed
    earlier and by that had forw_packet->if_incoming->soft_iface, set to NULL, the
    splat when trying to send/flush the OGM on eth3 occures.

    Signed-off-by: Linus Lüssing <linus.luessing@c0d3.blue>
    [sven@narfation.org: Reduced size of Oops message]
    Signed-off-by: Sven Eckelmann <sven@narfation.org>

:040000 040000 e5c20821beef13ab9a8f1fe8f6a4d621b78714cc ccb1e6f710d179b612ac7b47c39b9dd9662e0eb3 M      net
```

#2 - 11/02/2016 06:16 PM - Sven Eckelmann

The simple revert for the commit can be found at
https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/20161102171443.9491-1-sven@narfation.org/

#3 - 11/02/2016 06:35 PM - Julian Labus

Thanks for the quick reply. I just tested 2016.4 without commit bac7733d06fac28ce68a79bcdf88b2b265600cf2 and now I can remove interfaces from bat0 and re-add them.

#4 - 11/02/2016 08:48 PM - Jean-Jacques Sarton

On Kernel 4.8.4-200.fc24.x86_64 (Fedora) there are no problems with the ip commands stated by the reporter.
I don't have applied them mentioned patch, my batman_adv version is the original 2016.4
May be that the kernel version is important.

#5 - 11/02/2016 09:09 PM - Sven Eckelmann

@Jean: If you don't have this problem with the commands given by the poster then you don't have 2016.4 installed. Everything required to trigger it is

part of batman-adv, no compat code is involved and it is not a race condition.

The problem is described in the revert patch:

1. soft_iface is set for an hard_iface in batadv_hardif_enable_interface
2. soft_iface is then not set back to NULL in batadv_hardif_disable_interface
3. it is detected in batadv_softif_slave_add that "hard_iface->soft_iface" is not NULL and thus the ip link set dev eth1 master bat0 is rejected

Maybe your ip link isn't showing the same error message. But `ip link` will not show "master bat0" for the device you've used for the tests.

#6 - 11/03/2016 11:41 AM - Julian Labus

I get exactly the same results with Fedora 24 and Kernel 4.8.4. But I see why you may don't get the error. When I install batman-adv.ko via "make install" and do a "modprobe batman-adv" it always loads batman-adv 2016.3 which is included in the kernel. I had to load the 2016.4 module explicitly with "insmod ./batman-adv.ko".

I guess its because of the new install path that conflicts with depmod which excludes all paths with `build` in it.

https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/20161031072719.26286-1-sven@narfation.org/

#7 - 11/03/2016 11:33 PM - Linus Lüssing

@Jean-Jacques: You can verify the version of the loaded batman-adv kernel module via "$ dmesg" for instance.

Maybe Julian is right and you are accidentally using v2016.3?

#8 - 11/04/2016 01:46 AM - Linus Lüssing

@Sven: Simply removing the "|| hard_iface->soft_iface" check in batadv_softif_slave_add() seems wrong to me. Or at least I think there might still another issue somewhere deeper.

```
    What happened here is, that there was still a forw_packet (here: a BATMAN IV
    OGM) in the queue of eth3 with the forw_packet->if_incoming set to eth1 and the
    forw_packet->if_outgoing set to eth3.
```

I'm currently wondering, how this could actually happen. Or, it should not happen because the removal of eth1/eth3 in this case should have purged the forw_packet causing this splat.

When, for instance eth1 is turned off, any forw_packet which has either if_outgoing=<eth1> or if_incoming=<eth1> is removed from the queue. That check was present before and after the purging refactoring.

I had tested and seen this splat when testing with the "fix rare race condition on interface removal" patch. I'll retest.

#9 - 11/04/2016 04:38 AM - Linus Lüssing

I currently can't reproduce the splat with v2016.4.

Maybe all this fiddling with the hard-interfaces might have fixed it properly?

Feel free to revert my patch.

#10 - 11/04/2016 08:23 AM - Sven Eckelmann

- Status changed from New to Resolved

Thanks Linus. The patch was now applied.

#11 - 12/16/2016 09:40 AM - Sven Eckelmann

- Status changed from Resolved to Closed

Fix is part of the release 2016.5

#12 - 02/11/2017 08:32 AM - Sven Eckelmann

- Target version set to 2016.5