batman-adv - Bug #299

kernel crash when shutting down batman-adv in netns

10/24/2016 08:32 PM - Jean-Jacques Sarton

Status:	Closed	**Start date:**	10/24/2016
Priority:	High	**Due date:**	
Assignee:	Jean-Jacques Sarton	**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:	2016.4		

Description

With batman-adv 2016.3 the kernel crash if the pc ist shutted down and batman work within a network namespace.

The crash end in an endless loop.

History

#1 - 10/24/2016 09:48 PM - Sven Eckelmann

- Assignee set to Jean-Jacques Sarton

- Status changed from New to Feedback

- Subject changed from kernel crash to kernel crash when shutting down batman-adv in netns

First things first: 2016.3 only has the first (initial) patches for netns. The rest comes with 2016.4.

Just tried with Emulation_Debug (Linux 4.8) and following steps with current master:

```
insmod /host/batman-adv/build/net/batman-adv/batman-adv.ko

EMU1="ip netns exec emu1"
EMU2="ip netns exec emu2"

ip netns add emu1
ip netns add emu2

ip link add emu1-veth1 type veth peer name emu2-veth1
ip link set emu1-veth1 netns emu1
ip link set emu2-veth1 netns emu2

$EMU1 ip link set emu1-veth1 name veth1
$EMU2 ip link set emu2-veth1 name veth1

$EMU1 ip link set veth1 up
$EMU2 ip link set veth1 up

ip link add emu1-veth2 type veth peer name veth2
ip link set emu1-veth2 netns emu1
$EMU1 ip link set emu1-veth2 name veth2

$EMU1 ip link set veth2 up
ip link set veth2 up

$EMU1 batctl if add veth1
$EMU1 batctl if add veth2
$EMU1 ip link set bat0 up

$EMU2 batctl if add veth1
$EMU2 ip link set bat0 up

batctl if add veth2
ip link set bat0 up

sleep 10
batctl o

sleep 3
poweroff
```

This worked fine.

Then I did following with 2016.3 + 6b0485c758be31ff3d466644975eaedc54db6d17 (just to get netns moves of devices working with batman-adv) :

```
insmod /host/batman-adv/net/batman-adv/batman-adv.ko

EMU1="ip netns exec emu1"

ip netns add emu1

ip link add emu1-veth2 type veth peer name veth2
ip link set emu1-veth2 netns emu1
$EMU1 ip link set emu1-veth2 name veth2

$EMU1 ip link set veth2 up
ip link set veth2 up

$EMU1 batctl if add veth2
$EMU1 ip link set bat0 up

sleep 3
poweroff
```

This also worked fine

Can you please explain what to change in my little example to reproduce the problem? And what is the crash?

#2 - 10/25/2016 10:03 AM - Jean-Jacques Sarton

Distribution: Fedora 24
Kernel: 4.7.9-200.fc24.x86_64

The crash with my environment is allways reproducible.

The Layout for the network devices is as follow:

I use a veth pair for communication with the device put to the network name space (veth1 <-> eth0)

Within the network name space I have

```
[br-wan]            comunication with main system
   [eth0]

[br-client]
   [bat0]
      [mesh-vpn]      tap device from fastd
      [mesh0]         usb wifi virtual device as MP
   [wlan0]            same usb device as AP
   [eth1]             usb ethernet device
```

Debug fs is mounted for the name space.

Some applications are running withn the name space (fastd, alfred, alfred-vis, host_apd, thhtpd and a few others).

According to the informations one the monitot a page fault occur within sys_remove_file which is called from:

batadv_sysfs_remove_file_ns
batadv_softif_destroy vlan
batadv_softif_destroy_netlink

#3 - 10/25/2016 10:11 AM - Jean-Jacques Sarton

I don't use quemu, my implementation is very light.

#4 - 10/25/2016 10:28 AM - Sven Eckelmann

There is no such thing as batadv_sysfs_remove_file_ns. Maybe you are using a patched version of batman-adv?

Can you please test with the current master (because it contains the missing patches for the netns support)? And can you please post the complete log.

And can you please post the relevant changes from my example to reproduce the problem? Please don't post the complete guide to setup this freifunk node because it most likely contains a lot of redundant things not relevant for this problem.

PS: qemu is only used by me so I can test around on a virtual system without destroying my development environment. It is not really relevant for netns because the netns stuff is done inside the qemu and doesn't know that it is inside a virtual environment.

PPS: the optimizer of gcc can turn the kernel backtraces to garbage. So it can be useful to compile the kernel/modules with a less aggressive optimization step to make them readable:

```
make EXTRA_CFLAGS="-fno-inline -O1 -fno-optimize-sibling-calls" KERNELPATH=/home/sven/tmp/linux-next V=1
```

#5 - 10/25/2016 11:27 AM - Sven Eckelmann

Thought a little bit about what could cause this (assuming that you meant sysfs_remove_file_ns). And my best guess is that a bugfix for a leak is triggering a parallel removal of the bat0 device vlans. This conceptual problem was fixed using two commits in next:

- https://git.open-mesh.org/batman-adv.git/commit/a1f0a804cacda8ca0847cbac07adcbf95e9c24c1
- https://git.open-mesh.org/batman-adv.git/commit/0994697e926d9633adb4912036548b0bb1d110ea

So you can also first test the next branch (which is basically 2016.4 and doesn't contain the newest, not so well tested patches). But you should also know that 2016.4 is not providing debugfs in a non-init netns anymore. So you should also upgrade batctl to its current next branch (v2016.4) and alfred to its master branch (v2016.4).

#6 - 10/25/2016 12:43 PM - Jean-Jacques Sarton

I know that batman 2016.4 don't work with my actual code. I had downloaded it from kernel.org made a short test and reverted to 2016.3.
2016.4 seem to be promising bat I assume that we will have to wait a long time before ist is integrated within openwrt/led gluon.
How can I get the 2016 branch ? Git list only the master brancg.

#7 - 10/25/2016 01:19 PM - Sven Eckelmann

There is no 2016 branch. And v2016.4 has to be tagged (maybe Simon tag the next branch as v2016.4 in the coming days).

And the time it takes until 2016.4 will be available in LEDE/OpenWrt depends on the loudness of our screams towards Marek. Not sure why he didn't accept my pull request for 2016.3 yet. I hope he will accept 2016.4 faster in openwrt-routing

And I will be waiting for gluon until Matthias finished its LEDE port. But you can get a first glance at it by using https://github.com/FreifunkVogtland/gluon/tree/v2016.1.x-hwtest. Freifunk Vogtland is already using it on their nodes (supernodes are still running 2016.3)

Btw. why did you download anything from kernel.org (there is no final Linux release which includes batman-adv 2016.4)? Our stuff can be found under https://git.open-mesh.org/ and https://downloads.open-mesh.org/ (v2016.4 is not yet available as tarball).

Easiest way to download it:

```
git clone -b next https://git.open-mesh.org/batman-adv.git
make -C batman-adv
sudo make -C batman-adv install
reboot
....
batctl -v
```

The same can be done with the batctl repo (not sure where you install it on your system) and the master branch of alfred.

And what didn't work? Did it still crash?

#8 - 10/27/2016 09:48 AM - Jean-Jacques Sarton

batctl -v tell 2016.3-41-gb7c5cd4

The kernel.org version is 2016.3-63-g061486b
(strings ./net/batman-adv/batman-adv.ko | grep '2016\.')

For both I have no more a crash.

by the way I suspect that there is an other problem, see:
https://forum.freifunk.net/t/router-advertisement-an-absender-zurueck/13714

#9 - 10/27/2016 10:00 AM - Sven Eckelmann

- Status changed from Feedback to Resolved

Mark this as resolved. Ticket will be closed after release of batman-adv 2016.4

#10 - 10/28/2016 10:37 AM - Sven Eckelmann

- Status changed from Resolved to Closed

Simon released 2016.4 yesterday. openwrt-packages integration is pending (pull request was issued). Gluon integration will follow later (RFC will be prepared after openwrt-packages pull request was merged).

#11 - 02/11/2017 08:34 AM - Sven Eckelmann

- Target version set to 2016.4