

batman-adv - Bug #327

Possible Fragmentation Issue

03/03/2017 03:01 AM - Martin Weinelt

Status:	Closed	Start date:	03/03/2017
Priority:	Normal	Due date:	
Assignee:	Simon Wunderlich	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	2017.0.1		
Description			
HTTP download of a small file from the webserver stalls, while large file works flawlessly.			
Setup is as follows:			
<ul style="list-style-type: none">• Embedded Router, Gluon, batman-adv 2016.5 connected to Gateway via Ethernet (1500 MTU), speaking BATMAN IV• Gateway, batman-adv 2017.0 connected to Webserver via Ethernet (1500 MTU), reachable via Routing• Webserver ships Firmware manifests and updates			
Since upgrading to batman-adv 2017.0 on the gateway our autoupdater is unable to download its update manifest, a downgrade to 2016.5 resolved this.			
I've attached pcaps of the upper (bat0) and lower (bat0 slave interface) interfaces of both router and gateway, as well as the file that fails to download.			

History

#1 - 03/03/2017 03:23 AM - Martin Weinelt

Sorry, the captures are missing one path because of a different reverse path.

#2 - 03/03/2017 09:03 AM - Sven Eckelmann

- File *webserver_reply* added

Hm, it is rather unfortunate that the download direction is now captured. It looks in the [router-upper.pcap](#) like the transfer finished. The only odd thing is the missing packet at the end. There are roughly 1185 bytes missing at the end which were dropped from the socket and not re-requested by the client

Let's guess the size of the missing packet $1185 + 20$ (TCP) + 12 (TCP options) + 40 (IPv6) 1257 is hopefully smaller than your MTU on the gateway. I forgot now what your lower interface supports. I remember vaguely that it was also 1280 on fastd. $1257 + 10$ (unicast batadv) + 14 (inner ethernet) 1281 byte. This is exactly one byte larger than the allowed mtu on your gateway and should therefore be fragmented for your fastd setup.

Can you check how you've configured your lower interface on the gateway?. Is it set to 1280 ? Did you re-apply [Matthias patch](#) on our gw batman-adv version when you installed batman-adv 2017.0?

I've also read in #gluon that you've tested to use master and saw that it still happened. The theory from Matthias that the last fragment was too small and rejected by the ethernet code/hw/... seems not to be correct. At least it should not happen anymore with it because the packets are now equally sized.

So, now to the possible regression. The changes between 2016.5 and 2017.0 are rather small. It should be bisectable in 4/5 steps - any chance you

could do that to make the search for the regression easier?

I've also attached the expected reply from the webserver. Just to make it a little bit easier to reproduce and have less moving parts involved.

#3 - 03/03/2017 01:56 PM - Martin Weinelt

I downgraded all of our machines to 2016.5 yesterday night and the issue persisted, so it was rather triggered by the manifests filesize than the update to 2017.0.

Our gateway setup looks as follows:

ffda-transport is an untagged VLAN interface in ffda-bat that gets tagged in the host and put in a VLAN trunk, arrives on the switch, and is directly fed, also untagged, into the router for mesh on wan connectivity. We can therefore rule out fastd issues. The webserver is on the L2 on the services interface.

```
hexa@gw01:~$ ip -d l
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT group default qlen 1
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00 promiscuity 0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP mode DEFAULT group default qlen 1000
link/ether da:ff:00:00:01:01 brd ff:ff:ff:ff:ff:ff promiscuity 0
3: ffda-transport: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1280 qdisc pfifo_fast master ffda-bat state UP mode DEFAULT group default
qlen 1000
link/ether da:ff:61:00:01:05 brd ff:ff:ff:ff:ff:ff promiscuity 0
batadv_slave
4: services: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP mode DEFAULT group default qlen 1000
link/ether da:ff:00:00:01:06 brd ff:ff:ff:ff:ff:ff promiscuity 0
5: ffda-br: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP mode DEFAULT group default qlen 1000
link/ether da:ff:61:00:01:04 brd ff:ff:ff:ff:ff:ff promiscuity 0
bridge
6: ffda-bat: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue master ffda-br state UNKNOWN mode DEFAULT group default
qlen 1000
link/ether da:ff:61:00:01:02 brd ff:ff:ff:ff:ff:ff promiscuity 1
batadv
bridge_slave
7: ffda-vpn: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1280 qdisc pfifo_fast master ffda-bat state UNKNOWN mode DEFAULT group
default qlen 1000
link/ether da:ff:61:00:01:03 brd ff:ff:ff:ff:ff:ff promiscuity 0
tun
batadv_slave
```

```
root@salt:/home/hexa# salt 'gw01*' cmd.run 'batctl -m ffda-bat if
gw01.darmstadt.freifunk.net:
ffda-transport: active
ffda-vpn: active
```

All our gateways are currently running last nights HEAD, including the max frag size patch by Matthias and your fragmentation balancing patch. The issue persists.

```
root@salt:/home/hexa# salt 'gw*' cmd.run 'batctl -v'
gw04.darmstadt.freifunk.net:
batctl 2017.0-1-g4cb312c [batman-adv: 2017.0-8-g5283463]
gw03.darmstadt.freifunk.net:
batctl 2017.0-1-g4cb312c [batman-adv: 2017.0-8-g5283463]
gw02.darmstadt.freifunk.net:
batctl 2017.0-1-g4cb312c [batman-adv: 2017.0-8-g5283463]
gw05.darmstadt.freifunk.net:
batctl 2017.0-1-g4cb312c [batman-adv: 2017.0-8-g5283463]
gw01.darmstadt.freifunk.net:
batctl 2017.0-1-g4cb312c [batman-adv: 2017.0-8-g5283463]
gw06.darmstadt.freifunk.net:
batctl 2017.0-1-g4cb312c [batman-adv: 2017.0-8-g5283463]
```

It will try and do a proper capture of the traffic, I'll just have to get consistent routing path, not ecmp like it currently is.

#4 - 03/03/2017 10:59 PM - Martin Weinelt

And one day later I cannot reproduce the issue that I've seen for several hours and across reboots yesterday. :/

#5 - 03/03/2017 11:14 PM - Martin Weinelt

- File *router-upper.pcap* added
- File *router-lower.pcap* added
- File *gw-upper.pcap* added
- File *gw-lower.pcap* added

Anyway, here are the dumps of it working.

#6 - 03/03/2017 11:42 PM - Martin Weinelt

- File *router-upper.pcap* added
- File *router-lower.pcap* added
- File *gw-upper.pcap* added
- File *gw-lower.pcap* added

An it reappeared after downgrading (from HEAD) to 2017.0 release.

#7 - 03/03/2017 11:49 PM - Matthias Schiffer

So as suspected in the Gluon IRC earlier, the fragmentation code creates broken Ethernet frames <60byte for a certain range of lengths, which can be seen in the latest *gw-lower.pcap*. "batman-adv: Keep fragments equally sized" fixes this issue, it should be picked into maint.

I guess something with the test setup was broken when the issue seemed to occur even with the batman-adv master yesterday.

#8 - 03/04/2017 12:35 AM - Sven Eckelmann

- Assignee changed from *batman-adv developers* to *Sven Eckelmann*
- Status changed from *New* to *In Progress*

Ok, lets go through this. Usually the underlying device is responsible for padding an ethernet frame to its correct size (because it is the only thing which knows how big it has to be). What now easily can happen is that a 1280 byte fragment is created and a 40-something byte fragment. This fragment has also an ethernet header and is therefore 55 bytes or more. A 55-59 bytes fragment then gets padded by the underlying device to some larger ethernet frame (for example 60 bytes).

The receiving side will receive it and be unable to re-assemble anything because the fragment header is missing any per-fragment length information. And the combined length is larger then the "complete size" in the header -> rejected as bogus

I have to rewrite the commit message a little bit to make it more clear what it fixes. But who wants to be mentioned in the Reported-by line(s) of the patch?

#9 - 03/04/2017 12:41 AM - Martin Weinelt

You can add me as a reporter:

Martin Weinelt <martin@darmstadt.freifunk.net>

#10 - 03/04/2017 04:37 PM - Sven Eckelmann

- Target version set to 2017.0.1
- Assignee changed from Sven Eckelmann to Simon Wunderlich

Patch with new commit message can be found at <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/20170304153331.22420-2-sven@narfation.org/>

#11 - 03/04/2017 05:30 PM - Sven Eckelmann

Commit message was rewritten again and patch can now be found at <https://patchwork.open-mesh.org/project/b.a.t.m.a.n./patch/20170304162925.559-1-sven@narfation.org/>

#12 - 03/07/2017 05:04 PM - Sven Eckelmann

- Status changed from In Progress to Closed

Patch was applied and released as v2017.0.1

Files

gw-upper.pcap	6.35 KB	03/03/2017	Martin Weinelt
gw-lower.pcap	44.4 KB	03/03/2017	Martin Weinelt
router-lower.pcap	50 KB	03/03/2017	Martin Weinelt
router-upper.pcap	78.5 KB	03/03/2017	Martin Weinelt
experimental.manifest	67 KB	03/03/2017	Martin Weinelt
webserver_reply	67.2 KB	03/03/2017	Sven Eckelmann
gw-lower.pcap	89.1 KB	03/03/2017	Martin Weinelt
gw-upper.pcap	79.7 KB	03/03/2017	Martin Weinelt
router-lower.pcap	89.5 KB	03/03/2017	Martin Weinelt
router-upper.pcap	79.7 KB	03/03/2017	Martin Weinelt
gw-upper.pcap	101 KB	03/03/2017	Martin Weinelt
gw-lower.pcap	123 KB	03/03/2017	Martin Weinelt
router-lower.pcap	123 KB	03/03/2017	Martin Weinelt
router-upper.pcap	78.5 KB	03/03/2017	Martin Weinelt