

batman-adv - Bug #213

batman-adv gets confused if multiple interfaces sharing the same macaddr are bridged to multiple VLANs on the same batX interface

04/22/2015 01:30 AM - Alessandro Bolletta

Status:	Closed	Start date:	04/22/2015
Priority:	Normal	Due date:	
Assignee:	Simon Wunderlich	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	2015.2		
Description			
<p>How to reproduce this bug: if I create 2 VLANs on the top of a bat0 interface and then I bridge bat0.x with eth0.x and and bat0.y with eth0.y, and if eth0.x and eth0.y share the same MAC address, the batman-adv routing engine gets crazy.</p> <p>Infact, if I try through another host facing to bat0.x or/and bat0.y to do a batctl traceroute pointing to the MAC address of the host who shares the same MAC address, I can see that the routing path changes unexpectedly, getting right for a while and then randomly gets wrong, such as in a loop fashion.</p> <p>At the moment, in order to avoid this behaviour, the best workaround I found is having different MAC addresses on every interface bridged on a different vlan configured the same mesh cloud interface (let's assume bat0, for example). It's not so much practical, especially for who wants to configure (and then bridge with external interfaces) many VLAN interfaces on the batman-adv mesh.</p> <p>I'm going to categorize it as a "bug", since it's a batman-adv related problem and I think it should be great if it could be solved.</p>			

History

#1 - 04/22/2015 01:31 AM - Alessandro Bolletta

PS: i'm running on 2014.4.0

#2 - 04/22/2015 03:39 AM - Marek Lindner

- Assignee set to Simon Wunderlich

Thanks for opening the bug report. We will go through each component to see what might be going on.

Simon, should BLA be fine with this kind of setup ?

#3 - 04/22/2015 09:42 AM - Alessandro Bolletta

Also tried with BLA without success. It is the same bug I was talking about on IRC.

#4 - 04/22/2015 09:52 AM - Marek Lindner

Alessandro Bolletta wrote:

Also tried with BLA without success. It is the same bug I was talking about on IRC.

I know what this is about. The BLA question was for Simon, not for you. Everything is under control. ;)

#5 - 04/22/2015 10:20 AM - Alessandro Bolletta

Sorry for the intrusion but I just thought it could be helpful :)

#6 - 04/22/2015 12:57 PM - Simon Wunderlich

Before we try blaming BLA (which obviously is always a good idea ;]), I'd like to ask a few specific questions how to reproduce the problem - I've tried to reproduce it in my KVM setup but didn't see any looping or problems.

This is the setup I tried to reproduce:

- Create two hosts with the same setup as described below * Created br-lan10 which includes bat0.10 and eth1.10. MAC address was adopted from bat0.10 * Created br-lan20 which includes bat0.20 and eth1.20. MAC address was adopted from bat0.20 * br-lan10, br-lan20, bat0.10 and bat0.20 share the same MAC address, eth1.10 and eth1.20 share the same MAC address * batman-adv uses eth0 between the two hosts to connect. eth1 is actually not connected in this setup * Both hosts used different MAC addresses for eth0, eth1, and bat0 respectively * batman-adv 2014.4.0 was used * I can ping from br-lan10 of one device to the other, as well as br-lan20. Also batctl traceroute works

Does this setup resemble yours, or did I misunderstand something? I can't reproduce that problem here so far.

Another question: Do you happen to use the same MAC address among different hosts? That is not supposed to work, since it would be treated as roaming by batman-adv.

#7 - 04/22/2015 04:01 PM - Alessandro Bolletta

Well, the main thing that is missing on your setup is that I have an host (for example, a firewall) facing on those bridges and who shares the same MAC address on both interfaces that are facing the bridges: in this way I'm able to experience the problem I described before. If I set two different MAC addresses on that host, I can workaround the problem, but it is very uneasy to manage things like that, especially when bridged VLANs are many more than 2.

If I reproduce this setup on a common managed I2 switch, for example, I don't see this loop, since two VLANs should be two different broadcast domains and collisions should not exist between them (instead, in this batman-adv setup the collision seems to happen).

Also, another thing that's different on my setup is that the bridge interfaces adopt MAC addresses from the eth0.x interface, and not from the bat0.x, as in your setup.

Another question: Do you happen to use the same MAC address among different hosts? That is not supposed to work, since it would be treated as roaming by batman-adv.

No, I don't use the same MAC address on different hosts but as I said before I would use the same MAC address on different interfaces belonging to the same host.

#8 - 04/22/2015 04:05 PM - Alessandro Bolletta

I can create a visual diagram if you prefer :)

#9 - 04/22/2015 04:08 PM - Simon Wunderlich

We do support this kind of VLAN separation into different broadcast domains, at least since we introduced VLAN support into translation tables. I've also seen this setup working in practice. So yeah, we should find out where that problems comes from you experience here.

So you mean, I'm missing another host which is bridged on the Ethernet port? A diagram which illustrates your minimal setup to replicate would indeed be helpful. :)

Thanks!

#10 - 04/22/2015 04:23 PM - Alessandro Bolletta

- File *batmanadv-vlan-loop-graph.pdf* added

Here the graph. If i change the mac addr of one of the interfaces on the Host A, I solve the problem and traceroute gets ok.

#11 - 04/22/2015 07:34 PM - Simon Wunderlich

Thanks for the graph, that made it much clearer. I've tried to set up a test network accordingly, but still can't reproduce the problem. I've used 4 hosts connected like this to resemble your setup. On each connection, two VLANs 10 and 20 are configured:

Host1 <--- Ethernet---> Host2 <---Mesh---> Host3 <--- Ethernet ---> Host4

The configuration is the same as described in my earlier comment, please find some outputs below from host 1. I have no issues pinging from Host 1 to Host 4 through either VLAN. Host 4 also uses the same MAC on these VLANs.

So unfortunately it works fine[tm] over here. Therefore, a few more questions:

- Are the MAC addresses of your eth0 and your WiFi device you run your mesh on different in one device? * Is it possible that you have a bridged connection somewhere between your two VLANs? * Do you see "flapping" of MAC address you try to ping among the originators when doing "batctl tg"? * Am I missing something else? :)

Just as another note, if you do batctl traceroute it will not actually ping your Host A, but only the router which is the last hop in the mesh - that would be the left "batman-adv router" in your illustration if everything is working correctly.

Host 1 configuration (for references, the others look similar):

```
root@OpenWrt:/# brctl show
bridge name      bridge id          STP enabled      interfaces
br-lan10         8000.fef10000101  no               eth1.10
                 8000.fef10000101  no               bat0.10
br-lan20         8000.fef10000101  no               eth1.20
                 8000.fef10000101  no               bat0.20

root@OpenWrt:/# batctl if
eth0: active
root@OpenWrt:/# ifconfig -a | grep HWaddr
bat0      Link encap:Ethernet HWaddr 7A:6F:5F:D5:12:DB
bat0.10   Link encap:Ethernet HWaddr 7A:6F:5F:D5:12:DB
bat0.20   Link encap:Ethernet HWaddr 7A:6F:5F:D5:12:DB
br-lan10  Link encap:Ethernet HWaddr FE:F1:00:00:01:01
br-lan20  Link encap:Ethernet HWaddr FE:F1:00:00:01:01
eth0      Link encap:Ethernet HWaddr FE:F0:00:00:01:01
eth1      Link encap:Ethernet HWaddr FE:F1:00:00:01:01
eth1.10   Link encap:Ethernet HWaddr FE:F1:00:00:01:01
eth1.20   Link encap:Ethernet HWaddr FE:F1:00:00:01:01
root@OpenWrt:/# ip a l | grep inet.*br-lan
    inet 192.168.10.1/24 brd 192.168.10.255 scope global br-lan10
    inet 192.168.20.1/24 brd 192.168.20.255 scope global br-lan20
```

#12 - 04/22/2015 08:57 PM - Sven Eckelmann

It may also be interesting to know which kernel is used in this setup. Because it is already known that 2014.4.0 is not compatible with kernels < 2.6.39 when bridges are involved. Maybe there is a similar problem with VLAN.

There was already a bug in older kernels when using VLAN . It affected kernels < 3.8.0 and Simon worked around the bug with an hack for skb_share_check. But this bug was only affecting batman-adv when the slave device was a VLAN device.

#13 - 04/27/2015 07:01 PM - Alessandro Bolletta

The kernel version is 3.10.x

Maybe the bug may be reproduced if there are several intermediate mesh nodes between the source and the destination hosts.

I will try to reproduce it in my mesh net and I'll try to inspect it deeper.

ps. if you remember in the IRC channel marec said me that using the same MAC address in this environment is not supported by batman-adv, since batman-adv uses mac address to address the destination path. maybe I explained not so clearly my problem...

#14 - 04/27/2015 07:27 PM - Alessandro Bolletta

- *Are the MAC addresses of your eth0 and your WiFi device you run your mesh on different in one device? They are in two different devices, one handles batman-adv, the other one only bridges the WiFi connection through a bridge*

My question is just if your wlan0 has a different MAC than your eth0. I've seen that some (cheap) routers don't assign different addresses which sometimes leads to problems. The question is really about the MAC address only.

- *Do you see "flapping" of MAC address you try to ping among the originators when doing "batctl tg"? Yes, if I ping intermediate hosts that I get printed on the I2 traceroute, I can ping them successfully.*

My question was if the MAC address of your destination in the global translation table stays at the same originator or if it is moving around from one host to the other? You can check by checking "batctl tg" multiple times and grepping for the address you are pinged. It's a good thing that the rest of the mesh is stable though. ;)

Let me know how you want to proceed on that.

#15 - 05/17/2015 09:54 PM - Alessandro Bolletta

Ok, here's the problem I was talking about. Today it is getting down one of the hosts bridged into the mesh cloud again.

If I do a "watch" in order to update with 1second frequency the "batctl tr" command I get, randomly, this behaviour:

- this for a while:

```
Every 1s: batctl tr 00:15:5D:0A:01:21          2015-05-17 19:45:21
```

```
traceroute to 00:15:5D:0A:01:21 (d4:ca:6d:0b:97:8a), 50 hops max, 20 byte packets
```

```
1: d4:ca:6d:18:97:26 1.015 ms 0.659 ms 0.696 ms
```

```
2: d4:ca:6d:0b:97:8a 1.329 ms 1.389 ms 1.509 ms
```

- and then it goes back to the correct path, which is this:

Every 1s: batctl tr 00:15:5D:0A:01:21

2015-05-17 19:46:22

traceroute to 00:15:5D:0A:01:21 (d4:ca:6d:18:97:2c), 50 hops max, 20 byte packets

```
1: d4:ca:6d:18:97:26 0.876 ms 1.067 ms 0.712 ms
2: d4:ca:6d:0b:97:8a 1.538 ms 1.172 ms 1.098 ms
3: d4:ca:6d:54:fb:20 1.861 ms 2.235 ms 1.561 ms
4: d4:ca:6d:0b:a1:ee 2.743 ms 2.695 ms 2.734 ms
5: d4:ca:6d:18:97:2c 2.979 ms 2.953 ms 3.165 ms
```

I'm sure that there's no mac addr duplication on the network.
Why?

#16 - 05/18/2015 08:39 AM - Alessandro Bolletta

Same problem on another node this morning.

Yesterday, in order to solve the problem, I had to change the MAC address.

Today it's gonna be the same for the host with this mac addr:00:15:5d:ff:7f:0b

I get randomly this:

Every 1s: batctl tr 00:15:5d:ff:7f:0b

2015-05-18 06:36:35

traceroute to 00:15:5d:ff:7f:0b (d4:ca:6d:55:0a:70), 50 hops max, 20 byte packets

```
1: d4:ca:6d:18:97:26 2.486 ms 0.627 ms 0.801 ms
2: d4:ca:6d:0b:97:8a 2.096 ms 1.421 ms 1.231 ms
3: d4:ca:6d:54:fb:20 2.833 ms 3.310 ms 2.808 ms
4: d4:ca:6d:0b:97:8f 4.313 ms 4.487 ms 2.882 ms
5: d4:ca:6d:18:97:7b 4.009 ms 3.450 ms 8.085 ms
6: d4:ca:6d:55:0a:70 13.252 ms 11.158 ms 14.812 ms
```

and then this:

Every 1s: batctl tr 00:15:5d:ff:7f:0b

2015-05-18 06:38:44

traceroute to 00:15:5d:ff:7f:0b (d4:ca:6d:18:97:7b), 50 hops max, 20 byte packets

```
1: d4:ca:6d:18:97:26 0.686 ms 0.605 ms 0.627 ms
2: d4:ca:6d:0b:97:8a 2.007 ms 1.577 ms 1.491 ms
3: d4:ca:6d:54:fb:20 3.398 ms 2.827 ms 4.761 ms
4: d4:ca:6d:0b:97:8f 3.937 ms 2.583 ms 5.839 ms
5: d4:ca:6d:18:97:7b 5.203 ms 5.526 ms 5.883 ms
```

Obviously the right path is the first one.

#17 - 05/19/2015 12:14 PM - Simon Wunderlich

That address "00:15:5d:ff:7f:0b" seems to be from a cisco router within your network. It seems the address is announced by different hosts. I'm not sure if I interpret correctly, but it seems there is a flapping of the global translation table happening here.

When the problem happens again, could you please do:

```
watch "batctl tg | grep 00:15:5d:ff:7f:0b"
```

(replace the MAC as needed). It would be interesting to see if the "via" part switches from one batman originator to another.

Thanks

#18 - 07/10/2015 11:21 AM - Marek Lindner

Any update on this matter ?

#19 - 07/10/2015 02:00 PM - Simon Wunderlich

I'm working with Alessandro offline on this. So far, it seems that there is a translation table divergence between multiple hosts, which isn't resolved automatically - global and translation table do not match, although TTVN, CRC, etc are the same.

We are updating versions to the latest fixes to see if this resolves the problem, and will be digging further otherwise.

#20 - 01/29/2016 10:43 AM - Simon Wunderlich

- Status changed from New to Resolved

This issue has been fixed with release 2015.2

#21 - 01/29/2016 10:44 AM - Simon Wunderlich

- Status changed from Resolved to Closed

#22 - 02/11/2017 08:50 AM - Sven Eckelmann

- Target version set to 2015.2

Files

batmanadv-vlan-loop-graph.pdf	386 KB	04/22/2015	Alessandro Bolletta
-------------------------------	--------	------------	---------------------