(Some) Anomalies in Active Network Measurements

Robert Kisteleki, IETF116 MAPRG

CPE / Middlebox Behaviour

- A “real life” traceroute result ->
  - Caused by a particular network device
- It is not a unique case
- How would your analysis react to this?

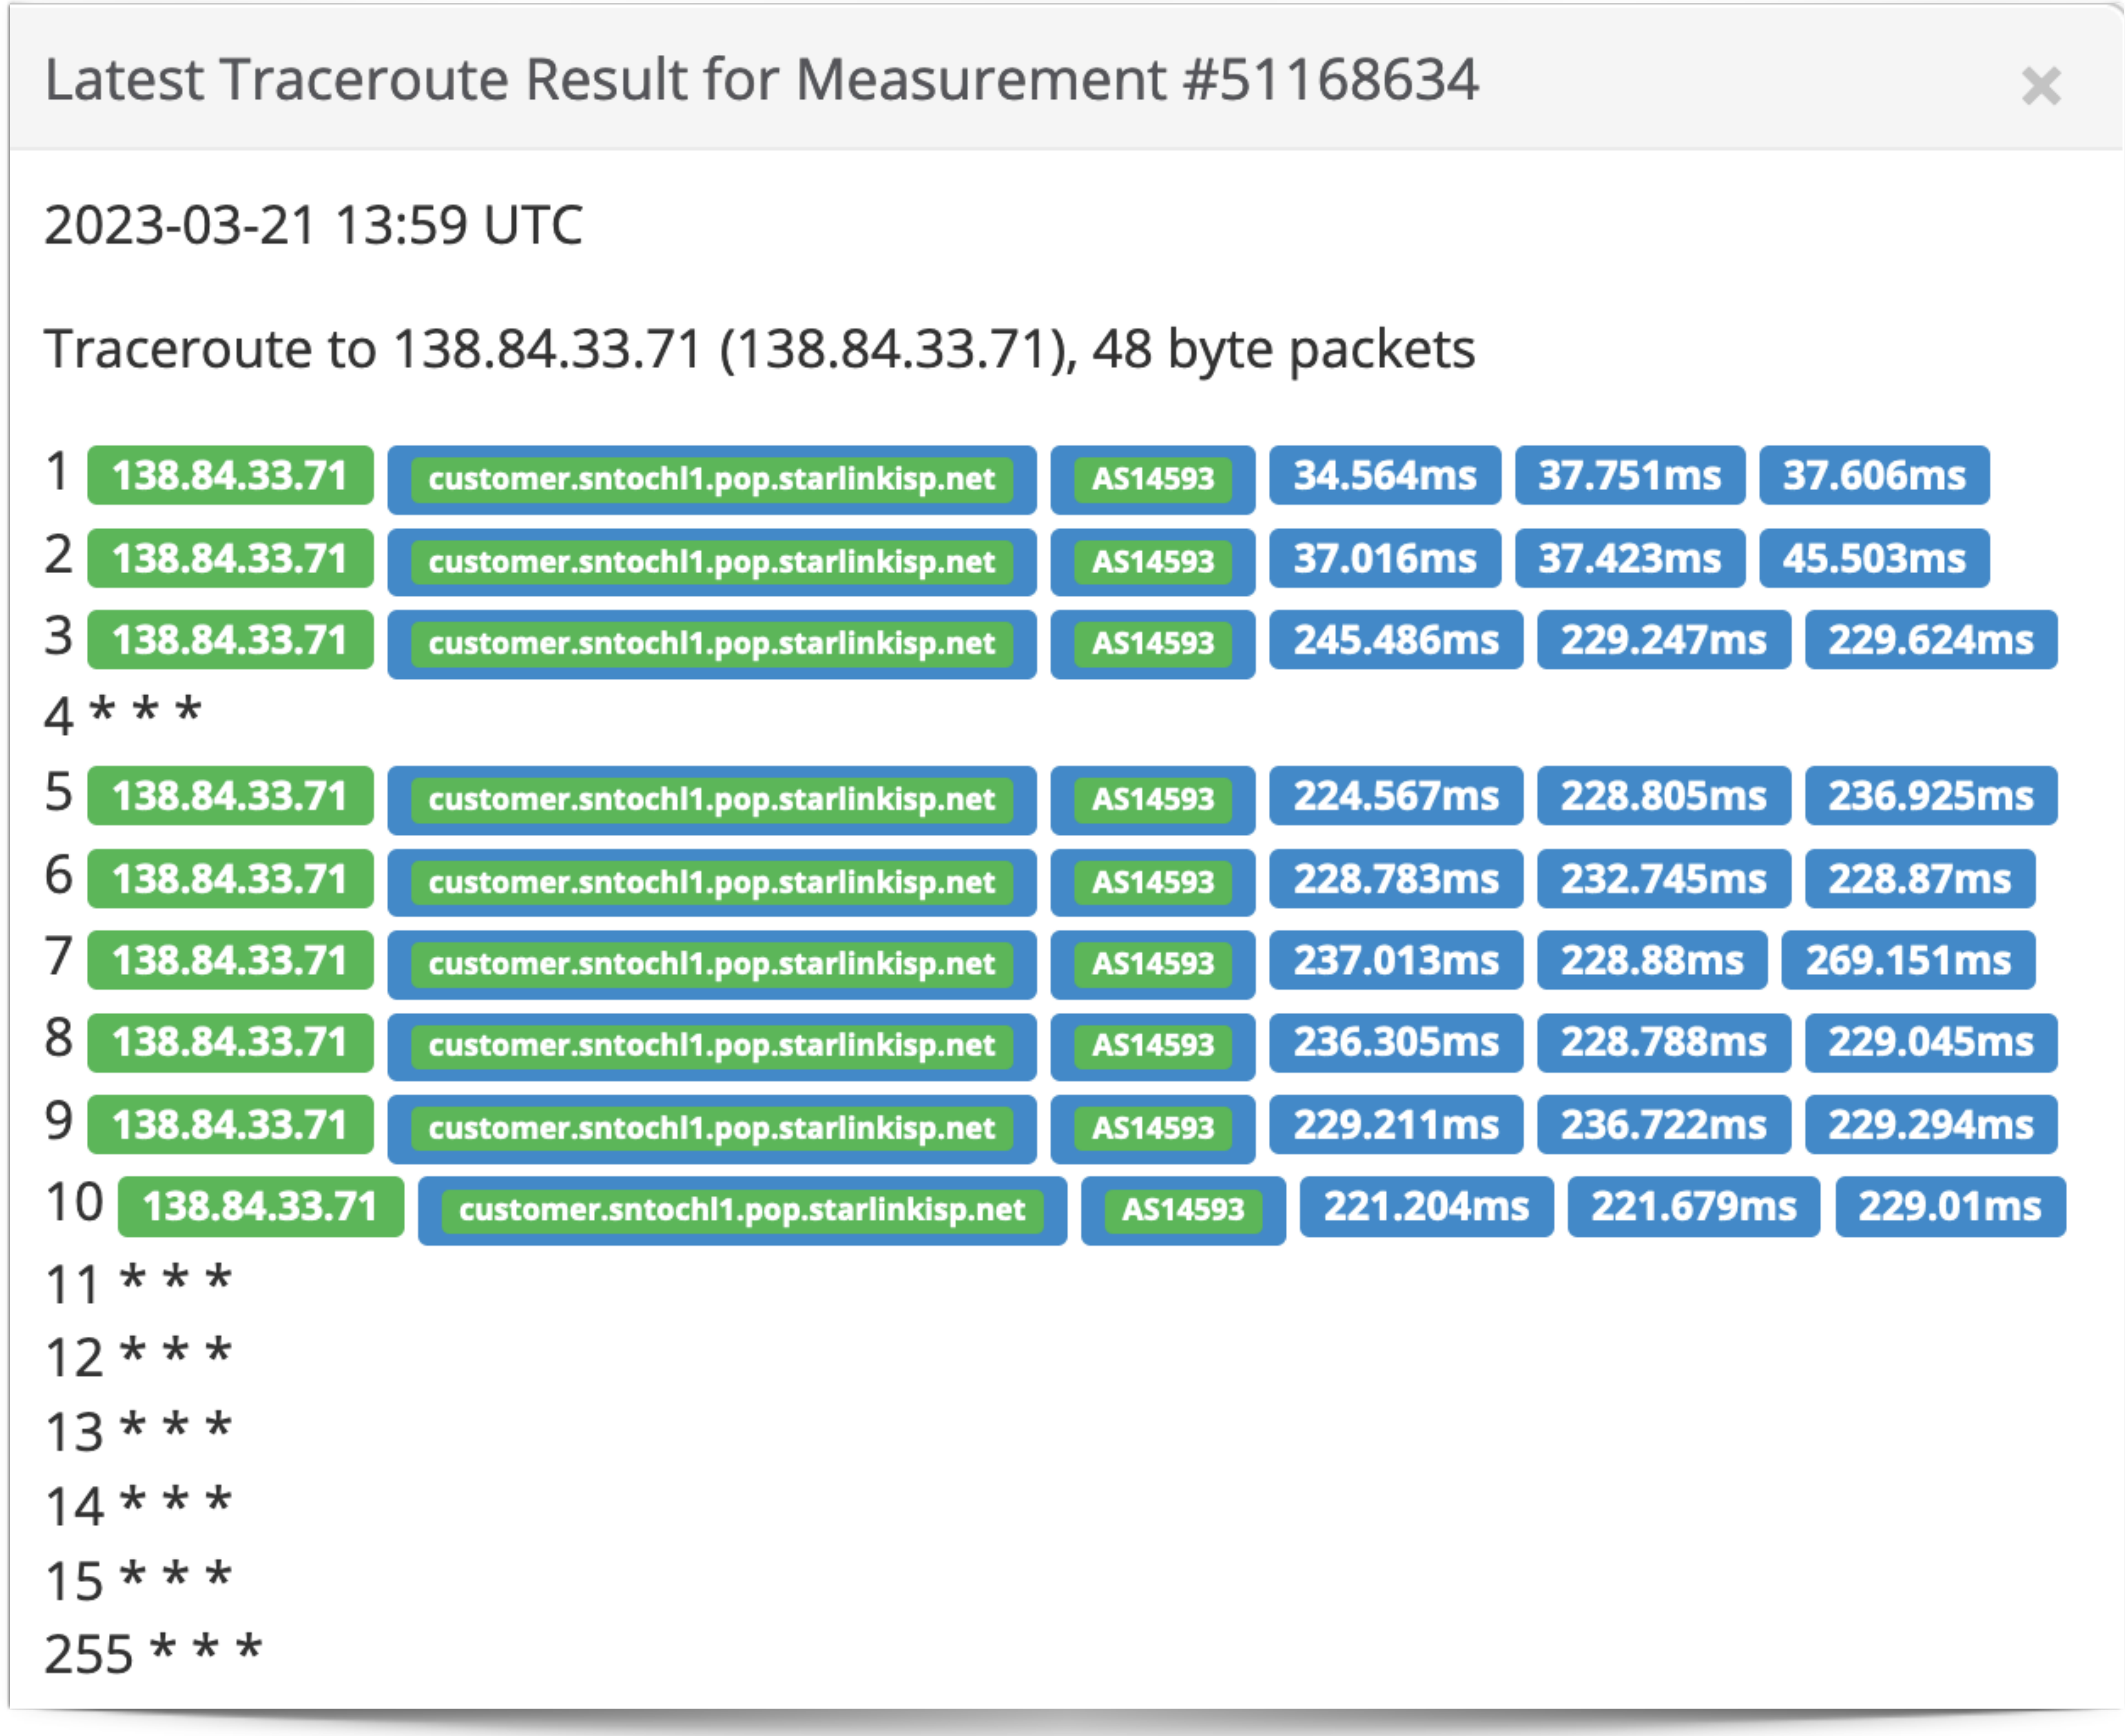

Latest Traceroute Result for Measurement #51168634

2023-03-21 13:59 UTC

Traceroute to 138.84.33.71 (138.84.33.71), 48 byte packets

```
1 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 34.564ms 37.751ms 37.606ms
2 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 37.016ms 37.423ms 45.503ms
3 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 245.486ms 229.247ms 229.624ms
4 * * *
5 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 224.567ms 228.805ms 236.925ms
6 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 228.783ms 232.745ms 228.87ms
7 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 237.013ms 228.88ms 269.151ms
8 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 236.305ms 228.788ms 229.045ms
9 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 229.211ms 236.722ms 229.294ms
10 138.84.33.71 customer.sntochl1.pop.starlinkisp.net AS14593 221.204ms 221.679ms 229.01ms
11 * * *
12 * * *
13 * * *
14 * * *
15 * * *
255 * * *
```

What IP addresses can be used? (1)

- Can you see packets from ::/128?
  - Yes you can!
- Although, in theory, you shouldn't:

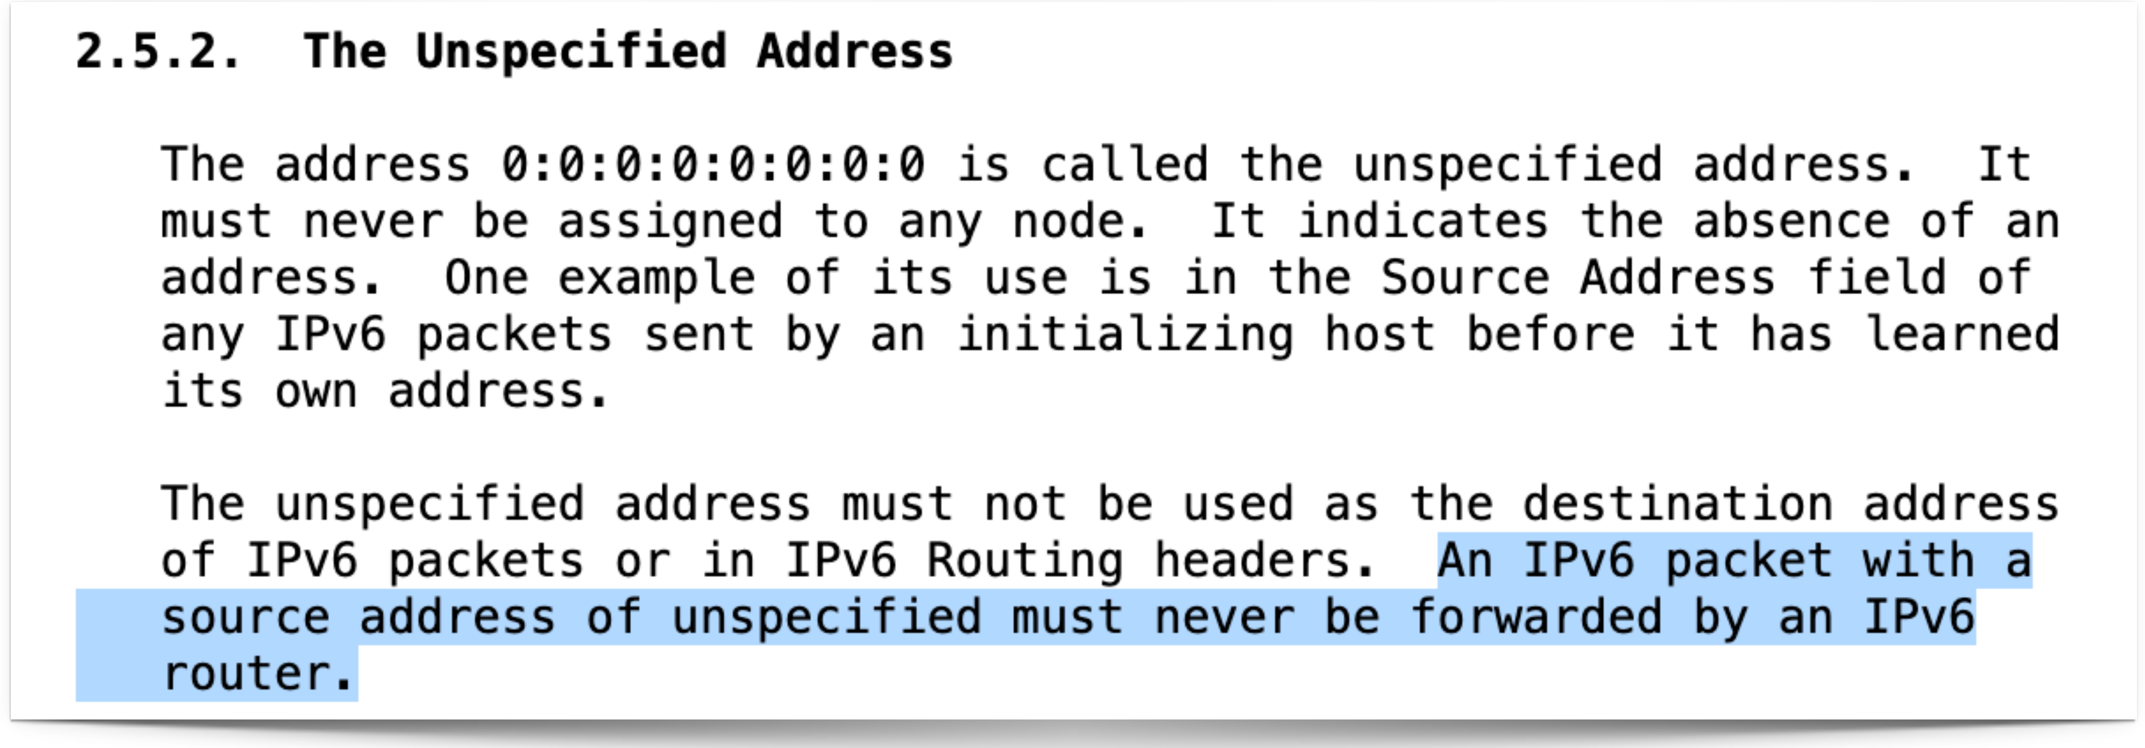

```
2.5.2.  The Unspecified Address

   The address 0:0:0:0:0:0:0:0 is called the unspecified address.  It
   must never be assigned to any node.  It indicates the absence of an
   address.  One example of its use is in the Source Address field of
   any IPv6 packets sent by an initializing host before it has learned
   its own address.

   The unspecified address must not be used as the destination address
   of IPv6 packets or in IPv6 Routing headers.  An IPv6 packet with a
   source address of unspecified must never be forwarded by an IPv6
   router.
```

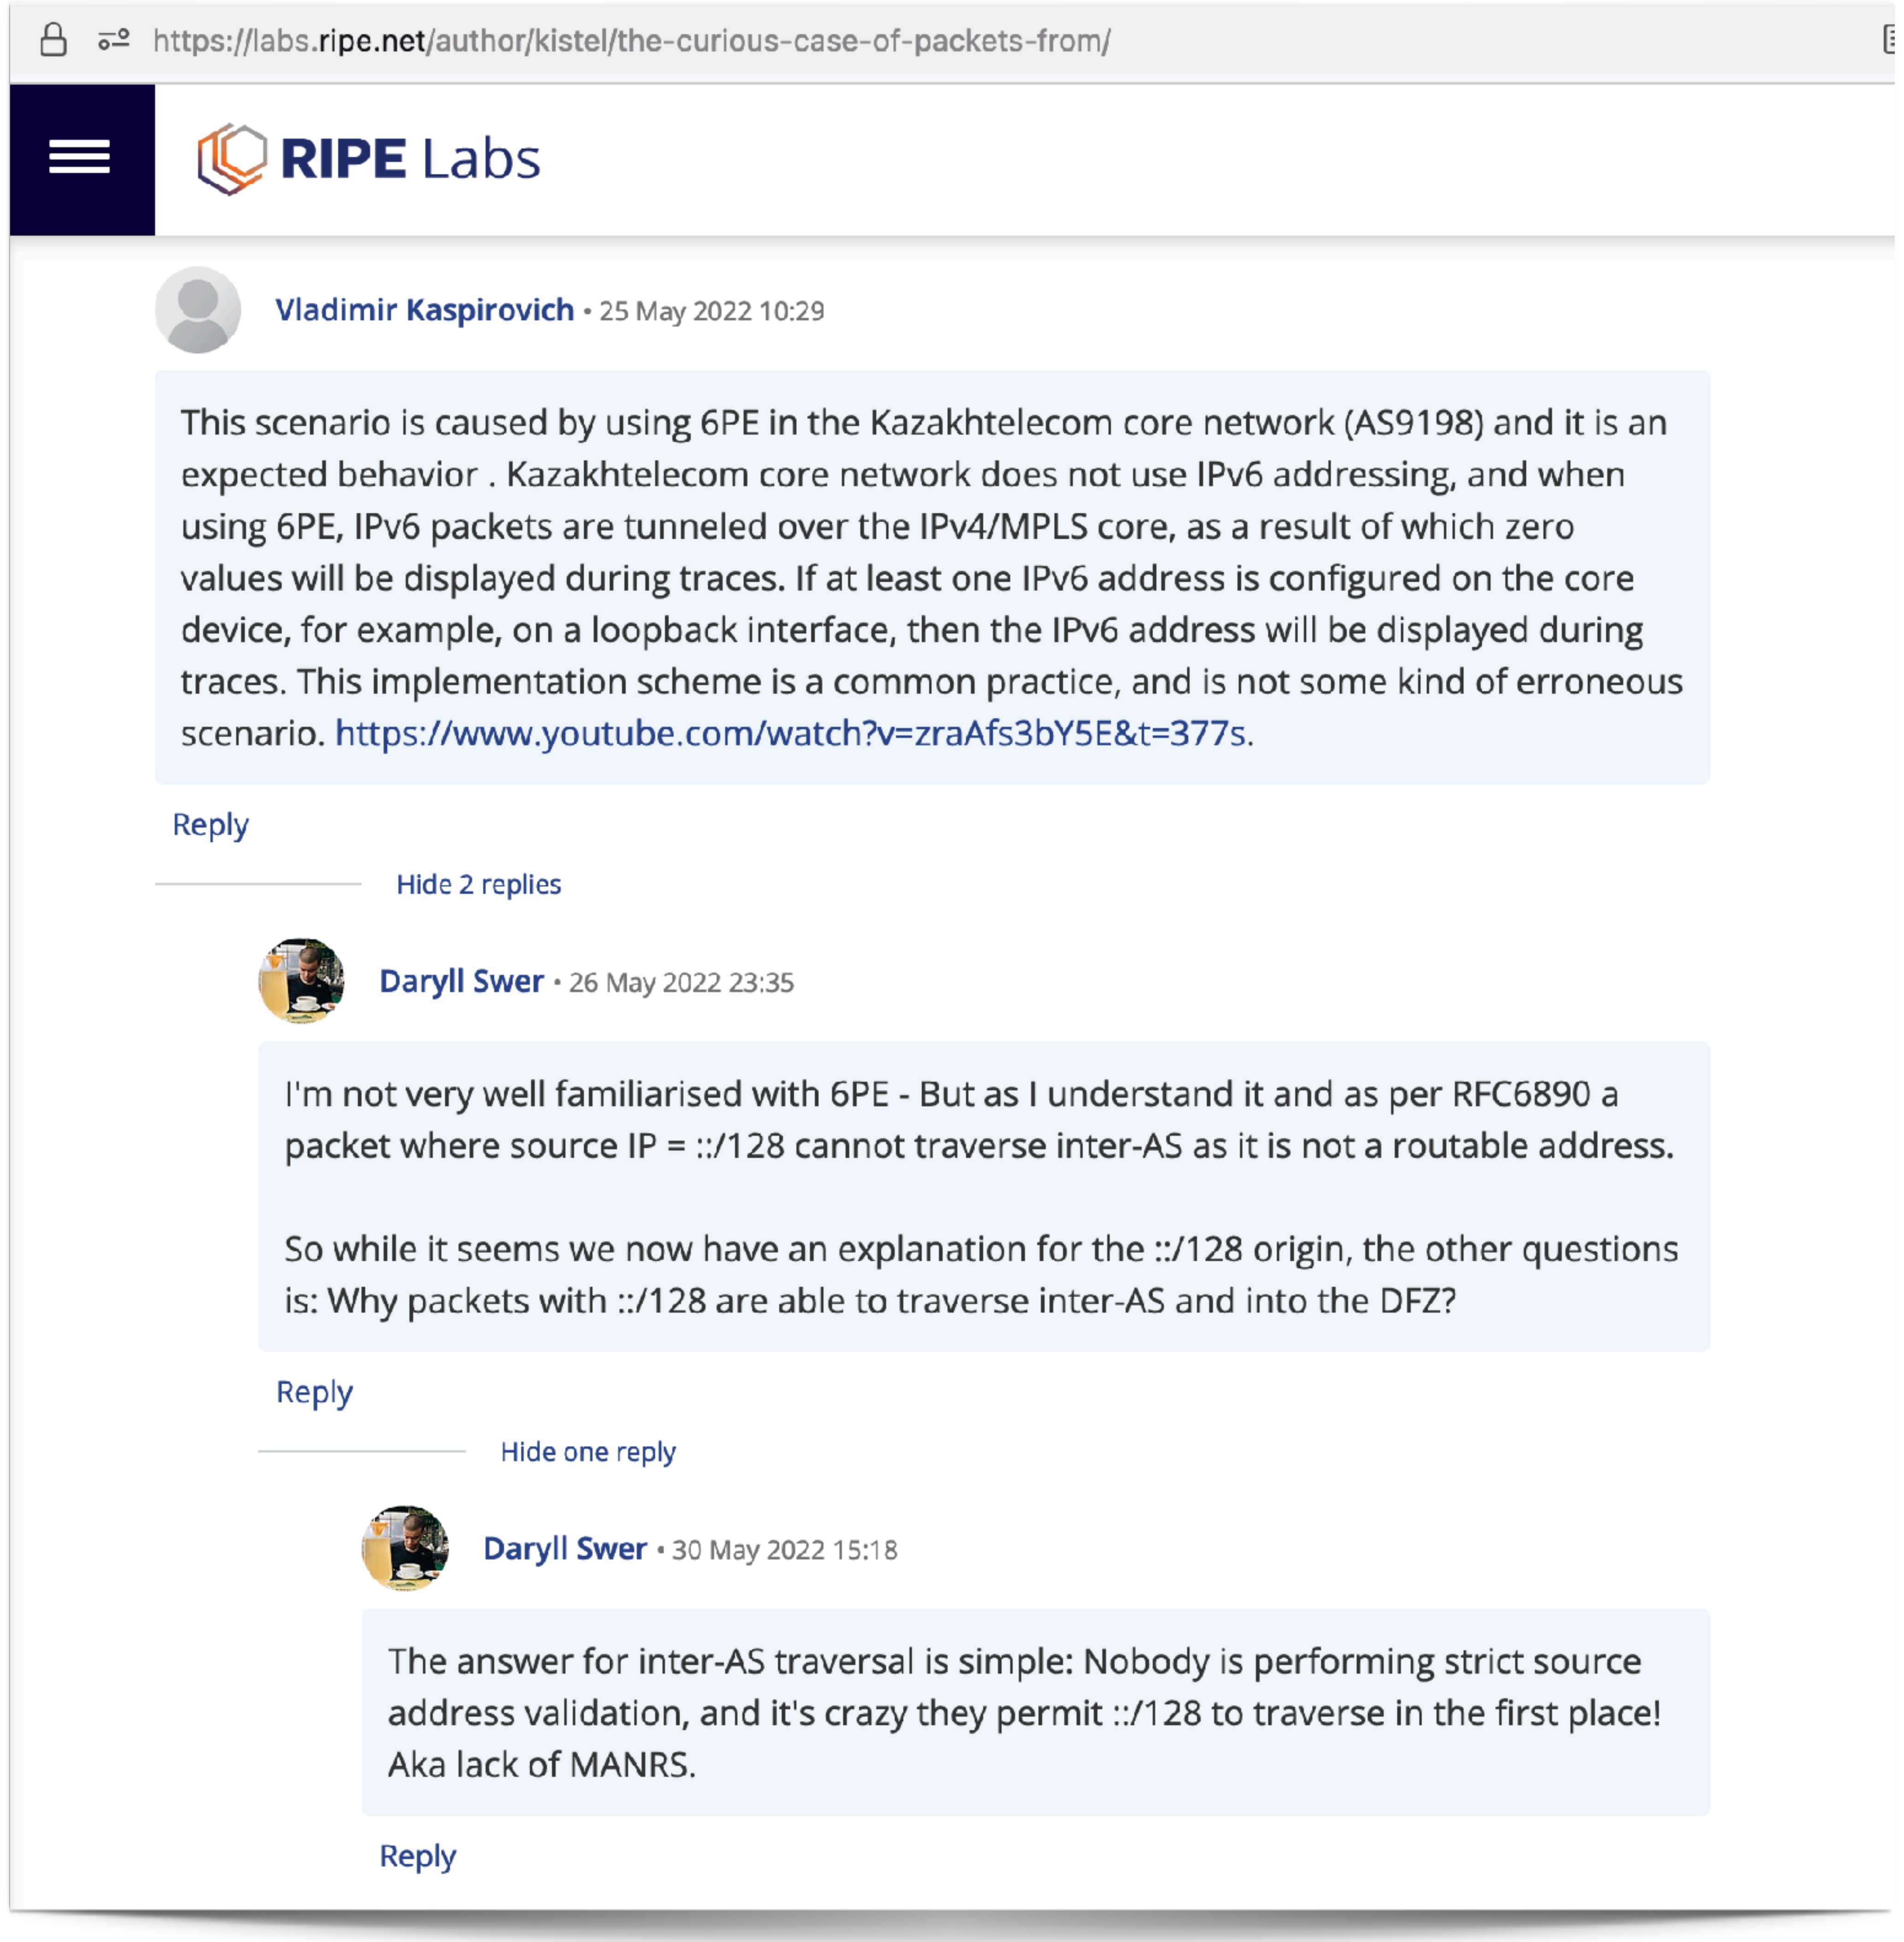

https://labs.ripe.net/author/kistel/the-curious-case-of-packets-from/

Vladimir Kaspirovich • 25 May 2022 10:29

This scenario is caused by using 6PE in the Kazakhtelecom core network (AS9198) and it is an expected behavior . Kazakhtelecom core network does not use IPv6 addressing, and when using 6PE, IPv6 packets are tunneled over the IPv4/MPLS core, as a result of which zero values will be displayed during traces. If at least one IPv6 address is configured on the core device, for example, on a loopback interface, then the IPv6 address will be displayed during traces. This implementation scheme is a common practice, and is not some kind of erroneous scenario. https://www.youtube.com/watch?v=zraAfs3bY5E&t=377s.

Daryll Swer • 26 May 2022 23:35

I'm not very well familiarised with 6PE - But as I understand it and as per RFC6890 a packet where source IP = ::/128 cannot traverse inter-AS as it is not a routable address.

So while it seems we now have an explanation for the ::/128 origin, the other questions is: Why packets with ::/128 are able to traverse inter-AS and into the DFZ?

Daryll Swer • 30 May 2022 15:18

The answer for inter-AS traversal is simple: Nobody is performing strict source address validation, and it's crazy they permit ::/128 to traverse in the first place! Aka lack of MANRS.

What IP addresses can be used? (2)

- How about IPv4 240/4?
  - That’s in use too.

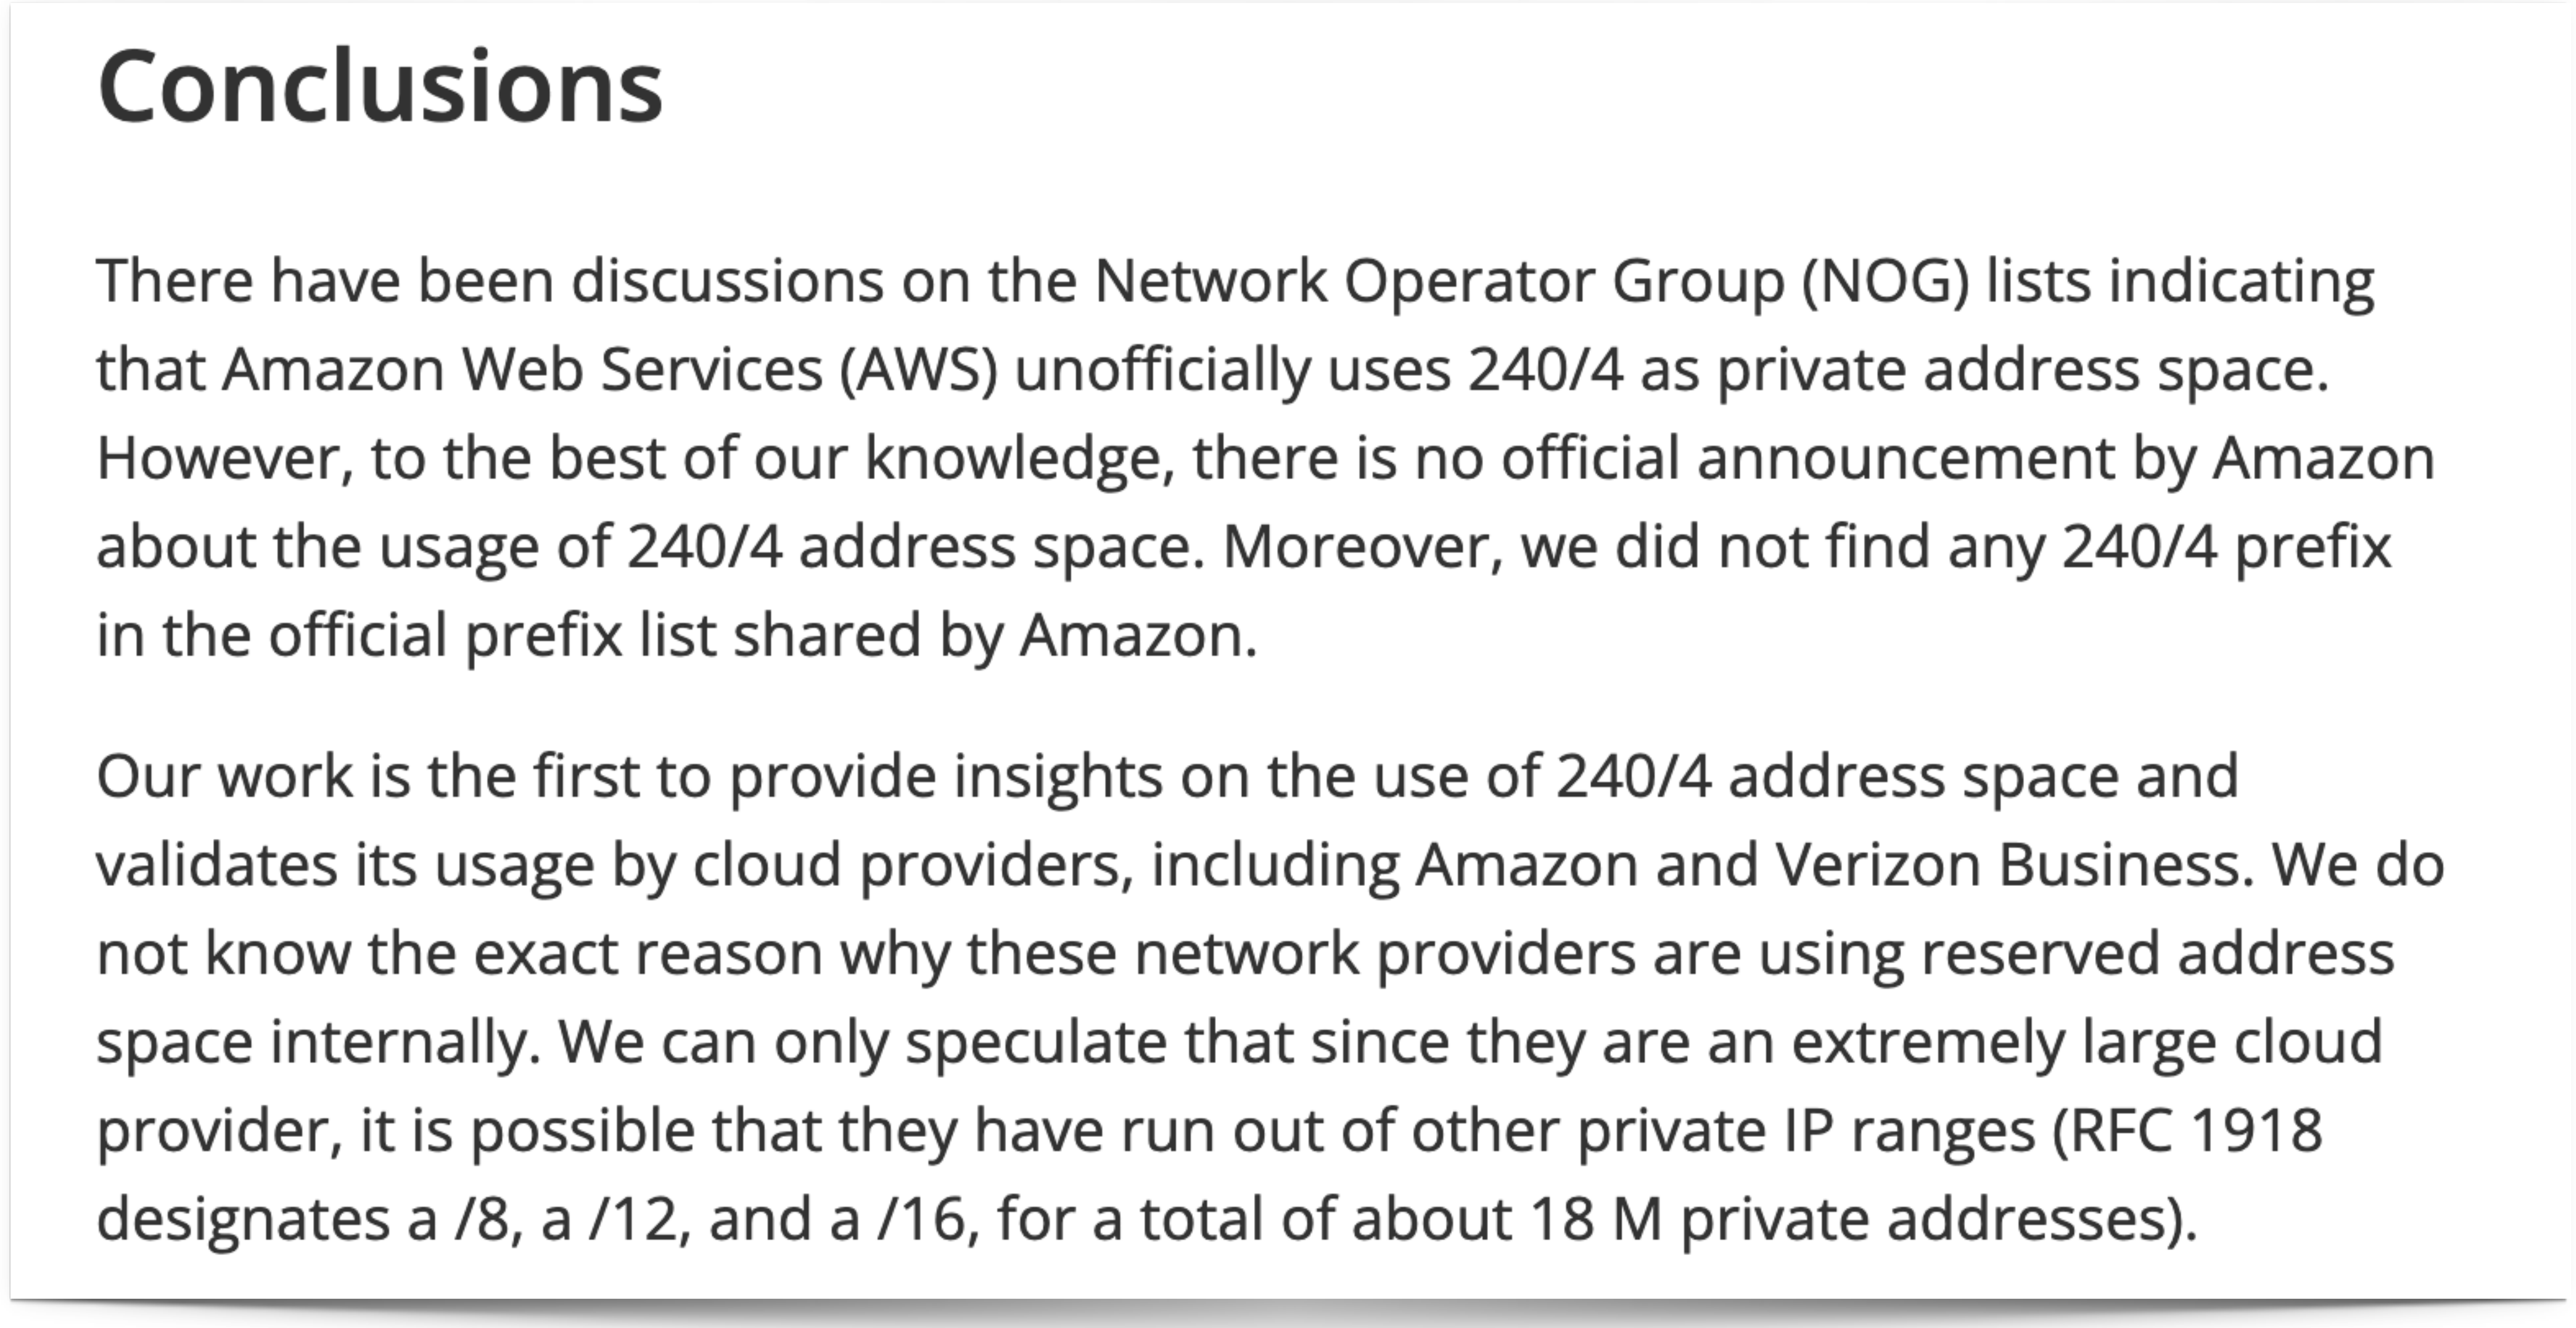

Conclusions

There have been discussions on the Network Operator Group (NOG) lists indicating that Amazon Web Services (AWS) unofficially uses 240/4 as private address space. However, to the best of our knowledge, there is no official announcement by Amazon about the usage of 240/4 address space. Moreover, we did not find any 240/4 prefix in the official prefix list shared by Amazon.

Our work is the first to provide insights on the use of 240/4 address space and validates its usage by cloud providers, including Amazon and Verizon Business. We do not know the exact reason why these network providers are using reserved address space internally. We can only speculate that since they are an extremely large cloud provider, it is possible that they have run out of other private IP ranges (RFC 1918 designates a /8, a /12, and a /16, for a total of about 18 M private addresses).

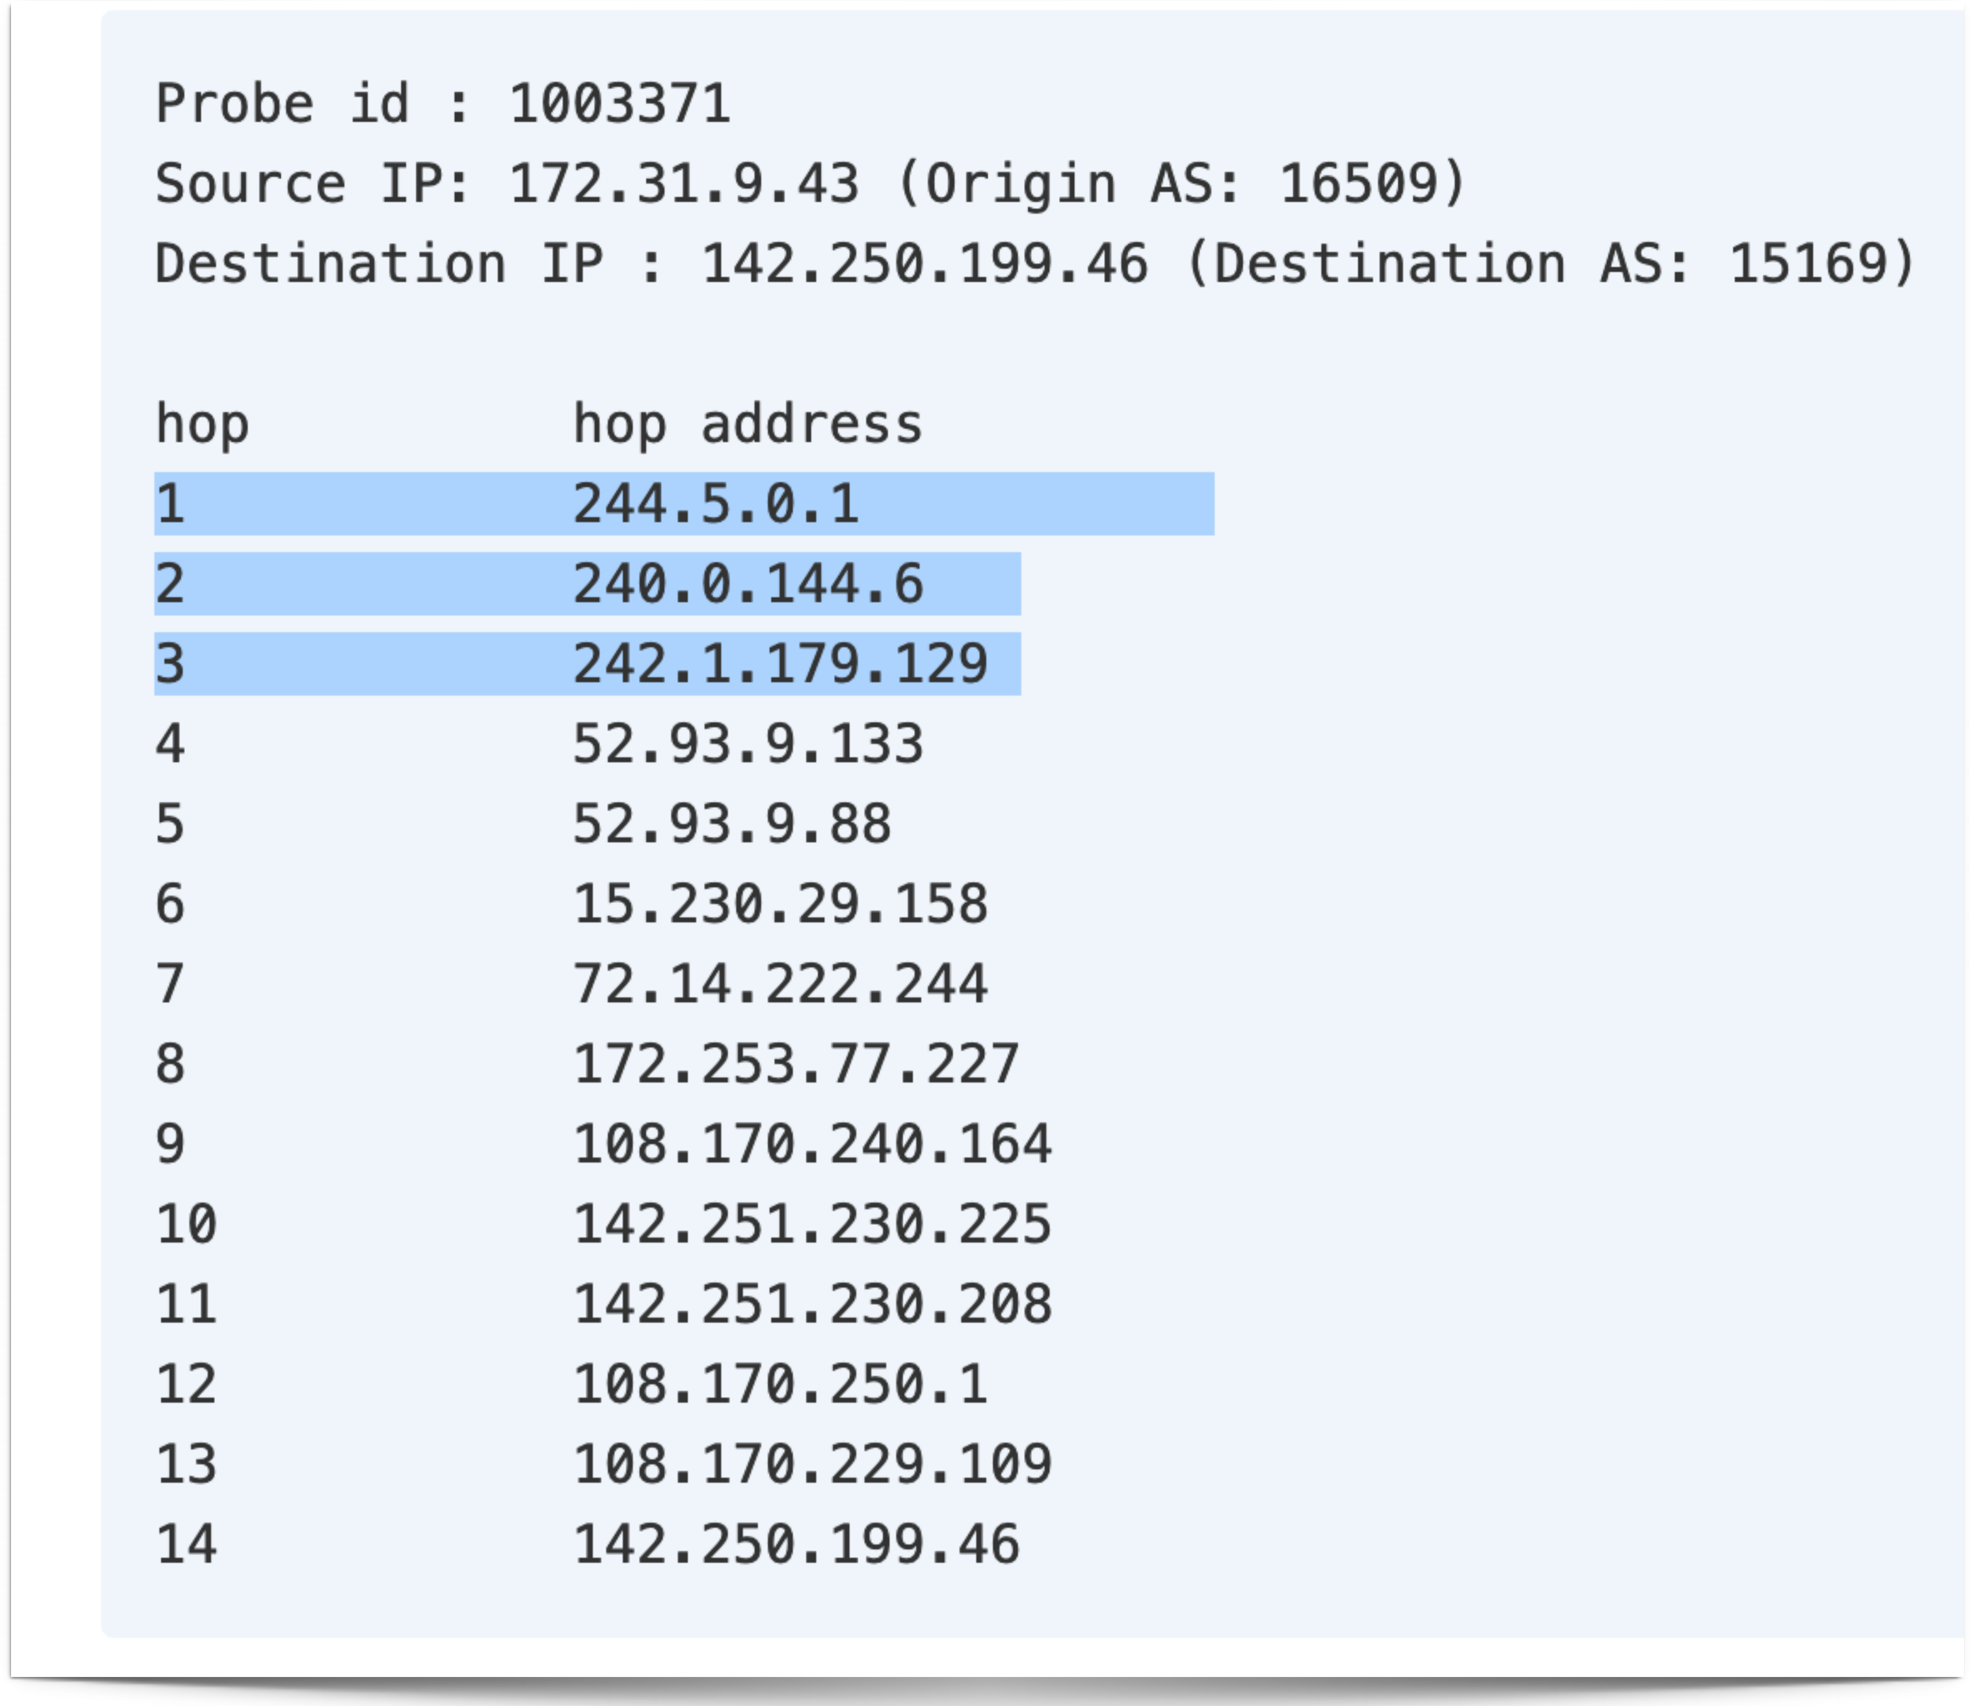

```
Probe id : 1003371
Source IP: 172.31.9.43 (Origin AS: 16509)
Destination IP : 142.250.199.46 (Destination AS: 15169)

hop         hop address
1           244.5.0.1
2           240.0.144.6
3           242.1.179.129
4           52.93.9.133
5           52.93.9.88
6           15.230.29.158
7           72.14.222.244
8           172.253.77.227
9           108.170.240.164
10          142.251.230.225
11          142.251.230.208
12          108.170.250.1
13          108.170.229.109
14          142.250.199.46
```

Correct Geolocation (i.e. Not An Anomaly)

- “Where is France?”

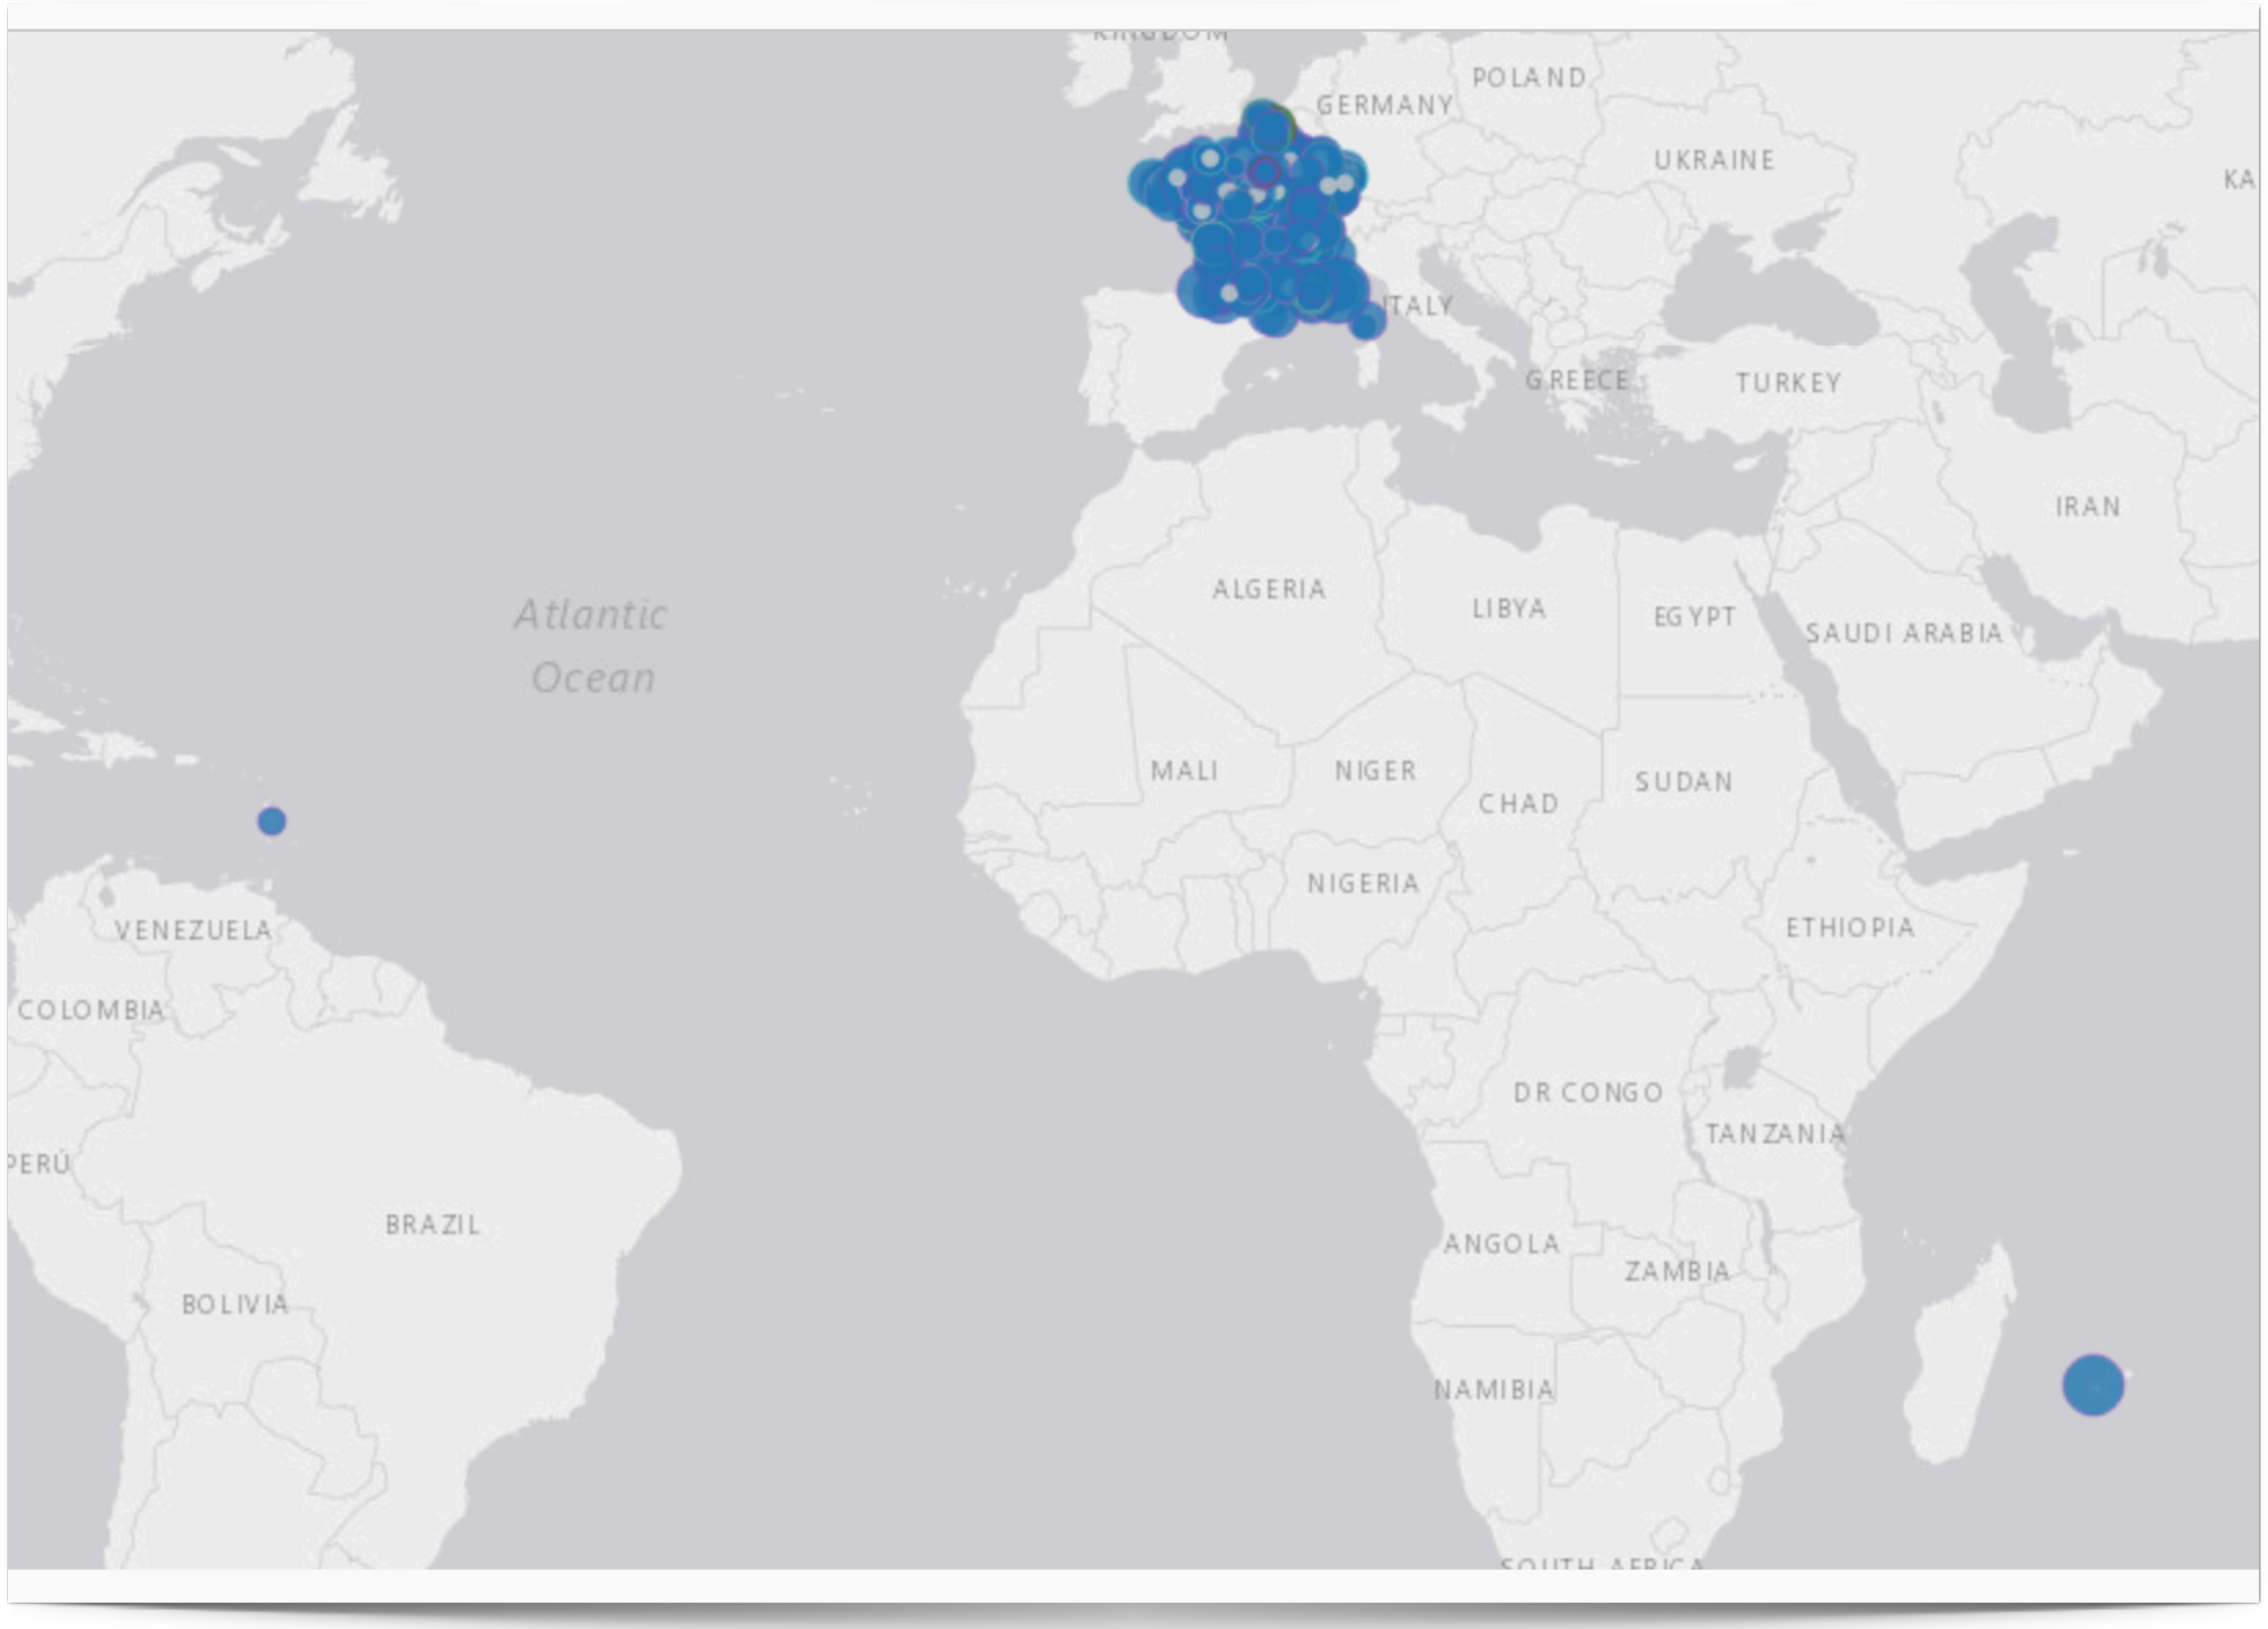

Incorrect Geolocation

- Example: probe in BR, with RTT to DE<13ms:

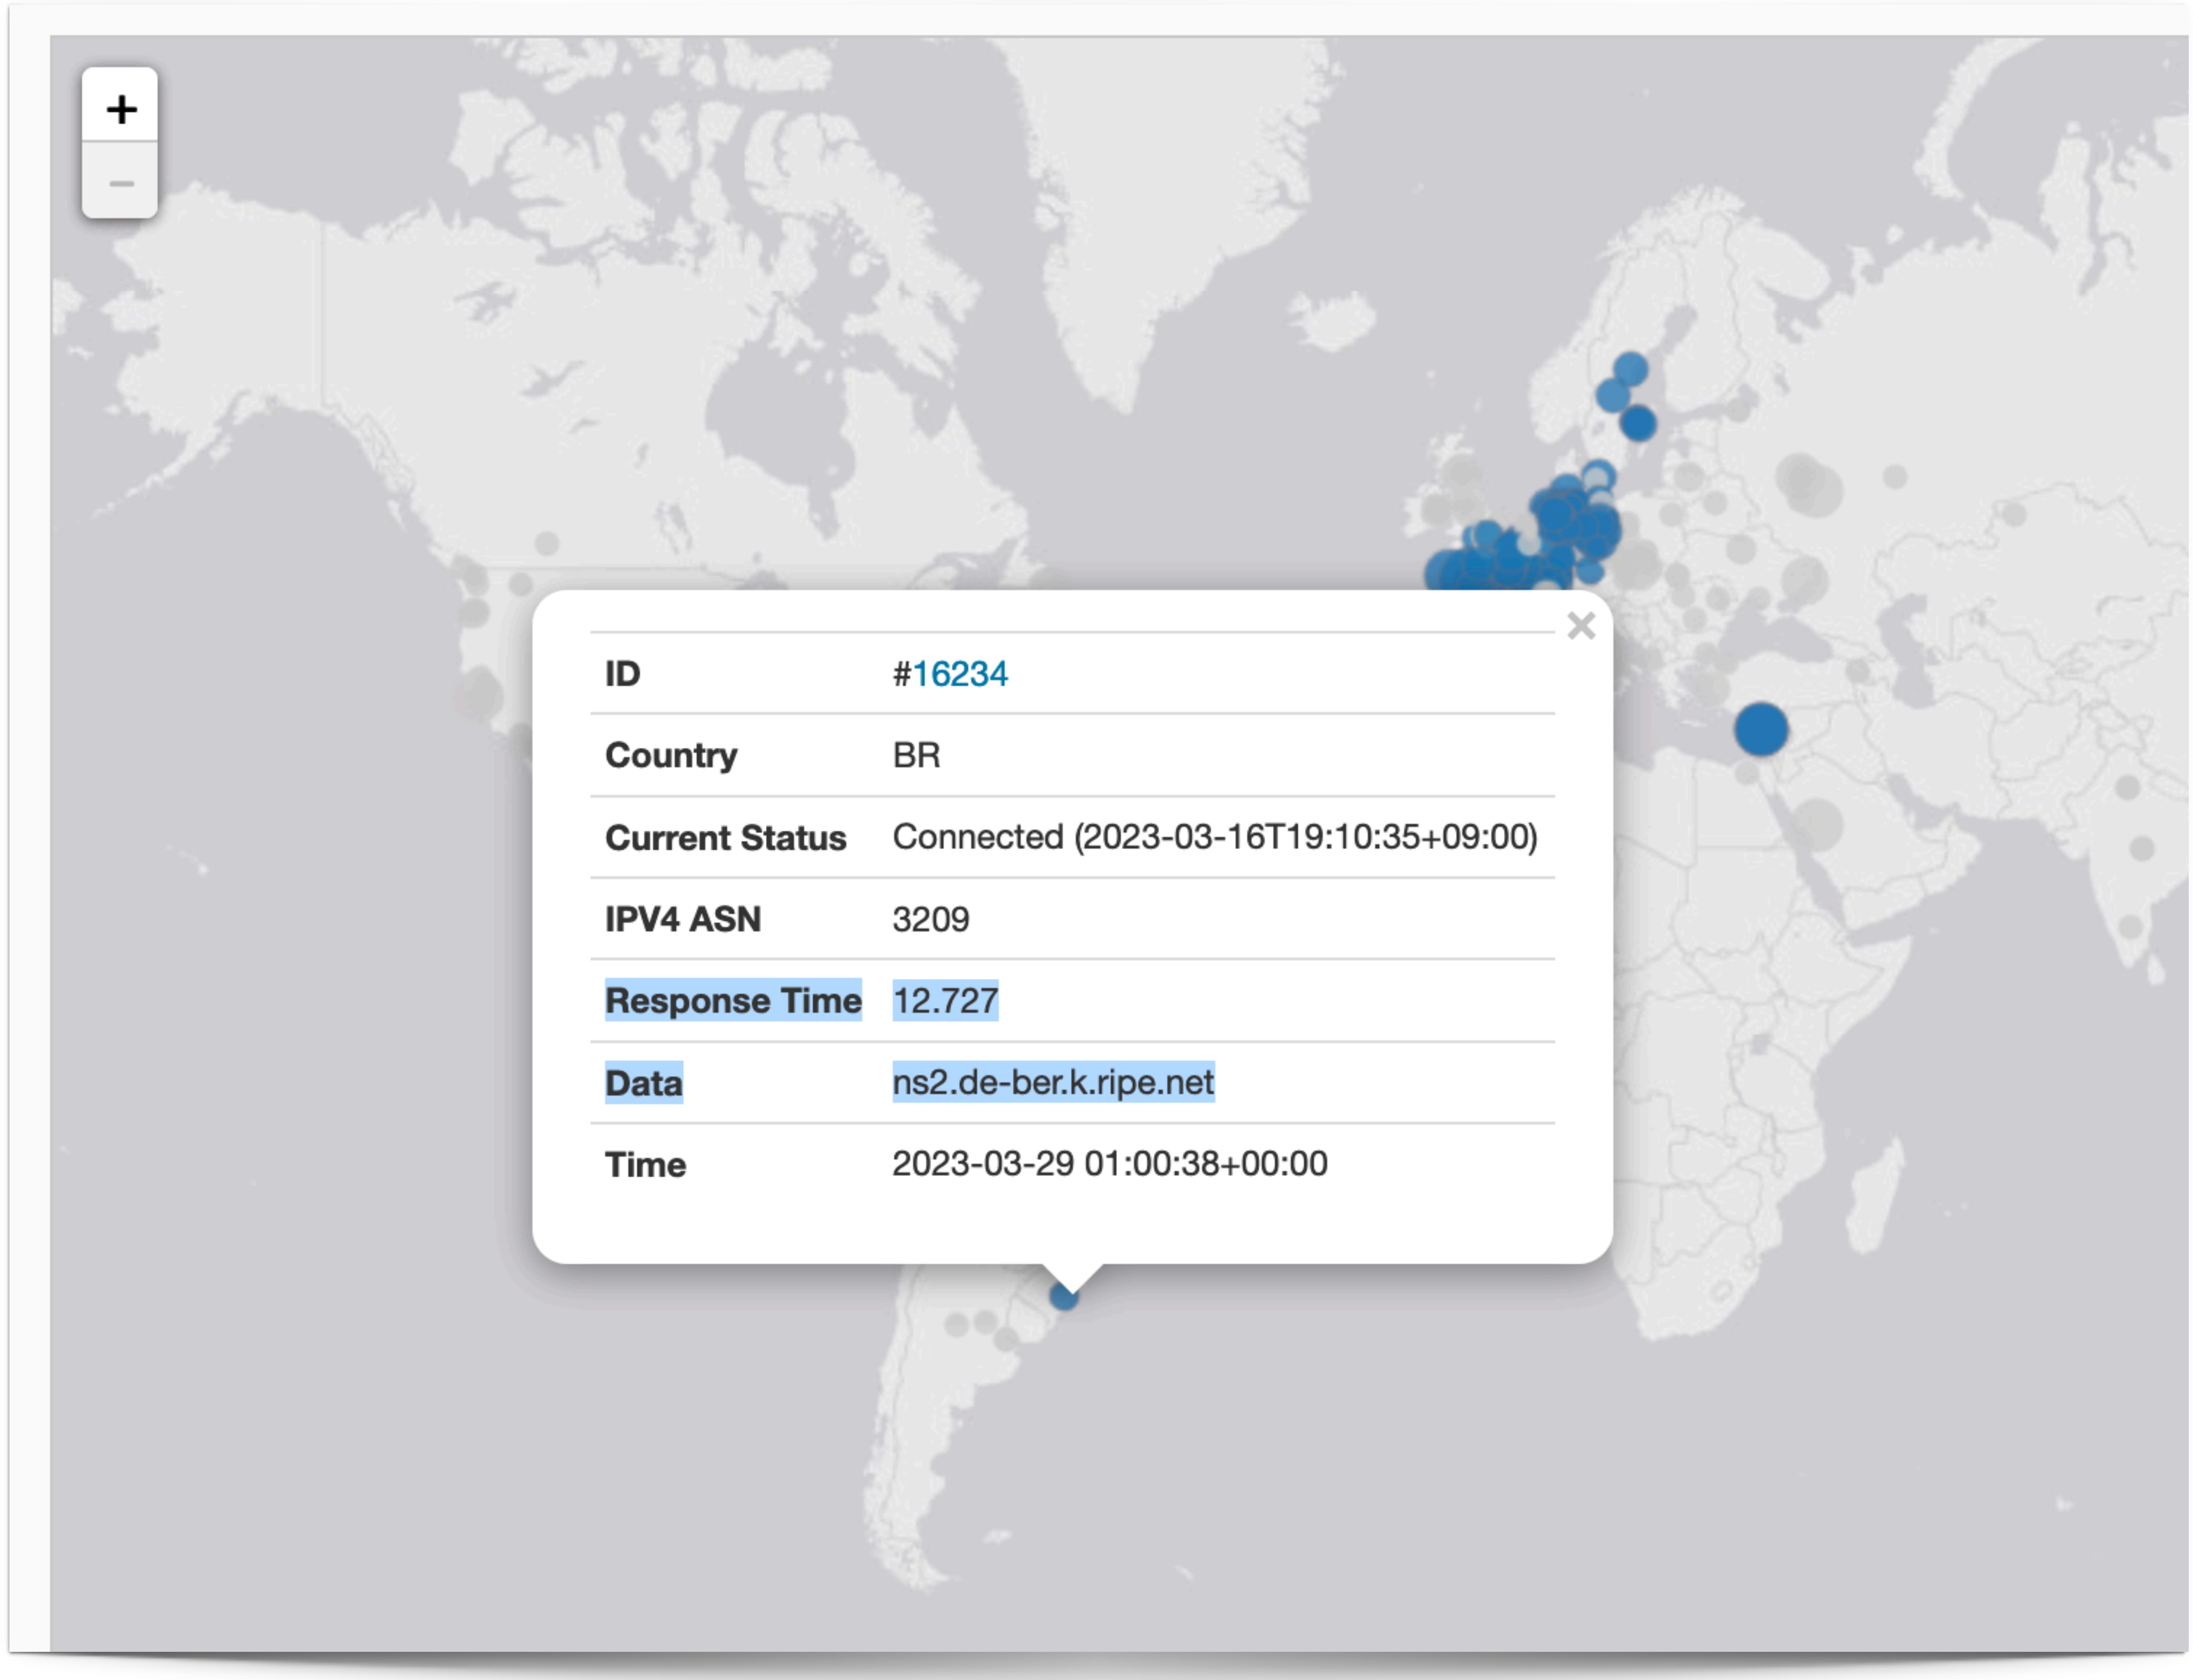

ID	#16234
Country	BR
Current Status	Connected (2023-03-16T19:10:35+09:00)
IPV4 ASN	3209
Response Time	12.727
Data	ns2.de-ber.k.ripe.net
Time	2023-03-29 01:00:38+00:00

Faster-Than-Light (FTL) packets

- Similar in spirit to the previous issue
- The geolocation is correct, yet results are too-good-to-be-true

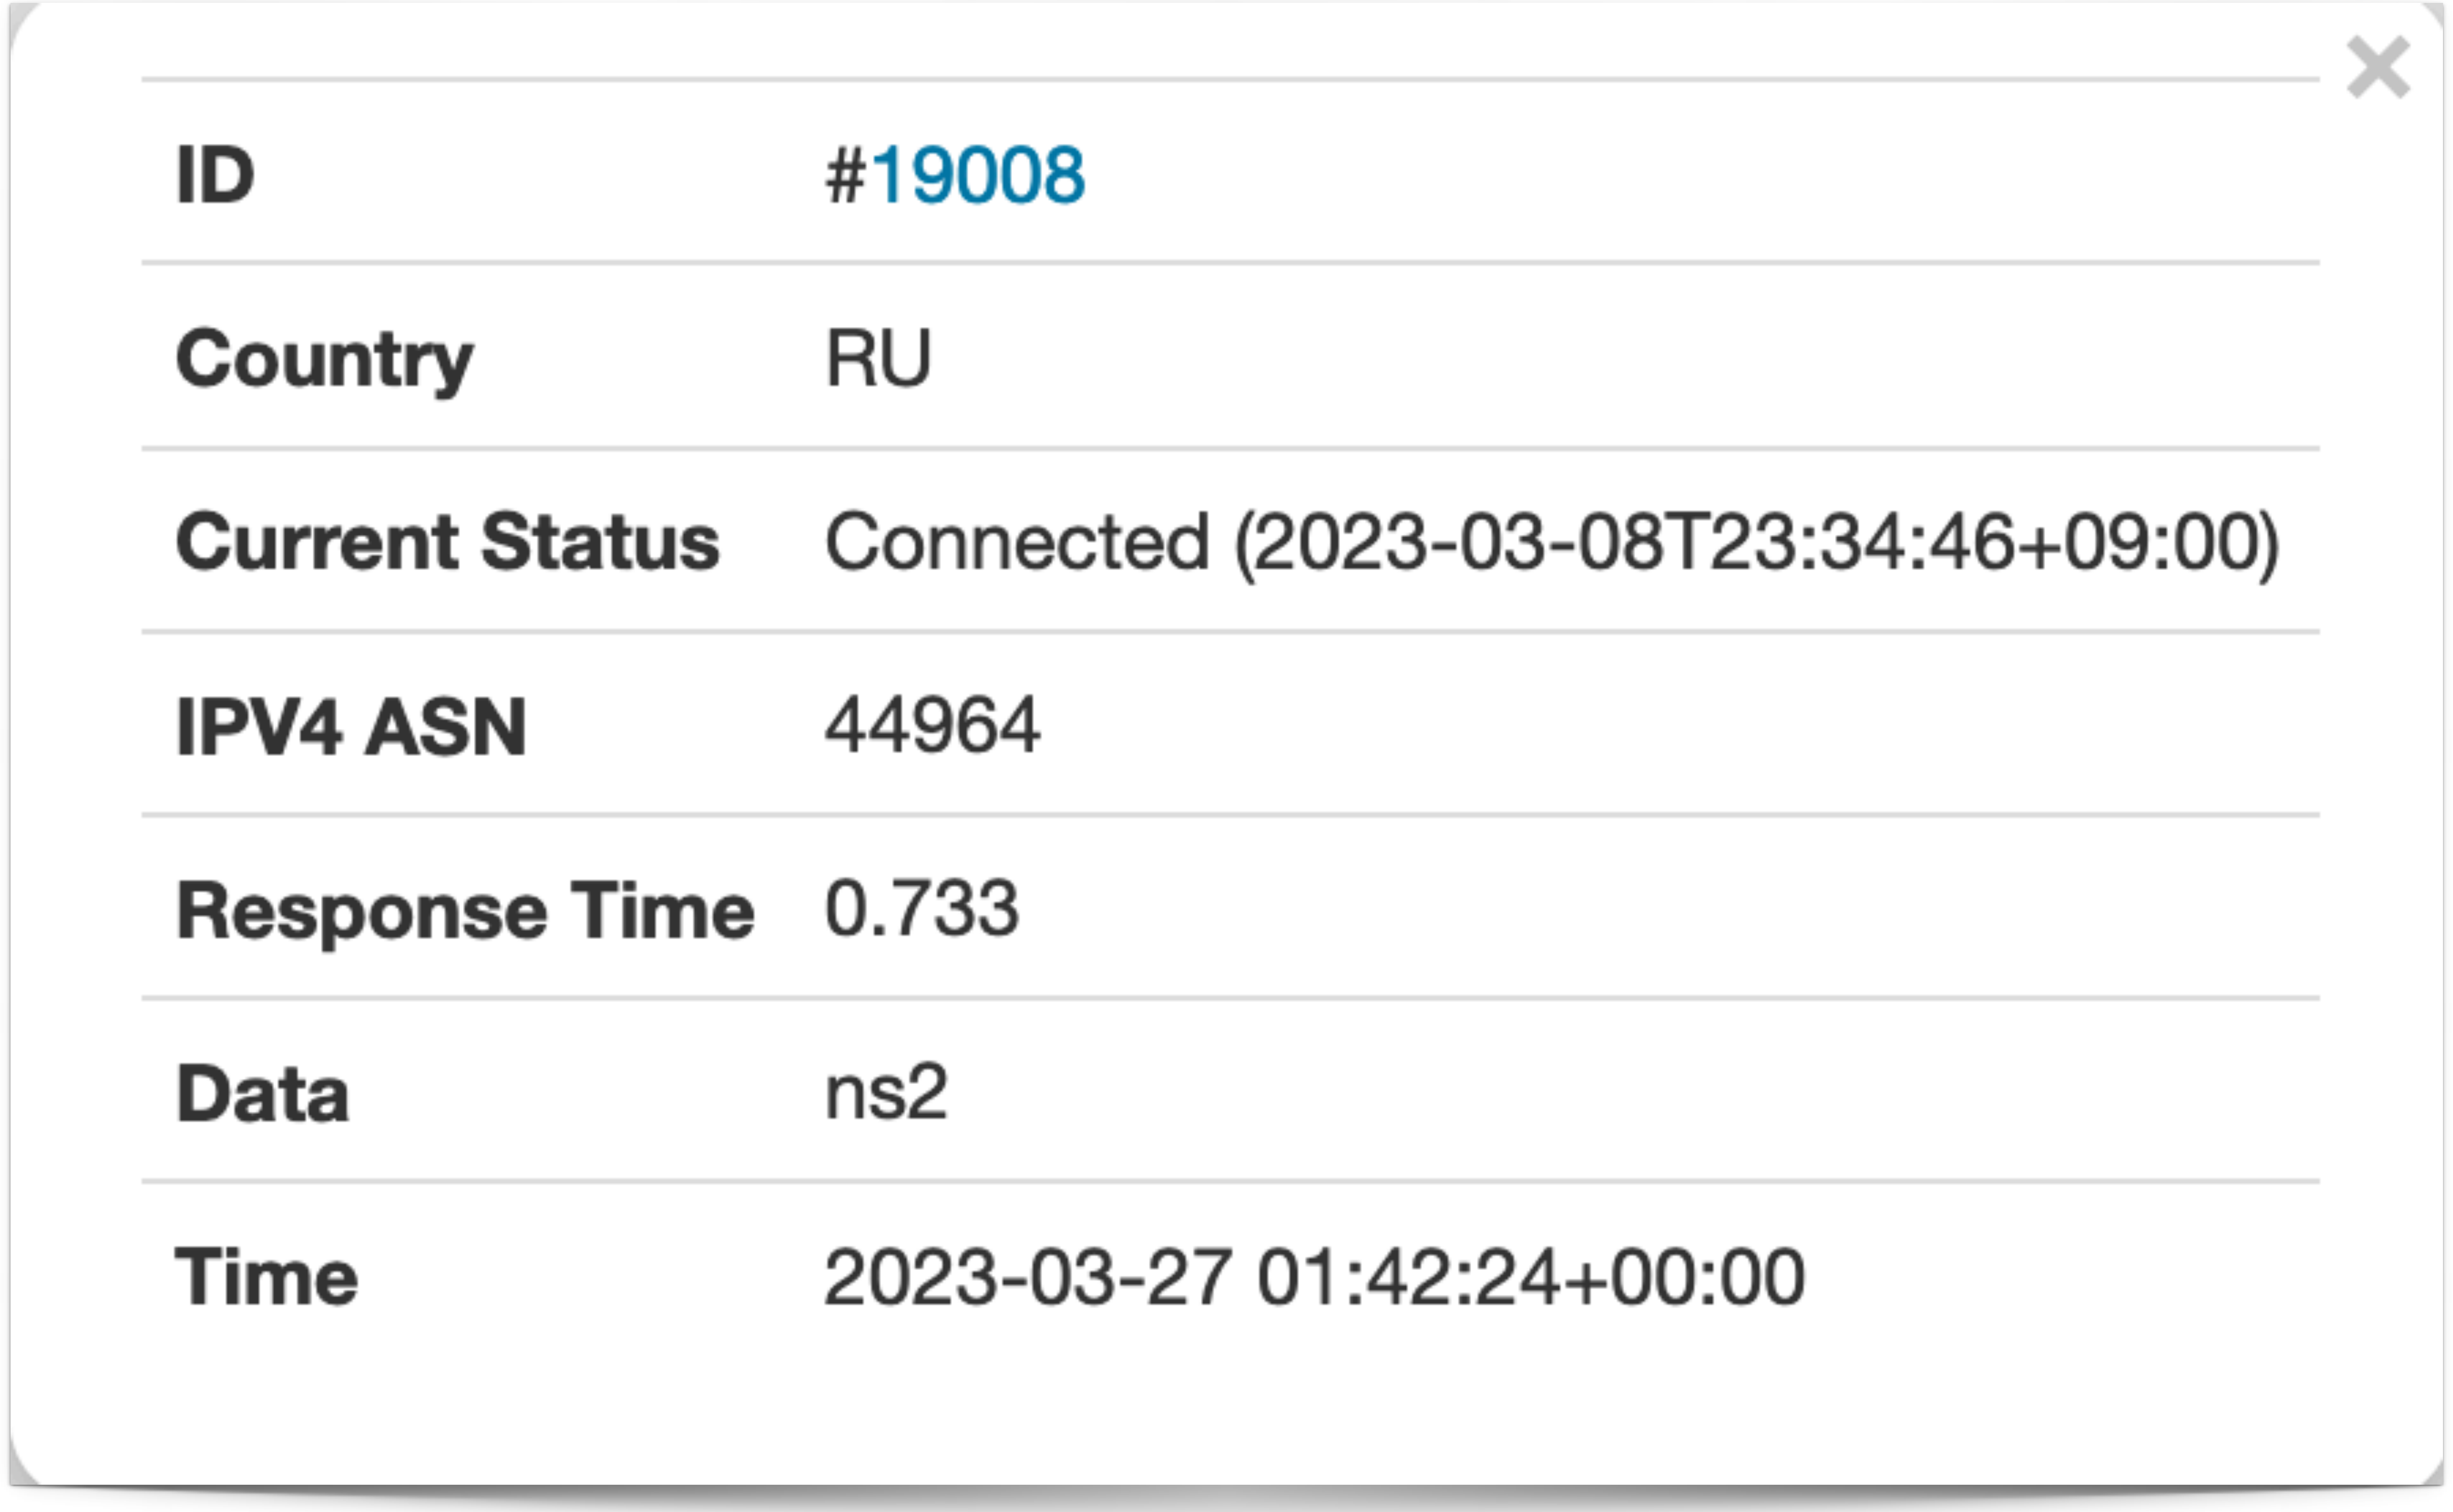

ID	#19008
Country	RU
Current Status	Connected (2023-03-08T23:34:46+09:00)
IPV4 ASN	44964
Response Time	0.733
Data	ns2
Time	2023-03-27 01:42:24+00:00

Conclusions

None yet 🙂