Analysis of probably redundant more specific announcements in BGP

Julien Gilon
University of Liège - CAIDA

Supervisors:
Benoit Donnet (University of Liège)
Matthew Luckie (UCSD/CAIDA)

March 2015

Plan

- The memory problem on routers
- What is a more specific? Why to use it?
- Methodology to infer borders
- First results
- What’s next?

In April 2014, some routing tables were about to reach the 500K entries. This alarms the scientific community which thought about the 512K limit on some routers…

> “The remainder of the prefixes (45%) shares the same origin AS and the same path. <…> I could make a wild guess and call these 45% of more specifics to be an act of senseless routing vandalism.”
>
> – Geoff Huston (APNIC)

Is this statement true?
Are the more specifics used for traffic engineering only?

Example

```
|203.181.248.168|7660|199.9.180.0/24|7660 22388 668 518|
|203.181.248.168|7660|199.9.181.0/24|7660 22388 668 518|
```

↓

```
|203.181.248.168|7660|199.9.180.0/23|7660 22388 668 518|
```

Source: route-view2 - 3/26/2015 19:00 UTC

Why?

Traffic engineering

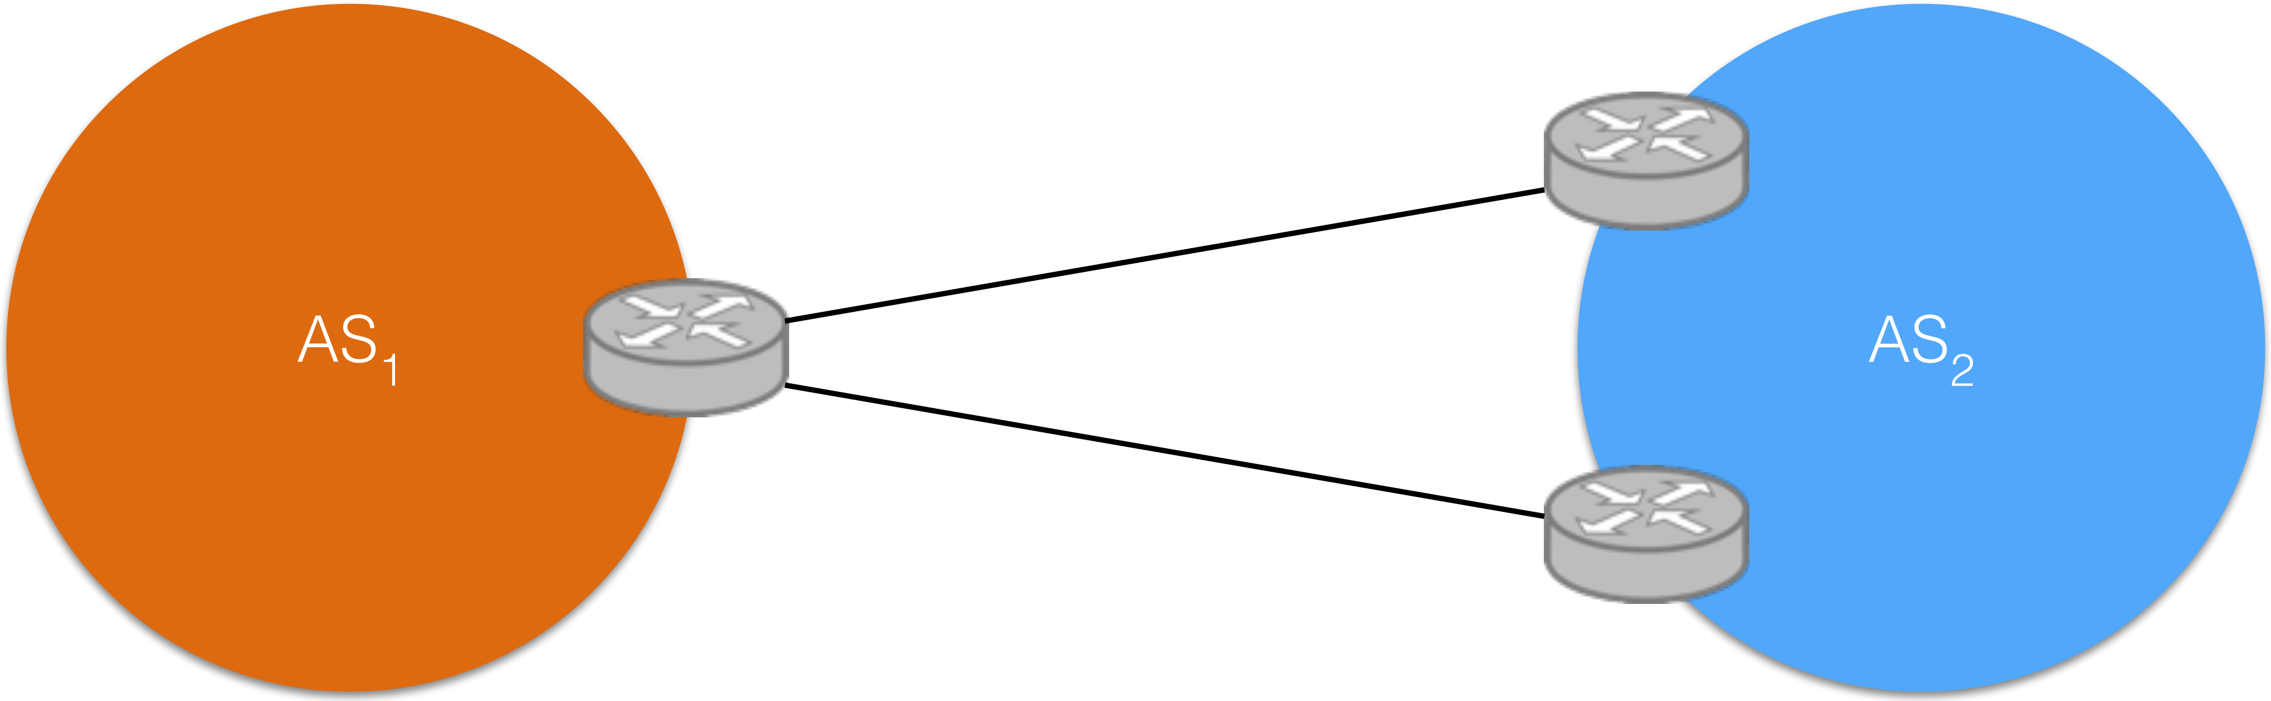

Why?

Security

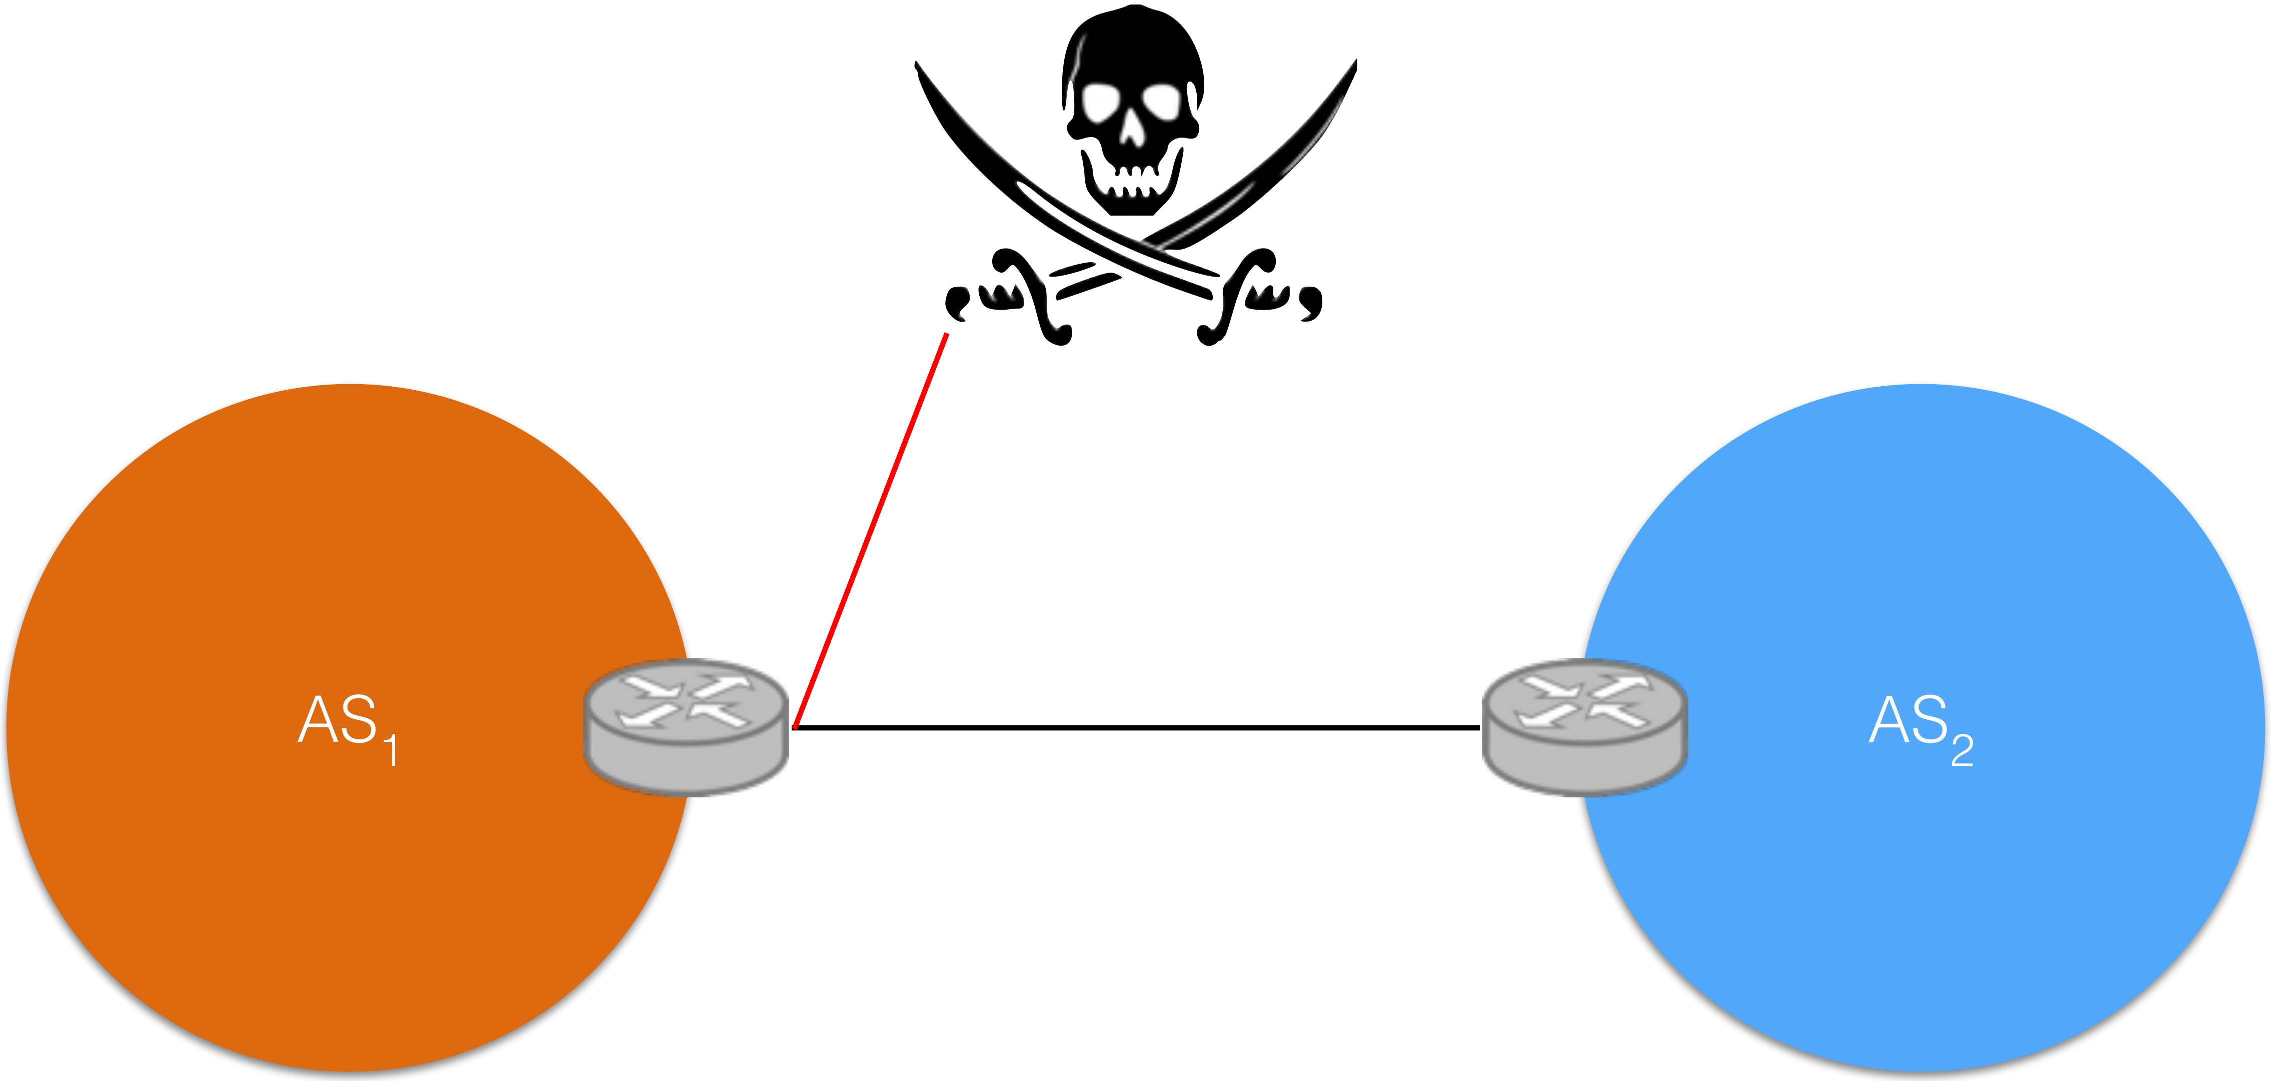

It is all about topology !

Methodology

- Find a BGP view can analyse with an ark monitor
- Find more specifics and aggregate them in group
- Launch Traceroute to reach announced prefixes
- Analyse and compare obtained traces

Find a BGP view

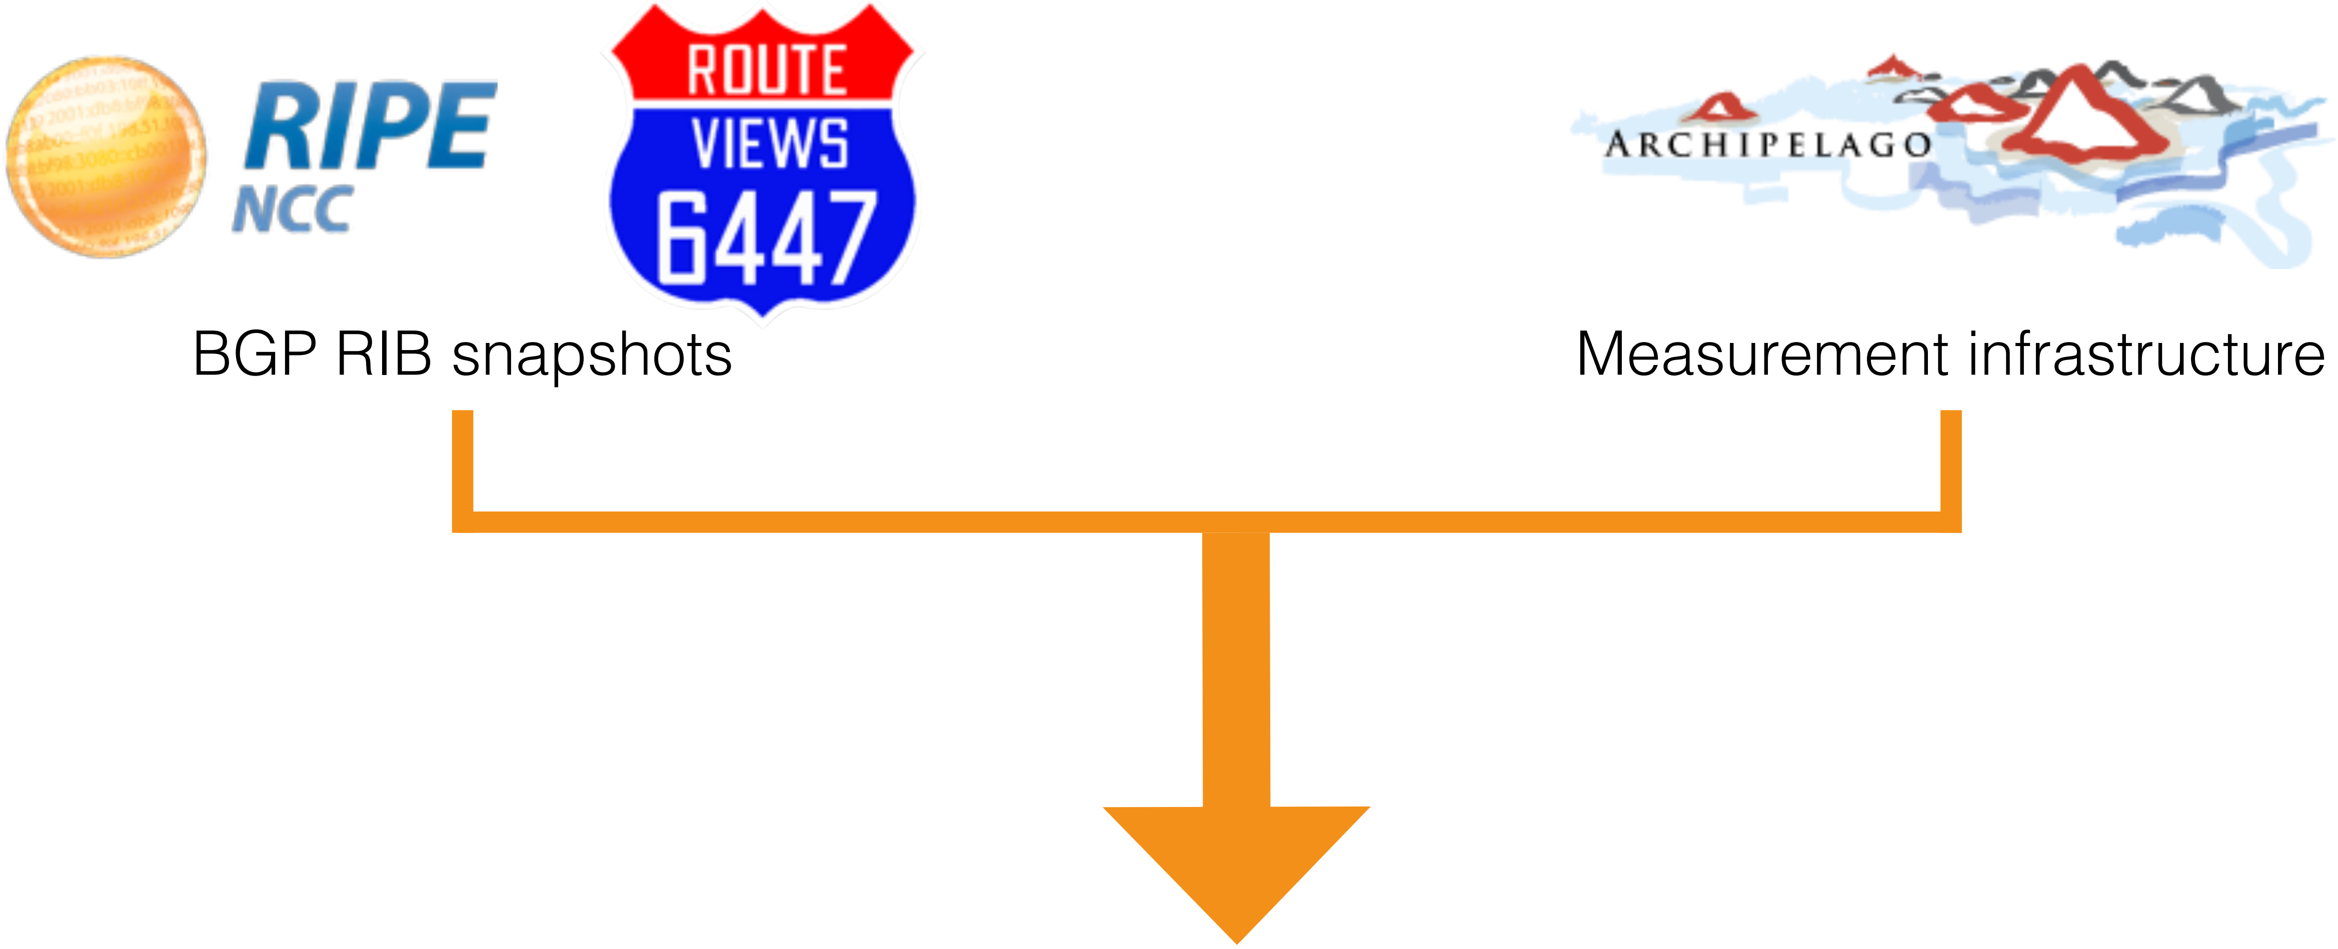

BGP RIB snapshots

Measurement infrastructure

Find entries coming from an AS where an Ark monitor is present!

Inferring the border

Provider-to-customer AS relationship

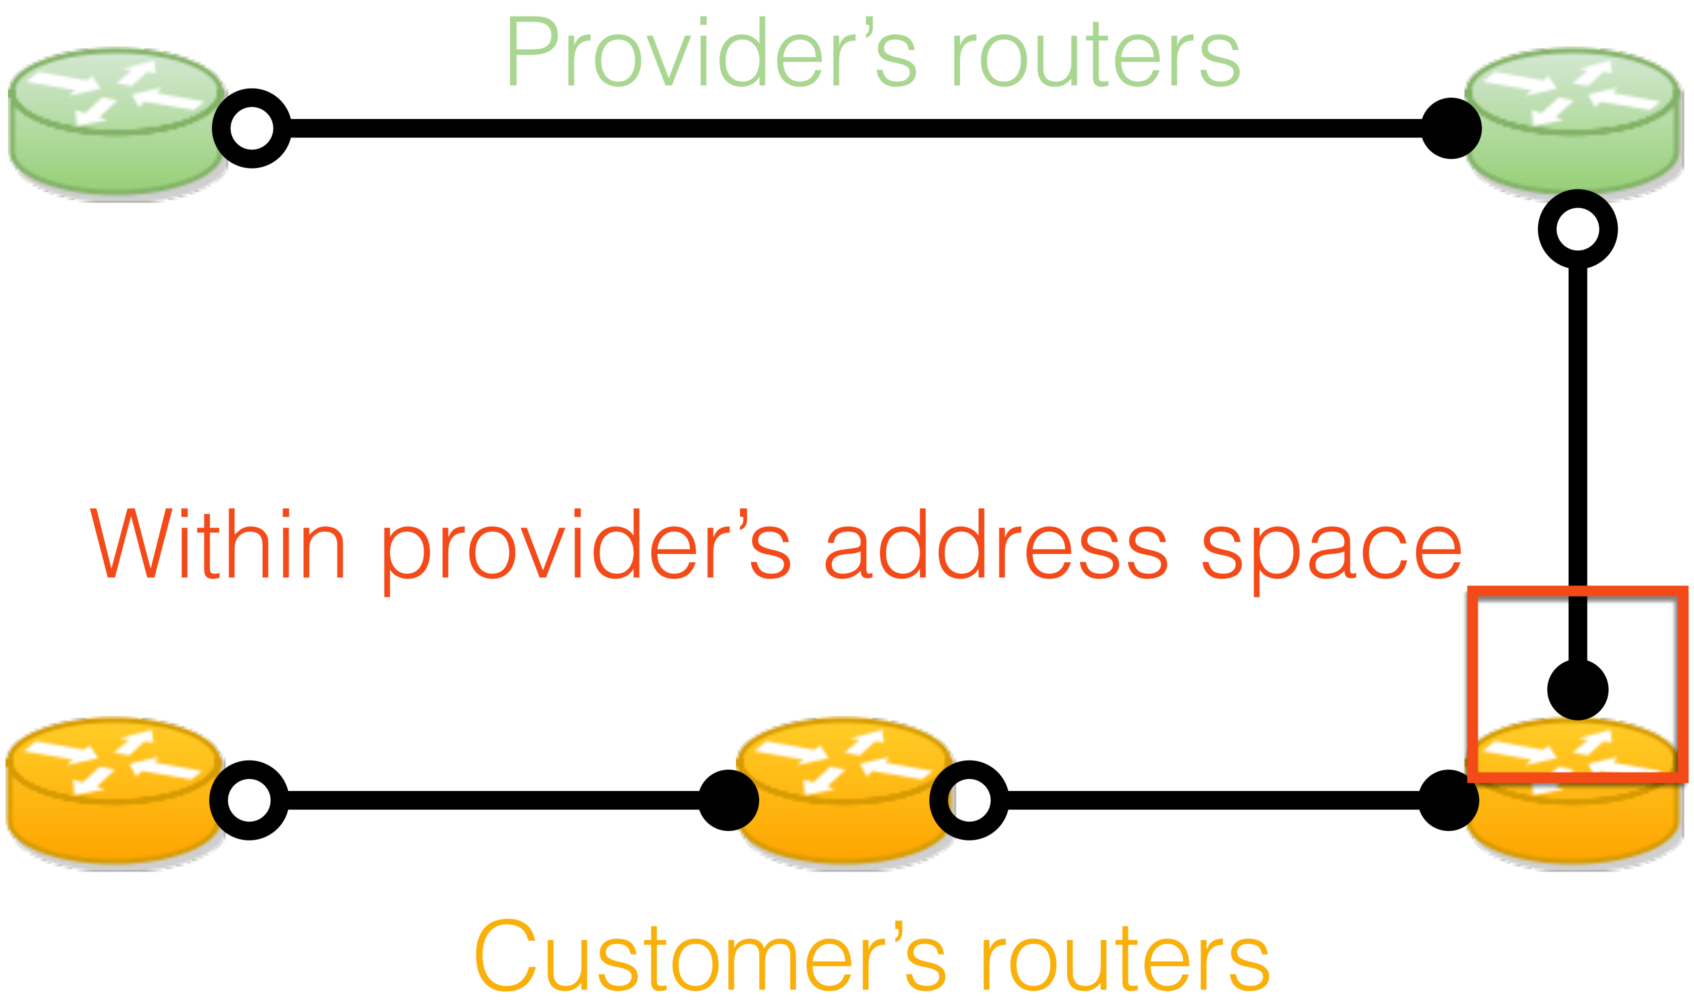

Provider's routers

Within provider's address space

Customer's routers

● : interfaces visible in traces | ○ : interfaces not visible in traces

Analysable border cases

One Traceroute that reaches the customer address space

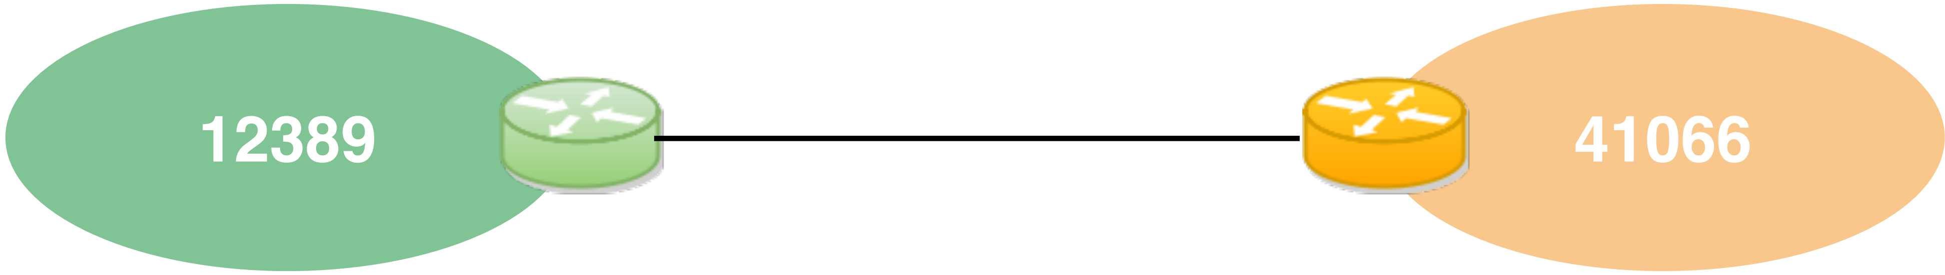

```
Group id: 23147 - 95.172.158.0/23 - ASPath: 2516|7660|-1 2516|12389|0 12389|41066|-1
```

```
HOP 1   ICMP<11,0>  203.181.248.60    7660
HOP 2   No result
HOP 3   ICMP<11,0>  203.181.102.129   2516
HOP 4   ICMP<11,0>  118.155.197.129   2516
HOP 5   ICMP<11,0>  106.187.6.158     2516
HOP 6   ICMP<11,0>  188.254.55.225    12389
HOP 7   ICMP<11,0>  95.167.91.146     12389
HOP 8   ICMP<11,0>  95.167.93.10      12389
HOP 9   ICMP<11,0>  188.254.20.182    12389
HOP 10  ICMP<11,0>  95.172.144.18      41066
```

```
HOP 1   ICMP<11,0>  203.181.248.60    7660
HOP 2   No result
HOP 3   ICMP<11,0>  203.181.102.129   2516
HOP 4   ICMP<11,0>  118.155.197.1     2516
HOP 5   ICMP<11,0>  106.187.6.154     2516
HOP 6   ICMP<11,0>  188.254.55.225    12389
HOP 7   ICMP<11,0>  95.167.91.146     12389
HOP 8   ICMP<11,0>  95.167.93.75      12389
HOP 9   ICMP<11,0>  217.65.88.22      12389
```

When the customer address space is reached

Analysable border cases

Multiple border interfaces which are not inferred to be aliases

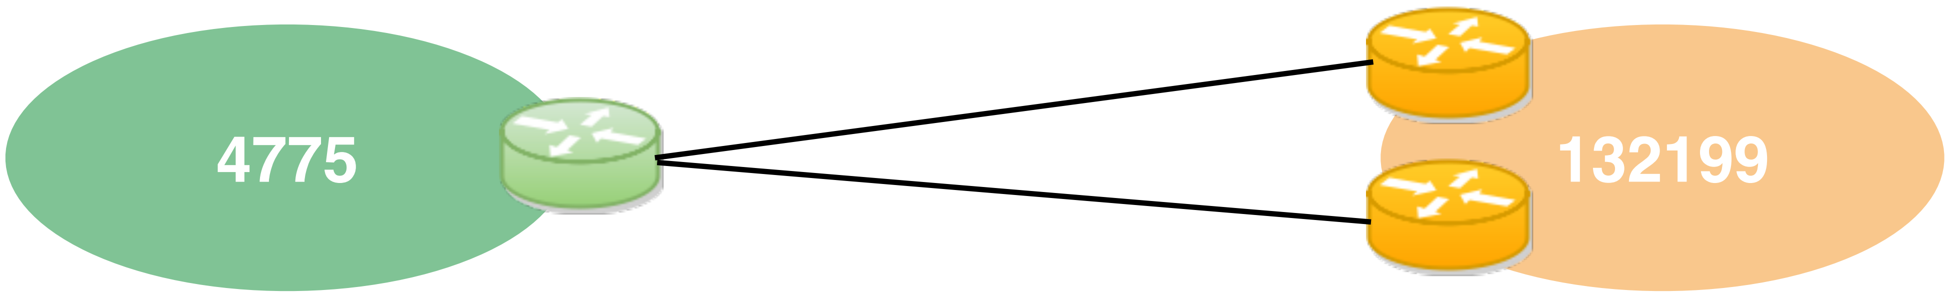

Group id: 16407 – 180.191.96.0/21 – ASPath: 2516|7660|-1 2516|4775|-1 4775|132199|-1

```
HOP 1  ICMP<11,0>  203.181.248.60   7660
HOP 2  No result
HOP 3  ICMP<11,0>  203.181.102.129  2516
HOP 4  ICMP<11,0>  118.155.197.142  2516
HOP 5  ICMP<11,0>  111.87.10.6      2516
HOP 6  ICMP<11,0>  120.28.10.77     4775
HOP 7  ICMP<11,0>  120.28.10.153    4775
HOP 8  ICMP<11,0>  222.127.123.238  132199
```

```
HOP 1  ICMP<11,0> 203.181.248.60   7660
HOP 2  No result
HOP 3  ICMP<11,0> 203.181.102.129  2516
HOP 4  ICMP<11,0> 118.155.197.14   2516
HOP 5  ICMP<11,0> 111.87.10.6      2516
HOP 6  ICMP<11,0> 120.28.0.101     4775
HOP 7  ICMP<11,0> 120.28.10.49     4775
HOP 8  ICMP<11,0> 120.28.10.6      4775
HOP 9  ICMP<11,0> 222.127.123.246  132199
```

When the customer address space is reached

Analysable border cases

Same border interface observed multiple times

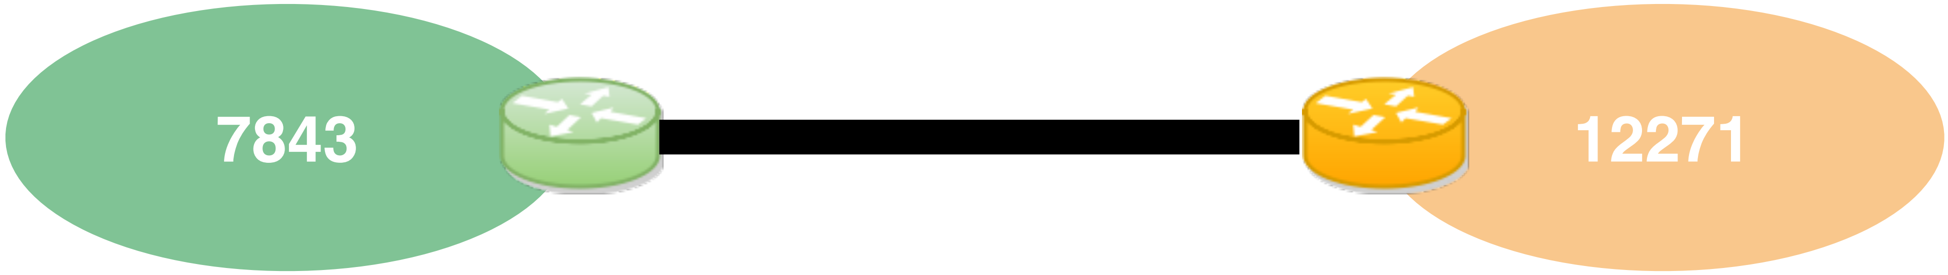

```
Group id: 21713 – 24.215.128.0/17 – ASPath: 2516|7660|-1 2516|7843|0 7843|12271|-1
 HOP 1  ICMP<11,0>  203.181.248.60   7660       HOP 1  ICMP<11,0>  203.181.248.60   7660
 HOP 2  No result                               HOP 2  No result
 HOP 3  ICMP<11,0>  203.181.102.129  2516       HOP 3  ICMP<11,0>  203.181.102.129  2516
 HOP 4  ICMP<11,0>  118.155.197.130  2516       HOP 4  ICMP<11,0>  118.155.197.130  2516
 HOP 5  ICMP<11,0>  203.181.100.154  2516       HOP 5  ICMP<11,0>  203.181.100.178  2516
 HOP 6  ICMP<11,0>  111.87.3.30      2516       HOP 6  ICMP<11,0>  124.211.34.134   2516
 HOP 7  ICMP<11,0>  107.14.16.113    7843       HOP 7  ICMP<11,0>  107.14.16.113    7843
 HOP 8  ICMP<11,0>  66.109.6.138     7843       HOP 8  ICMP<11,0>  66.109.6.138     7843
 HOP 9  ICMP<11,0>  66.109.6.15      7843       HOP 9  ICMP<11,0>  66.109.6.15      7843
 HOP 10 ICMP<11,0>  107.14.19.35     7843       HOP 10 ICMP<11,0>  66.109.6.24      7843
 HOP 11 ICMP<11,0>  107.14.17.172    7843       HOP 11 ICMP<11,0>  107.14.17.172    7843
 HOP 12 ICMP<11,0>  107.14.19.25     7843       HOP 12 ICMP<11,0>  107.14.19.25     7843
 HOP 13 ICMP<11,0>  68.173.198.25    12271      HOP 13 ICMP<11,0>  184.152.112.106  12271
                                                …
```

When the customer address space is reached

Analysable border cases

One border router observed multiple times over multiple interfaces (aliases)

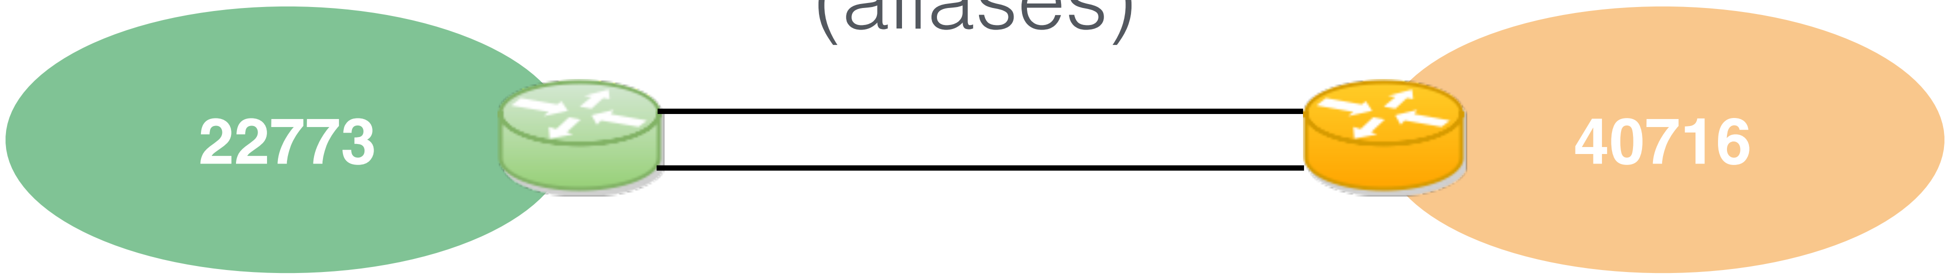

```
Group id: 24491 - 192.234.160.0/23 - ASPath: 2516|7660|-1 2516|22773|0 22773|40716|-1
```

```
HOP 1  ICMP<11,0>  203.181.248.60   7660
HOP 2  No result
HOP 3  ICMP<11,0>  203.181.102.129 2516
HOP 4  ICMP<11,0>  118.155.197.2    2516
HOP 5  ICMP<11,0>  203.181.100.118 2516
HOP 6  ICMP<11,0>  59.128.2.74      2516
HOP 7  ICMP<11,0>  124.215.192.210 2516
HOP 8  ICMP<11,0>  68.1.0.125      22773
HOP 9  ICMP<11,0>  68.13.8.150     22773
HOP 10 ICMP<11,0>  66.37.238.138   22773
HOP 11 ICMP<0,0>   192.234.160.3   40716
```

```
HOP 1  ICMP<11,0>  203.181.248.60   7660
HOP 2  No result
HOP 3  ICMP<11,0>  203.181.102.129  2516
HOP 4  ICMP<11,0>  118.155.197.130  2516
HOP 5  ICMP<11,0>  203.181.100.46   2516
HOP 6  ICMP<11,0>  59.128.2.74       2516
HOP 7  ICMP<11,0>  124.215.192.210   2516
HOP 8  ICMP<11,0>  68.1.0.127        22773
HOP 9  ICMP<11,0>  68.13.8.154       22773
HOP 10 ICMP<11,0>  66.37.238.198     22773
HOP 11 ICMP<0,0>   192.234.161.3     40716
```

When the customer address space is reached

Not analysable border cases

The provider address space is not observed in the Traceroute

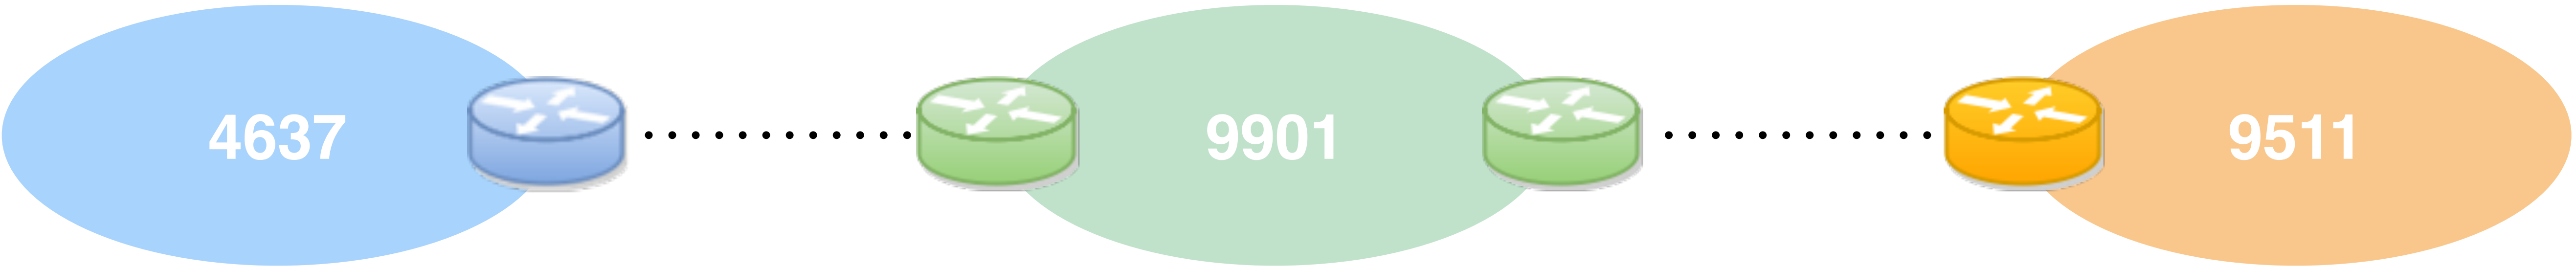

```
Group id: 43727 – 203.99.70.0/23 – ASPath: 2516|7660|-1 2516|4637|-1 4637|9901|-1 9901|9511|-1
HOP 1  ICMP<11,0>  203.181.248.60   7660
HOP 2  No result
HOP 3  ICMP<11,0>  203.181.102.129  2516
HOP 4  ICMP<11,0>  118.155.197.1    2516
HOP 5  ICMP<11,0>  106.187.6.138    2516
HOP 6  ICMP<11,0>  202.84.148.134   4637
HOP 7  ICMP<11,0>  202.84.148.230   4637
HOP 8  ICMP<11,0>  202.84.141.194   4637
HOP 9  ICMP<11,0>  202.84.220.189   4637
HOP 10 ICMP<11,0>  202.84.142.154   4637
HOP 11 ICMP<11,0>  134.159.174.38   4637
HOP 12 ICMP<11,0>  203.98.54.226    4768
HOP 13 ICMP<0,0>   203.99.71.29     9511
```

When the customer address space is reached

Not analysable border cases

We are missing hops where a border link could be

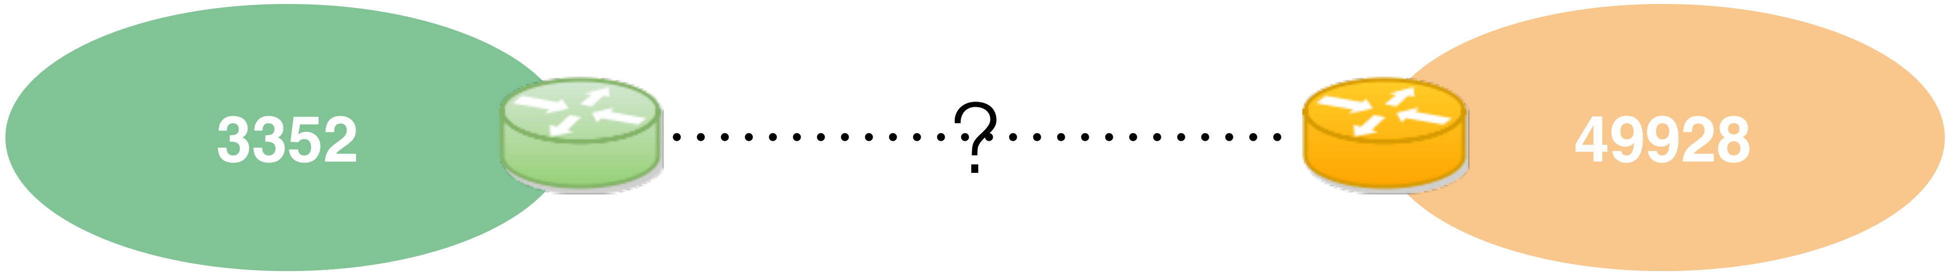

```
Group id: 51666– 213.99.54.0/23 – ASPath: 2516|7660|-1 12956|2516|-1 12956|3352|-1 3352|49928|-1
                           HOP 1  ICMP<11,0>  203.181.248.60   7660
                           HOP 2  No result
                           HOP 3  ICMP<11,0>  203.181.102.129  2516
                           HOP 4  ICMP<11,0>  118.155.197.130  2516
                           HOP 5  ICMP<11,0>  203.181.100.70   2516
                           HOP 6  ICMP<11,0>  124.211.34.130   2516
                           HOP 7  ICMP<11,0>  203.181.104.218  2516
                           HOP 8  ICMP<11,0>  213.140.49.6     12956
                           HOP 9  ICMP<11,0>  94.142.119.190   12956
                           HOP 10 ICMP<11,0>  213.140.37.45    12956
                           HOP 11 ICMP<11,0>  216.184.113.111  12956
                           HOP 12 ICMP<11,0>  213.0.190.38     3352
                           HOP 13 ICMP<11,0>  193.152.56.66    3352
                           HOP 14 No result
                           HOP 15 No result
                           HOP 16 ICMP<0,0>    213.99.54.55     49928
```

When the customer address space is reached

Not analysable border cases

The customer and the provider are under the same management

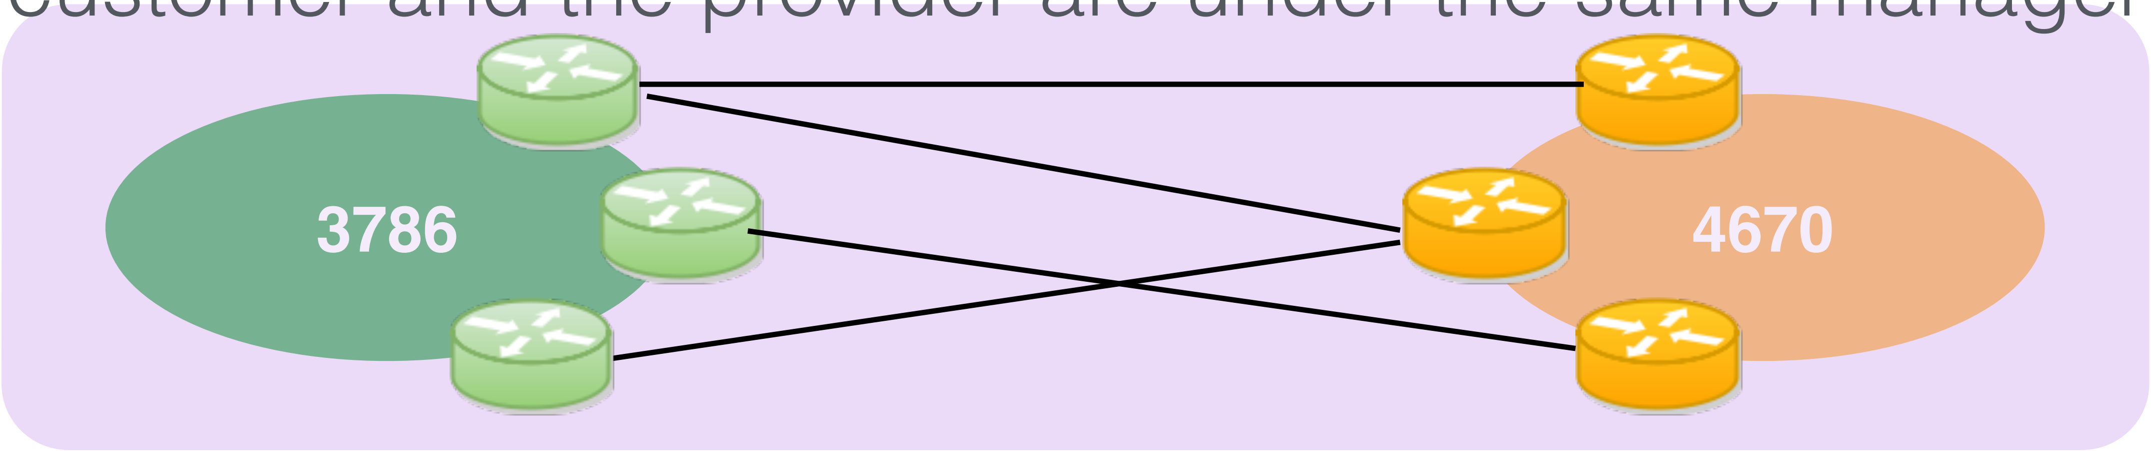

Group id: 13983 – 58.180.0.0/16 ASPath: 2516|7660|-1 2516|3786|0 3786|4670|-1

```
…
HOP 7  ICMP<11,0>  61.42.202.129    3786
HOP 8  ICMP<11,0>  1.208.107.218    3786
HOP 9  ICMP<11,0>  211.116.60.126   3786
HOP 10 ICMP<11,0>  202.30.147.209   4670
HOP 11 No result
HOP 12 No result
HOP 13 ICMP<11,0>  1.208.9.197      3786
HOP 14 ICMP<11,0>  1.208.9.150      3786
HOP 15 No result
HOP 16 No result
HOP 17 ICMP<0,0>  58.180.247.254    4670
```

When the customer address space is reached

Not analysable border cases

De-alias process doesn’t return an answer

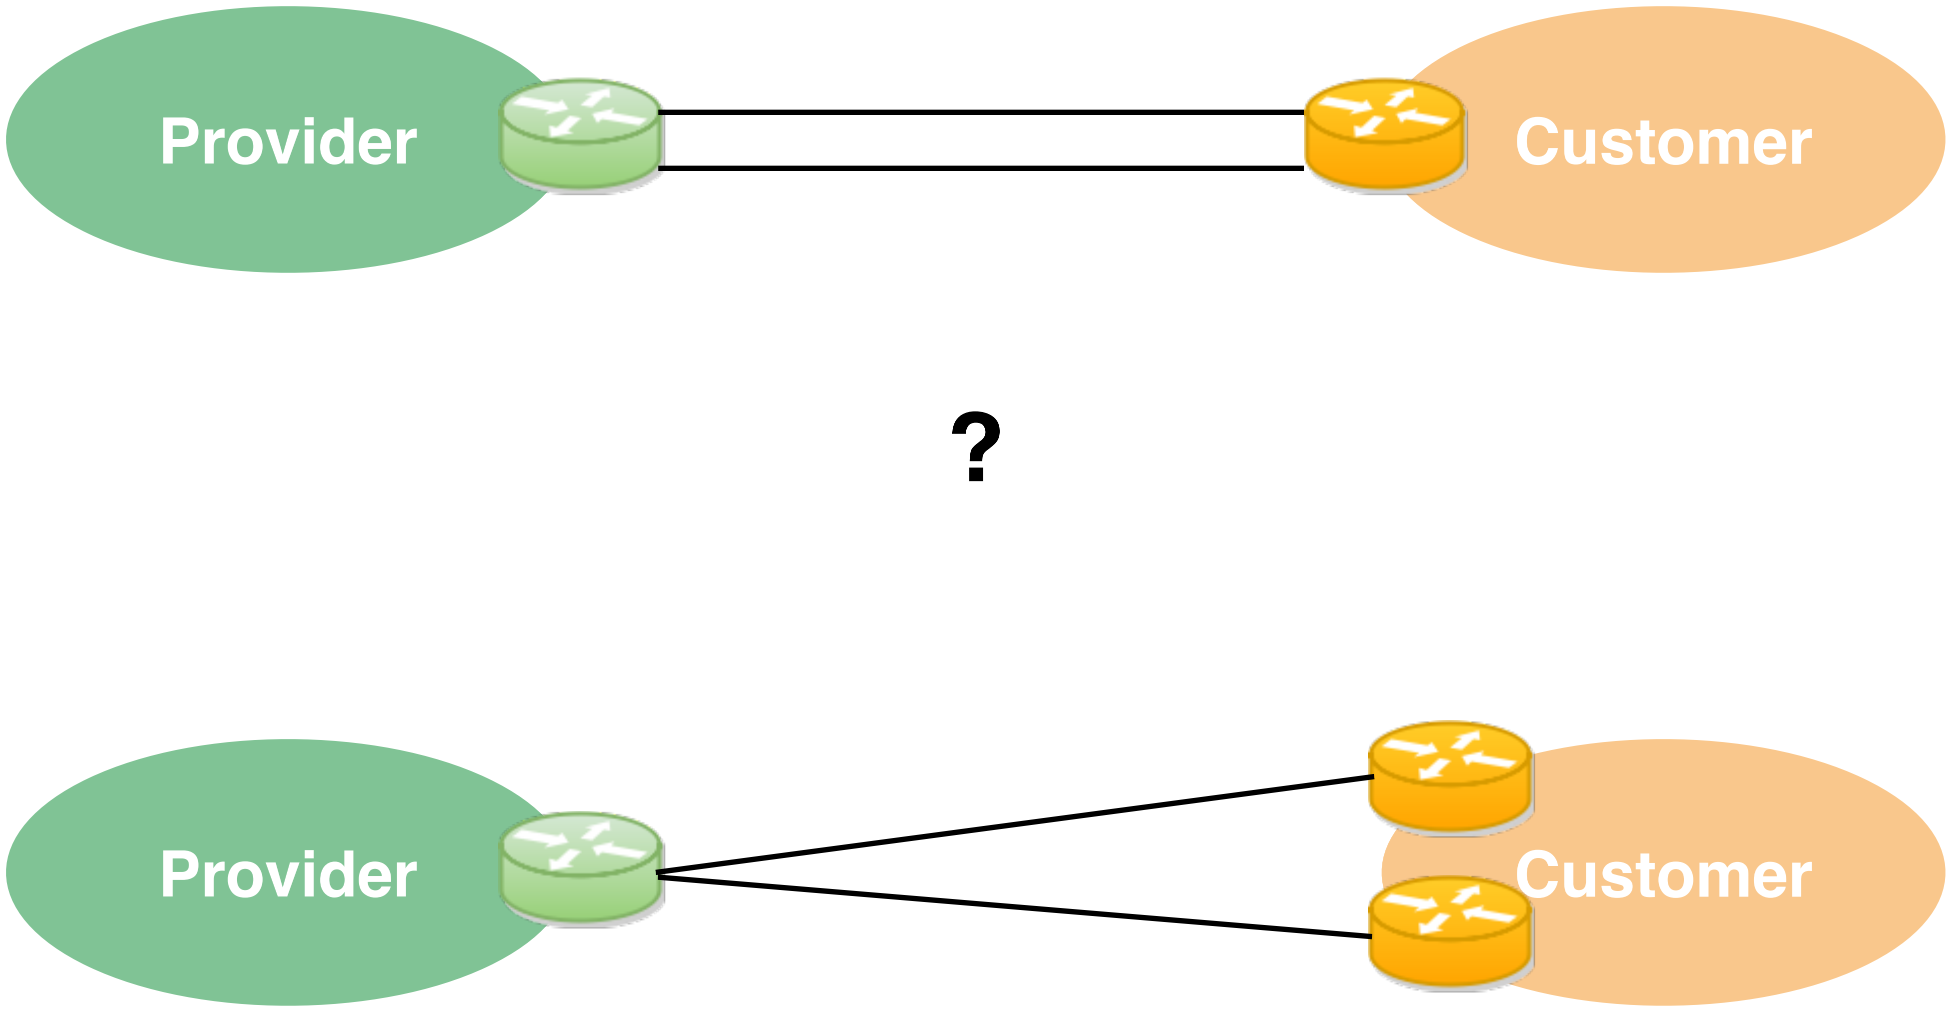

When the customer address space is reached

When the customer address space is not reached

Group’s traces stop :

- Without observing the provider address space
- At a random interface
- All at the same interface

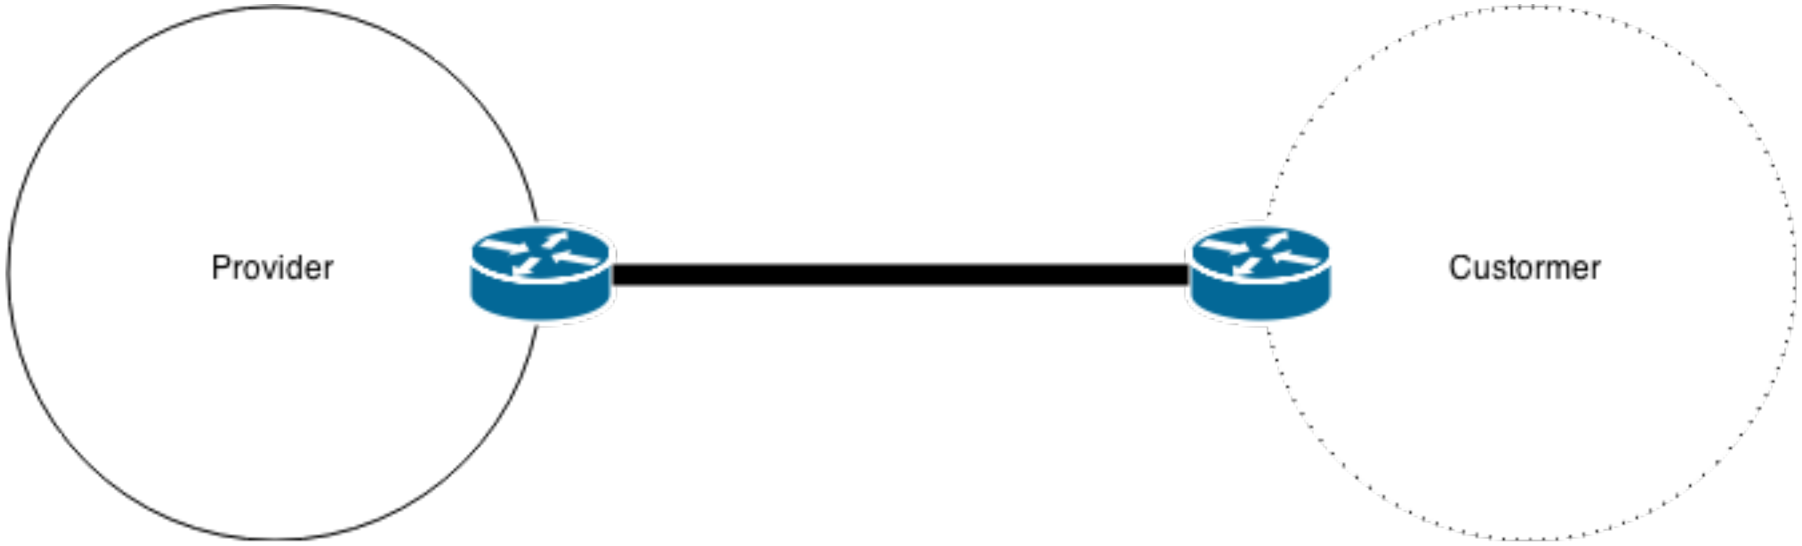

Statistics

Sample size: 5000 | Number of prefixes/group: 2

Not analysable border cases

1777 groups => 35%

Sample size: 5000 | Number of prefixes/group: 2

Analysable border cases

2030 groups => 40%

One Traceroute that reaches the customer address space:

617 groups => 30%

Multiple border interfaces which are not inferred to be aliases:

177 groups => 9%

One border router observed multiple times over multiple interfaces (aliases)

85 groups => 4%

Same border interface observed multiple times

1151 groups => 57%

Sample size: 5000 | Number of prefixes/group: 2

Customer address space not reached

1214 groups => 25%

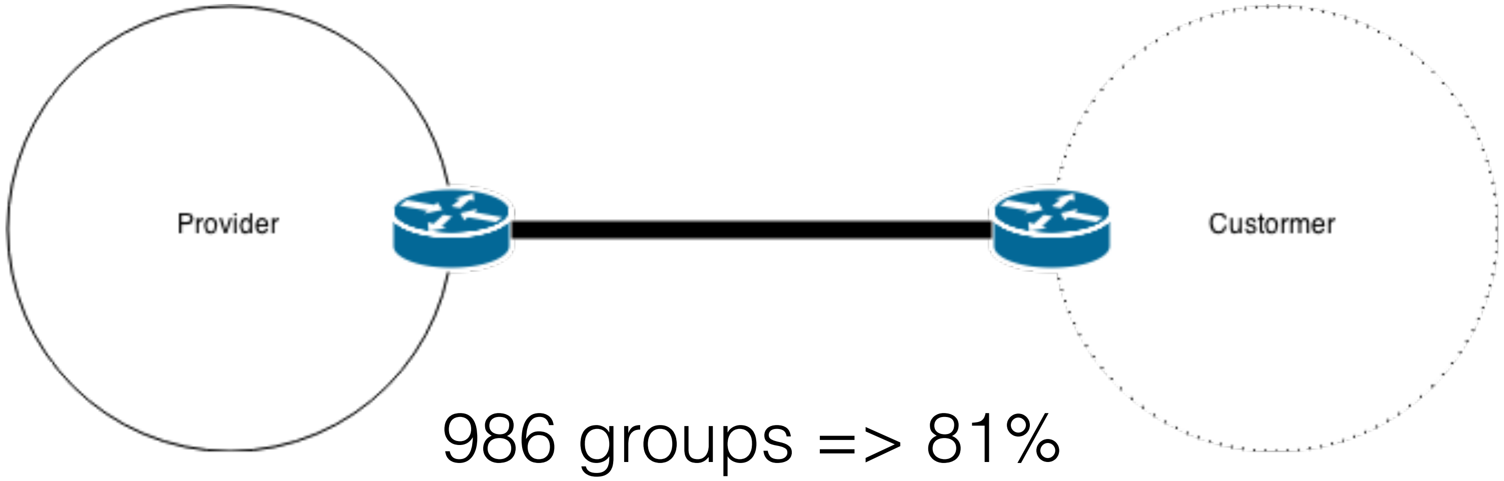

986 groups => 81%

Others: 228 groups => 19%

Sample size: 5000 | Number of prefixes/group: 2

What’s next?

- Analyse more samples with more prefixes/group
- Link the more specifics to real life problems (i.e. the 512K problem)

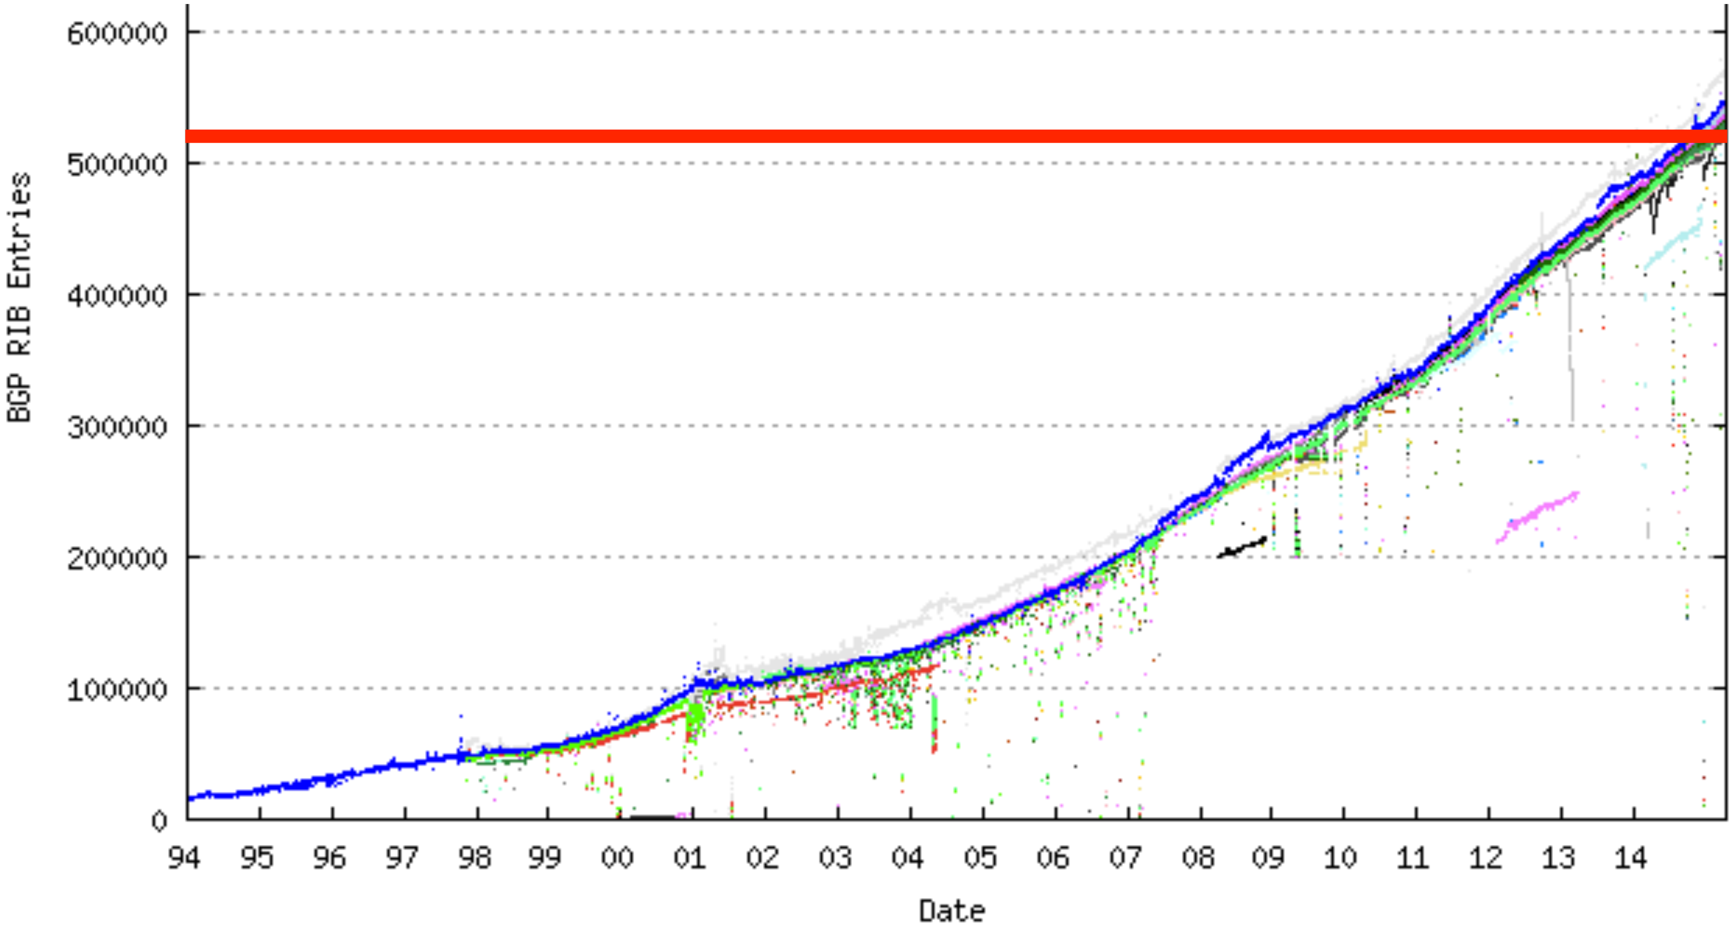

Source: bgp.potaroo.net/

Thank you