



RIPE NCC

RIPE NETWORK COORDINATION CENTRE

RIPE IPmap Geolocating Internet Infra-Structure with Inference Engines and Crowdsourcing

Victor Naumov
R&D
RIPE NCC

vnaumov@ripe.net ENOG15 | 2018



**But now for something
completely different first**

User to User Still Important



- The quality of end-user connections are often expressed and optimised in download speeds towards content providers.
- Instead, these sketches focus on peer-to-peer connections within a country.
- Fabric snapshots at points in time
- Trying to estimate an amount of ways the networks interconnect their users
 - Any “single point of failure”?
 - How much IXPs and Transit Providers are involved?

We call it sketches since it uses rough estimates data from many source

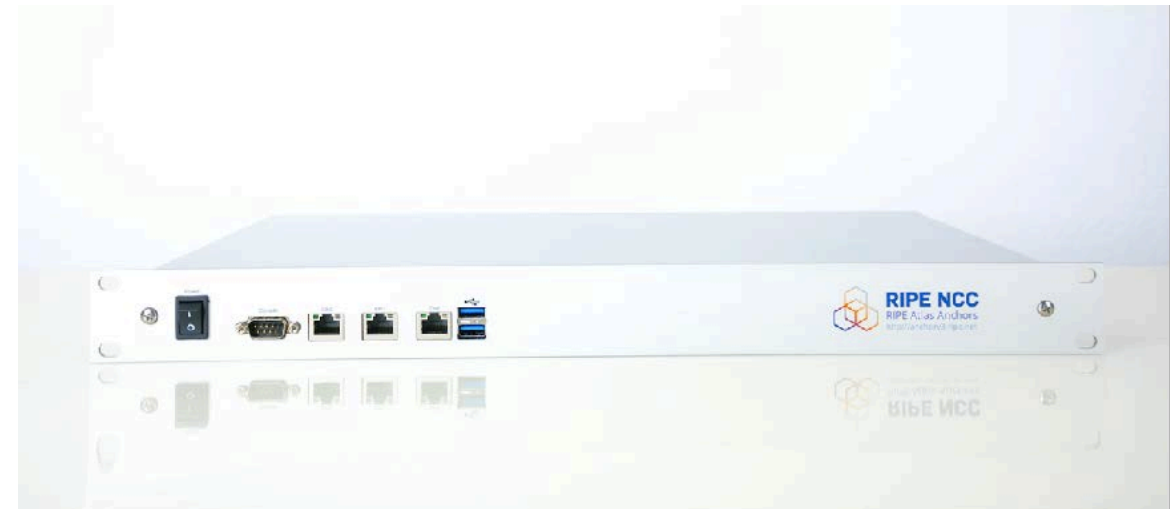
Data Sources: RIPE Atlas



- Measure the Internet!
 - By community, for community
- Wikipedia: https://en.wikipedia.org/wiki/RIPE_Atlas



Data Sources: RIPE Atlas



10,500+ probes connected (310+ Anchors)

6000 results collected per second

40,000+ user-defined measurements weekly

Five types of user-defined measurements available to probe hosts and RIPE NCC members: ping, traceroute, DNS, SSL, NTP, HTTP, WiFi

Data sources



- APNIC end-users per network estimation
 - <https://stats.labs.apnic.net/aspop/>
- IXP-country-jedi
 - <https://www.ripe.net/ixp-country-jedi/>
 - mesh traceroutes between RIPE Atlas probes in a country
 - measures and analyses (using IPMap)

Data sources



- CAIDA AS-to-ORG
 - <https://www.caida.org/data/as-organizations/>
- PeeringDB
 - <https://www.peeringdb.com/>
- RIPE IPmap
 - <https://ipmap.ripe.net/>
- RIPEStat
 - <https://stat.ripe.net/>



Show me the Sketch!

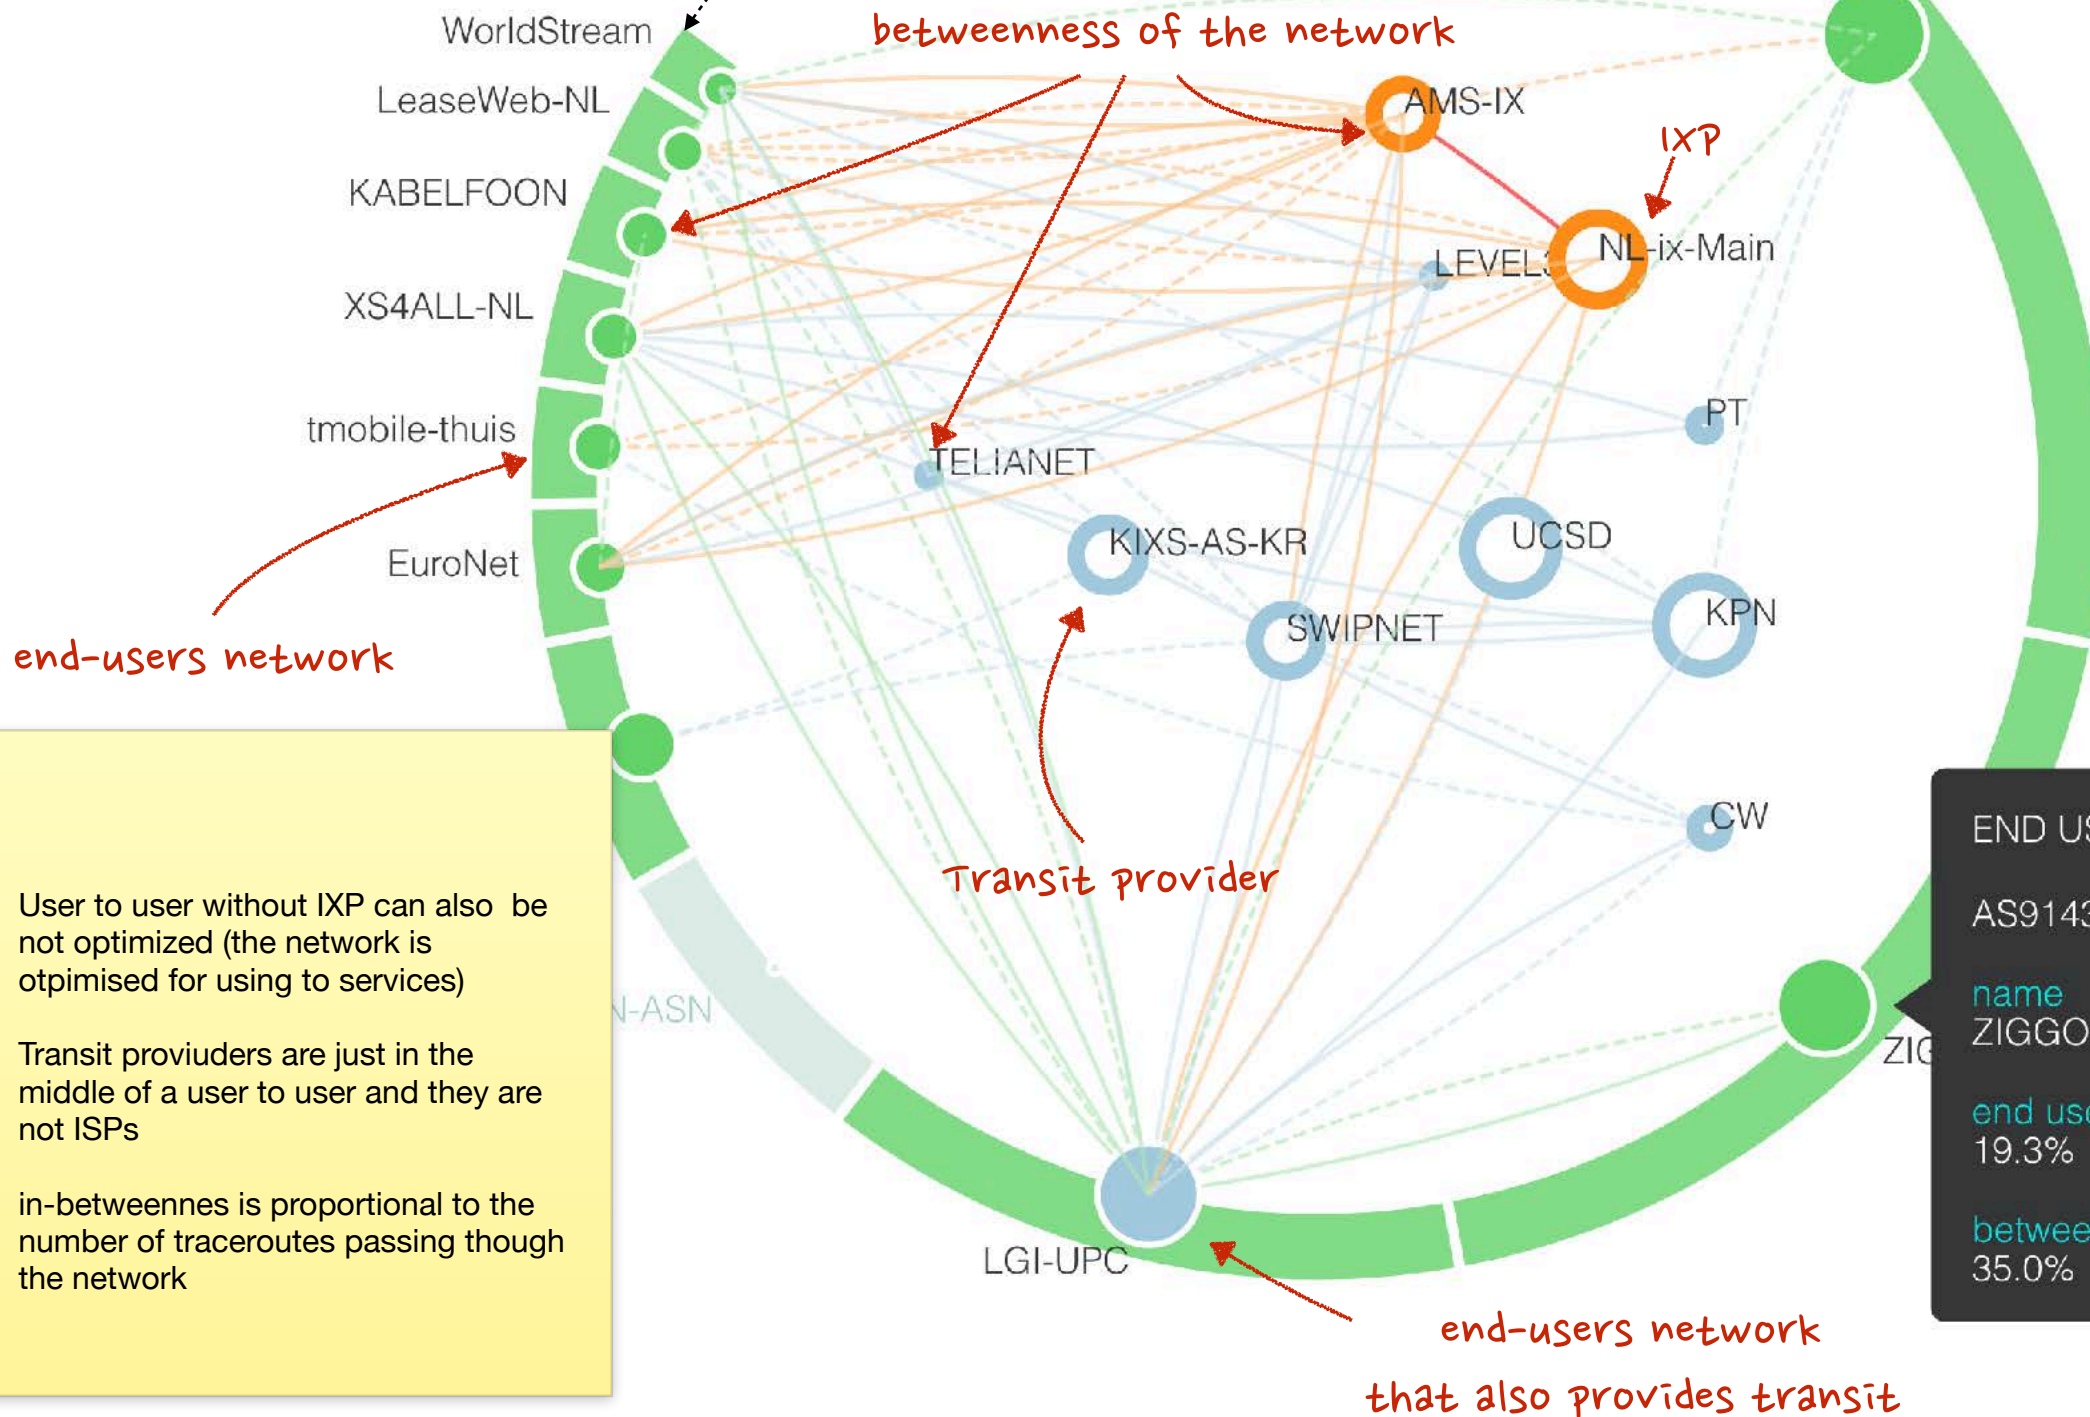
The Netherlands



Empty segment represents ASes
providing connectivity to < 1% users

the size of the circle indicates the in-
betweenness of the network

size of a segment is proportional
to amount of end-users
in the network



User to user without IXP can also be
not optimized (the network is
optimised for using to services)

Transit providers are just in the
middle of a user to user and they are
not ISPs

in-betweenness is proportional to the
number of traceroutes passing through
the network

Russian Federation



orange lines - end-users networks
connected through an IXP

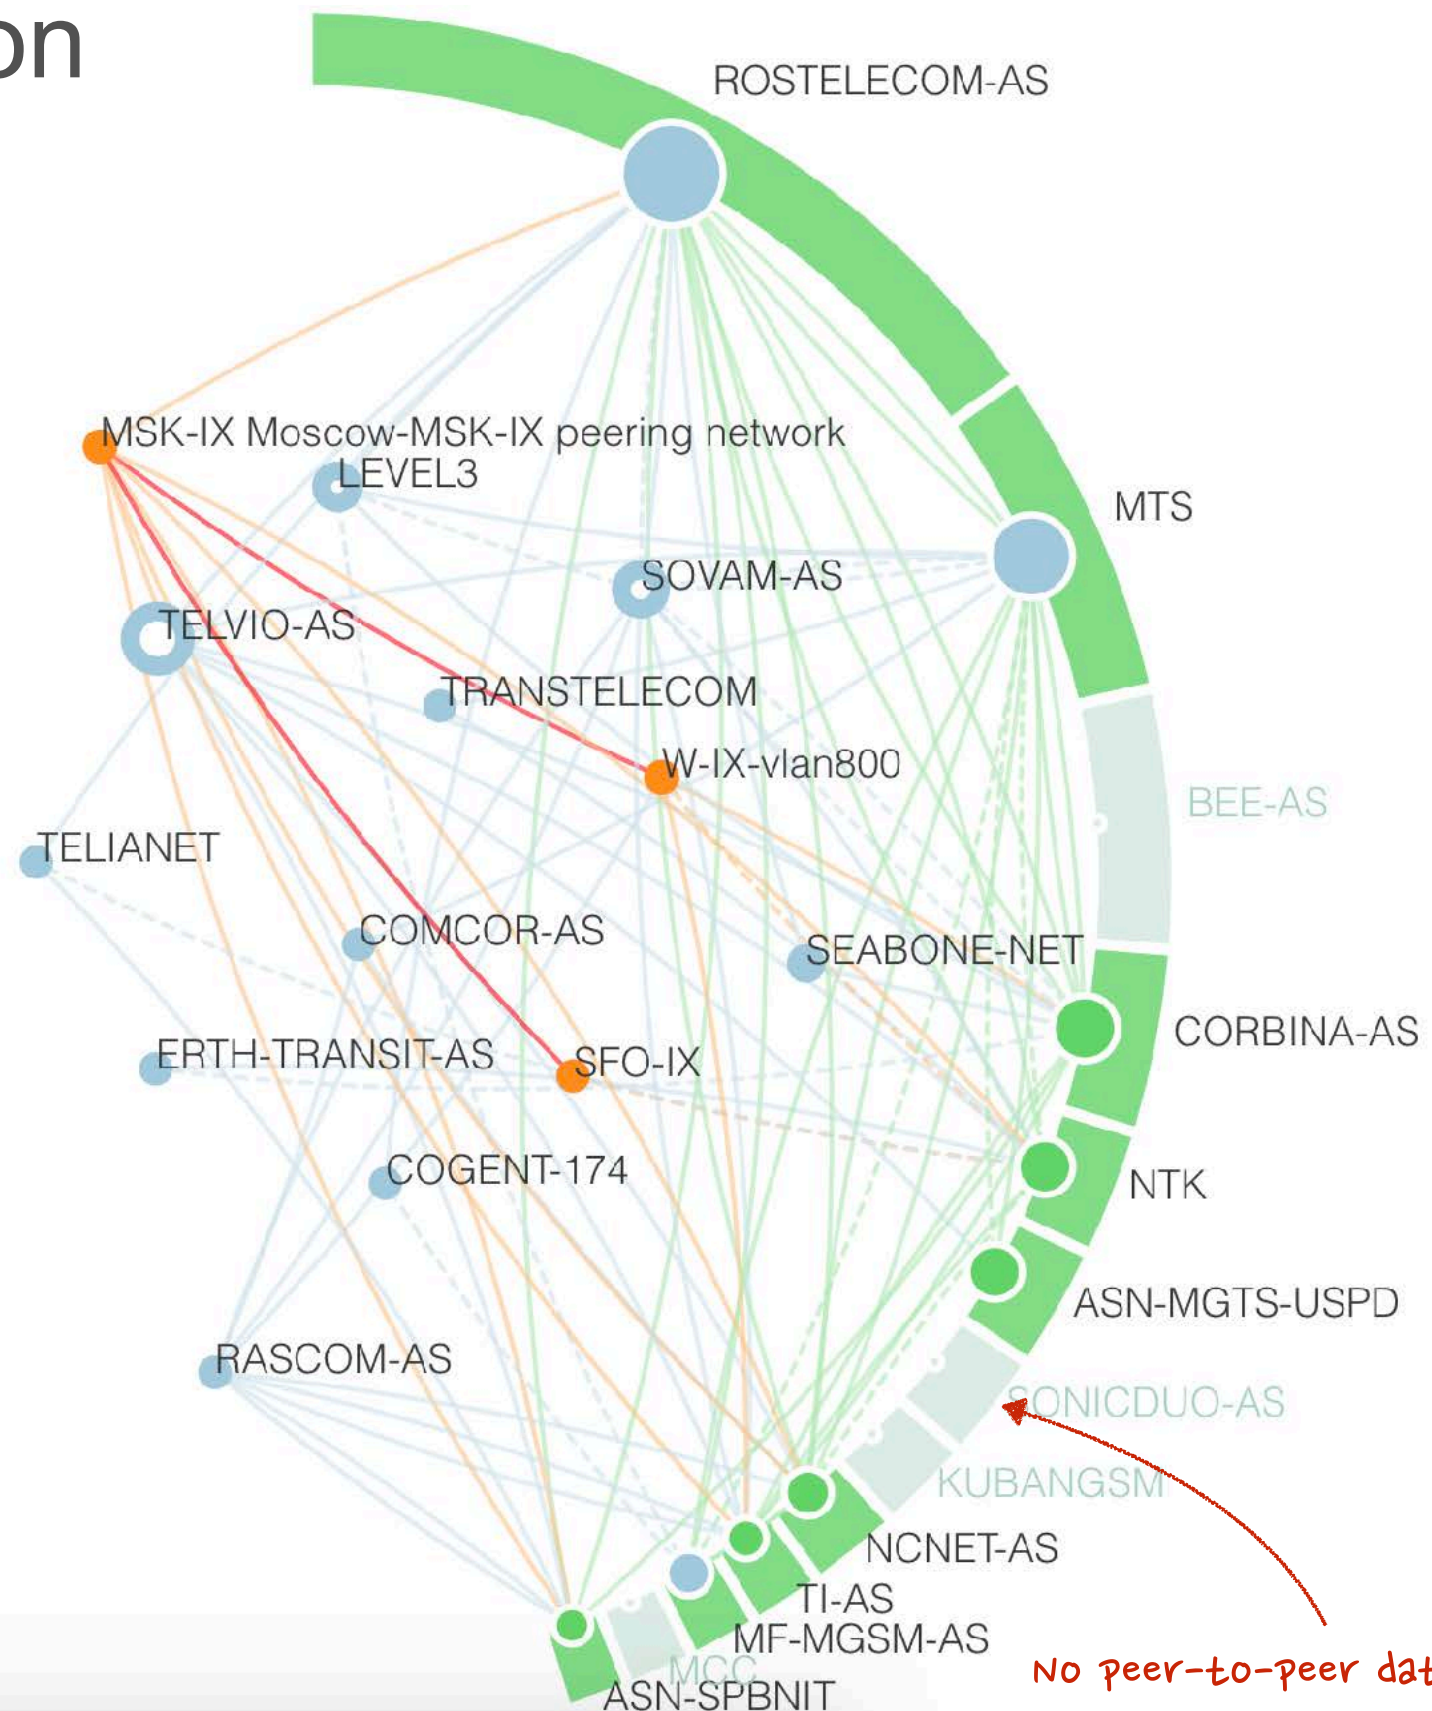
green lines - directly connected

blue lines - through a transit
network

dashed lines - we cannot fully map
this path

dashed lines - where dns lookup is
wrong or have stars in the trace route
outputs

red lines show paths between IXPs





Why Geolocation of infra-structure?

Reasons



**FOCUS ON
Infrastructure**

- Increased interest in IP geolocation
 - Content providers
 - From operators and researchers
 - The “geoloc” attribute is not a solution
- A unified geographical data format is needed
- IP geolocation is extremely difficult
 - Various approaches, some of them cannot be used singularly
 - Academia is working on it! Let's work together
 - A validation/feedback loop is needed
 - A unified geographical data format is needed

Different Research Approaches



- Triangulation a.k.a. trilateration
 - Paristech anycast
 - RIPE Atlas
- Reverse DNS based location inference
 - CAIDA DNS decoding database - DDec
- ‘Administrative’ analyses
 - PeeringDB
 - RIPE Database
- Verification/falsification procedures

Commercial Offerings



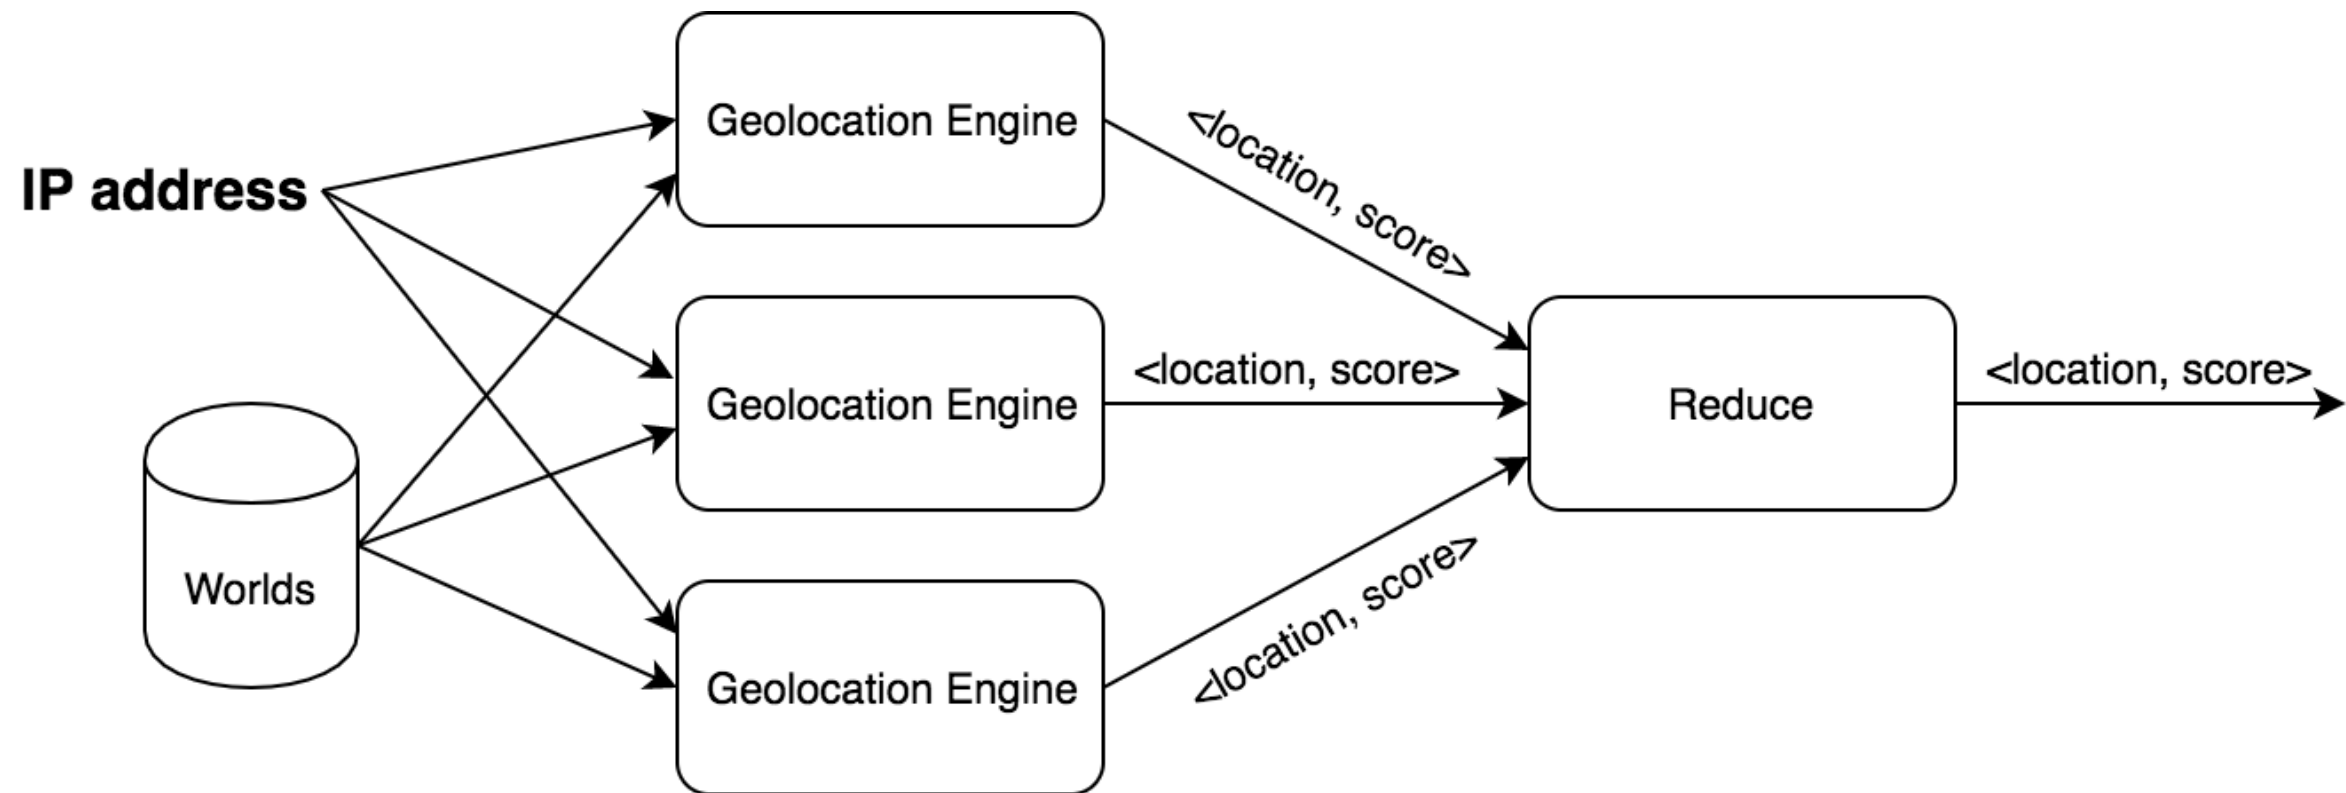
- Tend to concentrate on the end-user IP addresses
- Opaque methodology
- IPv6 address space largely ignored

Our Integration attempt



- Accumulate research efforts as Inference Engines
- Each engine is applicable only in some cases
- Each engine has a score factor
- Complete transparency in inference methodology

Inference Engines



What is ipmap.ripe.net



- A web application where you can query/correct the geolocation of an IP address
- An API where you can query for multiple IP addresses in bulk
- An API where you can correct/provide yourself the geolocation of an IP address
- A web application to visualize traceroutes geolocated on a map

Where is RIPE IPmap?



RIPE NCC
RIPE NETWORK COORDINATION CENTRE

RIPE Database (Whois) Website

Search IP Address or ASN

Manage IPs and ASNs > **Analyse** > Participate > Get Support > Publications > About Us >

You are here: Home > Analyse > Internet Measurements > RIPE Atlas > Measurements

RIPE Atlas <<

- About RIPE Atlas >
- Get Involved >
- Probes and Anchors >
- Measurements, Maps and Tools** v
- Measurements **1**
- Internet Maps
- Tools
- Resources >
- RIPE NCC Members

Measurements


Search by target Search... Any Status IPv4/v6 All types Of all time

Ping **Traceroute** **2** DNS HTTP SSL NTP WiFi Built-in Anchoring

ID	Type	Target	Description	Probes	Interval	Time (UTC)
9855715	Traceroute	185.3.64.1	traceroute_From_NL_POP-Chanrion-Monohome	1	one-off	2017-10-23 15:05 Never
9854358	Traceroute	se-sto-as199150.anchors.atlas.ripe.net	Calibration for anchoring measurement: IPv6 Traceroute for se-sto-as199150.anchors.atlas.ripe.net	4097	one-off	2017-10-23 14:25 2017-10-23 14:35
9854357	Traceroute	se-sto-as199150.anchors.atlas.ripe.net	Calibration for anchoring measurement: IPv4 Traceroute for se-sto-as199150.anchors.atlas.ripe.net	10120	one-off	2017-10-23 14:25 2017-10-23 14:35

Where is RIPE IPmap?



**RIPE NCC**
RIPE NETWORK COORDINATION CENTRE

RIPE Database (Whois) Website

Search IP Address or ASN

Manage IPs and ASNs > Analyse > Participate > Get Support > Publications > About Us

» You are here: Home > Analyse > Internet Measurements > RIPE Atlas > Measurements > Measurement #12317005

⚡ Traceroute measurement to wikipedia.org *here!*

General Information Probes Map TraceMON **IPmap (beta)** Results

General Information

ID	#12317005
Group ID	#12317005
Type	⚡ Traceroute
Public measurement?	Yes
Target	wikipedia.org
Resolve on Probe	No
This is a one-off measurement	
Timing	2018-04-26 12:25 - 2018-04-26 12:35
Costs	60 per result, 1200 per day
Response timeout	4000
Protocol	TCP

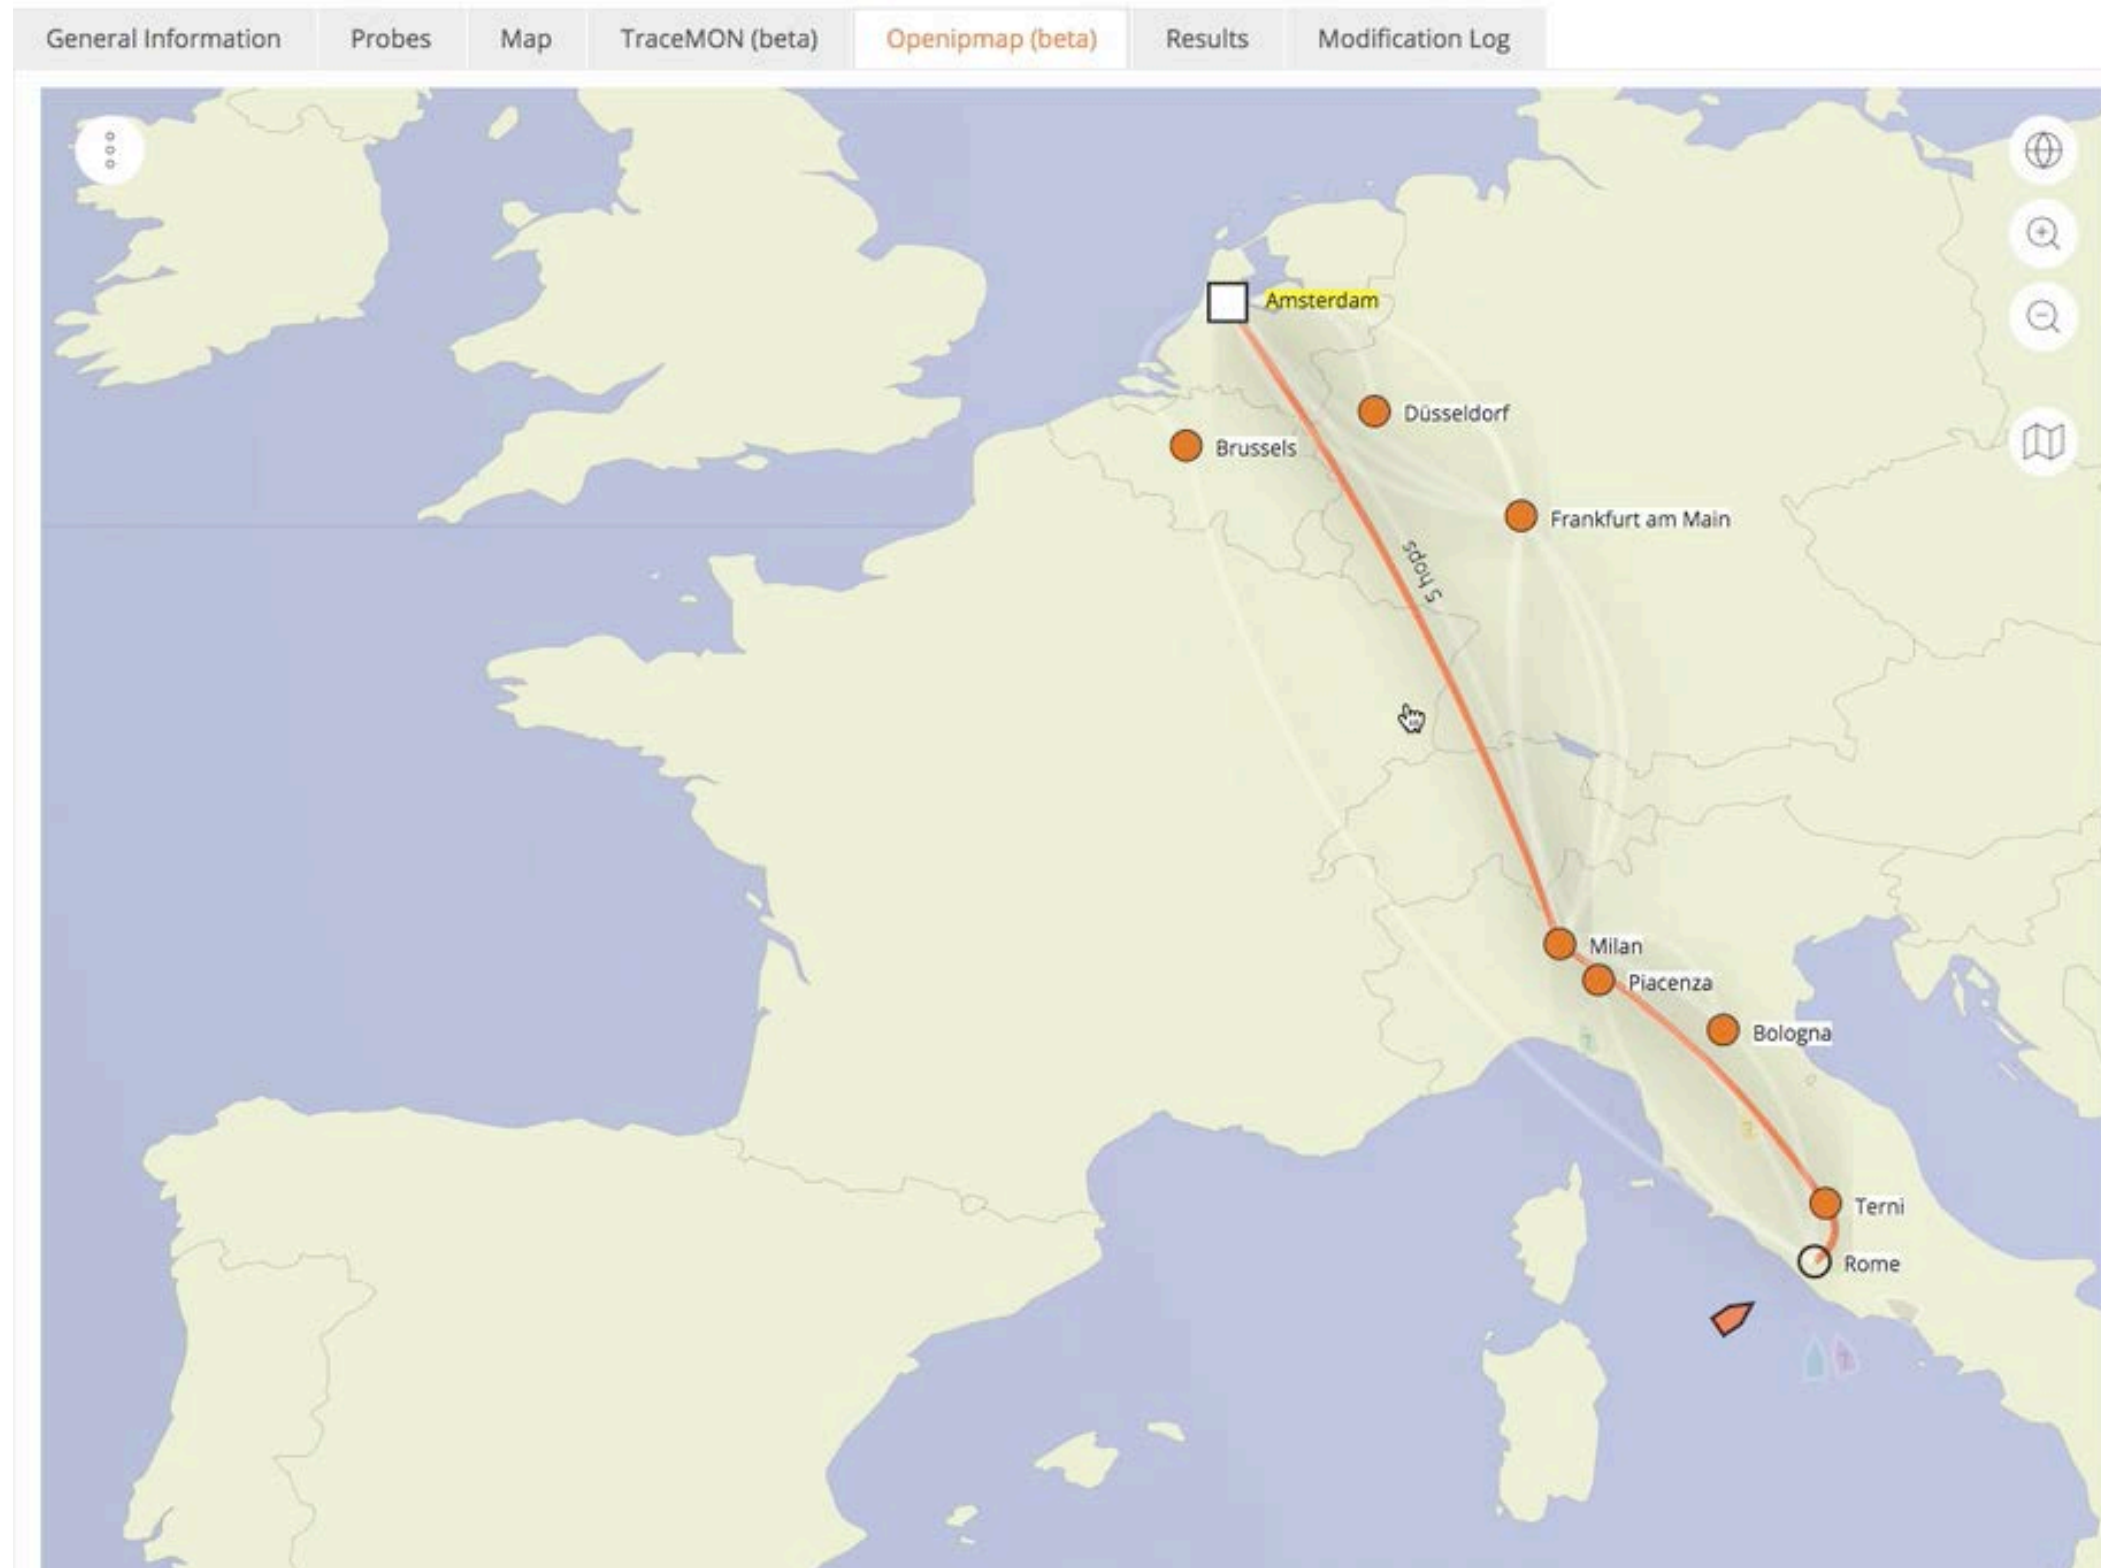
IPmap Demo



IPmap Demo



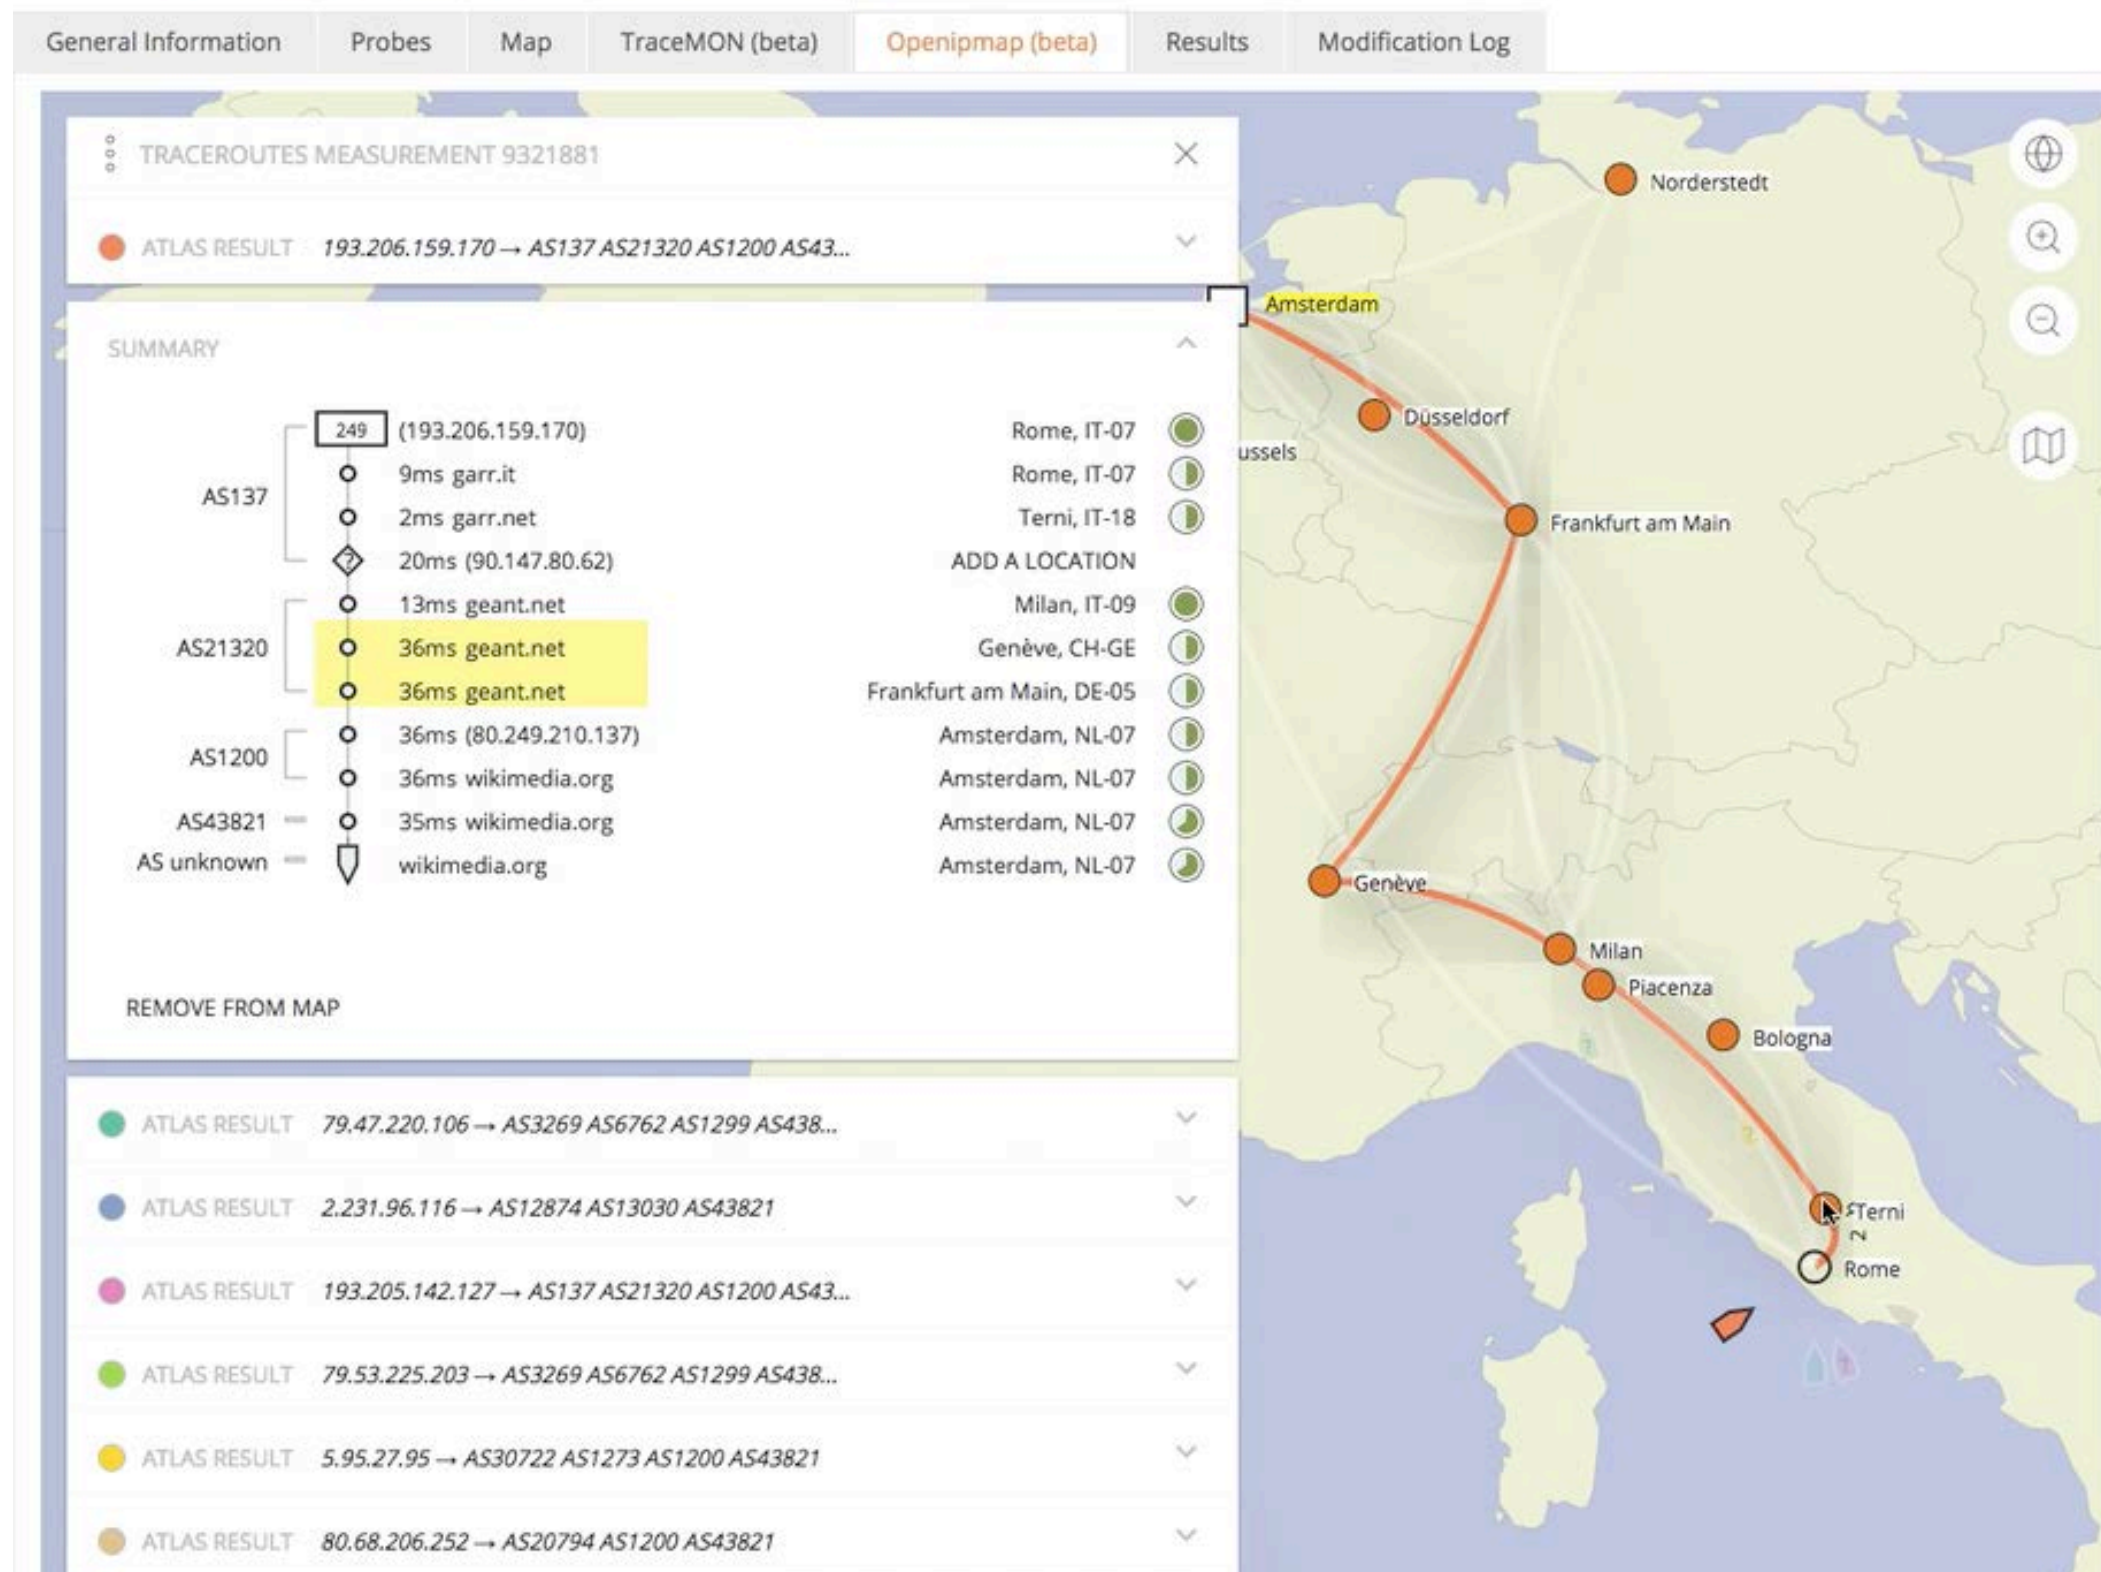
⚡ Traceroute measurement to wikipedia.org



IPmap Demo



⚡ Traceroute measurement to wikipedia.org





Geolocation API

<https://ipmap.ripe.net/api/v1>

/locate



<https://openipmap.ripe.net/api/v1/locate/83.163.50.165/best>

```
{
  "location": {
    "score": 145,
    "countryCodeAlpha3": "NLD",
    "countryCodeAlpha2": "NL",
    "cityPopulation": 147590,
    "stateAnsiCode": "07",
    "pointGeometry": "0101000020E61000005C72DC291D8C12401B81785DBF304A40",
    "cityNameAscii": "Haarlem",
    "stateIsoCode": "NL-07",
    "countryName": "Netherlands",
    "stateName": "North Holland",
    "longitude": 4.63683,
    "geonameId": 2755003,
    "latitude": 52.38084,
    "cityName": "Haarlem",
    "type": "city",
    "id": "HAARLEM-NL-07-U173CX8KTBR196ECJF92"
  },
  "meta": {
    "distribution": {
      "version": "17.9.18.1"
    },
    "service": {
      "version": "0.0.1"
    },
    "request": {
      "params": {
        "ip": "83.163.50.165"
      },
      "query": {}
    }
  }
}
```

*queries can be bundled with:

[https://ipmap.ripe.net/api/v1/locate/
all?resources=ip1,ip2,ip3..](https://ipmap.ripe.net/api/v1/locate/all?resources=ip1,ip2,ip3..)

/locate



<https://openipmap.ripe.net/api/v1/locate/83.163.50.165/partial>

```
{
  {
    "engine": "probeslocation",
    "description": "Probes location suggestor - based on user setting",
    "scoreFactor": 10,
    "locations": [ ... ] // 1 item
  },
  {
    "engine": "anycastparistech",
    "description": "Anycast engine - Paristech dataset",
    "scoreFactor": 10,
    "locations": []
  },
  {
    "engine": "crowdsourced",
    "description": "Crowdsourced engine",
    "scoreFactor": 9,
    "locations": []
  },
  {
    "engine": "triangulation",
    "description": "Triangulation engine (if empty try in 3 minutes, triangulation requires time)",
    "scoreFactor": 5,
    "locations": [ ... ] // 20 items
  }
],
"meta": {
  "distribution": {
    "version": "17.9.18.1"
  },
  "service": {
    "version": "0.0.1"
  },
  "request": {
    "params": {
      "ip": "83.163.50.165"
    },
    "query": {}
  }
}
```

/locate - Active geolocation



- If the IP has not been measured yet, a new Ping measurement starts
- Peering DB data and BGP data are used to reduce the locations probed
- Score based on RTT, only RTT <10ms are considered
- PeeringDB facilities and population bust the score
- A list of possible locations is returned
- We are working on it! (Contributions are welcome!)

That's why you need...



RIPE Atlas coverage!





What's new?

Short RIPE Atlas update

Some New Features and Other Info



- Measurement tagging/labeling
- Result archives
- RIPE Atlas (storage) timestamps
- DNS-over-TLS support (API only yet)
- Scaling up our central infrastructure
 - Working on introducing Elasticsearch

Anchor VMs




- An anchor is a probe and a willing target
 - Automatically measured and generates more credits
- Anchor as a Virtual Machine: almost the same
 - They can be installed where physical machines are a no-go
 - Can be “in the cloud” as well
- We're in the pilot phase as of now
 - Five VMs are up and running
 - With help from members of the community
- We'll report on this activity around end of Q2

Probes



- We stopped preparing new v3 probes
 - Remaining stock is being distributed
- We're testing v4 (NanoPi based)
 - Have a few to give out for field testing
 - Working on logistics and proper casing
- Still (always?) on the lookout for new devices



Measurement Tagging



- It allows you to:
 - Group together any number of your measurements
 - Aggregate results from multiple targets and time periods
 - Easily stop all of a campaign's ongoing measurements
 - Choice between collaboration (tags) and full control (labels)
 - Upcoming: support in various RIPE Atlas visualisations
- For more info:
 - https://labs.ripe.net/Members/chris_amin/ripe-atlas-measurement-tagging

References & Feedback



- RIPE Atlas
 - <https://atlas.ripe.net>
- Peer-to-peer sketches
 - <http://sg-pub.ripe.net/ixp-country-jedi/>
- RIPE IPmap
 - <https://ipmap.ripe.net/>
- RIPE Labs
 - <https://labs.ripe.net>



Questions

