



**RIPE**  
**NCC**



## OpenIPMap

---

Emile Aben  
[emile.aben@ripe.net](mailto:emile.aben@ripe.net)

## *Crowdsourcing Geolocation of IPs*

- Open and crowdsourced mapping of IPs and hostnames to geographical locations
- Emphasis on *infrastructure*, not eyeballs
- Think of OpenStreetMap for IPs



- Creating a database with:

IP



Geographic location

`2001:a18:0:404::4`

`158.64.58.244`

`lu-lux-as2602.anchors.atlas.ripe.net`

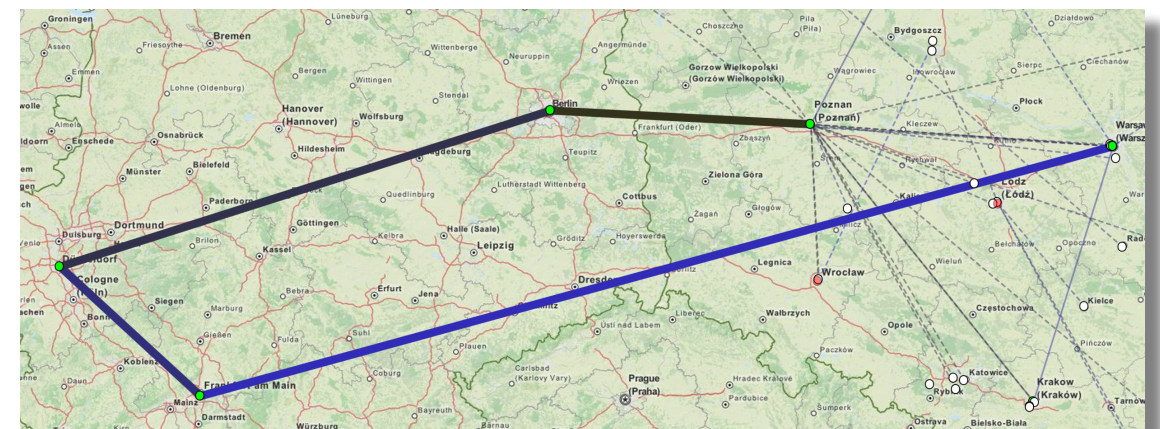




- Anyone can contribute
- Operators have high-quality knowledge
  - Authoritative knowledge of own network



- All Internet users can benefit
- Improves existing geoloc
- Huge potential for improving network diagnostics





- Users can visualise traceroutes
  - Compare changing routes
  - Visually identify outliers / strange routing
- Statistics/Analysis
  - Percentage of observed paths go outside a country?
  - Contribute to tools like RIPEstat and RIPE Atlas
  - Providing raw/bulk data to the community/researchers
- Potentially help answer strategic questions
  - Where to build out my network to/peer?
  - How to reduce latency/improve resiliency/security?

- No comprehensive and open service like this exists currently
  - Existing resources mainly focus on eyeballs, not **infrastructure**

IP address
213.242.117.109
141.136.103.165
62.115.45.201
129.250.4.78

- No comprehensive and open service like this exists currently
  - Existing resources mainly focus on eyeballs, not **infrastructure**

IP address	Geoloc
213.242.117.109	GB
141.136.103.165	FR
62.115.45.201	EU
129.250.4.78	Englewood, CO, US



- No comprehensive and open service like this exists currently
  - Existing resources mainly focus on eyeballs, not **infrastructure**

IP address	Geoloc	Hostname
213.242.117.109	GB	xe-10-2-0-130.bar1. <b>Warsaw</b> 1.Level3.net
141.136.103.165	FR	ae0-103. <b>waw</b> 20.ip4.tinet.net
62.115.45.201	EU	<b>war</b> -b1-link.telialia.net
129.250.4.78	Englewood, CO, US	xe-4-1.r00. <b>wrs</b> wpl01.pl.bb.gin.ntt.net

- No comprehensive and open service like this exists currently
  - Existing resources mainly focus on eyeballs, not **infrastructure**

IP address	Geoloc	Hostname
213.242.117.109	GB	xe-10-2-0-130.bar1. <b>Warsaw</b> 1.Level3.net
141.136.103.165	FR	ae0-103.waw20.ip4.tinet.net
62.115.45.201	EU	war-b1-link.telialia.net
129.250.4.78	Englewood, CO, US	xe-4-1.r00.wrswwpl01.pl.bb.gin.ntt.net

***Location=Warsaw,PL!***

- Crowdsourcing is a proven concept
- Anyone can contribute
  - No need to be authoritative (DNS LOC, registry)
  - It is possible to use a source confidence factor
- Operators have high-quality knowledge
  - e.g. Integrating with RIR authentication we can collect high confidence input
- OpenIPMap should be more global than RIPE NCC service region



- Import existing data with lower confidence
  - eg. populate with RIPE Database and CAIDA data
- Help crowdsourcing efforts with a suggestion engine
  - Artificial Intelligence
- Capture naming schemes
  - Using tags: simple, but potentially ambiguous

**esams in wikimedia.org = Amsterdam, NL**
  - Using regex: complex, but can capture everything

**$^([a-z]{3})\d+.*\backslash.1e100\backslash.net$  = IATA:\1**
  - Domain analysis tool

# We have a working prototype

msm:1664844 prb:15110 ts:2014-05-09T07:43:25.000Z					
1664844 15110 1399621405					
1	10.220.9.1		0.6 0.5 0.5		V
2	213.5.8.193		14.1 14.6 14.2		V
3	*				
4	78.152.54.1	r1.war1.pl.as5580.net	1.1 1.0 1.6	Warsaw,Masovian Voiv	V ok
5	78.152.34.93	eth1-1.r1.pra1.cz.atrato.net	14.4 14.7 14.5	parana,08,ar	V ok
6	78.152.34.97	eth3-6.edge1.fra6.de.atrato.net	21.7 43.1 30.9	frankfurtammain,05,de	V ok
7	78.152.44.246	eth11-1.core1.ams1.nl.atrato.net	28.1 27.8 27.9	amsterdam,07,nl	V ok
8	78.152.42.92	eth1-8.r1.ams2.nl.atrato.net	29.1 39.9 39.5	amsterdam,07,nl	V ok
9	195.69.147.140		28.9 28.8 28.8	Amsterdam,North Holl	V
10	190.93.244.243		28.7 28.5 28.7		V

- Implemented the basic functionality with some of the additional possibilities to populate the initial set of data

# We have a working prototype

msm:1664844 prb:15110 ts:2014-05-09T07:43:25.000Z

1664844|15110|1399621405

1	10.220.9.1		0.6 0.5 0.5		V
2	213.5.8.193		14.1 14.6 14.2		V
3	*				
4	78.152.54.1	r1.war1.pl.as5580.net	1.1 1.0 1.6	Warsaw,Masovian Voiv	V ok
5	78.152.34.93	eth1-1.r1.pra1.cz.atrato.net	14.4 14.7 14.5	parana,08,ar	V ok
6	78.152.34.97	eth3-6.edge1.fra6.de.atrato.net	21.7 43.1 30.9	frankfurtammain,05,de	V ok
7	78.152.44.246	eth11-1.core1.ams1.nl.atrato.net	28.1 27.8 27.9	amsterdam,07,nl	V ok
8	78.152.42.92	eth1-8.r1.ams2.nl.atrato.net	29.1 39.9 39.5	amsterdam,07,nl	V ok
9	195.69.147.140		28.9 28.8 28.8	Amsterdam,North Holl	V
10	190.93.244.243		28.7 28.5 28.7		V

*Crowdsourced: green*

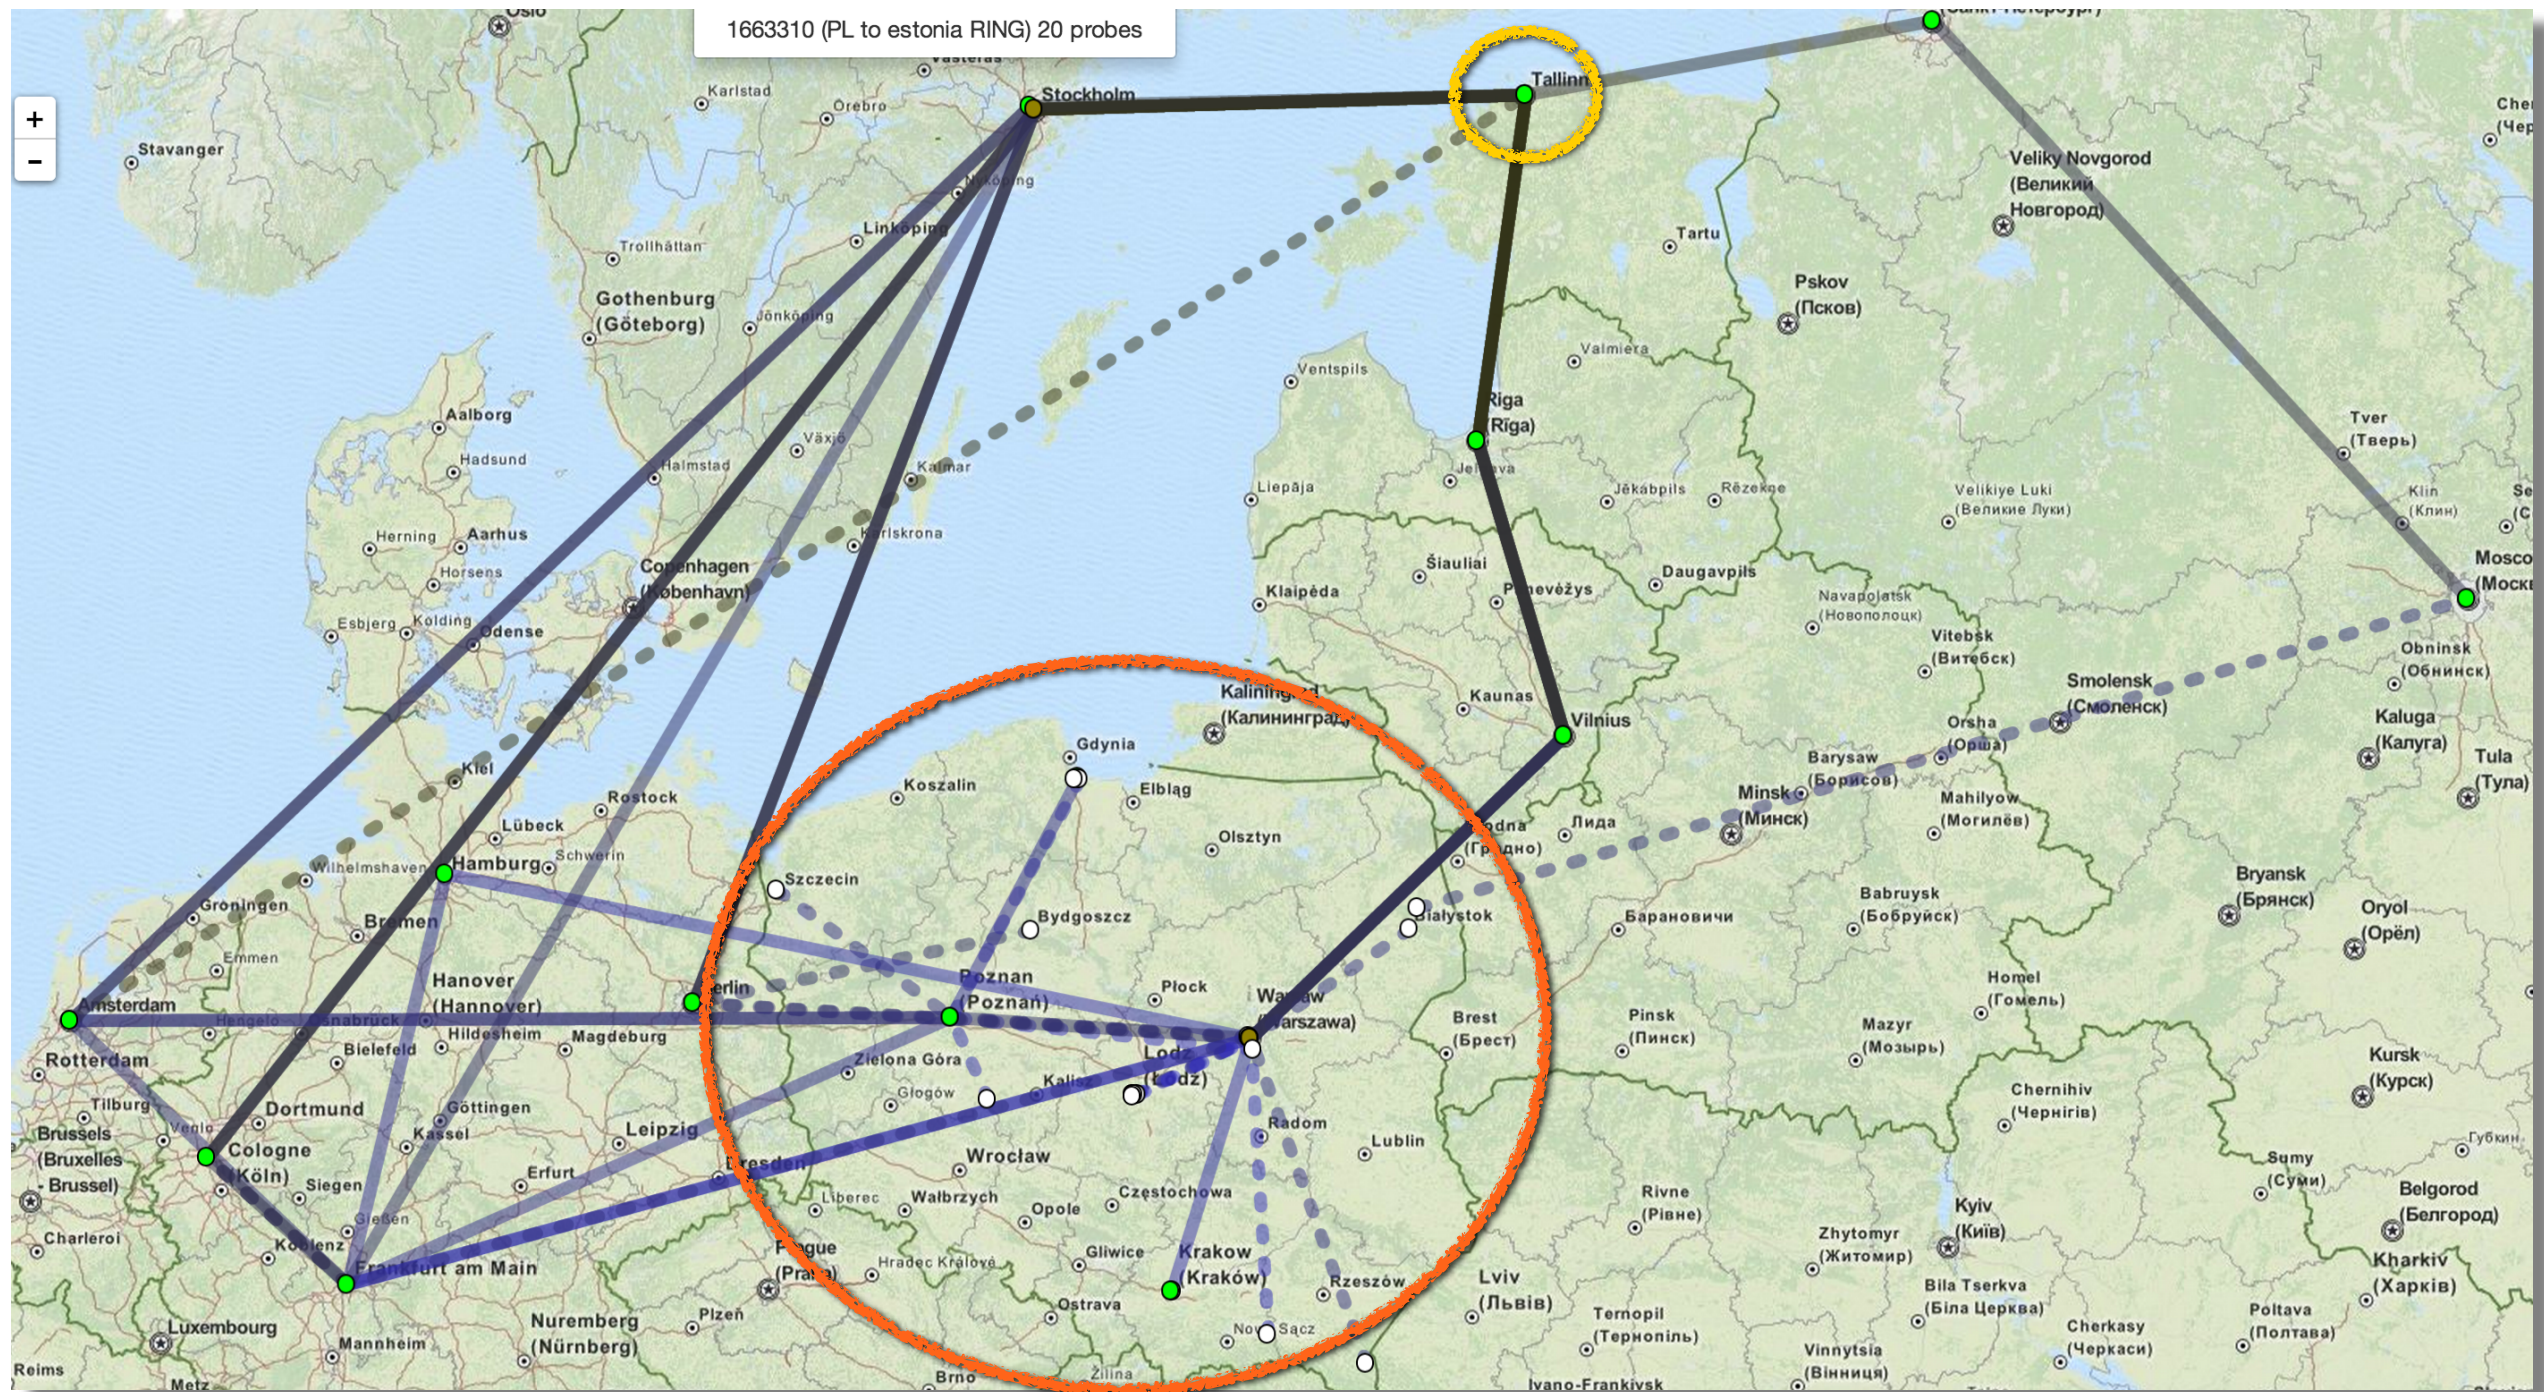
*Guesses: red*

- Implemented the basic functionality with some of the additional possibilities to populate the initial set of data



# We have a working prototype

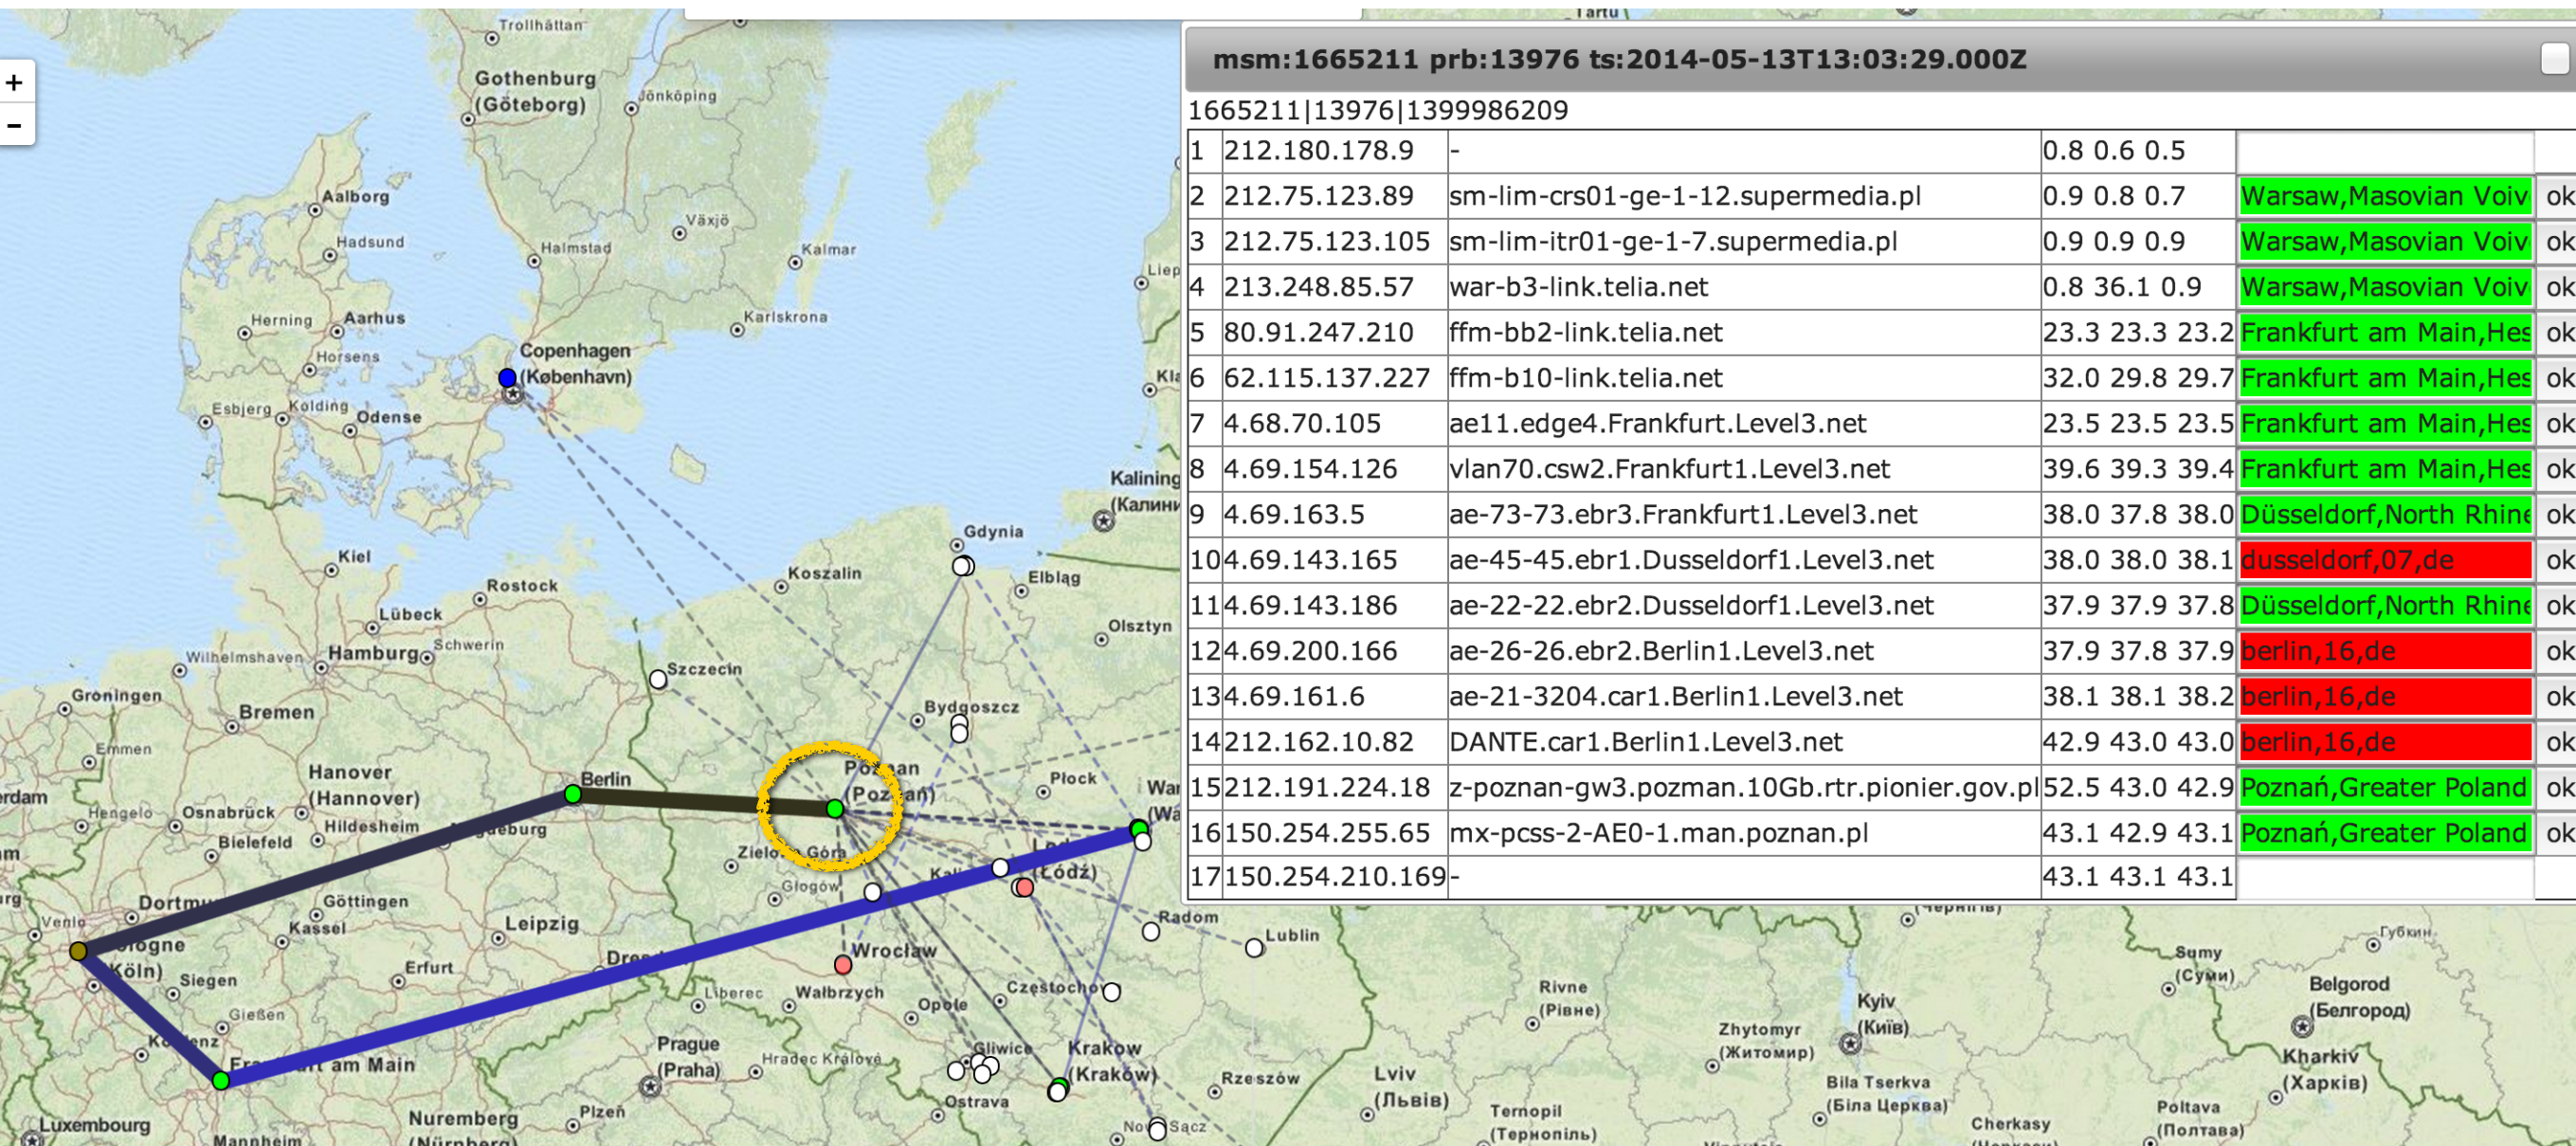
## From: Poland (30x) To: Tallinn, EE





# We have a working prototype

## From: Poland (30x) To: Poznan,PL



- Idea received a lot of support at RIPE 67 MAT WG
- Is this going in the right direction?
- Do you want to see this as a service in the future?
  - We can implement the service, publish the code and initial data and bootstrap it
  - Community decides who will operate such a global service
- Want to know more or help: Contact us!
  - To see the prototype
  - To participate in testing the prototype
- Updates on RIPE Labs: <https://labs.ripe.net/>





- Web interface
- API
- Bulk import. Format:
  - <resource>,<location>,<confidence>
  - 5.57.80.0/22,"London,GB",20
  - lu-lux-as2602.anchors.atlas.ripe.net,"Luxembourg,LU",90

- Web Interface
- API
- Bulk download
  - Also for Geoloc providers

- IP prefix / hostname:
  - *location*: User contributed location string: “Warsaw,PL”
  - *user*: User ID
  - *confidence*: 0-100%

Database adds (via geocoder):

- *canonical location*: Canonical version of location
- *lat, lon*: Coordinates

Not yet implemented:

- Update creates new object
- Anycast resources
- Granularities other than city

- User ID
- Authenticated with RIPE Access
  - Could do special treatment if user = resource holder
    - only allow confidence > 90% if user=resource holder?
- Other authentication?
  - PGP
  - OpenID

- Uses RIPE Atlas traceroutes
- Builds on underlying API to access crowd-sourced information
- Suggestion engine for when no crowd-sourced information is available
  - Based on IATA, CLLI, and city names



- Current focus on **city level** granularity
  - Conform to other geoIP databases
  - Backbone maps typically are city level
- PoP/datacenter level?
  - Is technically possible to implement if needed
- Others could be supported too
  - DNS LOC-style: point + radius
  - Arbitrary polygons?