



RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Multi-Approach Infrastructure Geolocation

Massimo Candela | 24 October 2017 | RIPE 75

Reasons



1. Increased interest in IP geolocation

- Content providers, researchers
- <https://www.ripe.net/manage-ips-and-asns/db/tools/geolocation-in-the-ripe-database>

2. A unified geographical data format is needed

- Across all our tools
- Enabling future geographical investigations
- Accurate information

3. IP geolocation is extremely difficult

- Various approaches, some of them not accurate enough to be used singularly
- Academia is working on it! Let's work together

**FOCUS ON
Infrastructure**

What's new: RIPE NCC's geo service



geo.ripe.net



```
▼ {  
  ▼ {  
    "url": "/locate",  
    "description": "Geolocation service. It provides geolocation according to a set of passive and active geolocation approaches"  
  },  
  ▼ {  
    "url": "/worlds",  
    "description": "Worlds dataset, providing standard geolocation format to all other services"  
  },  
  ▼ {  
    "url": "/crowdsourced",  
    "description": "Geolocation service based on crowdsourced information."  
  },  
  ▼ {  
    "url": "/peeringdb",  
    "description": "PeeringDB interface for geolocation purposes."  
  },  
  ▼ {  
    "url": "/triangulation",  
    "description": "Active geolocation service based on latency triangulation."  
  },  
  ▼ {  
    "url": "/anycast",  
    "description": "Anycast geolocation service based on active measurements"  
  }  
}
```

/locate



<https://geo.ripe.net/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://geo.ripe.net/locate/all?
resources=ip1,ip2,ip3](https://geo.ripe.net/locate/all?resources=ip1,ip2,ip3)

/locate



<https://geo.ripe.net/locate/83.163.50.165/>

```
{
  "locations": [
    {
      "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"
    },
    {
      "id": "AMSTERDAM-NL-07-U173ZQ2SF4C47GPE4JPJ",
      "type": "city",
      "cityName": "Amsterdam",
      "latitude": 52.37403,
      "geonameId": 2759794,
      "longitude": 4.88969,
      "stateName": "North Holland",
      "countryName": "Netherlands",
      "stateIsoCode": "NL-07",
      "cityNameAscii": "Amsterdam",
      "pointGeometry": "0101000020E61000001E5036E50A8F134087DC0C37E02F4A40",
      "stateAnsiCode": "07",
      "cityPopulation": 741636,
      "countryCodeAlpha2": "NL",
      "countryCodeAlpha3": "NLD",
    }
  ]
}
```

Multi-Approach Geolocation



Various engines contribute to geolocation

- Each of them is applicable only in some cases
- Some of them are used to remove false positives
 - Active geolocation is not accurate in many cases - but can clearly exclude locations
- Some engines are more accurate with a POST request
 - To provide additional data; e.g. a traceroute
- Each of them has a score factor
- You can select the engines you want to use
- They “confirm” each other

/locate



<https://geo.ripe.net/locate/83.163.50.165/partials>

```
▼ {
  ▼ "partials": [
    ▼ {
      "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"
    },
  },
}
```

Anycast GET vs. POST (8.8.8.8)



The POST body provides additional hints to the engine

GET

```
{
  "locations": [
    {
      "score": 105,
      "countryCodeAlpha3": "CZE",
      "countryCodeAlpha2": "CZ",
      "cityPopulation": 1165581,
      "stateAnsiCode": "52",
      "pointGeometry": "0101000020E61000000FEECEDA6DD72C40B35E0CE5440B4",
      "cityNameAscii": "Prague",
      "stateIsoCode": "CZ-52",
      "countryName": "Czechia",
      "stateName": "Praha",
      "longitude": 14.42076,
      "geonameId": 3067696,
      "latitude": 50.08804,
      "cityName": "Prague",
      "type": "city",
      "id": "PRAGUE-CZ-52-U2FKBNHZUP9HBDDQK374"
    },
    {
      "score": 100,
      "countryCodeAlpha3": "DEU",
      "countryCodeAlpha2": "DE",
      "cityPopulation": 1739117,
      "stateAnsiCode": "04",
      "pointGeometry": "0101000020E61000002BA4FCA4DA0724400B5EF415A4C94A",
      "cityNameAscii": "Hamburg",
      "stateIsoCode": "DE-04",
      "countryName": "Germany",
      "stateName": "Hamburg",
      "longitude": 10.01534,
      "geonameId": 2911298,
      "latitude": 53.57532,
      "cityName": "Hamburg",
      "type": "city",
      "id": "HAMBURG-DE-04-U1X0GCHFY7DB2FK5H8J2"
    },
    {
      "score": 100,
      "countryCodeAlpha3": "USA",
```

POST

* based on ParisTech
Anycast Dataset

```
{
  "locations": [
    {
      "score": 100,
      "countryCodeAlpha3": "DEU",
      "countryCodeAlpha2": "DE",
      "cityPopulation": 1739117,
      "stateAnsiCode": "04",
      "pointGeometry": "0101000020E61000002BA4FCA4DA0724400B5EF415A4C94A",
      "cityNameAscii": "Hamburg",
      "stateIsoCode": "DE-04",
      "countryName": "Germany",
      "stateName": "Hamburg",
      "longitude": 10.01534,
      "geonameId": 2911298,
      "latitude": 53.57532,
      "cityName": "Hamburg",
      "type": "city",
      "id": "HAMBURG-DE-04-U1X0GCHFY7DB2FK5H8J2"
    },
    {
      "id": "LIBEN-CZ-52-U2FM13UJ3GTSM0TBXWBN",
      "type": "city",
      "cityName": "Libeň",
      "latitude": 50.10819,
      "geonameId": 3071966,
      "longitude": 14.47457,
      "stateName": "Praha",
      "countryName": "Czechia",
      "stateIsoCode": "CZ-52",
      "cityNameAscii": "Liben",
      "pointGeometry": "0101000020E61000005053CBD6FAF22C408BE07F2E",
      "stateAnsiCode": "52",
      "cityPopulation": 31756,
      "countryCodeAlpha2": "CZ",
      "countryCodeAlpha3": "CZE",
      "score": 50
    },
    {
```


/locate - Active geolocation



1. 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!)

* Vasileios Giotsas RIPE 73 Hackathon project

/crowdsource



```
▼ {
  ▼ "actions": [
    ▼ {
      "method": "GET",
      "url": "('/:ip/all",
      "description": "Get all crowdsourced information for a given IP address"
    },
    ▼ {
      "method": "GET",
      "url": "('/:ip",
      "description": "Get best location based on crowdsourced information for the given IP address"
    },
    ▼ {
      "method": "POST",
      "url": "('/:ip",
      "description": "Crowdsource a geolocation for the specified IP address",
      ▼ "inputExamples": [
        ▼ {
          "latitude": 1.1,
          "longitude": 1.1
        },
        ▼ {
          "id": "HAARLEM-NL-07-U173CX8KTBR196ECJF92",
          "type": "city"
        },
        ▼ {
          "cityName": "Amsterdam",
          "country": "NL",
          "state": null
        }
      ]
    }
  ],
  ▼ "meta": {
```

Extra: /peeringdb



```
▼ {
  "method": "GET",
  "url": "/ixps",
  "description": "List of all IXPs"
},
▼ {
  "method": "GET",
  "url": "/ixps/:id",
  "description": "Specific IXP information"
},
▼ {
  "method": "GET",
  "url": "/ixps/:id/peerings",
  "description": "Provides all the ASNs in peering in a specified IXP"
},
▼ {
  "method": "GET",
  "url": "/ixps/peerings/:asn",
  "description": "It shows all the IXPs an AS is peering with"
},
▼ {
  "method": "GET",
  "url": "/facilities",
  "description": "Get all facilities available"
},
▼ {
  "method": "GET",
  "url": "/facilities/:id",
  "description": "Information about a specific facility"
},
▼ {
  "method": "GET",
  "url": "/facilities/customers/:asn",
  "description": "Get all the facilities in which the AS is present"
},
▼ {
  "method": "GET",
  "url": "/facilities/:id/customers/",
  "description": "Get all the ASNs hosted in a specific facility"
}
```

**Refer to the official PeeringDB for
up-to-date data**

**Update your PeeringDB
information!**

<https://peeringdb.com/>

Future Work



1. Introduce new geolocation engines
 - Increase research collaborations
 - Integrate RIR data
 - Reverse DNS engine
 - Third-party services
2. Define and publish some KPI for service evaluation
 - We already collect metadata
3. Define user rating policy for crowdsourcing
4. Define incentive for contribution (Atlas credits?)

RIPE NCC's geo service



geo.ripe.net



Questions

mcandela@ripe.net



Do you want to include your
geolocation engine in our collection?
Contact me!