The Remote Peering Jedi

A portal in the remote peering ecosystem

RIPE 73
24 - 28 OCT 2016
MADRID, SPAIN

Vasileios Giotsas, UCSD/CAIDA, vgiotsas@caida.org
Petros Gigis, ICS-FORTH/UOC, gkigkis@ics.forth.gr
Alexandros Milolidakis, ICS-FORTH/UOC, alexmil@ics.forth.gr
Eric Nguyen Duy, AMS-IX, eric.nguyenduy@ams-ix.net
Marios Isaakidis, UCL, marios.isaakidis.15@ucl.ac.uk
Edwards Mukasa, NFT Consult, edwards.mukasa@nftconsult.com

caida
FORTH Foundation for Research & Technology - Hellas
amsix amsterdam internet exchange
UCL
NFT CONSULT
RIPE

Motivation

- Knowing which peers allows better-informed peer selection process and higher transparency.
- Operational concerns:
  - Added latency and troubleshooting complexity
  - Routing inefficiencies
  - Invisible Layer-2 intermediaries
  - Network economics and business models

Castro, Ignacio, Juan Camilo Cardona, Sergey Gorinsky, Pierre Francois.
"Remote peering: More peering without internet flattening."
ACM CoNEXT 2014.

Methodology

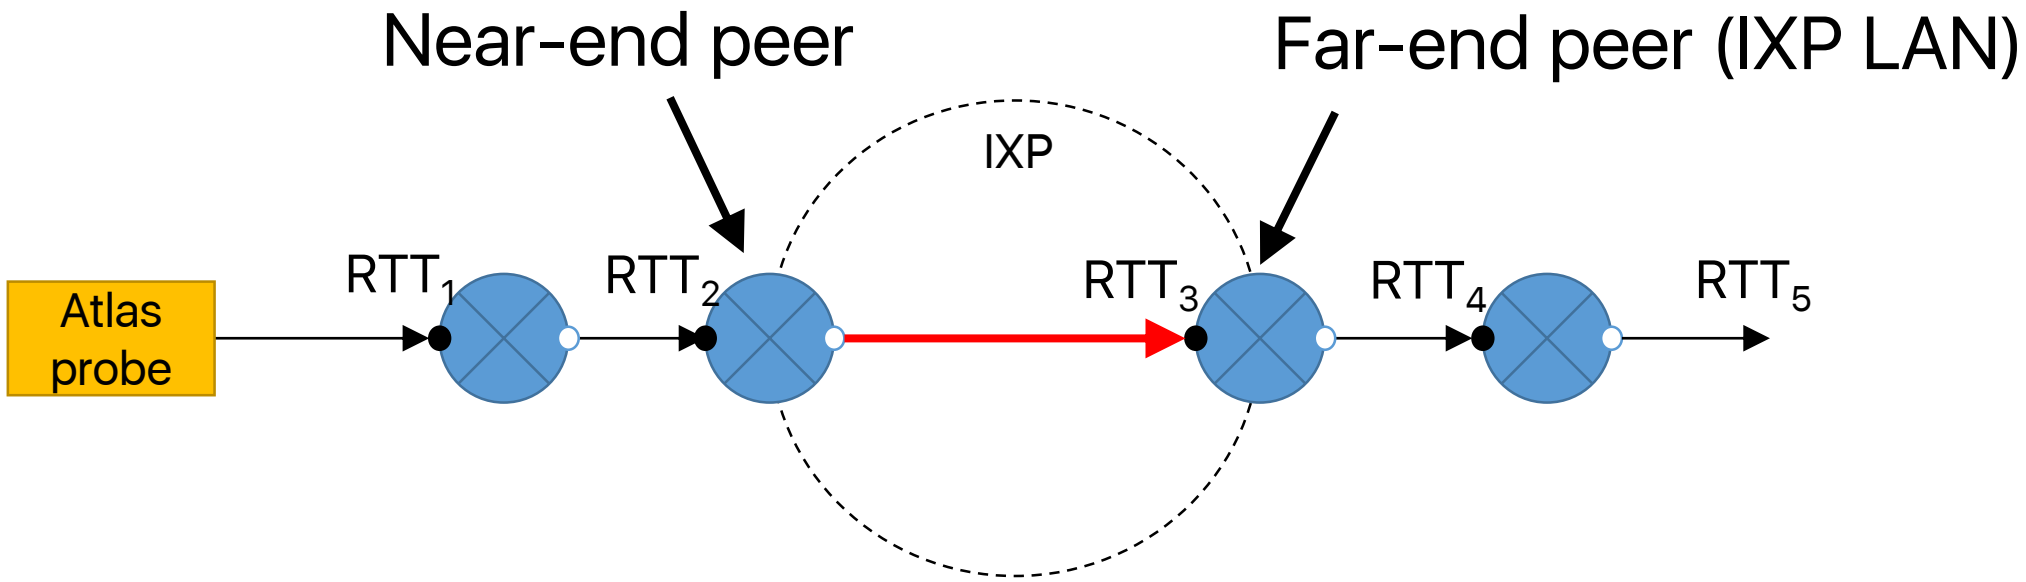

- Parse traceroute paths and detect IXP hops according to traIXroute
- Calculate the RTT between the IXP far-end and near-end peers:

RTT of IXP link: $RTT_3 - RTT_2$

Robust RTT estimations

- Latency estimation from traceroute can be noisy

BUT

- RIPE Atlas offers a massive corpus of traceroute paths from diverse vantage points
  - Multiple observations allow us to remove outliers and de-noise

- For every pair of near-end IP, far-end IXP we require at least 50 paths from which we calculate the median RTT difference.
- Take bottom 50% of lower percentile of RTTs, infer remote peering if `Median_RTT_diff  ≥ 20 ms`

Results

~60-75% ≤ 10 ms (definitely local peers)

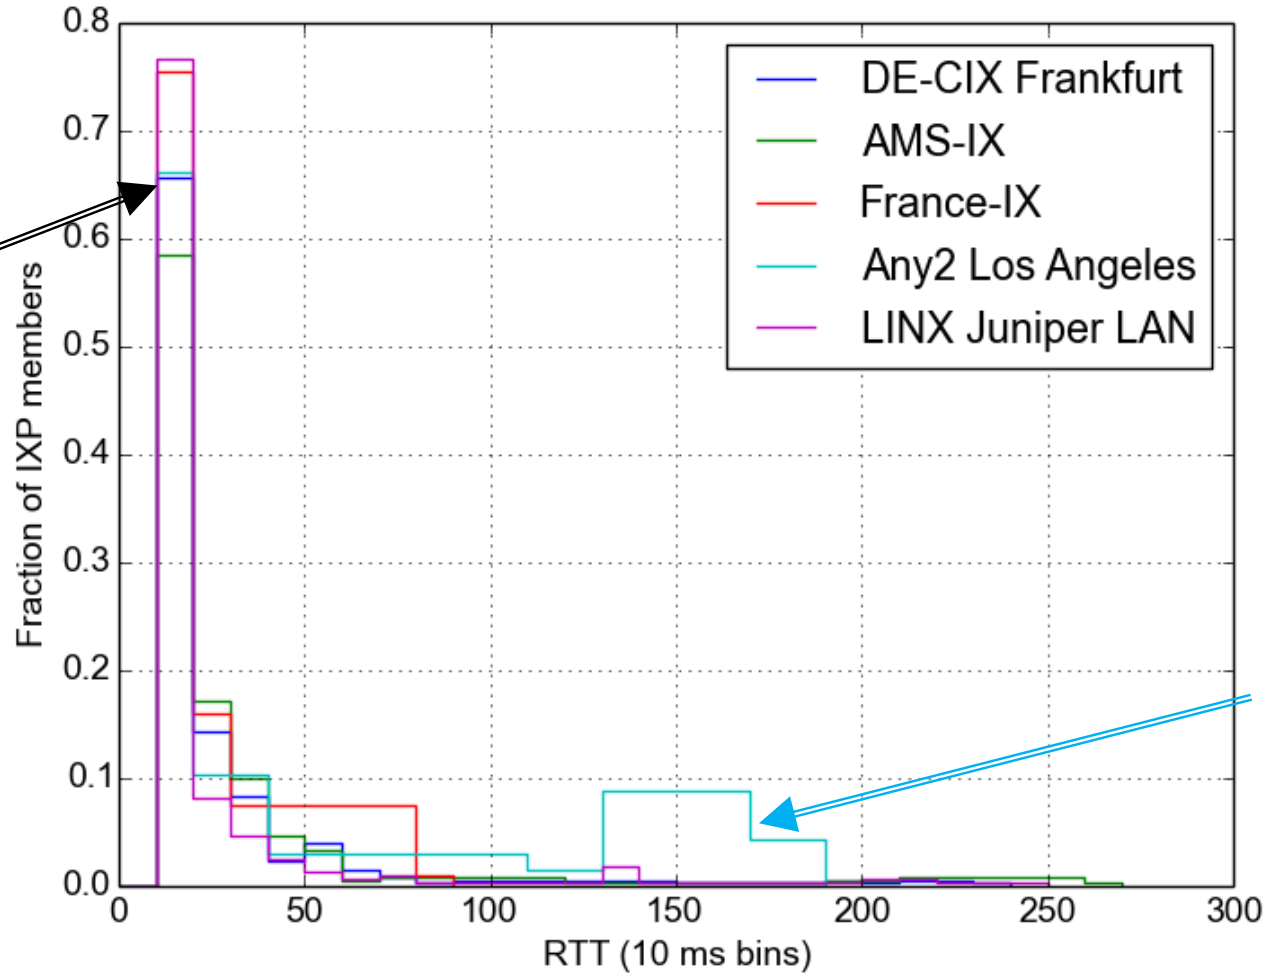

Trans-pacific remote peers of Any2 Los Angeles

Validation

- We collected validation data for the latencies from three large IXPs (RTT_{IXP}) and compared it against the RTTs estimated through Atlas (RTT_{ATLAS}):
  - AMS-IX (ARPing from inside the IXP)
  - DE-CIX (Ping from inside the IXP)
  - France-IX (Ping from inside the IXP)

- True positive if $RTT_{IXP} \geq 20ms$ and $RTT_{ATLAS} \geq 20ms$
  - France-IX: 99%
  - DE-CIX: 99%
  - AMS-IX: 97%

Remote peerings used to access large CDNs

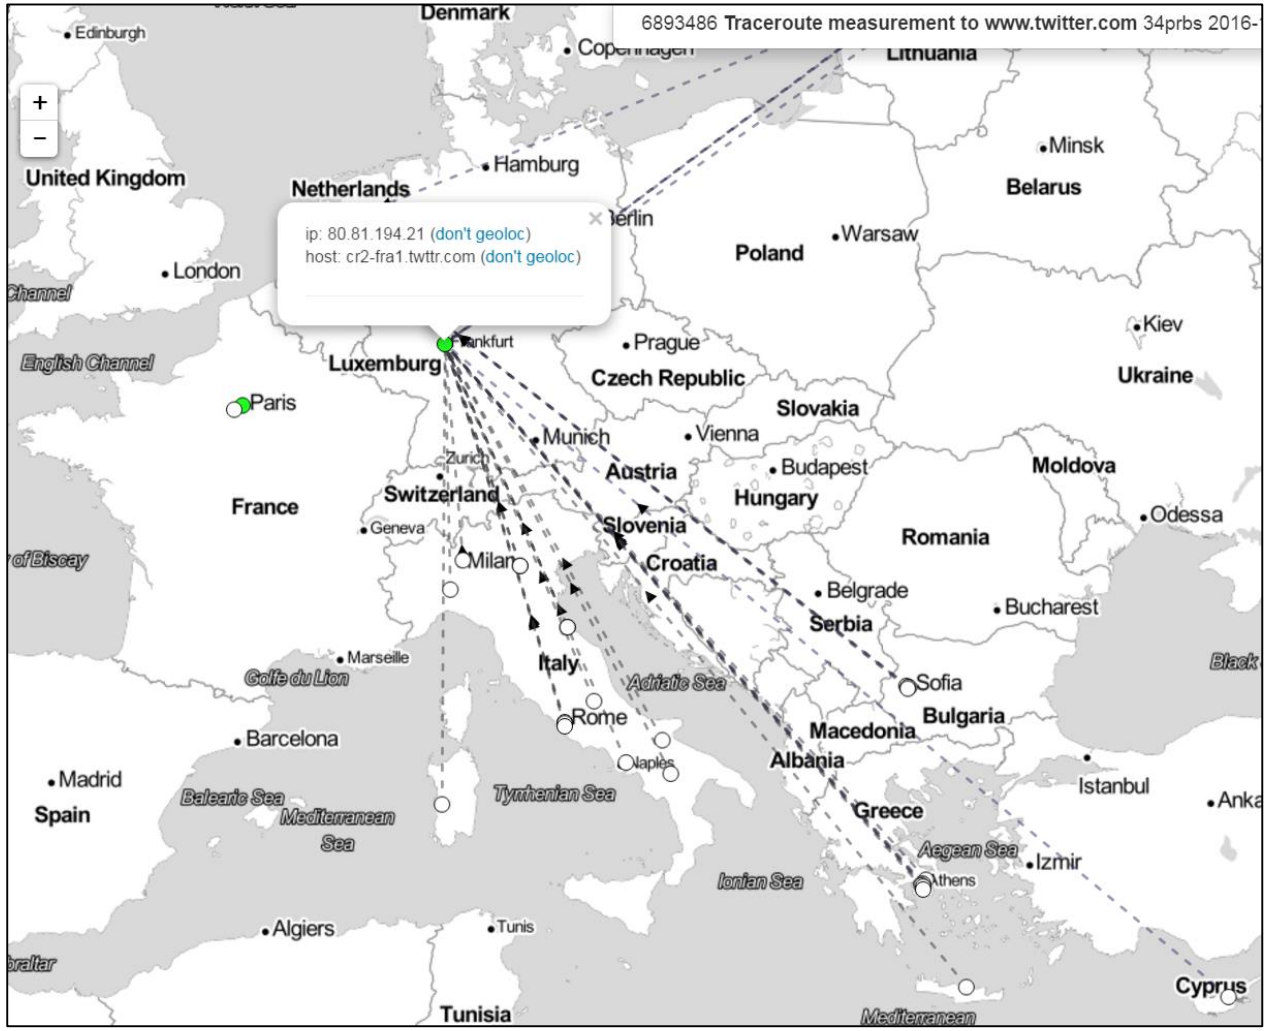

Traceroutes to twitter.com and reddit.com from top-10 remote peers

RIPE 73
24 - 28 OCT 2016
MADRID, SPAIN

The usual suspects

- Remote peers tend to peer remotely at multiple IXPs

Autonomous System	Location	Remote presences
AS20485 (TransTelekom)	Moscow RU	AMS-IX, DE-CIX, LINX, France-IX, PLIX, Equinix Ashburn
AS8262 (Evolink)	Sofia BG	AMS-IX, DE-CIX, LINX, France-IX
AS31042 (Serbia Broadband)	Belgrad RS	AMS-IX, DE-CIX, LINX, France-IX
AS7713 (Telin)	Hong Kong HK	AMS-IX, DE-CIX, LINX, Any2 LA
AS52320 (GlobeNet)	Miami FL US	AMS-IX, DE-CIX, LINX, Equinix Ashburn
AS12578 (LatTelecom)	Riga LV	AMS-IX, DE-CIX, LINX,MSK-IX
AS1267 (Wind)	Milan IT	AMS-IX, DE-CIX, LINX
AS8866 (VivaComm)	Sofia BG	AMS-IX, DE-CIX, LINX
AS45352 (IPSERVERONE)	Singapore SG	AMS-IX, DE-CIX, LINX
AS6866 (CYTA)	Nicosia CY	AMS-IX, DE-CIX, LINX

Interpreting the facility information provided in PeeringDB

- Detecting remote peering at IXPs provides an indirect way to interpret the facility information provided by ASes in PeeringDB
- Percentage of remote peers claiming to have local presence at the IXP:
  - AMS-IX: 16%
  - LINX: 20%
  - DE-CIX: 25%
- ASes may record facility presence not to indicate tenancy but their availability for private interconnections over the facility.
- ASes may record inaccurate information by mistake or to appear more appealing peers.

Where are the city-level locations of the remote peers?

Presence-informed RTT Geo-location

- Most of the available accurate geo-location methods can resolve only a subset of the remote peering IPs:
  - **OpenIP Map**: self-reported data, covers only a subset of the IPs
  - **DNS-based geo-location**: cannot be applied to addresses without reverse DNS record
- Other geo-location methodologies not available or too error prone:
  - **Trilateration**: high complexity, errors for regions with many large metro areas close to each other (e.g. West/Central Europe).
  - **Geo-location databases**: Very inaccurate for router geo-location

- **Key intuition**: reduce the problem space by exploiting the fact that the IPs of a given AS can be where the AS has presence.

Presence-informed RTT Geolocation: Methodology

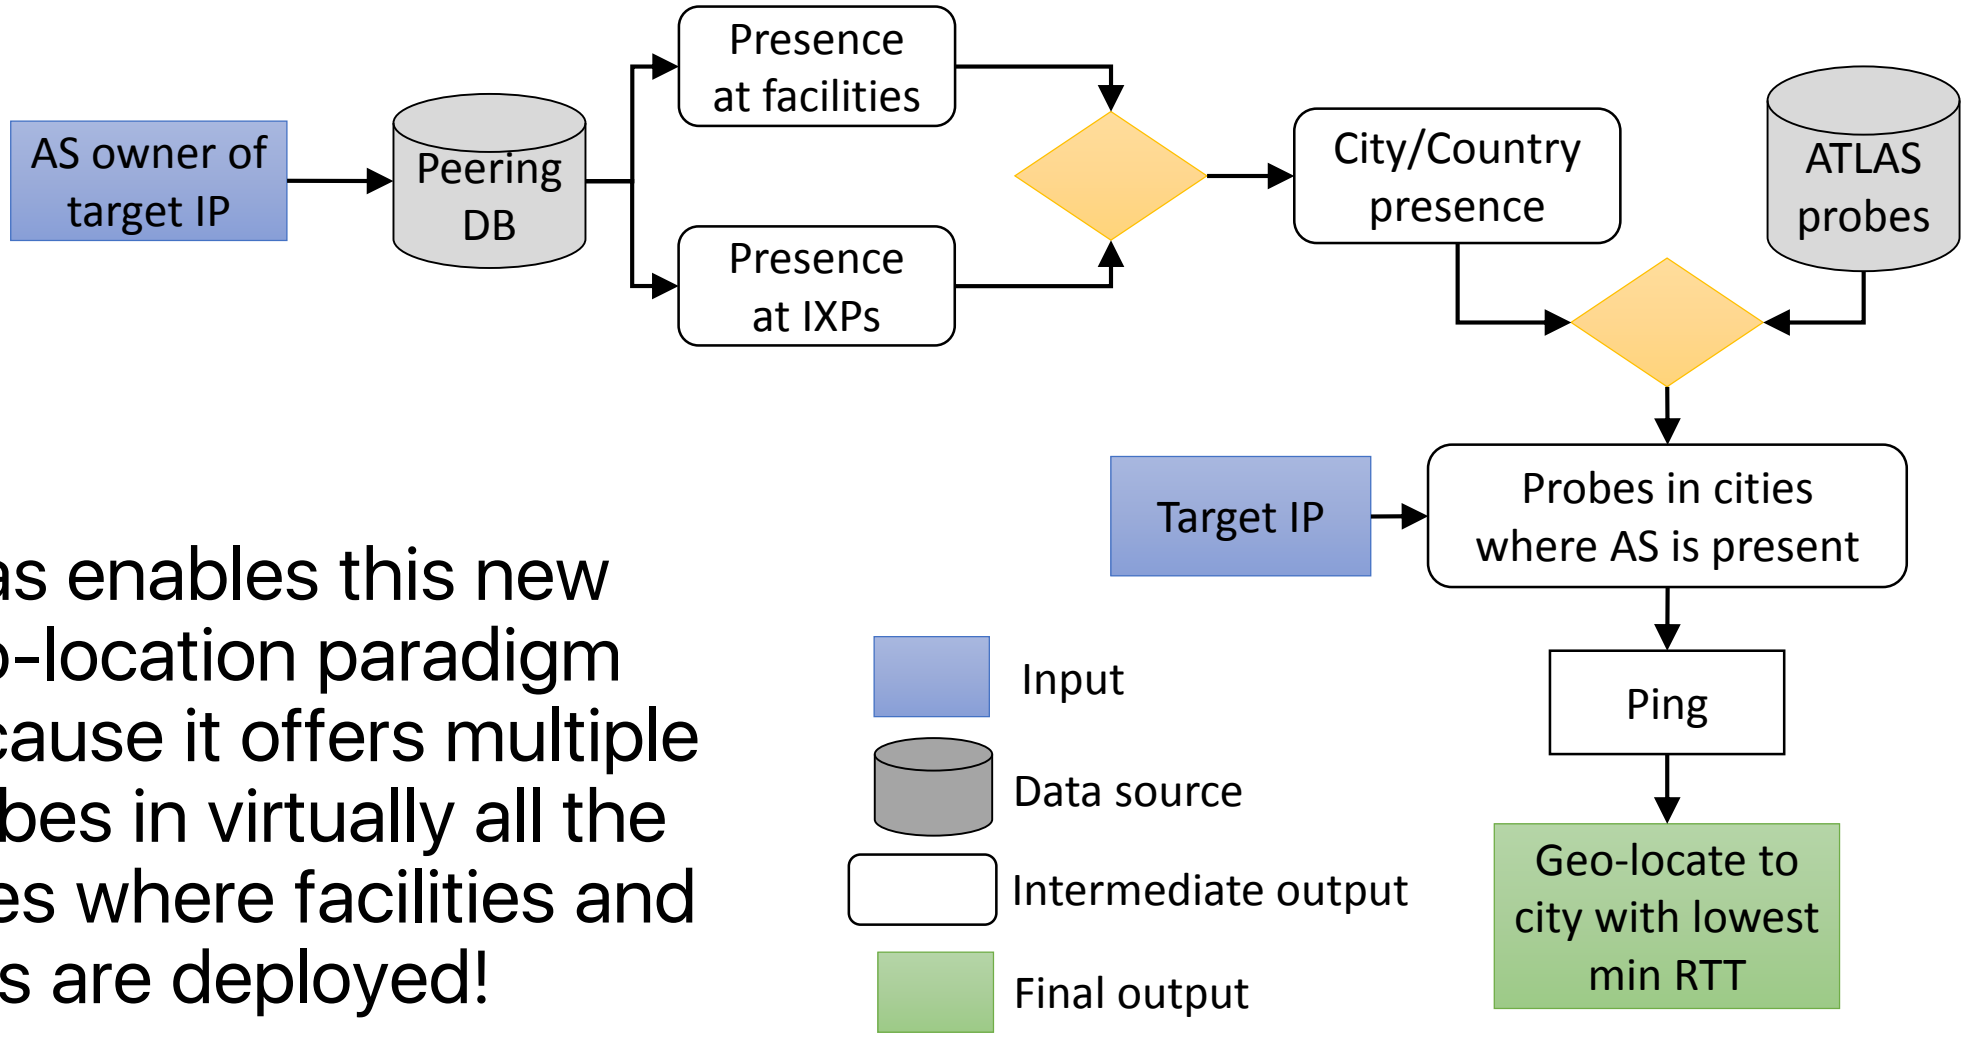

Atlas enables this new geo-location paradigm because it offers multiple probes in virtually all the cities where facilities and IXPs are deployed!

Presence-informed RTT Geolocation: Example

✱[4] Presence-informed RTT geolocation

General Information | Probes | Map | Results

Probe	ASN (IPv4)	ASN (IPv6)	Time (UTC)	RTT	Packe
6198	52048	52048	2016-10-23 02:27	1.900	0.0%
4005	31519		2016-10-23 02:27	3.513	0.0%
14175	24651		2016-10-23 02:27	3.536	0.0%
21477	43956		2016-10-23 02:27	4.309	0.0%
21252	8285		2016-10-23 02:27	7.613	0.0%
21458	56588				0.0%
20035	29518				0.0%
3375	3249		2016-10-23 02:27	11.724	0.0%
1333	8473		2016-10-23 02:27	14.696	0.0%
13982	60415		2016-10-23 02:27	15.092	0.0%
28562	35807		2016-10-23 02:27	15.931	0.0%
2583	39651		2016-10-23 02:27	17.523	0.0%
11443	3216		2016-10-23 02:27	20.879	0.0%
12846	13110		2016-10-23 02:27	21.413	0.0%
19936	2586		2016-10-23 02:27	21.564	0.0%
20059	8334	6939	2016-10-23 02:27	22.586	0.0%
10234	44050		2016-10-23 02:27	23.870	0.0%
13251	31514		2016-10-23 02:27	24.439	0.0%
20904	42148		2016-10-23 02:27	25.559	0.0%
1223	39102		2016-10-23 02:27	25.645	0.0%
24588	3301	1299	2016-10-23 02:27	26.966	0.0%
27644	34867		2016-10-23 02:27	28.429	0.0%
11629	15683		2016-10-23 02:27	30.324	0.0%
4534	2586		2016-10-23 02:27	30.663	0.0%

Geolocation: Riga, Latvia

26742	61156		2016-10-23 02:27	32.504	0.0%
25647	42610		2016-10-23 02:27	33.462	0.0%
20190	58054		2016-10-23 02:27	33.487	0.0%
25780	58010	58010	2016-10-23 02:27	33.642	0.0%
4012	15895		2016-10-23 02:27	35.392	0.0%
25207	20853	8664	2016-10-23 02:27	35.548	0.0%
14746	2119		2016-10-23 02:27	36.219	0.0%
6019	3333	3333	2016-10-23 02:27	37.852	0.0%
12638	42610		2016-10-23 02:27	38.706	0.0%
6220	34106	34106	2016-10-23 02:28	38.994	0.0%
25626	37100	37100	2016-10-23 02:27	40.655	0.0%
27207	43139		2016-10-23 02:27	42.506	0.0%
2296	1103	1103	2016-10-23 02:27	45.207	0.0%
4378	28792		2016-10-23 02:27	46.497	0.0%
3622	50343	50343	2016-10-23 02:27	47.348	0.0%
325	6830		2016-10-23 02:27	49.091	0.0%
14639	13037		2016-10-23 02:27	50.998	0.0%
17491	29535		2016-10-23 02:27	51.210	0.0%
2177	20712	20712	2016-10-23 02:27	53.083	0.0%
246	6830		2016-10-23 02:27	53.748	0.0%
16736	5617	5617	2016-10-23 02:27	56.662	0.0%
1343	3320		2016-10-23 02:27	57.640	0.0%
28796	34156		2016-10-23 02:27	64.494	0.0%
3193	5089		2016-10-23 02:27	67.737	0.0%
28569	5617	43447	2016-10-23 02:27	70.004	0.0%
2114	6830		2016-10-23 02:27		

Presence-informed RTT Geolocation: Example

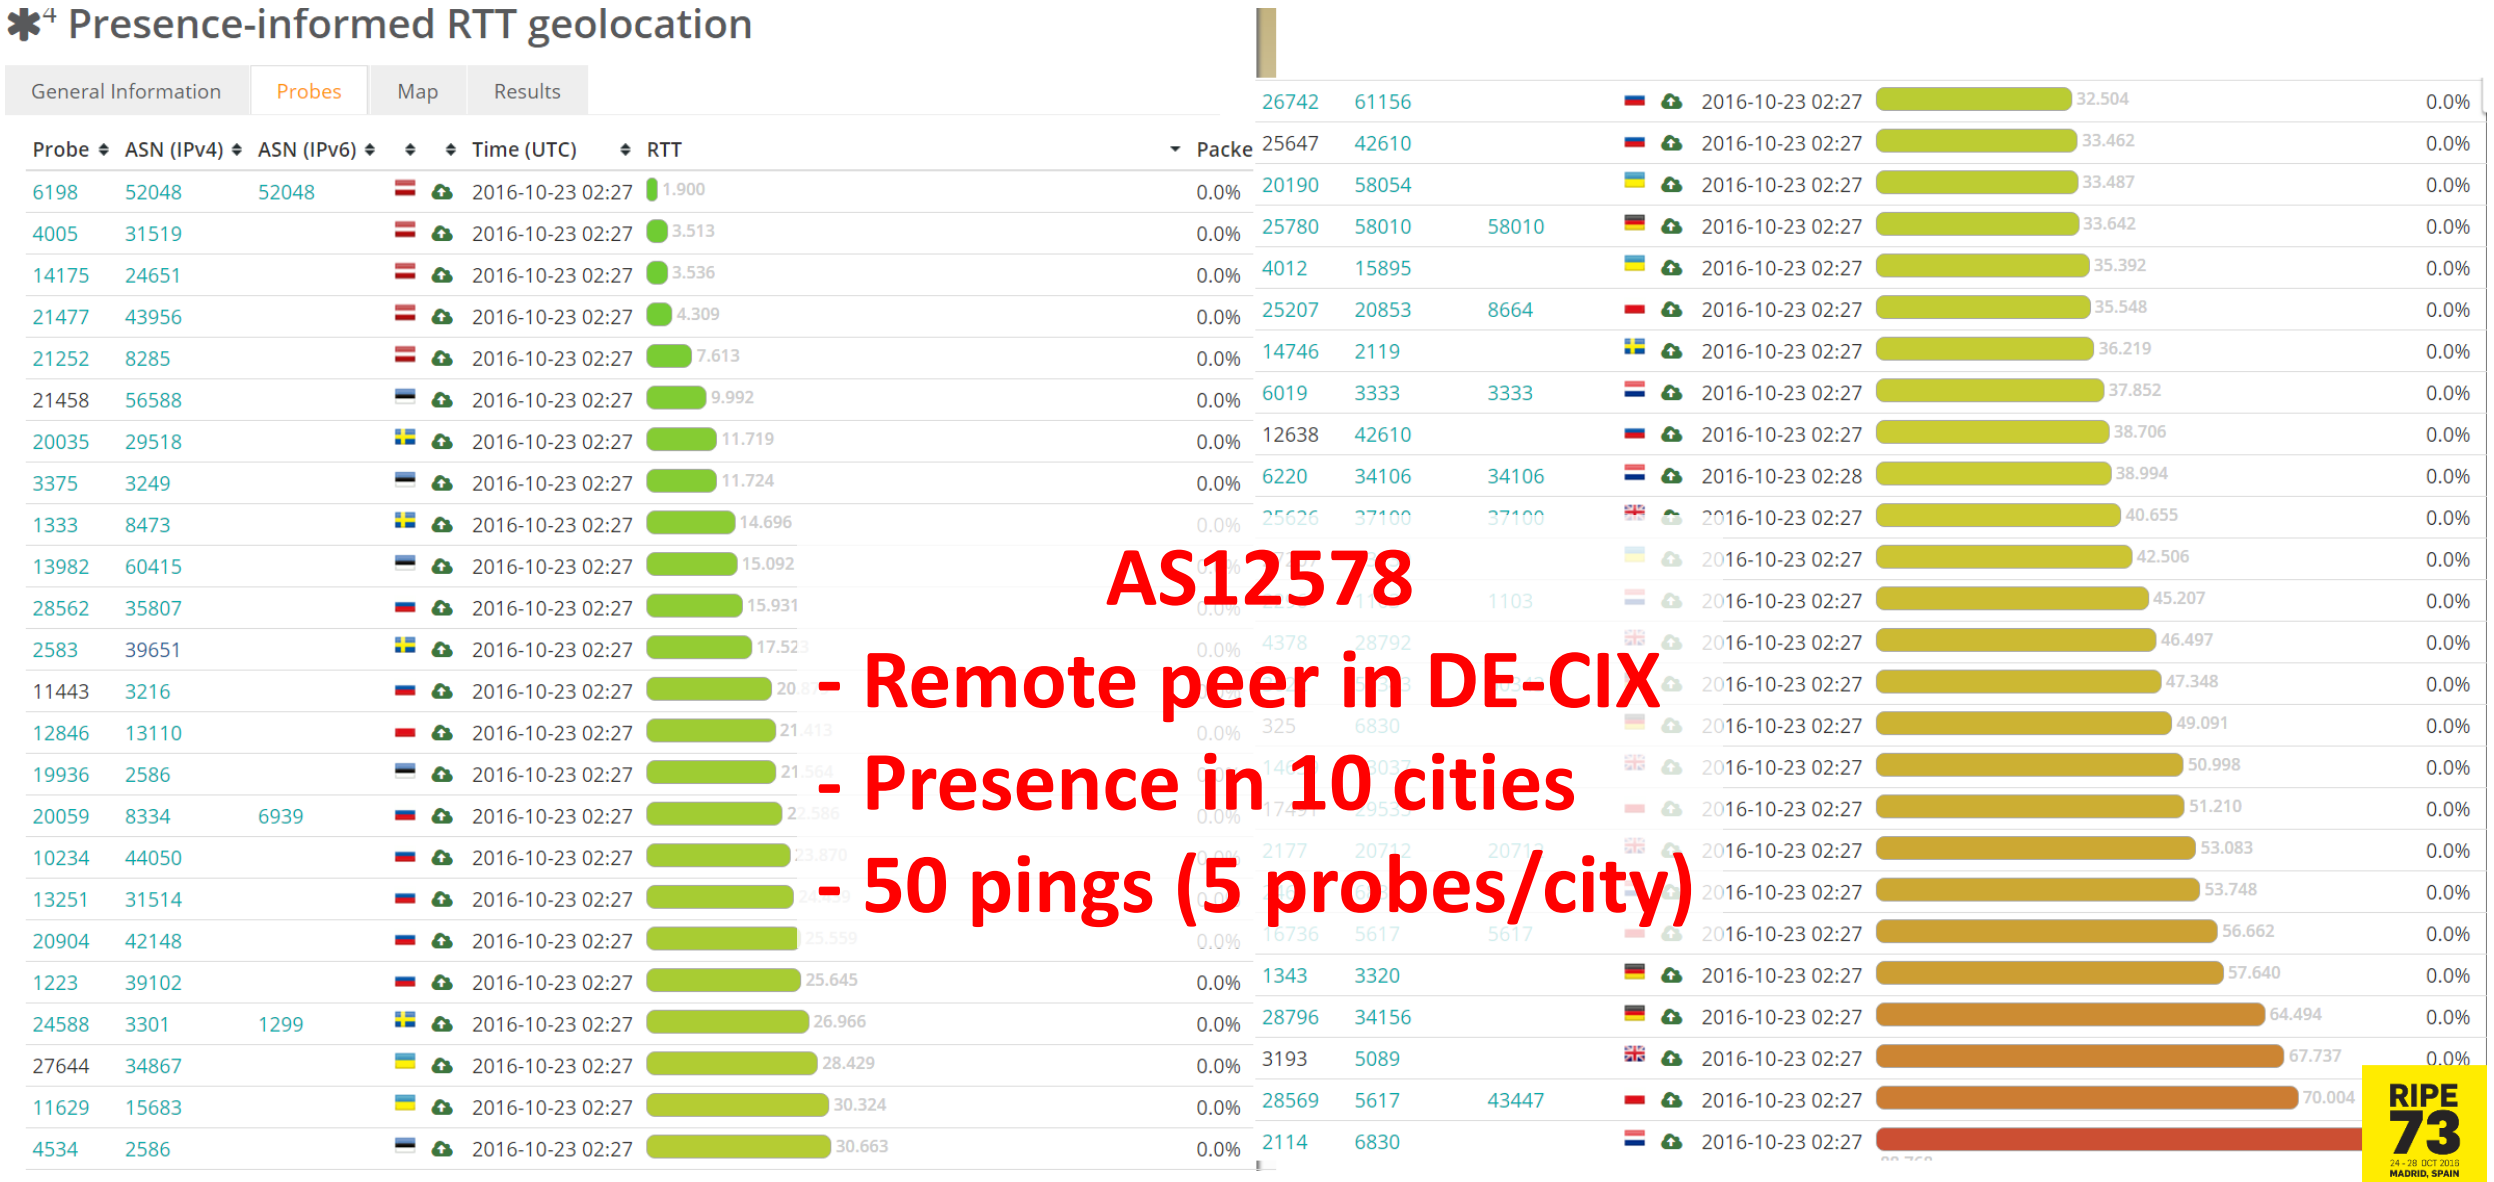

AS12578

- Remote peer in DE-CIX

- Presence in 10 cities

- 50 pings (5 probes/city)

RIPE 73
24 - 28 OCT 2016
MADRID, SPAIN

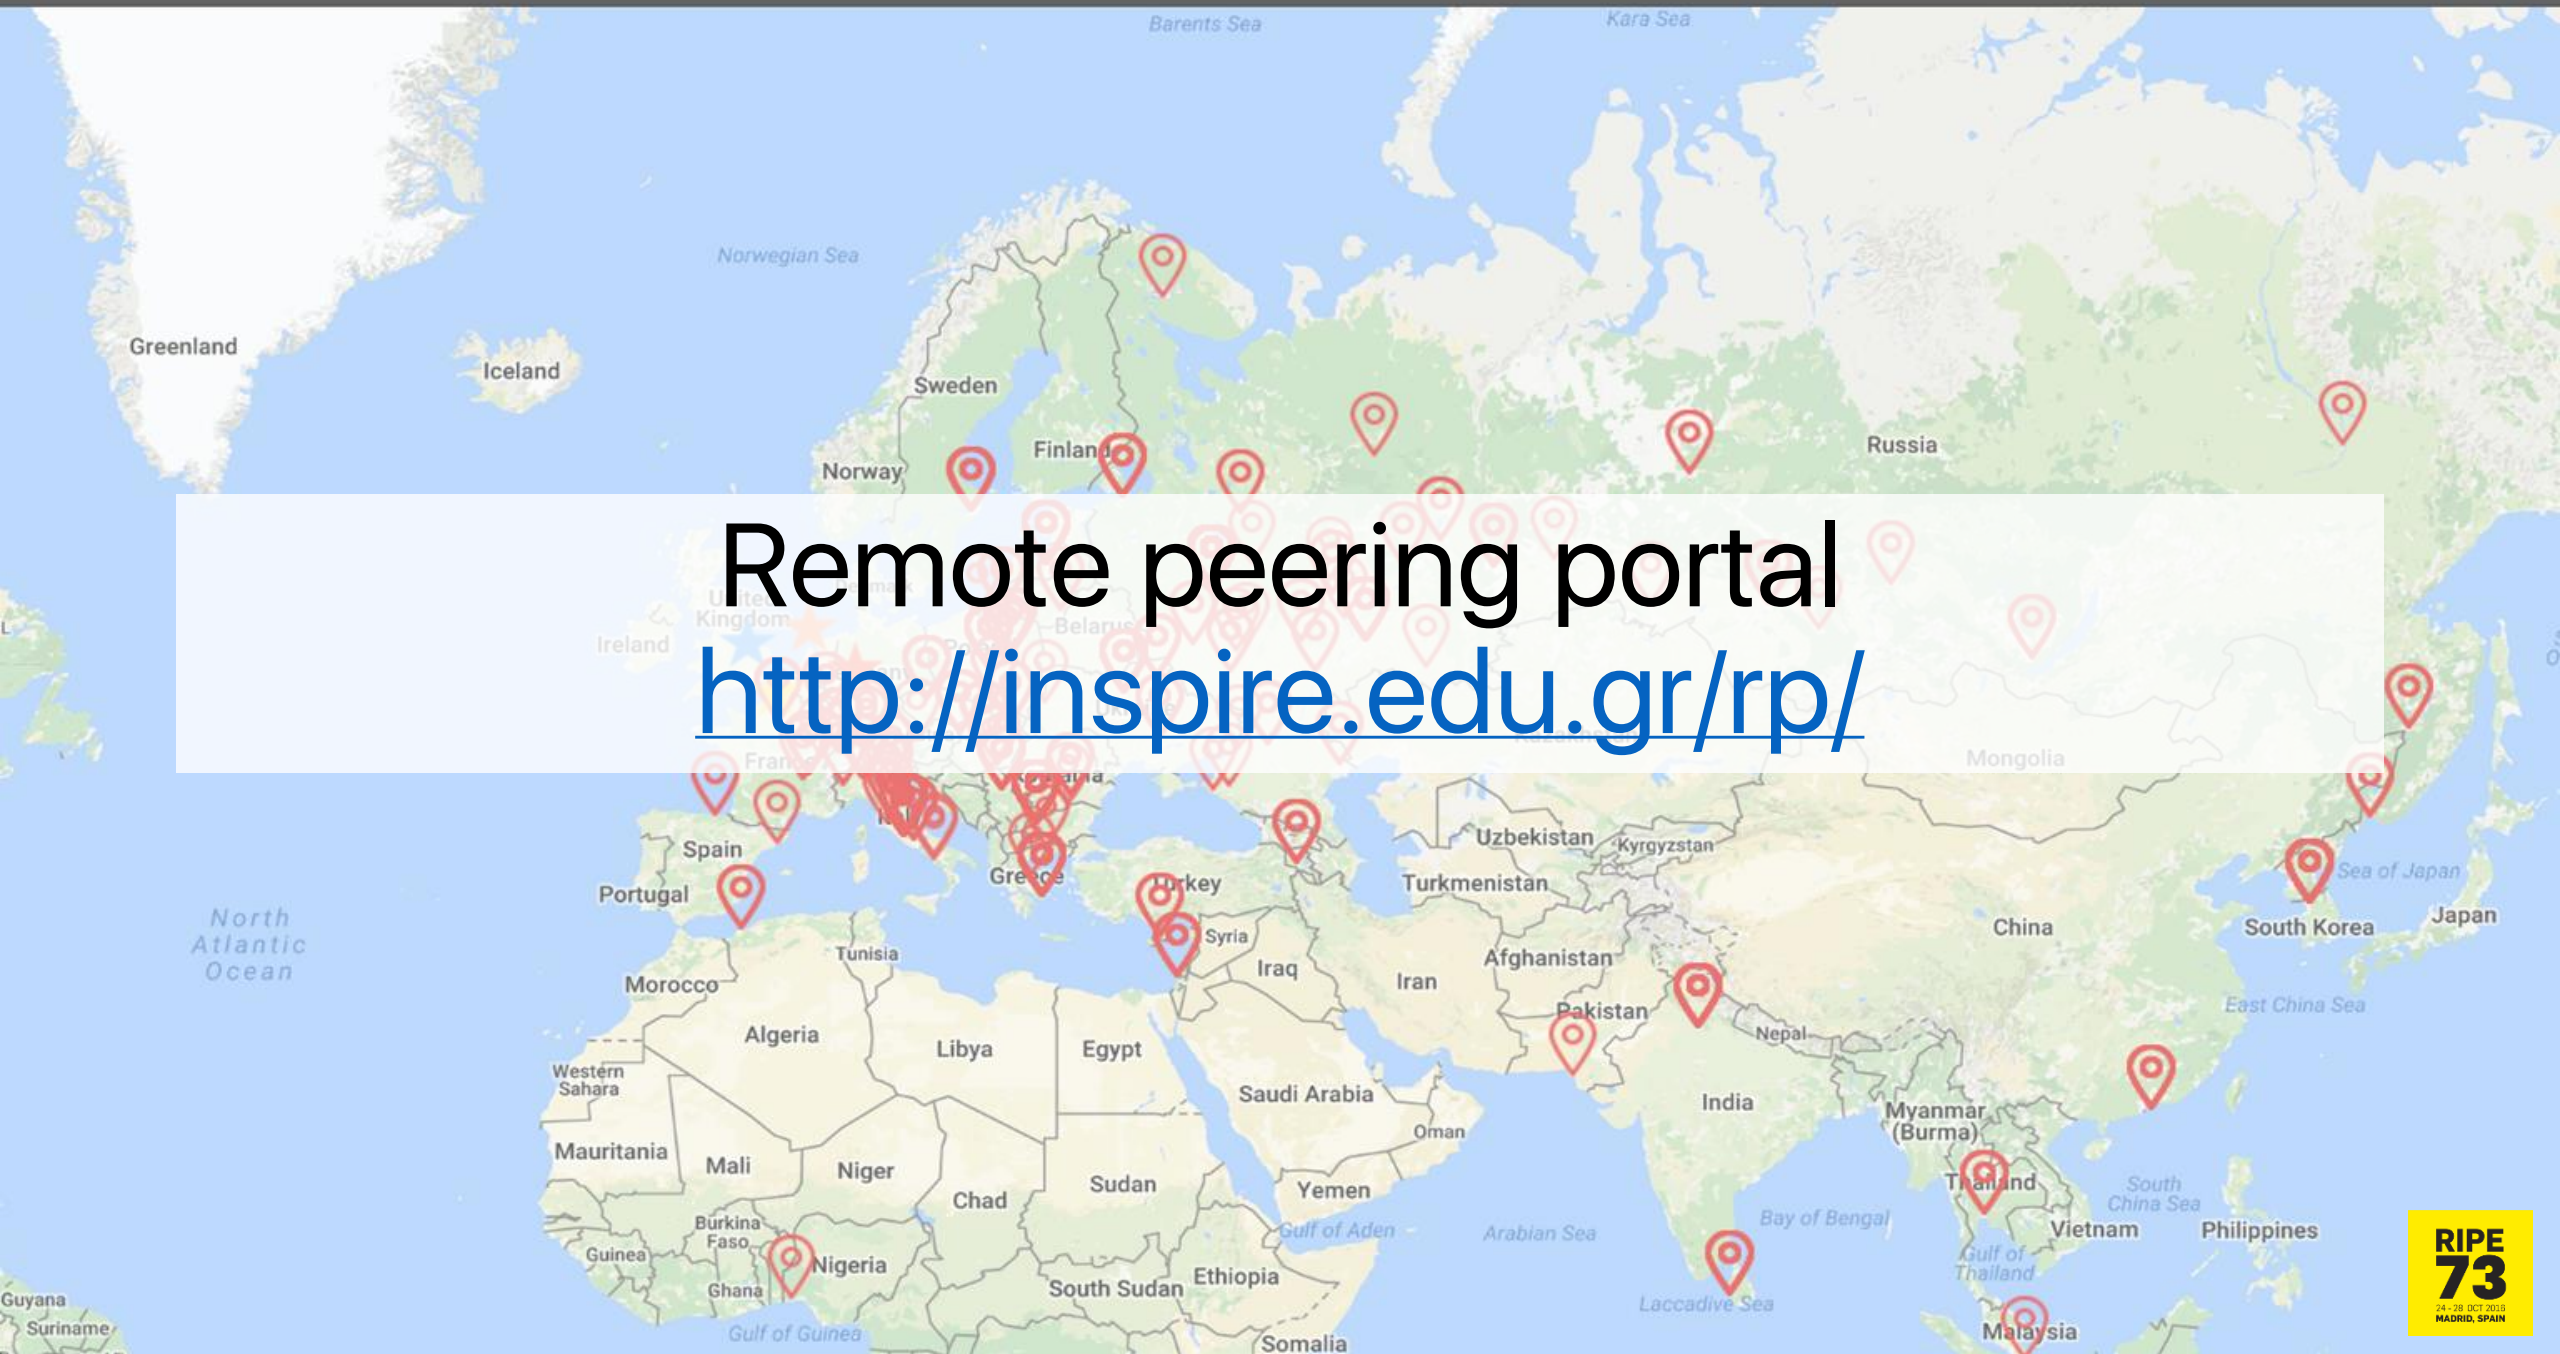

Remote peering portal

http://inspire.edu.gr/rp/

RIPE 73
24 - 28 OCT 2016
MADRID, SPAIN