RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Large Scale Internet Measurement Infrastructures

Challenges and Opportunities

Emile Aben | 2018-11 | Internet Conference

Key Internet Statistics

- Size: 510,072,000 km2
- Population: 4,208,571,287 (55.2% of world)
- Addresses:
  - IPv4: 4,294,967,296 (many unusable)
  - IPv6: 340,282,366,920,938,463,463,374,607,431,768,211,456
- Subunits:
  - Networks (ASNs): 62,408
  - IPv4 prefixes: 692k - 752k
  - IPv6 prefixes: 56k - 62k

How is it all Interconnected?

- Large vs small networks
  - AS4143 (106,187,232 IPv4 addresses)
  - 15,099 networks announcing a /24 (256 IPv4 addresses)
- Types of networks (user, content, transit, …)
- Means of interconnect: direct or IXP

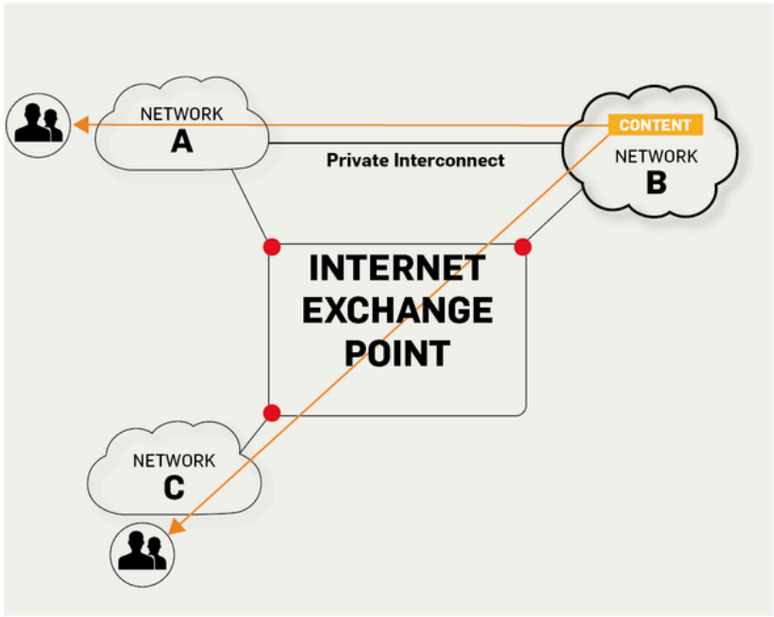

Why do we Measure?

- "That which is measured improves. That which is measured and reported improves exponentially." -- Karl Pearson

What is RIPE NCC?

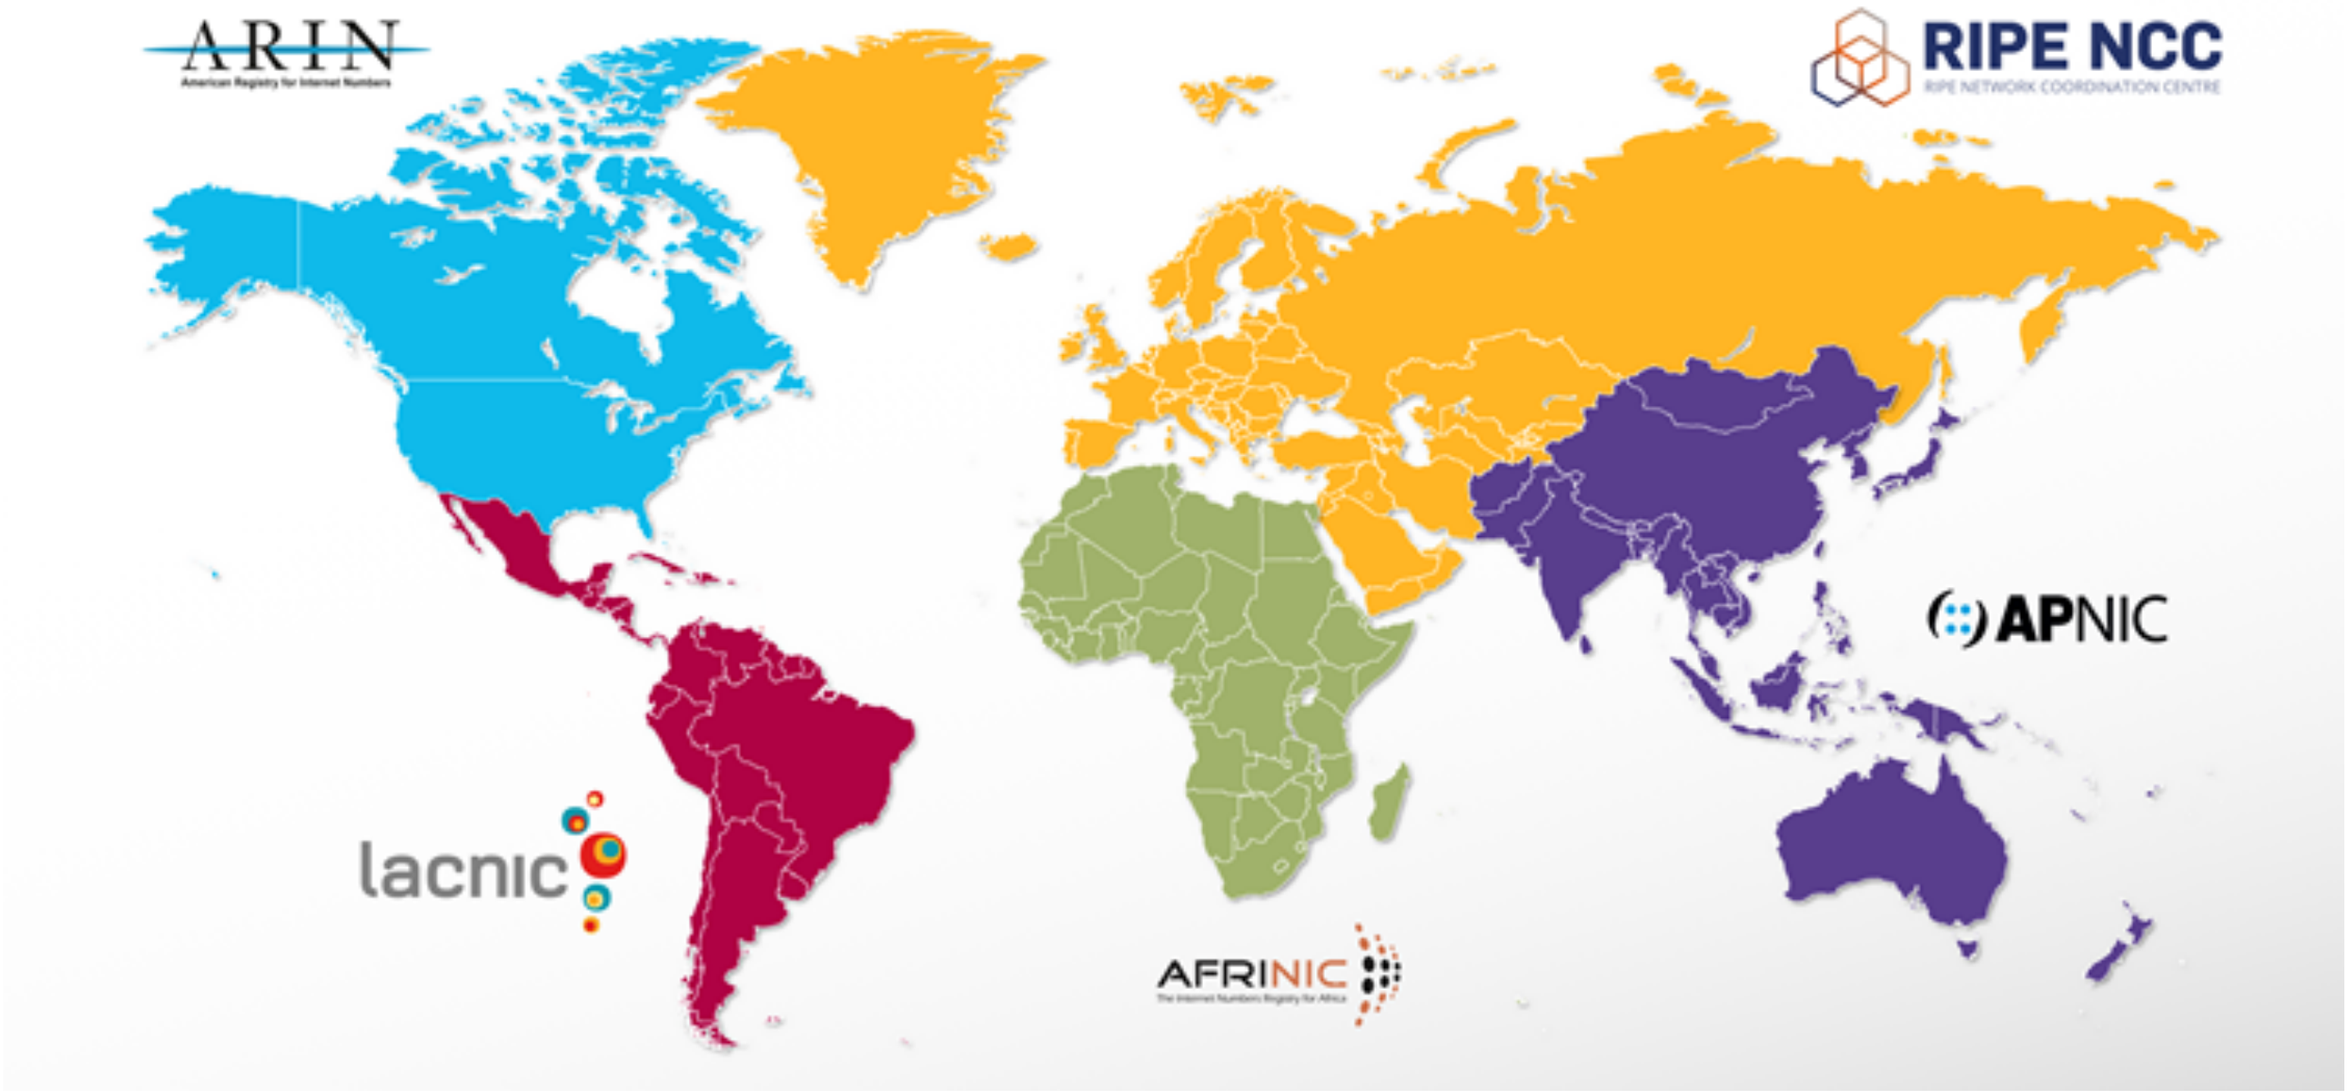

RIPE Atlas


NORTH
AMERICA
SOUTH
AMERICA
EUROPE
ASIA
AFRICA
AUSTRALIA
Atlantic
Ocean
Pacific
Ocean
Pacific
Ocean
Indian
Ocean
Leaflet | Tiles © Esri — Esri, DeLorme, NAVTEQ, TomTom, Intermap, iPC, USGS, FAO, NPS, NRCAN, GeoBase, Kadaster NL, Ordnance Survey, Esri Japan, METI, Esri China (Hong Kong), and the GIS User Community

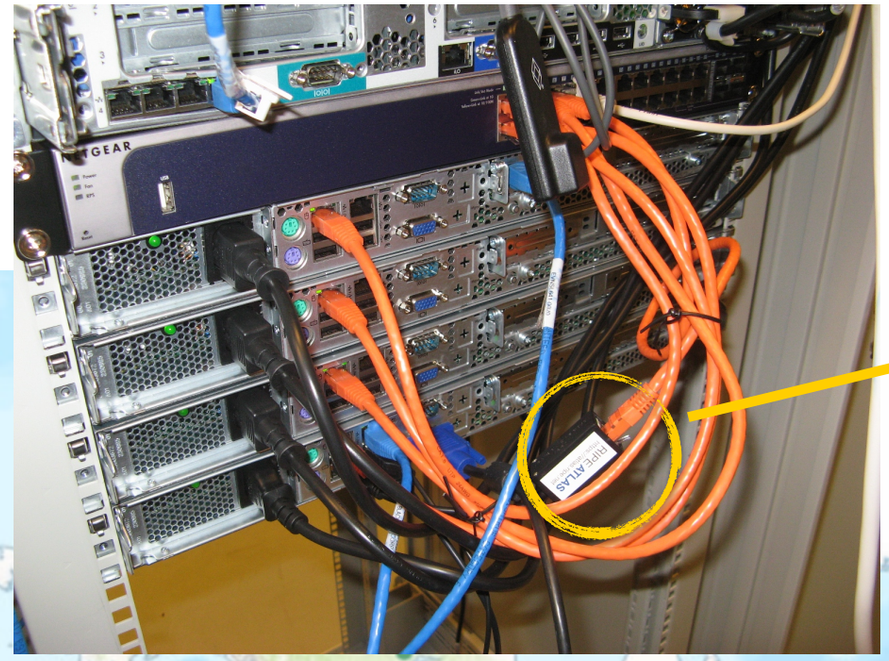
NETGEAR
RIPE ATLAS

NORTH AMERICA
SOUTH AMERICA
AFRICA
ASIA
AUSTRALIA
Atlantic Ocean
Pacific Ocean
Pacific Ocean
Indian Ocean
Leaflet | Tiles © Esri — Esri, DeLorme, NAVTEQ, TomTom, Intermap, iPC, USGS, FAO, NPS, NRCAN, GeoBase, Kadaster NL, Ordnance Survey, Esri Japan, METI, Esri China (Hong Kong), and the GIS User Community

RIPE Atlas

RIPE Atlas – Wikipedia, the free encyclopedia

https en.wikipedia.org/wiki/RIPE_Atlas

Reader

Apple iCloud Facebook Twitter Wikipedia Yahoo! News Popular

Becha 0 Talk Sandbox Preferences Beta Watchlist Contributions Log out

WIKIPEDIA
The Free Encyclopedia

Article Talk Read Edit source Edit More Search

Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia
Wikipedia store

Interaction
Help
About Wikipedia
Community portal
Recent changes
Contact page

RIPE Atlas

From Wikipedia, the free encyclopedia

RIPE Atlas is a global, open, distributed Internet measurement platform, consisting of thousands of measurement devices that measure Internet connectivity in real time.

Contents [hide]

1 History
2 Technical details
3 Community
4 Research papers
5 Similar projects
6 References
7 External links
8 Categories

What is RIPE Atlas?

- Goal: Improve Internet through measurements

- Probes hosted by volunteers

"For the community, by the community"

- Data publicly available

https://atlas.ripe.net/

RIPE Atlas Probes

- Regular probes (version 1,2,3)
  - Small form factor boxes

- RIPE Atlas anchors
  - 1U servers (Soekris)

- New: Virtual Anchors

Most Popular Features

- Six types of measurements: ping, traceroute, DNS, SSL/TLS, NTP and HTTP (to anchors)
- APIs and CLI tools to start measurements and get results
- Built-in Measurements
  - DNS Root servers
  - Probes to Anchors
- User-Defined Measurements
  - Probe owners can schedule, not only on their own probes!

RIPE RIS

Internet “Air Traffic Control”

- BGP makes Interdomain routing work
- Border routers: Routers that receive explicit routing information for all of the Internet
- How do we monitor this “air traffic control”?
- Listen in on this “chatter” at interesting places in the Internet?

Routing Data (RIS)

- 20 BGP collectors
- 600+ peers (150+ “full-feed”)
- Most at IXPs

Example RRC06

RRC06 -- DIX-IE, Tokyo, Japan -- Peer List

Status	ASN	Description	Address	IPv4 prefixes	IPv6 prefixes
Up	AS42	WOODYNET-1 - WoodyNet	202.249.2.183	69	0
Up	AS2497	IIJ Internet Initiative Japan Inc.	2001:200:0:fe00::9c1:0	0	56653
Up	AS2497	IIJ Internet Initiative Japan Inc.	202.249.2.169	721553	0
Up	AS3856	PCH-AS - Packet Clearing House	202.249.2.184	1	0
Up	AS4777	APNIC-NSPIXP2-AS Asia Pacific Network Information Centre	2001:200:0:fe00::12a9:0	0	56717
Up	AS4777	APNIC-NSPIXP2-AS Asia Pacific Network Information Centre	202.249.2.20	726451	0
Up	AS25152	K-ROOT-SERVER - Reseaux IP Europeens Network Coordination Centre (RIPE NCC)	2001:200:0:fe00::6249:0	0	59300
Up	AS25152	K-ROOT-SERVER - Reseaux IP Europeens Network Coordination Centre (RIPE NCC)	202.249.2.185	735164	0

IPv4 full tables: 3
IPv6 full tables: 3
Total peerings: 8
Data source time: 2018-11-14T08:00:00

http://www.ris.ripe.net/peerlist/rrc06.shtml

Challenges and Opportunities

RIPE Atlas Bias

Collaborator: Petros Gigis, George Michaelson

Challenge: How biased is RIPE Atlas

- Opportunistic distribution of Internet vantage points
- Intuitively biased towards: netops, techie, Europe
- Can we measure this bias?
  - Identify redundancies - i.e. multiple probes with very similar view of the Internet
  - Identify where probes are missing

Challenge: RIPE Atlas Bias

Coverage

IPv4 ASNs covered
3555 (5.676%)

IPv6 ASNs covered
1439 (8.944%)

Number of countries covered
180 (91.837%)

ASN (v4)	Probes	ASN (v6)	Probes	Prefix (v4)	Probes	Prefix (v6)	Probes	Country	Probes
3320	341	6939	250	73.0.0.0/8	130	2001:470::/32	249	Germany	1411
7922	341	3320	247	79.192.0.0/10	61	2003::/19	247	United States of America	1131
6830	317	7922	194	87.128.0.0/10	52	2601::/20	174	France	857
3215	223	3215	138	84.128.0.0/10	50	2a01:e00::/26	100	United Kingdom	597
12322	210	12322	116	78.192.0.0/11	46	2002::/16	45	Netherlands	535
3265	103	3265	96	93.192.0.0/10	45	2a02:908::/33	38	Russia	463
701	94	6830	73	91.0.0.0/10	44	2001:980::/29	36	Italy	288
31334	87	8881	57	83.160.0.0/14	36	2001:8b0::/32	35	Switzerland	273
33915	81	31334	51	88.176.0.0/12	35	2001:980::/32	34	Czechia	259
5089	78	20712	37	109.190.0.0/16	34	2a02:a000::/26	34	Canada	222

https://atlas.ripe.net/results/maps/network-coverage/

Opportunity: Probes vs Eyeballs

Select Date: 18/11/2018

Showing data for **18/11/2018**

Details for : **Japan (JP)** | View Japan on RIPEstat

87.85

Total Internet Users: **115111595**

Internet Users in networks with RIPE Atlas probes: **101125674**

Internet users coverage is estimated using percentage of IPv4 Public probes.

IPv4 Public Probes >= 3 | 3 > IPv4 Public Probes > 1

Search:

Network (ASN)	Network Name	Estimated User Population % ▼	IPv4 Public Probes	IPv4 Private Probes	IPv4 Total Probes	IPv6 Public Probes	IPv6 Private Probes	IPv6 Total Probes	More
4713	OCN	18.74	8	9	17	4	2	6	View
2516	KDDI	16.68	15	1	16	15	1	16	View
17676	GIGAINFRA	15.95	11	2	13	8	2	10	View
2527	SO-NET	4.46	9	2	11	6	1	7	View
9605	DOCOMO	4.19	2	0	2	0	0	0	View
17511	K-OPTICOM	2.91	3	0	3	1	0	1	View
9824	JTCL-JP-AS	2.85	3	1	4	0	0	0	View
2518	BIGLOBE	2.72	2	1	3	1	1	2	View
17506	UCOM	2.32	6	1	7	0	0	0	View
2497	IIJ	2.15	13	2	15	2	1	3	View
2519	VECTANT	1.96	1	0	1	0	0	0	View
2510	INFOWEB	1.86	1	1	2	0	0	0	View
2514	INFOSPHERE	1.84	7	0	7	1	0	1	View
4685	ASAHI-NET	1.47	3	2	5	4	2	6	View
18126	CTCX	1.21	1	1	2	0	0	0	View
10010	TOKAI	1.19	3	1	4	3	0	3	View

http://sg-pub.ripe.net/petros/population_coverage/country.html?name=JP

Select Date: 18/11/2018

Showing data for **18/11/2018**

Details for : **South Korea (KR)** | View South Korea on RIPEstat

58.35

Total Internet Users: **43274132**

Internet Users in networks with RIPE Atlas probes: **25251463**

Internet users coverage is estimated using percentage of IPv4 Public probes.

IPv4 Public Probes >= 3

3 > IPv4 Public Probes > 1

Search:

Network (ASN)	Network Name	Estimated User Population %	IPv4 Public Probes	IPv4 Private Probes	IPv4 Total Probes	IPv6 Public Probes	IPv6 Private Probes	IPv6 Total Probes	More
4766	KIXS-AS-KR	41.19	2	1	3	0	0	0	View
9318	SKB-AS	17.09	2	0	2	0	0	0	View
9644	SKTELECOM-NET-AS	10.73	0	0	0	0	0	0	Apply for a probe
17858	POWERVIS-AS-KR	10.57	0	0	0	0	0	0	Apply for a probe
17853	LGTELECOM-AS-KR	6.18	0	0	0	0	0	0	Apply for a probe
3786	LGDACOM	3.89	0	0	0	0	0	0	Apply for a probe

http://sg-pub.ripe.net/petros/population_coverage/country.html?name=KR

RIPE Atlas Vantage Point Selection

Collaborators: Thomas Holtenbach, Cristel Pelsser, Randy Bush, Laurent Vanbever

Challenge: Where to Measure From?

- 10k RIPE Atlas probes
- Select by: Country, ASN

- 10 probes from Uruguay vs.10 probes from Ukraine?

- Other ways to find diversity/uniformity in measurement sources?

Probe Similarity

- How (dis)similar are RIPE Atlas vantage points?
- Goal: Optimise what you discover with as little vantage points as possible
- Means: Pairwise similarity metric

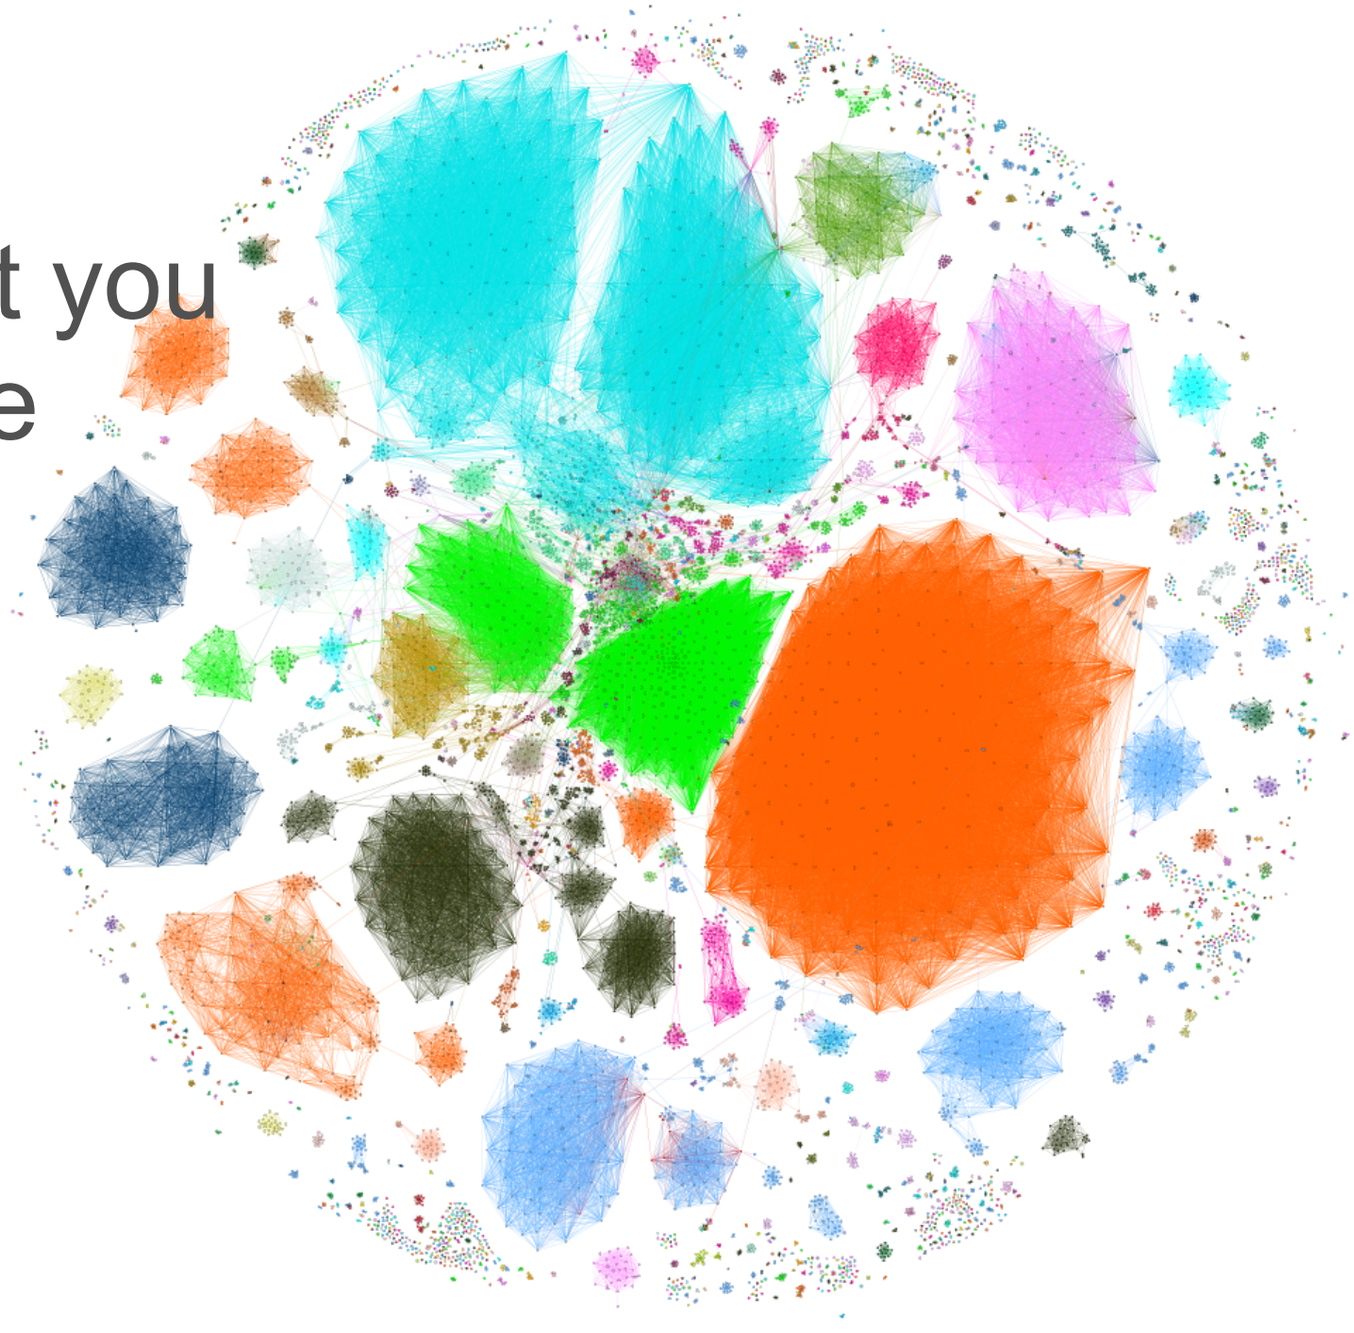

Probe Similarity

- 25% more IPs seen when picking dissimilar probes vs default

10% of probes have “identical IPv4 twin(s)”

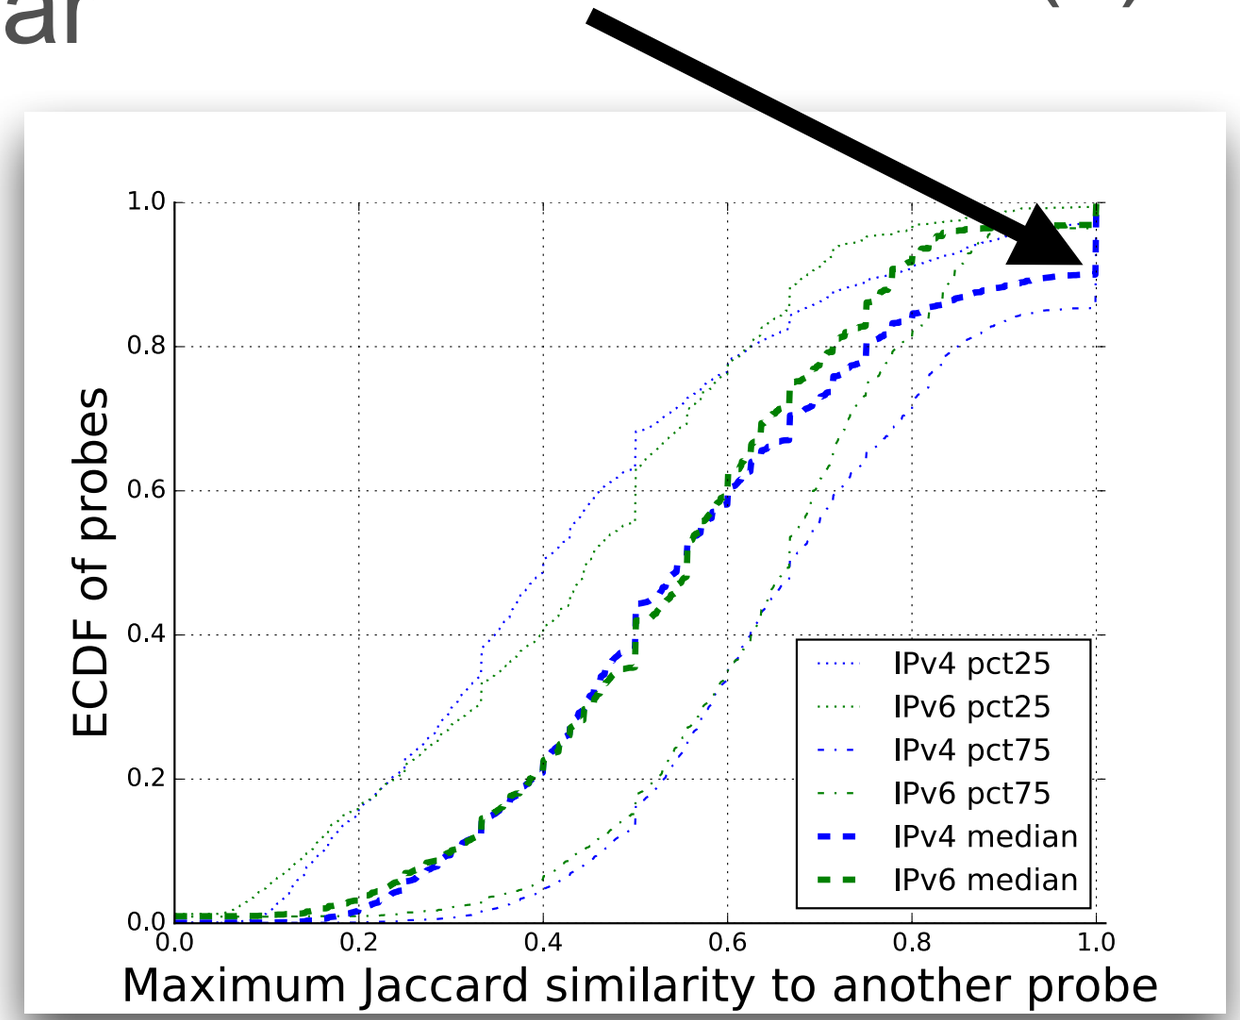

- Other similarities:
  - DNS similarity
  - RTT
  - Outages (physical infrastructure/electricity)

RIPE Atlas Bulk Data Analysis

Collaborator: Petros Gigis

Bulk Data Downloads

traceroute-2018-11-25T0000.bz2	26-Nov-2018 05:55	628M
traceroute-2018-11-25T0100.bz2	26-Nov-2018 05:53	626M
traceroute-2018-11-25T0200.bz2	26-Nov-2018 05:53	629M
traceroute-2018-11-25T0300.bz2	26-Nov-2018 05:53	626M
traceroute-2018-11-25T0400.bz2	26-Nov-2018 05:53	631M
traceroute-2018-11-25T0500.bz2	26-Nov-2018 05:53	628M

https://labs.ripe.net/Members/petros_gigis/announcing-daily-ripe-atlas-data-archives

https://data-store.ripe.net/datasets/atlas-daily-dumps

Challenge: Big Data Analysis

- Using all of RIPE Atlas traceroutes
- What's the bias?

Large IXP Outages

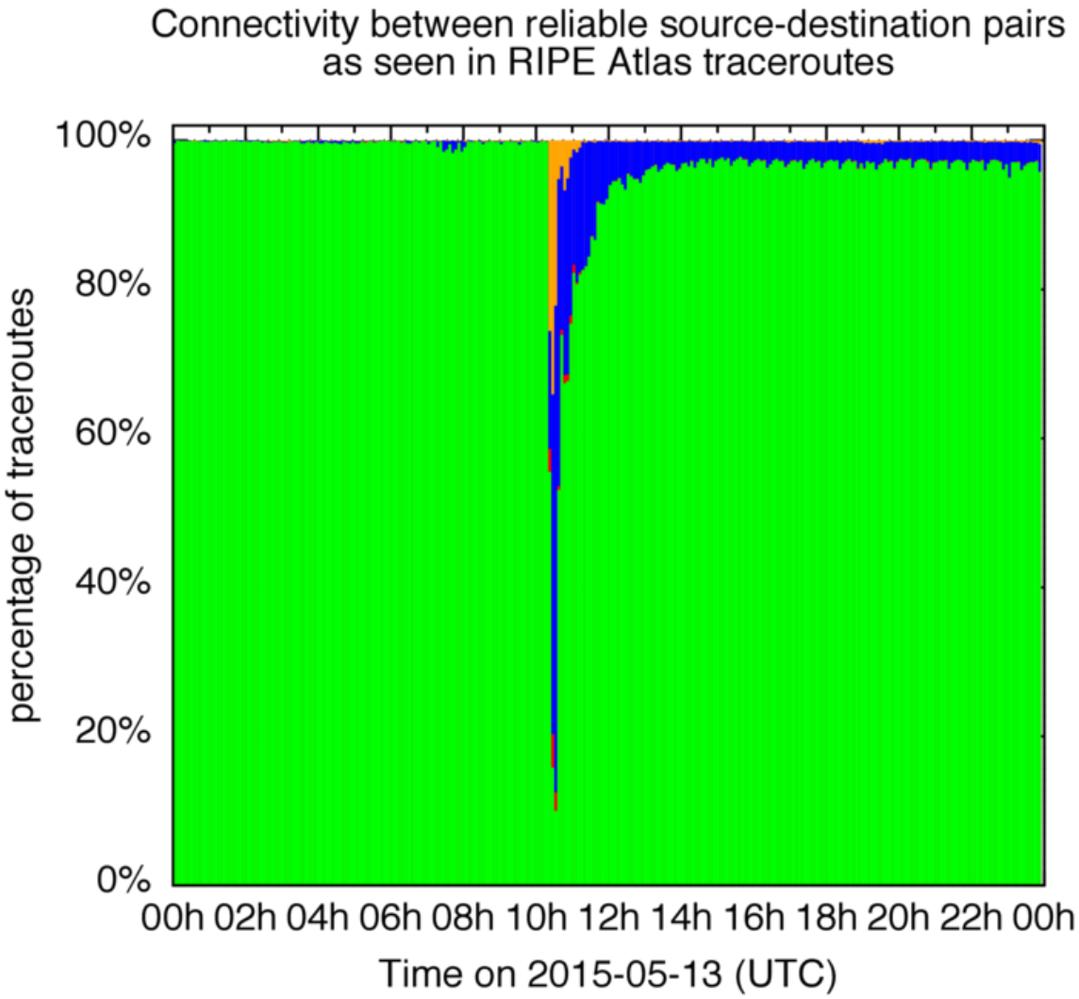

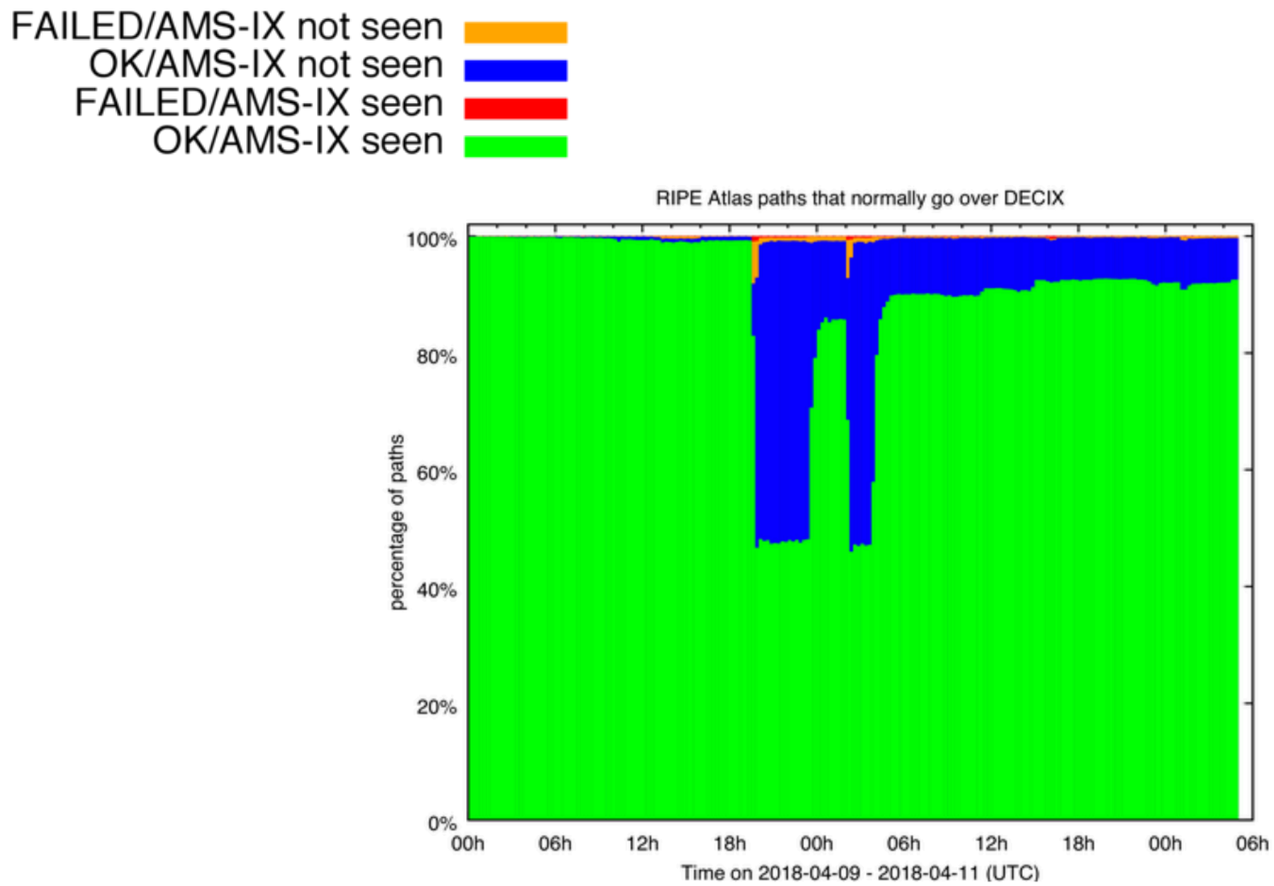

https://labs.ripe.net/Members/emileaben/does-the-internet-route-around-damage
https://labs.ripe.net/Members/emileaben/does-the-internet-route-around-damage-in-2018

Countering Bias

Infrastructure	Date	Sources (RIPE Atlas probes)	Destinations (IP addresses)	Source-Destination pairs (Internet paths)
AMS-IX	2015-05-12	3,289	753	14,431

- Useful step up from many other analysis already
  - but not enough!

- Challenge: per AS / per eyeballs / per IXP member share / Other?

RIPE Atlas: Traceroute IP Geolocation

In collaboration with CAIDA

Challenge: IP Geolocation

- Often thought of as a “solved problem”
- Edge geolocation is generally “ok”
- Infrastructure (routers, servers) IP geolocation much worse then edge

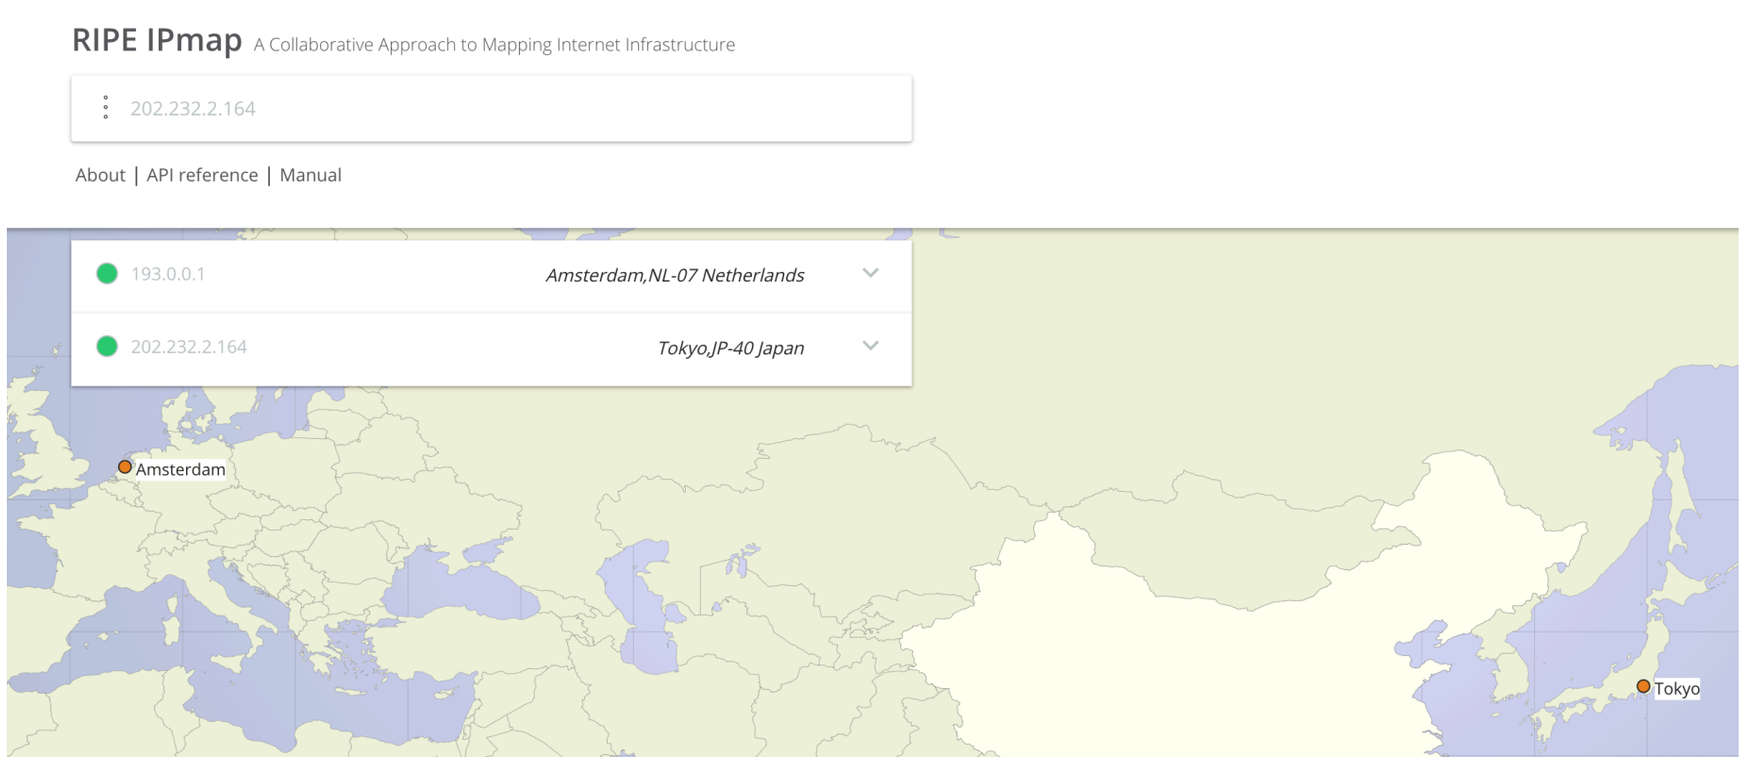

Opportunity: IPmap

- IPmap: collaborative approach to IP geolocation: https://ipmap.ripe.net/
- Multiple ‘engines’, e.g.
  - Crowdsourcing
  - RTT triangulation with RIPE Atlas
  - Anycast
- Open challenge: How to best combine multiple engine results?

RIPE Atlas: Interconnect with a Region

Collaborators: Petros Gigis, Daniele Arena, George Michaelson

Opportunity: Characterise Interconnect

- Does Internet traffic stay local?
  - default: local = within country borders
  - other: sets of countries, cities, custom sets of probes
- Do IXPs help?
- IXP-country-jedi is a measurement method and set of visualisations that provide insight

https://www.ripe.net/ixp-country-jedi

Opportunity: ixp-country-jedi

- Mesh of traceroutes within a country
- Example viz: IXP-country matrix
  - green: stays in country

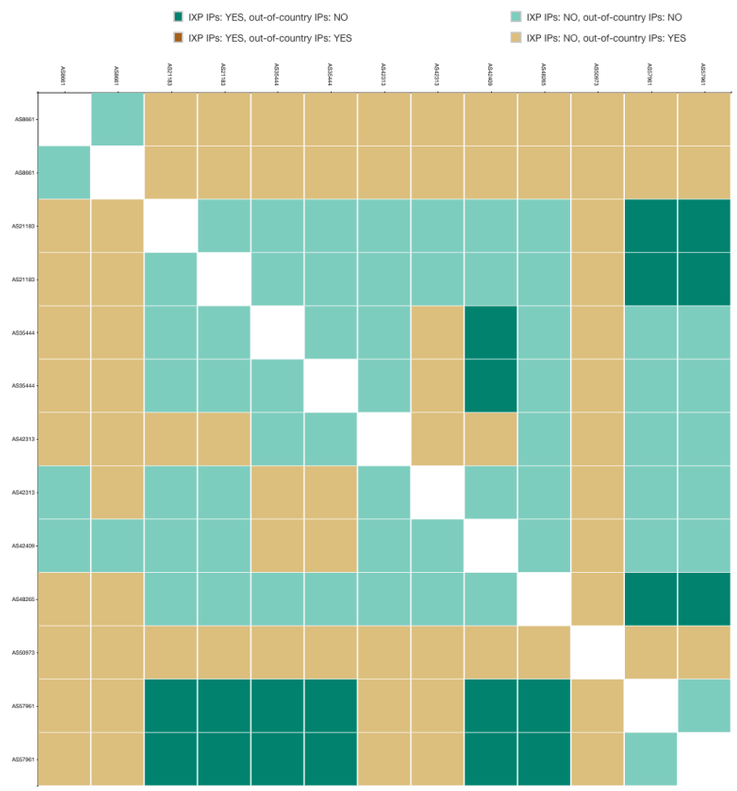

Albania

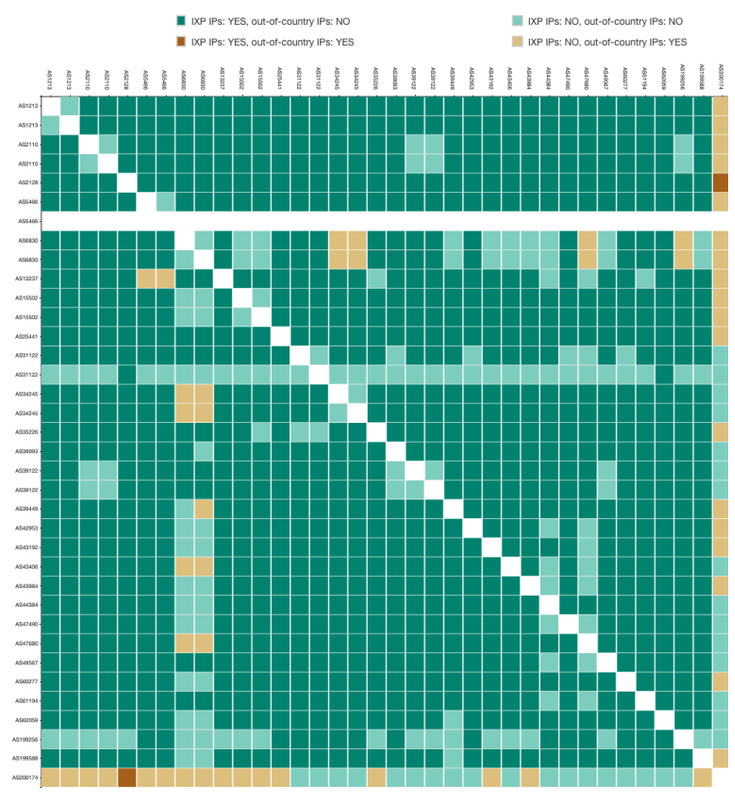

Ireland

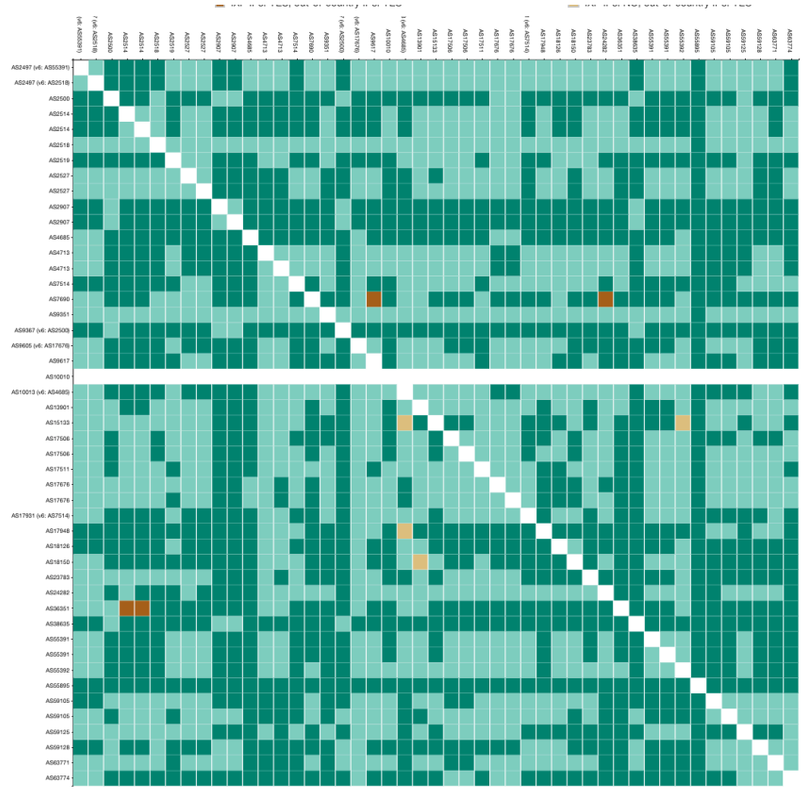

Japan

IXP-country-jedi: Geo Viz

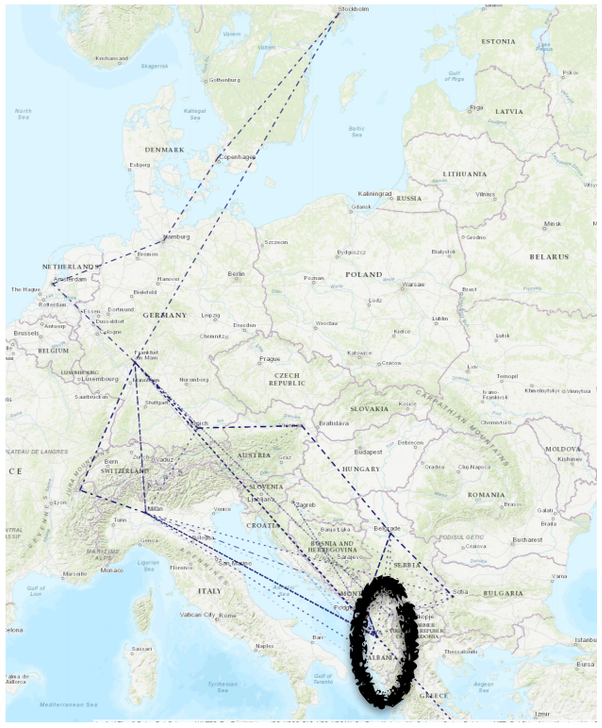

Albania

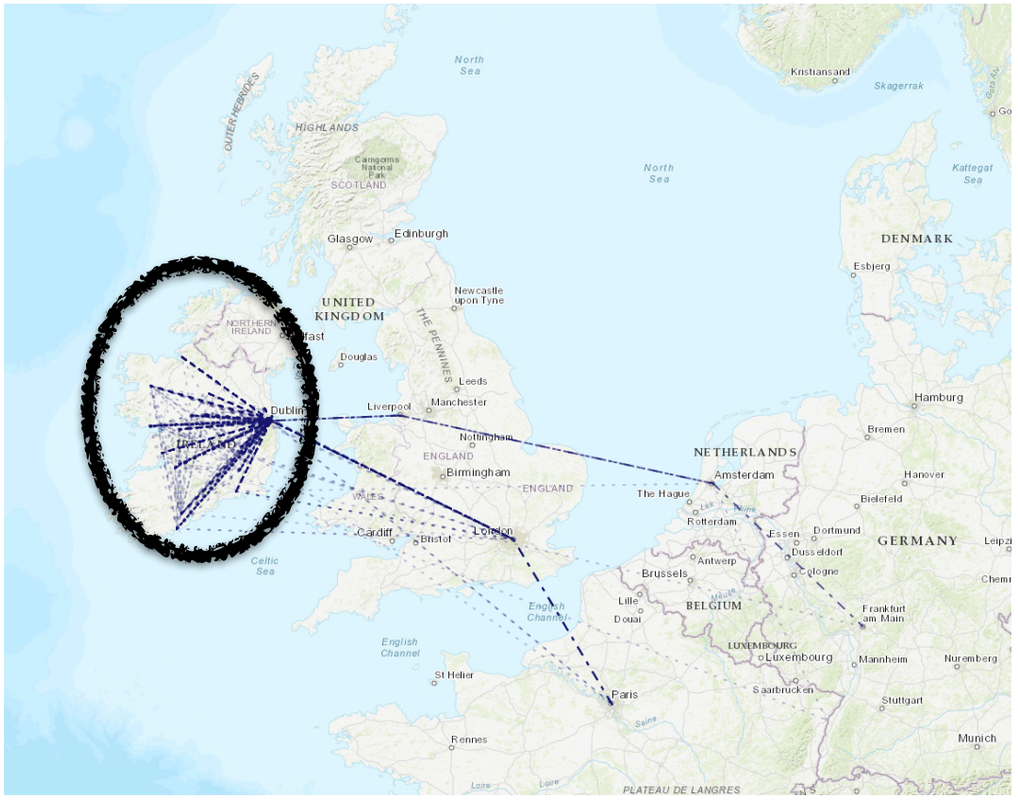

Ireland

Japan

IXP-country-jedi: Geo Viz

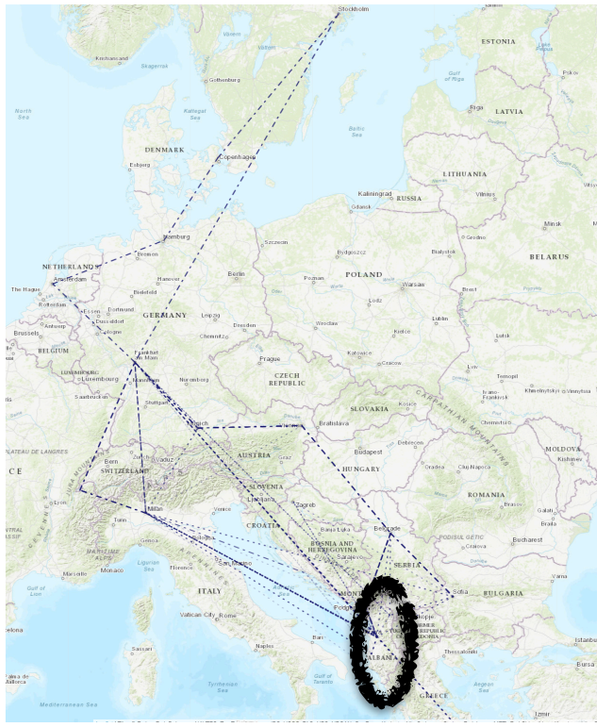

Albania

Close

IPv4 Traceroute:

AS(v4) Path:
1) AS7690
2) _BBIX Tokyo
3) AS6939
4) AS10026
5) AS9824

v4

msm_id:16422576 prb_id:33483 dst:122.197.14.15 ts:2018-10-02 11:37:03 -00:00

```
0: ()  172.16.114.1  (0.493 0.59 1.584) ||
1: (AS7690)  p219001.ip.mirai.ne.jp  (2.02 2.184 2.474) ||
2: (AS7690)  rs2-gi3.i
3: (AS7713)  218.100.6.74  (5.602 5.633 5.698) |Tokyo,T?ky?,JP|
4: (AS6939)  100ge11-1.core1.sea1.he.net  (87.003 87.184 98.605) ||
5: (AS6939)  100ge11-1.core1.sjc2.he.net  (106.135 106.187 107.661) |San Jose,California,US|
6: (AS6939)  pacnet.10gigabitethernet2-2.core1.sjc2.he.net  (112.124 114.468 132.919) |San Jose,Californi
7: (AS10026)  te-0-3-0-1.wr2.osa.asianetcom.net  (133.576 137.844 142.715) |Osaka,?saka,JP|
8: (AS*) Uknown (*) (* * *) |*|
9: (AS10026)  te0-1-0-3.gw3.osa1.10026.telstraglobal.net  (121.913 123.779) ||
10: (AS10026)  tnn-0003-tnn-0008-tnn-0014.10026.telstraglobal.net  (120.6 134.357 135.
11: (AS9824)  c3-be13.kb-dc.zaq.ad.jp  (121.
12: (AS*) Uknown (*) (* * *) |*|
13: (AS9824)  110-134-112-249.rev.home.ne.jp  (12
14: ()  10.1.12.229  (134.432 138.152 139.501) ||
15: ()  10.1.12.230  (122.333 127.849 132.449) ||
16: ()  10.1.195.239  (127.159 138.572 141.572) ||
17: (AS*) Uknown (*) (* * *) |*|
18: (AS*) Uknown (*) (* * *) |*|
19: (AS*) Uknown (*) (* * *) |*|
20: (AS*) Uknown (*) (* * *) |*|
21: (AS*) Uknown (*) (* * *) |*|
22: (AS*) Uknown (*) (* * *) |*|
```

Grey bar: max min RTT of all hops.
Light-blue bar: min RTT of the hop.

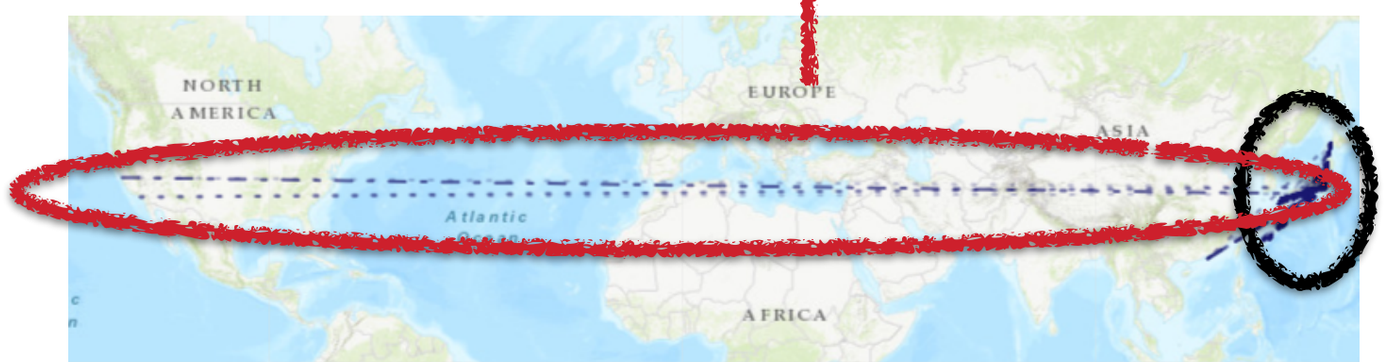

Japan

How Are End Users Interconnected?

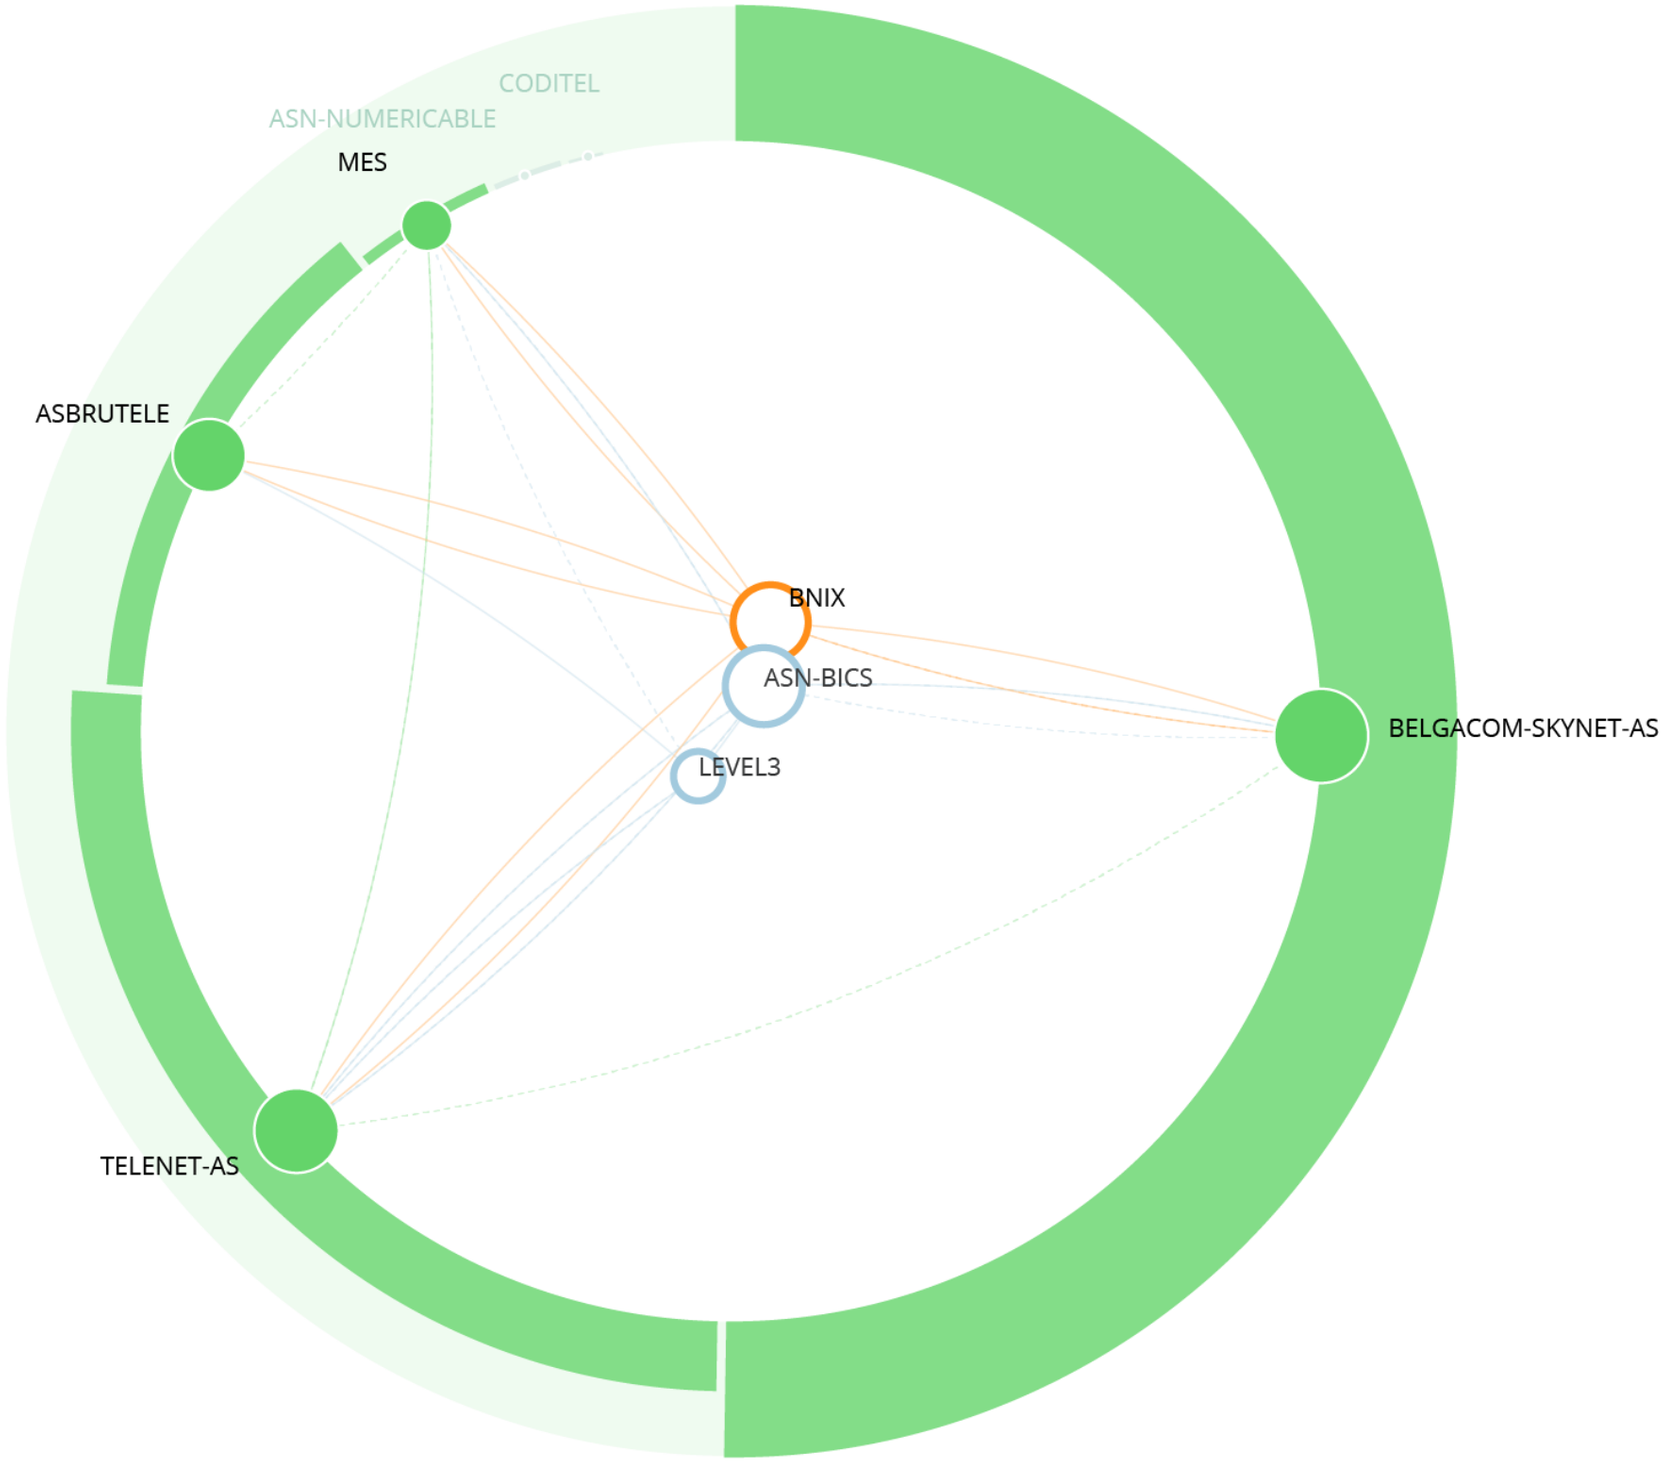

emile.aben@ripe.net | Internet Conference

http://sg-pub.ripe.net/ixp-country-jedi/

How Are End Users Interconnected?

Belgium

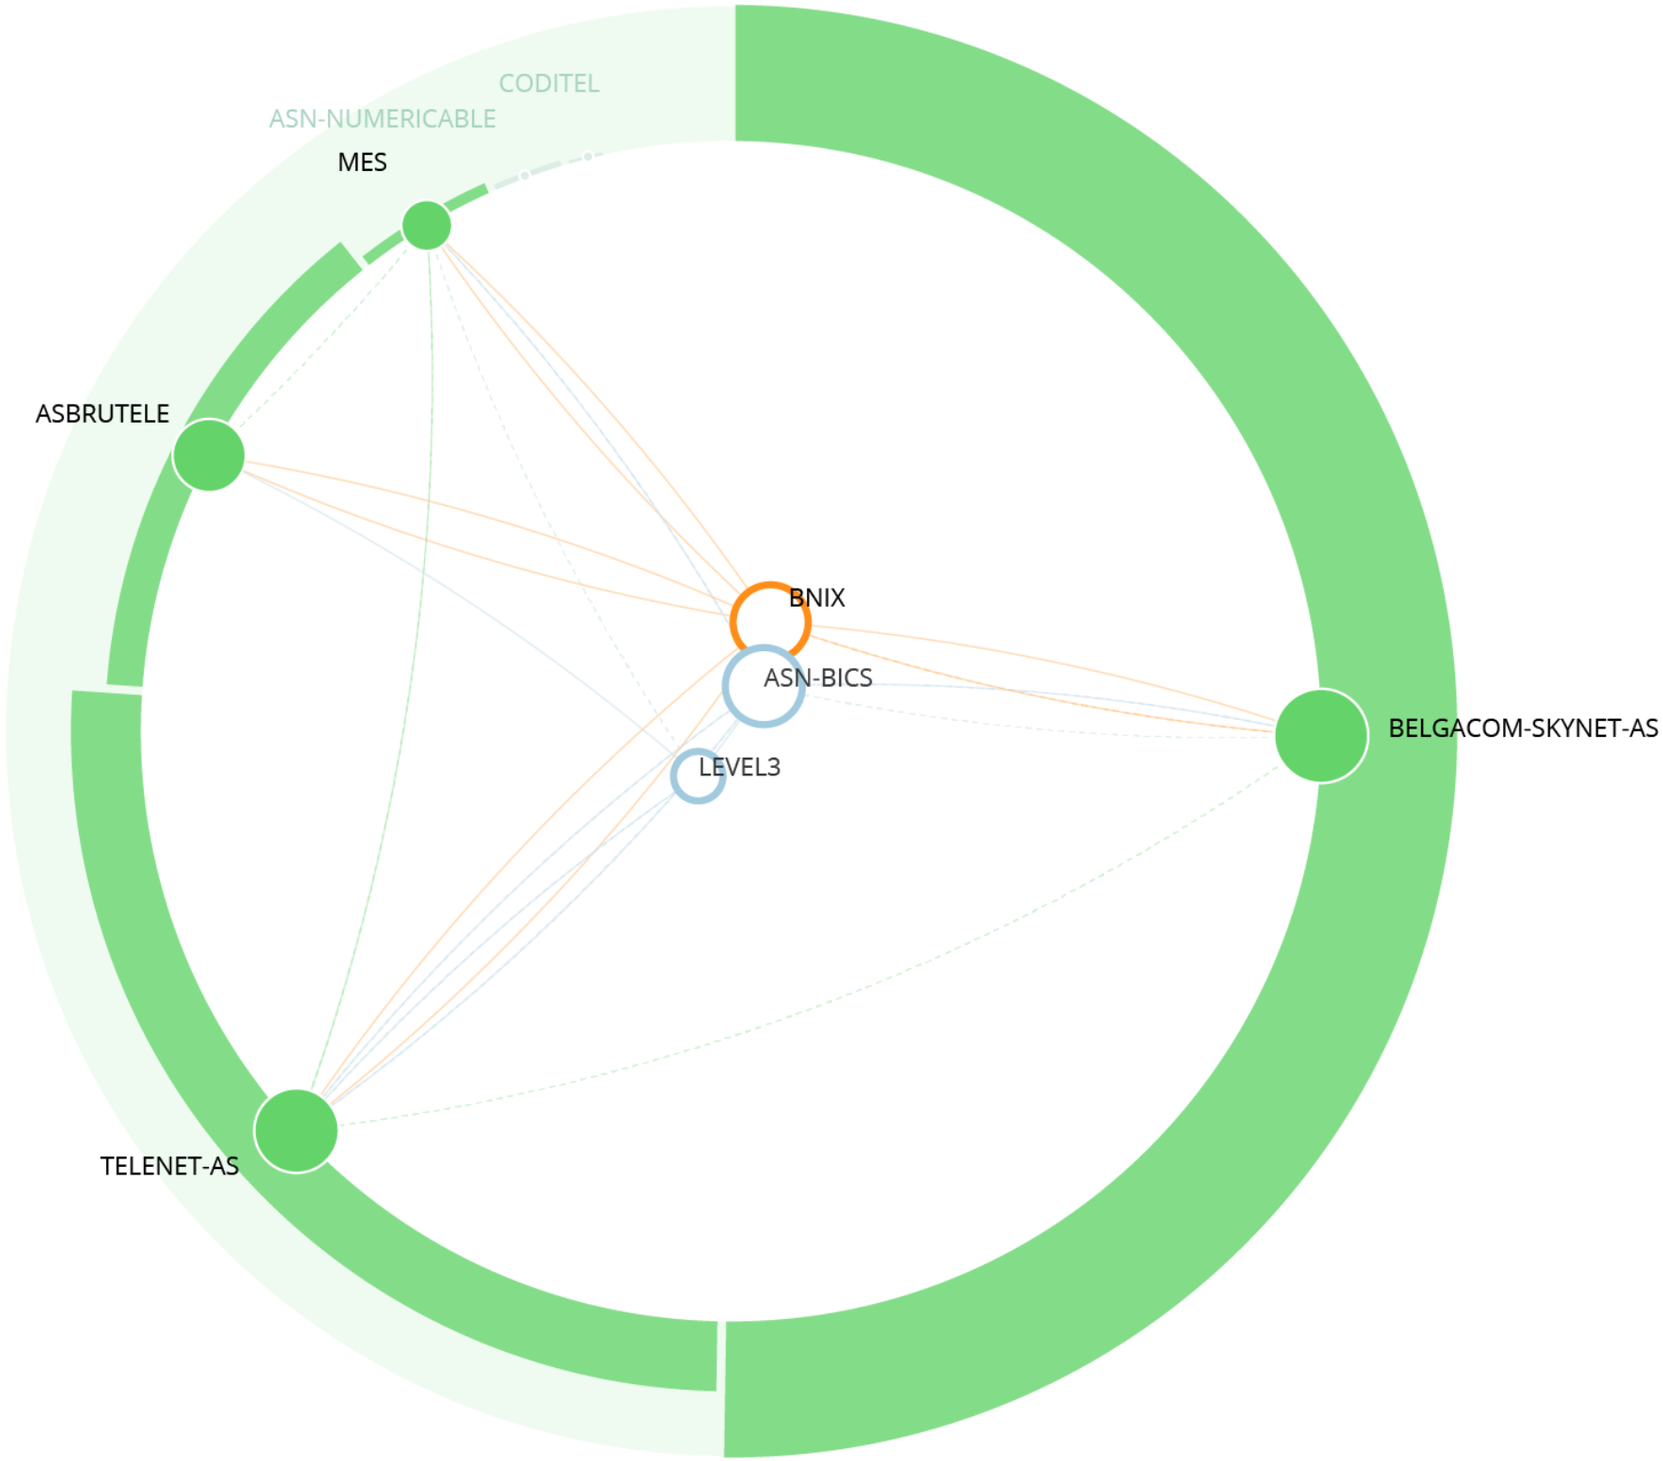

http://sg-pub.ripe.net/ixp-country-jedi/

How Are End Users Interconnected?

Belgium

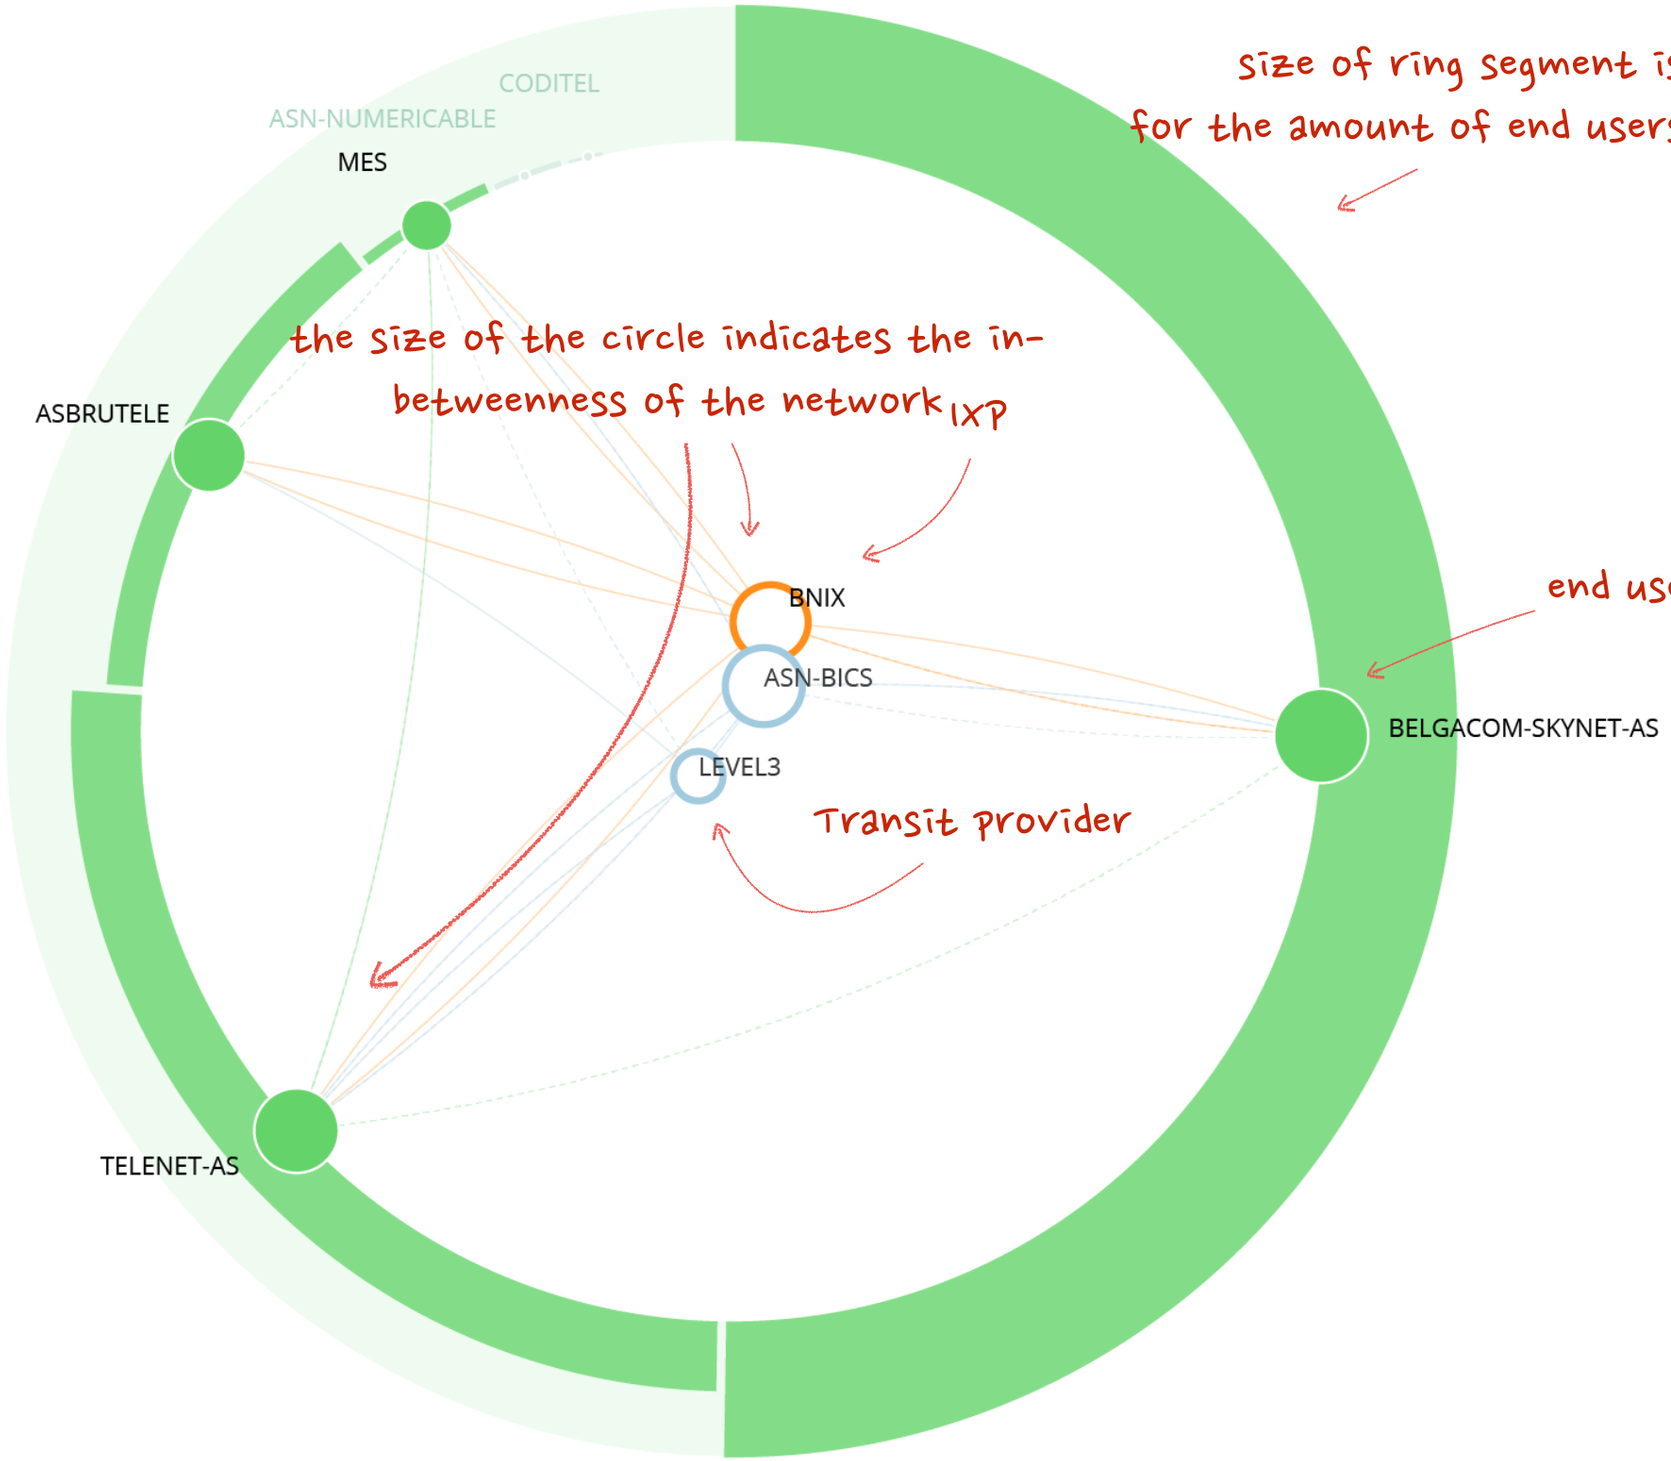

emile.aben@ripe.net | Internet Conference

http://sg-pub.ripe.net/ixp-country-jedi/

How Are End User Interconnected?

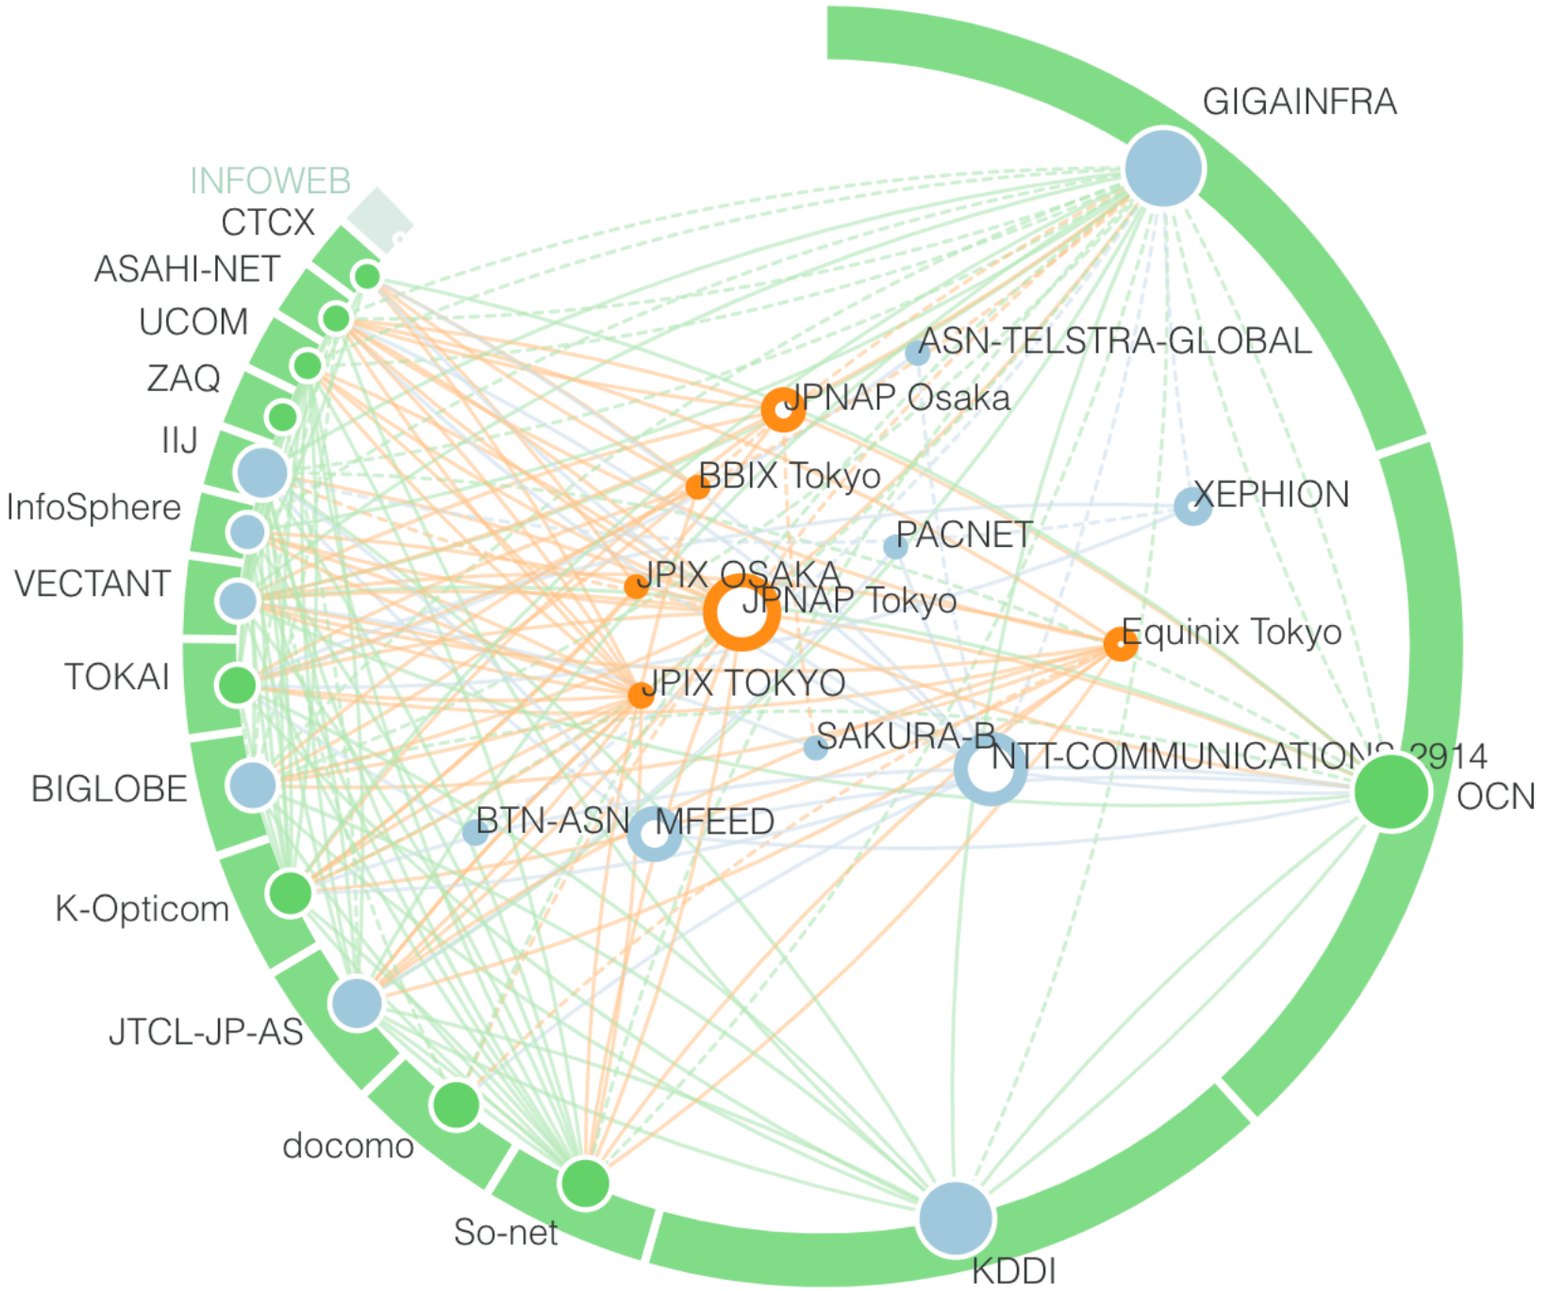

http://sg-pub.ripe.net/ixp-country-jedi/jp/2018/10/01

RIPE Atlas: Ethics in Measurement

Ethical Design Decisions

- Low, cheap barrier of entry
- Active measurements only
  - No passive measurements
  - Probes do not observe user traffic
- Data, API, source code, tools: free and open
- Set of measurement types limited
  - In order to prevent putting probe hosts at risk

https://www.ripe.net/about-us/press-centre/publications/presentations/2017/ethics-in-technology-nluug-najaar-2017

https://fosdem.org/2017/interviews/vesna-manojlovic/

Ongoing Moral Dilemmas

- 2013: Opening-up source code
- 2014: Keeping “non-public” measurements
- 2015: Not allowing HTTP measurements to random targets
- 2016: Security audit and pen-testing

Strong Community Involvement

- Active mailing list (ripe-atlas@ripe.net)
  - Passionate discussions
- HTTP measurements only towards RIPE Atlas anchors
  - https://labs.ripe.net/Members/kistel/ethics-of-ripe-atlas-measurements
- Responsible disclosure (bug reports)

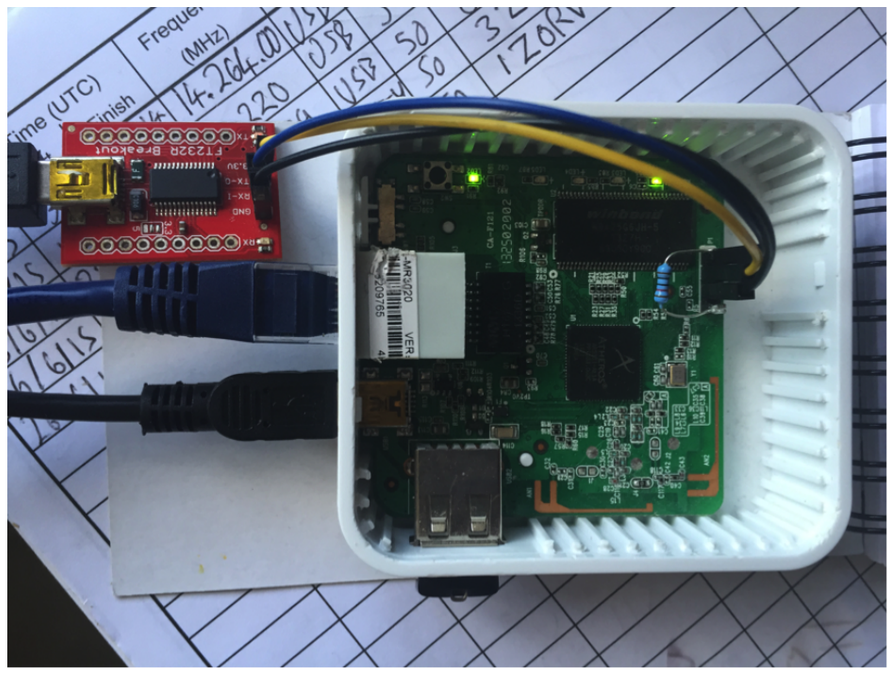

RIPE Atlas: DNS Characterisation

Willem Toorop

Opportunity: DNS resolution

- NLNetLabs project: DNSThought
- Started at a hackathon we organised

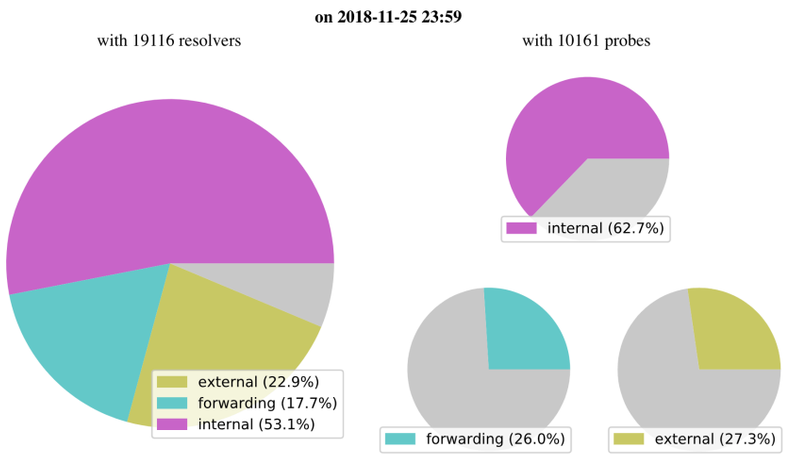

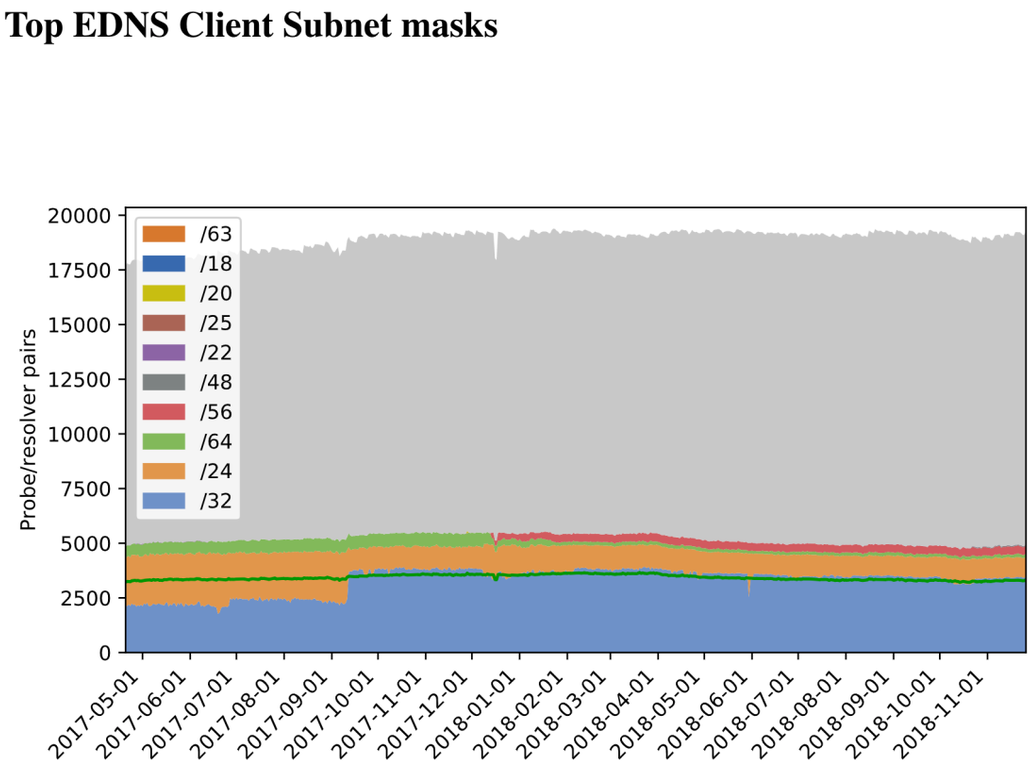

https://indico.dns-oarc.net/event/29/contributions/654/attachments/633/1041/dnsthought-oarc29.pdf

https://dnsthought.nlnetlabs.nl/

RIPE RIS: Zombie Routes?

Collaborator: Romain Fontugne

Challenge: Is A Route Withdrawn?

- Zombie Routes: Routing table entries for routes that are withdrawn from origin

Challenge: Is A Route Withdrawn?

- Zombie Routes: Routing table entries for routes that are withdrawn from origin

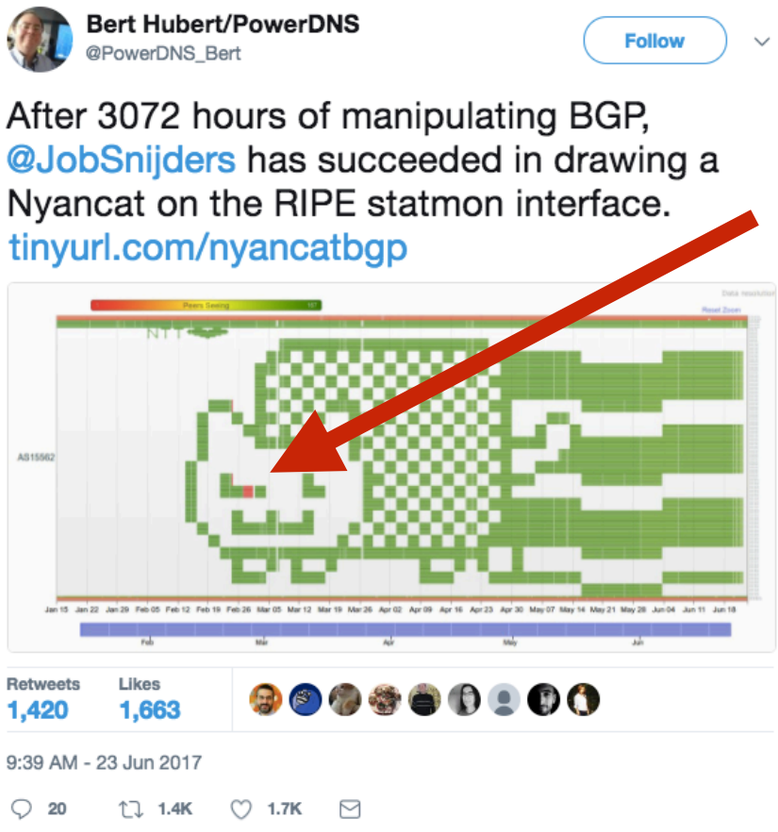

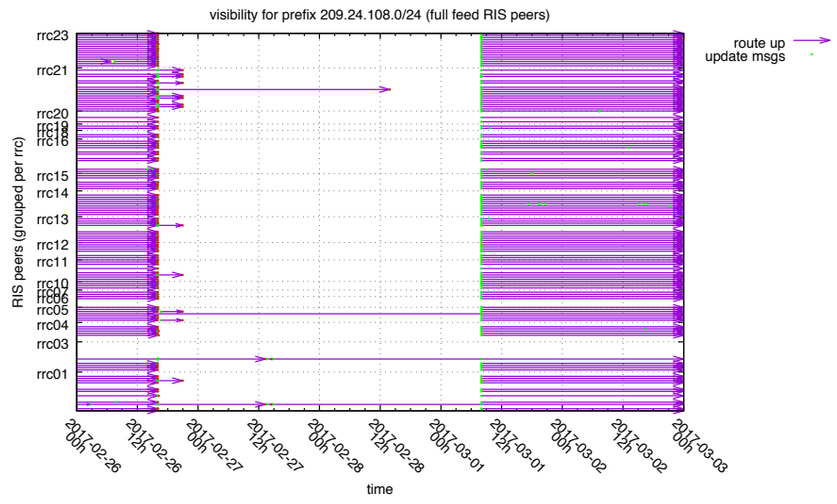

Example: Long Lived

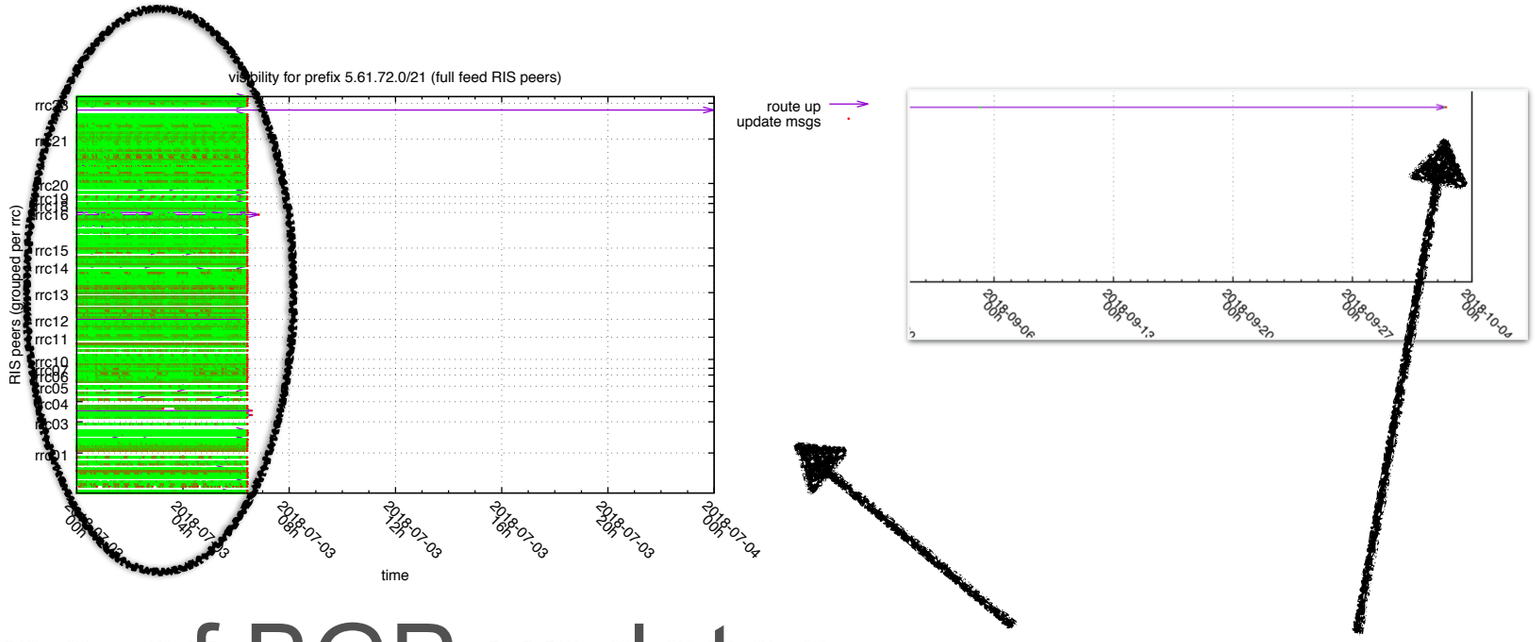

Tons of BGP updates

Route totally withdrawn only after manual intervention

3 Months!

Confusing if you want to know: Is this routed publicly?

BGP Beacons

Current RIS Routing Beacons

Note: IPv6 beacons and anchors are now being announced from the RRCs. See below for details.

All RRCs:

- Announcements at 00:00, 04:00, 08:00, 12:00, 16:00, 20:00 (UTC)
- Withdrawals at 02:00, 06:00, 10:00, 14:00, 18:00, 22:00 (UTC)

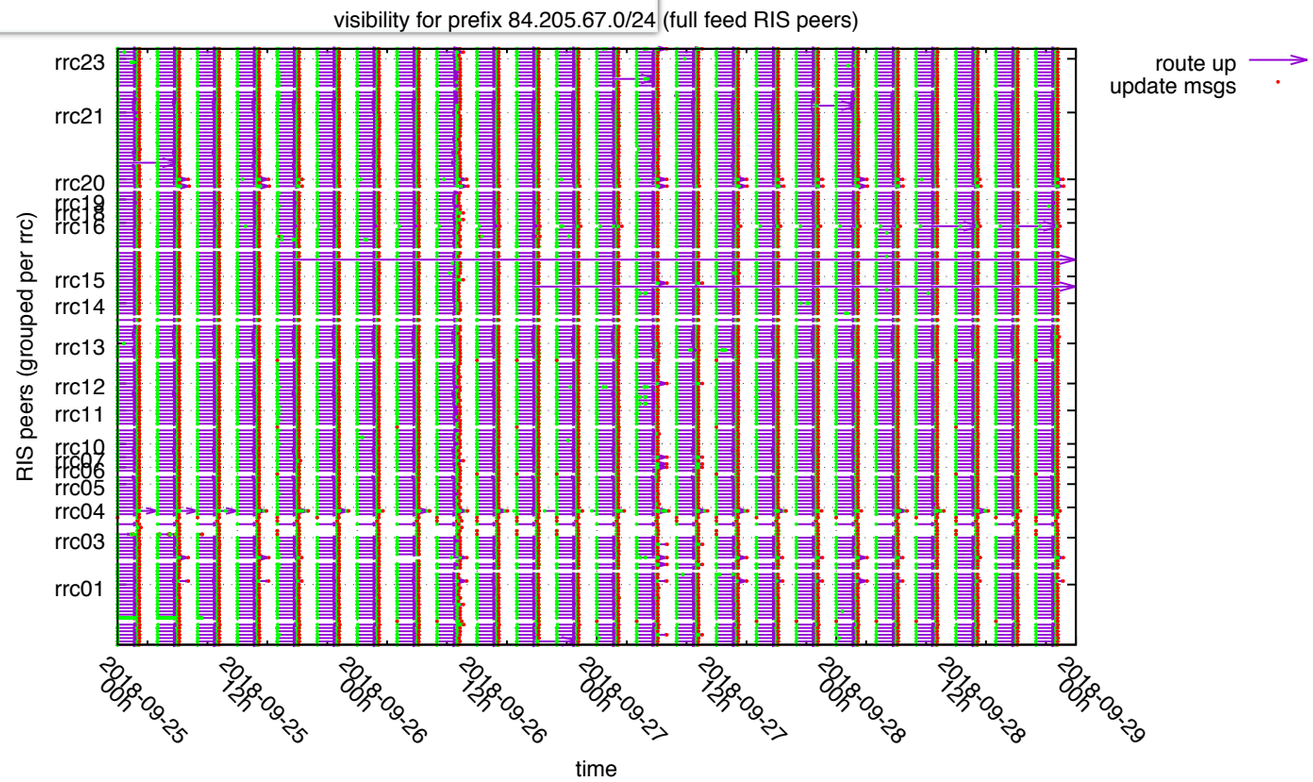

BGP Beacons

Current RIS Routing Beacons

Note: IPv6 beacons and anchors are now being announced from the RRCs. See below for details.

All RRCs:

- Announcements at 00:00, 04:00, 08:00, 12:00, 16:00, 20:00 (UTC)
- Withdrawals at 02:00, 06:00, 10:00, 14:00, 18:00, 22:00 (UTC)

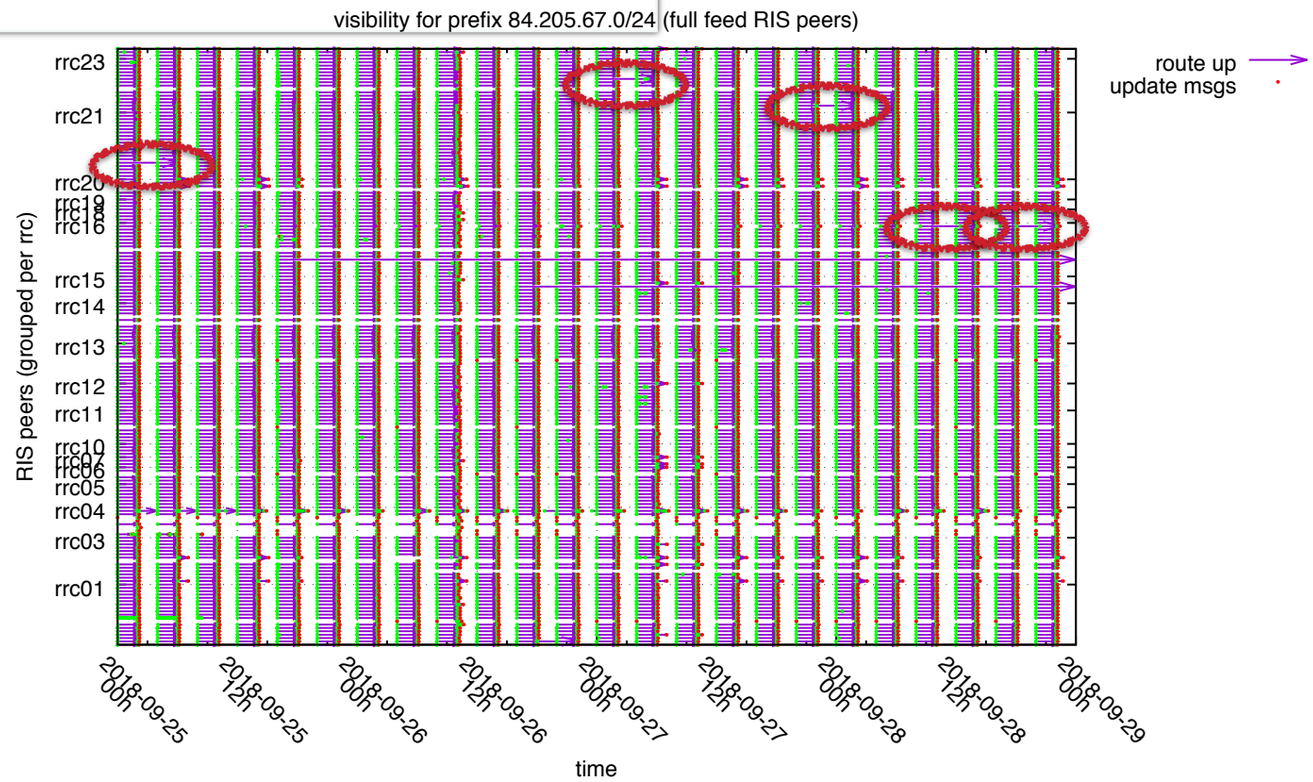

BGP Beacons

Current RIS Routing Beacons

Note: IPv6 beacons and anchors are now being announced from the RRCs. See below for details.

All RRCs:

- Announcements at 00:00, 04:00, 08:00, 12:00, 16:00, 20:00 (UTC)
- Withdrawals at 02:00, 06:00, 10:00, 14:00, 18:00, 22:00 (UTC)

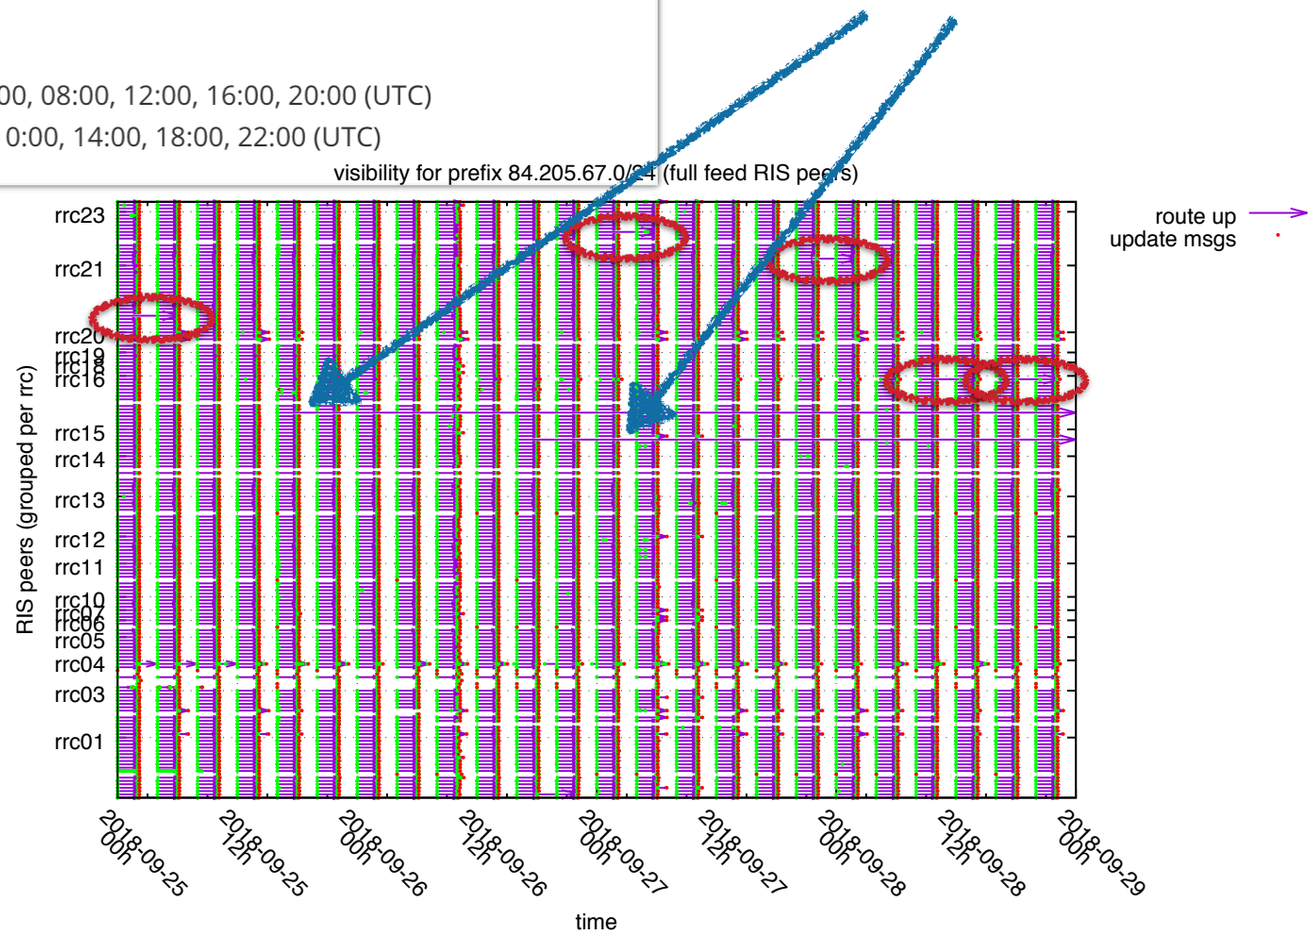

BGP Zombies Conclusions

- Made us look hard and long at our infrastructure
- Existence is commonly known among netops
  - Hard to debug
- One vendor identified by operator, but needs proof of incorrect behaviour

RIPE RIS: Bias?

Collaborator: Romain Fontugne

How Important is network X for you?

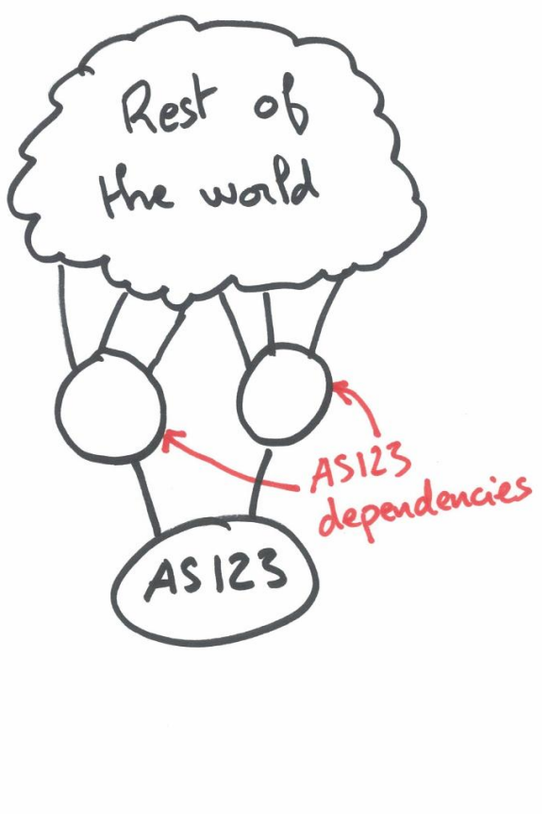

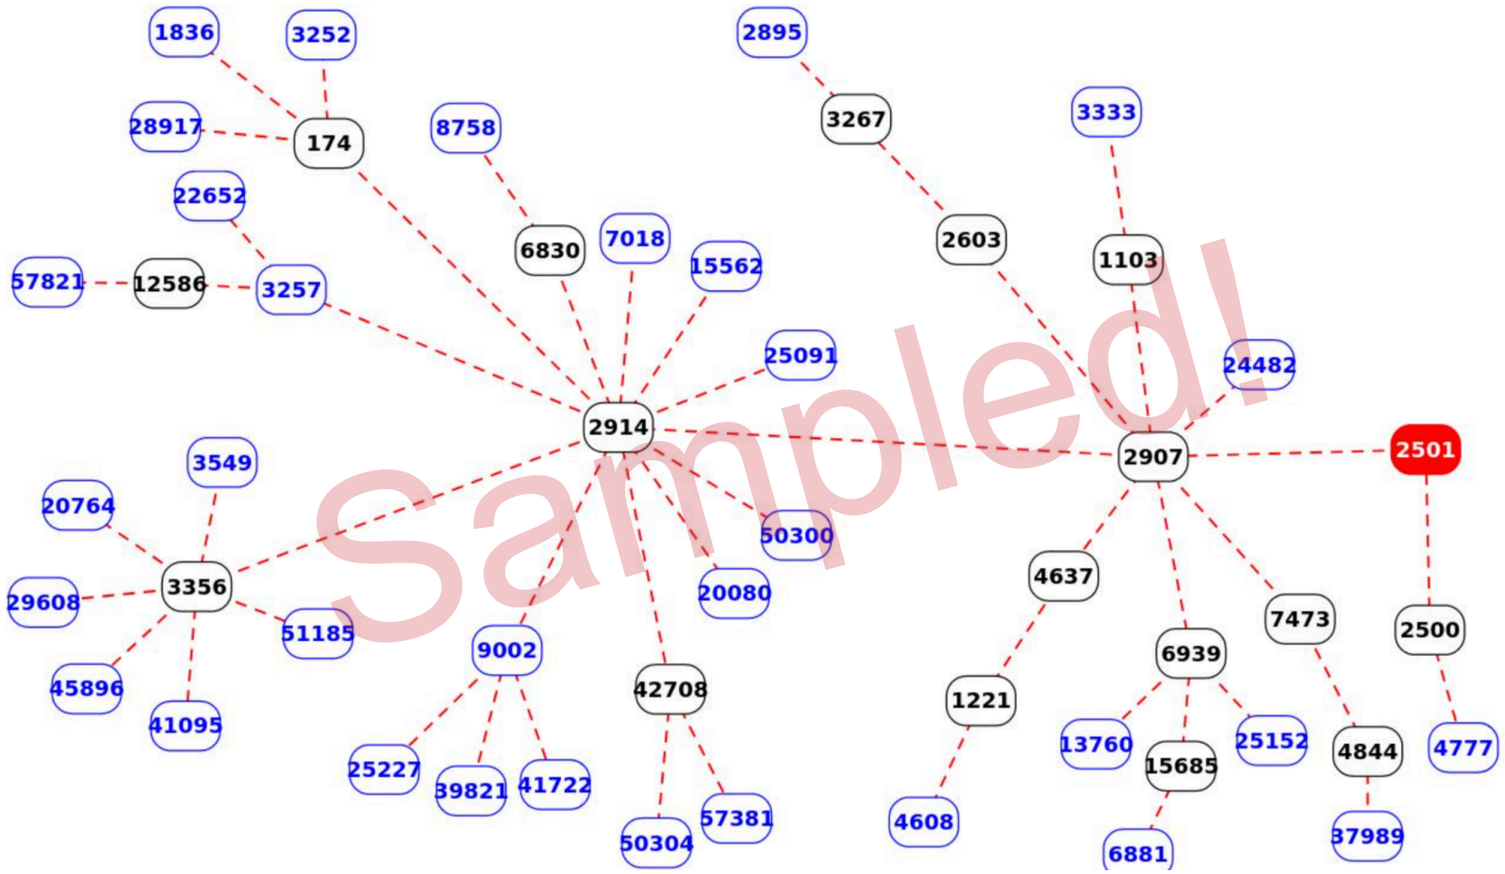

Sampled!

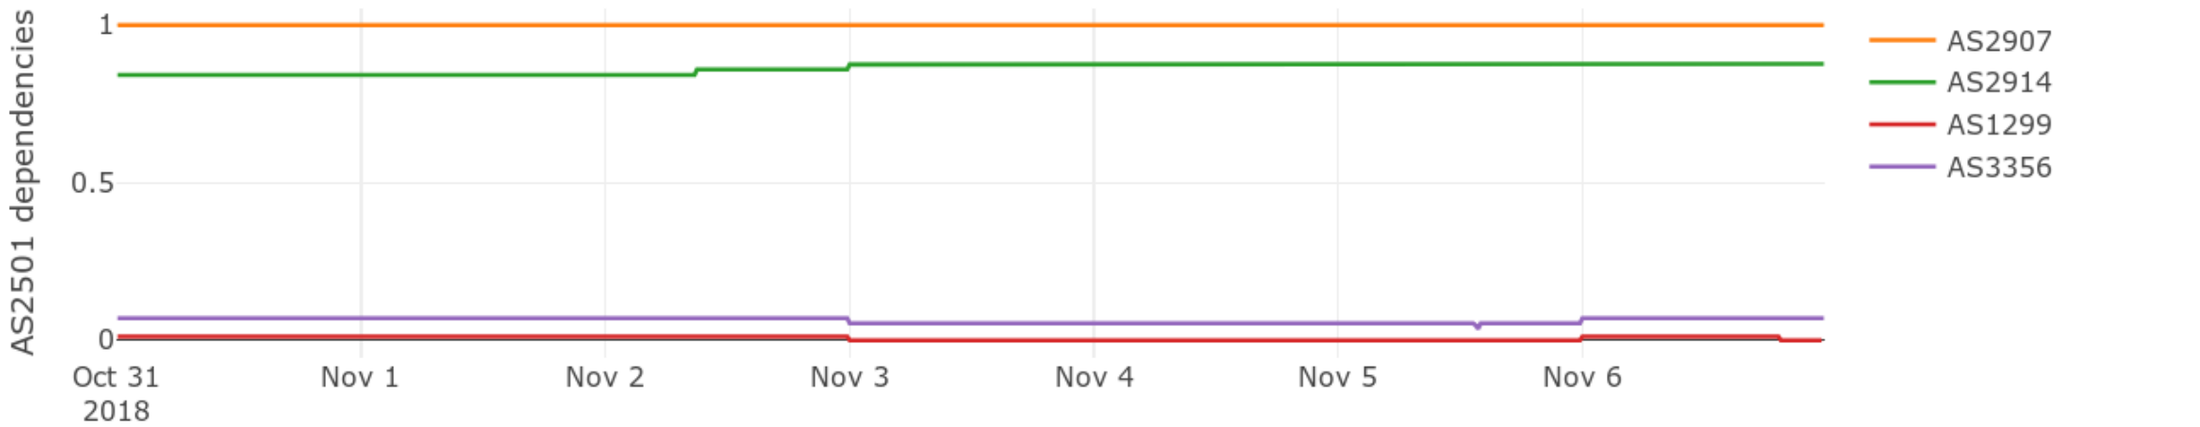

Counter Bias with AS Hegemony

- Remove bias with trimmed averaging
  - “AS Hegemony” scores

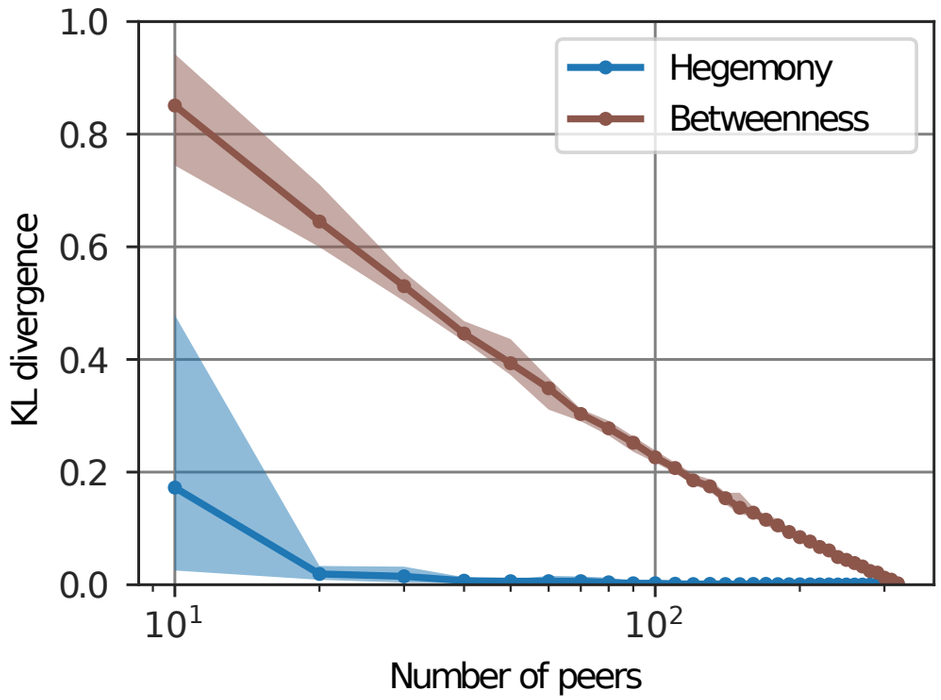

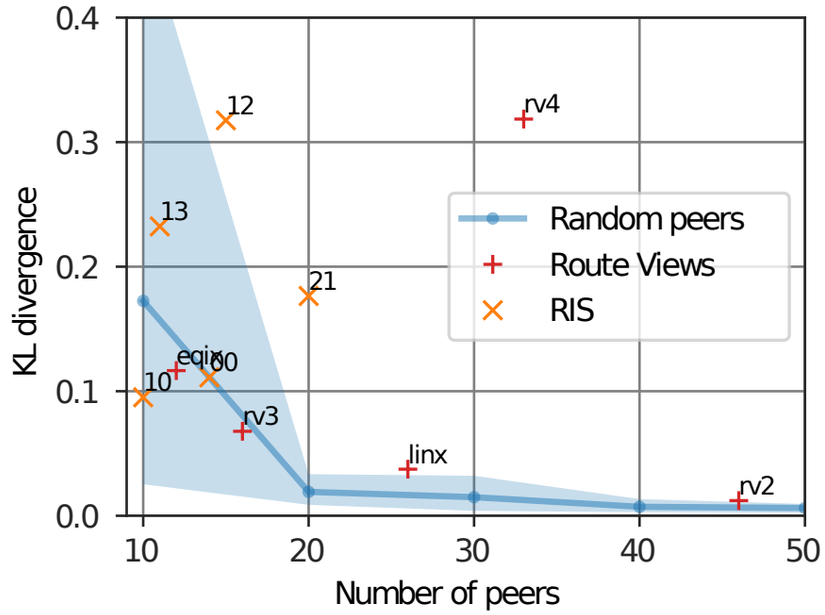

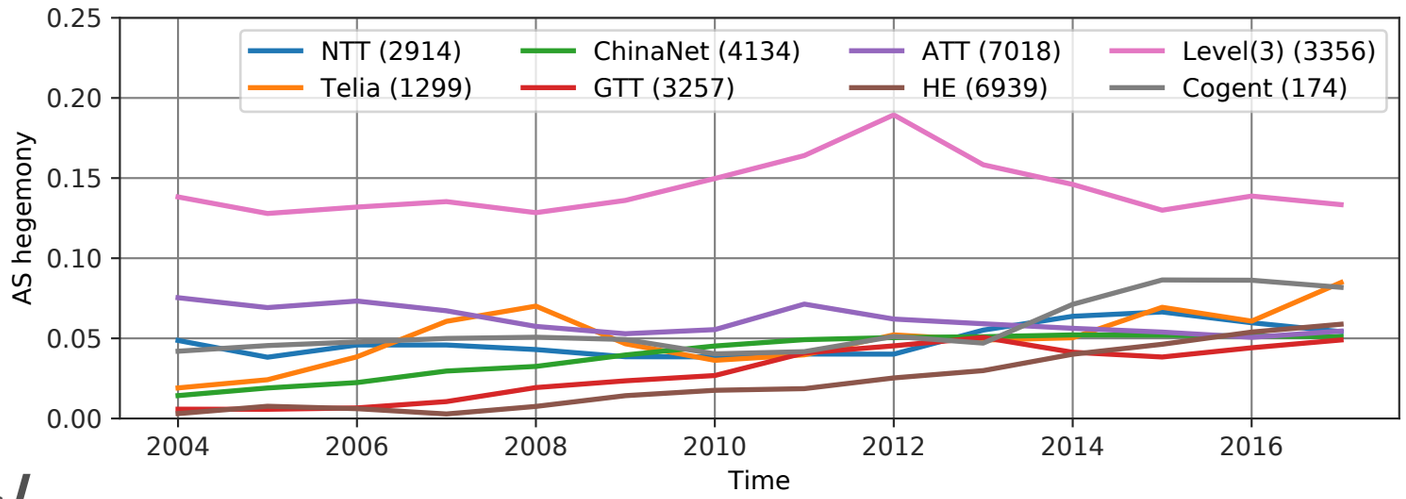

https://ihr.iijlab.net/ihr/

RIPE RIS: Outages

Collaborator: Collin Anderson

RIS Outage Detection

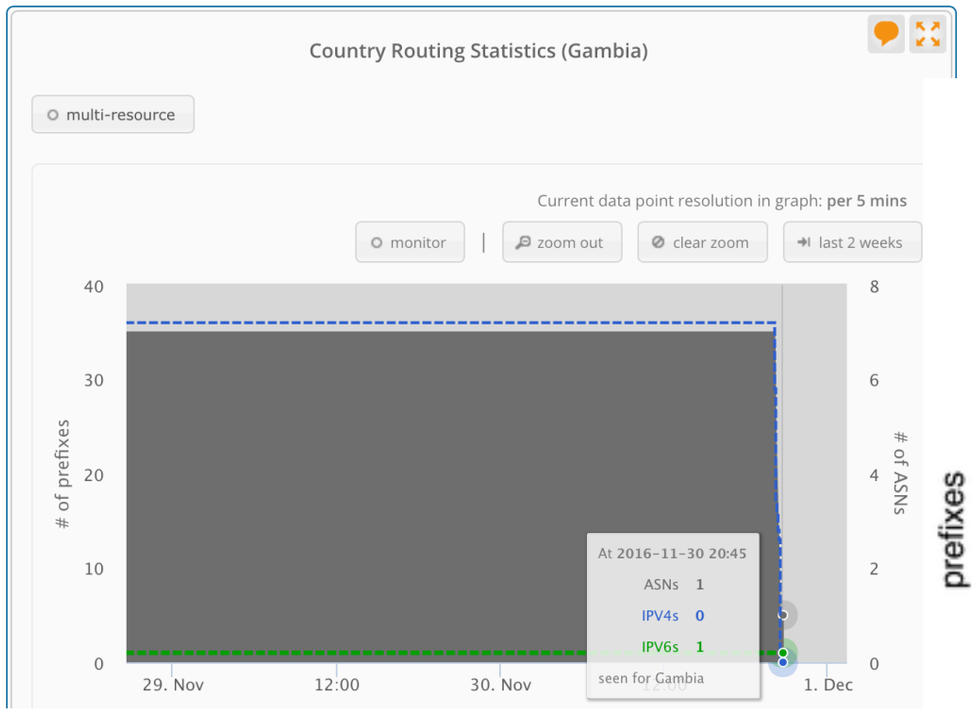

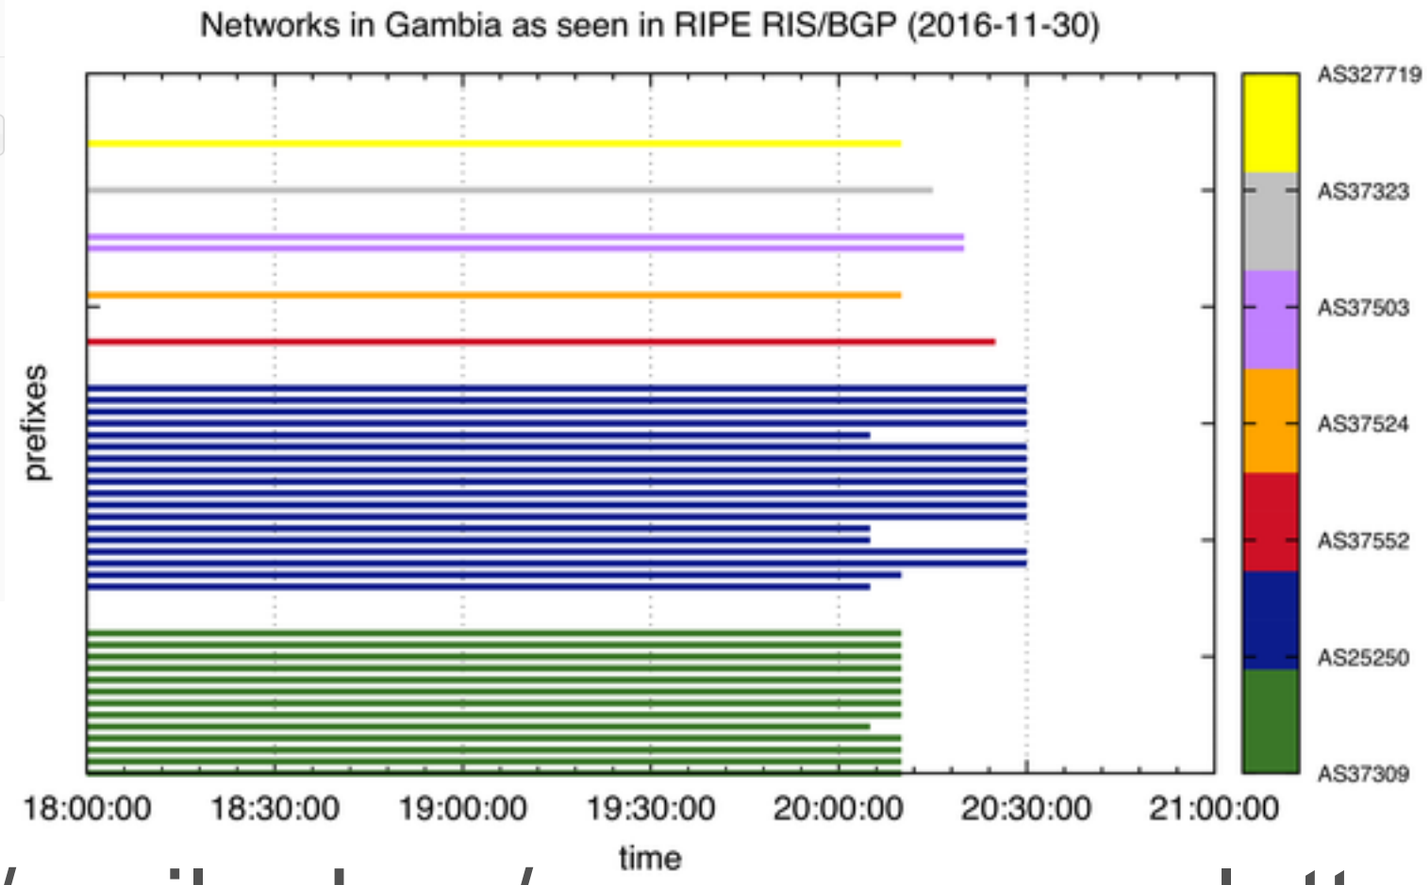

https://github.com/emileaben/resource-gnuplotter

https://labs.ripe.net/Members/emileaben/internet-access-disruption-in-the-gambia-2016

Even better: Combine with Atlas!

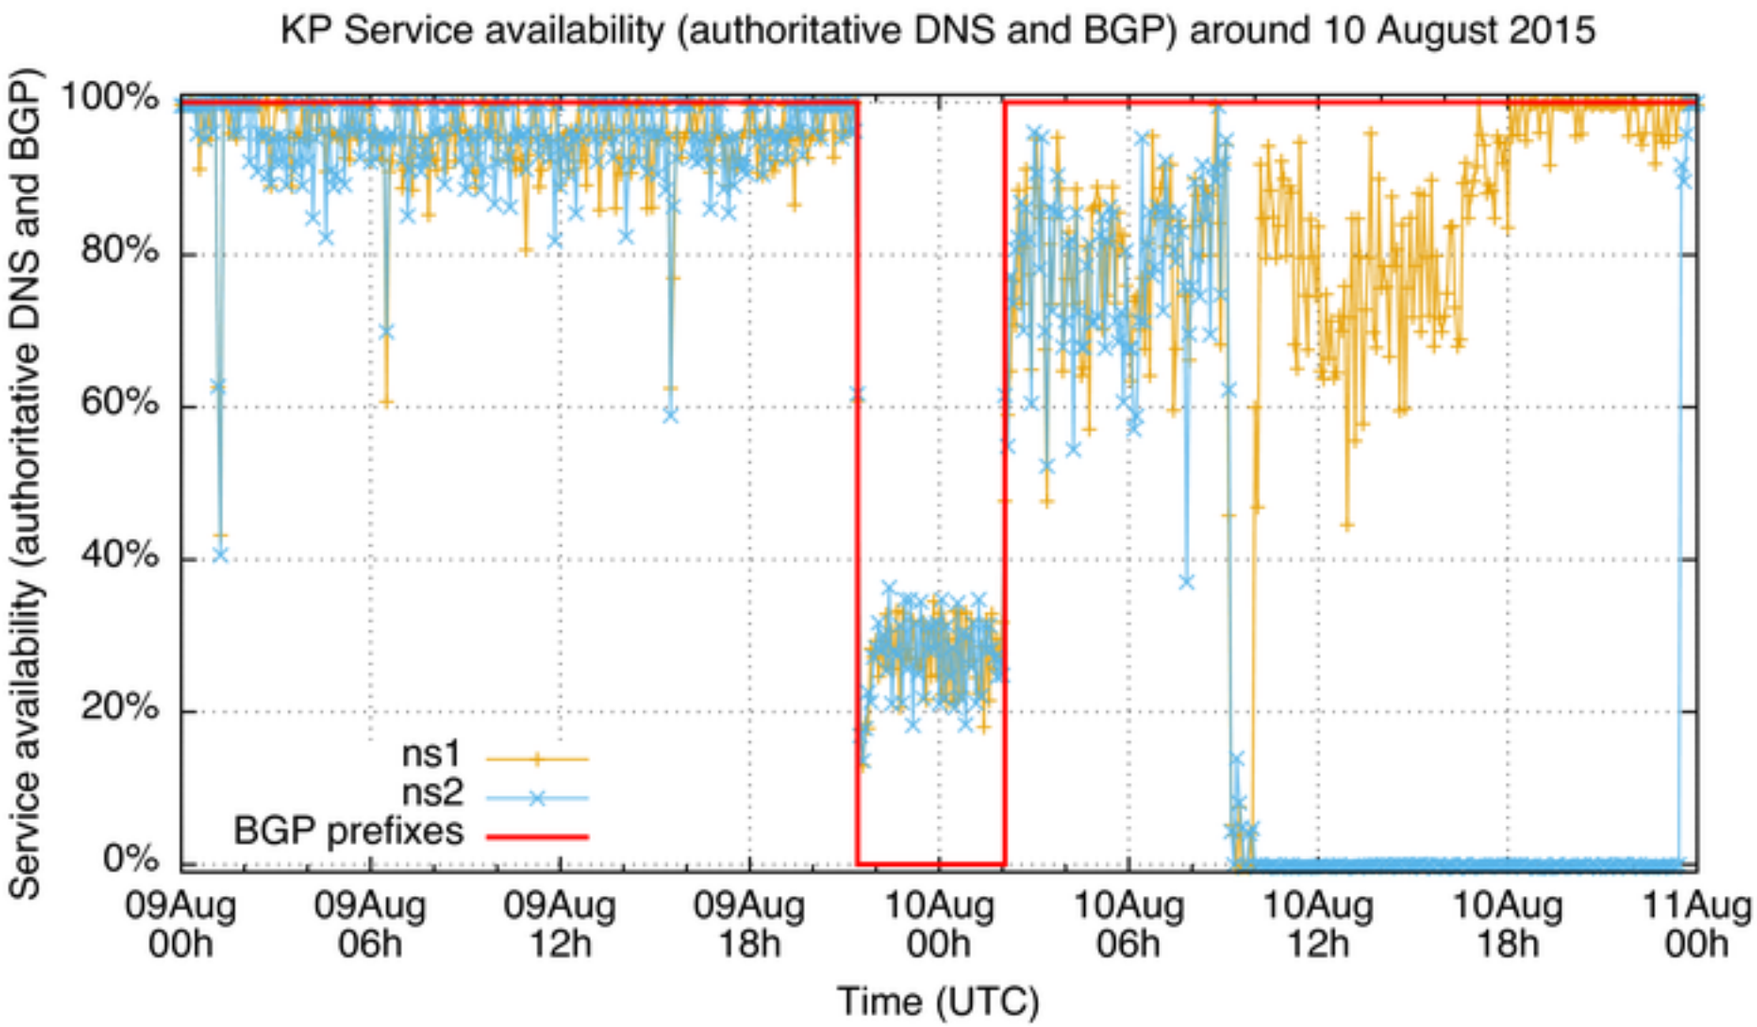

https://labs.ripe.net/Members/emileaben/the-internet-in-north-korea-hanging-by-a-single-thread

RIPE RIS: Tracking Technology Deployment

Longitudinal view

- Tracking IPv6: https://v6asns.ripe.net/

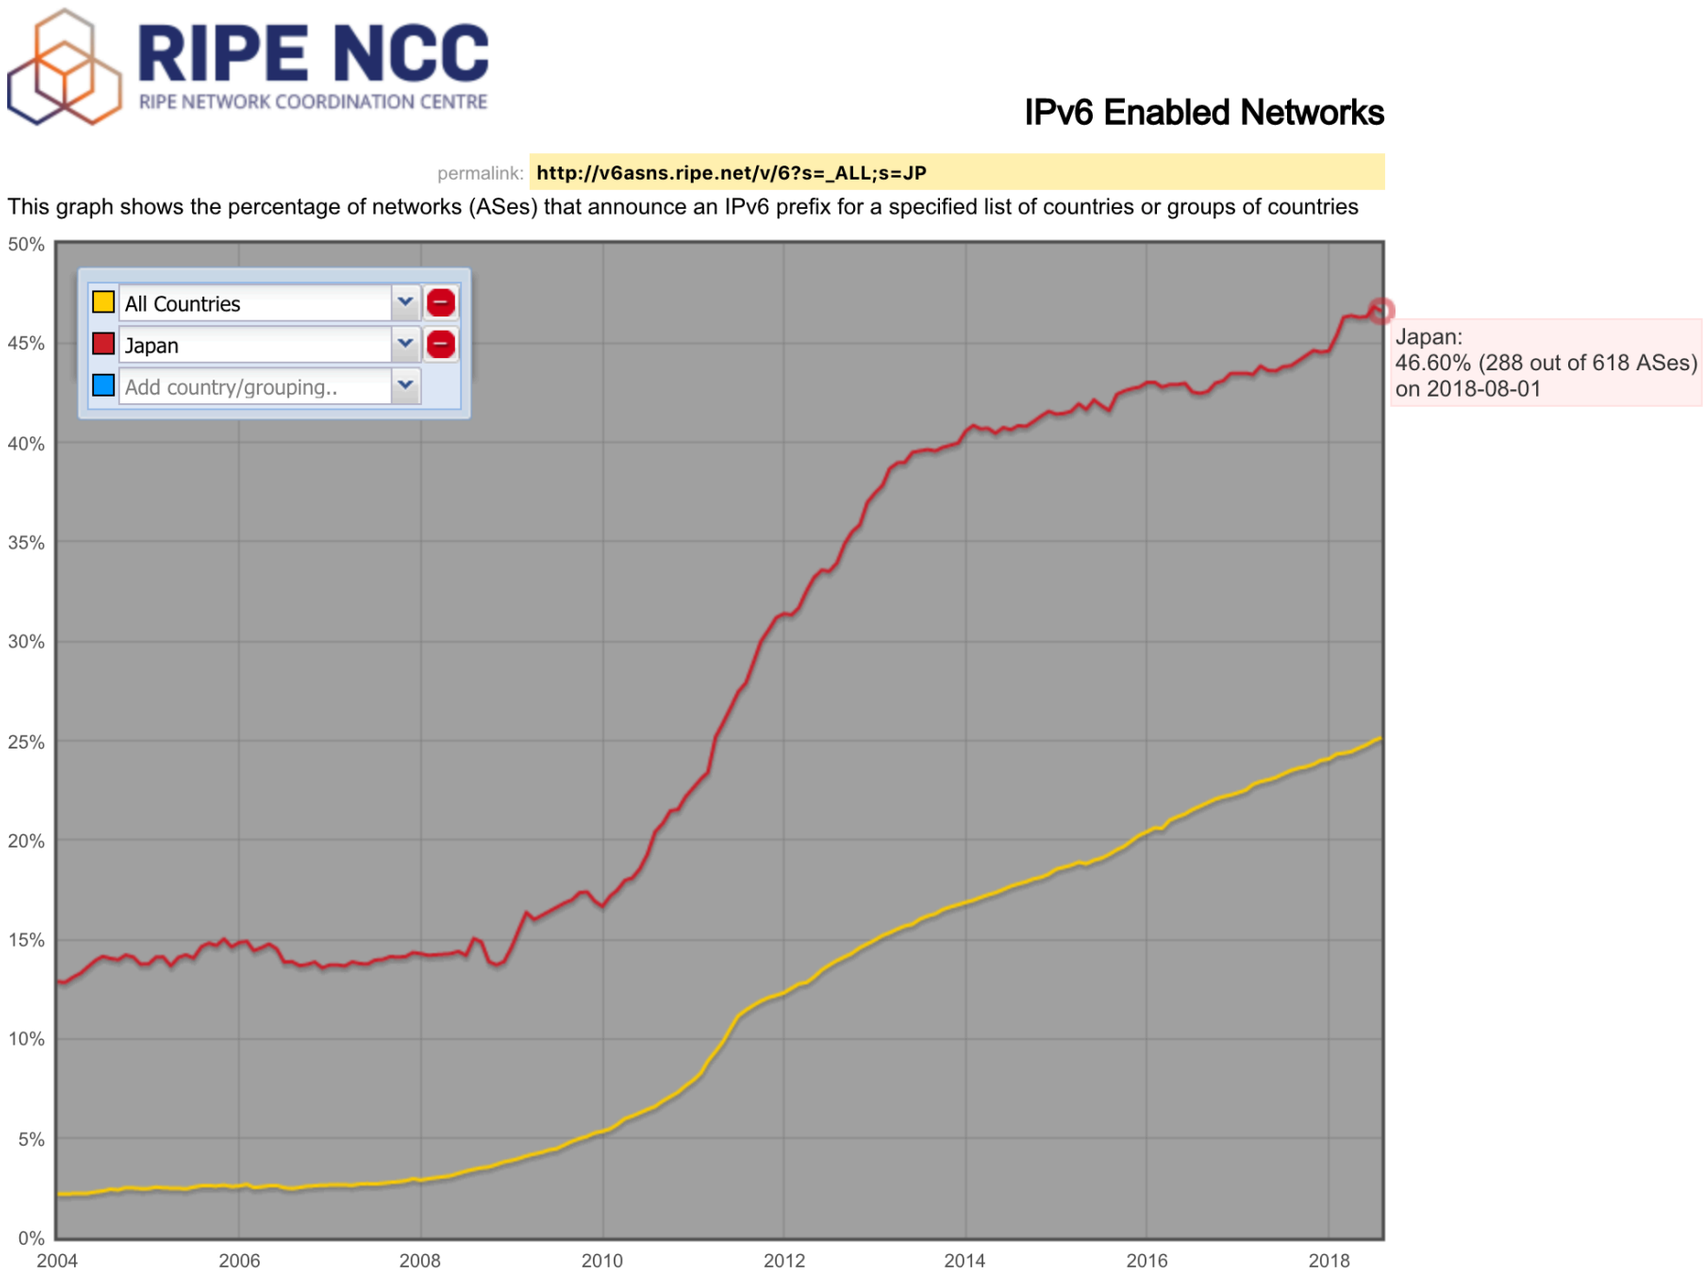

Tracking Large BGP Communities

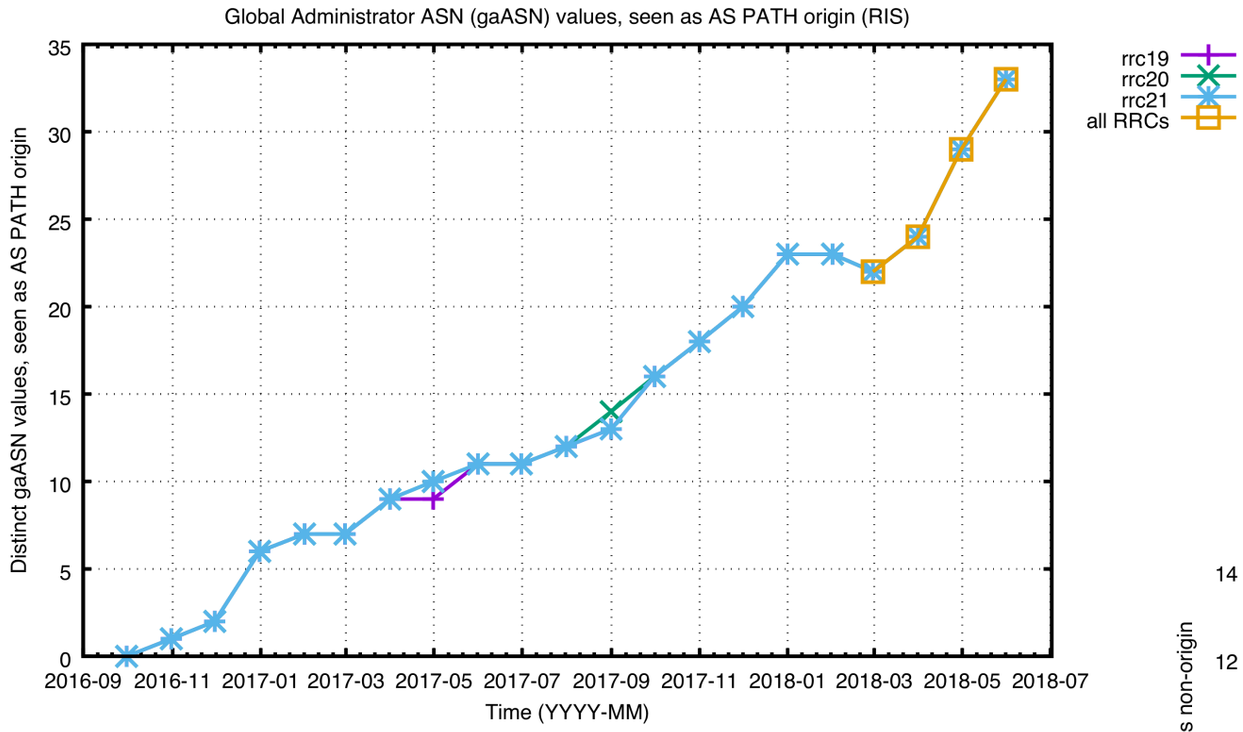

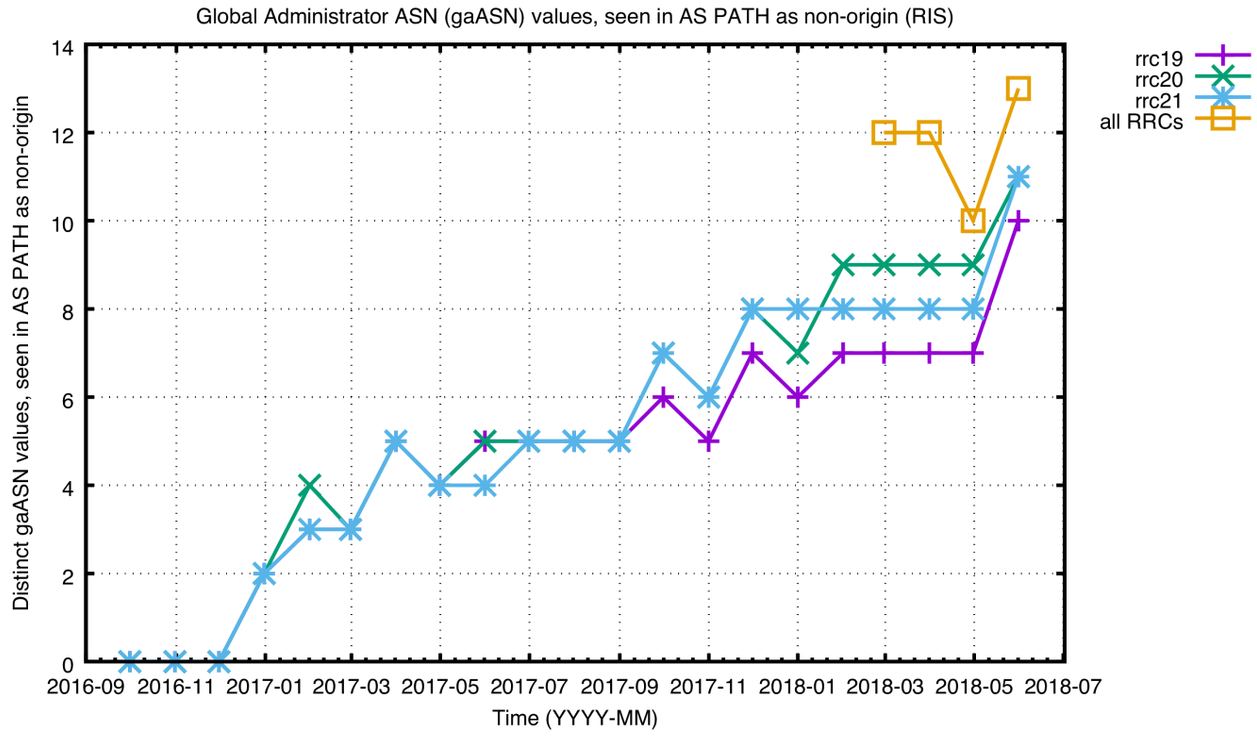

https://labs.ripe.net/Members/emileaben/bgp-large-communities-uptake-by-the-community-at-large

But … The Main Challenge/Opportunity?

Main Challenge/Opportunity

- Many challenges and opportunities around RIPE Atlas and RIPE RIS

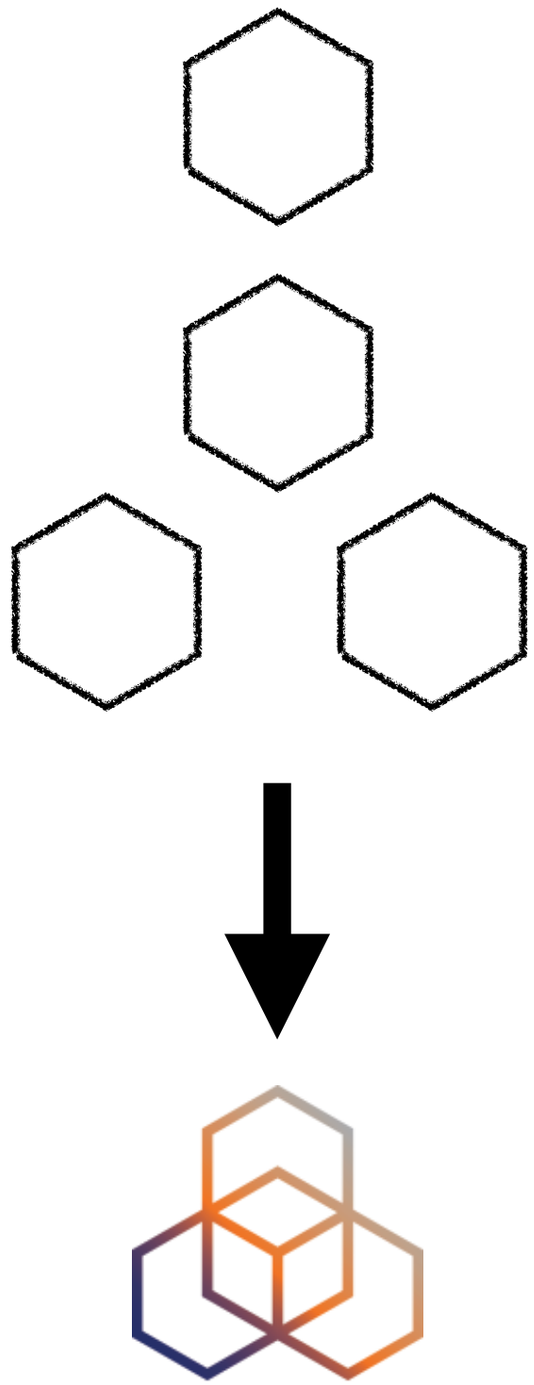

- Main Opportunity:

Collaboration

- We are open to collaborate

Collaboration Possibilities

- We can help:
  - Understand/analyse data
  - Schedule (large scale) measurements
  - Collaborate on projects
- Open Code (Github) / Open Data
- Communication with Internet community
  - https://labs.ripe.net/ / RIPE meetings / RACI
- Hackathons: https://labs.ripe.net/hackathons
- Students/Internships

Questions ?

emile.aben@ripe.net

not a typo

twitter: @meileaben
mastodon: @meileaben@vis.social

not a typo