

The Remote Peering Jedi A portal in the remote peering ecosystem

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











Motivation

- Knowing which peers allows better-informed peer selection process and higher transparency.
- Operational concerns:
 - $_{\odot}$ Added latency and troubleshooting complexity
 - \circ Routing inefficiencies
 - \circ Invisible Layer-2 intermediaries
 - $_{\odot}\,\text{Network}$ 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 tralXroute
- 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

 \circ 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



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)

 $\odot\,\text{DE-CIX}$ (Ping from inside the IXP)

• France-IX (Ping from inside the IXP)

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



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



Presence-informed RTT Geolocation: Methodology





Presence-informed RTT Geolocation: Example

★⁴ Presence-informed RTT geolocation

General	Information	Probes	Мар	Results			26742	61156		- 0	2016-10-23 02:27		32.504	0.0%
Probe 🕯	ASN (IPv4) ♦	ASN (IPv6) 🖨	\$	Time (UTC)	♦ RTT	- Packe	25647	42610		- 0	2016-10-23 02:27		33.462	0.0%
6198	52048	52048	= 6	2016-10-23 02:23	7 🔋 1.900	0.0%	20190	58054		- 0	2016-10-23 02:27		33.487	0.0%
4005	31519		= 4	2016-10-23 02:2	7 🛑 3.513	0.0%	25780	58010	58010	= 0	2016-10-23 02:27		33.642	0.0%
14175	24651		= 4	2016-10-23 02:2	7 🔵 3.536	0.0%	4012	15895		- 6	2016-10-23 02:27		35.392	0.0%
21477	43956		= 4	2016-10-23 02:2	7 🛑 4.309	0.0%	25207	20853	8664	- 0	2016-10-23 02:27		35.548	0.0%
21252	8285		= 4	2016-10-23 02:23	7	0.0%	14746	2119		Ξ Δ	2016-10-23 02:27		36.219	0.0%
21458	56588	Ge	olo	ocation	: Riga. Latvia	0.0%	6019	3333	3333	= 0	2016-10-23 02:27		37.852	0.0%
20035	29518		= 4	2016-10-23 02:21	7	0.0%	12638	42610		= 0	2016-10-23 02:27		38.706	0.0%
3375	3249		= 4	2016-10-23 02:2	7 11.724	0.0%	6220	34106	34106	= 0	2016-10-23 02:28		38.994	0.0%
1333	8473		•• 4	2016-10-23 02:23	7	0.0%	25626	37100	37100	叢 🔒	2016-10-23 02:27		40.655	0.0%
13982	60415		= 4	2016-10-23 02:2	7	0.0%	27207	43139		= 0	2016-10-23 02:27		42.506	0.0%
28562	35807		- 4	2016-10-23 02:2	7 15.931	0.0%	2296	1103	1103	= 0	2016-10-23 02:27		45.207	0.0%
2583	39651		- 4	2016-10-23 02:2	7	0.0%	4378	28792		₩ 🔒	2016-10-23 02:27		46.497	0.0%
11443	3216		- 4	2016-10-23 02:2	20.879	0.0%	3622	50343	50343	= 0	2016-10-23 02:27		47.348	0.0%
12846	13110		- 4	2016-10-23 02:2	21.413	0.0%	325	6830		= 🏻	2016-10-23 02:27		49.091	0.0%
19936	2586		= (2016-10-23 02:2	21.564	0.0%	14639	13037		₩ 🔒	2016-10-23 02:27		50.998	0.0%
20059	8334	6939	- 4	2016-10-23 02:2	22.586	0.0%	17491	29535		- 0	2016-10-23 02:27		51.210	0.0%
10234	44050		- 4	2016-10-23 02:2	23.870	0.0%	2177	20712	20712	₩ 🔒	2016-10-23 02:27		53.083	0.0%
13251	31514		- 4	2016-10-23 02:2	24.439	0.0%	246	6830		= 0	2016-10-23 02:27		53.748	0.0%
20904	42148		- 4	2016-10-23 02:2	25.559	0.0%	16736	5617	5617	- 0	2016-10-23 02:27		56.662	0.0%
1223	39102		- 4	2016-10-23 02:2	25.645	0.0%	1343	3320		Ξ Δ	2016-10-23 02:27		57.640	0.0%
24588	3301	1299	- 4	2016-10-23 02:2	26.966	0.0%	28796	34156		= 0	2016-10-23 02:27		64.494	0.0%
27644	34867		= 4	2016-10-23 02:2	28.429	0.0%	3193	5089		# 🔒	2016-10-23 02:27		67.737	0.0%
11629	15683		- 4	2016-10-23 02:2	30.324	0.0%	28569	5617	43447	- 0	2016-10-23 02:27		70.004	RIPE
4534	2586		= (2016-10-23 02:23	30.663	0.0%	2114	6830		= 0	2016-10-23 02:27	20.752		24-28 DCT 2016
							-							MADDID CDAIN

The Remote Peering Jedi - A portal in the remote peering ecosystem



The Remote Peering Jedi - Home About



BACKUP SLIDES

Remote peerings used to access large CDNs



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



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:
 - \circ AMS-IX: 16%
 - o LINX: 20%
 - \circ 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.



Presence-informed RTT Geo-location

- Most of the available accurate geo-location methods can resolve only a subset of the remote peering IPs:
 - o **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).
 - o 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: Example

★⁴ Presence-informed RTT geolocation

General Informati	on Probes	Мар	Results				26742	61156		- 0	2016-10-23 02:27		32.504	0.0%
Probe 🗢 ASN (IP)	(4) + ASN (IPv6)	* * ·	• Time (UTC)	♦ RTT		✓ Pac	ke ²⁵⁶⁴⁷	42610		- 6	2016-10-23 02:27		33.462	0.0%
6198 52048	52048	= 0	2016-10-23 02:	27 🔋 1.900		0.00	6 20190	58054		= 🏠	2016-10-23 02:27		33.487	0.0%
4005 31519		= 0	2016-10-23 02:	27 3.513		0.00	6 25780	58010	58010	= 0	2016-10-23 02:27		33.642	0.0%
14175 24651		= 0	2016-10-23 02:	27 3.536		0.09	6 4012	15895		= 🏻	2016-10-23 02:27		35.392	0.0%
21477 43956		= 0	2016-10-23 02:	27 🛑 4.309		0.09	6 25207	20853	8664	- 0	2016-10-23 02:27		35.548	0.0%
21252 8285		= 0	2016-10-23 02:	27 7.613		0.09	6 14746	2119		i 6	2016-10-23 02:27		36.219	0.0%
21458 56588		= 0	2016-10-23 02:	27 9.992		0.09	6019	3333	3333	= 6	2016-10-23 02:27		37.852	0.0%
20035 29518		i a	2016-10-23 02:	27 11.719		0.00	6 12638	42610		- 6	2016-10-23 02:27		38.706	0.0%
3375 3249		= 0	2016-10-23 02:	27 11.724		0.09	6220	34106	34106	= 0	2016-10-23 02:28		38.994	0.0%
1333 8473		H 🔒	2016-10-23 02:	27 14.6	96		6 25626	37100	37100	# o	2016-10-23 02:27		40.655	0.0%
13982 60415		= 0	2016-10-23 02:	27 15.0	92	Δς1	25	78		- 6	2016-10-23 02:27		42.506	0.0%
28562 35807		= 0	2016-10-23 02:	27 15.9	931	0.04	69	11		= 6	2016-10-23 02:27		45.207	0.0%
2583 39651		iii (1	2016-10-23 02:	27 17	.523	0.09	6 4378	28792			2016-10-23 02:27		46.497	0.0%
11443 3216		= 0	2016-10-23 02:	27	20.87 - Ke	mote pe	er	5 3 3	DE-C	.Ið	2016-10-23 02:27		47.348	0.0%
12846 13110		- 6	2016-10-23 02:	27	21.413	0.09	6 325	6830			2016-10-23 02:27		49.091	0.0%
19936 2586		- 0	2016-10-23 02:	27	21.564 Pre	esence i	1402	0 °C	ities		2016-10-23 02:27		50.998	0.0%
20059 8334	6939	= 0	2016-10-23 02:	27	22.586	0.04	6 17491	29535			2016-10-23 02:27		51.210	0.0%
10234 44050		= 0	2016-10-23 02:	27		ninge /E	6 2177	20712	2071		2016-10-23 02:27		53.083	0.0%
13251 31514		- 0	2016-10-23 02:	27		hing? (a		UD	es/c	ILY	2016-10-23 02:27		53.748	0.0%
20904 42148		= 4	2016-10-23 02:	27	25.559	0.09	6 16736	5617		- 6	2016-10-23 02:27		56.662	0.0%
1223 39102		- 4	2016-10-23 02:	27	25.645	0.00	6 1343	3320			2016-10-23 02:27		57.640	0.0%
24588 3301	1299	ι ο	2016-10-23 02:	27	26.966	0.00	6 28796	34156			2016-10-23 02:27		64.494	0.0%
27644 34867		– 6	2016-10-23 02:	27	28.429	0.00	6 3193	5089	10.1.17	••• ••	2016-10-23 02:27		67.737	
11629 15683		– 6	2016-10-23 02:	27	30.324	0.00	6 28569	5617	43447	- 0	2016-10-23 02:27		70.00	72
4534 2586		= 0	2016-10-23 02:	27	30.663	0.09	6 2114	6830		- 6	2016-10-23 02:27	00.750		24-28 DCT 2016