ImpROVement

Kevin Klercq
Willem Toorop
Koen van Hove

Yet Another Explanation of BGP

- Network of networks
- Interconnect them
- Organization = Autonomous system
- Announce your prefixes to your neighbors
- Not safe

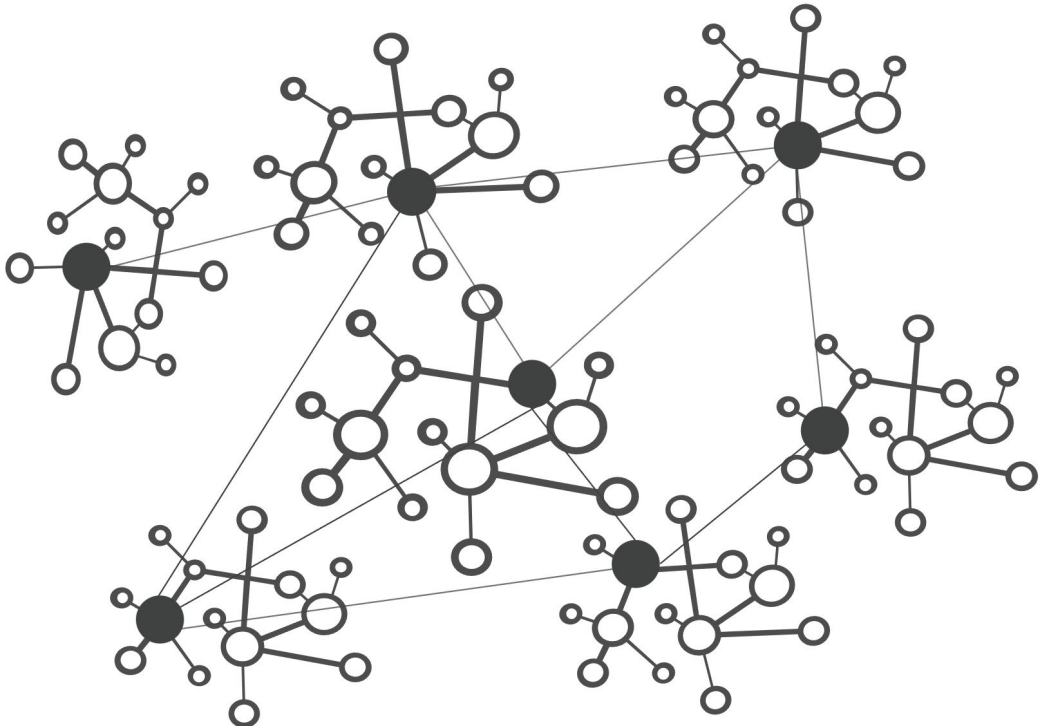

Arr Matey, This Be a BGP Hijack

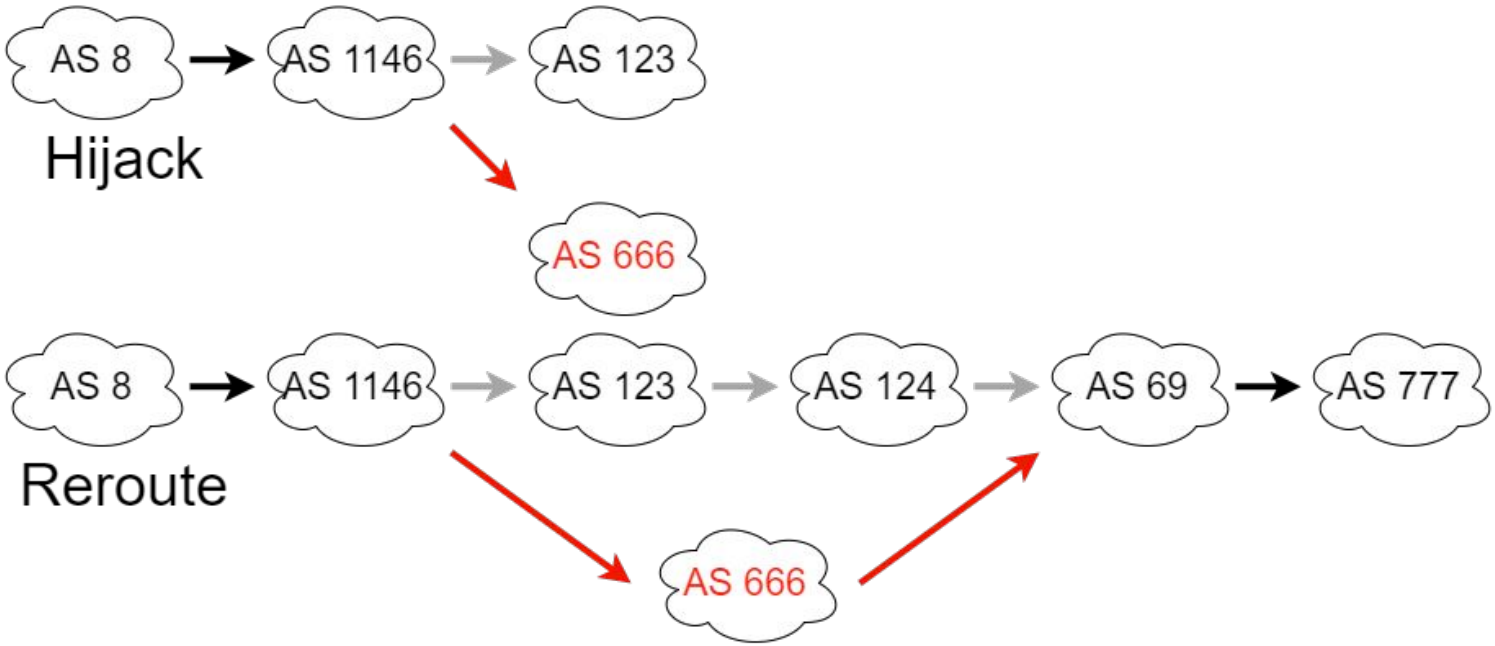

Hmm… What were those ROAs again?

- Resources (ASNs and prefixes) handed out by IANA to the 5 regional internet registries
- RIRs hand out to organizations

- Valid: ROA exists and everything is good
- Unknown: No ROA exists
- Invalid: ROA exists and received advertisement violates existing ROA (invalid prefix length, wrong AS)

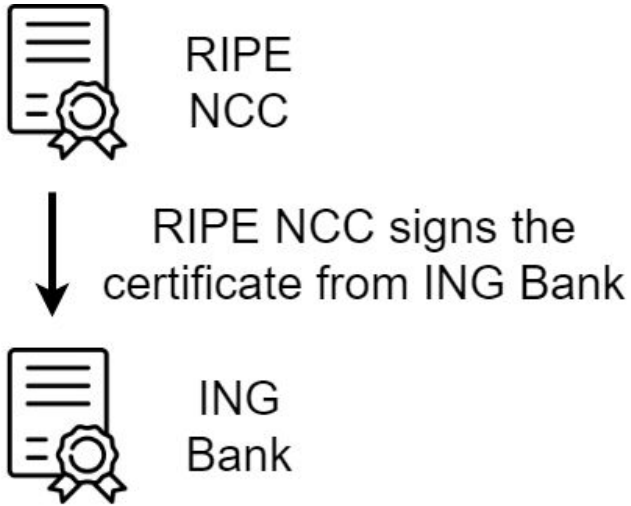

Can we determine which ASes really should start doing ROV?

Now For The Fun Part

- Two announcements. /23 and /24 (subnet of /23).
- /23 is valid and less specific
- /24 is invalid and more specific
- /23 is anycasted (Vultr, 30 locations), /24 in Amsterdam (ColoClue)

Key Observations

- Every router makes its own routing decisions
- A router that does ROV only routes both IP-addresses equally
- A router that does **not** do ROV likely routes both IP-addresses differently

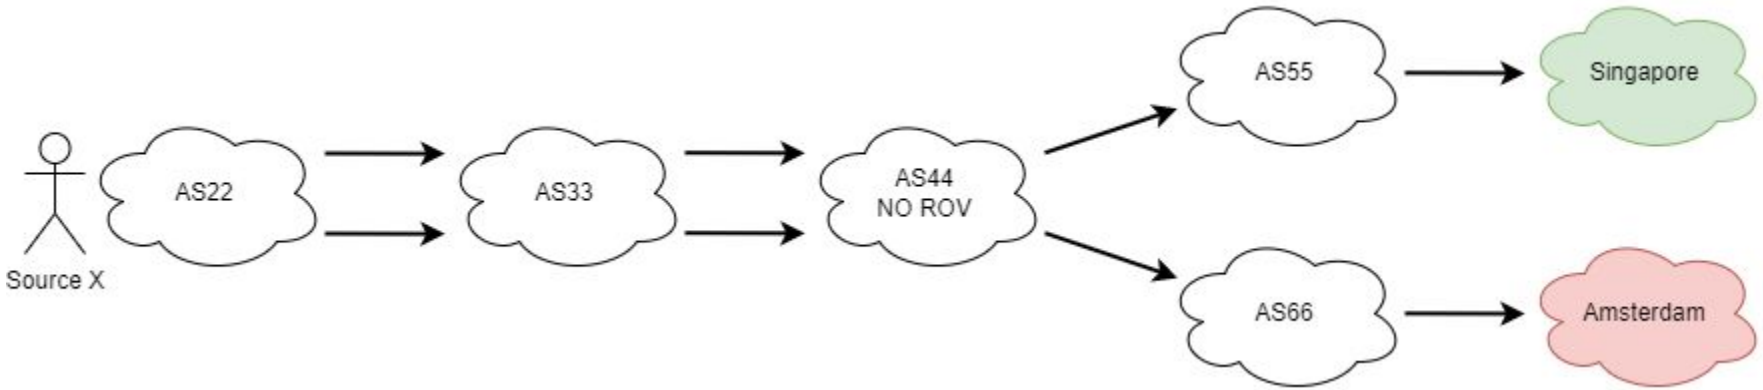

So, What Are The Results?

- Measurements on 2023-03-10 on RIPE Atlas with 12115 probes (Thank you Emile💖)
- Caveat: we only see the first AS on the path that misdirects.

	msm id	IPv6 equivalent
`dig @185.49.142.6 rpkitest.nlnetlabs.nl TXT +nsid`	50791569	50791565
`traceroute 185.49.142.6`	50791571	50791567
`traceroute 185.49.143.6`	50791572	50791568
`dig @185.49.143.6 rpkitest.nlnetlabs.nl TXT +nsid`	*50791570*	*50791566*

43%

... of IPv4 probes ended up at the invalid

(2288 out of 5285)

48%

... of IPv6 probes ended up at the invalid

(5539 out of 11442)

The impact if the top N ASes that currently do not do ROV would do ROV

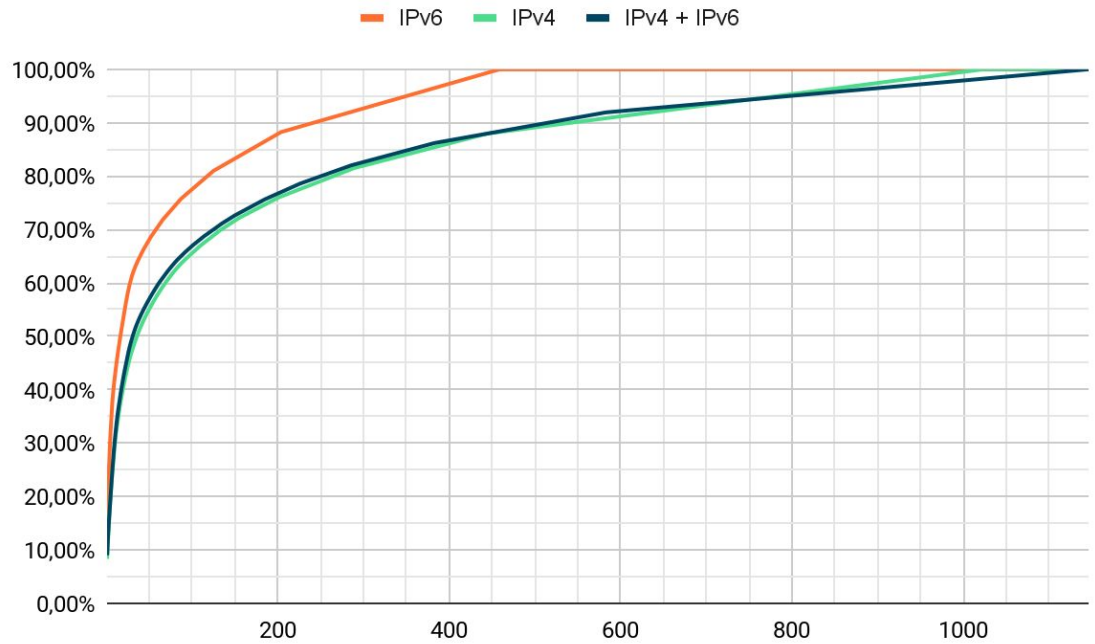

And the winner is...

🥇	**Telecom Italia (AS 6762)**	9%
🥈	**Free SAS (AS 12322)**	3.5%
🥉	**Vodafone (AS 3209)**	2.7%

IPv6

🥇	**Telecom Italia (AS 6762)**	10.5%
🥈	**Free SAS (AS 12322)**	9.3%
🥉	**NOS COMUNICACOES (AS 2860)**	4.8%

IPv4

🥇	**Telecom Italia (AS 6762)**	8.3%
🥈	**Superonline A.S. (AS 34984)**	3.3%
🥉	**PCCW Global, Inc. (AS 3491)**	3.1%

Percentage of total number of resolvers having affinity with a certain POP

Coloured area reaching the invalid

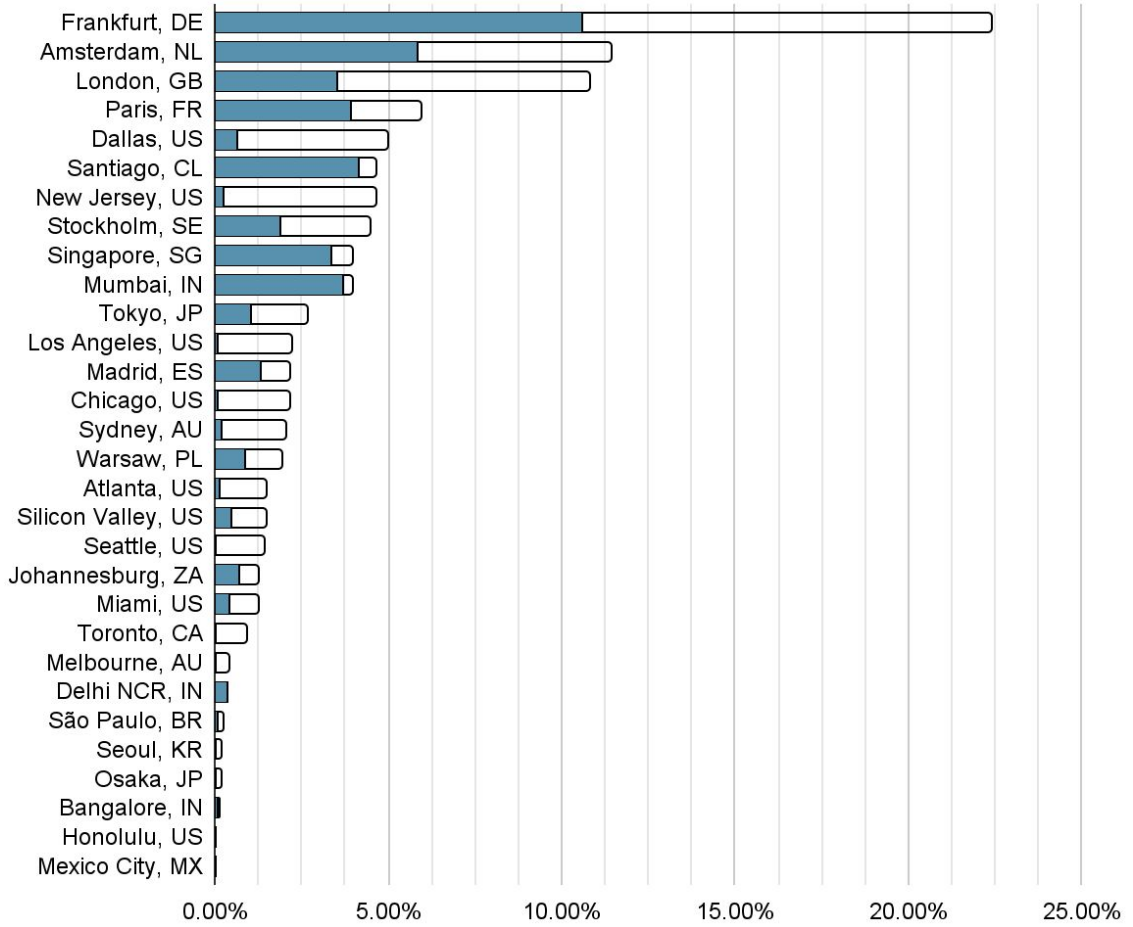

Traffic divergence per hop

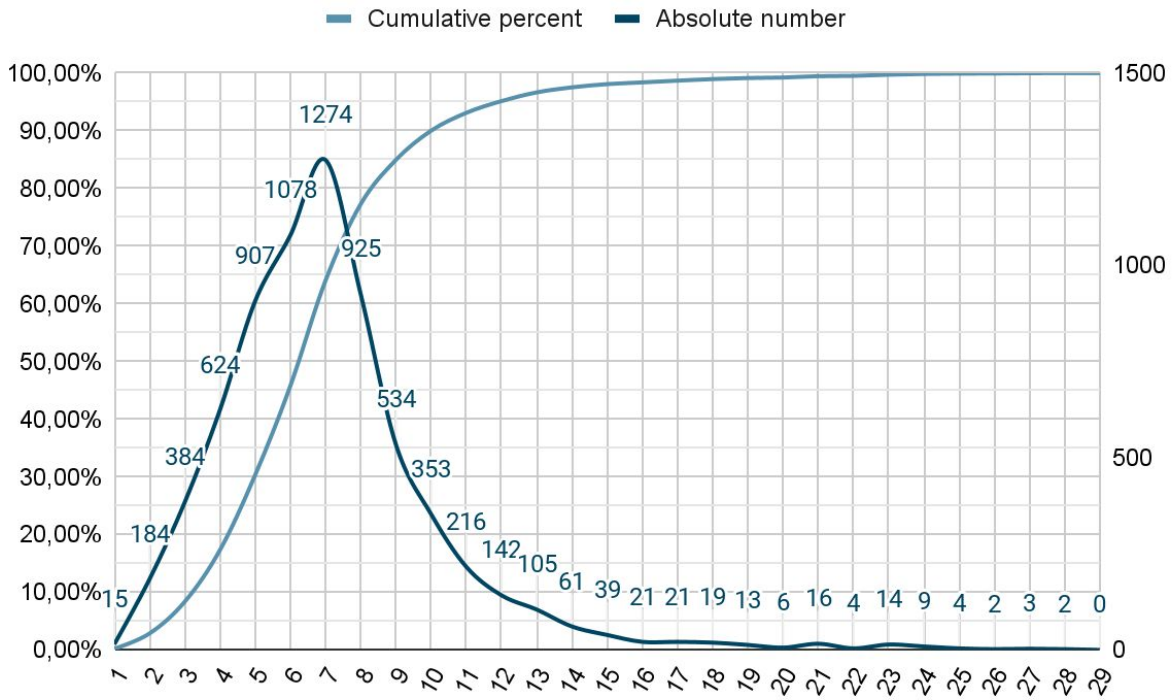

Cumulative diverging hop relative to valid path length

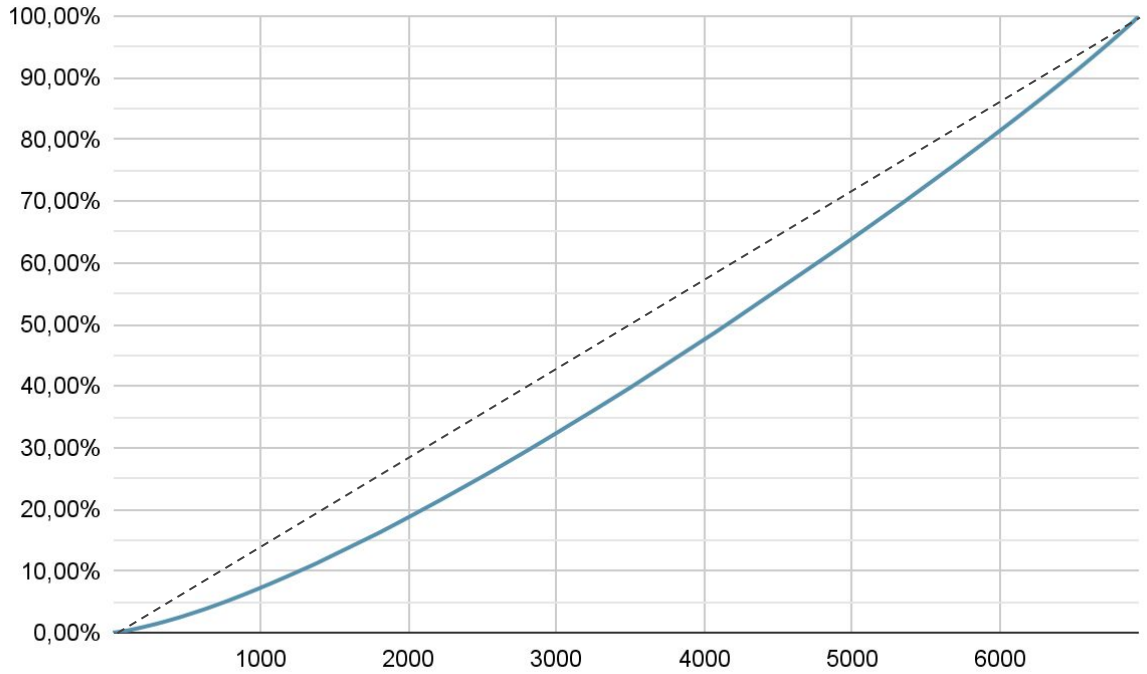

Conclusions

- A small group of organisations can have a big impact on routing security
- This can happen anywhere on the path
- Some POPs fare better than others

Percentage of total number of resolvers having affinity with a certain POP (IPv6 vs IPv4)

Coloured area reaching the invalid

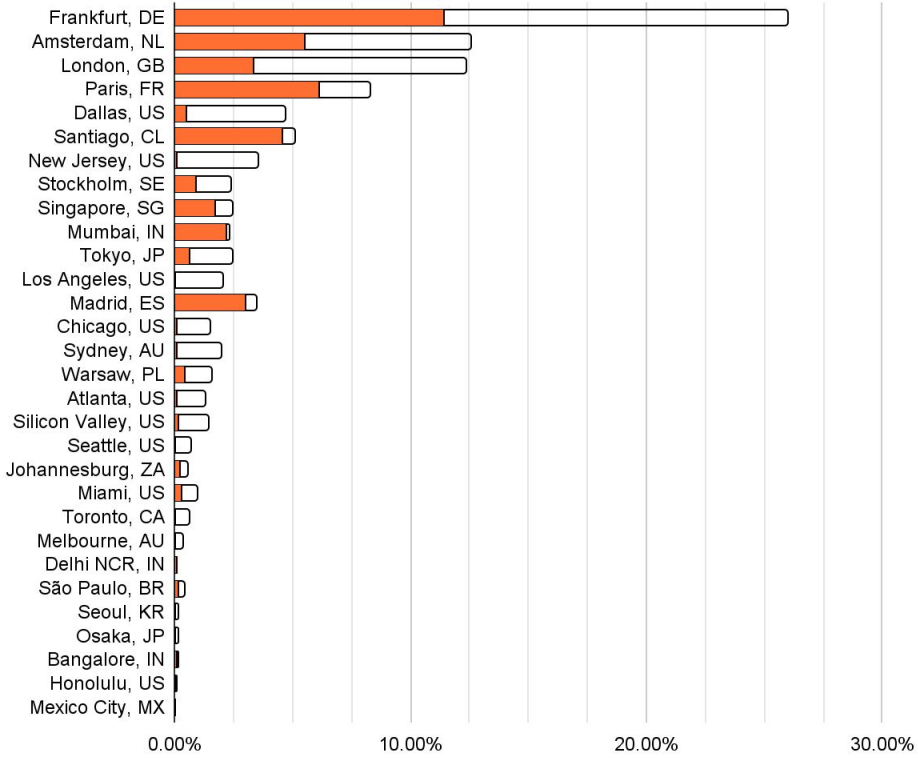

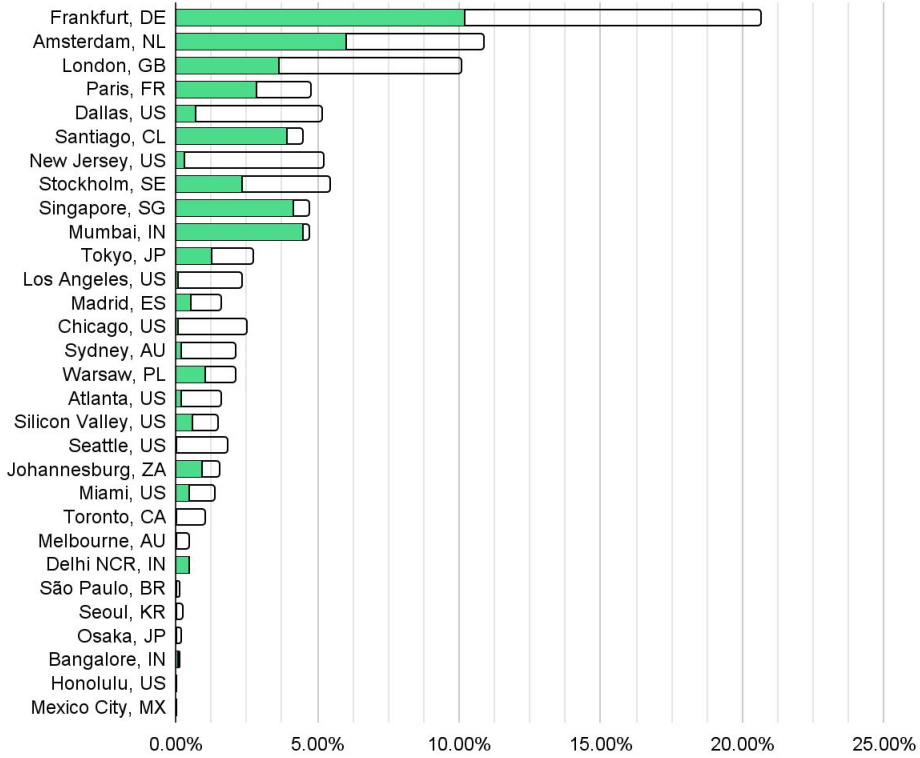

Percentage of traffic for a specific POP ending up at the invalid (sorted by most common POP)

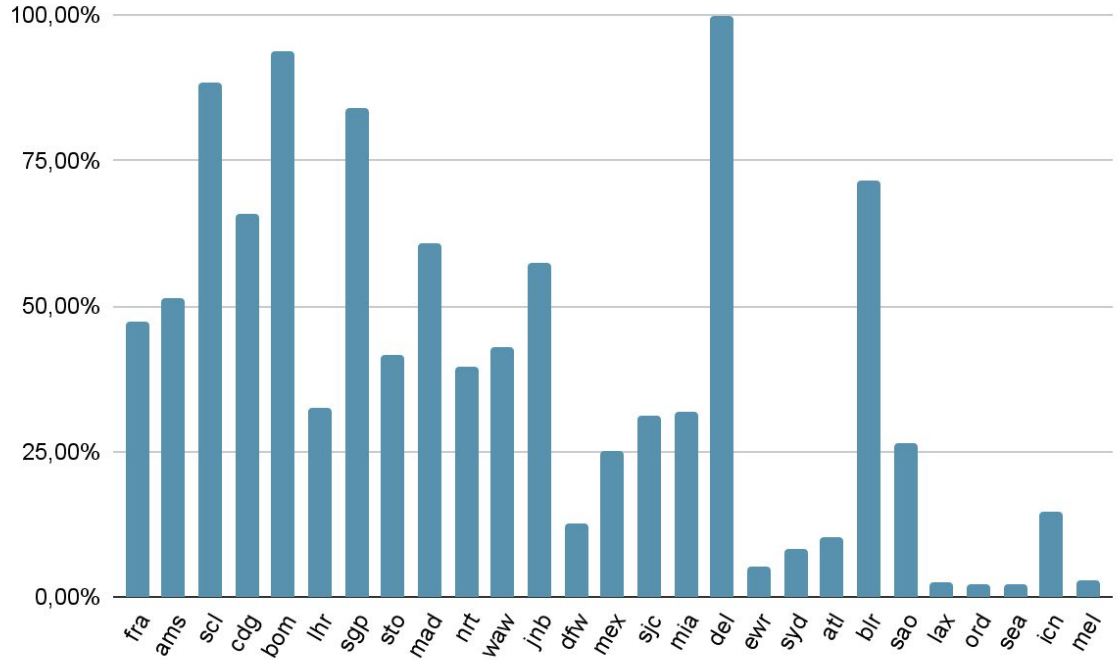

fra: Frankfurt, DE
ams: Amsterdam, NL
scl: Santiago, CL
cdg: Paris, FR
bom: Mumbai, IN
lhr: London, GB
sgp: Singapore, SG
sto: Stockholm, SE
mad: Madrid, ES
nrt: Tokyo, JP
waw: Warsaw, PL
jnb: Johannesburg, ZA
dfw: Dallas, US
sjc: Silicon Valley, US
mia: Miami, US
del: Delhi NCR, IN
ewr: New Jersey, US
syd: Sydney, AU
atl: Atlanta, US
blr: Bangalore, IN
sao: São Paulo, BR
lax: Los Angeles, US
ord: Chicago, US
sea: Seattle, US
icn: Seoul, KR
mel: Melbourne, AU
mex: Mexico City, MX