Quick notes from DITL 2009

George Michaelson

APNIC

ggm@apnic.net

Madrid: the coolest city in the world!

APNIC

- Regional Internet Registy (RIR) for the Asia-Pacific region
- provides allocation and registration services which support the operation of the Internet globally. It is a not-for-profit, membership-based organisation whose members include Internet Service Providers, National Internet Registries, and similar organisations. APNIC represents the Asia Pacific region, comprising 56 economies.

APNIC's DNS

- RIR's are the delegation point for in-addr, ip6 .arpa.
  - APNIC serves the Asia-Pacific, has secondary servers for some other RIR on an as-needs basis to provide NS diversity in AP region.
- DNS @ APNIC, two 'flavours'
  - The 'NS' hosts
    - APNIC's primary NS for its in-addr.arpa obligation
    - The entire Asia-Pacific managed IP address space
  - The 'SEC' hosts
    - Secondary NS for the other RIR (AfriNIC, LacNIC, RIPE)
    - A range of ccTLD, other forward namespaces of interest
- 3 locations: Brisbane, Tokyo, Hong Kong
  - Co Located, 100mbit switching fabric, good local connectivity

APNIC DNS data collection

- On-host TCPdump data collection since 2002
  - Sample based (15min interval)
  - Could not scale, affecting core service reliability
- Deployed new collection system for DITL 2008
  - passive tap of each DNS server, process locally and export abstracts
- Result: Continuous capture of APNIC DNS
  - 3 days back archive held at each location
    - Was designed to be more, but data grew faster than expected.

Presenting on PTR data only

- Problem: DNS servers take a lot of 'bogus' queries
  - Misconfigured clients can be promiscuous in what they ask, who they ask
  - Spurious domains like .local, malformed queries
  - DNS-internal queries like SOA refresh
  - Forward domain queries in ccTLD space
- Focus on PTR requests on all servers
  - Only tracking the in-addr.arpa and ip6.arpa queries at this time.
- Permits economy-to-economy measurement
  - Src IP address to in-addr.arpa or ip6.arpa value.

Data capture setup

- Datacomm SINGLEstream passive tap
  - Copper or fiber, 100mbit, 2 feeds of tapped net
  - Dual-redundant power capable
  - 1 packet switch-time to unpowered
    - Completely ‘fail safe’ to the tapped link

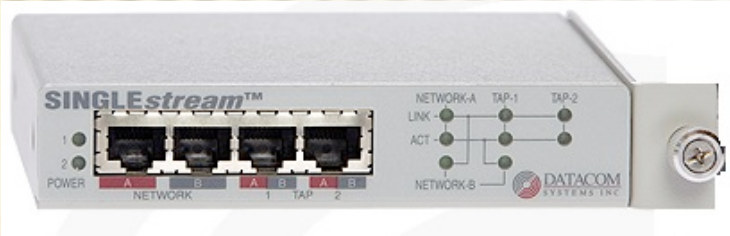

- Fed to data collection host
  - Dell 860, extra ethernet card added
    - Two gig-E dedicated to 2 distinct tapped DNS servers
    - Redhat EL5, pcap based capture mechanism
- Using DSC, dnscap from ‘the measurement factory’ & ISC

Day In The Life

- Continuous packet capture of DNS servers, IX, other places of interest
- Organized by CAIDA/OARC
- Provides resource for longer term analysis
  - Opportunity for retrospective/review of data
- First collection 2006
  - 4 DNS participants, selected campus/local IX
- Fourth event (March 29-April2) just completed
  - 37 participants, ~190 nodes of collection
  - Of the order 4Tb data (!)
- APNIC contributing since 2008 from all operated DNS servers
  - Due to secondaries at the other RIR, this only represents a subset of APNIC NS serve for its own domains.

Participants

- afilias apnic arin arl as112-gf
- brave caida camel cira cogent
- cznic everydns icann iis isc
- isi lacnic level3 namex nasa
- nethelp niccl nixcz nominet nrcca
- oarc orsnb pktpush qwest regbr
- ripe switch ultradns uninett uniroma2
- verisign wide

AP DiTL Data Capture (gb)

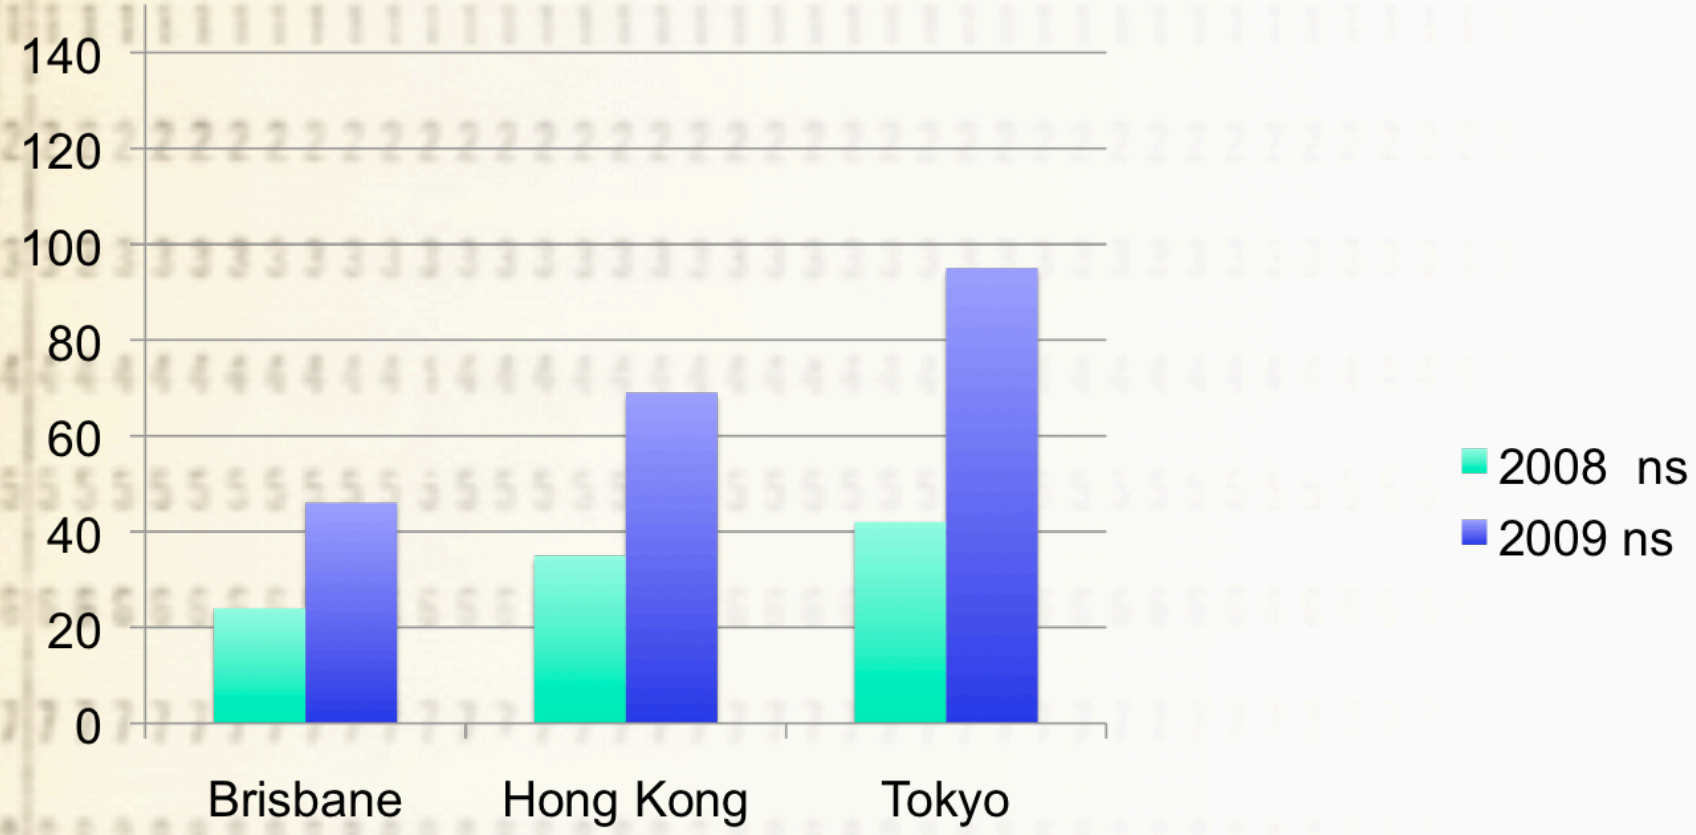

Average Query rate q/sec

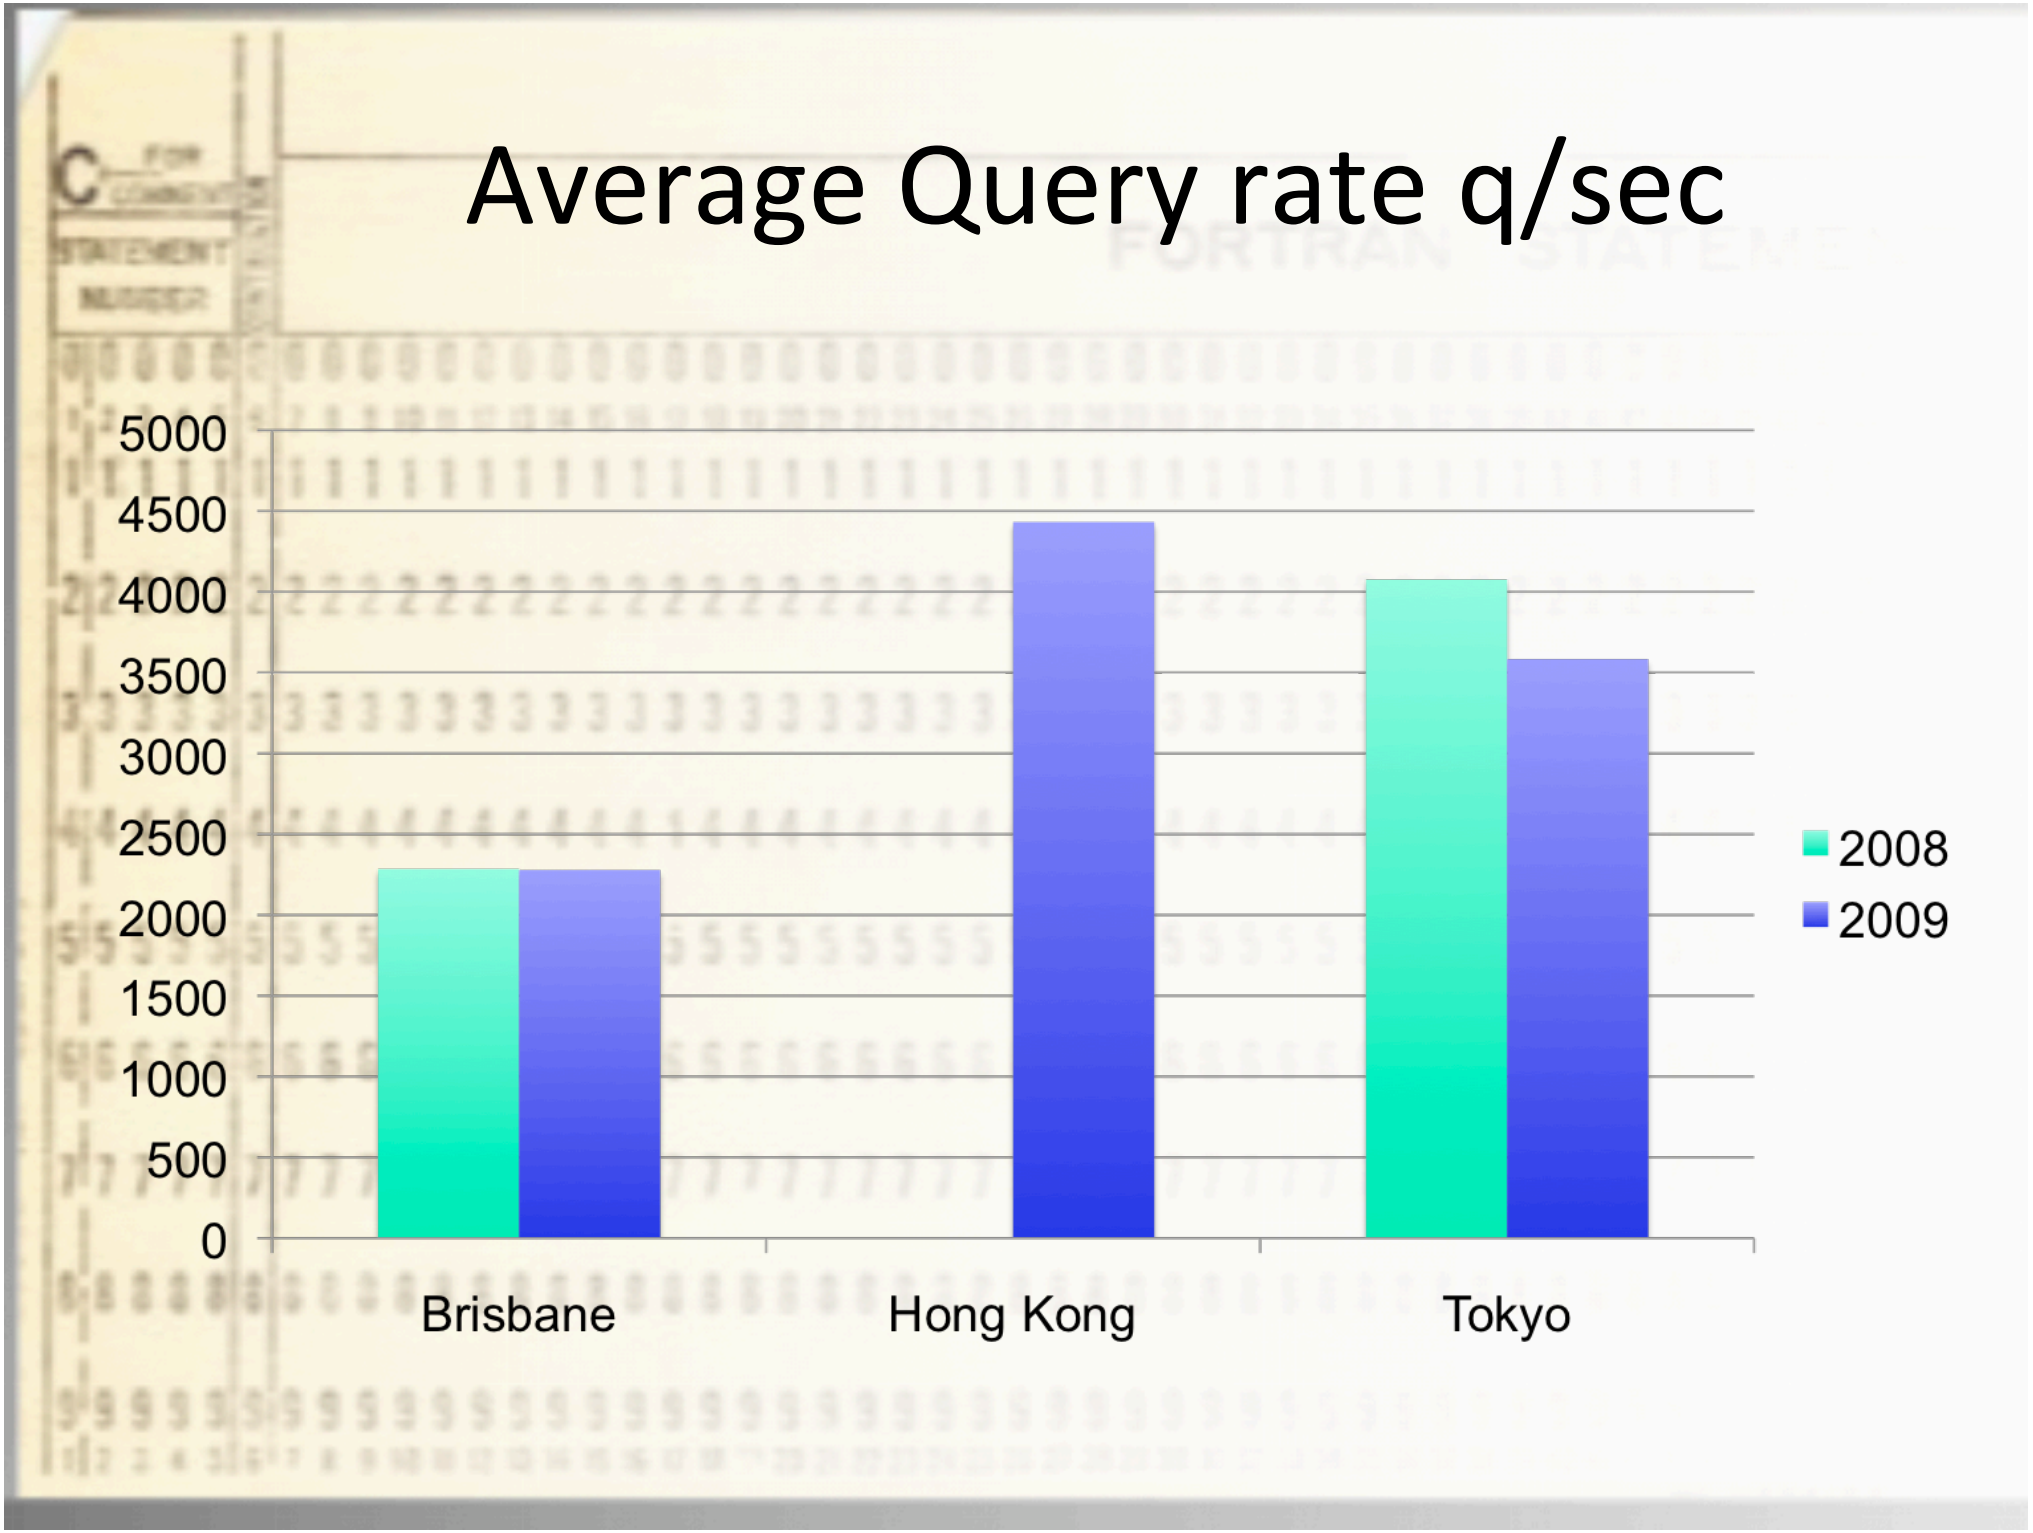

Brief quiz

- If you had DNS in 2008…
- Would you use the same IP address to do DNS in 2009?
  - (I would: I don't change my resolver that much)
- How many unique IP addresses seen in 2008 do you expect to see in 2009?
  - (I expected to see a lot. The majority in fact)
- ……..

Unique Ips in 24h

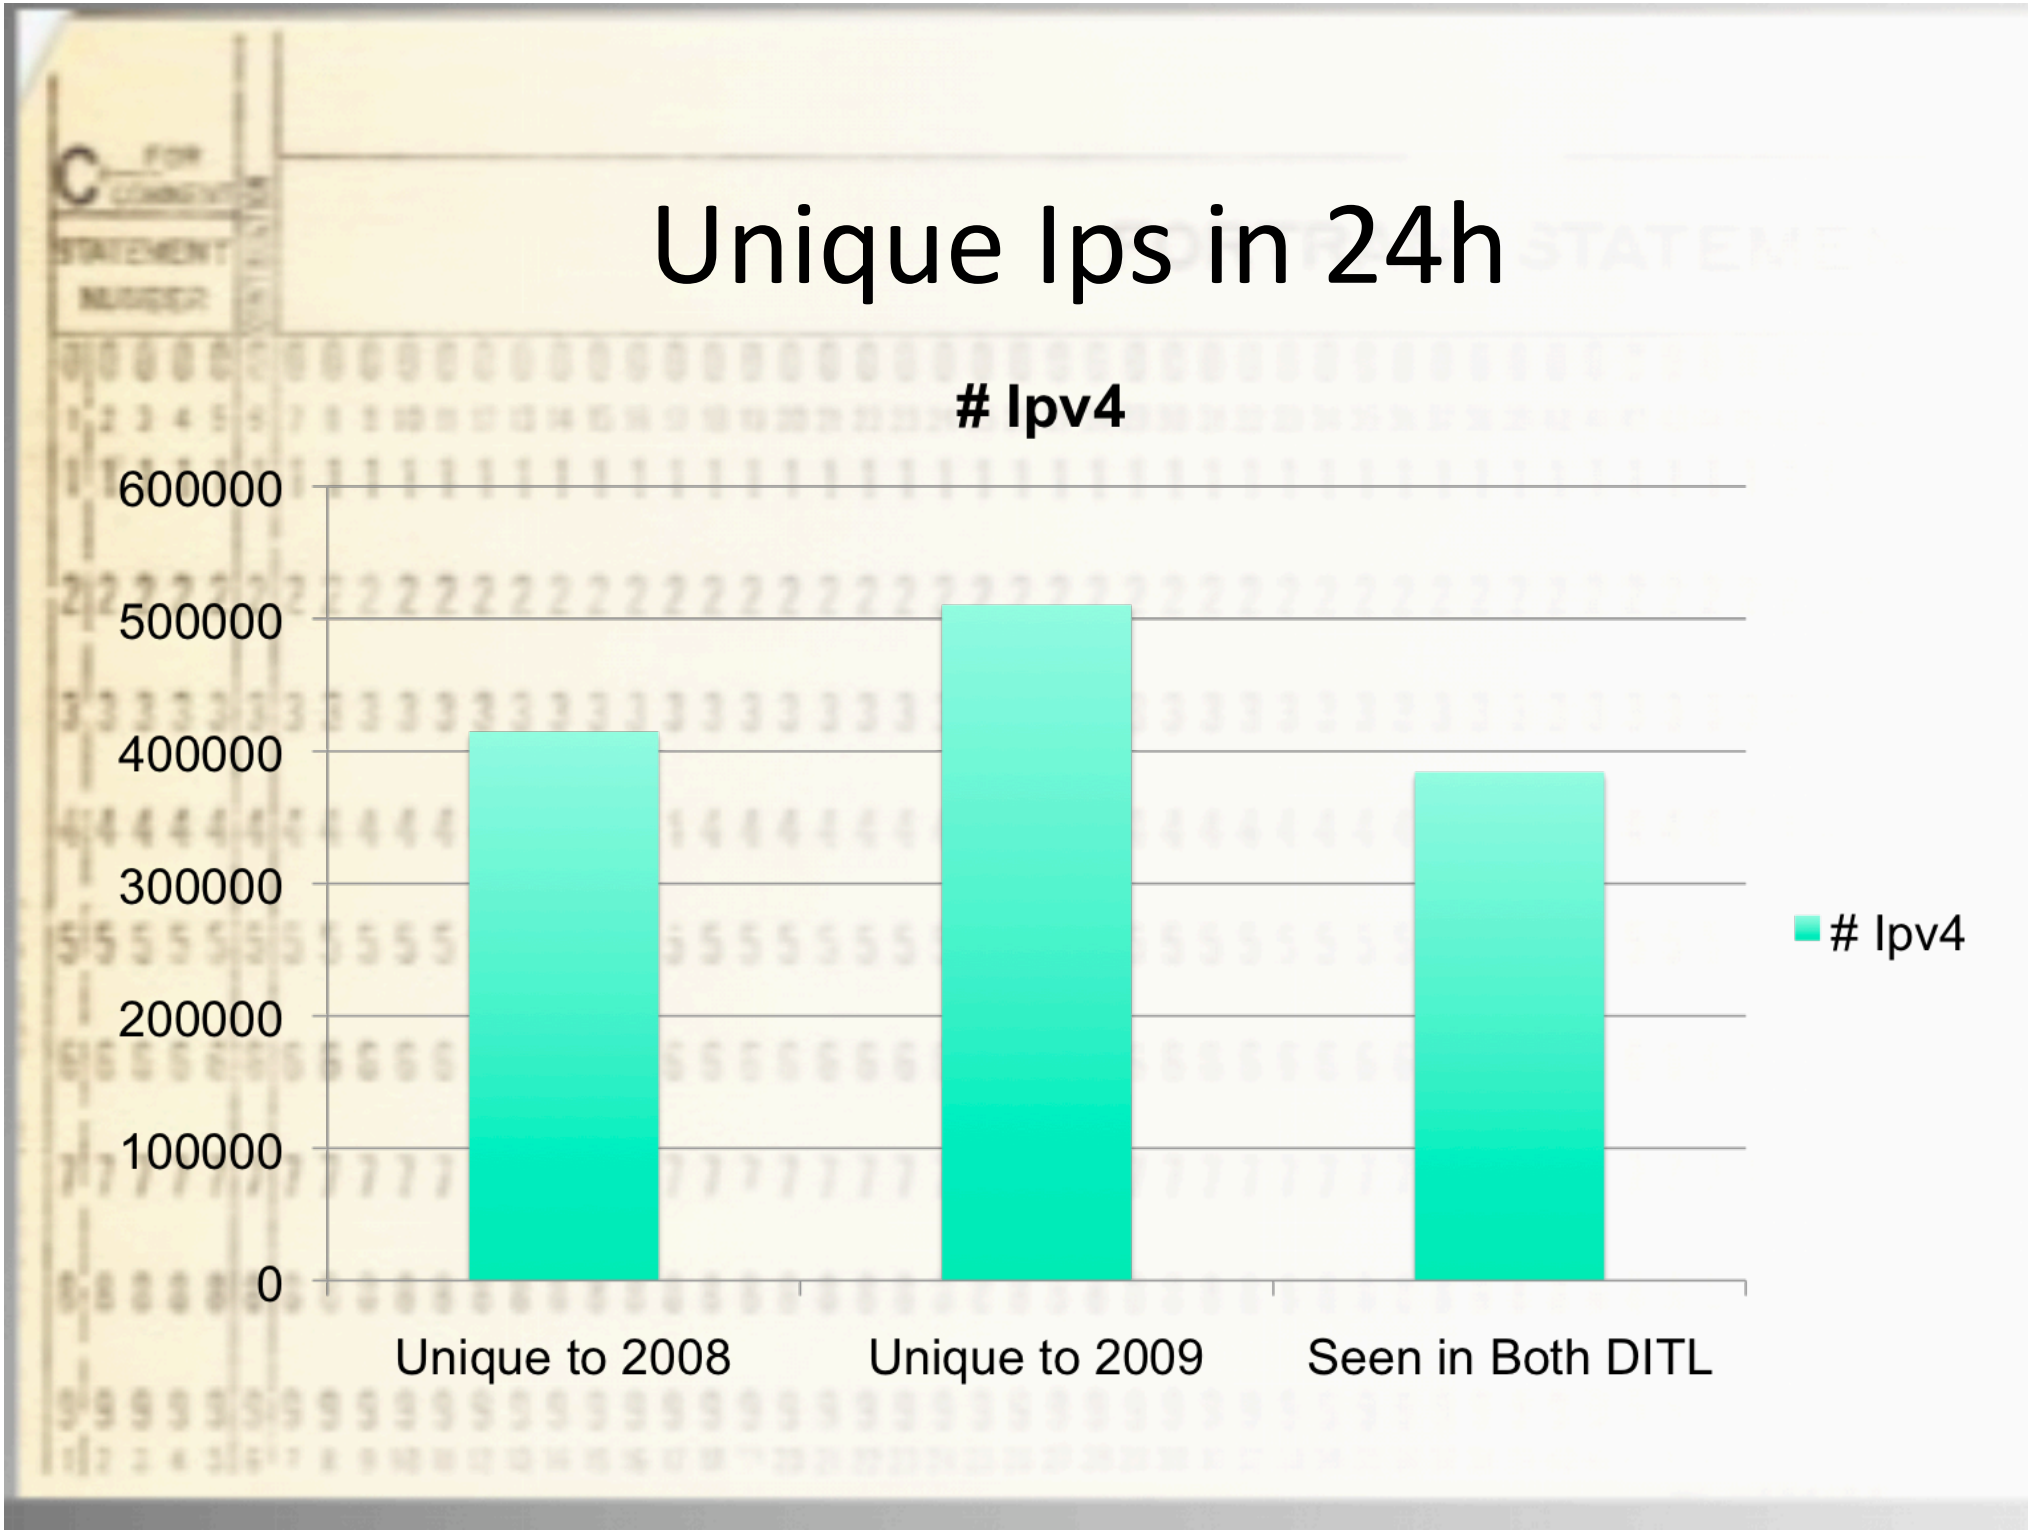

Not a lot of Address re-use

- Slightly less than 1/3 of the IP addresses seen, were seen the year before.
- Seems counter-intuitive: infrastructure DNS is believed to be machine driven, and from company/internal DNS servers, resolvers
  - Which are expected to be on stable IP addresses
- For further study

Brief quiz

- If a DNS server queries for reverse-DNS…
- Would you not expect it to query for a lot of reverse DNS?
  - (I would: applications which do reverse seem to do a lot)
- What sort of curve-shape of #lookups do you expect?
  - (I expected to see a lot of lookups from most hosts. The majority in fact)
- …….

How many times do Ips query?

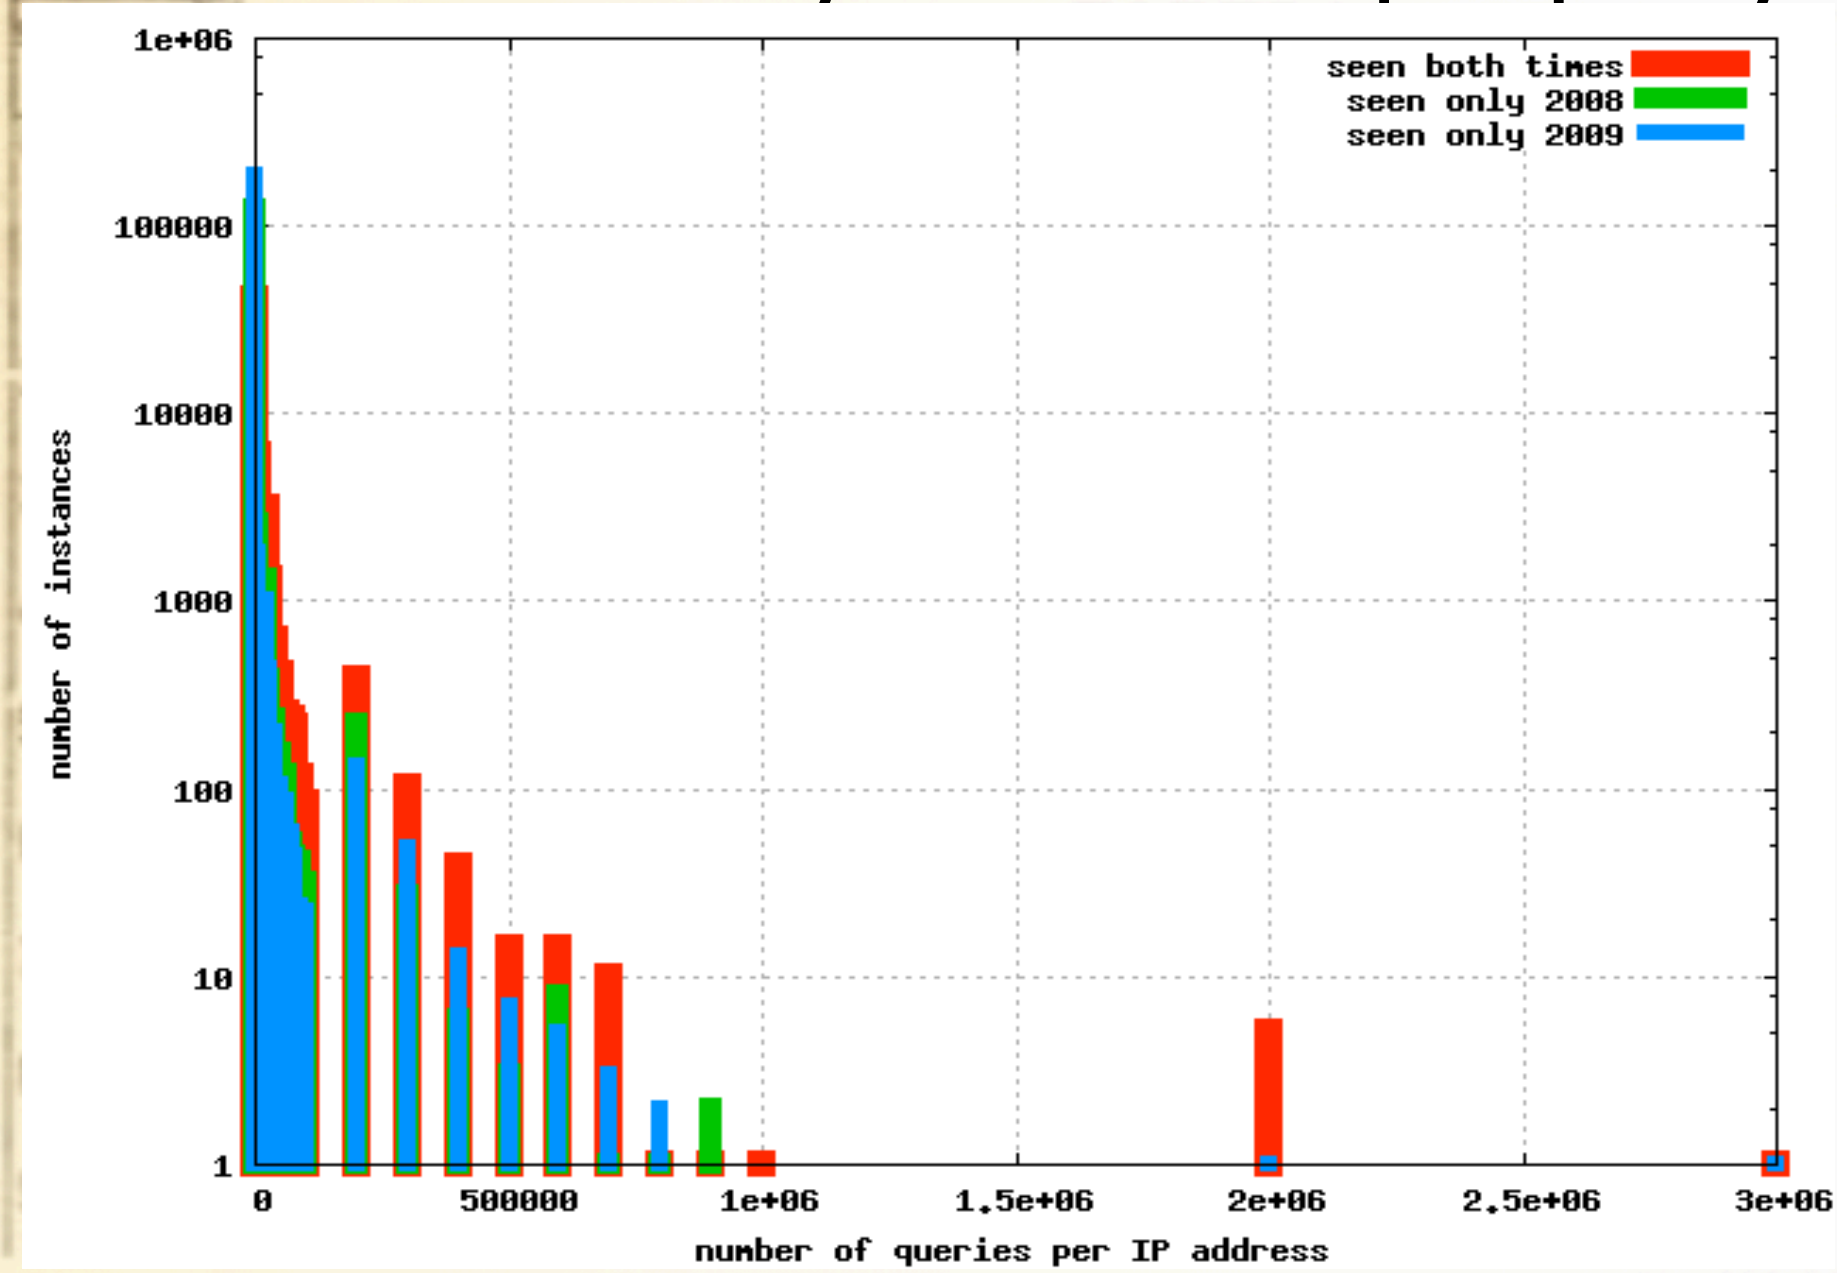

This is strange…

- The majority of seen IP addresses do 1, or a few queries. Only a very few addresses to hundreds of thousands, and <10 do millions.
- PTR: 'infrastructure' DNS.
  - If its infrastructure, why so much volatility in the IP addresses doing DNS querying?
  - Expected to see far more persistent IP addresses across 2008/2009.
- End-user boxes doing reverse-DNS?
  - Firewalls, probe-tests, other applications?
  - For further study.
- Suggests the 'real' count of infrastructure resolvers hitting APNIC is lower than thought
  - <millions. Most hits from 'singletons'

V4/V6

- Some Infrastructure DNS now flows over V6
  - Some even flows over tunneling technology
- Might indicate V6 uptake
- Freenet 6rd suggests deploying 6to4 internally can encourage uptake
- Signs of Increased V6 usage
  - But not enough to head off a problem in the context of V4 exhaustion.. Yet.
- Rather pretty 10:100:1000:10000 ratio.

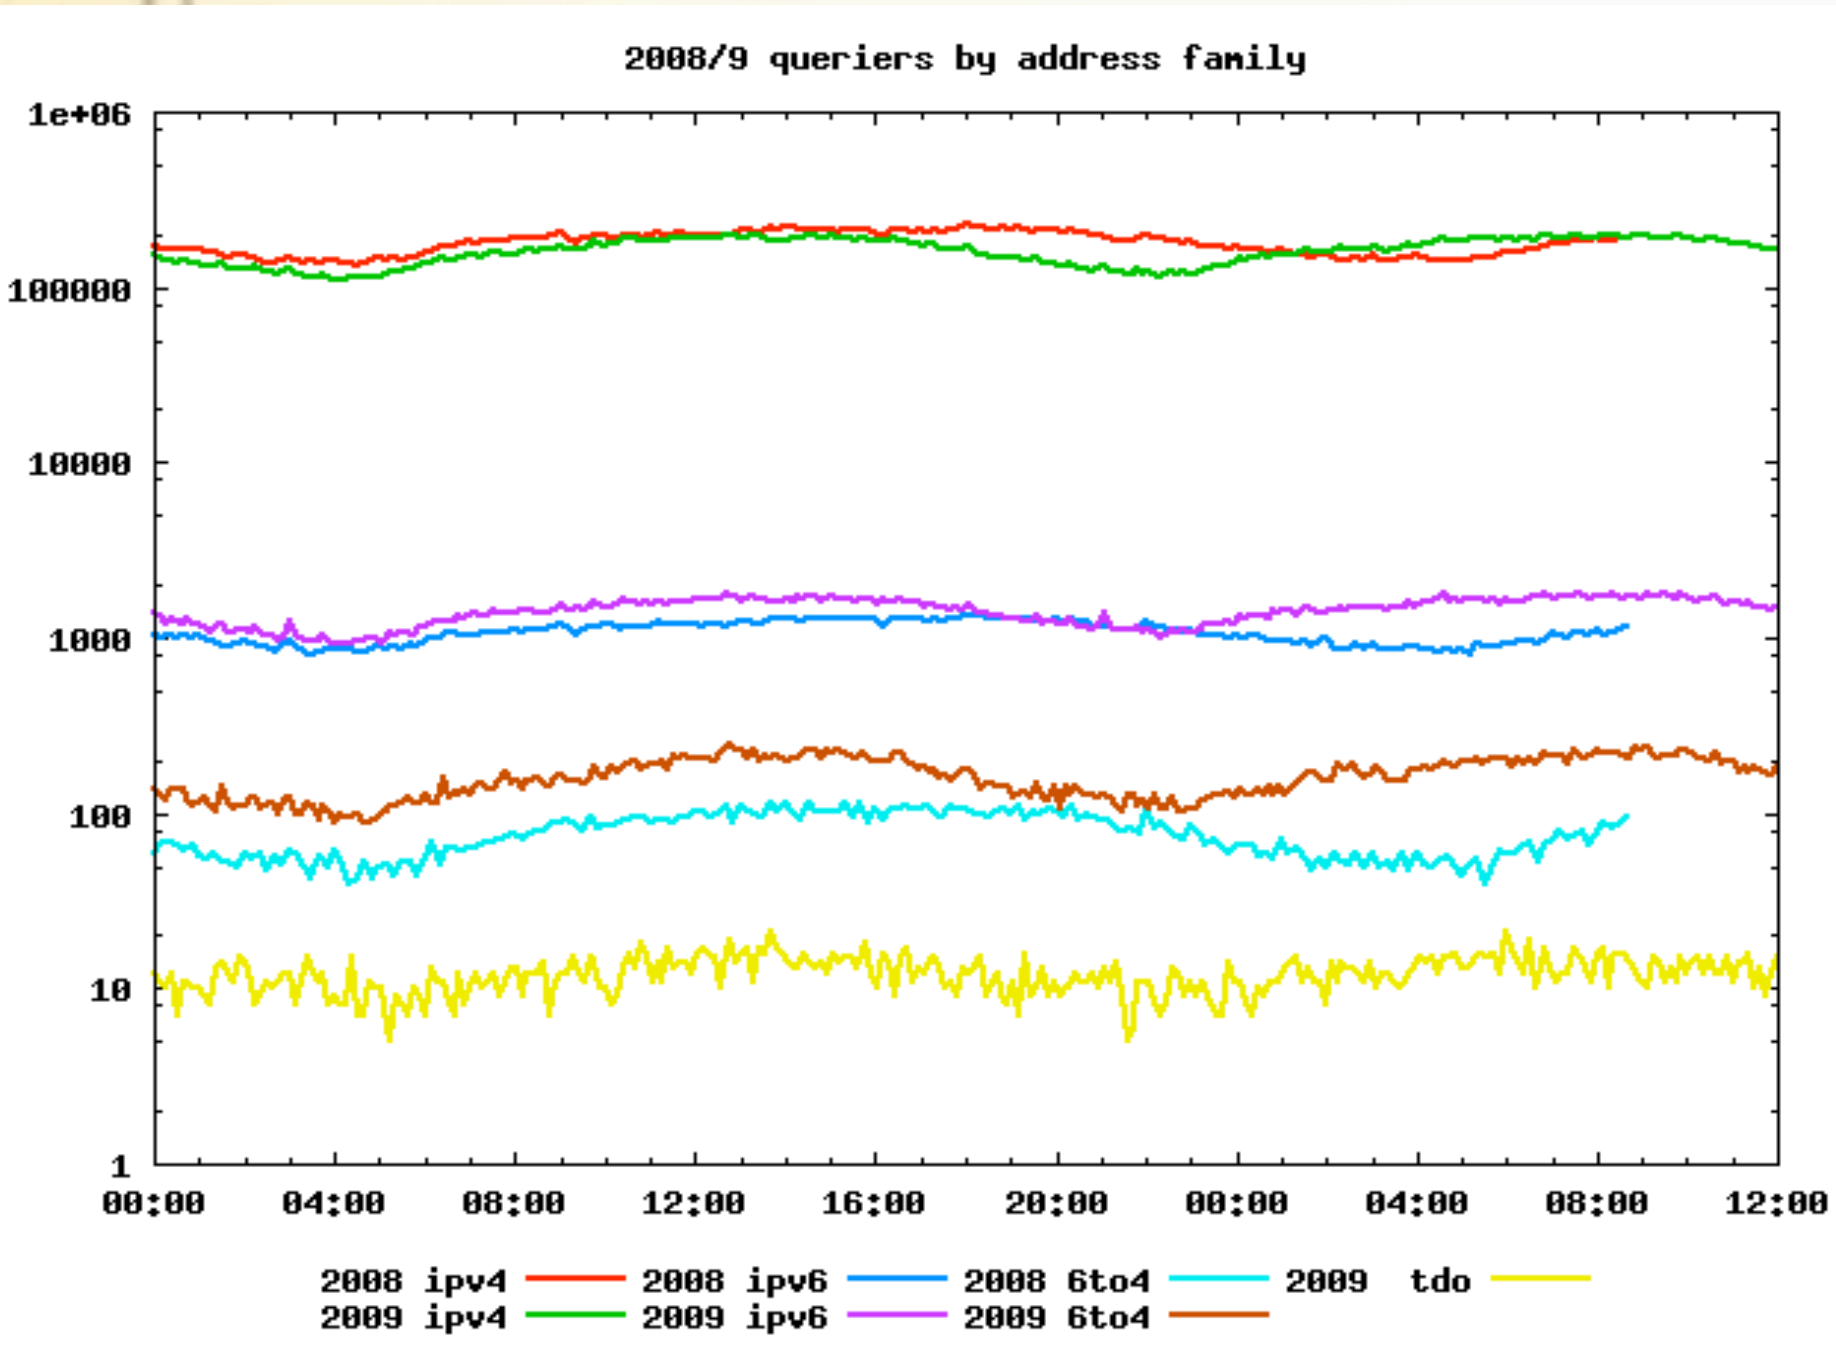
2008/9 queriers by address family
1e+06
100000
10000
1000
100
10
1
00:00
04:00
08:00
12:00
16:00
20:00
00:00
04:00
08:00
12:00
2008 ipv4
2008 ipv6
2008 6to4
2009
tdo
2009 ipv4
2009 ipv6
2009 6to4

Tunneled V6 for DNS?

- Strong evidence the Teredo DNS is p2p
  - Clients embed DNS resolver, do reverse-DNS on display of peer sets (N.Ward, Google-IPv6 workshop)
- Not a good choice for service dependency!
- 6to4 very likely to be combination of
  - Linux/FreeBSD
  - Mac, eg airport @home and other OSX 6to4

Is it any different in ES?

- No reason to assume it is or it isn't
- Little traffic from ES hits A-P regional DNS servers
  - You have RIPE-NCC hosted alternates at *far* shorter RTT on your doorstep
- Use a 563 prefix filterlist, python radix-tree filter.

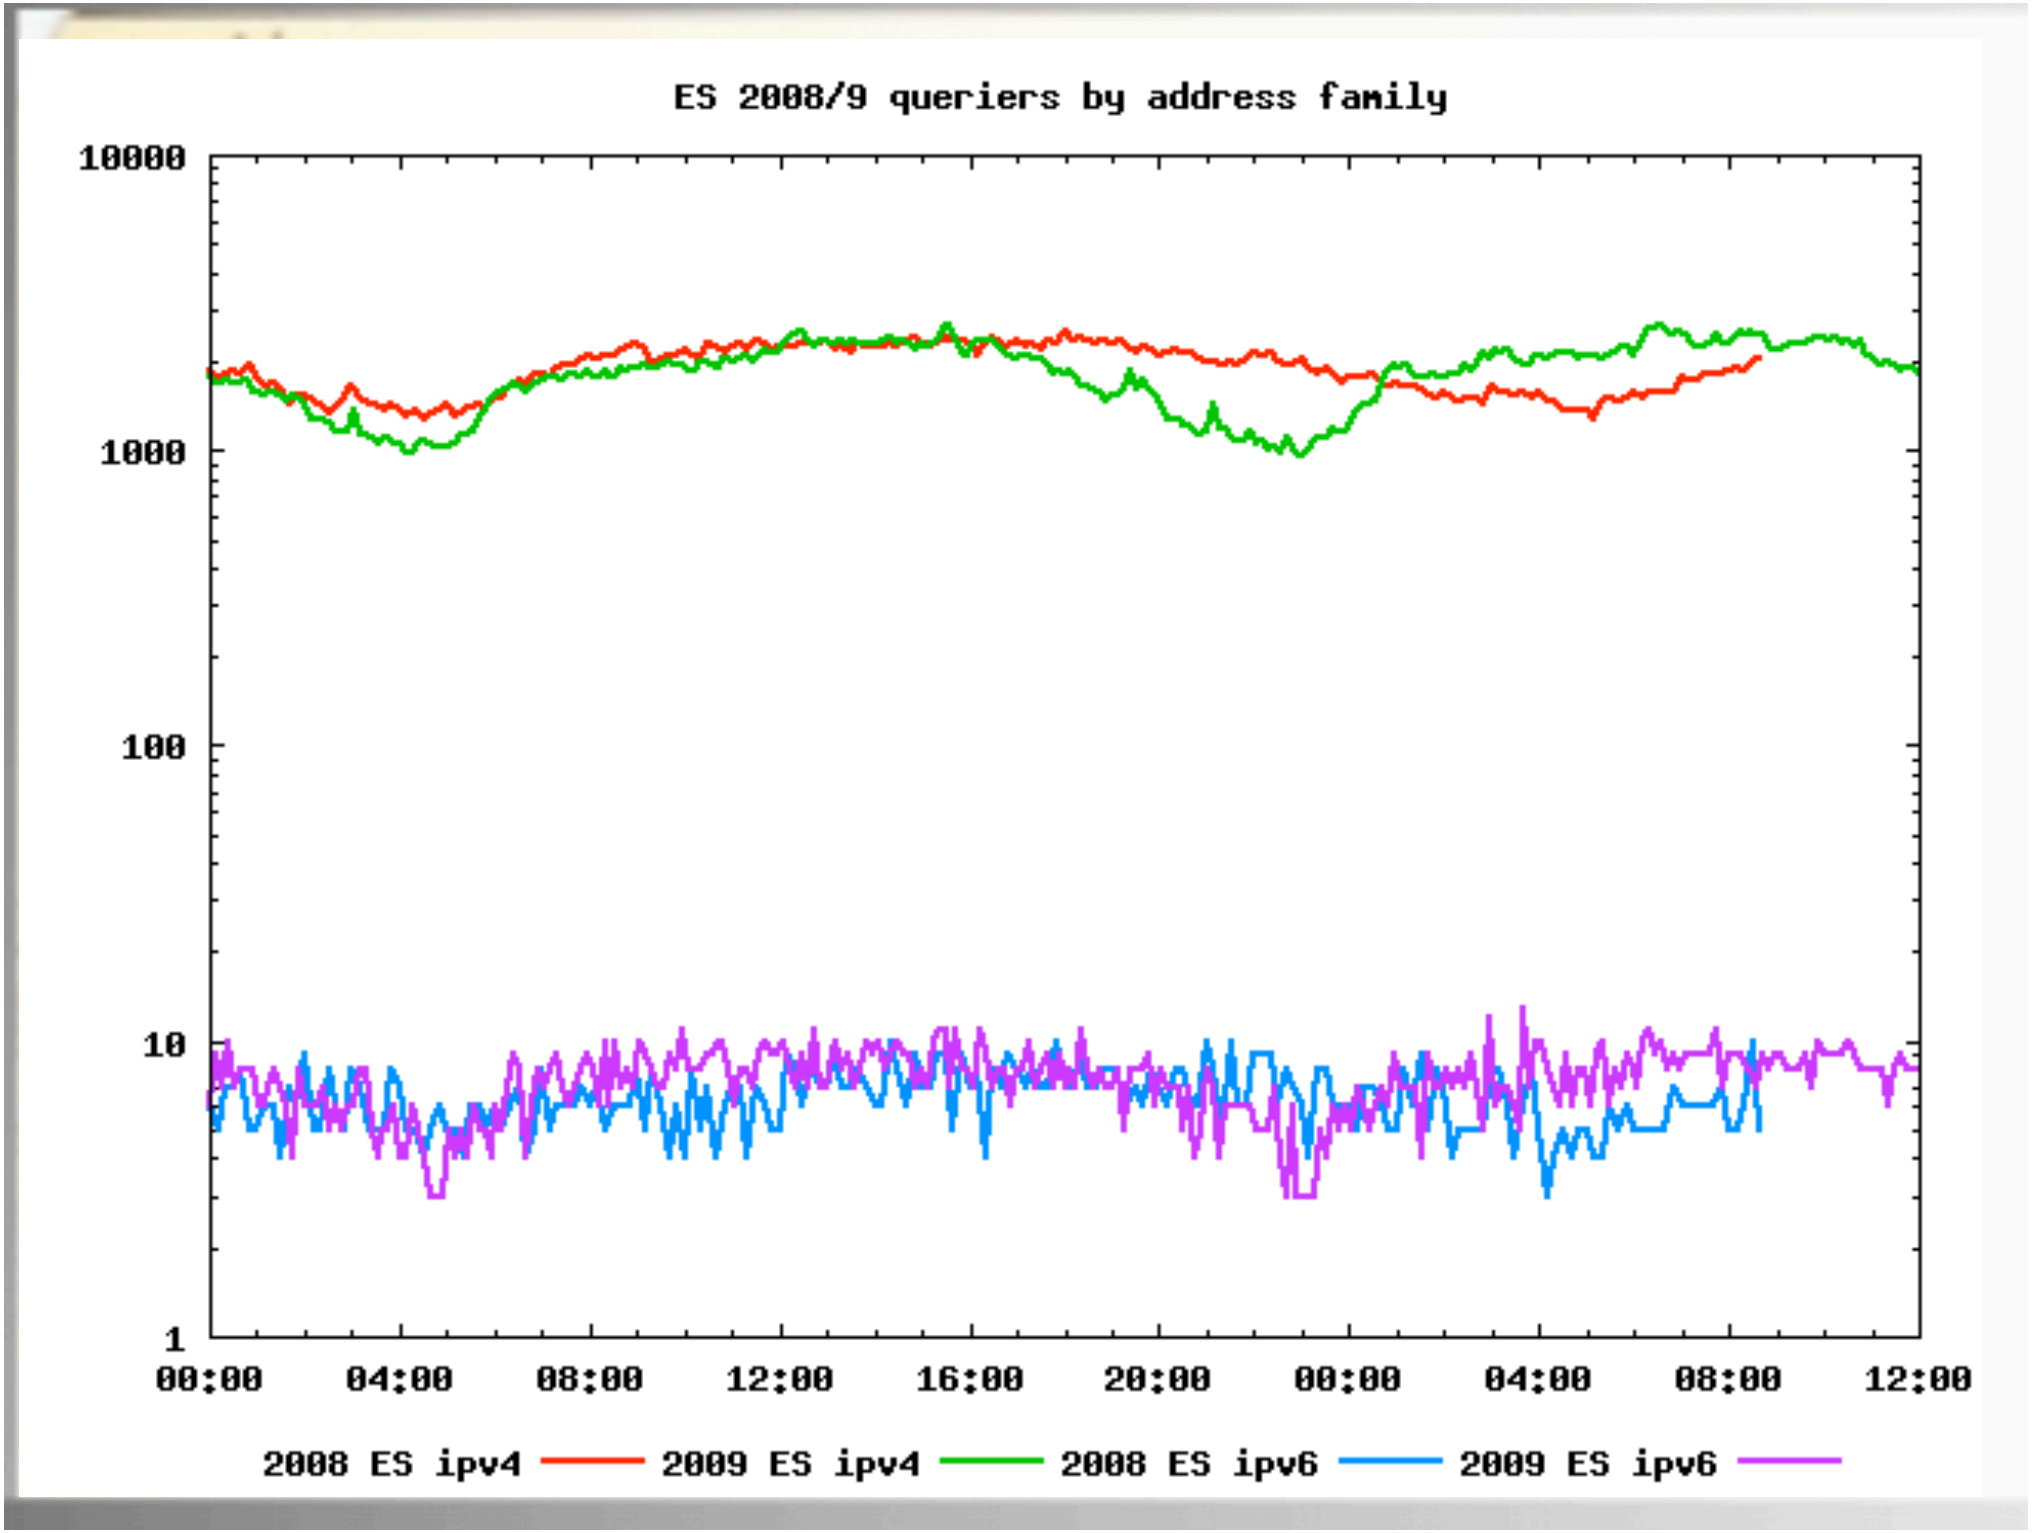
ES 2008/9 queriers by address family
10000
1000
100
10
1
00:00
04:00
08:00
12:00
16:00
20:00
00:00
04:00
08:00
12:00
2008 ES ipv4
2009 ES ipv4
2008 ES ipv6
2009 ES ipv6

ES DNS v4/v6

- No noticeable increase 2008-2009.
- Ratios stayed the same 1:100 v6 to v4
  - No 6to4/teredo in filter list (they derive from the V4 address, but are ‘mangled’) at this time
  - (indication from global picture is that 6to4/teredo are of the order 10% and 1% of total V6)
- Ratio of ES v4 to RoW v4, 1.12%. V6, 0.46%
  - Is this plausible, given size of ES address blocks?
  - (using other data, I expect 2.66%/0.012%)
  - This suggests ES is ½ size in V4, 40x in V6
    - V6 tunnels distort.
  - Per-economy metrics for further study.

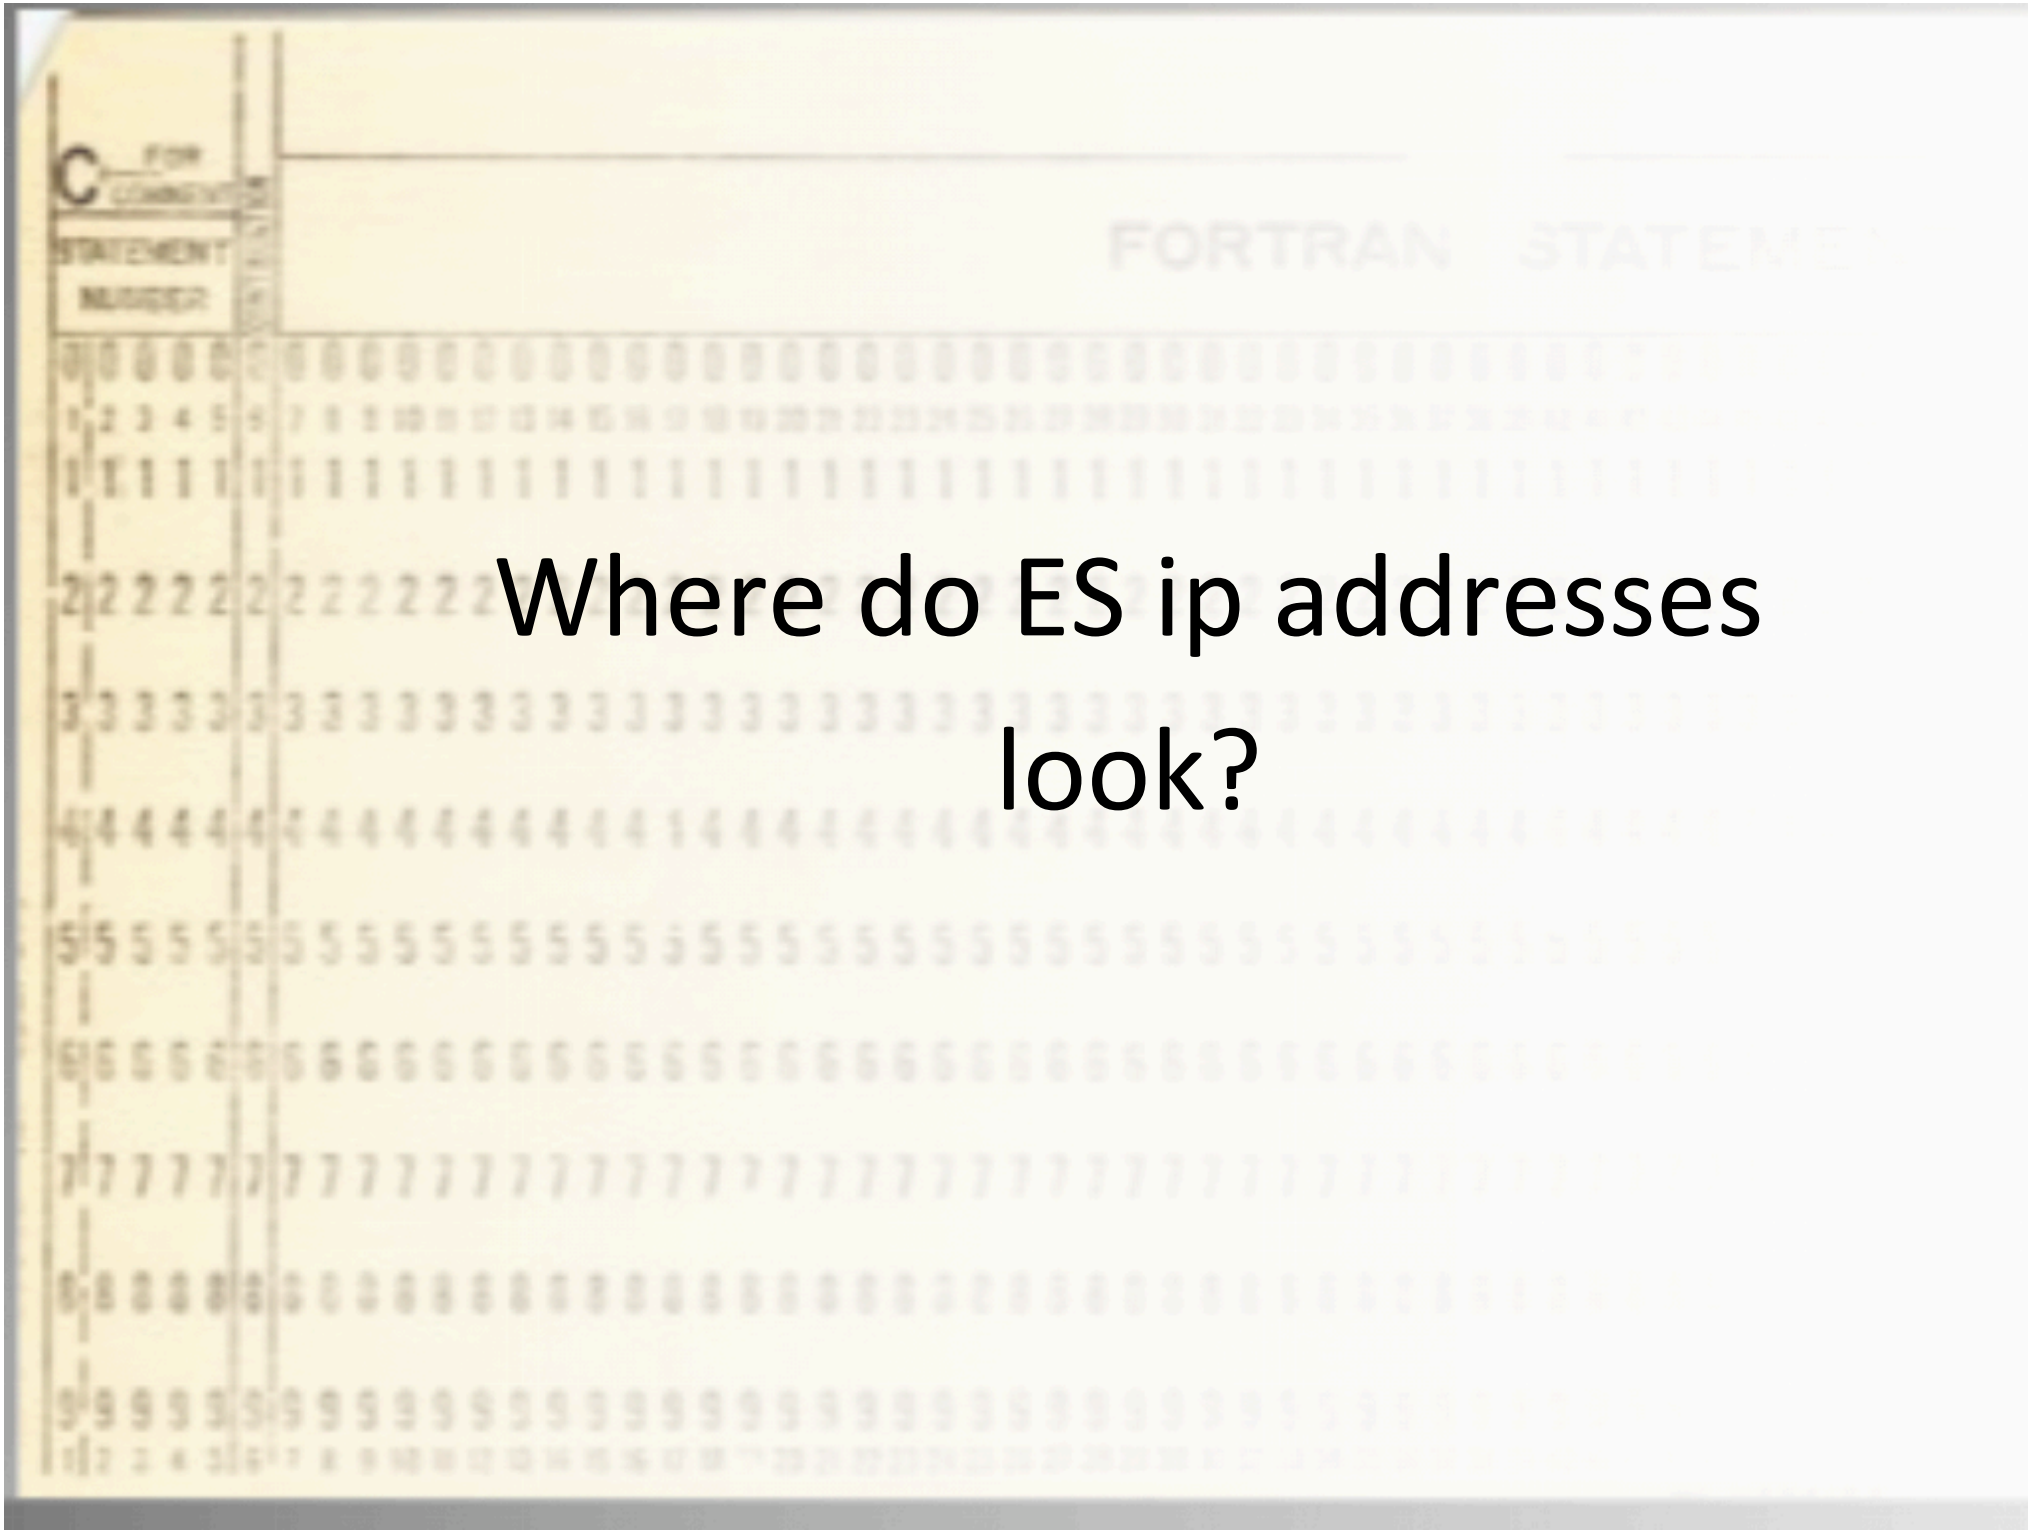
Where do ES ip addresses look?

Where in Asia-Pacific does ES look?

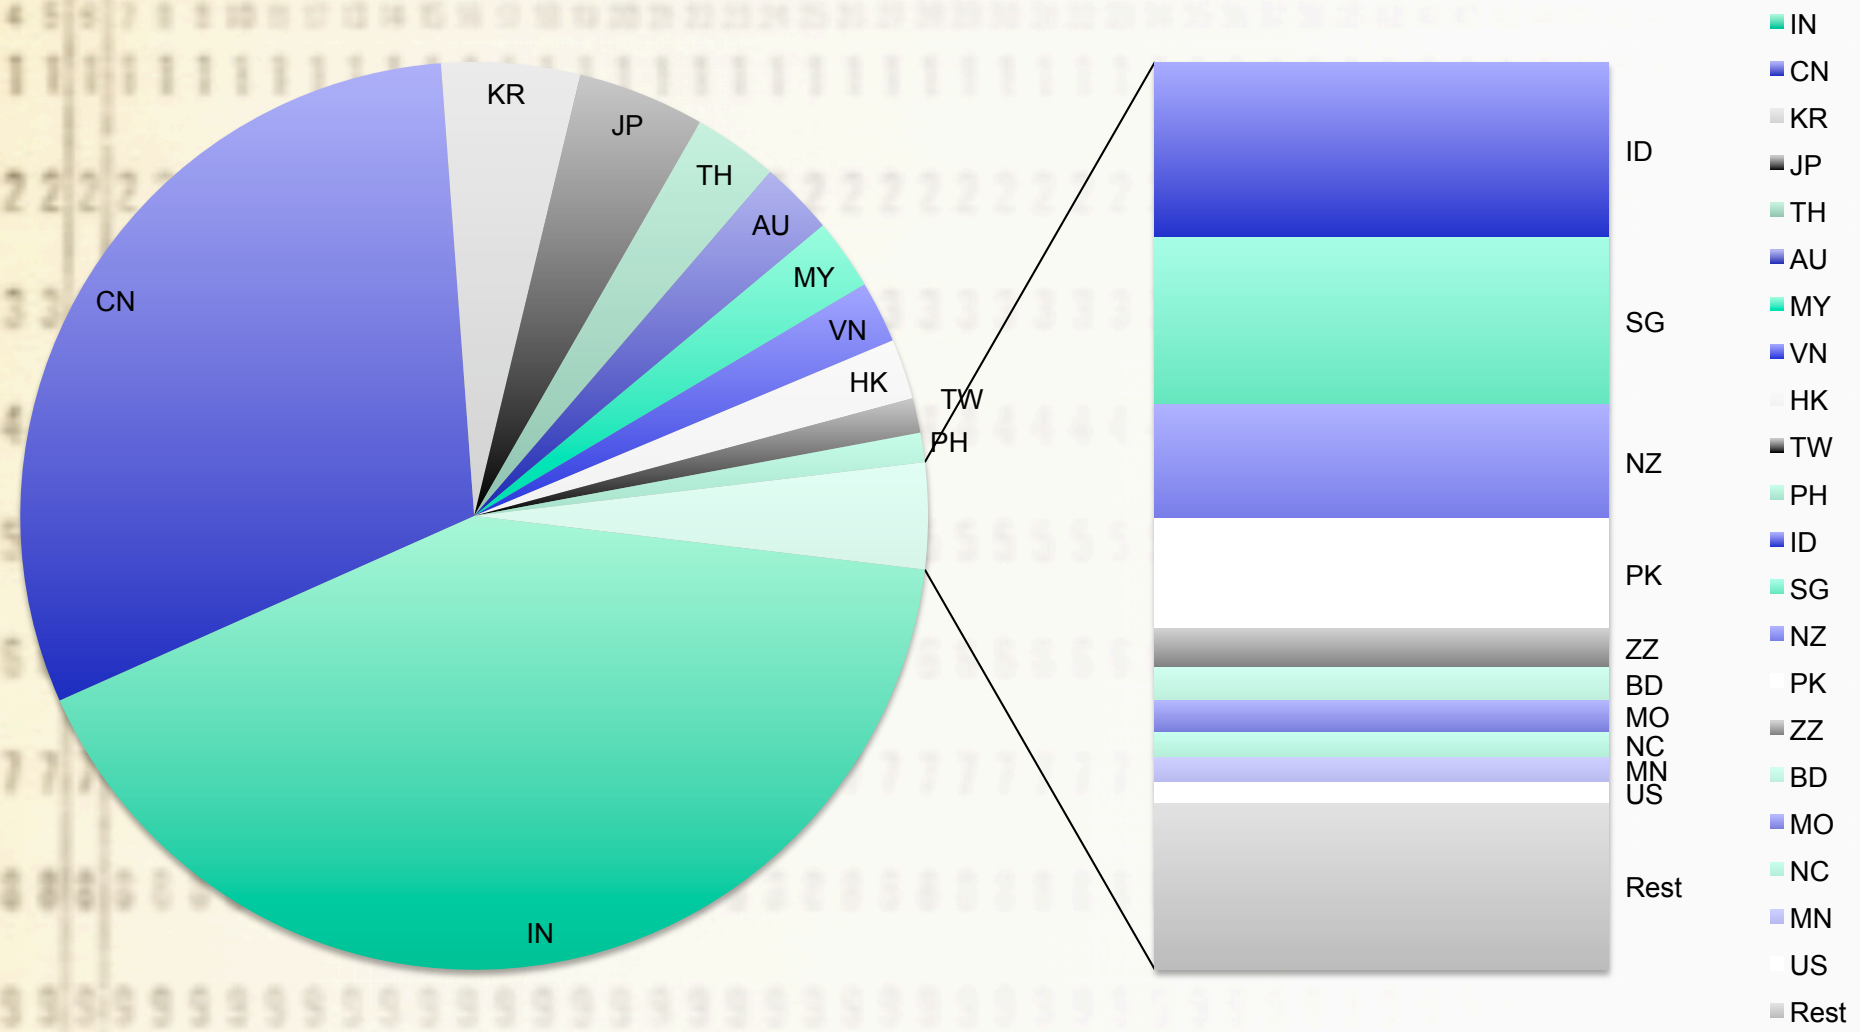

That’s a lot of IN/CN lookups

- Mailserver spam checks?
  - Is India a big spam source?
- For further study.

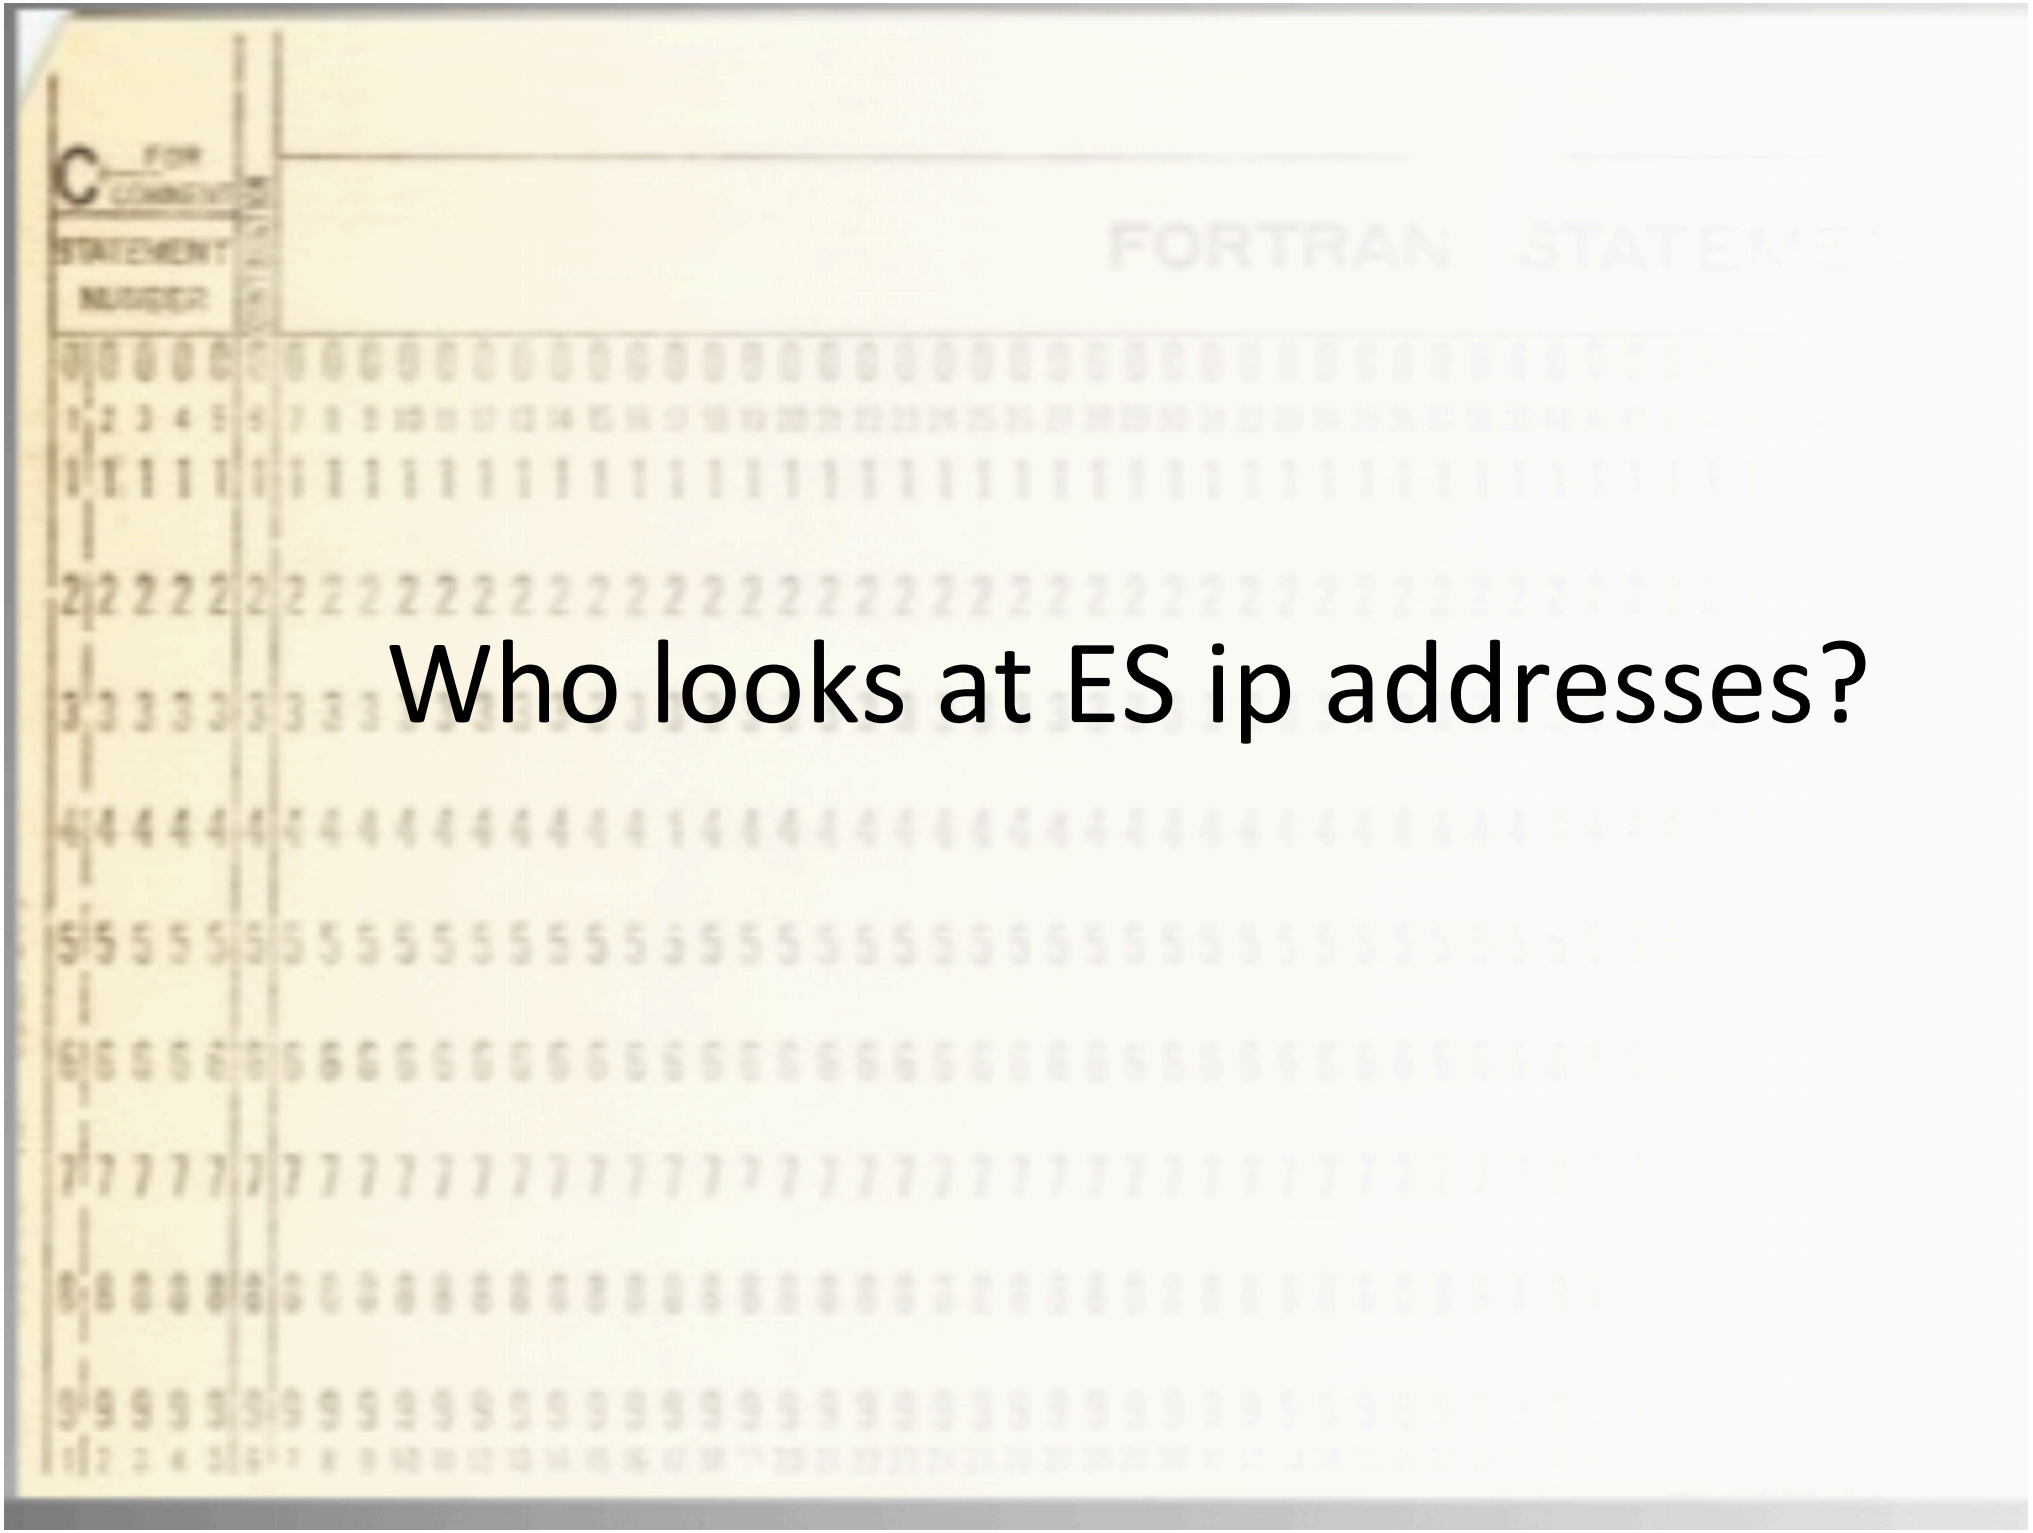
Who looks at ES ip addresses?

Who looks at ES?

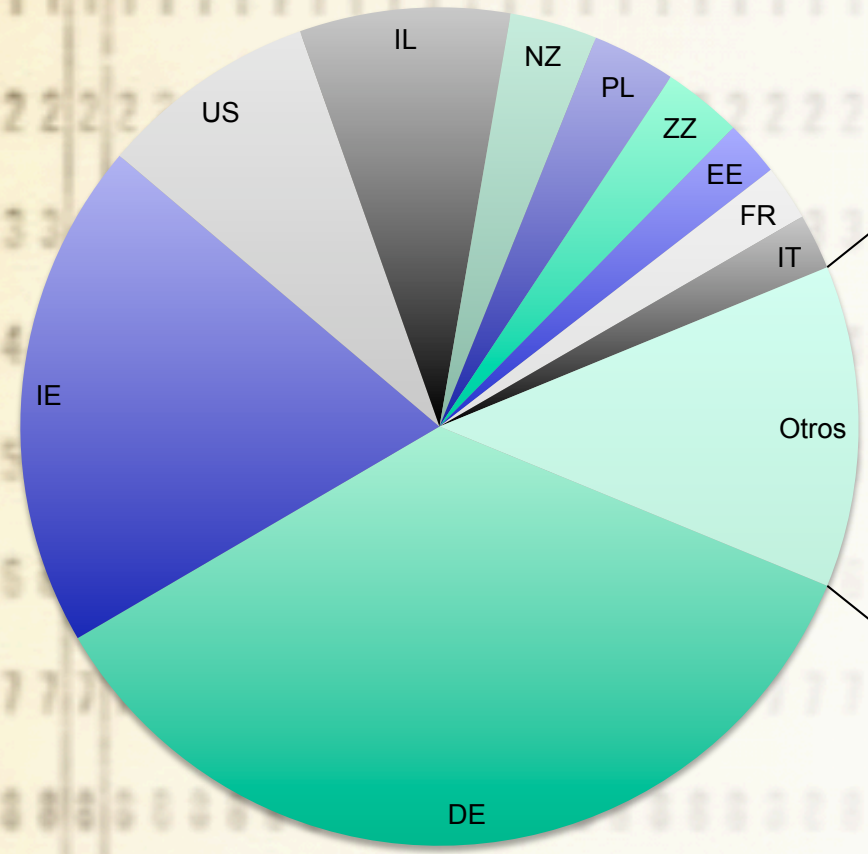

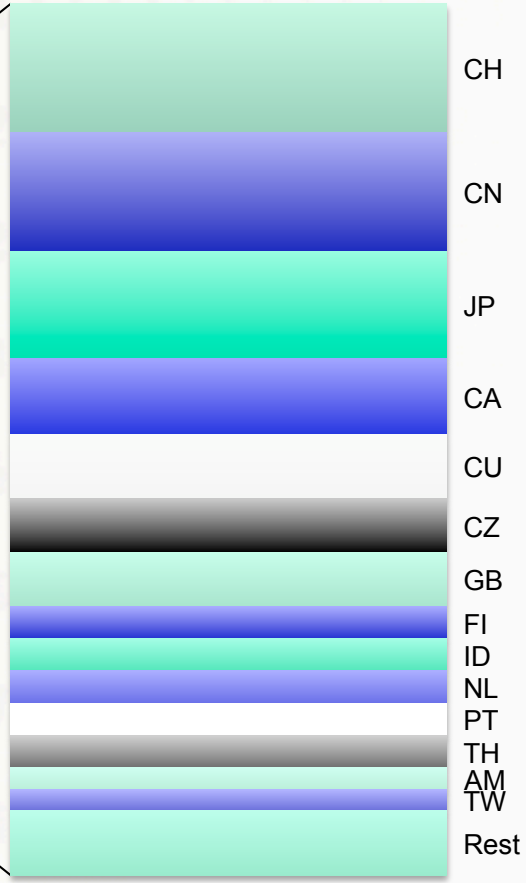

DE
IE
US
IL
NZ
PL
ZZ
EE
FR
IT
CH
CN
JP
CA
CU
CZ
GB
FI
ID
NL
PT

Its not the Asia-Pacific!

- Even noting the RTT, Many EU located economies use A-P located DNS servers to resolve PTR queries.
- Interesting to speculate if the lookup ratios reflect traffic, other measures of inter-economy dataflow
- For further study

Lessons learned 2008-2009

- 2008: 1hour captures
  - Huge risks if capture failed
  - Harder to upload to OARC (serialized)
  - 2009: 10 minute captures, parallel upload
- 2008: ran capture hosts on localtime
  - …but NTP was broken (2+hr offset) ☹
  - 2009: ran capture hosts on UTC, NTP checked!
- 2008: full capture, query + response
  - 2009: unable to capture responses on sec3
    - Too much data. Need to rethink what the value is in reply

Conclusions

- Infrastructure DNS is very odd.
  - More volatility in the query IP address than expected
  - Use of Teredo, other tunnels increasing
  - Use of IPv6 increasing
  - Some indications day-on-day comparison 2008/9 that V4 is not increasing significantly
  - Per economy, results can be confusing

PARIS 6e
Macintosh SE/30