

Measuring TCP Connection Establishment Times of Dual-Stacked Web Services [1]

Vaibhav Bajpai

NMRG Workshop, Zürich

Contributors

Vaibhav Bajpai, Jürgen Schönwälder

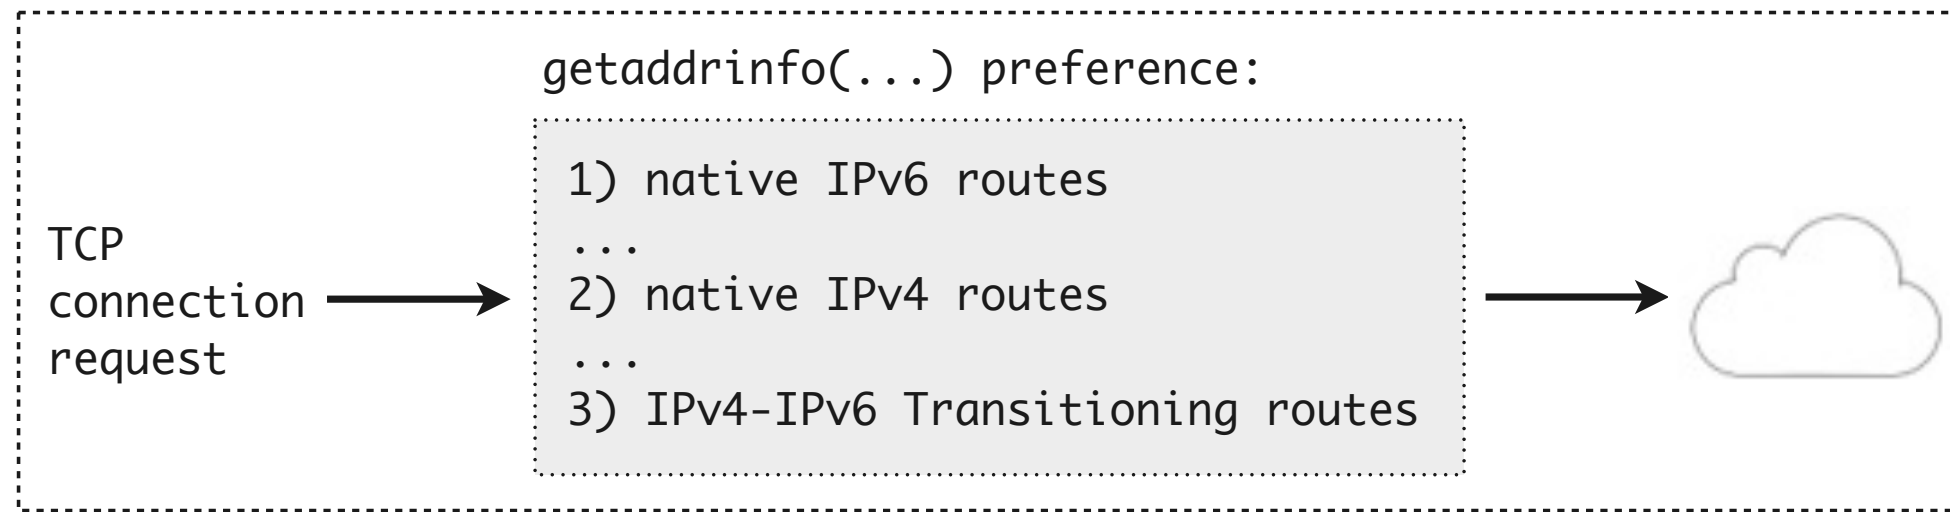
*Computer Networks and Distributed Systems (CNDS)
Jacobs University Bremen, Germany*

October 2013

Supported by:
Leone Project: <http://leone-project.eu>

Motivation

getaddrinfo(...) Behavior



- Returns a list of endpoints in an order that prioritizes an IPv6-upgrade path.
- The order is dictated by [\[RFC 6724\]](#) and `/etc/gai.conf`

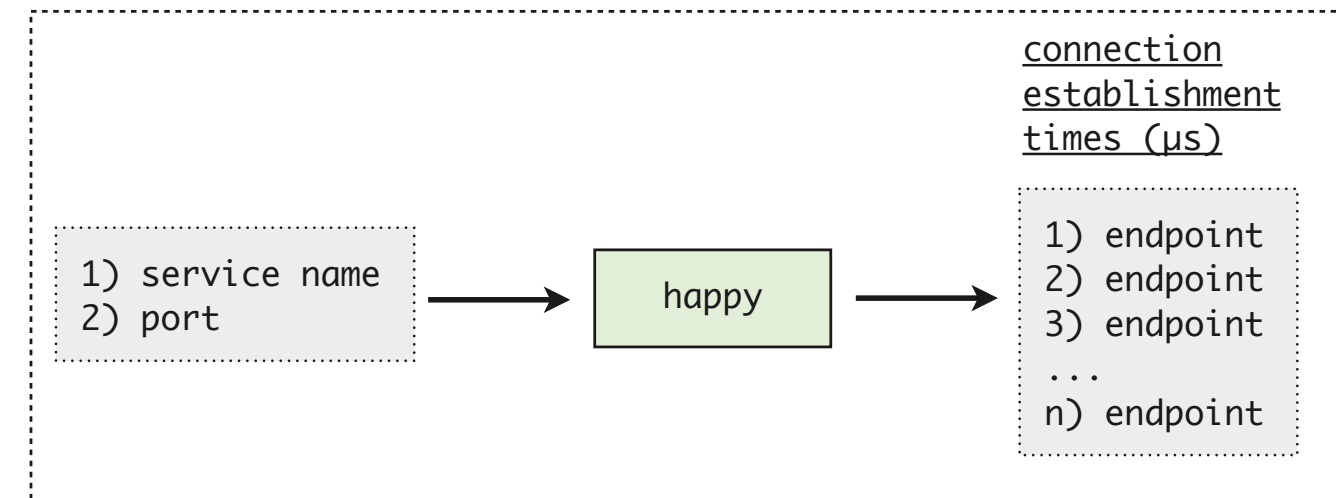
Research Question

- How does the preference given to IPv6 *impacts* the experience of dual-stacked users?

Methodology

Metric and Implementation

- Uses `getaddrinfo(...)` to resolve service names.
- Uses non-blocking TCP `connect(...)` calls.
- DNS resolution time is not accounted.
- Capability to read multiple service names as arguments.
- Capability to read service names list from a file.
- File locking capability.
- Applies a delay between `connect(...)` to avoid SYN floods.
- Capability to produce both human-readable and CSV output.
- Cross-compiled for OpenWrt platform. Currently running from SamKnows probes.



<http://happy.vaibhavbajpai.com>

```
$ ./happy -q 1 -m www.google.com www.facebook.com
HAPPY.0;1360681039;OK;www.google.com;80;173.194.69.105;8626
HAPPY.0;1360681039;OK;www.google.com;80;2a00:1450:4008:c01::69;8884
```

Selection of Web Services

How to compile a dual-stacked service names list?

- Hurricane Electric (HE) maintains a top 100 dual-stacked service names list.
<http://bgp.he.net/ipv6-progress-report.cgi>
 - HE uses top 1M service names list from Alexa Top Sites (ATS).
 - HE does *not* follow CNAMEs.
- Amazon has made the ATS top 1M service names list public.
<http://s3.amazonaws.com/alexastatic/top-1m.csv.zip>
 - Prepared a custom top 100 dual-stacked service names list.
 - Explicitly follow CNAMEs.
 - Prepend a `www` to each service name and cross-check any AAAA response.

Measurement Export

How to export results from SamKnows probes?

- REST-based export over HTTP [\[draft-bagnulo-lmap-http\]](#)
 - SamKnows probe already has curl and some sort of simplistic scripting possibility.

What should be the parameters of a measurement cycle?

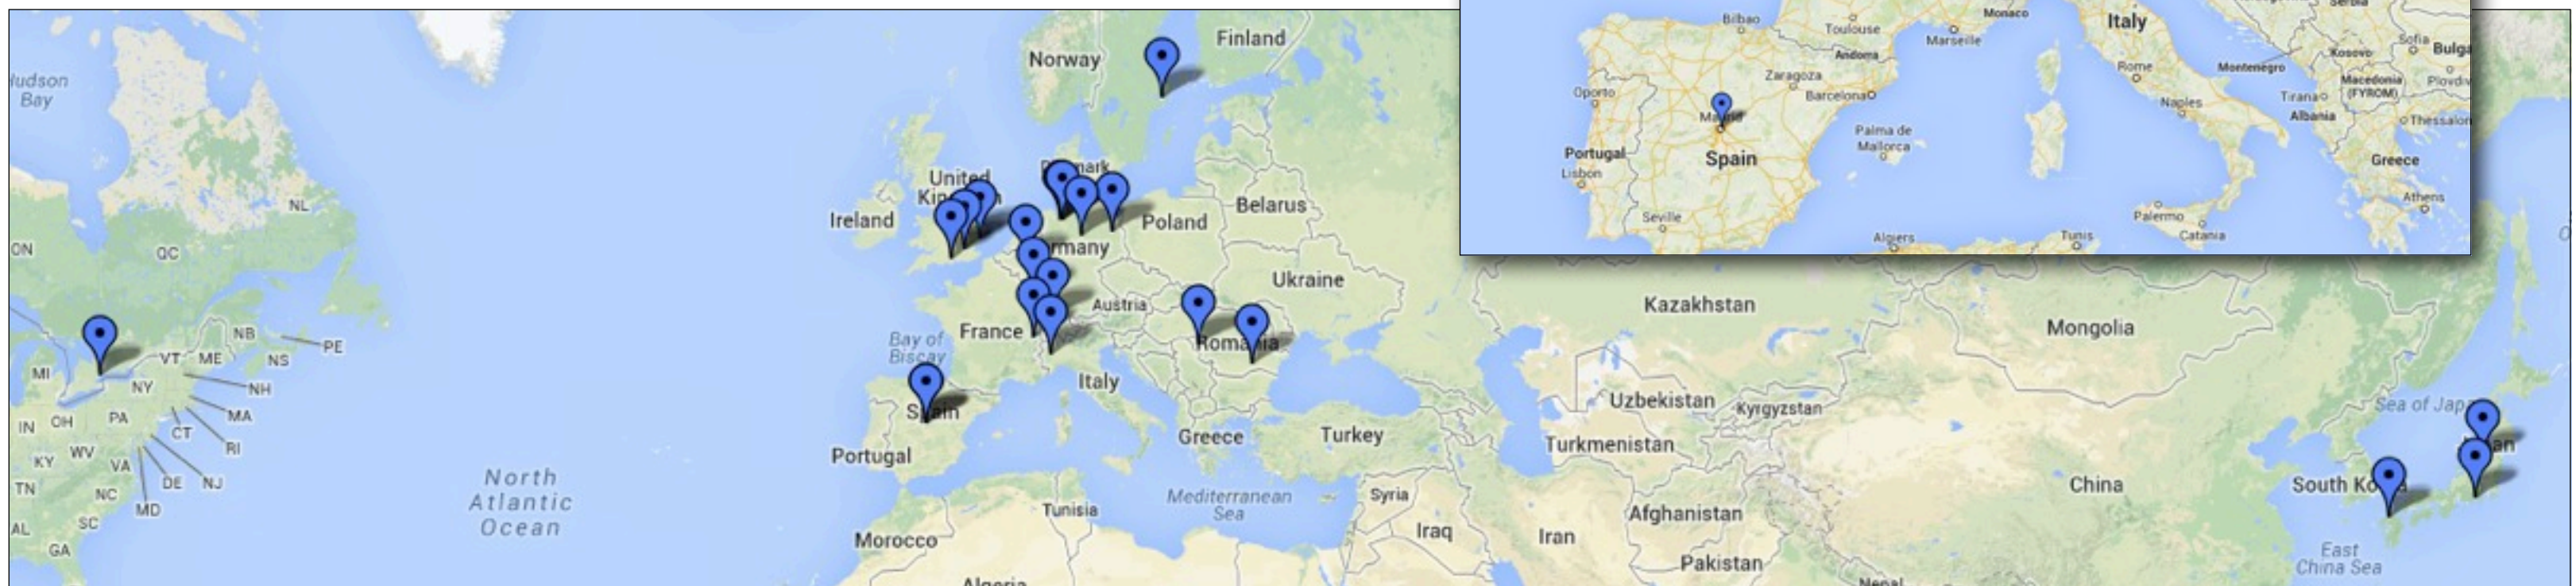
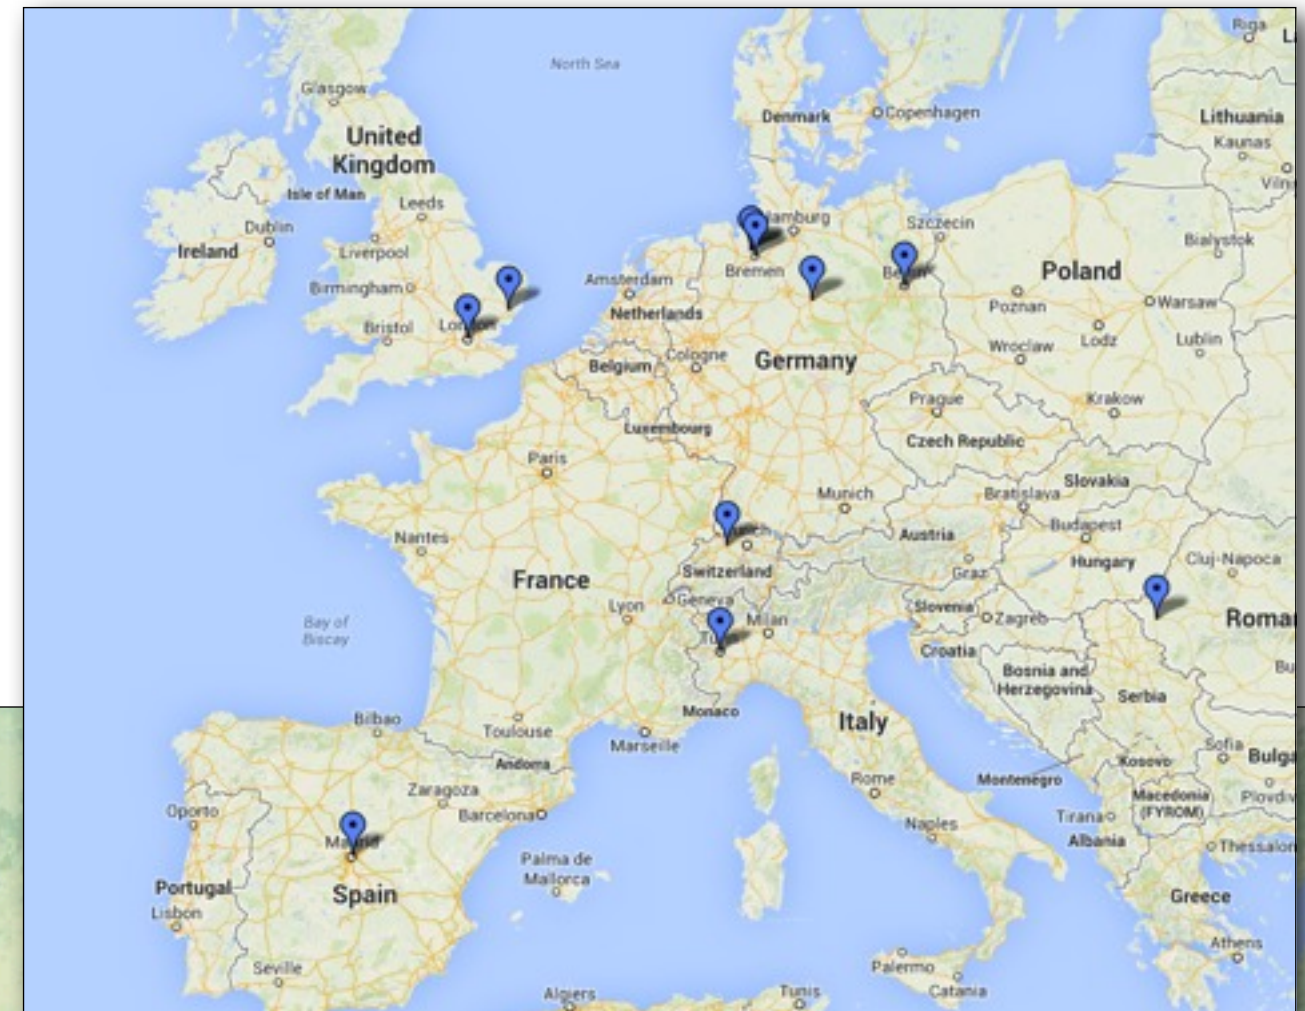
- Execute the test every 10 minutes.
- Report the results once in every 1 hour.

Measurement Trials

From where to run the measurement test?

Provider (IPv4, IPv6)	Location
(Deutsche Telekom AG, AS3320), (-)	Bremen
(Kabel Deutschland, AS31334), (-)	Bremen
(BT Spain, AS8903), (-)	Madrid
...	...

(-) means the IPv6 provider and AS are same as that for IPv4.




Whitelisting and Blacklisting

Are access to certain web services blocked administratively?

- AAAA prefix whitelisting [RFC 6589]
 - Google IPv6 used to be an opt-in service.
- Google IPv6 blacklist [googleipv6.vaibhavbajpai.com]
 - The policy has changed since the World IPv6 Launch Day.



Geolocation of announced blacklisted prefixes by Google over IPv6



Country-based distribution of blacklisted prefix for Google IPv6 services.

Related Work

How is our measurement different from [RFC 6556]?

- We do *not* account DNS in connection establishment time.
 - avoid input parameters that may *bias* the measurement (slow resolvers)
- Our testbed configuration is *active* rather than passive.
 - measurement test *actively* measures time taken to establish the TCP connection.
- Our testbed setup is designed for a *uncontrolled* environment.
 - does *not* require network path configuration changes.

Related Work


How is our measurement different from [\[RFC 6948\]](#)?

- Longer and *newer* measurement cycles.
 - [\[RFC 6948\]](#): May 25, 2011 – July 11, 2011
 - We are running the measurement since Mar 10, 2013 – Present.
- Measurement from a wider deployed vantage point
 - 3 MAs deployed somewhere in Finland, Sweden and Canada in [\[RFC 6948\]](#).
 - 14 MAs deployed across EU, more upcoming ...
- We do *not* measure the amount of AAAA entries within 1M ATS.
 - [\[RFC 6948\]](#) noticed around 300 (within top 10K ATS) services were dual stacked.
 - [\[RFC 6948\]](#) noticed around 30 (within top 100 ATS) services were dual stacked.
 - We take top 1M ATS and filter the top 100 dual-stacked services.

Preliminary Results

Measuring Raw Performance

How does the performance (mean) of IPv6 compare to that of IPv4?



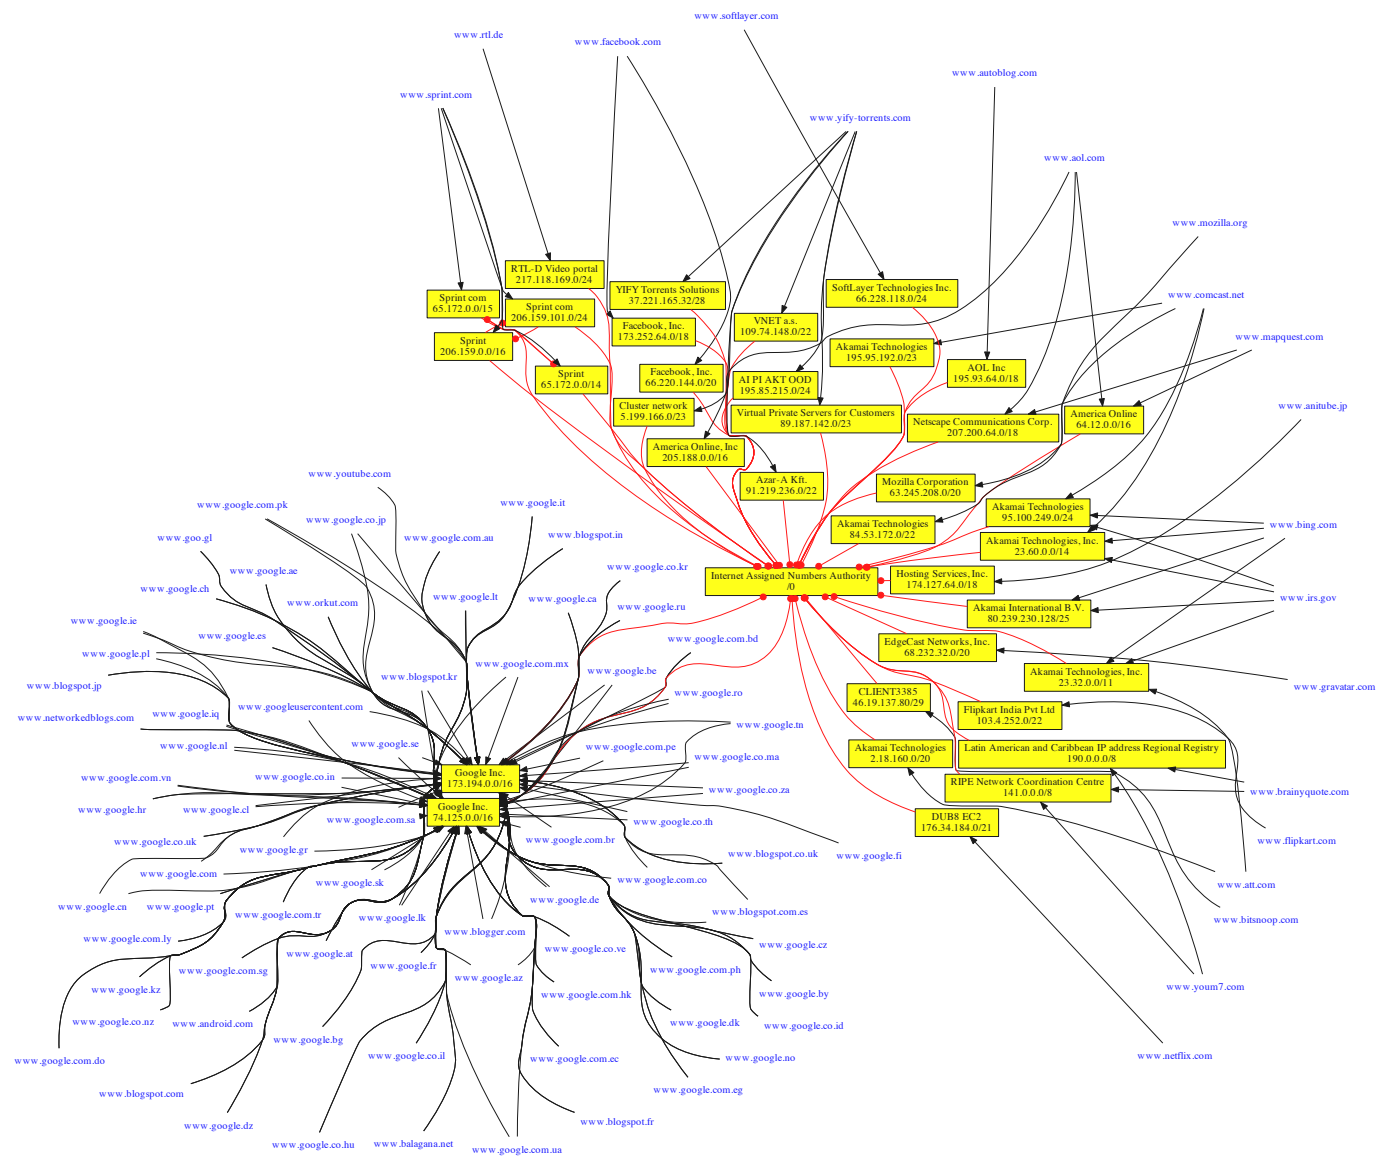
Service Clusters

To what extent do web services centralize on CDNs?

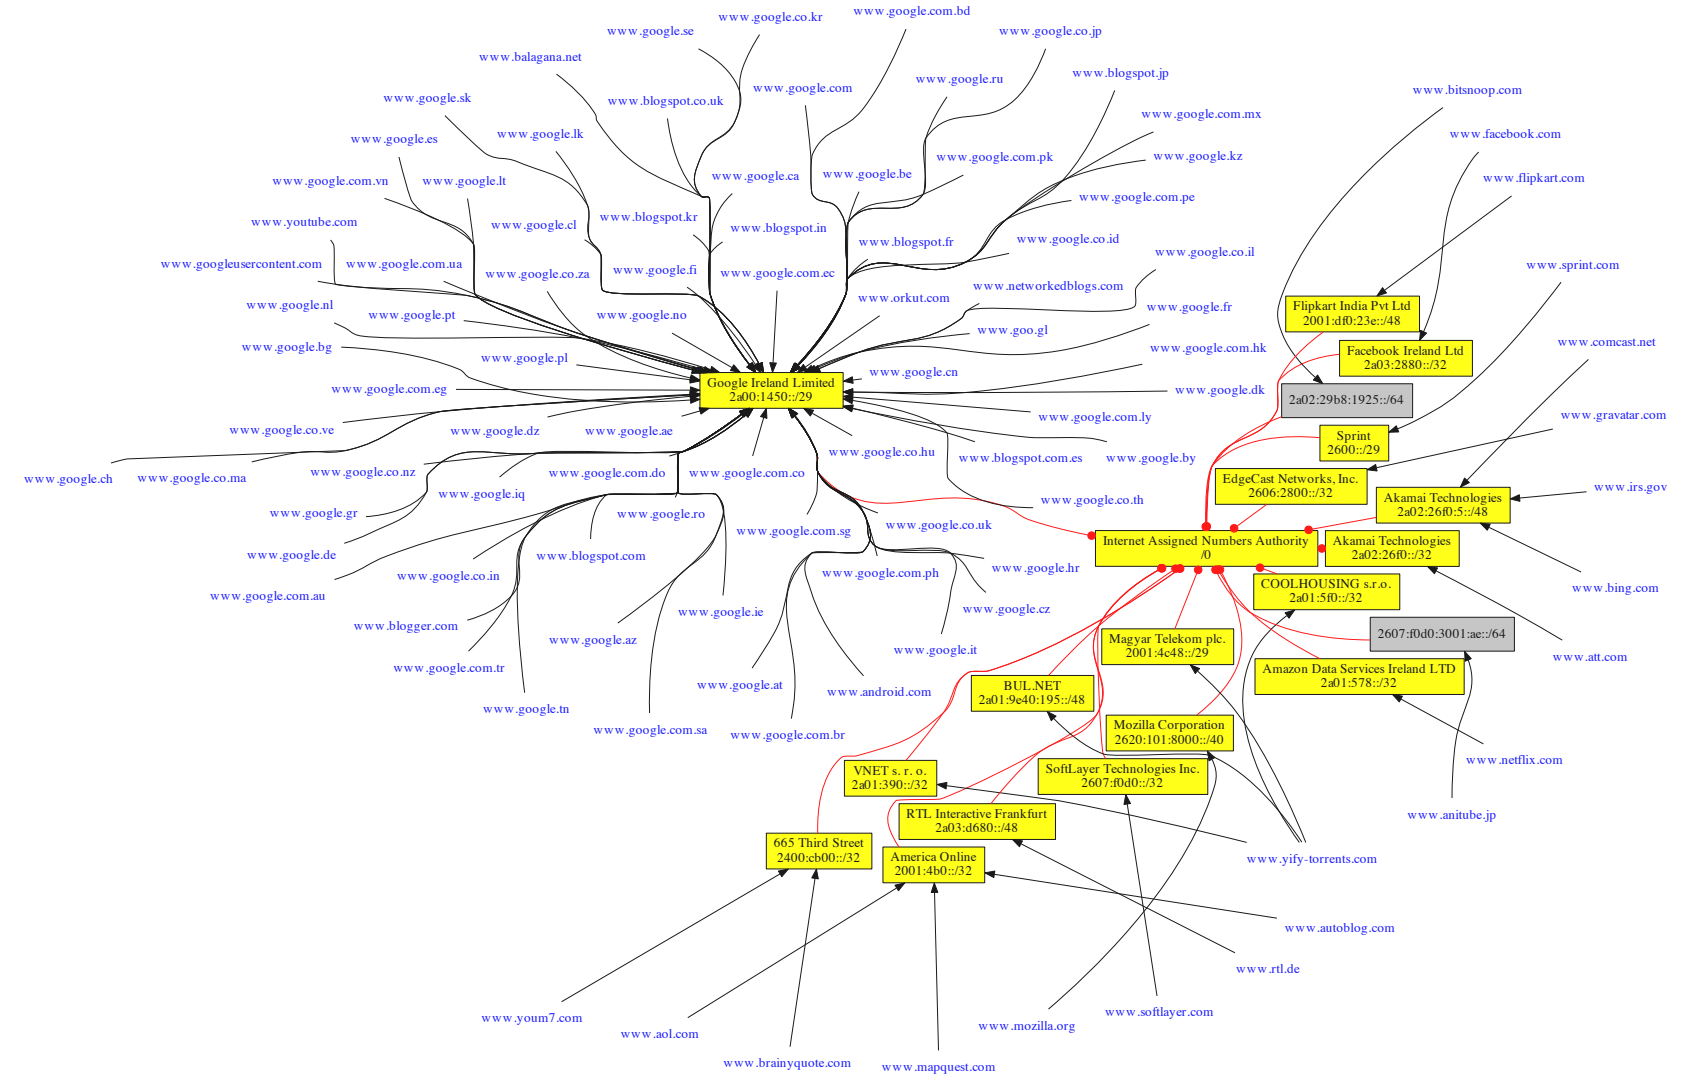
- WHOIS-based clusters
 - For each endpoint, send a REST call to whois.arin.net
 - Forward the REST call to whois.ripe.net if endpoint assigned by a different registry.
 - whois.ripe.net can also delegate the requests to APNIC and AFRINIC.
 - grab the (allocated prefix, holder organization, regional registry).

Service Clusters

To what extent do web services centralize on CDNs?



IPv4 Aggregation Cloud



IPv6 Aggregation Cloud

Service Clusters

How to better aggregate service name clusters?

- whois-based clusters are coarse-grained.
The owner can slice the allocated blocks and announce them from different ASes.
- BGP-based clusters may not be accurate.
The location of the route collector is different from that of the probe's location.
- RSD-based clusters require known AS topology graph [3]
RSD metric measures AS-path similarity to a destination prefix from all source ASes.
- Statistical clustering techniques
k-means clustering on observed TCP connection establishment means and their variation.
- Forward-path based clusters.
Use traceroute as a metric to capture the forward-path directly from the MA.

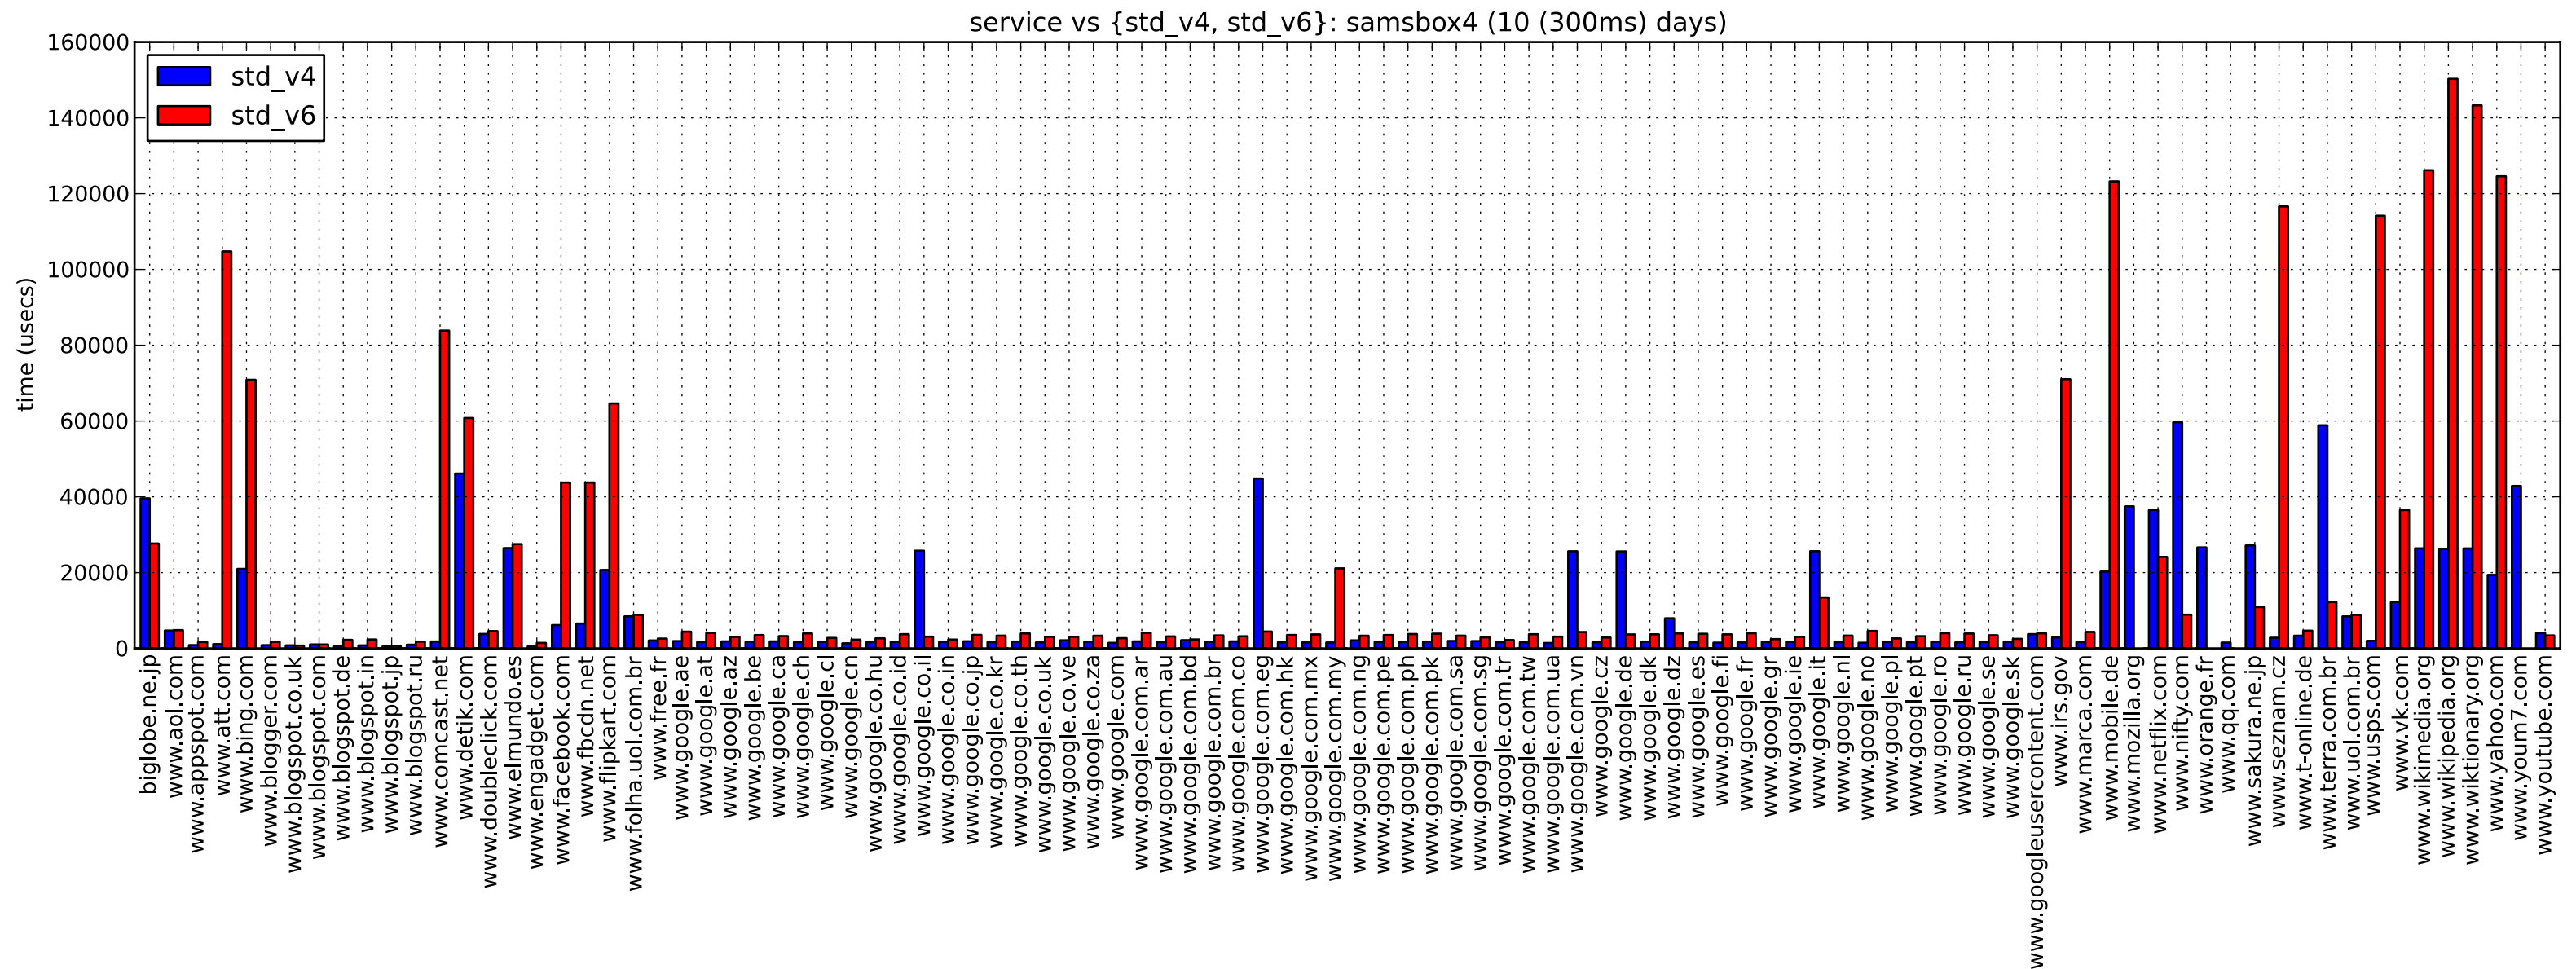
traceroute (mtr) on SamKnows probes.

- IPv4 and IPv6 support with -4 and -6 flags.
- Measures latency to each hop.
- Reverse DNS lookup to each hop. It can be disabled via --no-dns flag.
- ANS lookup to each hop endpoint via --aslookup flag
- Result generated in SamKnows compatible CSV format with --csv flag.
- Skip the service name on DNS resolution error, log to stderr and move on in the list.
- Capability to read multiple service names list as arguments.
- Capability to read service names list from a file with a --filename flag.
- Capability to lock file writes.
- Cross-compiled for OpenWrt platform. Currently running from SamKnows probes.

```
$ sudo ./mtr -c 1 --no-dns -6 --csv --aslookup www.facebook.com
MTR.0.84+git:ce36701d;1367775366;OK;www.facebook.com;1;2001:638:709:3000::1;AS680;4049
...
MTR.0.84+git:ce36701d;1367775366;OK;www.facebook.com;12;2620:0:1cff:dead:beef::97;AS32934;105817
MTR.0.84+git:ce36701d;1367775366;OK;www.facebook.com;13;2620:0:1cff:dead:beef::1329;AS32934;106714
MTR.0.84+git:ce36701d;1367775366;OK;www.facebook.com;14;2a03:2880:2110:cf01:face:b00c:0:9;AS32934;105930
```


Measuring Raw Performance

How does the performance (variation) of IPv6 compare to that of IPv4?



Preliminary Insights

- Higher connection times and variations over IPv6.
- A number of disparate services (bing, comcast, irs) show similar performances.
- whois data reveals they resolve to same RIR allocated blocks owned by a CDN.
- whois aggregation clouds reveal many services centralize at Google and Akamai CDNs.

How are our measurement results different from [\[RFC 6948\]](#)?

- We noticed significantly *higher* TCP connection setup delay differences.
 - Generally slower over IPv6.
 - Multiple services were twice as slow over IPv6 when compared to IPv4.
- We noticed significantly *lower* TCP connection setup failure rates.
 - We witnessed 1% of service failure rates, as opposed to 20% witnessed in [\[RFC 6948\]](#).

Further Reading

- [1] V. Bajpai, *et al.*, Measuring TCP Connection Establishment Times of Dual-Stacked Web Services Conference on Network and Service Management (CNSM), 2013.
- [2] ———, Understanding the Impact of Network Infrastructure Changes using Large-Scale Measurement Platforms, Conference on Autonomous Infrastructure, Management and Security (AIMS), 2013: http://dx.doi.org/10.1007/978-3-642-30633-4_19

- Measuring the Effects of Happy Eyeballs:
 - IETF 87 Talk, July 2013: <http://www.ietf.org/proceedings/87/slides/slides-87-v6ops-8.pdf>
 - Internet Draft: July 2013: <http://tools.ietf.org/html/draft-bajpai-happy-01>
 - RIPE Labs Technical Article, June 2013: https://labs.ripe.net/Members/vaibhav_bajpai/evaluating-the-effectiveness-of-happy-eyeballs
 - RIPE66 Talk, May 2013: <https://ripe66.ripe.net/archives/video/1208>

References

- [3] G. Gürsun, et al., Routing State Distance: A Path-Based Metric for Network Analysis, ACM Conference on Internet Measurement Conference (IMC), 2012.