

Understanding the Impact of Network Infrastructure Changes using Large-Scale Measurement Platforms

Vaibhav Bajpai and Jürgen Schönwälder

{v.bajpai, j.schoenwaelder}@jacobs-university.de

AIMS 2013, Barcelona

Computer Networks and Distributed Systems
Jacobs University Bremen
Bremen, Germany

June 2013

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

Introduction

- Large-Scale Broadband Measurement Use Case [[draft-linsner-lmap-use-cases-02](#)].
 - Internet Service Provider (ISP)
 - Identify, isolate and fix problems in the access network.
 - Evaluate the Quality of Experience (QoE) of the user.
 - Benchmark and look into competitor insights.
 - Consumers
 - Does the ISP service adhere to the service level agreements (SLA)s?
 - Diagnose impaired components in the private network.
 - Regulators
 - Need datasets to compare multiple broadband providers.
 - Frame better policies to help regulate the broadband industry:
<http://www.fcc.gov/measuring-broadband-america>
<http://maps.ofcom.org.uk/broadband>

State of the Art

- **Early Studies**
 - Inject packet trains to infer broadband link characteristics [Dischinger-IMC-2007].
- **Software-based Solutions:**
 - Speedtest.net a flash-tool to measure broadband throughput: <http://www.speedtest.net>.
 - DIMES, a software agent that performs ping and traceroute measurements [Shavitt-CCR-2005].
 - Glasnost, a Java-based applet that detects ISP-enforced traffic shaping [Dischinger-NSDI-2010].
 - Netalyzr, a Java-based applet that performs DNS, NAT, HTTP, IPv6-based tests. [Kreibich-IMC-2010].
 - Fathom, a Firefox-extension to Netalyzr [Dhawan-IMC-2012].
- **Large-Scale Measurement Platforms:**
 - SamKnows and BISmark <http://www.samknows.com>
 - RIPE Atlas <http://atlas.ripe.net>
 - Google's Measurement Lab (M-Lab) <http://www.measurementlab.net>
 - CAIDA's Archipelago (Ark) <http://www.caida.org/projects/ark>

State of the Art

- LMAP and IPPM Standardization

- Large Scale Measurement of Access Network Performance (LMAP) Birds of a Feather (BOF) at IETF 86.

- Control and Report Protocol candidates

- [\[draft-schoenw-lmap-netconf-00\]](#)

- [\[draft-bagnulo-lmap-ipfix-01\]](#)

- [\[draft-seedorf-lmap-lmap-alto-00\]](#)

- Data Model candidates

- [\[draft-schoenw-lmap-yang-00\]](#)

- IP Performance Metrics (IPPM) charter revision.

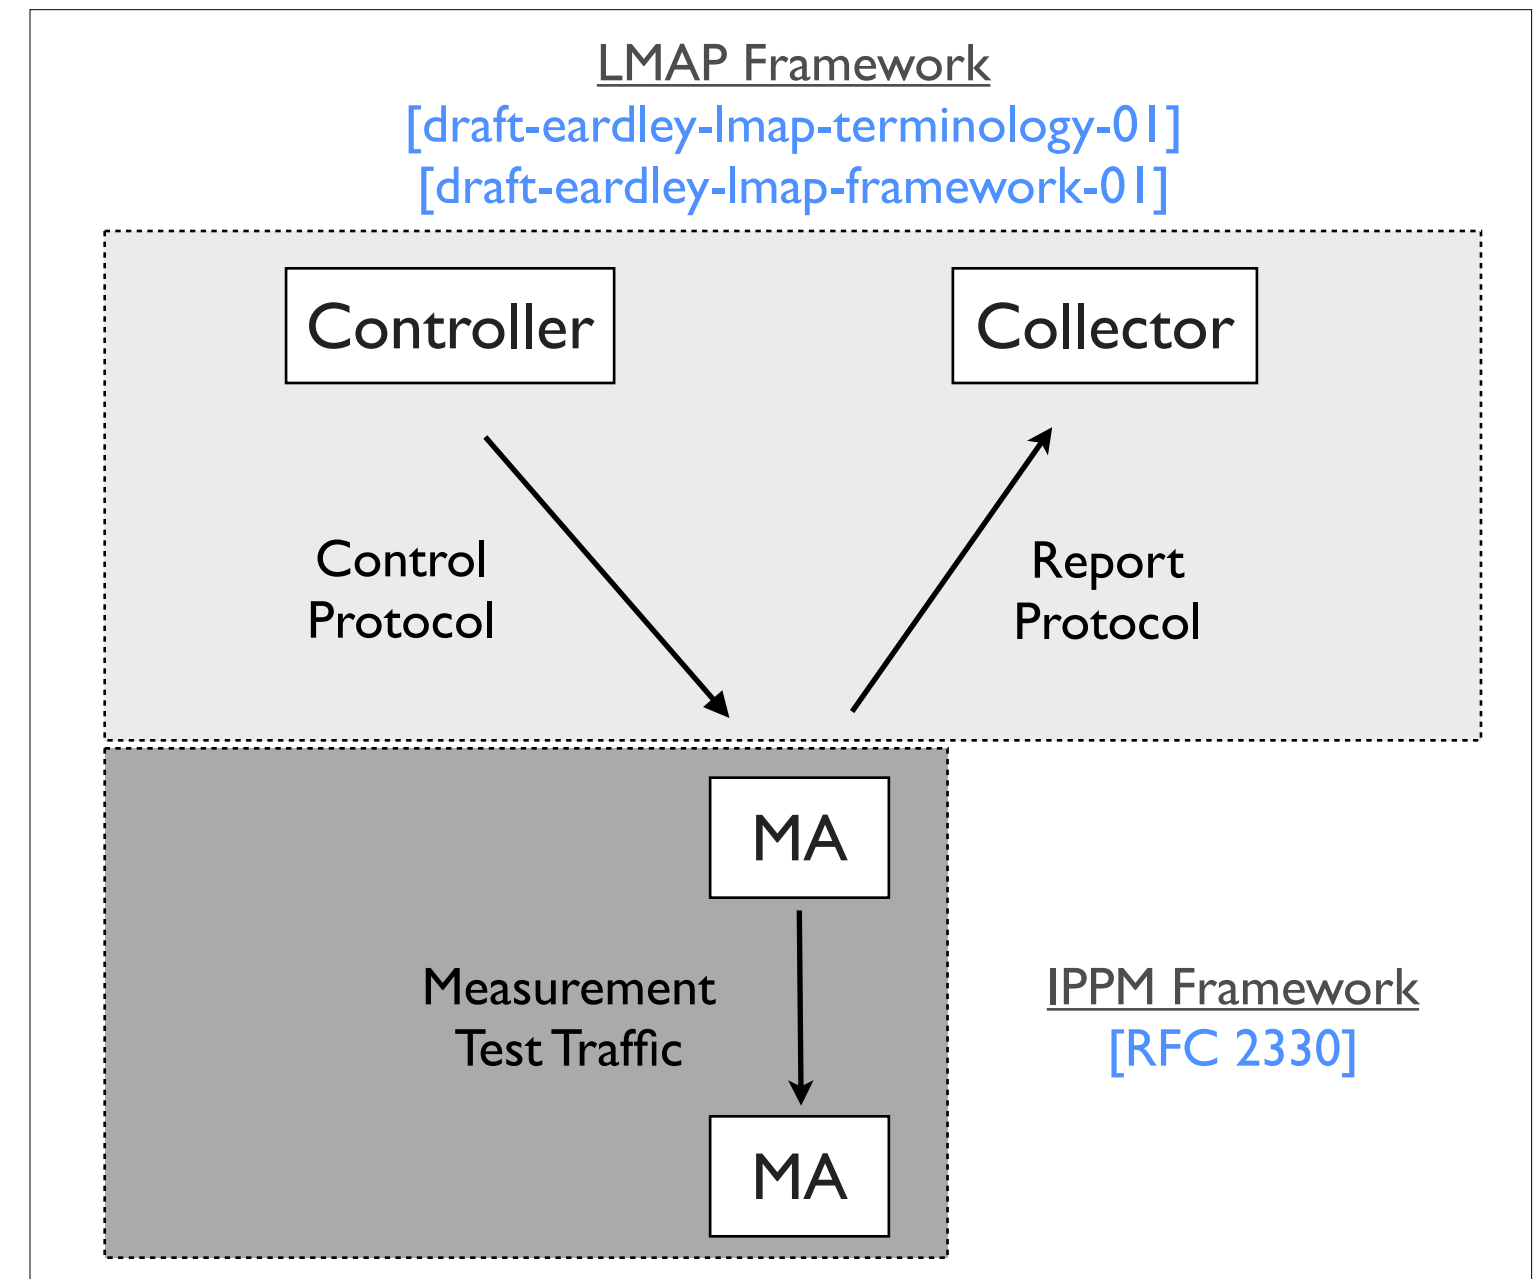
- Registry for commonly-used metrics

- [\[draft-bagnulo-ippm-new-registry-00\]](#)

- [\[draft-bagnulo-ippm-new-registry-independent-00\]](#)

- Regulatory Implications

- Standards body collaboration: IETF + BBF + IEEE



Research Statement

- Understanding the Impact of Network Infrastructure Changes using Large-Scale Measurement Platforms

- Measuring broadband performance from residential gateway.
- Helping regulators sketch better policy decisions.

Goals of Earlier Studies

- Understanding the Impact of Network Infrastructure Changes using Large-Scale Measurement Platforms

- Study IPv6 transition [Bajpai-AIMS-2012].
 - Can we identify a Carrier-Grade NAT (CGNAT) from a residential gateway?
 - Can we identify multiple layers of NATs from a residential gateway?
- Measure today's IPv6 network.
 - ~~Measure IPv6 adoption?~~ [Dhamdhere-IMC-2012] [Allman-SIGMETRICS-2013] [Colitti-PAM-2010]
<http://www.google.com/ipv6/statistics.html>
<http://bgp.he.net/ipv6-progress-report.cgi>
 - How does the performance of IPv6 compare to that of IPv4?
- Study the blend of network centralization and decentralization
 - To what extent do web services centralize on Content Delivery Networks (CDNs)?
 - To what extent does web experience depend on Regionalization?

Extending the Goal

Approach

- Requirements?

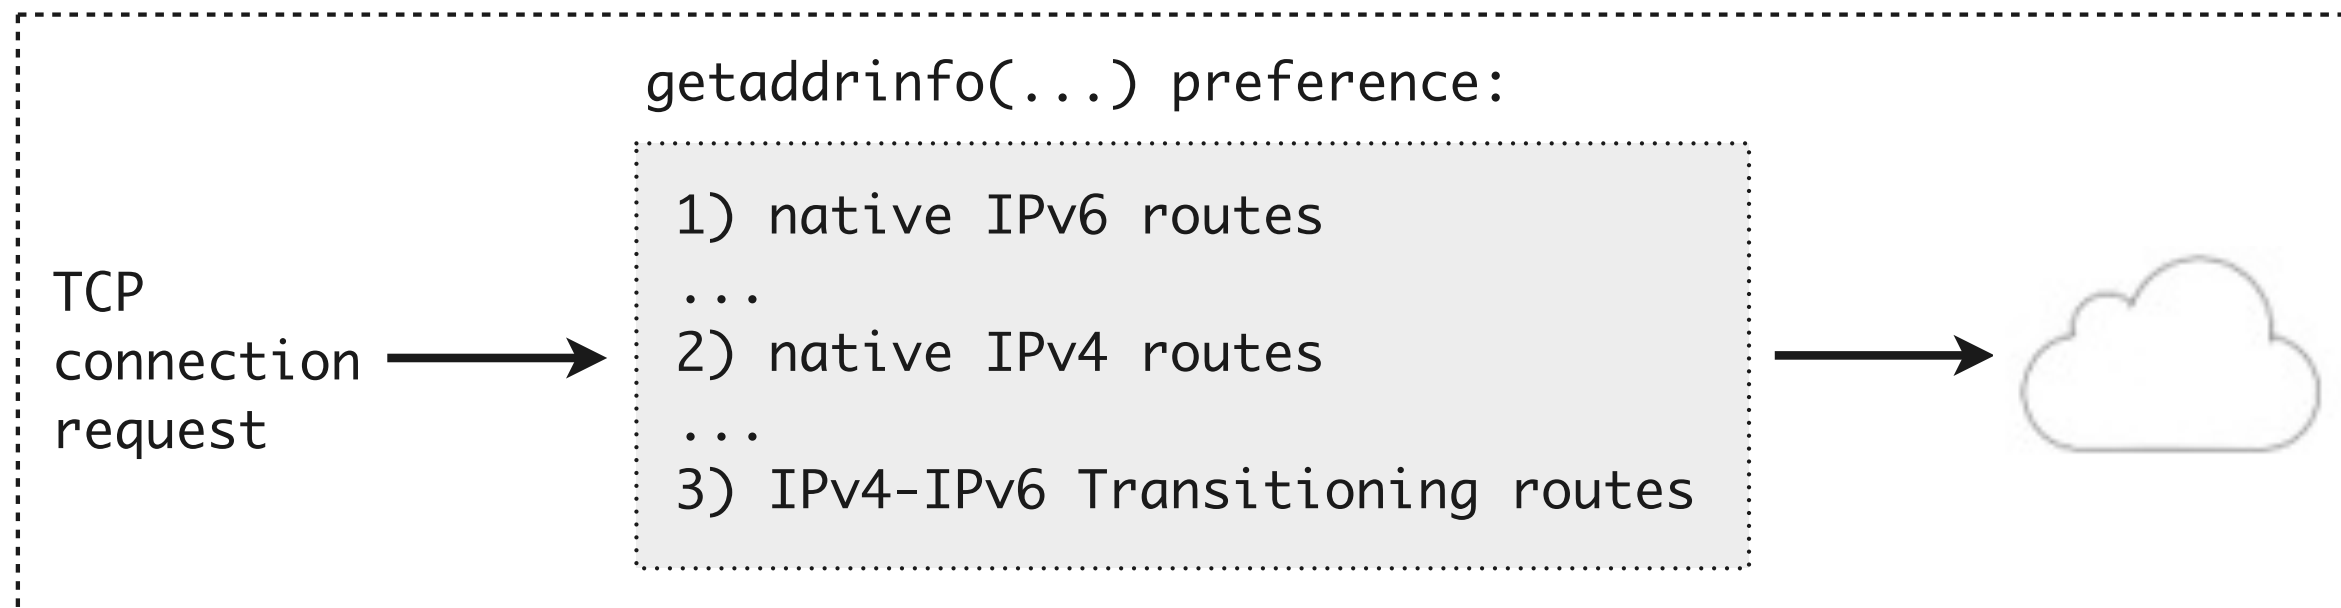
- Access to a large-scale measurement platform.
 - SamKnows and Jacobs University are partners of the Leone Consortium <http://www.leone-project.eu>.
- Address allocations from Regional Internet Registries (RIR).
- Publicly available BGP data from route collectors.

- Work Flow

- Define metrics targeted to our research questions.
- Implement measurement tests that adhere to the metric definition.
- Deploy measurement tests on a large-scale measurement platform.
- Conglomerate measurement results from multiple Measurement Agents (MA)s.
- Correlate measurement results with data from RIRs and route collectors.
- Prepare data analysis tools that can mine this multidimensional data.
- Uncover the insights to answer the research questions.

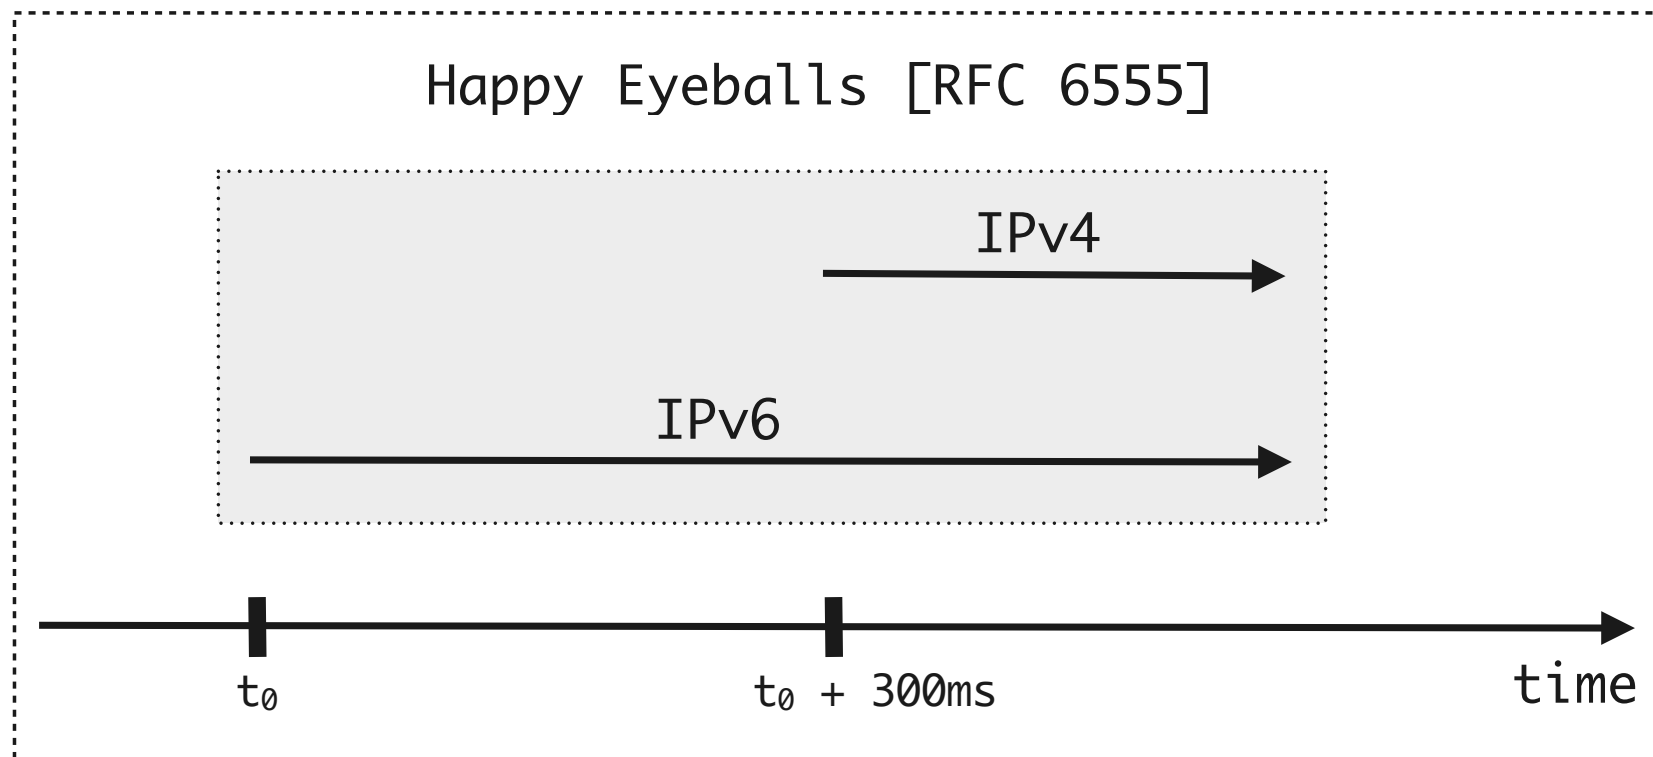
Preliminary Work | Measuring Happy Eyeballs

- `getaddrinfo(...)` behavior:
 - Returns a list of endpoints in an order that prioritizes IPv6-upgrade path.
 - The order is dictated by [\[RFC 6724\]](#) and `/etc/gai.conf`
 - If the IPv6 connectivity is broken, an application remains unresponsive in the order of seconds.



Preliminary Work | Measuring Happy Eyeballs

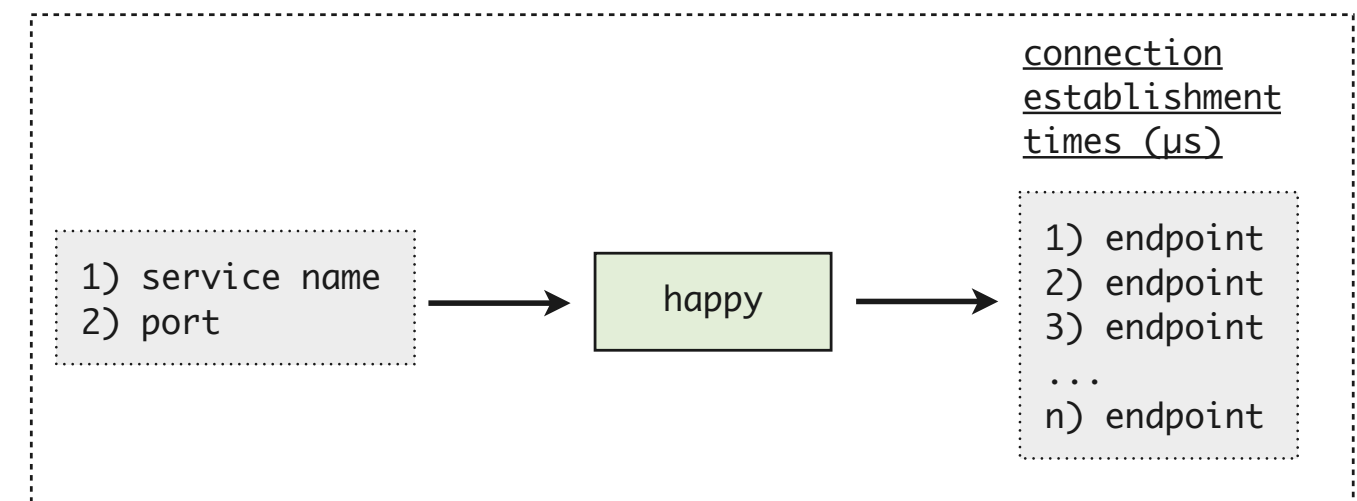
- Happy Eyeballs Algorithm [\[RFC 6555\]](#):
 - Initiate a TCP connect(...) with the first endpoint, give it 300ms.
 - Switch over with a TCP connect(...) to a different address family otherwise.
 - The competition runs fair after 300ms.



Preliminary Work | Measuring Happy Eyeballs

- Metrics and Implementation

- Uses `getaddrinfo(...)` to resolve service names.
- Uses non-blocking TCP `connect(...)` calls.
- Applies a delay between `connect(...)` to avoid SYN floods.
- Service name resolution time is not accounted.
- Capability to produce both human-readable and CSV output.
- Capability to read multiple service names as arguments.
- Capability to read service names list from a file.
- File locking capability.
- 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
HAPPY.0;1360681039;OK;www.facebook.com;80;2a03:2880:10:6f01:face:b00c::8;170855
HAPPY.0;1360681039;OK;www.facebook.com;80;31.13.72.39;26665
```

Preliminary Work | Measuring Happy Eyeballs

- 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/alexa-static/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.

Preliminary Work | Measuring Happy Eyeballs

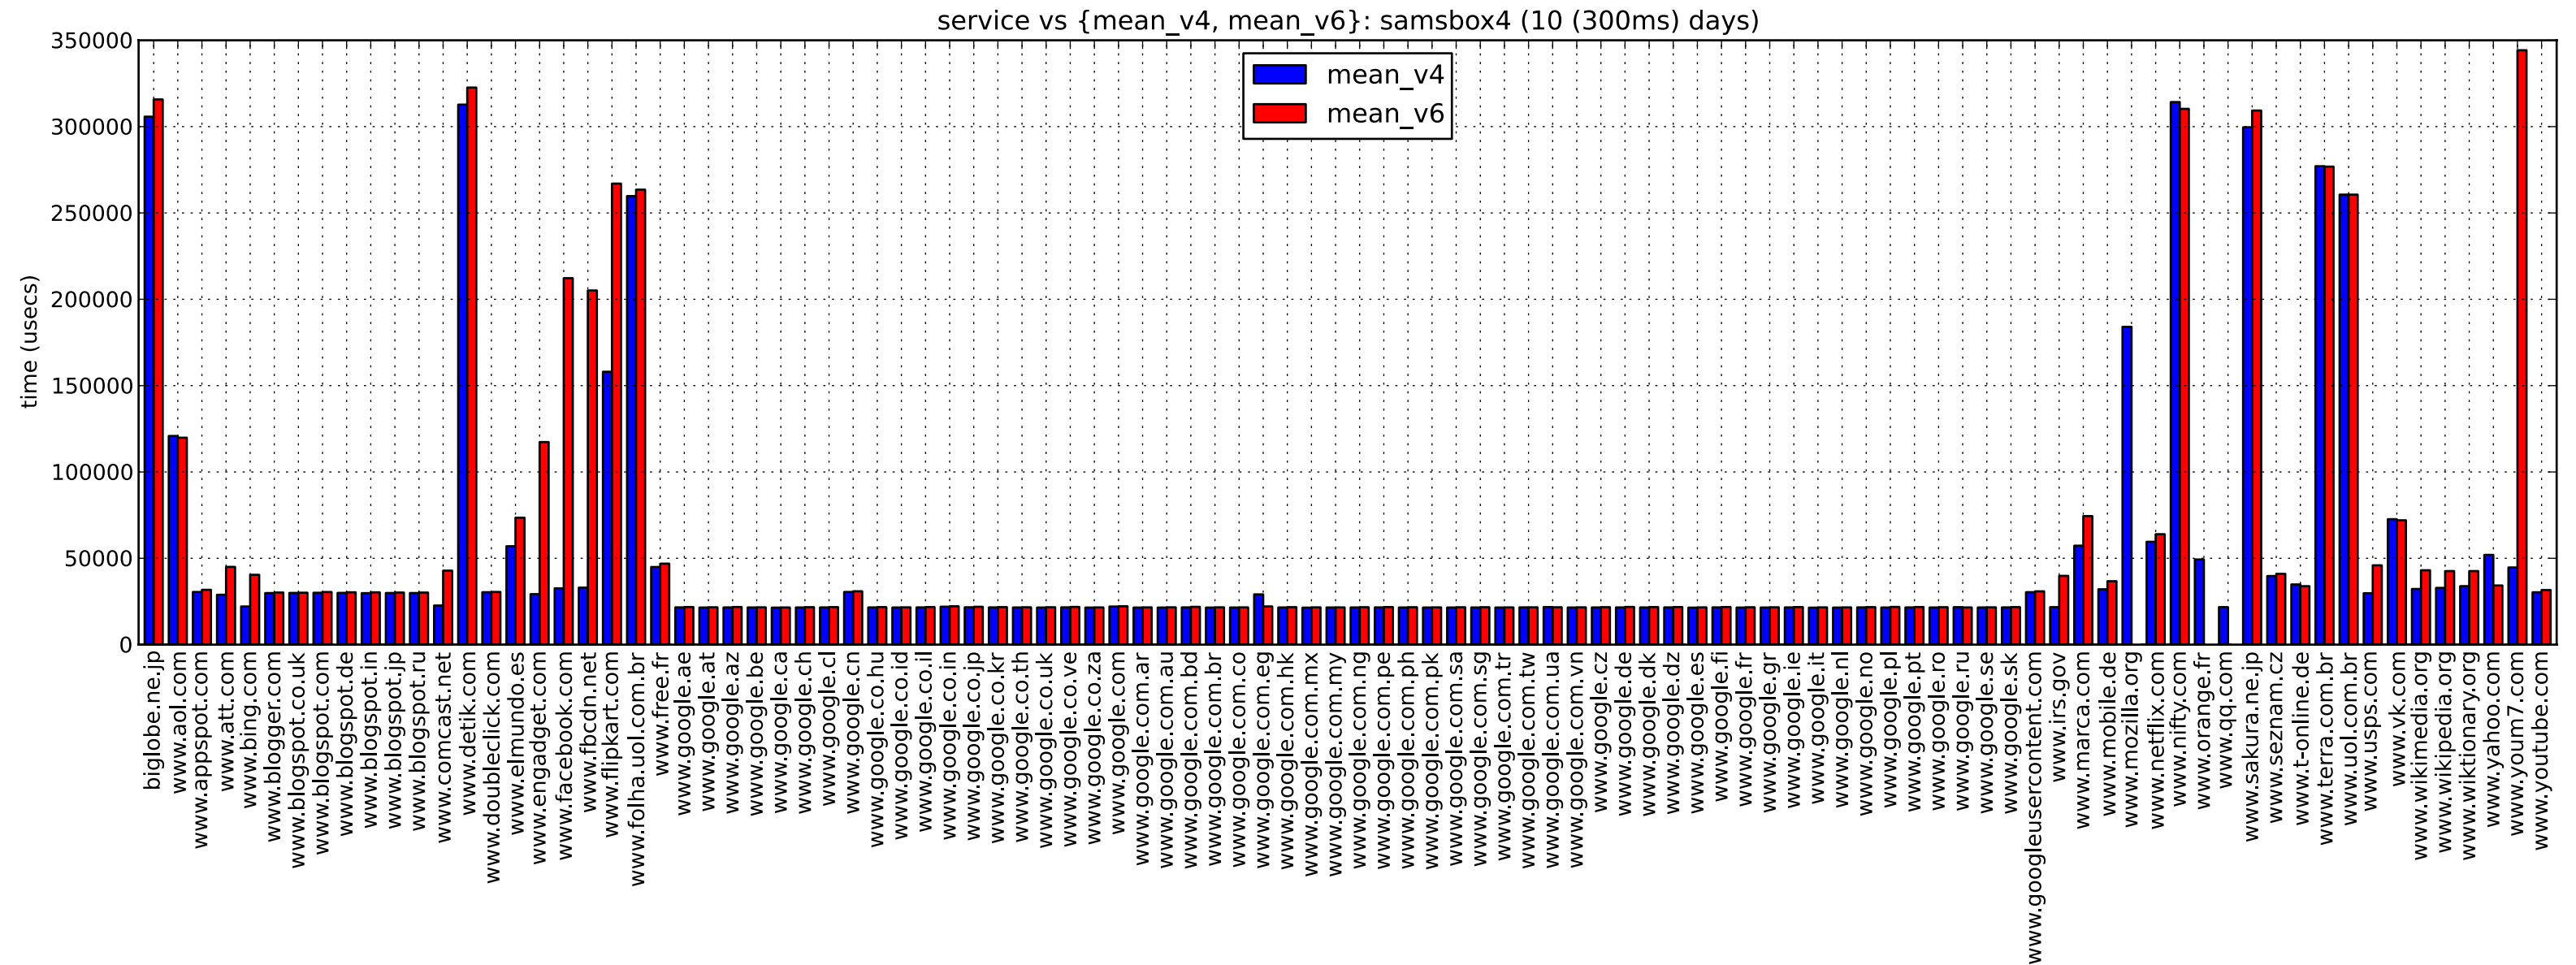
- From *where* to run the measurement test?

Provider (IPv4, IPv6)	Location	Platform
(dfn, AS680), (-)	Jacobs University Bremen	SamKnows
(Kabel Deutschland, AS31334), (HE, AS6939)	Bremen	SamKnows
(Gaertner Datensystems GmbH, AS24956), (-)	Braunschweig	SamKnows
(Deutsche Telekom AG, AS3320), (-)	Bremen	SamKnows
(British Sky Broadcasting Limited, AS5607), (-)	London	SamKnows
(Telekom Italia, AS3269), (-)	Torino	SamKnows
(BT Spain, AS8903), (-)	Madrid	SamKnows
(ROEDUNET, AS2614), (-)	Timisoara	SamKnows
(LambdaNet Communications, AS13237), (Teredo)	Berlin	GNU/Linux
(dfn, AS680), (-)	Jacobs University Bremen	Mac OS X

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

Preliminary Work | Measuring Happy Eyeballs

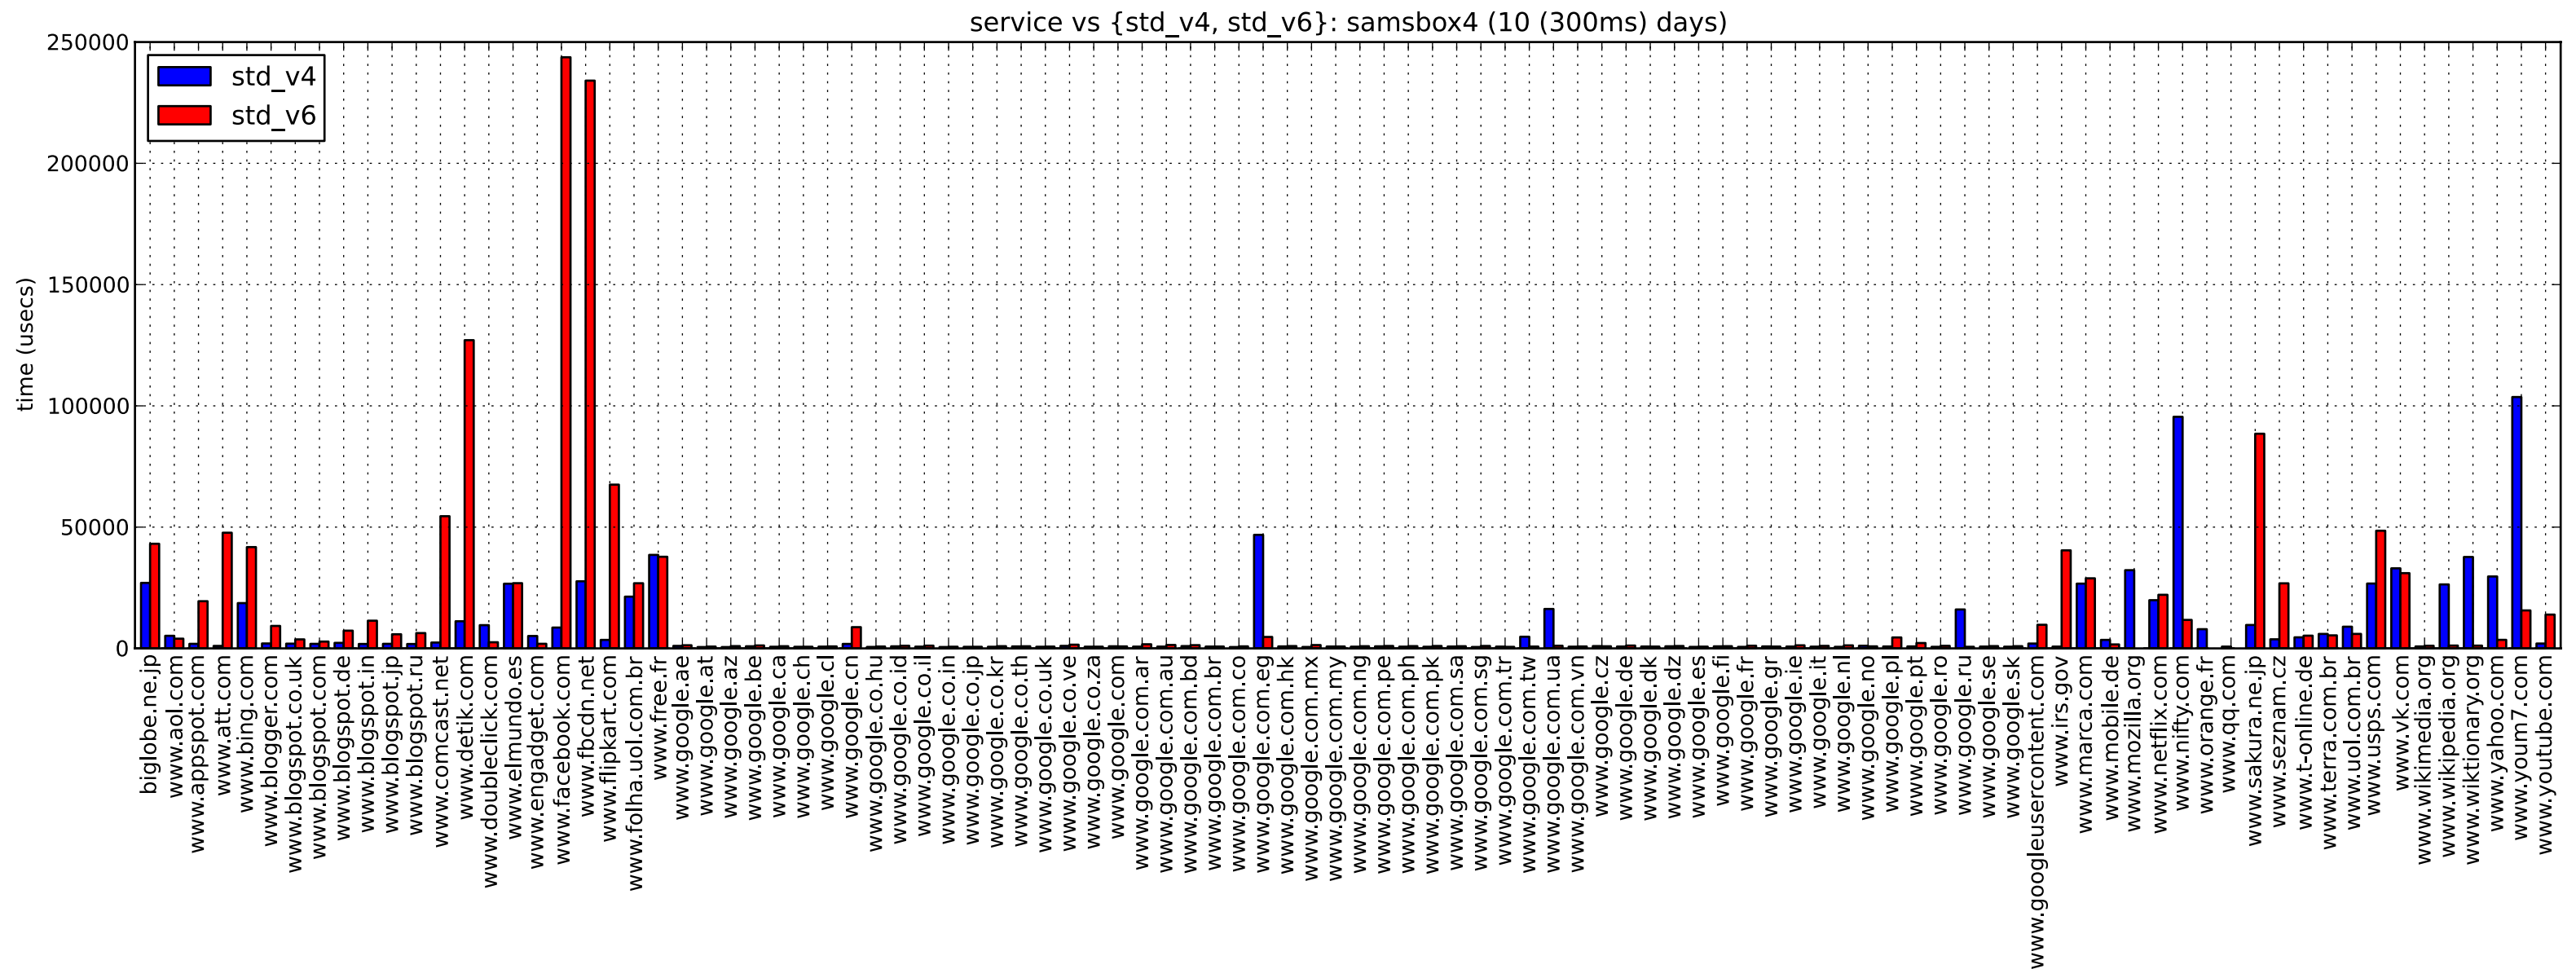
- How does the *performance (mean)* of IPv6 compare to that of IPv4?



- Native IPv4 and IPv6 connectivity via DTAG - Deutsche Telekom AG [AS 3320]

Preliminary Work | Measuring Happy Eyeballs

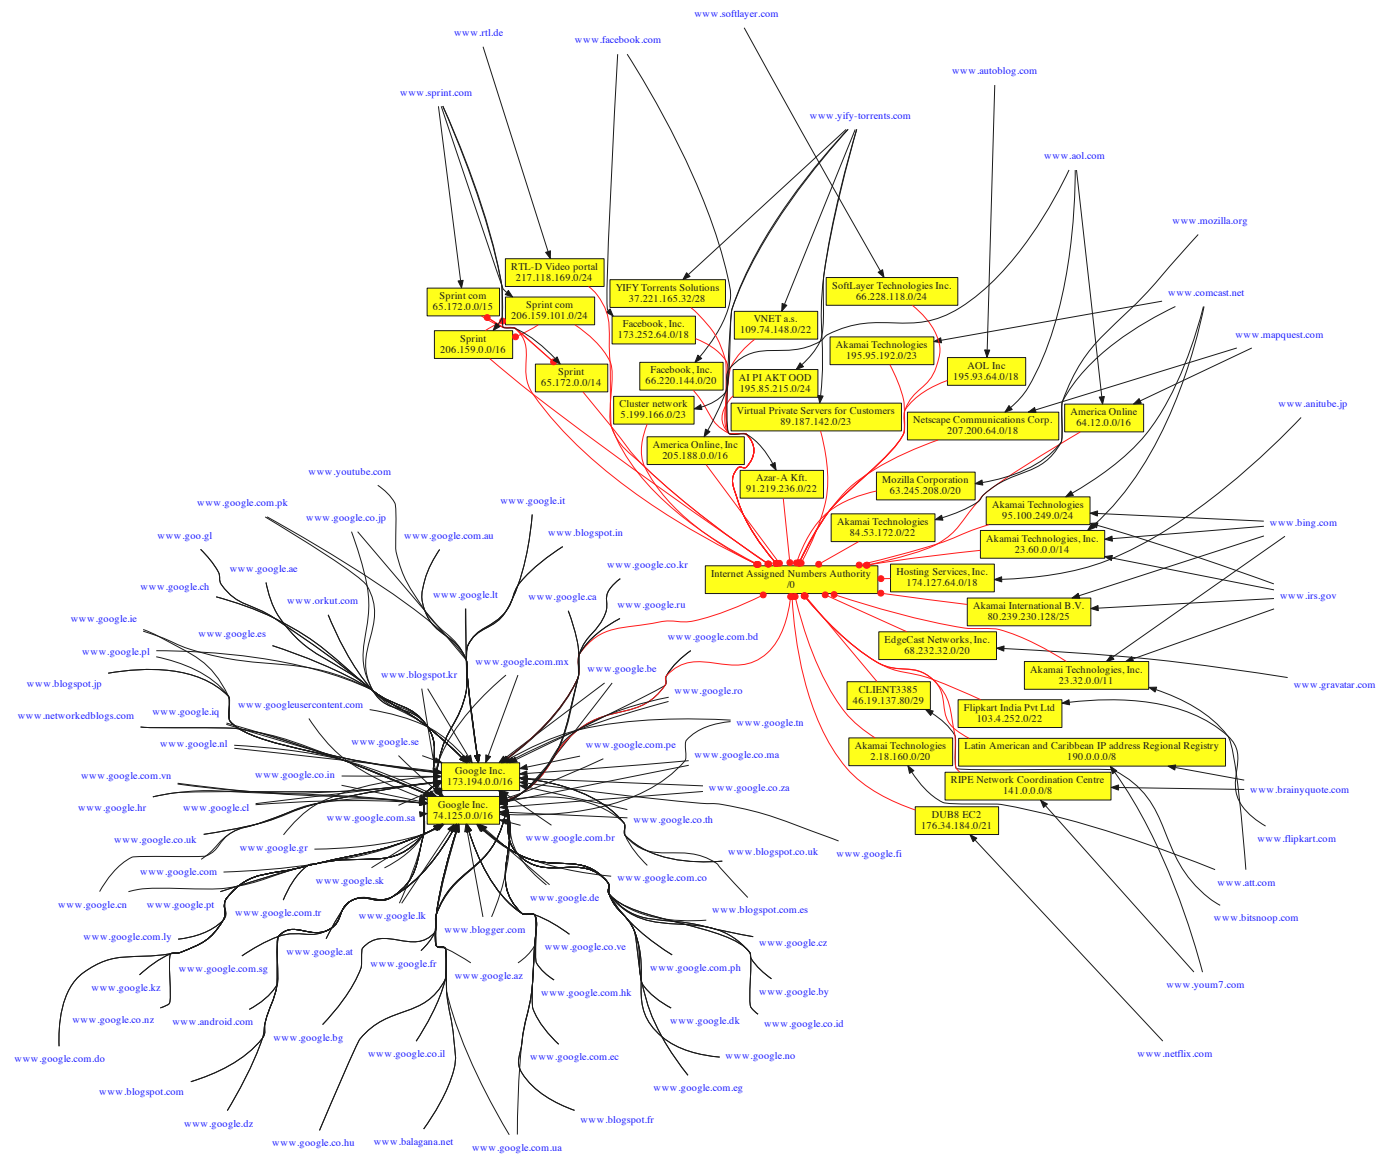
- How does the *performance (variation)* of IPv6 compare to that of IPv4?



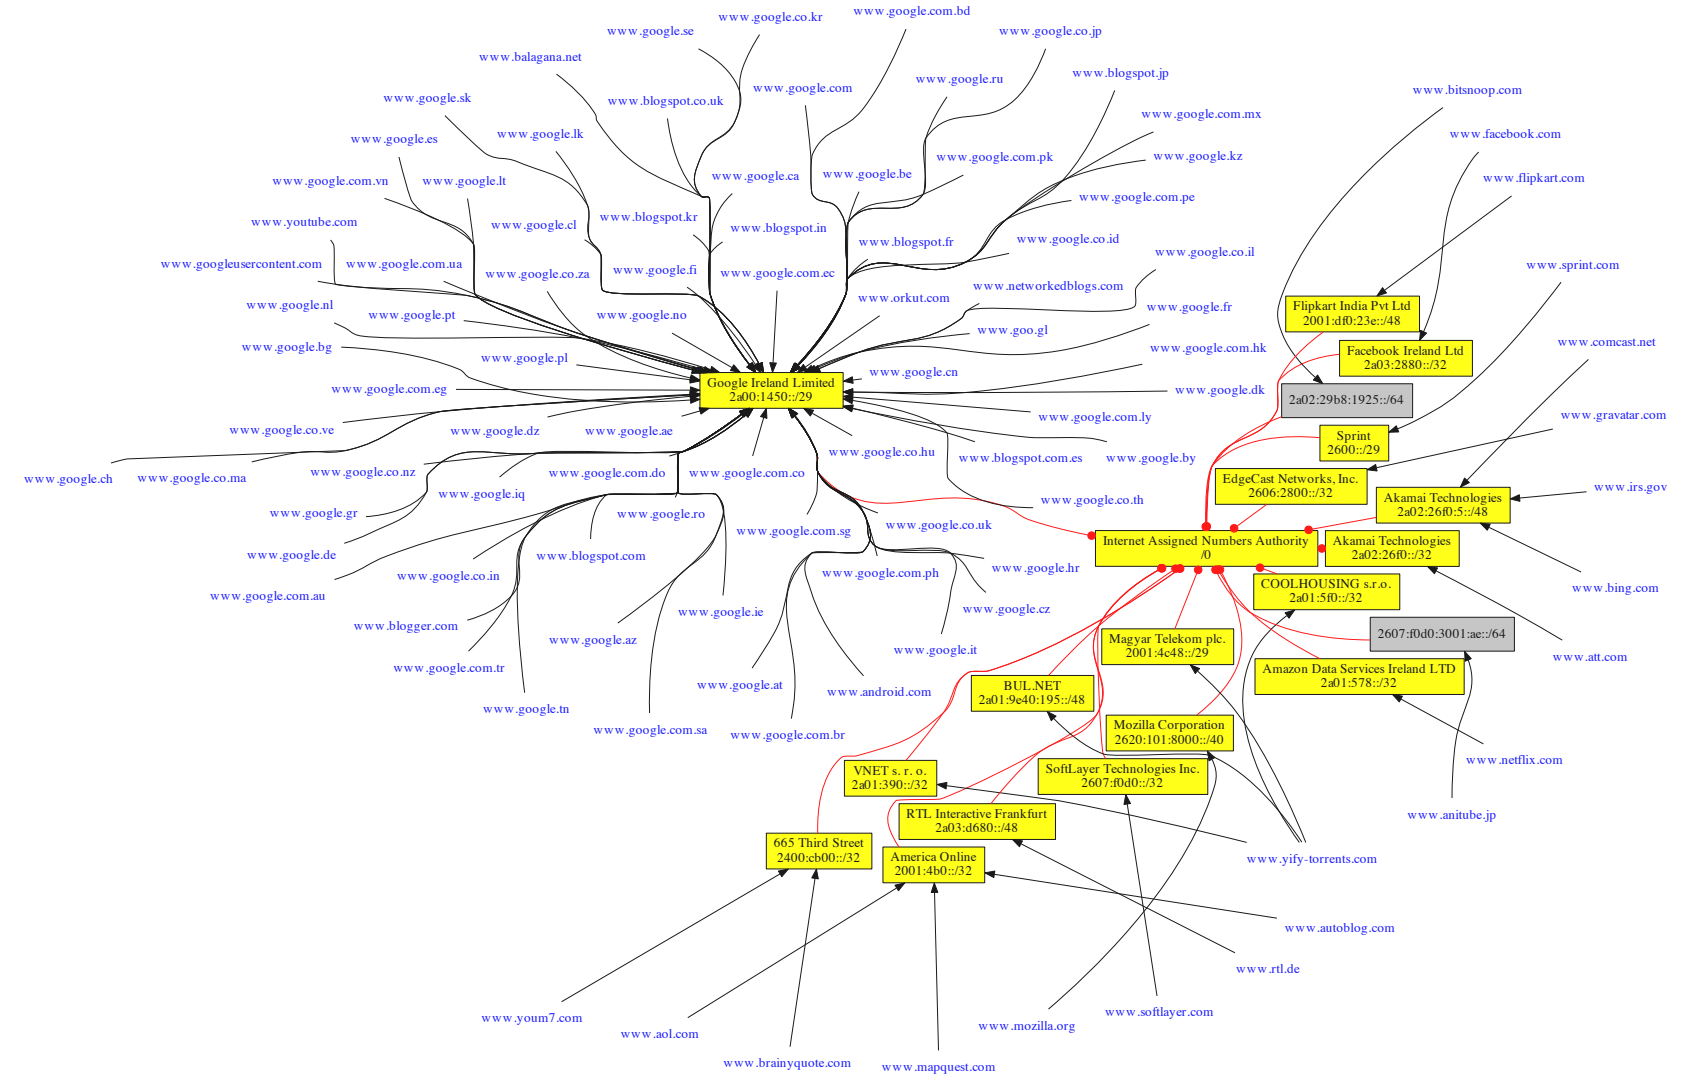
- Native IPv4 and IPv6 connectivity via DTAG - Deutsche Telekom AG [AS 3320]

Preliminary Work | Measuring Happy Eyeballs

- Do major portion of the web services centralize on CDNs?



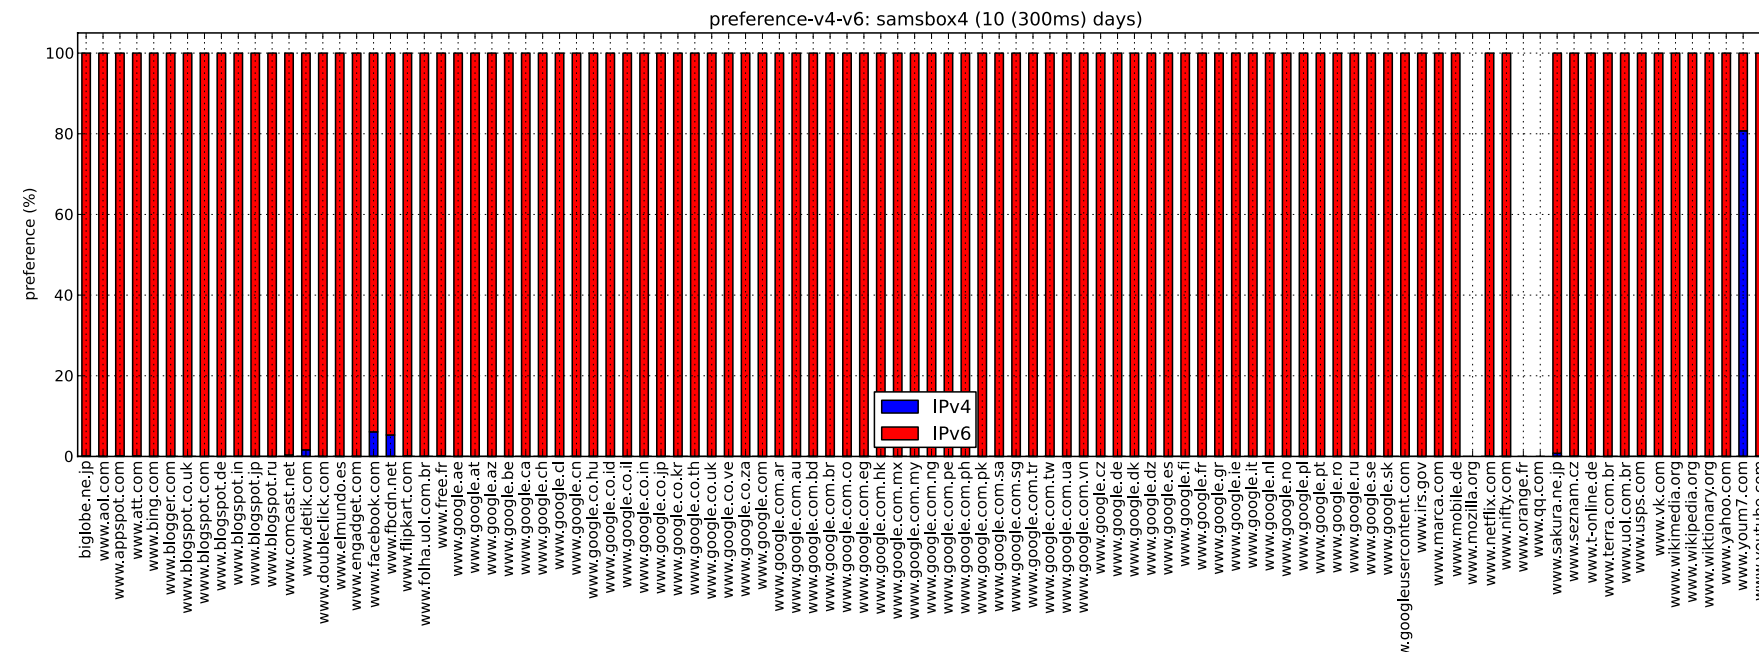
IPv4 Aggregation Cloud



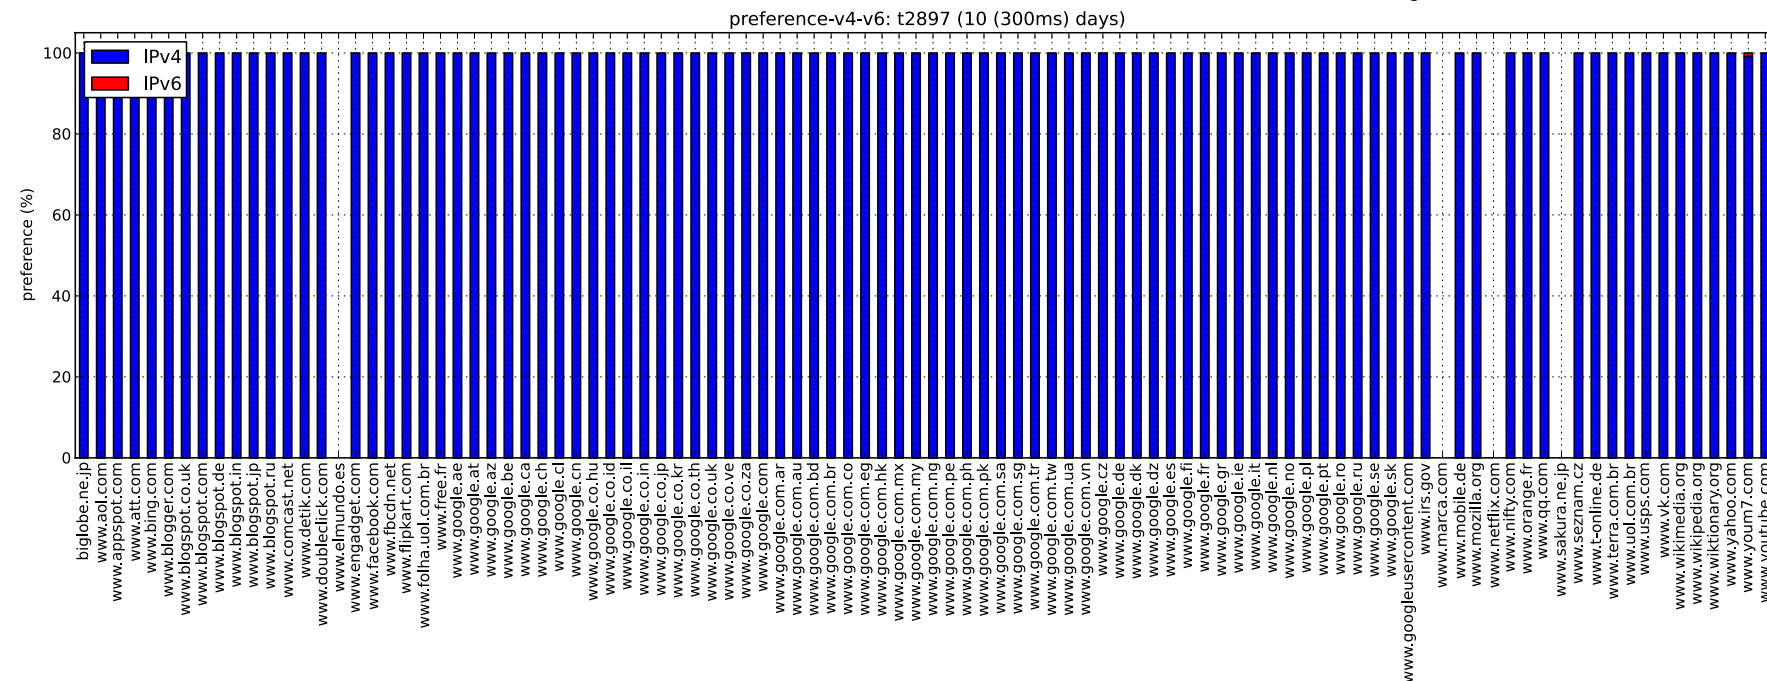
IPv6 Aggregation Cloud

Preliminary Work | Measuring Happy Eyeballs

- To what extent is IPv6 preferred when connecting to a dual-stacked service?



Native IPv4 and IPv6 connectivity via DTAG - Deutsche Telekom AG [AS 3320]



IPv4 connectivity via LambdaNet Communications [AS 13237]. IPv6 connectivity via Teredo.

Preliminary Work | Measuring Happy Eyeballs

- Data Analysis 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.
 - IPv4 and IPv6 whois aggregation clouds reveal many services centralize at Google and Akamai CDNs.
 - Measurement Agent (MA) will never use Teredo IPv6 unless IPv4 connectivity is broken.
 - A 300ms advantage leaves a MA 1% chance to prefer IPv4.

Research Questions

- Study IPv6 transition [[Bajpai-AIMS-2012](#)].
 - Can we identify a Carrier-Grade NAT (CGNAT) from a residential gateway?
 - Can we identify multiple layers of NATs from a residential gateway?
- Measure today's IPv6 network.
 - How does the performance of IPv6 compare to that of IPv4?
- Study the blend of network centralization and decentralization
 - To what extent do web services centralize on Content Delivery Networks (CDNs)?
 - To what extent does web experience depend on Regionalization?

Dissemination:

- Technical Article: Evaluating the Effectiveness of Happy Eyeballs, RIPE Labs, June 2013: https://labs.ripe.net/Members/vaibhav_bajpai/evaluating-the-effectiveness-of-happy-eyeballs
- Publication: PhD Workshop Paper, AIMS, June 2013
- Tutorial: Large Scale Measurement Platforms, AIMS, June 2013
- Invited Talk: Measuring the Effectiveness of Happy Eyeballs, RIPE 66, May 2013: <http://ripe66.ripe.net/archives/video/1208/>

References

- [1] M. Dischinger, et al., Characterizing Residential Broadband Networks, ACM Conference on Internet Measurement Conference (IMC), 2007.
- [2] Y. Shavitt, et al., DIMES: Let the Internet Measure Itself, ACM Computer Communications Review (CCR), 2005.
- [3] M. Dischinger, et al., Glasnost: Enabling End Users to Detect Traffic Differentiation, USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2010
- [4] C. Kreibich, et al., Netalyzr: Illuminating the Edge Network, ACM Conference on Internet Measurement Conference (IMC), 2010
- [5] M. Dhawan, et al., Fathom: A Browser-based Network Measurement Platform, ACM Conference on Internet Measurement Conference (IMC), 2012

References

- [6] V. Bajpai, et al., Flow-based Identification of Failures caused by IPv6 Transition Mechanisms, 6th Conference on Autonomous Infrastructure, Management and Security (AIMS), 2012
- [7] A. Dhamdhere, et al., Measuring the Deployment of IPv6: Topology, Routing and Performance, ACM Conference on Internet Measurement Conference (IMC), 2012
- [8] M. Allman, et al., Assessing IPv6 Adoption, ACM Special Interest Group (SIG) for Computer Systems Performance Evaluation (SIGMETRICS), 2013
- [9] L. Colitti, et al., Evaluating IPv6 Adoption in the Internet, Proceedings of Passive and Active Measurements (PAM), 2010