Measuring the Effects of Happy Eyeballs

Vaibhav Bajpai
Jacobs University, Bremen

RIPE 73
Madrid, Spain

Published at:
Applied Networking Research Workshop
July 2016: https://goo.gl/ZiIcQV

Joint work with

Jürgen Schönwälder
Jacobs University, Bremen

October 2016

Supported by:
Flamingo Project: flamingo-project.eu
Leone Project: leone-project.eu

- HE timer (300 ms) was chosen (2012) when *broken* IPv6 connectivity was prevalent.
  - Largely attributed to *failures* caused by Teredo [1] and 6to4 relays [2].
  - Even in situations where relays work, Teredo / 6to4 add *noticeable* latency [3, 4].
- These transition mechanisms have *declined* over the years due to efforts such as –

  2013 Microsoft *stopped* Teredo on Windows and *deactivated* public Teredo servers [5].
  2015 The 6to4 anycast prefix has been *obsoleted* [6].
- Consequentely, failure rates over IPv6 [7] have *dropped* significantly –

	Overall	Native
2011	40%	5.3%
2015	3.5%	2%

Fragmentation of HE is visible in browser implementations today –

- 2011 Chrome uses 300 ms [8]. [since v11]
- 2011 Safari uses history of witnessed latencies [9]. [since OS X 10.7]
- 2012 Opera uses parallel TCP connections [10]. [since v12.10]
- 2012 Firefox uses parallel TCP connections [11]. [since v15]

  Firefox `[network.http.fast-fallback-to-IPv4=false]` uses 250 ms.
- 2015 Safari uses 25 ms + history of witnessed latencies [12]. [since OS X 10.11 / iOS 9]

These HE timer values are arbitrarily chosen. What is the *right* timer value?

We measure against ALEXA top 10K websites for 3 years (2013 - 2016)

---

1. TCP connect times to websites over IPv6 have considerably *improved* over time.
2. 18% of websites are *faster* over IPv6 with 91% being at most 1 ms slower (May '16).
3. HE (300 ms) makes 99% of websites prefer IPv6 more than 98% of the time.
4. Slower IPv6 connections are preferred in ∼90% of the cases.
5. Lowering HE (150 ms) gives a margin benefit of 10% and retains same preference levels.

Methodology[1]

[1] Please see previous work [13] for a more detailed description of our methodology

Methodology | Metrics and Implementation

- Uses `getaddrinfo(…)` to resolve service names.
- Uses non-blocking TCP `connect(…)` calls.
- DNS resolution time is not accounted.
- Can read multiple service names as arguments.
- Can read service names list from a file.
- File locking capability.
- Sets a delay between `connect(…)`; avoids SYN floods.
- Can produce both human-readable & CSV output.
- Cross-compiled for OpenWrt; Running on SamKnows.

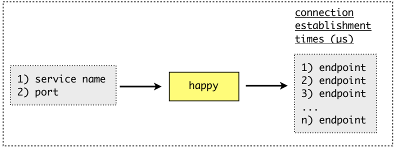

happy.vaibhavbajpai.com.

```
% 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
```

- We use the ALEXA top 10K websites as measurement targets [13].

1. `www.google.com`
2. `www.facebook.com`
3. `www.youtube.com`
4. `www.yahoo.com`
5. `www.wikipedia.org`
6. `www.qq.com`
7. `www.blogspot.com`
8. ...

The happy test repeats every hour.

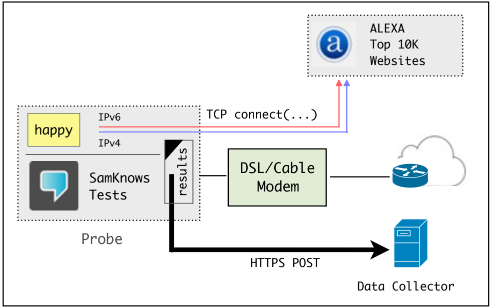

Methodology | Measurement Trial

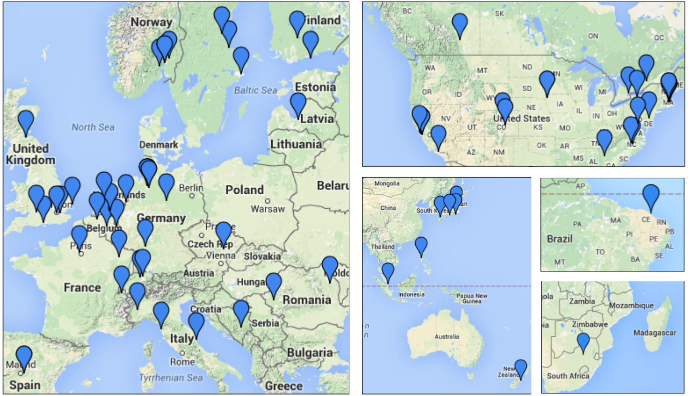

NETWORK TYPE	#
RESIDENTIAL	55
NREN / RESEARCH	11
BUSINESS / DATACENTER	09
OPERATOR LAB	04
IXP	01

RIR	#
RIPE	42
ARIN	29
APNIC	07
AFRINIC	01
LACNIC	01

We measure from 80 dual-stacked SamKnows [14] probes.

Data Analysis

[2013 - 2016]

$$\Delta s_a(u) = t_4(u) - t_6(u)$$

where $t(u)$ is the time taken to establish TCP connection to website u.

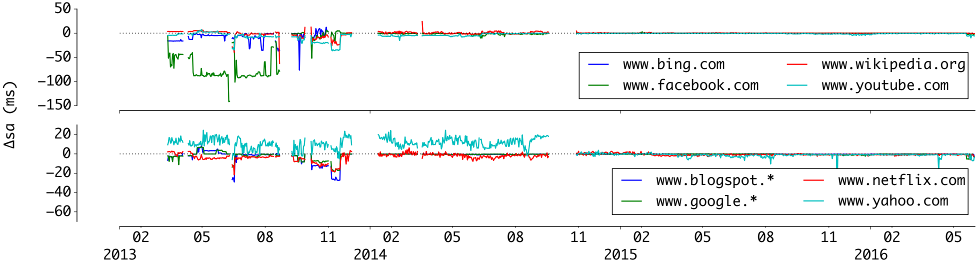

- TCP connect times to popular websites over IPv6 have *considerably* improved over time.

ALEXA top 10K websites (as of May 2016):

- 18% are *faster* over IPv6.
- 91% of the rest are at most 1 ms slower.
- 3% are at least 10 ms slower.
- 1% are at least 100 ms slower.

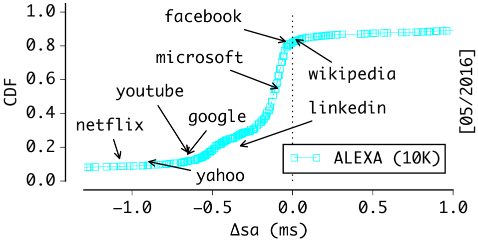

$$\Delta s_a(u) = t_4(u) - t_6(u)$$

Data Analysis | Preference

- Only ~1% of samples above HE timer value > 300 ms

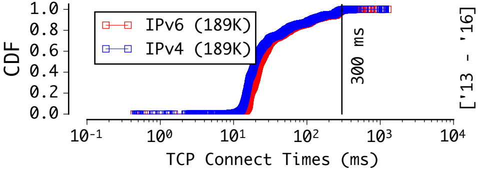

- A 300 ms HE timer value leaves 2% chance for IPv4.
- 99% of top 10K ALEXA prefer IPv6 98% of time.

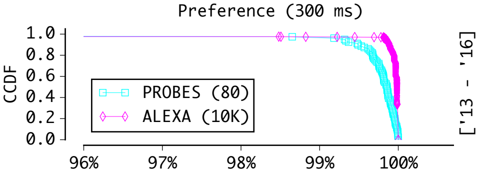

Data Analysis | Slowness

Samples where HE *prefers* IPv6 –

- HE prefers slower IPv6 connections 90% of the time.
- Absolute difference is not that far apart from IPv4
  - 30% – at least 1 ms slower.
  - 7% – at least 10 ms slower.

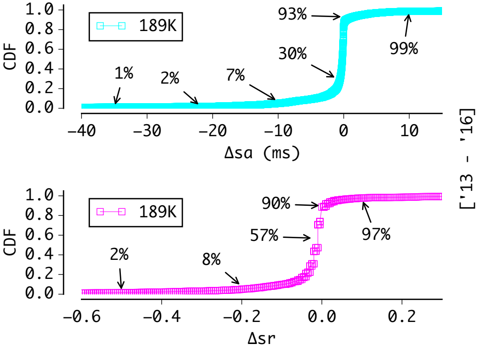

$$\Delta s_a(u) = t_4(u) - t_6(u)$$

$$\Delta s_r(u) = \frac{t_4(u) - t_6(u)}{t_4(u)}$$

Can a lower HE timer provide same preference over IPv6 but not penalise IPv4 when it's faster?

Are we ready to disable HE entirely?

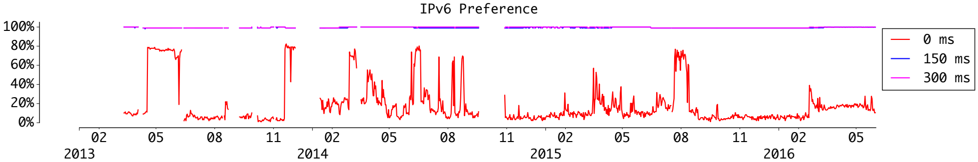

- 18% of ALEXA top 10K websites are faster (see slide 17) over IPv6 today.
- Parallel TCP connections[2] (HE with 0 ms timer) will *hamper* IPv6 preference.
- HE timer today still should give IPv6 a *fair* chance to succeed.

[2]such as used by Firefox and Opera today

- We control two[3] parameters and lower the HE timer value.
- Each data point is the 1th percentile preference towards ALEXA 10K websites.

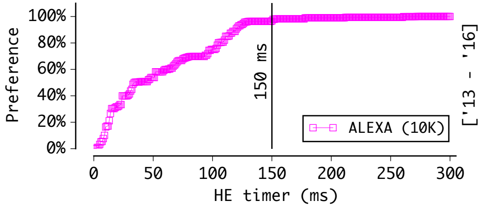

- Lowering to 150 ms retains preference levels over IPv6.
- We get margin benefit of 10% (18.9K) because timer cuts early.

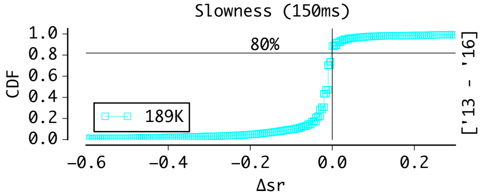

[3] 99% ALEXA top 10K websites prefer IPv6 connections 98.6% of the time

Limitations

1. The comparison reflects the performance as seen over TCP port 80 only.
2. The measurements cover ALEXA top 10K websites only.
3. The results are biased by our vantage points (centered largely around EU, US and JP).

Takeway

1. TCP connect times to websites over IPv6 have considerably *improved* over time.
2. 18% of websites are *faster* over IPv6 with 91% being at most 1 ms slower (May '16).
3. HE (300 ms) makes 99% of websites prefer IPv6 more than 98% of the time.
4. Slower IPv6 connections are preferred in ∼90% of the cases.
5. Lowering HE (150 ms) gives a margin benefit of 10% and retains same preference levels.

---

www.vaibhavbajpai.com

v.bajpai@jacobs-university.de | @bajpaivaibhav

Appendix

References

[1] C. Huitema, “Teredo: Tunneling IPv6 over UDP through Network NATs,” RFC 4380, Internet Engineering Task Force, Feb. 2006, https://tools.ietf.org/html/rfc4380.

[2] B. Carpenter and K. Moore, “Connection of IPv6 Domains via IPv4 Clouds,” RFC 3056, Internet Engineering Task Force, Feb. 2001, https://tools.ietf.org/html/rfc3056.

[3] S. Zander, L. L. H. Andrew, G. J. Armitage, G. Huston, and G. Michaelson, “Investigating the IPv6 Teredo Tunnelling Capability and Performance of Internet Clients,” ser. Computer Communication Review (CCR) ’12, 2012, pp. 13–20. [Online]. Available: http://doi.acm.org/10.1145/2378956.2378959

[4] L. Colitti, S. H. Gunderson, E. Kline, and T. Refice, “Evaluating IPv6 Adoption in the Internet,” ser. Passive and Active Measurement Conference (PAM) ’10, 2010, pp. 141–150. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-12334-4_15

[5] “Christopher Palmer - Teredo at Microsoft: Present and Future,” http://goo.gl/9I65Wy, [Online; accessed 10-February-2016].

[6] O. Troan and B. Carpenter, “Deprecating the Anycast Prefix for 6to4 Relay Routers,” RFC 7526, Internet Engineering Task Force, May 2015, https://tools.ietf.org/html/rfc7526.

[7] “Geoff Huston - Measuring IPv6 Performance,” https://goo.gl/n78W1t, [Online; accessed 10-February-2016].

[8] “Google Chrome - Revision 85934: Add a fallback socket connect() for IPv6.” https://goo.gl/nPhilZ, [Online; accessed 25-January-2016].

[9] J. Graessley, “Apple - Lion and IPv6,” http://goo.gl/uAPlV8, [Online; accessed 25-January-2016].

[10] “Opera 12.10 - Changelog,” http://goo.gl/MGsn4K, [Online; accessed 25-Jan-2016].

[11] “Mozilla Firefox 15 - Release Notes,” http://goo.gl/hA15eu, [Online; accessed 25-January-2016].

[12] D. Schinazi, “Apple and IPv6 - Happy Eyeballs,” https://goo.gl/1nzMs6, [Online; accessed 25-January-2016].

[13] V. Bajpai and J. Schönwälder, “IPv4 versus IPv6 - who connects faster?” ser. IFIP NETWORKING ’15, 2015, pp. 1–9. [Online]. Available: http://dx.doi.org/10.1109/IFIPNetworking.2015.7145323

[14] ——, “A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts,” ser. IEEE Communications Surveys and Tutorials (COMST) ’15, 2015, pp. 1313–1341. [Online]. Available: http://dx.doi.org/10.1109/COMST.2015.2418435

[15] D. Thaler, R. Draves, A. Matsumoto, and T. Chown, “Default Address Selection for Internet Protocol Version 6 (IPv6),” RFC 6724 (Proposed Standard), Internet Engineering Task Force, Sep. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6724.txt

[16] "Teemu Savolainen - Experiences of host behavior in broken IPv6 networks," http://goo.gl/4NnRiH, [Online; accessed 25-January-2016].

[17] P. Richter, M. Allman, R. Bush, and V. Paxson, "A Primer on IPv4 Scarcity," ser. Computer Communication Review (CCR), vol. 45, no. 2. New York, NY, USA: ACM, Apr. 2015, pp. 21–31. [Online]. Available: http://doi.acm.org/10.1145/2766330.2766335

[18] "Internet Society - World IPv6 Launch," http://www.worldipv6launch.org, [Online; accessed 11-January-2016].

[19] "Google - IPv6 Adoption Statistics," http://www.google.com/intl/en/ipv6/statistics.html, [Online; accessed 11-January-2016].

[20] "Lorenzo Colitti - Google no longer returning AAAA records?" https://goo.gl/6Z7gZM, [Online; accessed 11-January-2016].

[21] "Emile Aben - Hampering Eyeballs: Observations on Two Happy Eyeballs Implementations," https://goo.gl/3xVUIO, [Online; accessed 10-February-2016].

[22] "Geoff Huston - Dual Stack Esotropia," http://goo.gl/N1qUib, [Online; accessed 10-February-2016].

[23] "Geoff Huston - Bemused Eyeballs: Tailoring Dual Stack Applications for a CGN Environment," http://goo.gl/LMPc4h, [Online; accessed 10-February-2016].

[24] F. Baker, "Testing Eyeball Happiness," RFC 6556, Internet Engineering Task Force, 2012, https://tools.ietf.org/html/rfc6556.

[25] S. Zander, L. L. H. Andrew, G. J. Armitage, G. Huston, and G. Michaelson, "Mitigating Sampling Error when Measuring Internet Client IPv6 Capabilities," ser. Internet Measurement Conference (IMC) '12, 2012, pp. 87–100. [Online]. Available: http://doi.acm.org/10.1145/2398776.2398787

[26] V. Bajpai and J. Schönwälder, "Measuring the Effects of Happy Eyeballs," Internet Engineering Task Force, Internet-Draft draft-bajpai-happy-01, Jul. 2013, work in Progress. [Online]. Available: http://tools.ietf.org/html/draft-bajpai-happy-01

[27] S. Ahsan, V. Bajpai, J. Ott, and J. Schönwälder, "Measuring YouTube from Dual-Stacked Hosts," ser. Passive and Active Measurement Conference (PAM) '15, 2015, pp. 249–261. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-15509-8_19

Introduction | getaddrinfo(...) behavior

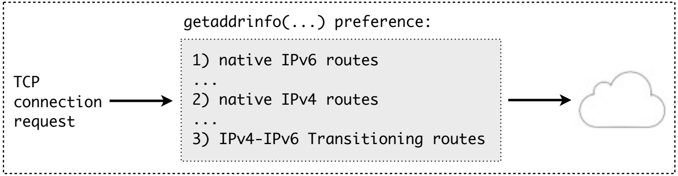

- returns a list of endpoints in an order that prioritizes an IPv6-upgrade path.
- The order is prescribed by RFC 6724 [15] and `/etc/gai.conf`
- Iterating sequentially over the list of IP endpoints has repercussions –
  - Broken IPv6 connectivity makes apps stall for *several* seconds before trying IPv4.
  - Studies have reported [16] browser connection timeouts in the order of 20 seconds.

HE helps *prevent* bad QoE in situations where IPv6 connectivity is broken.

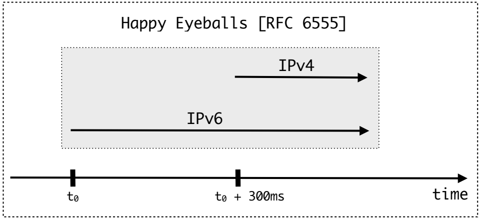

Design Goals –

- Honor the destination address selection policy [RFC 6724] [15].
- Quickly fallback to IPv4 when IPv6 connectivity is broken.
- Give a *fair* chance for IPv6 to succeed.

IPv6 landscape has changed today –

- 4/5 RIRs have *exhausted* available pool of IPv4 address space [17].

APNIC	Apr′11
RIPE	Sep′12
LACNIC	Jun′14
ARIN	Sep′15

- Large IPv6 broadband rollouts[4] since World IPv6 Launch Day in 2012 [18].
- IPv6 global adoption at ∼12.2% (native) with Teredo / 6to4 at ∼0.01% [19] (July 2016)
- Google over IPv6 (whitelist) program *replaced* by a Google IPv6 blacklist [13].
- Google will not return AAAA to resolvers where latency over IPv6 > 100 ms worse [20].

[4] Comcast, Deutsche Telekom AG, AT&T, Verizon Wireless, T-Mobile USA

The *effects* of HE (300 ms) on the QoE of a dual-stacked user remains *largely* unclear.

We want to know –

- *In what percentage of cases HE makes a bad decision of choosing IPv6 when it's slower?*
- *In such situations what is the amount of imposition (in terms of latency impact) a dual-stacked user has to pay as a result of the high HE timer (300 ms) value?*

Applications apply HE not only where IPv6 is broken, but also when IPv6 is comparable.

Related Work

2011 - 2012 Studies [21, 22, 23] have analyzed HE implementations.

- Chrome reduces degraded user experience when IPv6 is broken.
- Firefox `[network.http.fast-fallback-to-IPv4=false]` behaves similar to Chrome.
- Safari prefers IPv4 even when IPv6 connectivity is similar (*hampering eyeballs*).

These studies are dated. HE implementations have *changed* with time (see slide 7).

2012 Baker [24] describes HE metrics and testbed configurations.

2012 Zander [25] showed that 75% of the connection attempts preferred[5] IPv6.

2013 We [26] showed that HE never prefers IPv6 using Teredo.

2015 We [27] showed that HE prefers YouTube over IPv6 even when IPv4 performs better.

[5] In this work, we show that this preference has increased to 98% today