A Longitudinal View of Dual-stacked Websites — Failures, Latency and Happy Eyeballs

Vaibhav Bajpai
Technische Universität München

NetSys 2019 (Invited Talk)
Garching bei München, Germany

IEEE/ACM Transactions on Networking
March 2019
https://doi.org/10.1109/TNET.2019.2895165

March 20, 2019

Motivation

- Literature focus *largely* on IPv6 adoption.
- Very little work on measuring IPv6 performance.

Do users experience benefit (or suffer) from web content delivery over IPv6? What factors contribute to the difference? How has content delivery over IPv6 evolved over years?

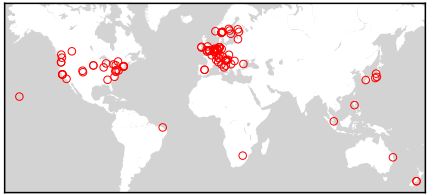

∼100 dual-stacked SamKnows probes (∼66 origin ASes)

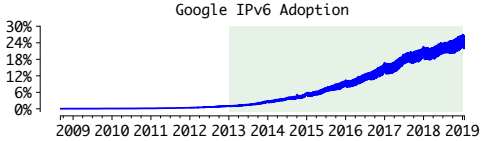

shaded region represents the duration of the longitudinal study.

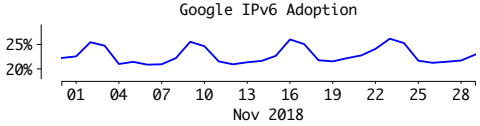

- ¼ of connections to Google are made over IPv6.
- IPv6 penetration more in home deployments.

Research Contributions

- Penetration of Content Delivery Networks (CDNs)
- Latency, Failures and Happy Eyeballs
- IP Path Lengths and Content Cache Deployments

This is the first study to provide a longitudinal view (6 years) of web content delivery over IPv6.

- Relevance:
  - Network operators in *early* stages of IPv6 deployment.
  - Content providers to see how their *service delivery* over IPv6 compares to IPv4.
  - Drive related *standards* work in the IETF.

Content Delivery Networks (CDNs)

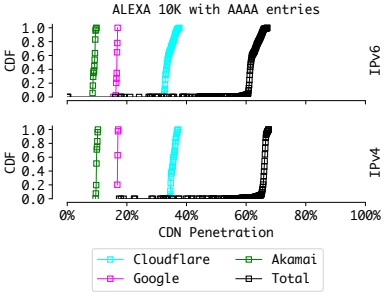

- CDNs play a *leading* role in technology adoption.
- Leading players:
  1. Cloudflare (∼35%)
  2. Google (∼16%) and
  3. Akamai (∼9%)

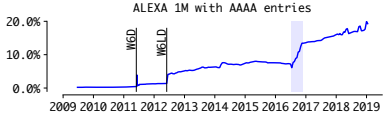

- ⅕ of 1M websites announce AAAA entries (2019).
- Cloudflare added AAAA entries for all websites [1].
- Cloudflare's impact (shaded) > W6D (or W6LD).

CDNs serve (2018) more than ½ of 10K websites.

Complete Failures

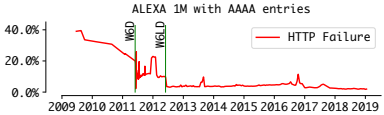

- Failures reduced from 40% (2009) to 2% (2019).

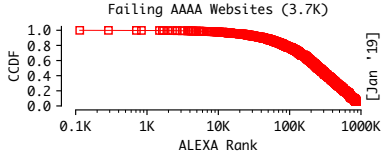

- 88% failing websites rank > 100K.
- 1% rank < 10K, five websites rank < 300.

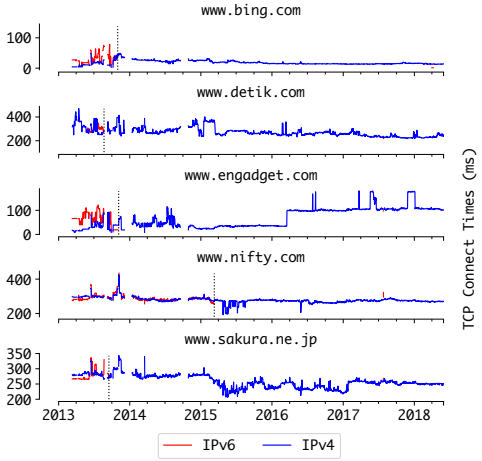

Failures in content delivery over IPv6 have reduced to 2%

Partial Failures

ALEXA top 100 websites with AAAA entries.

- 27% show some rate of failure over IPv6.
- 9% exhibit more than 50% failures over IPv6.

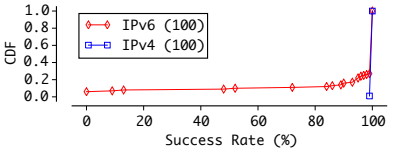

- Limiting to root webpage can lead to overestimation of IPv6 adoption numbers
- Unclear whether websites with partial failures can be deemed *IPv6-ready*
- ISOC now supporting [2] development of *tools* that identify such partial failures

#	Webpage	Success Rate (%) IPv6(↓)	Success Rate (%) IPv4	W6LD
01	www.bing.com	0	100	✓
02	www.detik.com	0	100	✓
03	www.engadget.com	0	100	✓
04	www.nifty.com	0	100	
05	www.qq.com	0	100	
06	www.sakura.ne.jp	0	100	
07	www.flipkart.com	09	99	✓
08	www.folha.uol.com.br	13	100	
09	www.aol.com	48	100	✓
10	www.comcast.net	52	100	✓
11	www.yahoo.com	72	100	✓
12	www.mozilla.org	84	100	✓
13	www.orange.fr	86	100	✓
14	www.seznam.cz	89	100	✓
15	www.mobile.de	90	100	✓
16	www.wikimedia.org	90	100	
17	www.t-online.de	93	100	✓
18	www.free.fr	95	100	
19	www.usps.com	95	100	
20	www.vk.com	95	100	✓
21	www.wikipedia.org	95	100	✓
22	www.wiktionary.org	95	100	
23	www.elmundo.es	96	100	✓
24	www.uol.com.br	96	100	✓
25	www.marca.com	97	100	✓
26	www.terra.com.br	98	100	✓
27	www.youm7.com	99	100	

Partial Failures | Root Cause Analysis

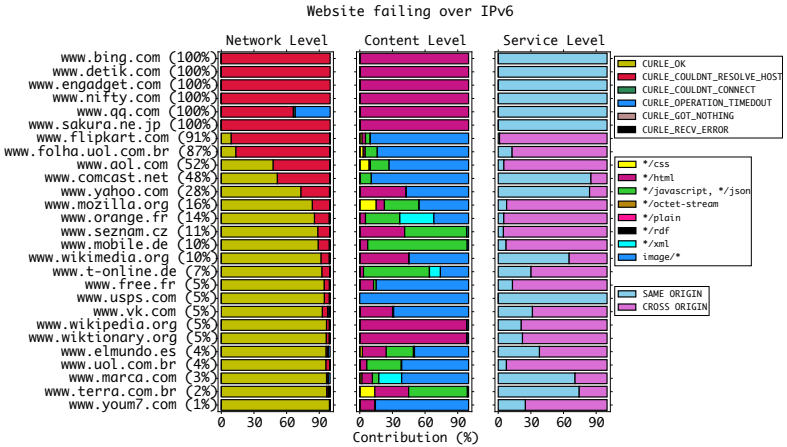

- Failures due to DNS resolution error on `image/*`, `*/javascript`, `*/json` and `*/css` content.

- Failures *silently* exist; clients do not notice them due to IPv4 fallback.
- Identification of operational issues relevant for upcoming IPv6-only networks

Partial Failures | Root Cause Analysis

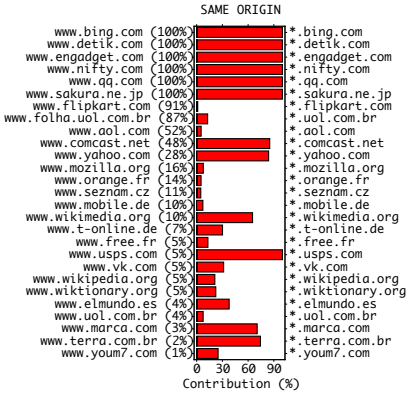

- 12% of websites have more than 50% webpage elements that belong to the same origin source and fail over IPv6.

- CDN infrastructure does not have IPv6 turned on by default for all same-origin webpage elements.

#	Webpage	Same Origin (↓)
01	www.bing.com	100%
02	www.detik.com	100%
03	www.engadget.com	100%
04	www.nifty.com	100%
05	www.usps.com	100%
06	www.qq.com	100%
07	www.sakura.ne.jp	100%
08	www.comcast.net	85%
09	www.yahoo.com	83%
10	www.terra.com.br	74%
11	www.marca.com	70%
12	www.wikimedia.org	65%
13	www.elmundo.es	37%
14	www.vk.com	31%
15	www.t-online.de	30%
16	www.youm7.com	24%
17	www.wiktionary.org	22%
18	www.wikipedia.org	22%
19	www.free.fr	13%
20	www.folha.uol.com.br	12%
21	www.mozilla.org	7%
22	www.uol.com.br	7%
23	www.mobile.de	7%
24	www.aol.com	5%
25	www.orange.fr	5%
26	www.seznam.cz	4%
27	www.flipkart.com	1%

Partial Failures | Root Cause Analysis

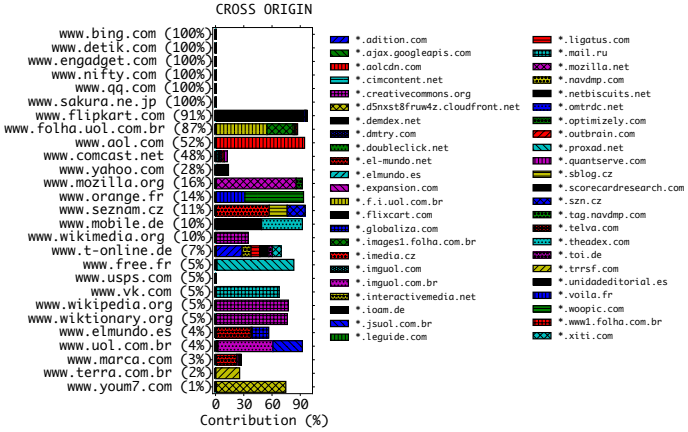

- Third-party advertisements (*.doubleclick.net)
- Analytics (*.scorecardresearch.com, *.quantserve.com)
- User-centric content (*.facebook.com, *.ajax.googleapis.com)
- Static content (*.wikimedia.org, *.creativecommons.org)

- Enabling IPv6 on few cross cross-origin sources (creativecommons.org, doubleclick.net) will help reduce partial failure of multiple websites.

IPv6 Preference

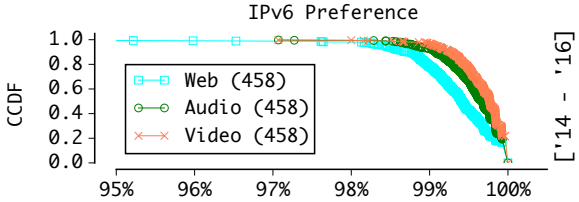

- RFC 6724 [3] makes apps prefer connections made over IPv6.
- RFC 6555 [4] allows apps to fallback to IPv4 when IPv6 connectivity is bad.
- TCP connections over IPv6 are preferred at least 97% of the time.

Clients prefer web and video content over IPv6

Latency

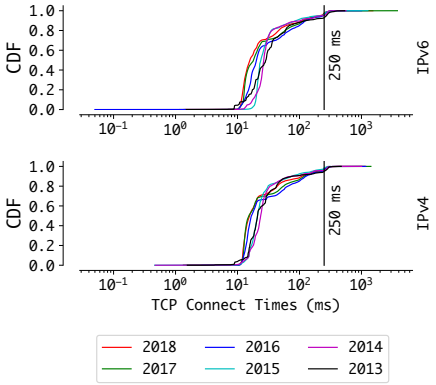

Latency (2013 → 2018) reduced by:

- 29% (21 ms → 15 ms) over IPv4
- 57% (28 ms → 16 ms) over IPv6

- Latency over IPv6 has reduced by more than half in past six years (2013 - 2018).

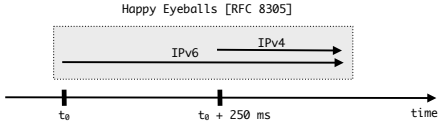

- Only 3% of samples above HE timer > 250 ms
- Clients prefer IPv6 when slower in 81% cases.
- HE timer can be reduced to 150 ms.

- RFC 6555 should have used 150 ms timer.
- Measurements should inform IETF work.

Latency

- 50% samples used to 5 ms or more slower over IPv6 (2013); reduced to 8% (2018).

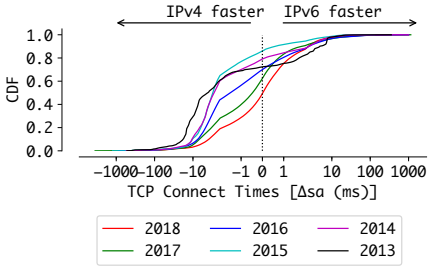

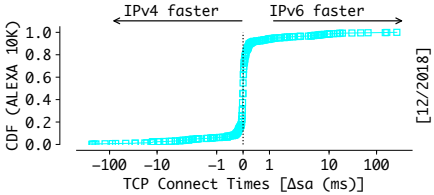

- 56% websites are *faster* over IPv6.
- 95% of the rest are at most 1 ms slower.
- 2% are at least 25 ms slower.

- More than half of ALEXA 10K websites are faster over IPv6 (2018)

How beneficial are the content caches?

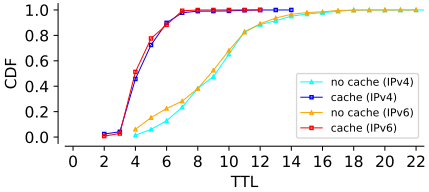

- 90% caches are reachable within 6 IP hops.

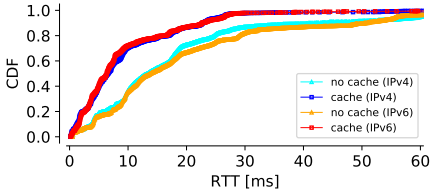

- Caches are reachable within 20 ms.
- Caches reduce latencies by:
  - IPv4: 25 ms → 17 ms; ⅓ improvement
  - IPv6: 29 ms → 16 ms; ½ improvement

- GGC caches reduce IP path lengths by ½ over both AF.
- GGC caches reduce latencies over IPv6 by ½ and over IPv4 by ⅓.

Dual-stacked content caches

- Caches always reduce IP path lengths.
- Path lengths differ when caches not hit.

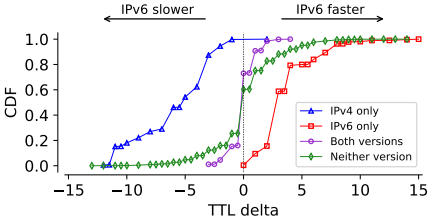

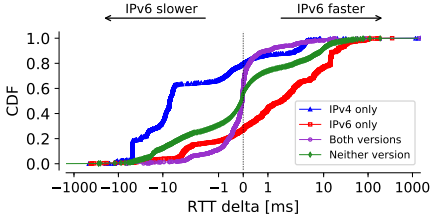

- Dual-stacked: ∼80% within 1 ms.
- IPv4-only: ∼80% faster over IPv4.
- IPv6-only: ¼ still slower over IPv6.

IPv6 performance degrades when content caches are not dual-stacked.

Takeaway

- ¼ of connections to Google are made (mostly from home networks) over IPv6.
- CDNs serve more than ½ of ALEXA 10K websites (2018).
- IPv6 failures reduced to 2%. IPv6 latency reduced by more than half in six years.
- More than ½ of ALEXA 10K websites are faster over IPv6 (2018).
- Caches can be reached in six IP hops; reduce IP path lengths by ½ over both AF.
- Caches reduce latencies over IPv6 by ½ and over IPv4 by ⅓.
- Latency over IPv6 degrades when caches are not dual-stacked.
- Reproducibility Considerations:
  https://github.com/vbajpai/2018-ton-v6-longitudinal-websites

www.vaibhavbajpai.com

bajpaiv@in.tum.de | @bajpaivaibhav

References

[1] "98.01% of sites on Cloudflare now use IPv6," https://blog.cloudflare.com/98-percent-ipv6, [Online; accessed 15-Apr-2017].

[2] "NAT64 Check," nat64check.ipv6-lab.net, [Accessed 15-Apr-2017].

[3] D. Thaler, R. Draves, A. Matsumoto, and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)," RFC 6724, Internet Engineering Task Force, Sep. 2012, https://tools.ietf.org/html/rfc6724.

[4] D. Wing and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts," RFC 6555, IETF, 2012, https://tools.ietf.org/html/rfc6555.