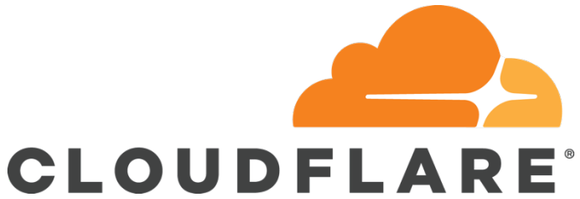

IPv6 @ Cloudflare (plus related items)

ENOG14 Minsk Belarus – October/2017

Martin J. Levy @ Cloudflare

// Personal Introduction

Martin J. Levy @ Cloudflare

MY HISTORY

A dedicated IPv6 evangelist. Long time TCP/IP developer/programmer, network operator, peering expert, IETF member, NANOG member, and IP networking development/strategy expert.

MY TERSE RESUME

Bell Labs (New Jersey) – Unix for Unix's sake, TCP/IP (1982/1983)

Random startups and ISPs (Bay Area)

Concentric/XO (Bay Area) – IP backbone and hosting

Telecom Italia (Rome & Miami) – Global IP backbone

Hurricane Electric (Bay Area) – Global IPv4/IPv6 backbone

Cloudflare (Bay Area) – Global CDN, DDoS, DNS, Security

CLOUDFLARE

// The Punchline!

// The Punchline!

At Cloudflare, IPv6 is always on!

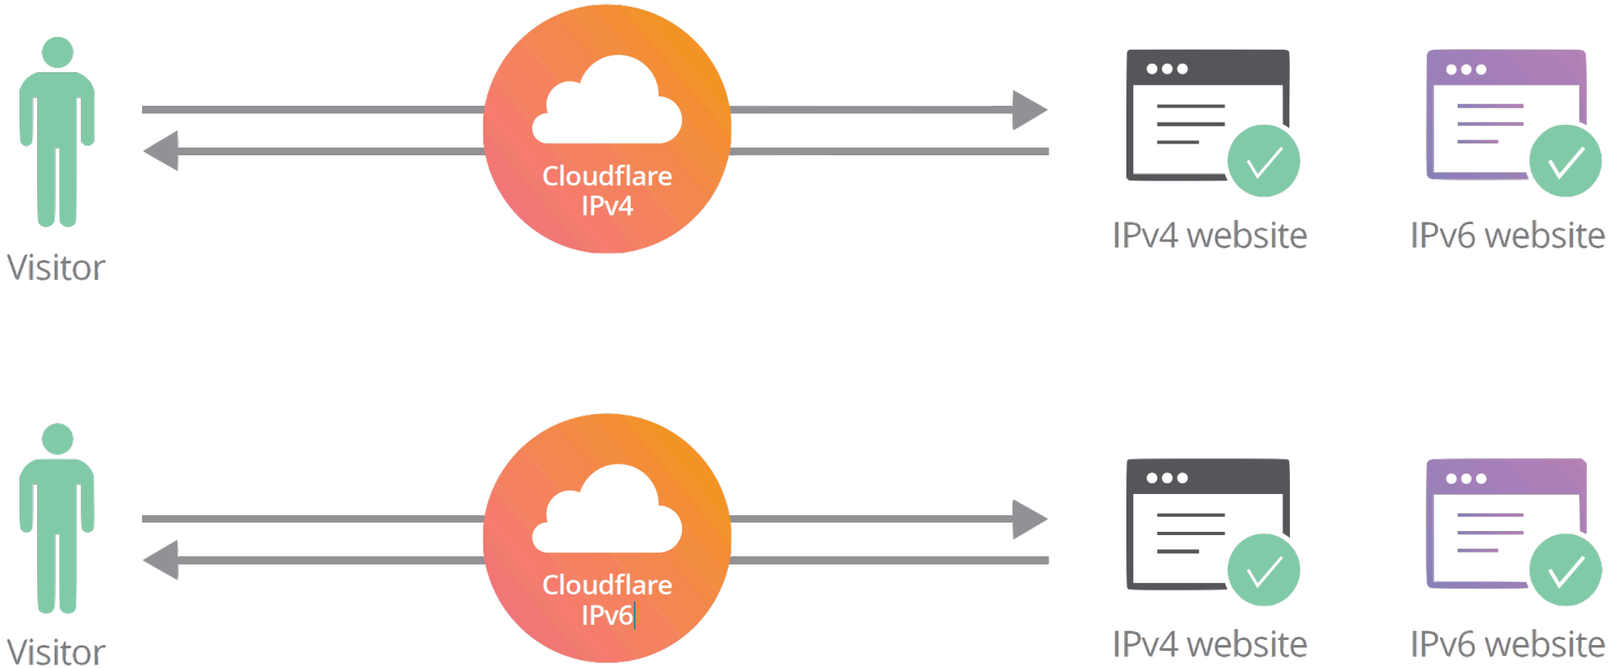

CLOUDFLARE

At Cloudflare, DNSSEC always!

// The Punchline!

CLOUDFLARE

Cloudflare provides performance, security, reliability,
and insights to anything connected to the Internet.

// Introduction to Cloudflare

AS13335 / Cloudflare's Global Anycast Network

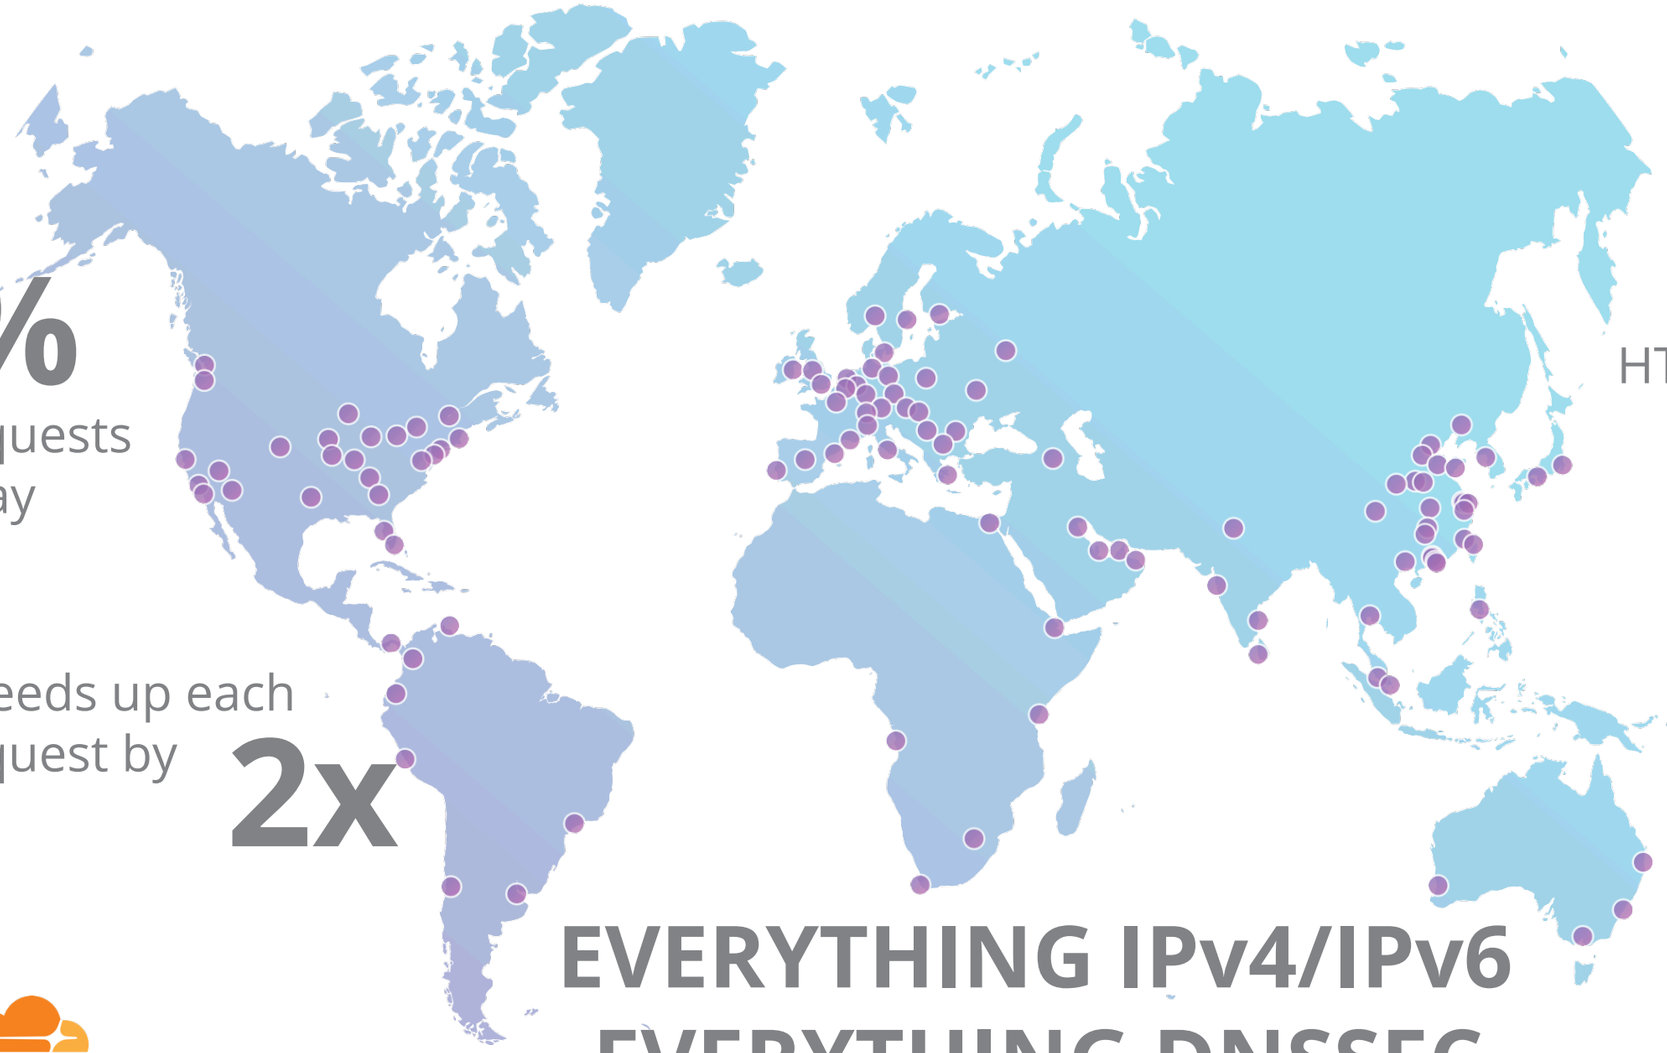

10%
Internet requests everyday

Speeds up each request by **2x**

5M
HTTP requests/second

Data centers globally
116

1.2M
DNS requests/second

EVERYTHING IPv4/IPv6
EVERYTHING DNSSEC

CLOUDFLARE

Cloudflare’s benefits

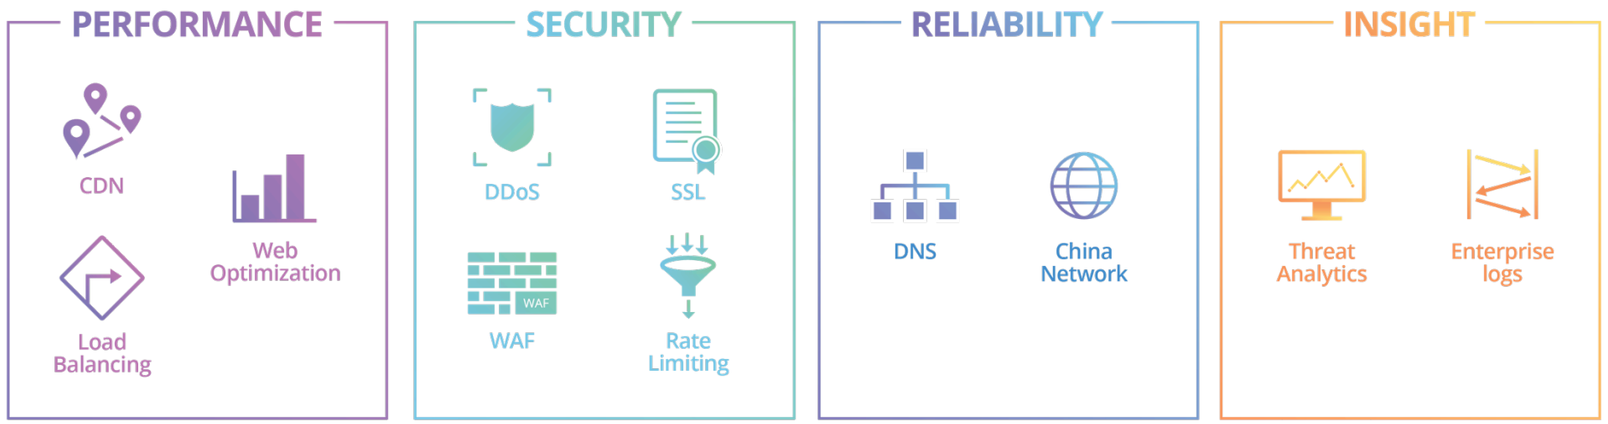

CLOUDFLARE

Performance

CDN
Moving content physically closer to visitors with our CDN.

Website Optimization
Cloudflare lets you automatically enable the latest in web technologies.

DNS
Cloudflare is one of the fastest managed DNS providers in the world.

SSL
Modern SSL isn't just for security—it can actually improve the performance of your website.

Dedicated SSL Certificates
With a few clicks within the Cloudflare dashboard, you can easily and quickly issue new certificates, securely generate private keys and more.

Load Balancing
Cloudflare Load Balancing provides load balancing, geo-steering, monitoring and failover for your Internet facing infrastructure enhancing service availability.

CLOUDFLARE

Security

DDoS Protection
Our enterprise-class DDoS protection network has 20 times more capacity than the largest DDoS attack ever recorded.

WAF
Our web application firewall benefits from the collective intelligence of our entire network.

SSL
HTTPS is a must-have for modern websites, and Cloudflare makes it easy to configure SSL.

Secure Registrar
Registering your domain through Cloudflare is the most secure way to protect your trademark from domain hijacking.

Dedicated SSL Certificates
With a few clicks within the CloudFlare dashboard, you can easily and quickly issue new certificates, securely generate private keys and more.

Rate Limiting
Rate Limiting gives you granular controls to detect bad traffic, customized rulesets to ensure that your legitimate visitors are not impacted, and insights to improve your security posture as attacks evolve.

CLOUDFLARE

Reliability

DNS
Cloudflare's DNS service is powered by the same 102 data center network that powers our DDoS and CDN services. This not only improves DNS resolution times, but also makes DNS-related attacks and outages a thing of the past.

China Network
Cloudflare's China service optimizes Internet connections in mainland China, dramatically improving the viewing experience for visitors in China.

Predictable Bandwidth Costs
We believe that you should never be surprised by your monthly bill. Our flat-rate pricing structure makes your CDN and DDoS bandwidth expenses predictable.

CLOUDFLARE

Insight

Enterprise Logs
For enterprise customers, we can provide consolidated logs from around the world. These are very rich, containing detailed information about every request and response.

Threat Analytics
When we identify requests that are threats, we log them and block them. That means we not only protect your site, but also provide insight into the malicious activity we're seeing.

Rate Limiting
Rate Limiting gives you granular controls to detect bad traffic, customized rulesets to ensure that your legitimate visitors are not impacted, and insights to improve your security posture as attacks evolve.

CLOUDFLARE

A few of our Technology customers

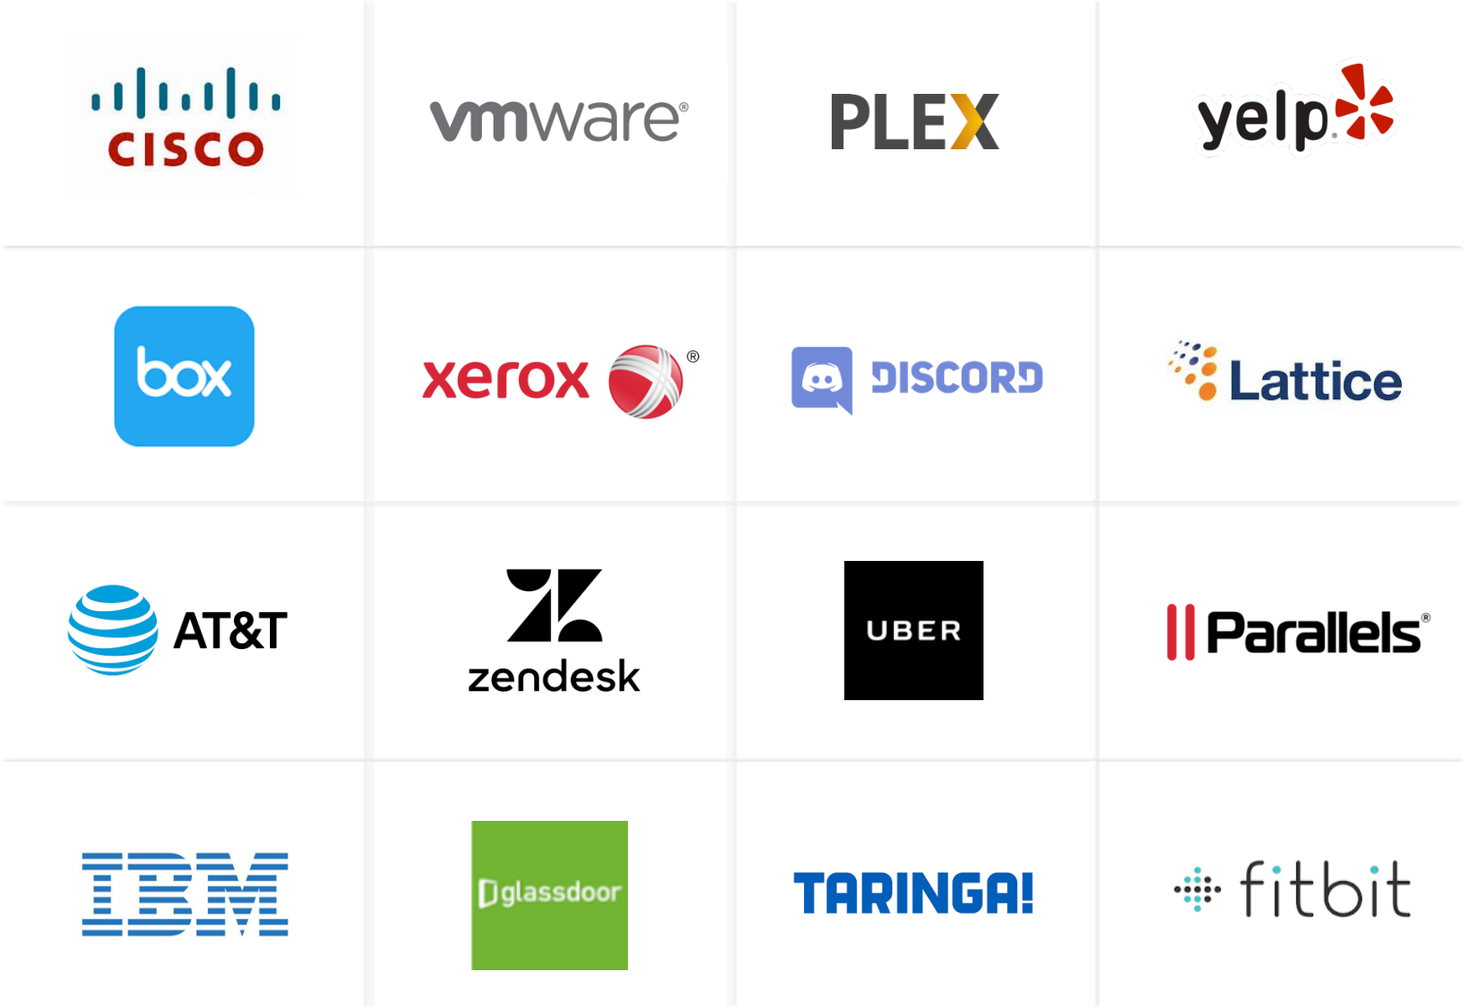

CLOUDFLARE

Cloudflare has a solid history of giving back to the community, both in open-source software, IETF protocol development, network services, etc.

// Now Down to the Technical Parts ...

The Technical Part

1. Backstory behind the IPv6 switch at Cloudflare
2. Why we removed the switch!
3. DNSSEC at Cloudflare
4. A serious discussion about DNS in a v6 world

IPv6 @ Cloudflare is so 2606:4700::5ca1:ab1e:6810:4737

Cloudflare can be a “bridge” to IPv6

CLOUDFLARE

Cloudflare can be a “bridge” to IPv6

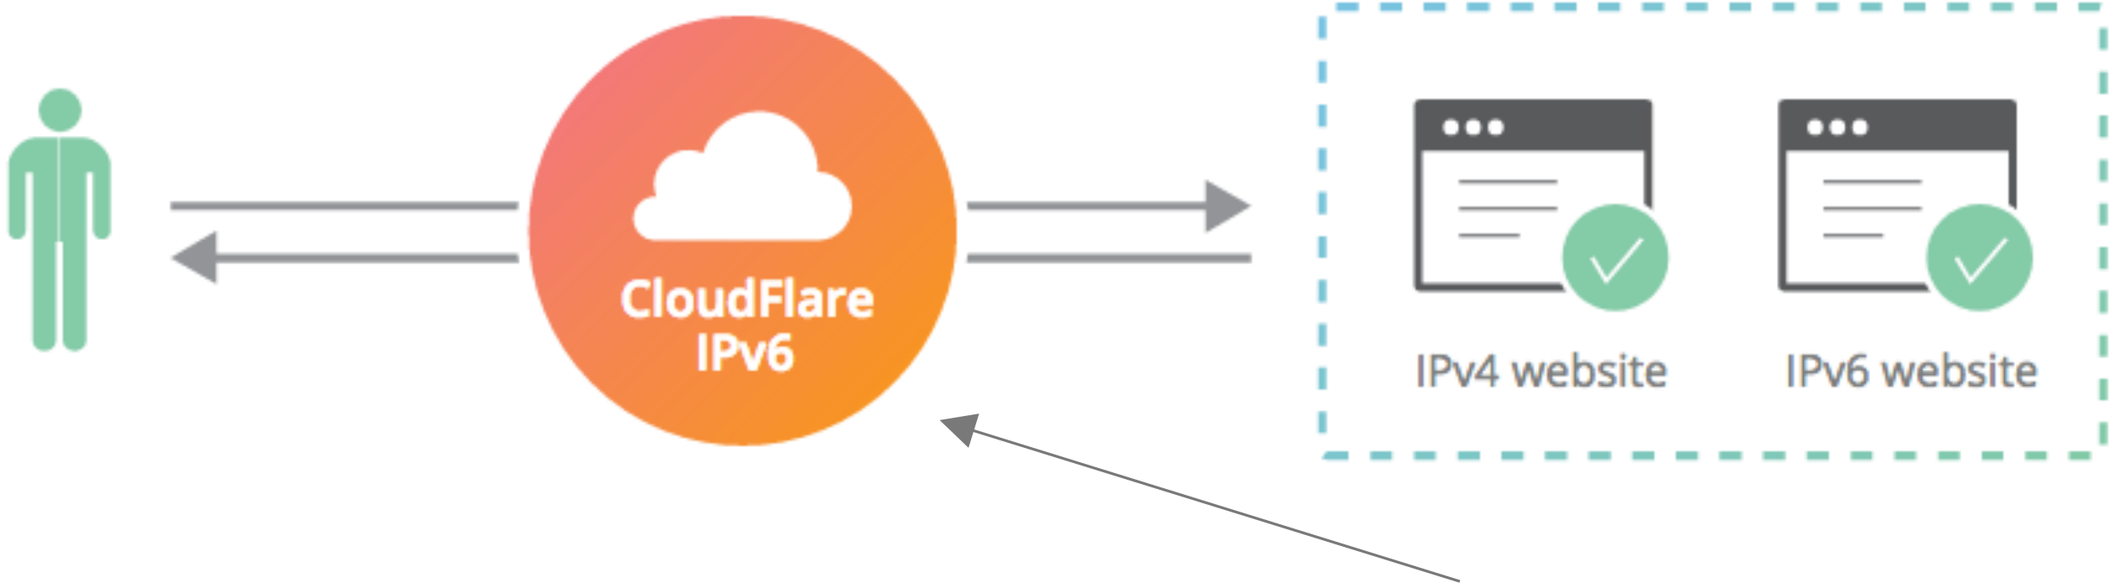

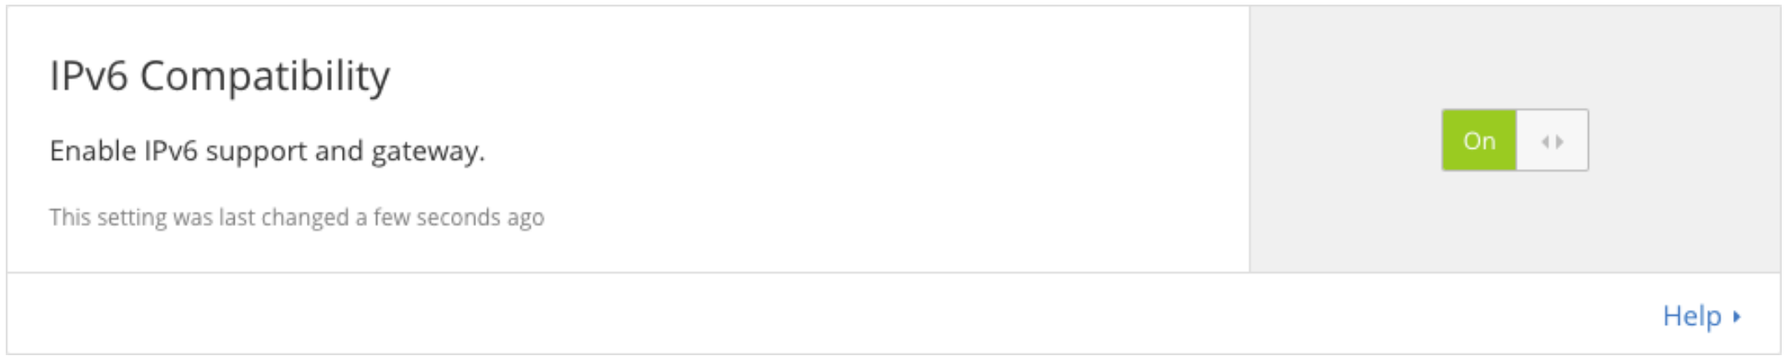

CLOUDFLARE

Cloudflare can be a “bridge” to IPv6

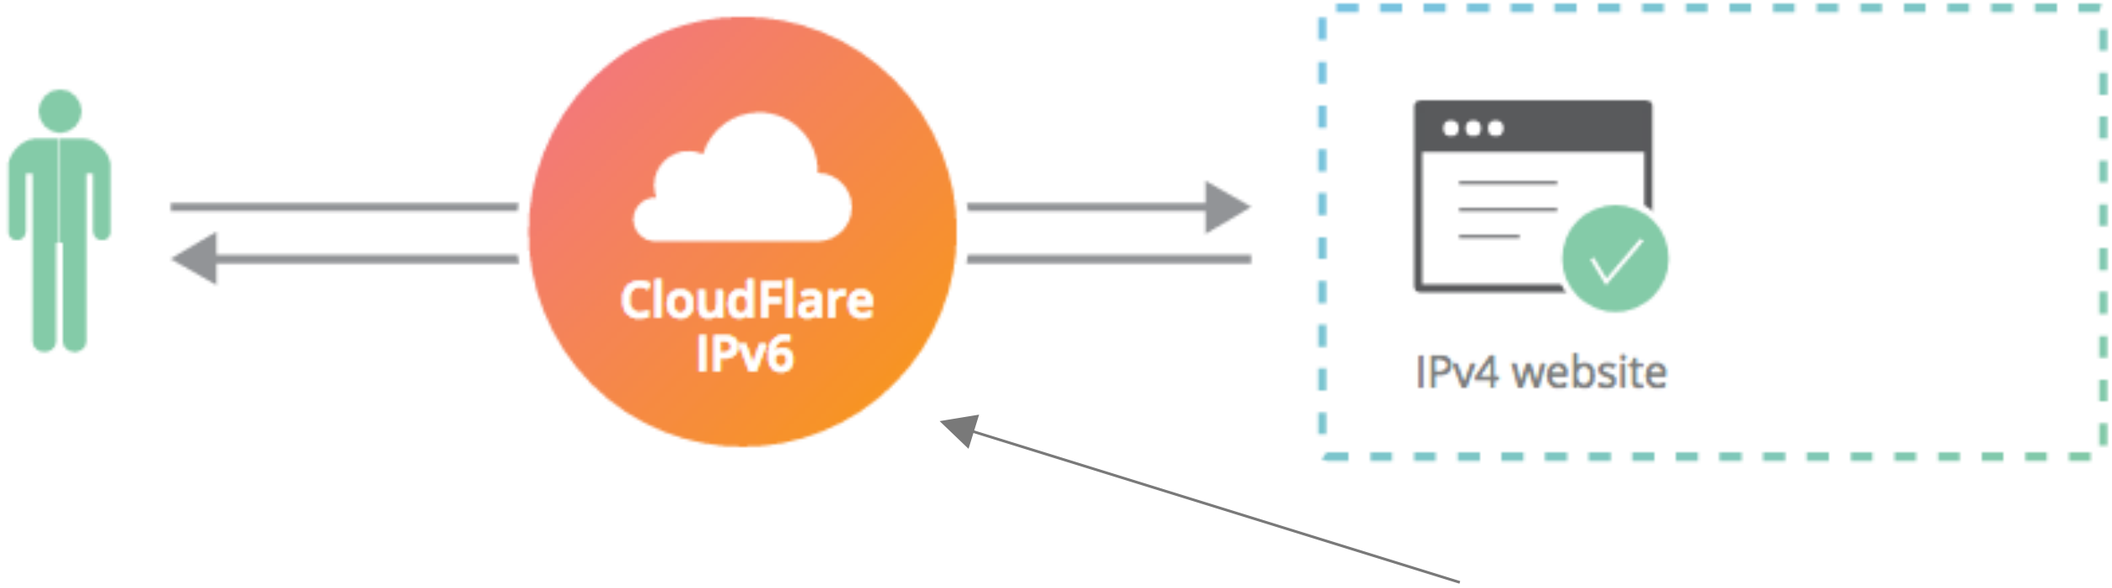

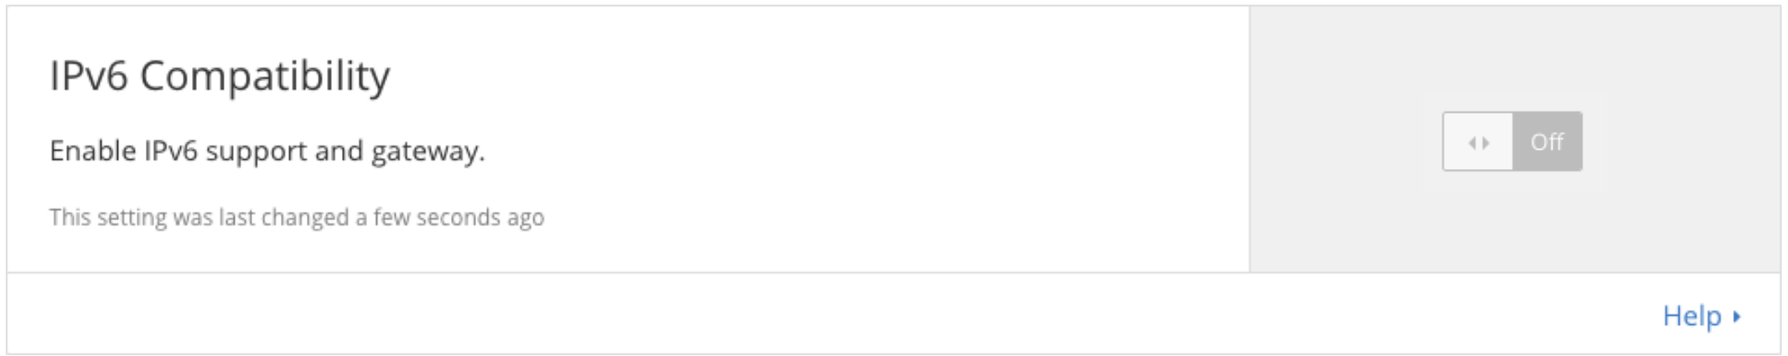

CLOUDFLARE

Cloudflare can be a “bridge” to IPv6

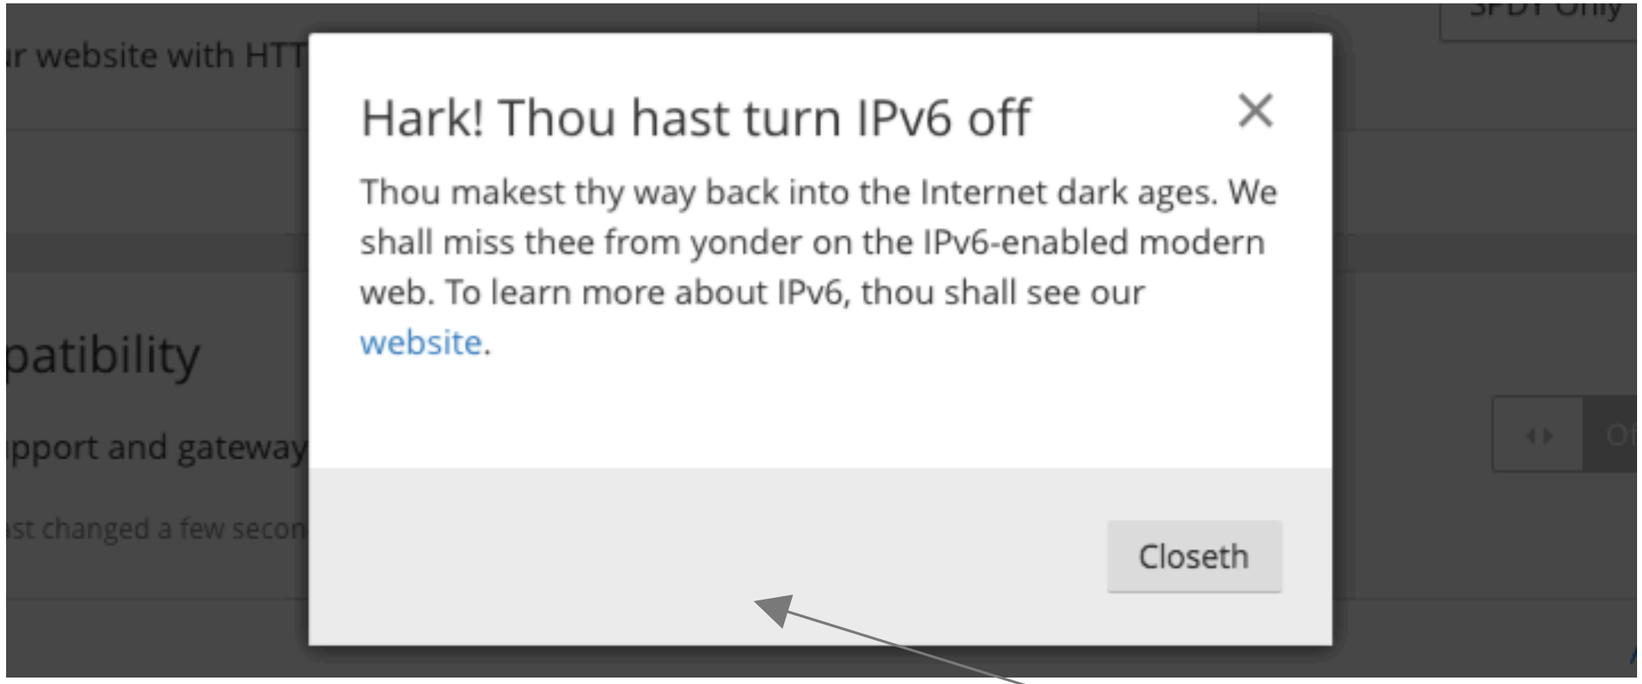

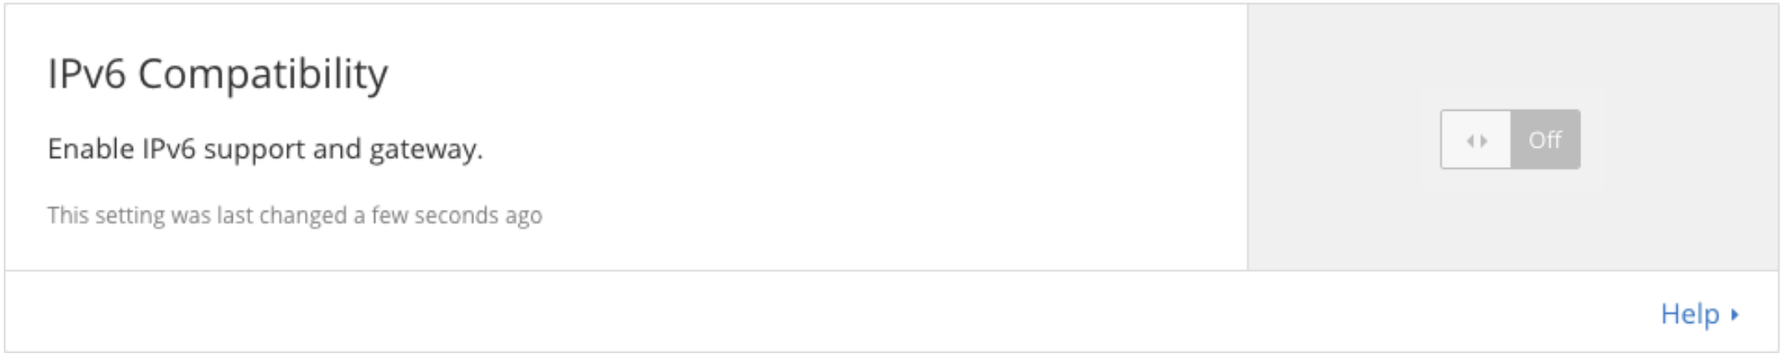

CLOUDFLARE

Five plus years of having the IPv6 switch in our system.
The default was “off”.

// Flipping the Switch!

Flipping the Switch on Every Domain/Zone

- At the time, we had nearly five million zones on Cloudflare
- If the user had never touched the IPv6 switch; then flip it on!
- Slow start; then running faster (around ~100,000 zones per day)
- A few months to finish process

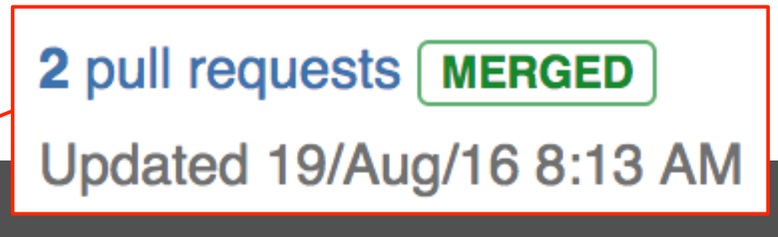


```
for zone in all_zones:
    if zone.ipv6.value == False:
        if zone.ipv6.date == None:
            zone.ipv6.value = True
            zone.ipv6.date = Now()
    sleep()
```

CLOUDFLARE

People (Some You May Know) Noticed!

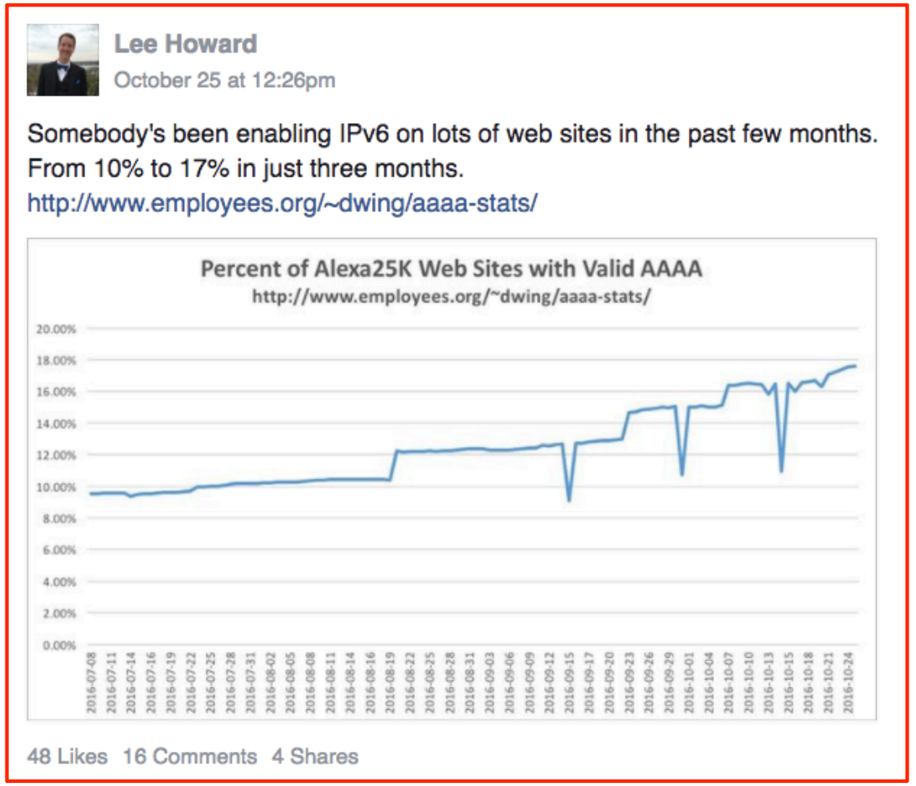

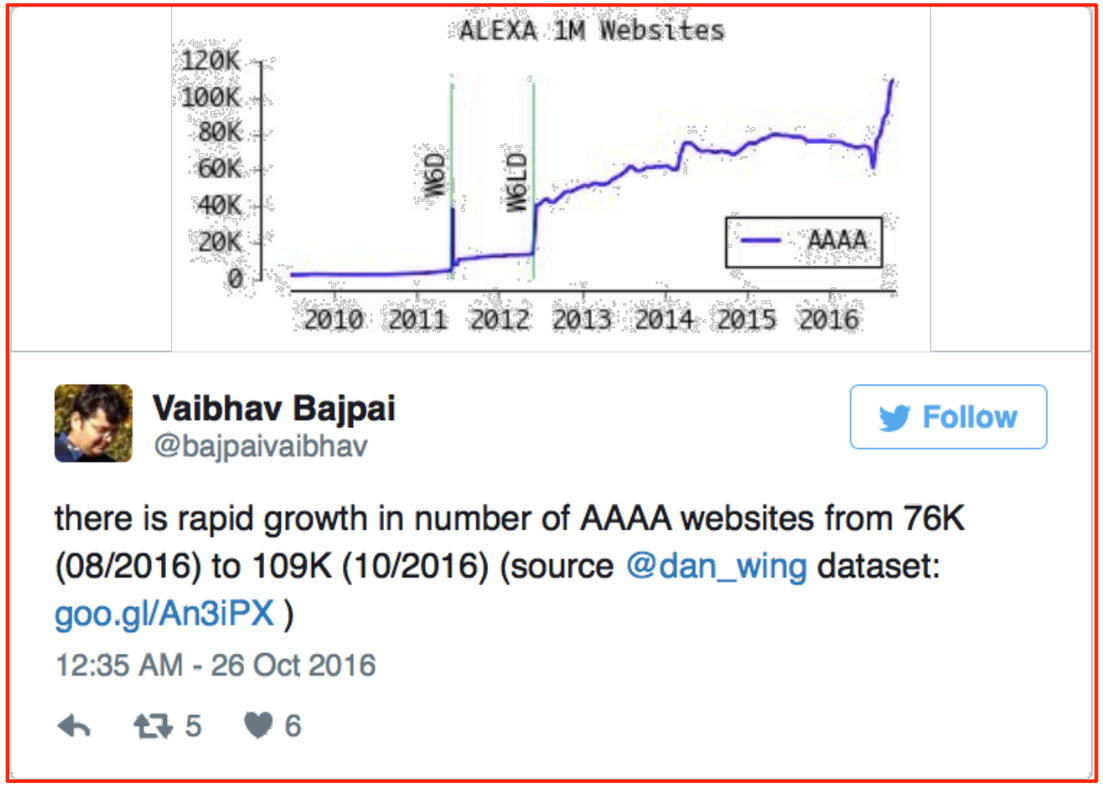

CLOUDFLARE

Eric Vyncke's graph is it's full glory!

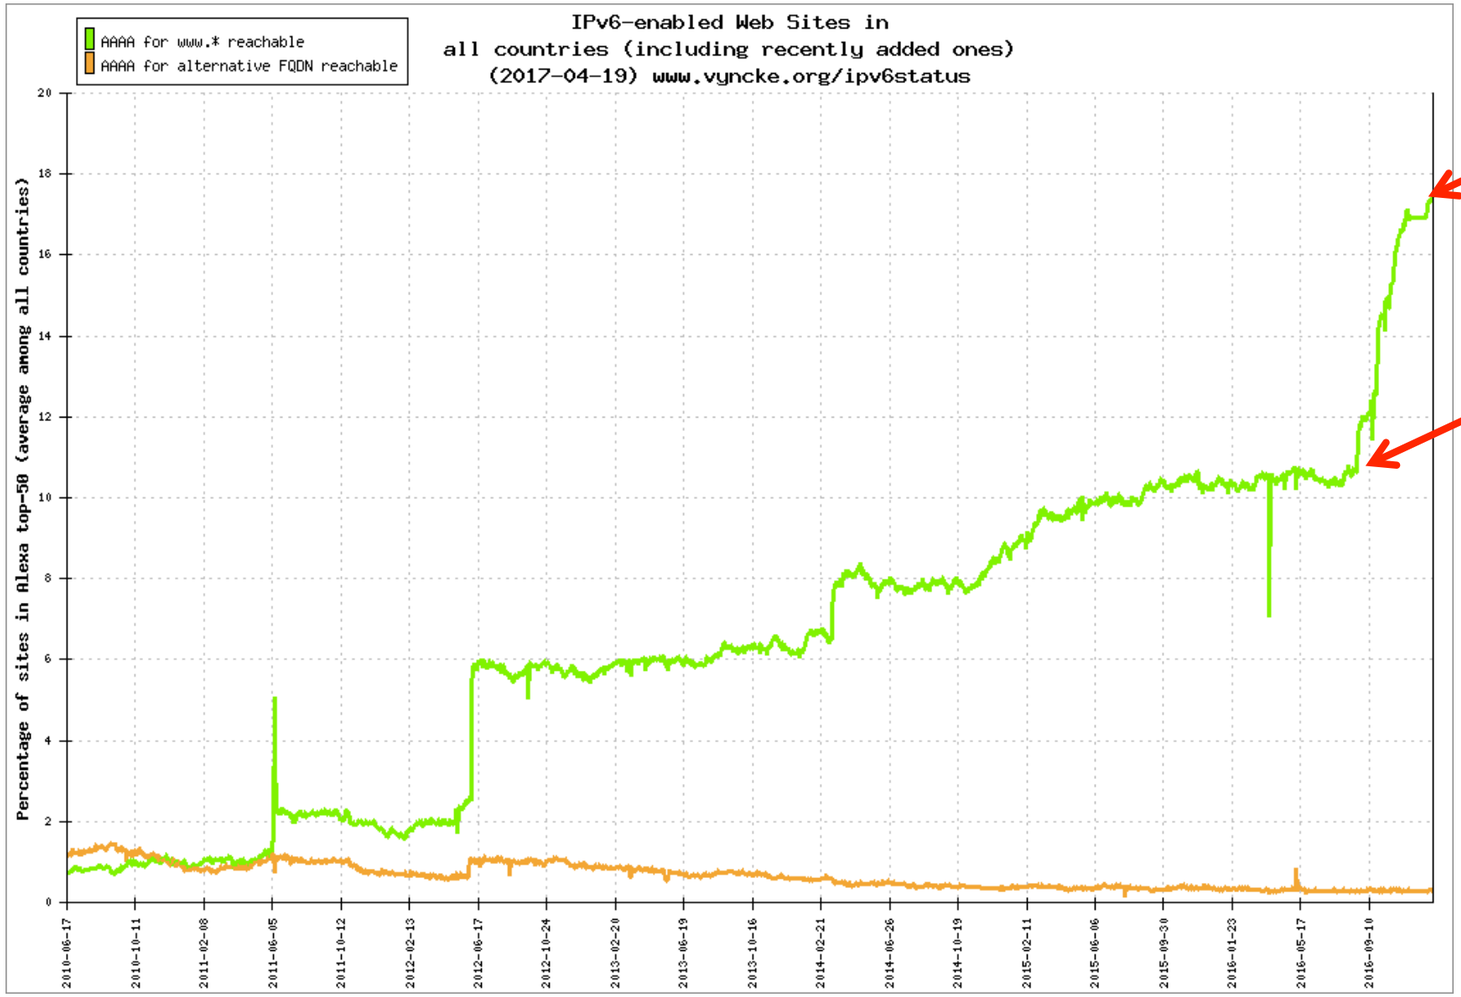

Cloudflare hits 98.01%

Cloudflare starts process

CLOUDFLARE

https://www.vyncke.org/ipv6status/
https://blog.cloudflare.com/98-percent-ipv6/

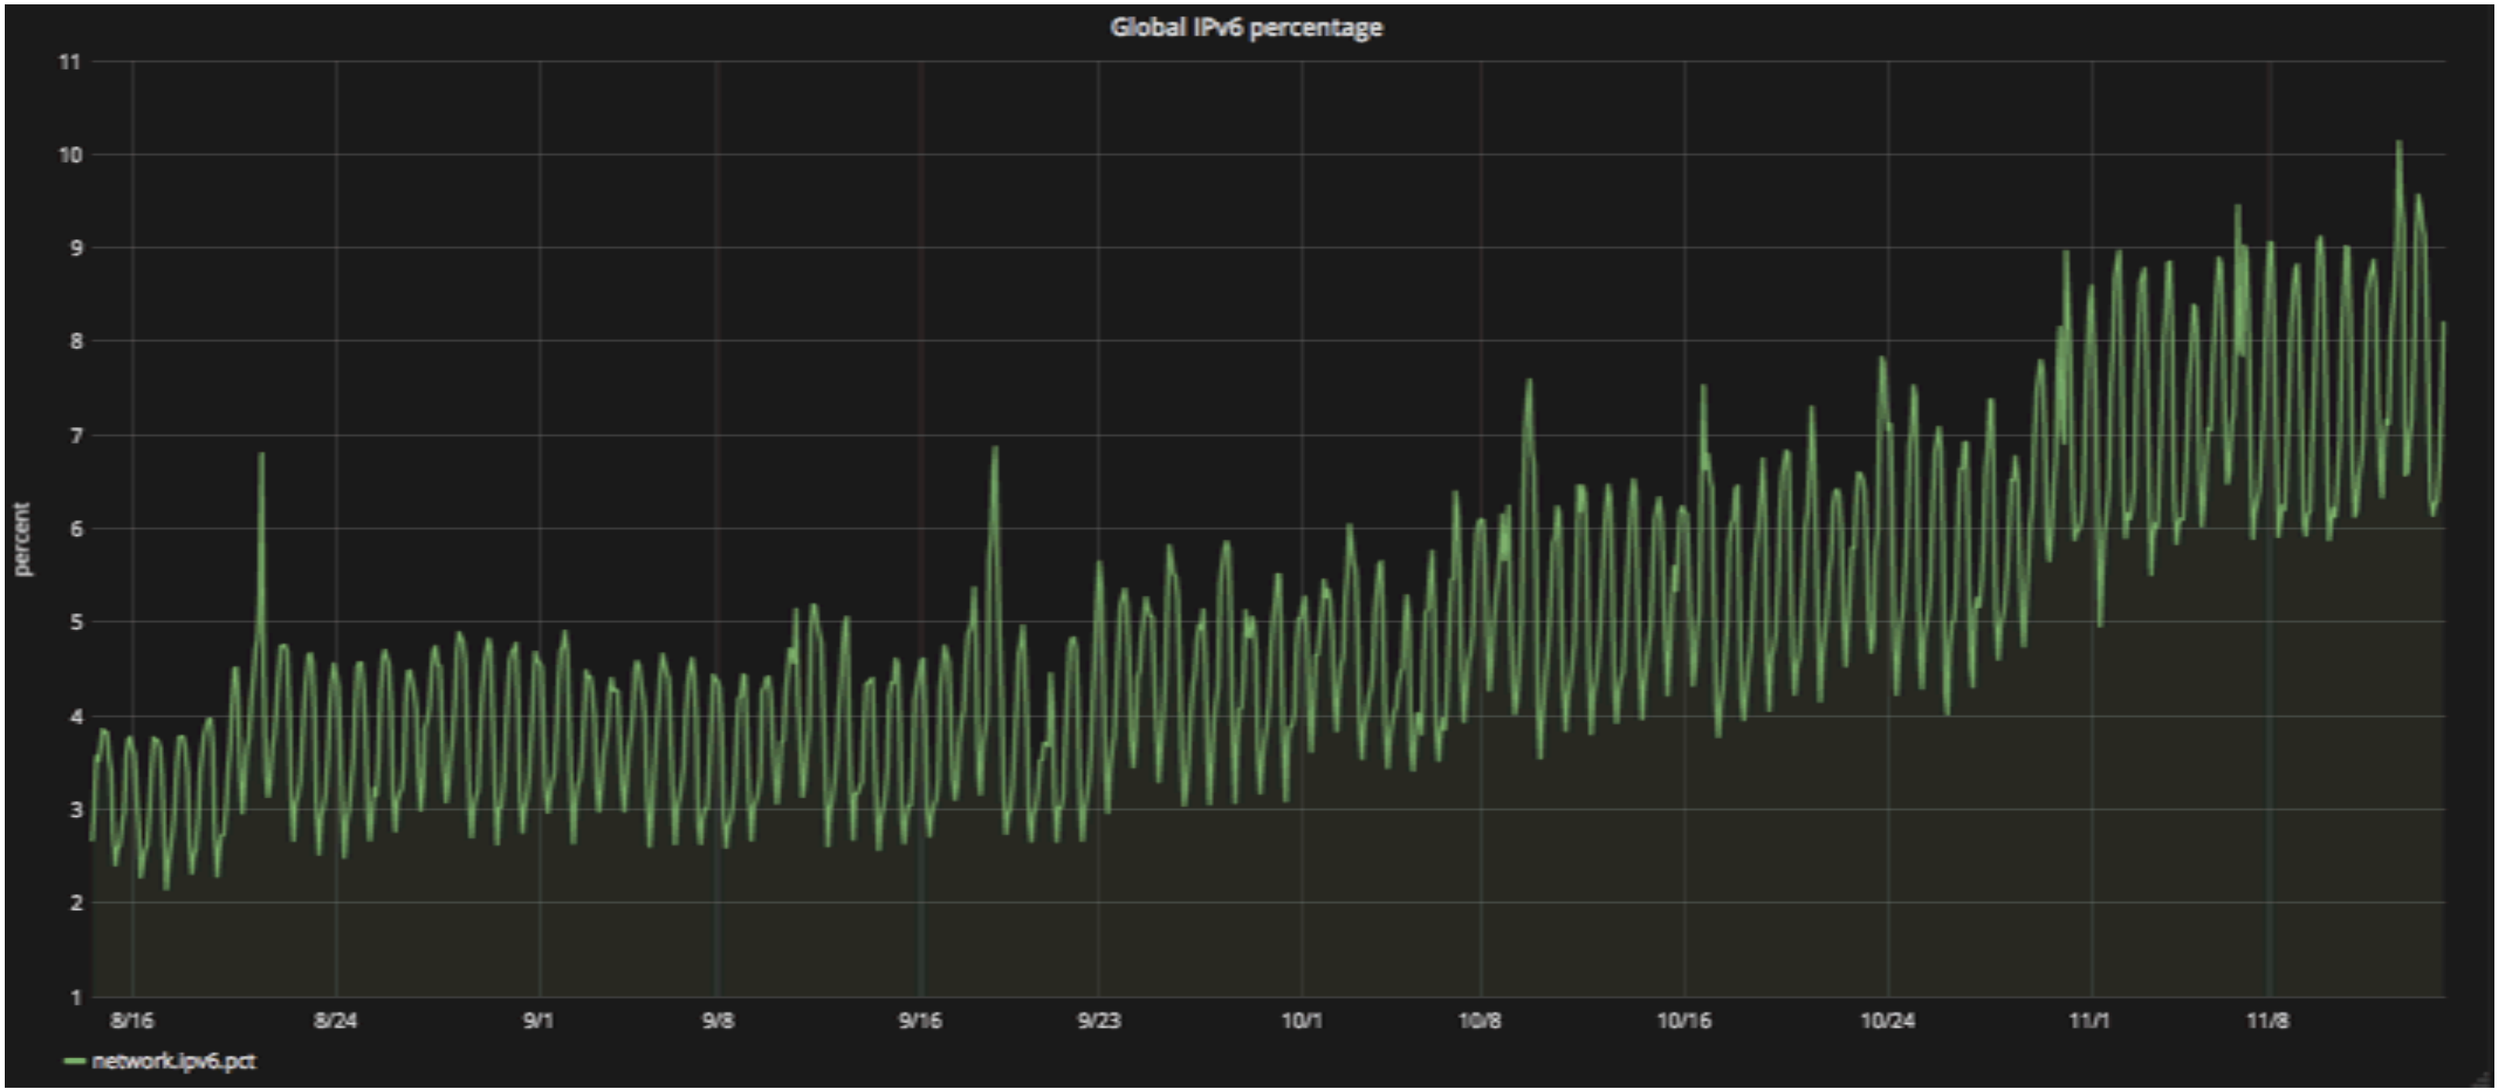
Global IPv6 percentage
11
10
9
8
7
6
5
4
3
2
1
percent
8/16
8/24
9/1
9/8
9/16
9/23
10/1
10/8
10/16
10/24
11/1
11/8
network.ipv6.pct

CLOUDFLARE

// Removing the Switch

The Disable IPv6 Switch Goes Away!

Before:

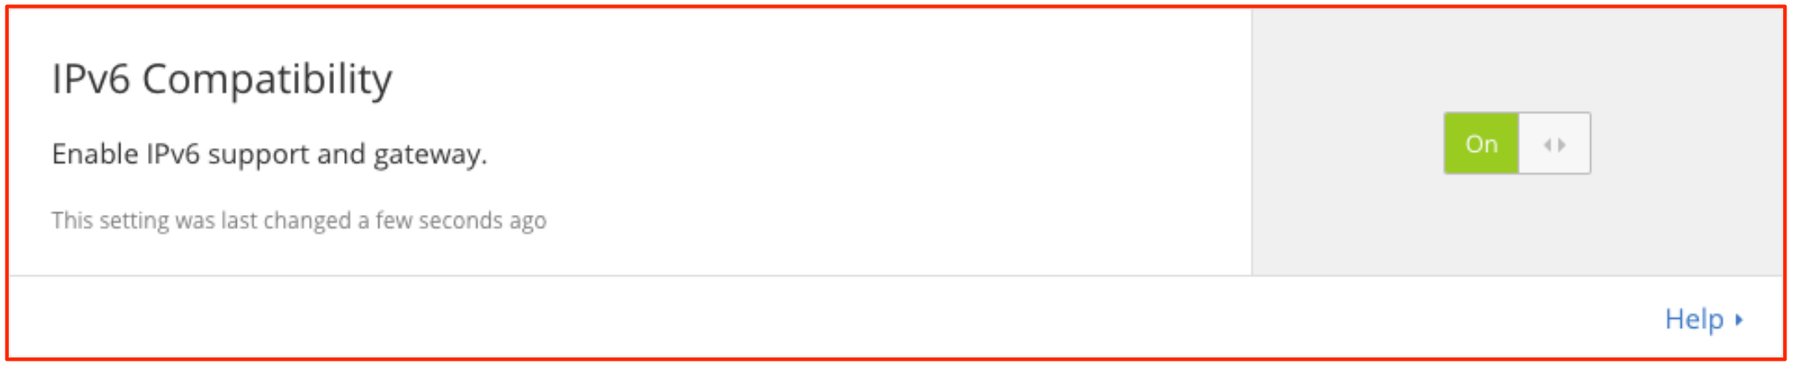

After:

… IPv6 is on by default (and unchangeable) for the vast majority** of accounts!

** *high paying accounts still get the switch*

CLOUDFLARE

// Who and What is Driving IPv6?

Percentage of IPv6 vs. Bandwidth per Network

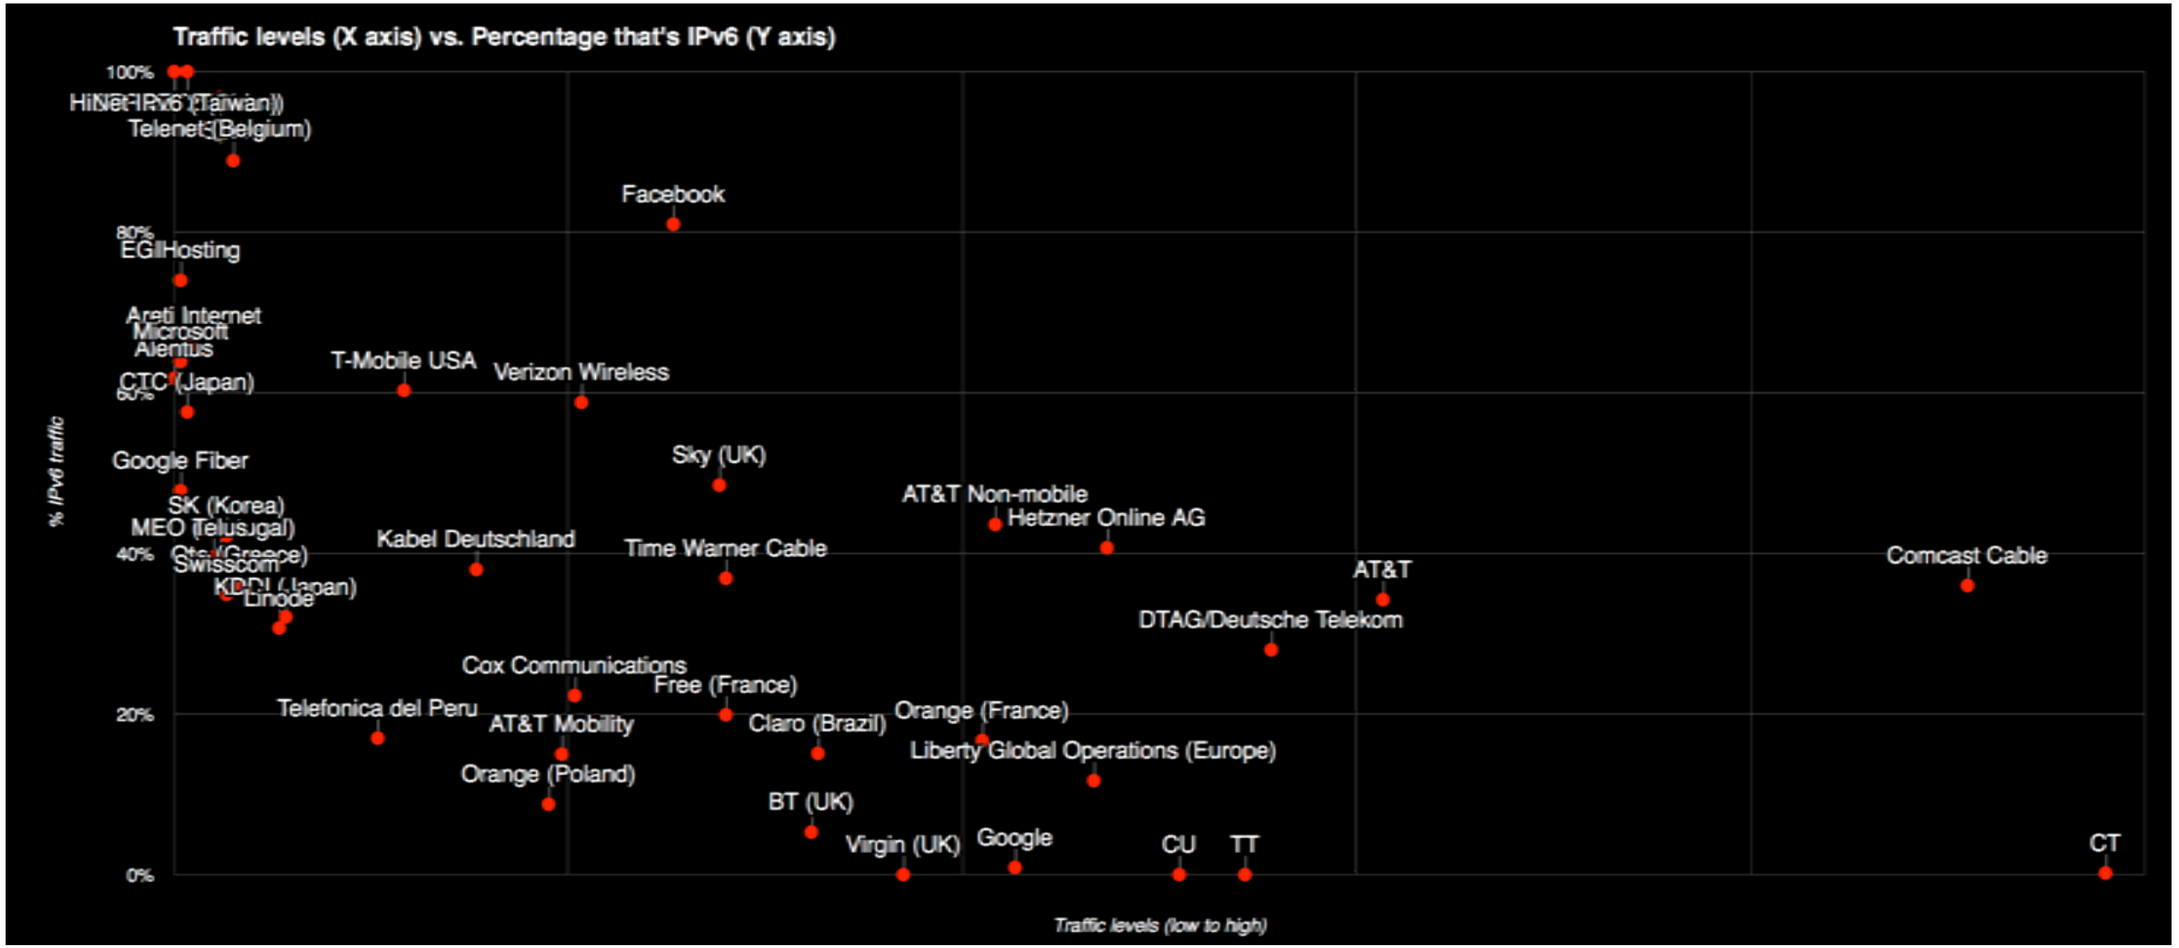

CLOUDFLARE

IPv6 by Device Type // iOS vs Android

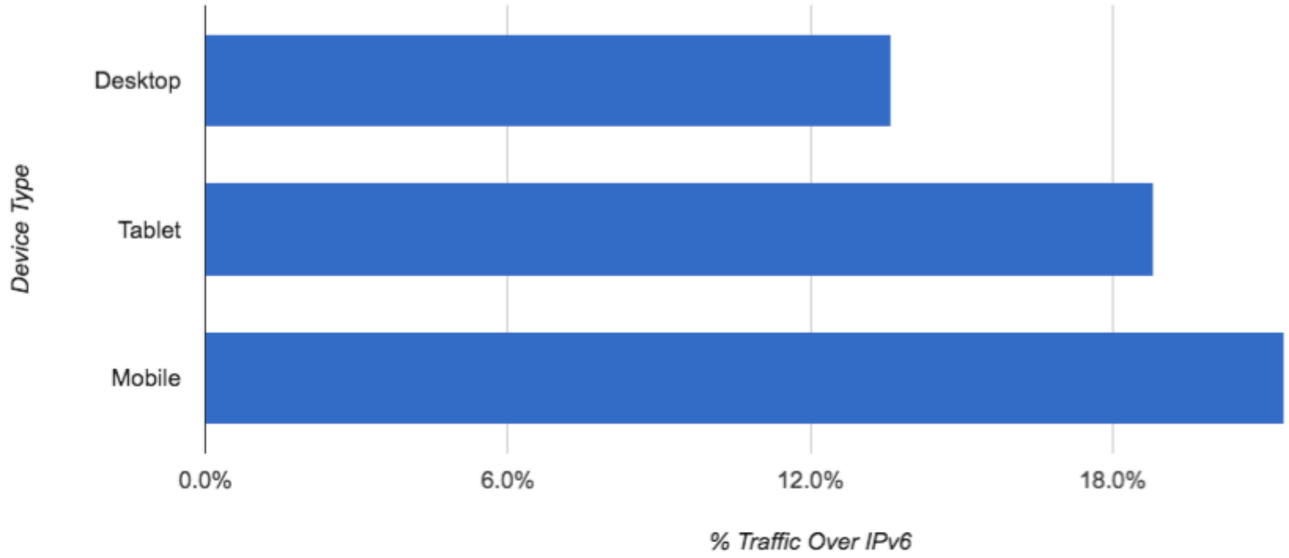

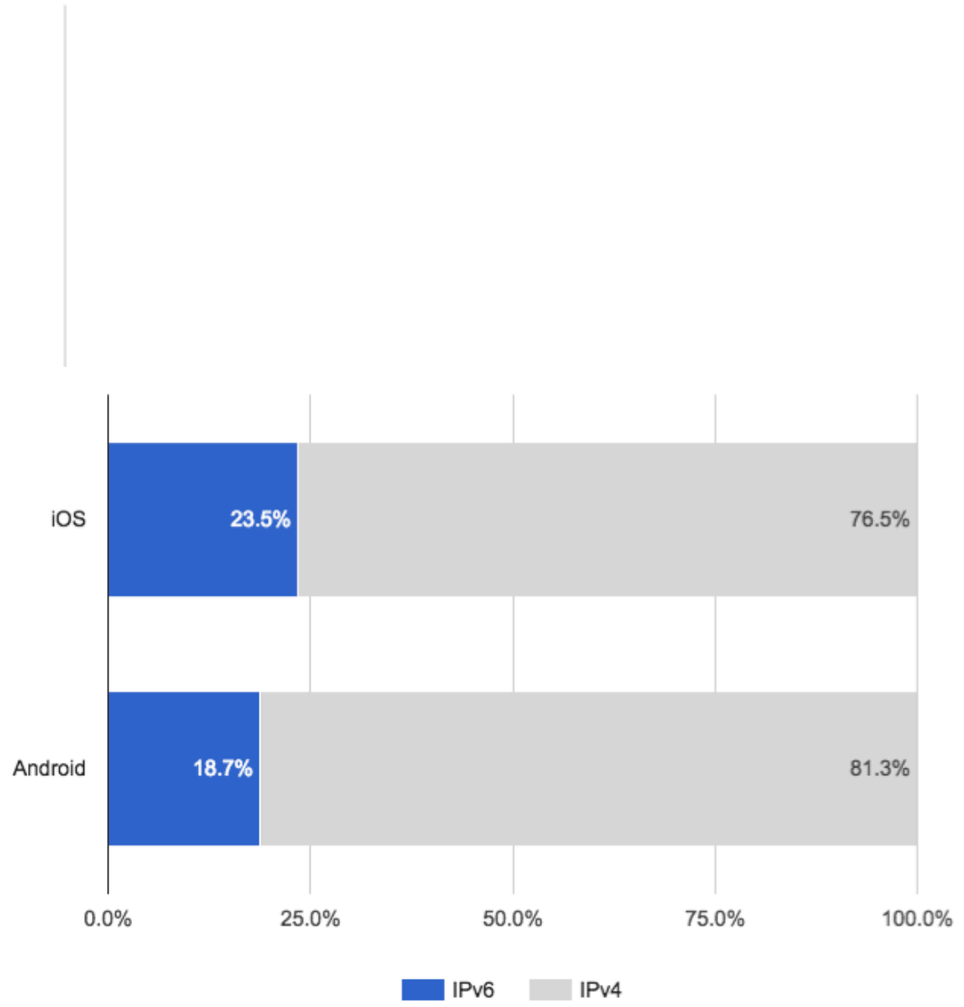

CLOUDFLARE

Windows and IPv6

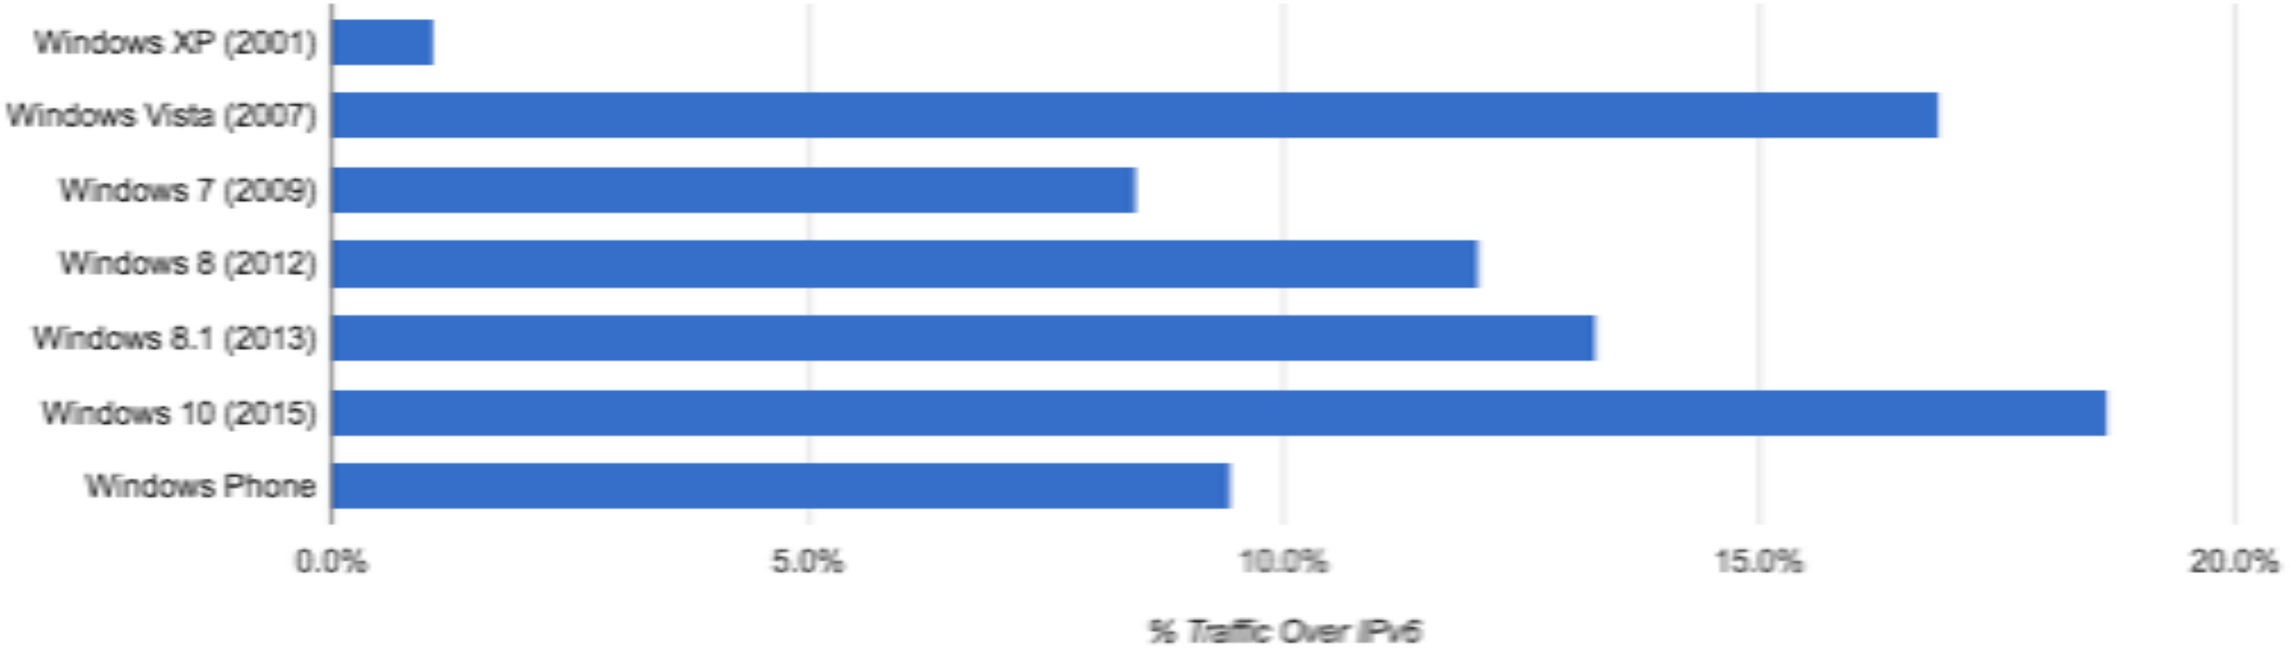

CLOUDFLARE

IPv6 Global Map (AAAA queries)

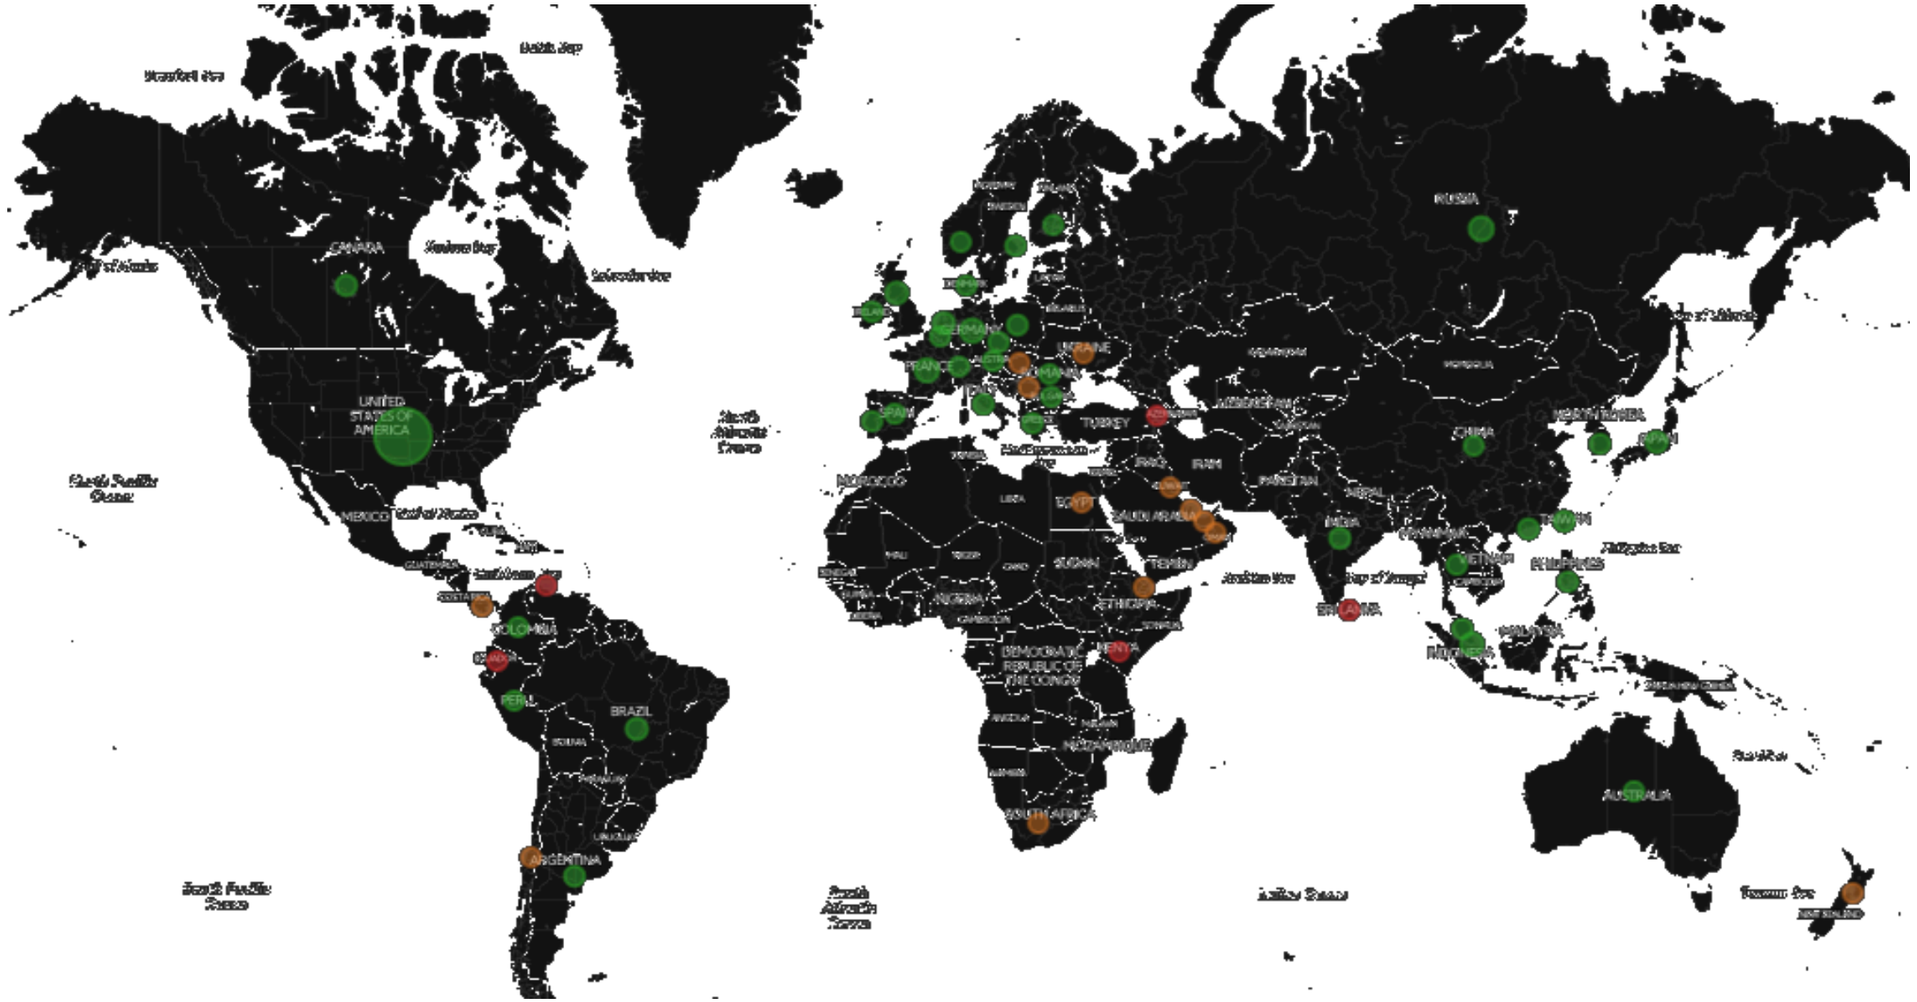

CLOUDFLARE

// Cloudflare and DNSSEC - done!

DNSSEC for every zone (at web scale!)

- Scalable // DNSSEC for entire Cloudflare customer base
- Simple // make it easy to consume

- DNSSEC shouldn’t be for power users only! It should be for everyone!
- DNS & DNSSEC software structure for this large scale deployment
- Cloudflare wrote our own DNSSEC systems (scale & speed dictated this)
- Cloudflare uses modern crypto and sign-on-the-fly at the edge

CLOUDFLARE

Solving speed (and size): ECDSA P-256

- ECDSA (Elliptic Curve Digital Signature Algorithm) P-256 signatures
  - > 3x faster than RSA1024
- Measured on OpenSSL 1.0.2 on our servers
- We (Vlad Krasnov) ported OpenSSL ASM to Go
- 21x speedup for the sign: https://go-review.googlesource.com/#/c/8968/
- Bonus: small signatures, small keys, modern crypto!
- Supported by most validators, working on registrars

CLOUDFLARE

https://tools.ietf.org/html/rfc6605

Solving speed (and size): ECDSA P-256

RSA: 1181 BYTES

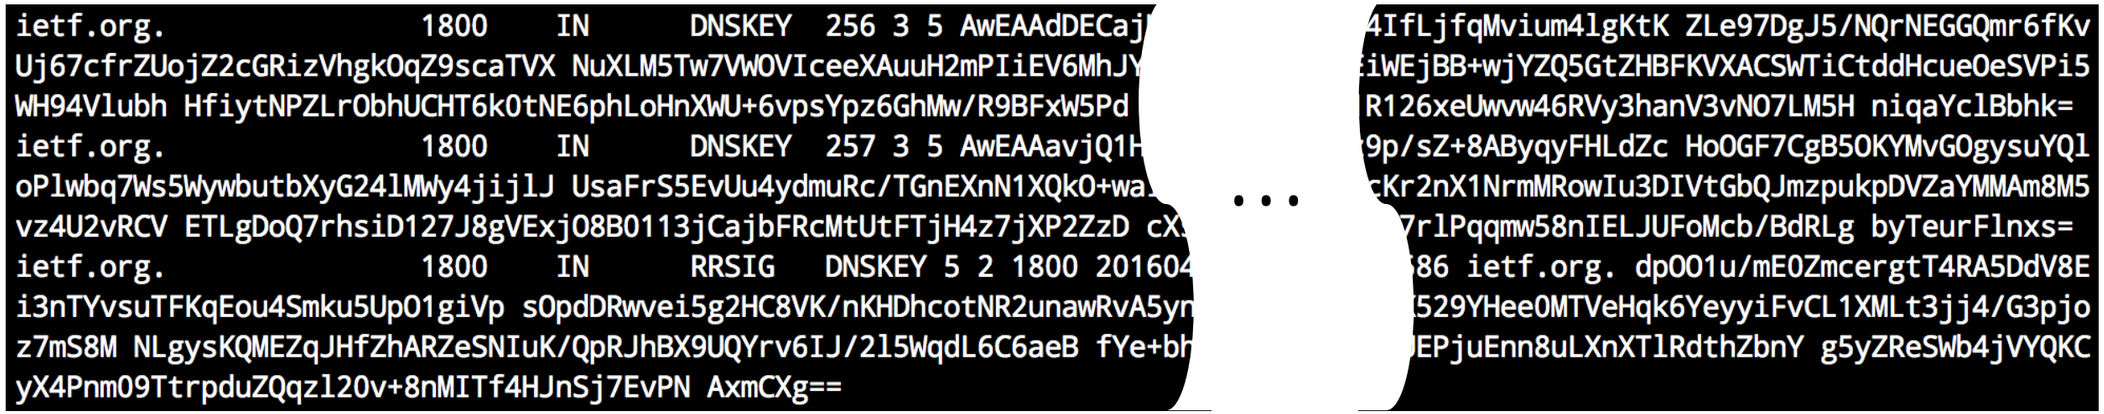

ECDSA: 305 BYTES

CLOUDFLARE

Solving speed (and size): ECDSA P-256

- Standard Go crypto:

BenchmarkSingleSignECDSA	832,295 ns/op
BenchmarkSingleSignRSA	6,003,261 ns/op

- Go with Vlad's changes:

BenchmarkSingleSignECDSA	60,806 ns/op
BenchmarkSingleSignRSA	3,124,274 ns/op

CLOUDFLARE

https://blog.cloudflare.com/go-crypto-bridging-the-performance-gap/

// TLS 1.3

The handshake (TLS 1.2 ECDHE)!

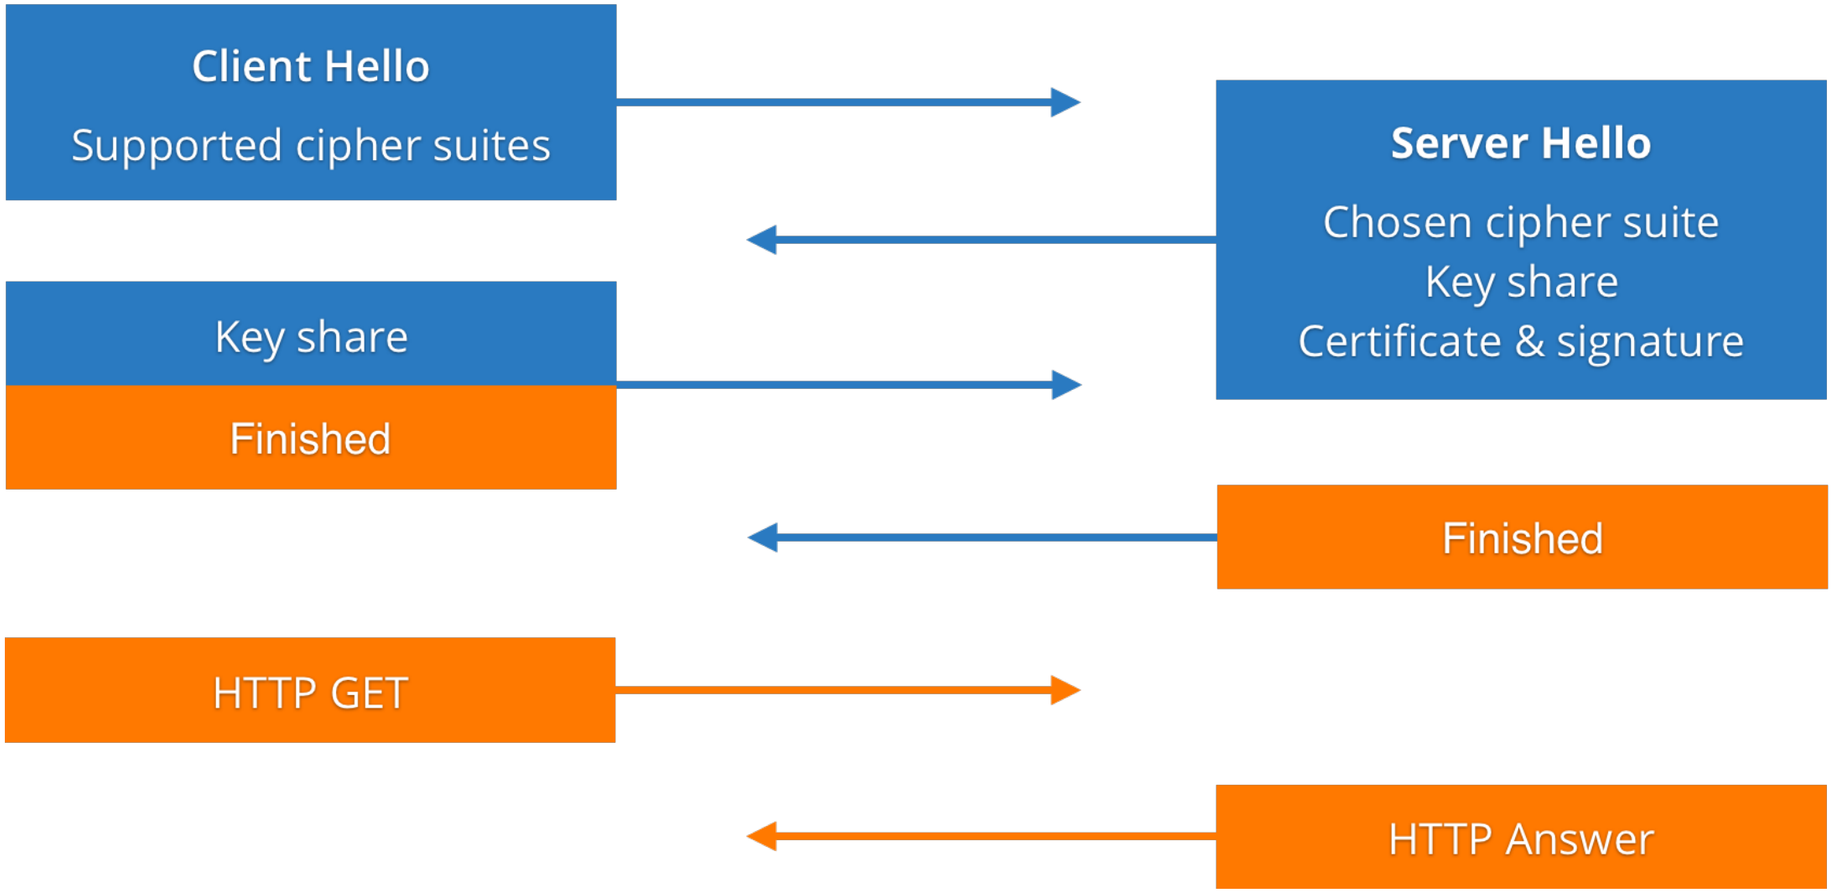

CLOUDFLARE

TLS 1.3

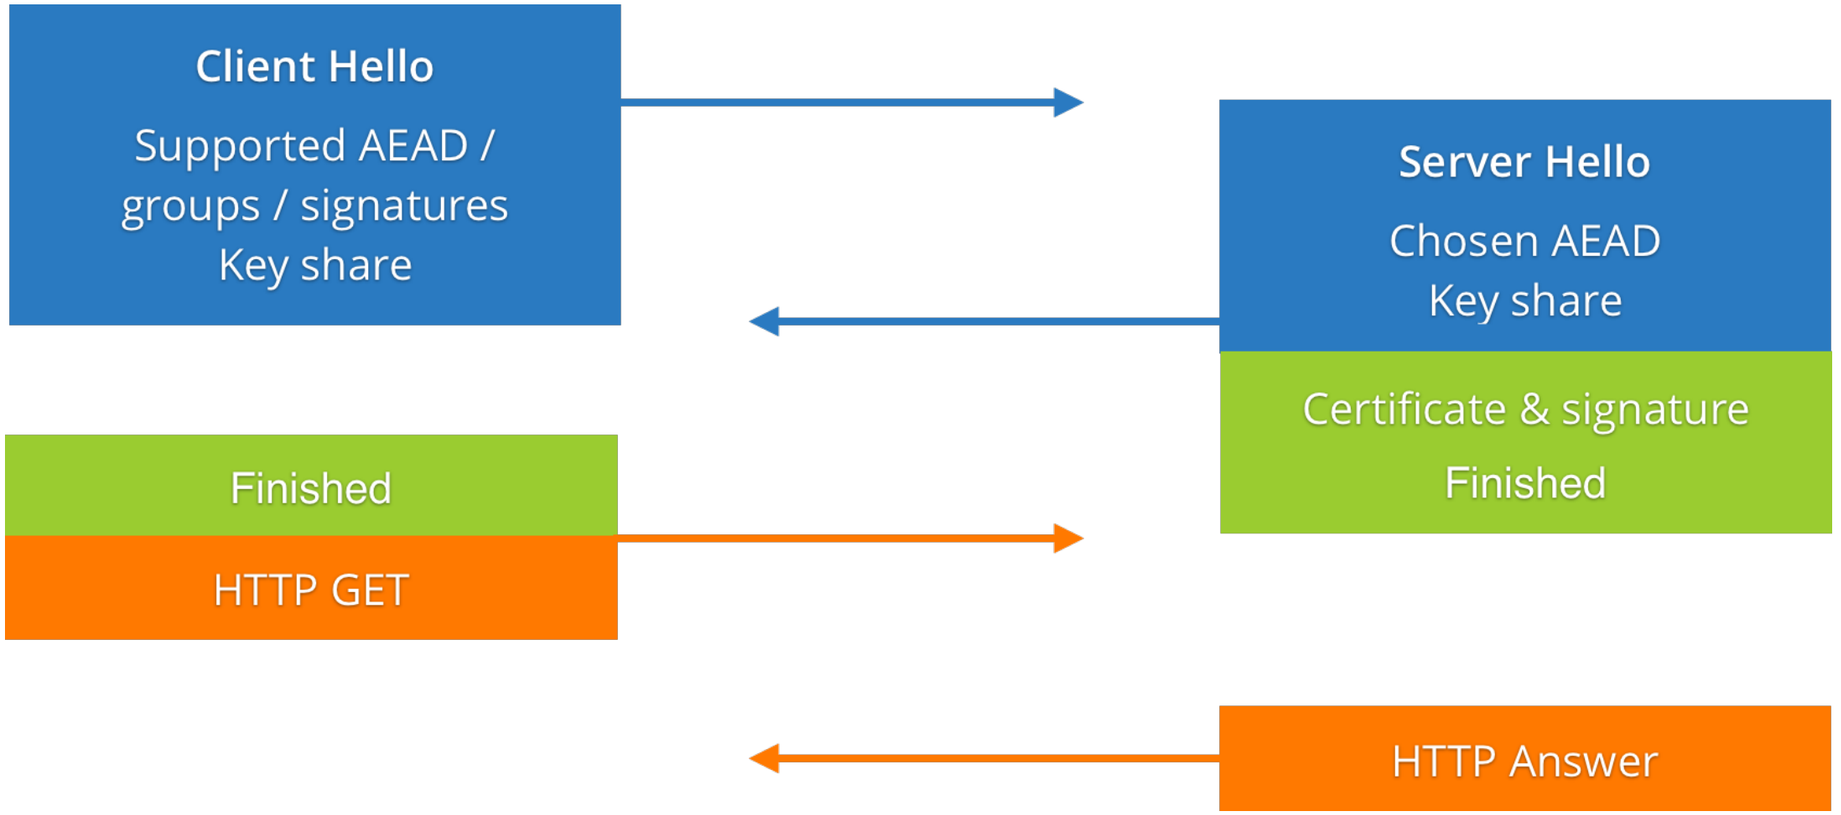

CLOUDFLARE

TLS 1.3 (0-RTT w/ ECDHE)

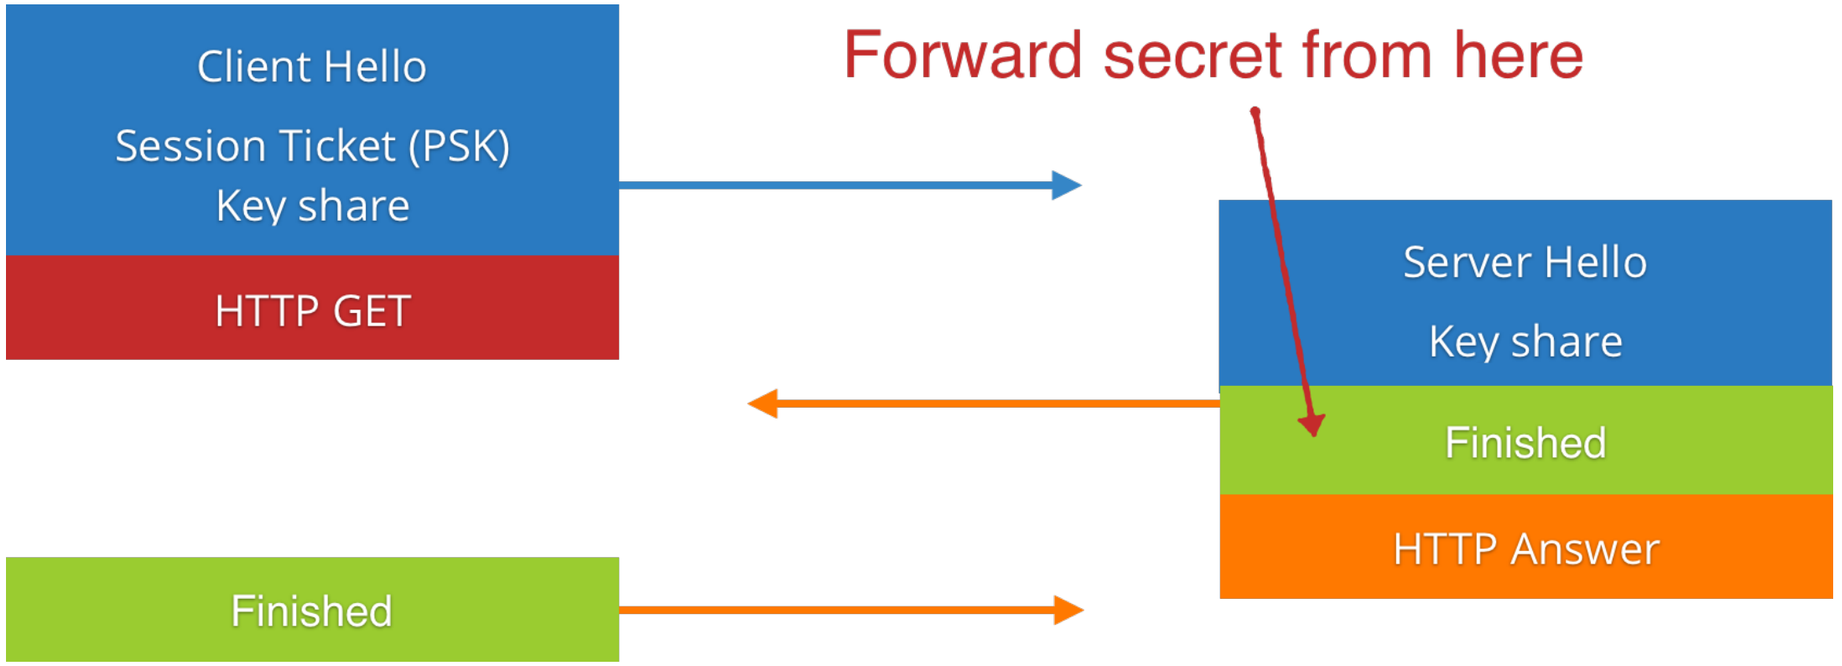

https://blog.cloudflare.com/introducing-0-rtt/

CLOUDFLARE

What's been taken out in TLS/1.3

- TLS/1.3 removed the following:
  - RC4
  - 3DES
  - AES-CBC
  - MD5 & SHA1
  - Compression
  - RSA-PKCS1-1.5
  - Renegotiation
- Sometimes it's good to clean-up!

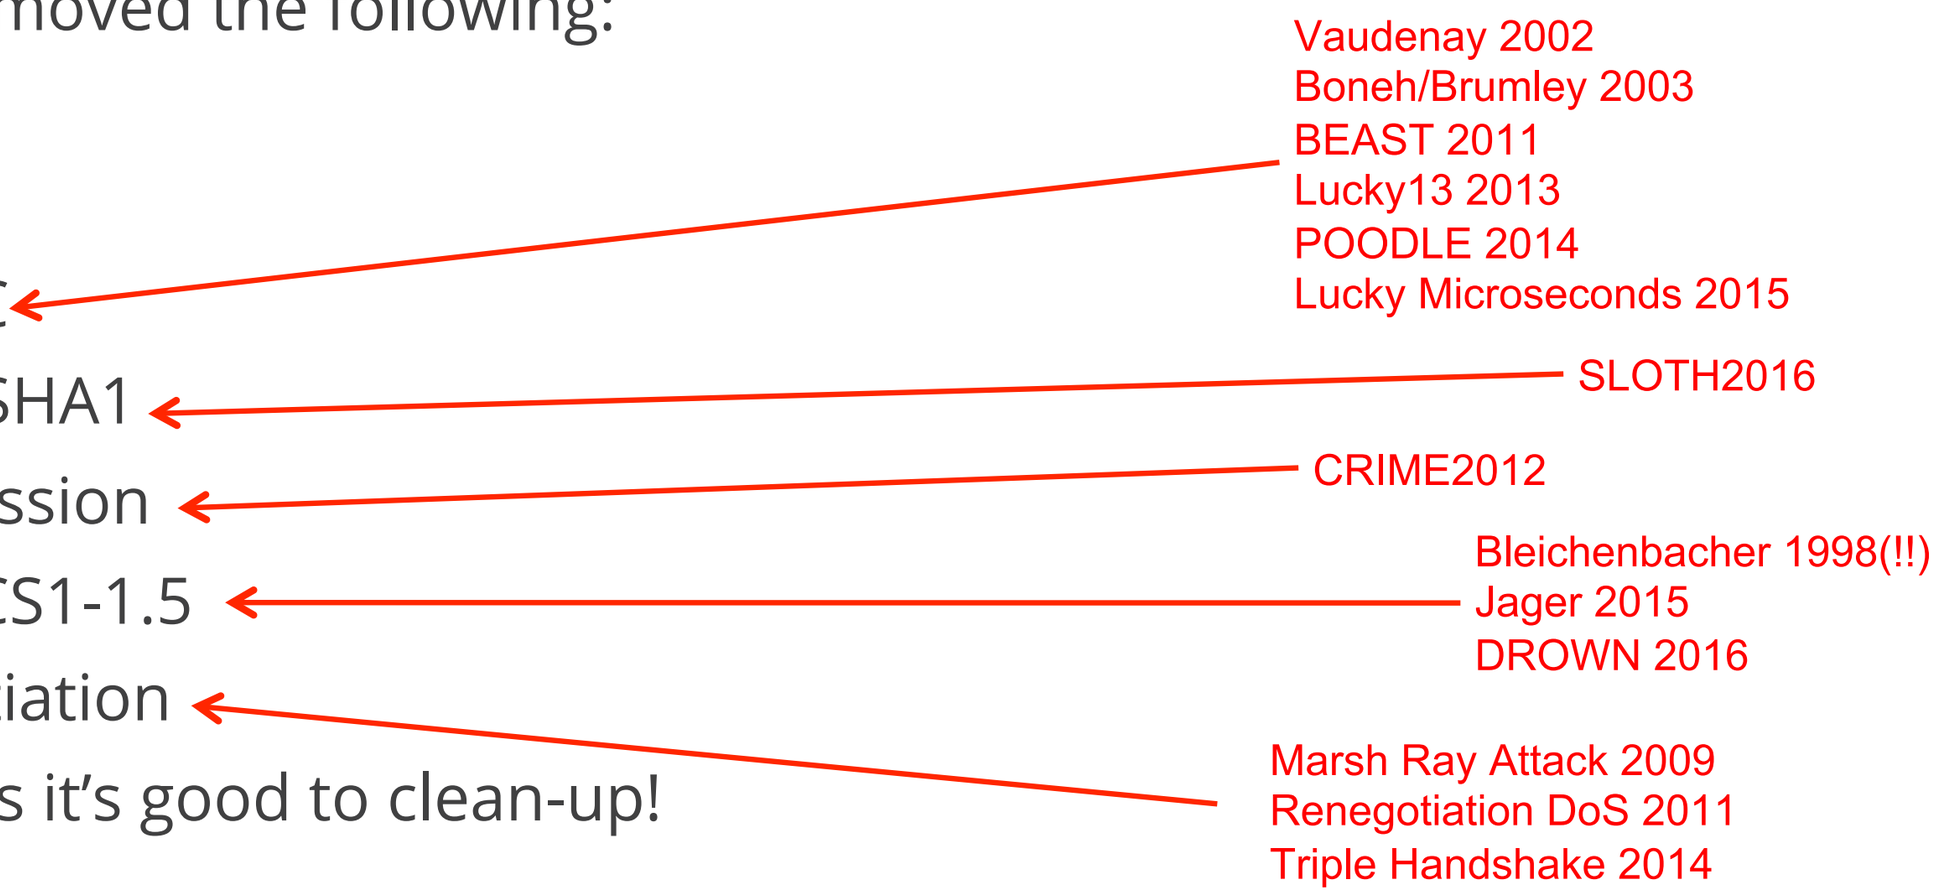

CLOUDFLARE

There’s always someone …

To: IETF TLS 1.3 Working Group Members

My name is Andrew Kennedy and I work at BITS, the technology policy division of the Financial Services Roundtable (http://www.fsroundtable.org/bits). My organization represents approximately 100 of the top 150 US-based financial services companies including banks, insurance, consumer finance, and asset management firms.
[...]
Deprecation of the RSA key exchange in TLS 1.3 will cause significant problems for financial institutions, almost all of whom are running TLS internally and have significant, security-critical investments in out-of-band TLS decryption.

Hi Andrew,

My view concerning your request: no.

Rationale: We're trying to build a more secure internet.

Meta-level comment:
You're a bit late to the party. We're metaphorically speaking at the stage of emptying the ash trays and hunting for the not quite empty beer cans.
More exactly, we are at draft 15 and RSA key transport disappeared from the spec about a dozen drafts ago. I know the banking industry is usually a bit slow off the mark, but this takes the biscuit.
Cheers,
Kenny

CLOUDFLARE

// What’s next for IPv6? Fix DNS!

A & AAAA Records - How Silly is this in 2017?

- Separate A & AAAA records
- In a happy-eyeball environment we still need two DNS queries (before any TCP connection can be instigated)

Query for A record

Header	QR AA RCODE=NOERROR
Question	www.example.com IN A
Answer	www.example.com. IN A 192.0.2.1
Authority	<empty>
Additional	<empty>

Query for AAAA record

Header	QR AA RCODE=NOERROR
Question	www.example.com IN AAAA
Answer	www.example.com. IN AAAA 2001:db8::1
Authority	<empty>
Additional	<empty>

CLOUDFLARE

AAAA For Free (When Doing an A Query)!

Cloudflare proposed solution:

1. A + AAAA in new meta-query
2. Resolver asks for A or AAAA
3. If positive answer, the resolver then checks AAAA + A meta-query
4. Resolver remembers whether authoritative server supports meta-query for future queries
5. Resolver adds both A and AAAA to cache

Working code (an IETF must!)

```
$ dig cloudflare.com @ns1.cloudflare.com -t TYPE65535 +short
198.41.215.162
198.41.214.162
2400:cb00:2048:1::c629:d6a2
2400:cb00:2048:1::c629:d7a2
$
```

This is live - try it with any domain on Cloudflare.

```
$ dig taylorswift.com @ashley.ns.cloudflare.com -t TYPE65535 +short
104.16.193.61
104.16.194.61
104.16.191.61
104.16.192.61
104.16.195.61
2400:cb00:2048:1::6810:c33d
2400:cb00:2048:1::6810:c13d
2400:cb00:2048:1::6810:bf3d
2400:cb00:2048:1::6810:c23d
2400:cb00:2048:1::6810:c03d
$
```

CLOUDFLARE

IETF draft – pick one, any one (maybe ours?)

https://tools.ietf.org/html/draft-vavrusa-dnsop-aaaa-for-free-00
https://tools.ietf.org/html/draft-yao-dnsop-accompanying-questions-02
https://tools.ietf.org/html/draft-bellis-dnsext-multi-qtypes-03

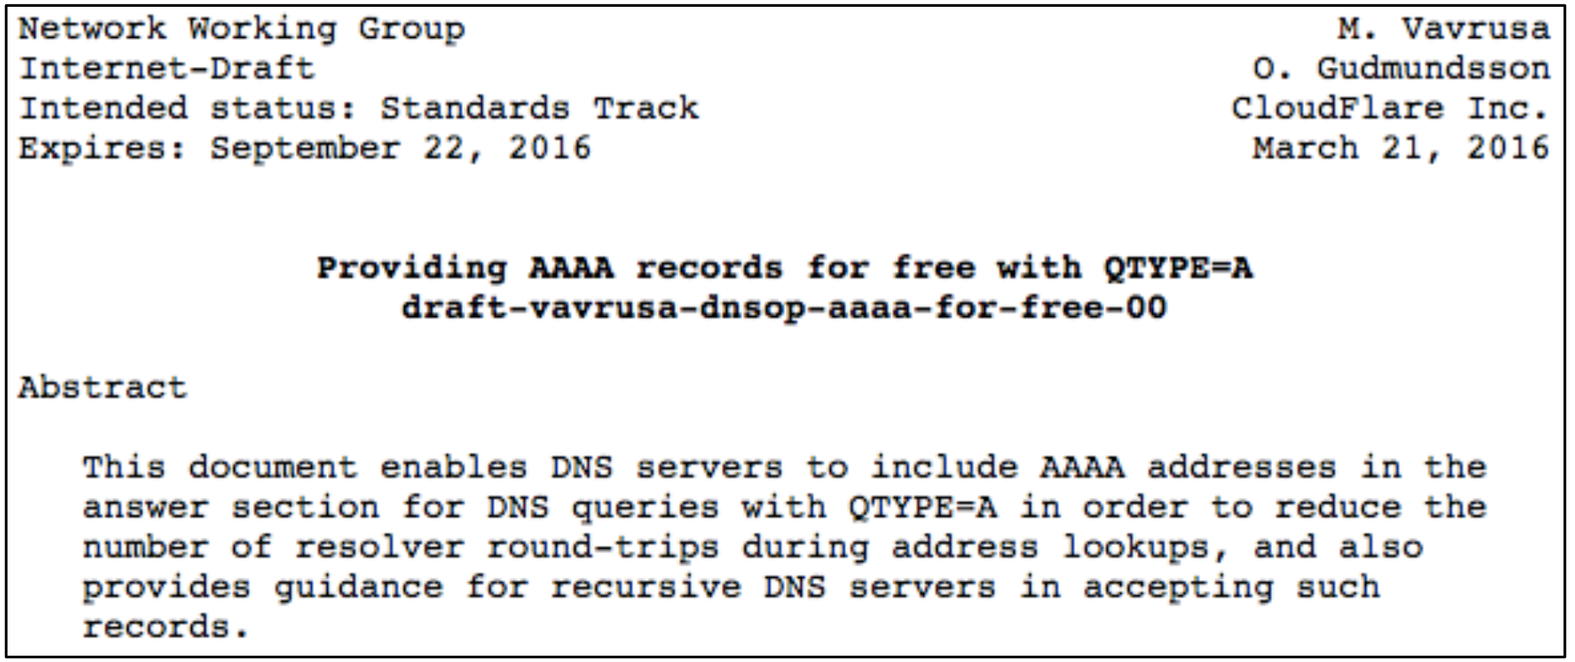

Network Working Group — M. Vavrusa
Internet-Draft — O. Gudmundsson
Intended status: Standards Track — CloudFlare Inc.
Expires: September 22, 2016 — March 21, 2016

Providing AAAA records for free with QTYPE=A
draft-vavrusa-dnsop-aaaa-for-free-00

Abstract

This document enables DNS servers to include AAAA addresses in the answer section for DNS queries with QTYPE=A in order to reduce the number of resolver round-trips during address lookups, and also provides guidance for recursive DNS servers in accepting such records.

https://tools.ietf.org/html/draft-vavrusa-dnsop-aaaa-for-free-00

CLOUDFLARE

Thank you!

martin@cloudflare.com

@mahtin / @cloudflare