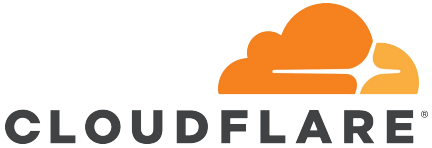

DoH and DoT experience

Ólafur Guðmundsson

Marek Vavrusa

Announced April 1st 2018

Our mission: to help build a better Internet.

We use 1.1.1.1 and 1.0.0.1 (easy to remember) for our resolver.

Provided to Cloudflare by APNIC for both joint research and this service.

We focused on privacy!

We knew we would spend a lot of time cleaning up the global Internet to make 1.1.1.1 work!

CLOUDFLARE

1.1.1.1

DNS resolver, 1.1.1.1, is served by Cloudflare's Global Anycast Network.

The Cloudflare network (DNS, DDoS, CDN, WAF, more)

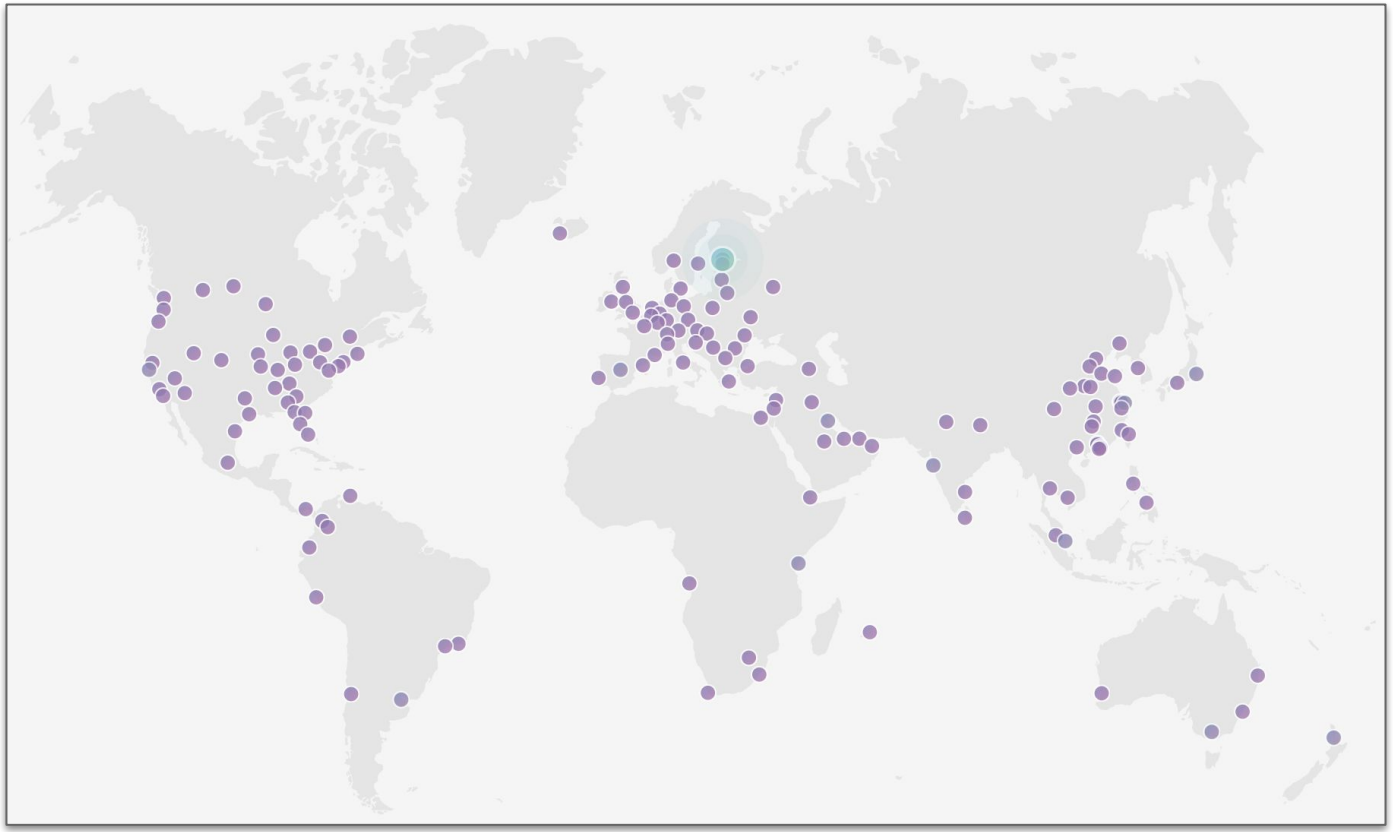

151+

Data centers globally

151+

DNS resolver locations

151+

DNS authoritative locations

1.1.1.1 design goals

DNS and privacy!

DNS itself is a 35-year-old protocol (and it's showing its age). It was never designed with privacy or security in mind.

DNS inherently is unencrypted so it leaks data to anyone who's monitoring your network connection.

We focused on privacy:

- **Query Minimization RFC7816**
- **Aggressive negative answers RFC8198**
- **No Client Subnet on queries**

- **DNS-over-TLS (Transport Layer Security) RFC7858**
- **DNS-over-HTTPS protocol DoH (draft-ietf-doh-dns-over-https)**

CLOUDFLARE

1.1.1.1

In 2014, we decided to enable https encryption for free for all our customers (we doubled the size of the encrypted web).

In 2017, we made DDoS mitigation free & unmetered across all our plans.

Data Policy

- **We don’t store client IP addresses never, ever!**
- **We only use query logs for things that improve DNS resolver performance.**
- **After obfuscation, APNIC research gets access to data (under our joint agreement).**

- **Cloudflare never stores any information in logs that identifies end user.**
    - **All log records are deleted within 24 hours.**
- **We will continue to abide by our privacy policy and ensure that no user data is sold to advertisers or used to target consumers.**

CLOUDFLARE

1.1.1.1

All log records deleted within 24 hours

Aggregations is only on traffic stats based of AS#

1.1.1.1 DoT and DoH implementations

DoT: DNS over TLS

Open Port on Firewalls
Tell DDoS systems about it

Knot Resolver is the engine behind 1.1.1.1

Provides DNS over TLS by default

- **Uses GnuTLS for diversity**
- **Latest version only supports TLS 1.3 draft 28**
  - **Not compatible with OpenSSL GA**
- **“Long-lived” connections supported (tens of seconds)**
- **Session resumption supported**

CLOUDFLARE

DoH: DNS over HTTPS

Marek add more

Added as a Lua module to the resolver, fronted by the NGINX

- **Terminates HTTPS and forwards to resolver over local socket**
- **Clears PII from standard logs due to privacy policy**
- **Added JSON format support (compatible with Google Public DNS)**

CLOUDFLARE

Challenges

Tools supporting DNS over TLS

- **kdig supports DNS over TLS (using GnuTLS)**
- **getdns (Stubby)**
- **Unbound**
- **Android P**
- **Tenta (browser)**
- **...**
- **https://dnsprivacy.org/wiki/display/DP/DNS+Privacy+Clients**

Bugs discovered
Bugs fixed fast
Tools added support fast

Support for TLS 1.3 is lagging

CLOUDFLARE

Ripe Atlas: Epic Fail

Does not support modern crypto (ECC):

- old crypto library

ECC is faster
ECC keys are smaller

“Tools” are hostages of crypto libraries

Bind has crypto library flexibility

CLOUDFLARE

DoH: Support

Browsers: Firefox, Chrome

Tools: Curl

Phones: Android P

Services: Cloudflare, Google , PowerDNS ....

IETF Interop London: Success

8+ different implementations in one room

-- issues raised and resolved

-- ID’s reflected experience the day after

Being in the same room and having access to all participants is a time safer

Protocols become better

Tools get better

CLOUDFLARE

Usage

No performance issues: DoH

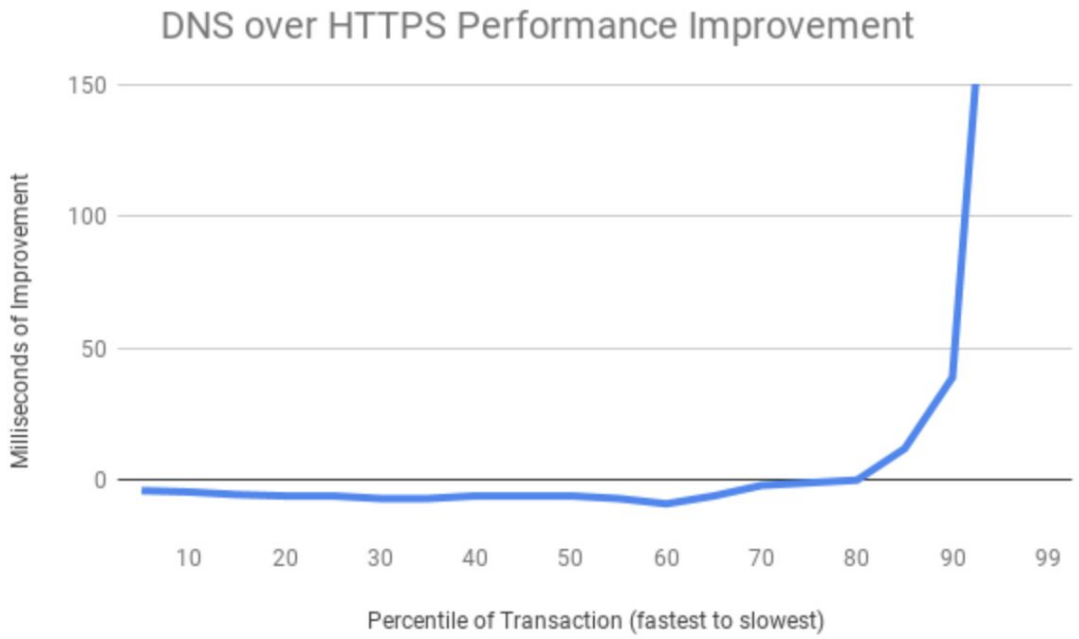

Proximity matters

Cache sharing matters more

Burts mask setup overhead

CLOUDFLARE

Better performance: DoT

1. TCP is better than UDP
   a. dealing with retransmissions
   b. Buffer size
   c. Less middlebox interference
2. TLS overhead get amortized over many queries
3. TLS session resumption lowers cost
4. Great for busy recursor to Authority

Unbound experiment max answers over UDP and TCP full cache

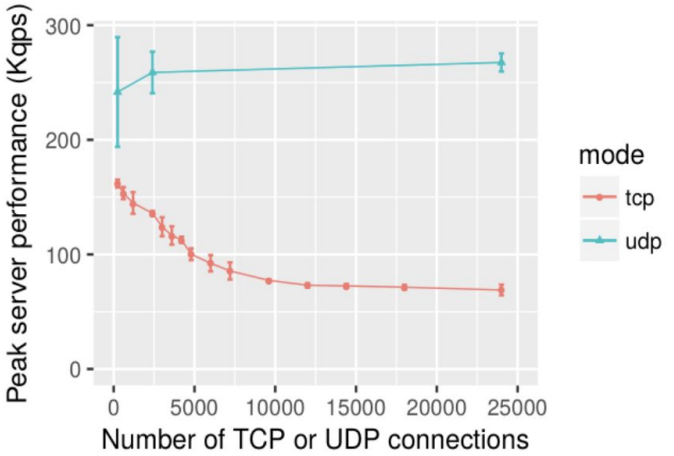

CLOUDFLARE

Connection reuse

We do not collect this info

Artifact of privacy policy

Experimenting with DoT to few Authorities

CLOUDFLARE

Q/A