Akamai

Experience the Edge

DNS HTTP RR

a bright new future?

Ralf Weber

What is the HTTPS RR?

The best thing since sliced bread ;-)

- Big improvement to DNS lookups
  - Old (at least 5 lookups):

```
www.akamai.com. 20 IN CNAME www.akamai.com.edgekey.net.
www.akamai.com.edgekey.net. 20 IN CNAME www.akamai.com.edgekey.net.globalredir.akadns.net.
www.akamai.com.edgekey.net.globalredir.akadns.net. 20 IN CNAME e1699.dscx.akamaiedge.net.
e1699.dscx.akamaiedge.net. 20 IN A 23.75.246.178
e1699.dscx.akamaiedge.net. 20 IN AAAA 2a02:26f0:3100:39e::6a3
e1699.dscx.akamaiedge.net. 20 IN AAAA 2a02:26f0:3100:393::6a3
```

  - New (down to 1 lookup):

```
www.akamai.com. 20 IN HTTPS 1 e1699.dscx.akamaiedge.net. alpn="h2” ipv4hint="23.75.246.178”
ipv6hint=“2a02:26f0:3100:39e::6a3,2a02:26f0:3100:393::6a3"
```

  - Also allows CNAME at APEX
    - No more ANAME, BNAME, ..

```
akamai.com. 20 IN HTTPS 0 www.akamai.com.edgekey.net.
```

Photo: Fran Hogan
https://commons.wikimedia.org/wiki/File:Fresh_made_bread_05.jpg

Why HTTPS Now?

- DNS community tried before
  - SRV (RFC2052, RFC2782)
  - draft-bellis-dnsop-http-record
  - draft-ietf-dnsop-aname
  - draft-yao-dnsext-bname

But…

It doesn't matter what you put in the DNS until it gets implemented and used

Why HTTPS now?

- SVCB (Service Binding) and HTTPS
  - Originated from the Alt-Svc HTTP header (RFC7838)
  - Authors:
    - Ben Schwartz (Google)
    - Mike Bishop (Akamai)
    - Erik Nygren (Akamai)
  - Already implemented by Apple
  - Already used - *lots* of traffic

Overview of Traffic Impact

All

A

AAAA

HTTPS

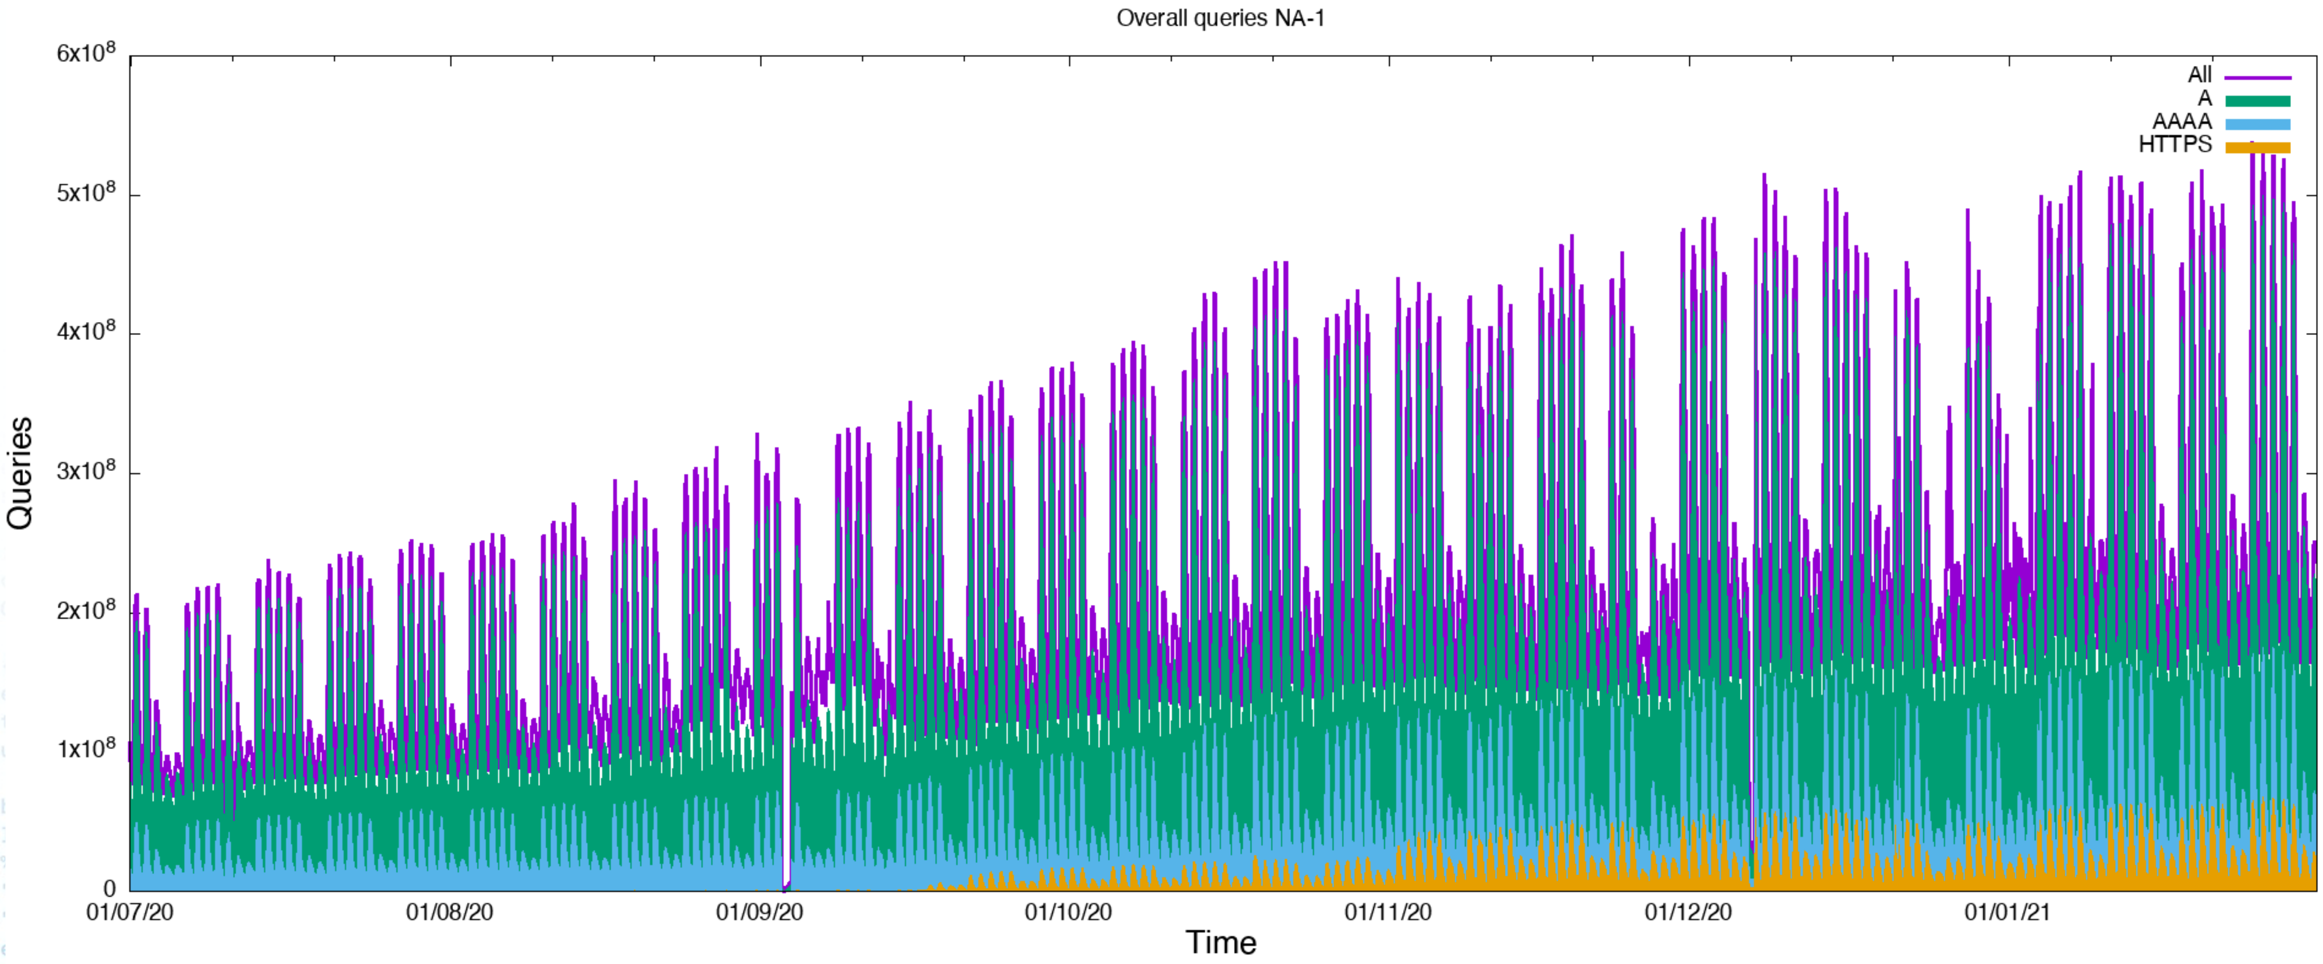

Regional Impacts

All
A
AAAA
HTTPS

South America

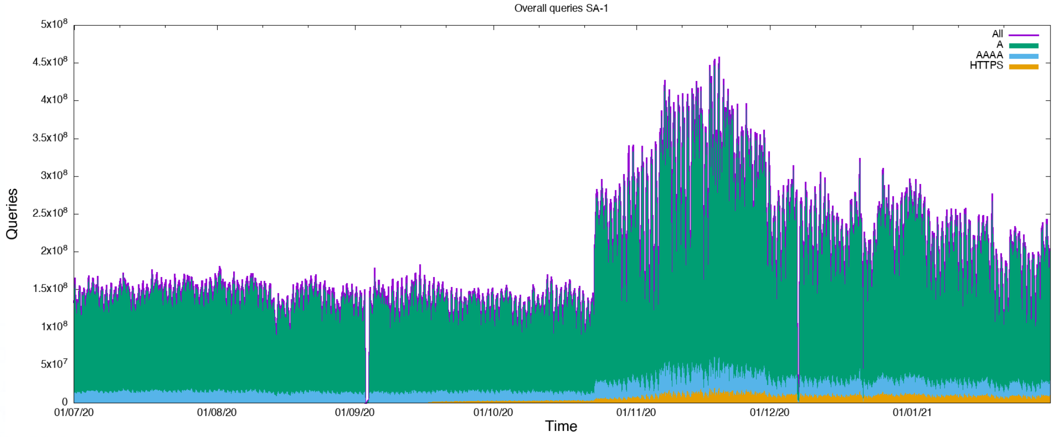

North America

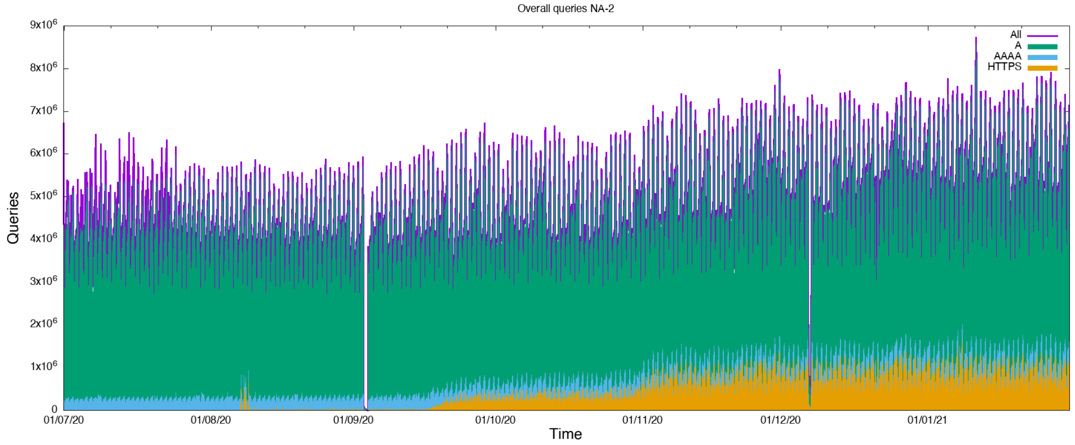

Asia Pacific

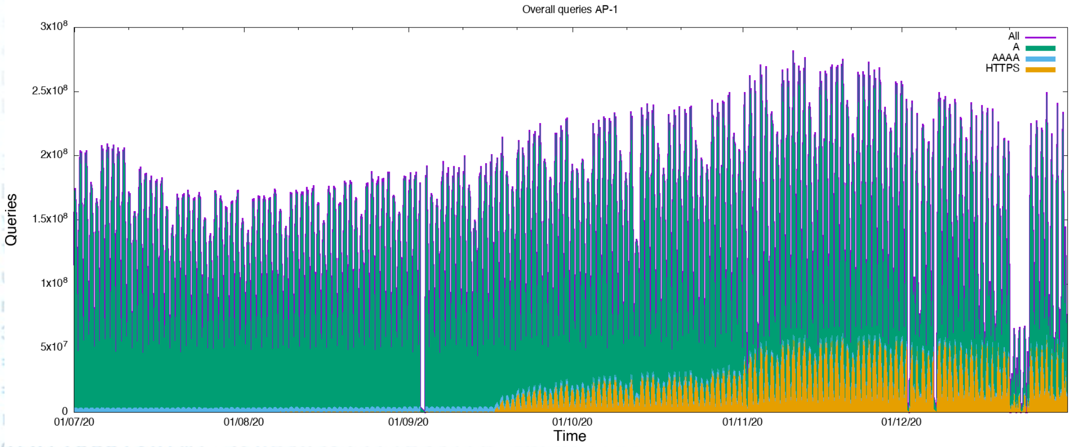

Europe

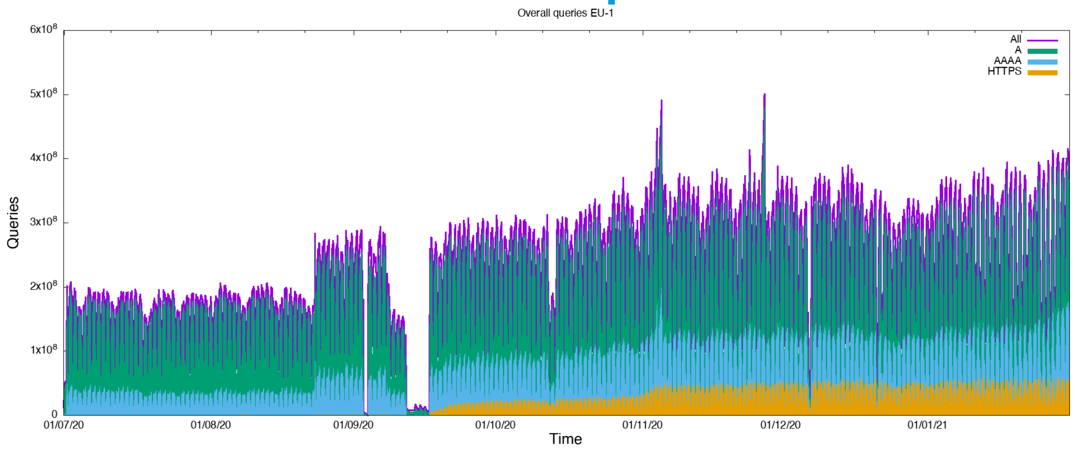

When Did It Actually Start?

All
A
AAAA
HTTPS

South America

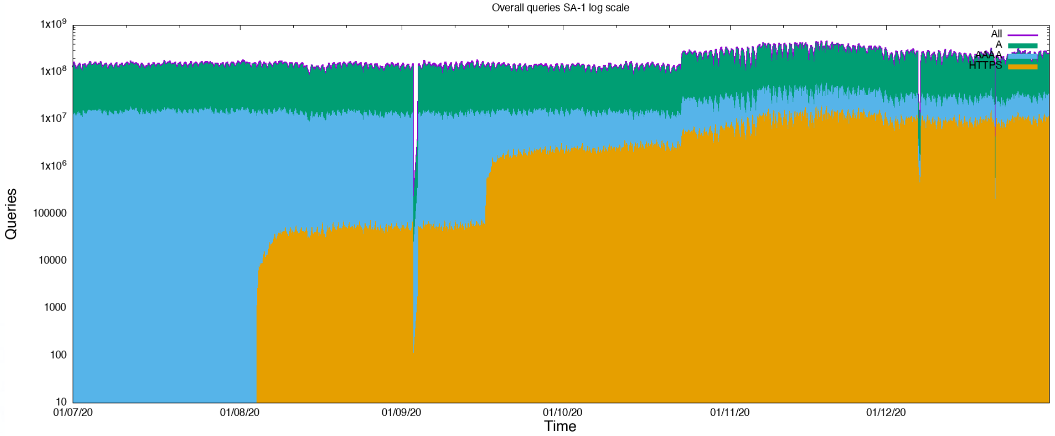

North America

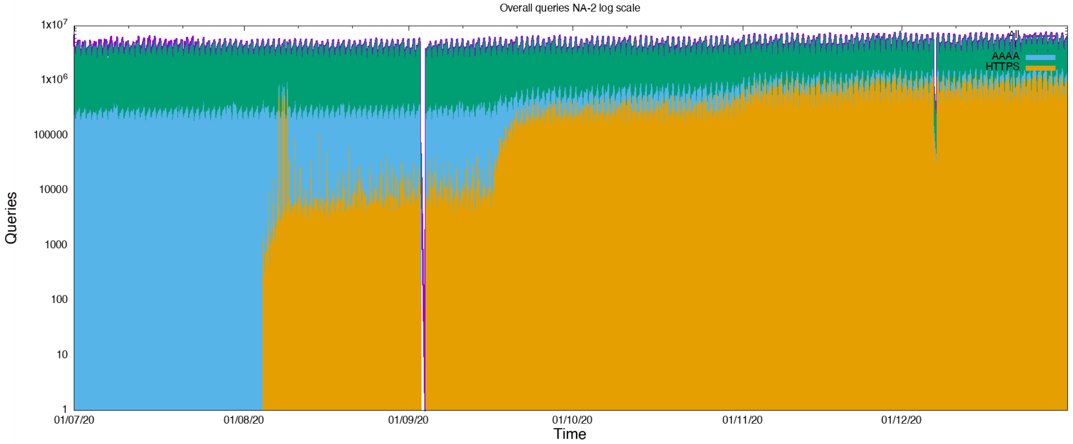

Asia Pacific

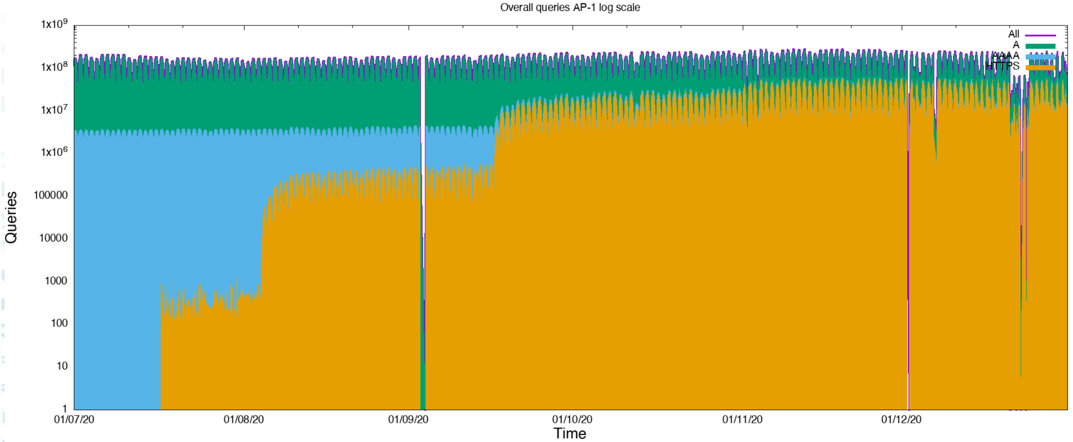

Europe

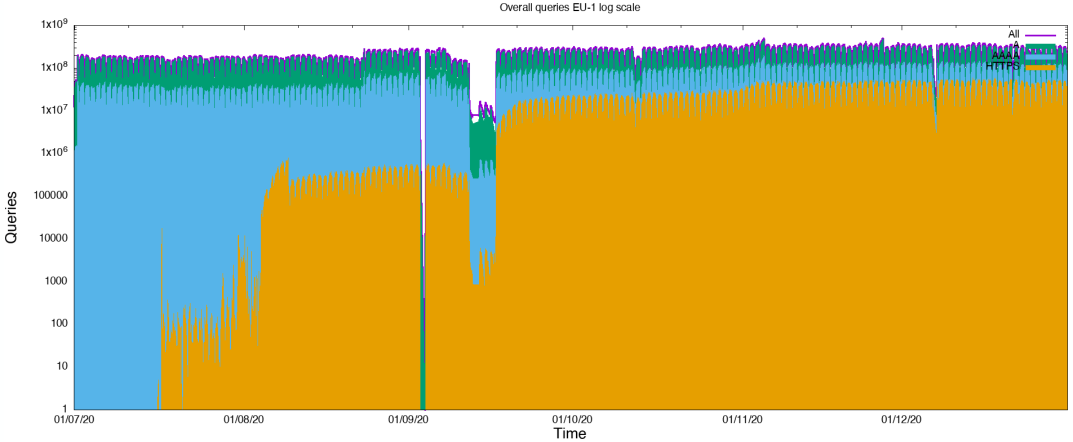

Did It Go Wrong ?

- What signals “wrong” for DNS?
  - NXDomain just means somethings not there
  - NOTIMP is rarely (never) given out by resolvers
  - SERVFAIL means something has gone wrong, but we have no idea what
  - Hopefully extended DNS errors give a better picture in the future
- For now SERVFAIL is the signal
  - Graphs show percentage of SERVFAIL against all queries per query type
    - HTTPS
    - A
    - AAAA
    - All

Percentage Good vs Bad Queries

All 🟪
A 🟩
AAAA 🟦
HTTPS 🟨

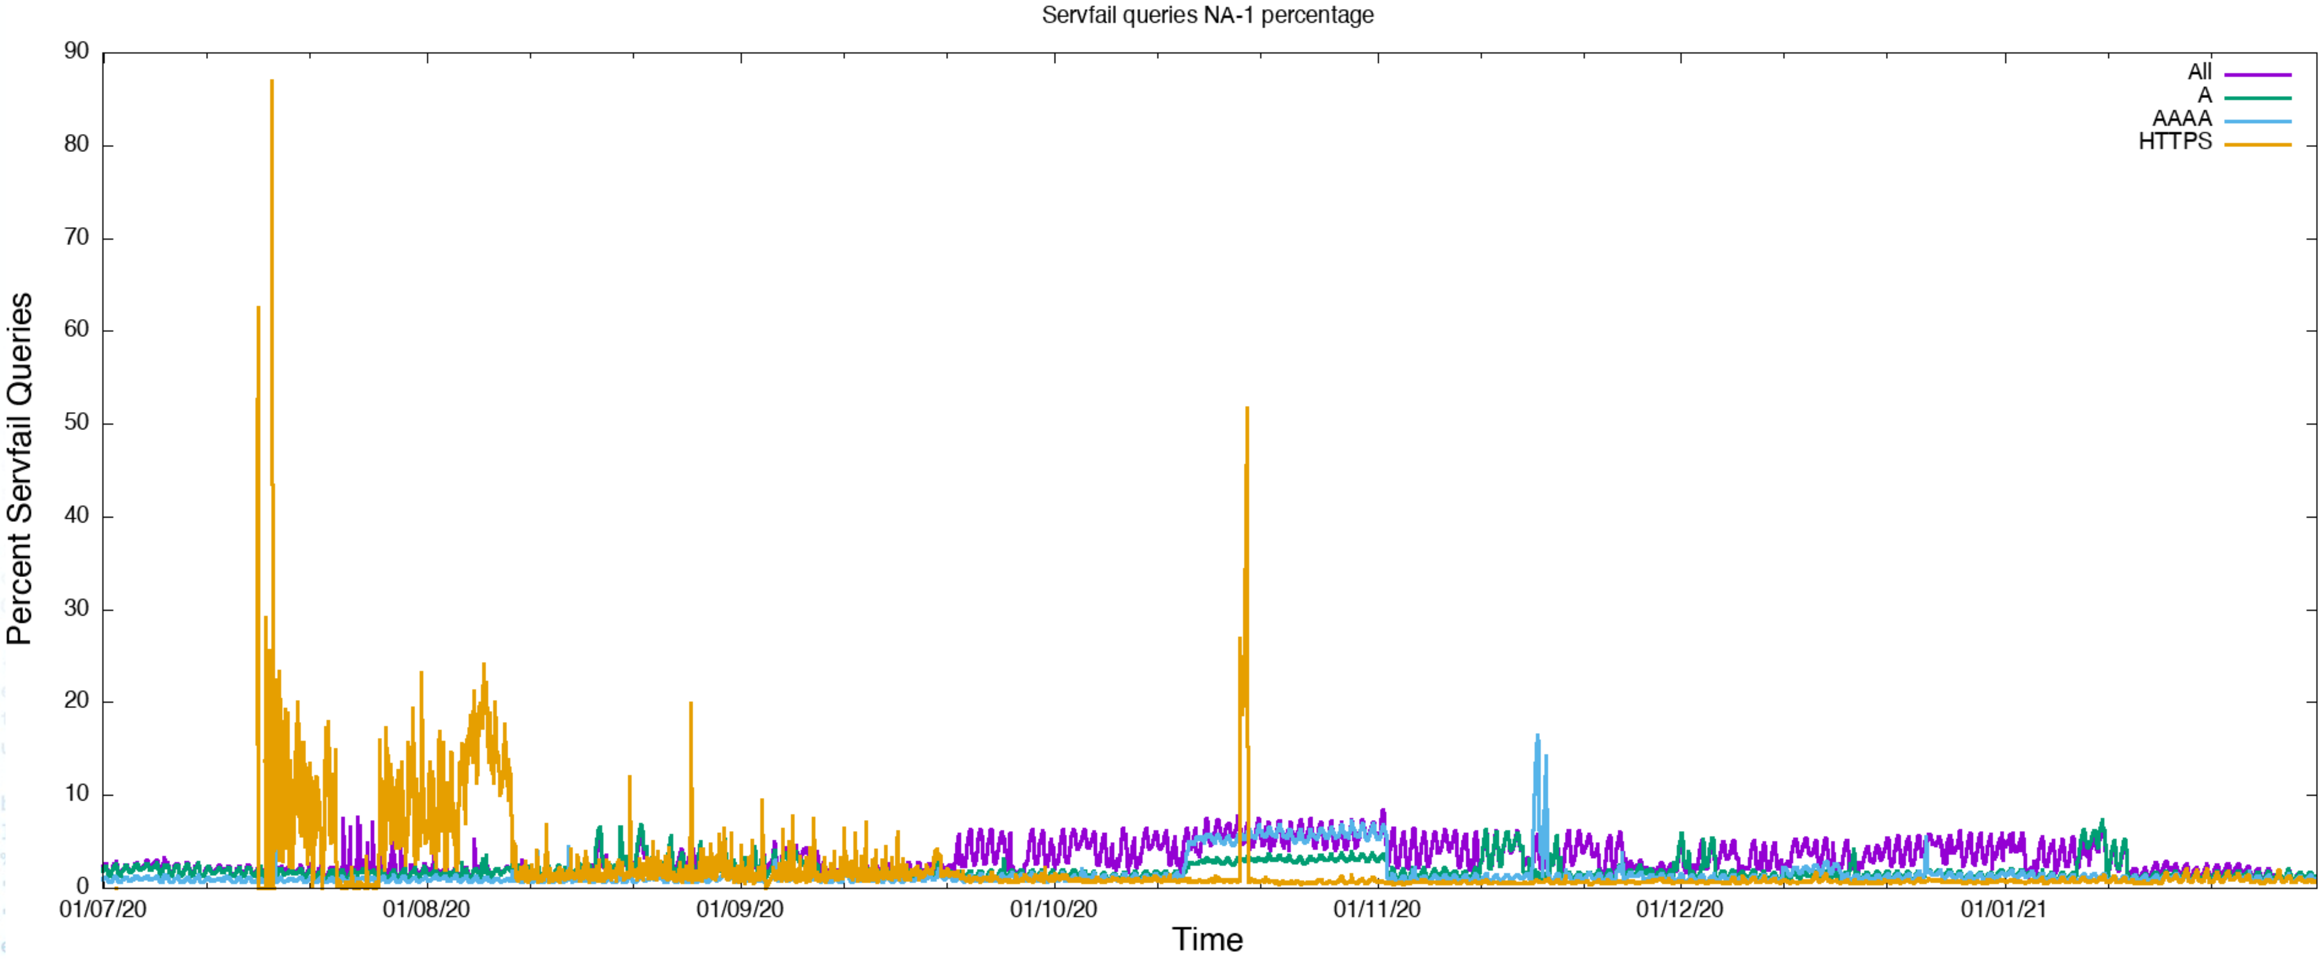

Regional Percentage Good vs Bad Queries

All ▬
A ▬
AAAA ▬
HTTPS ▬

South America

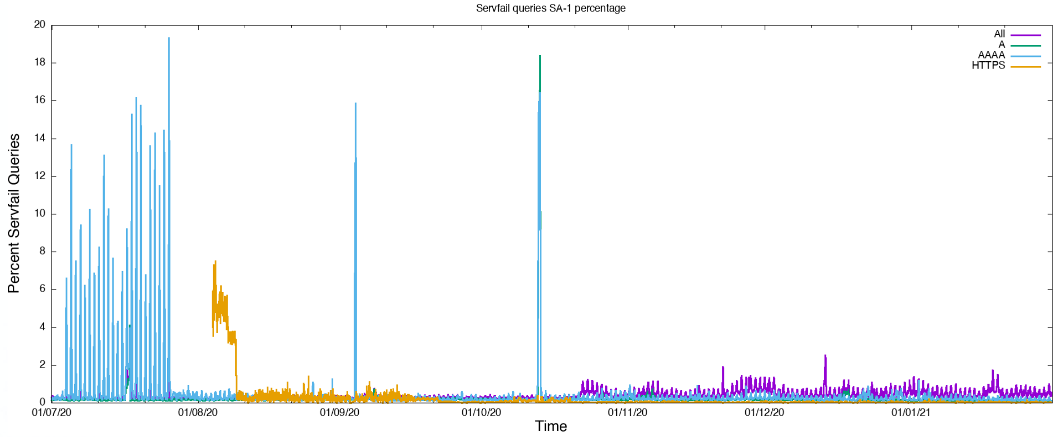

North America

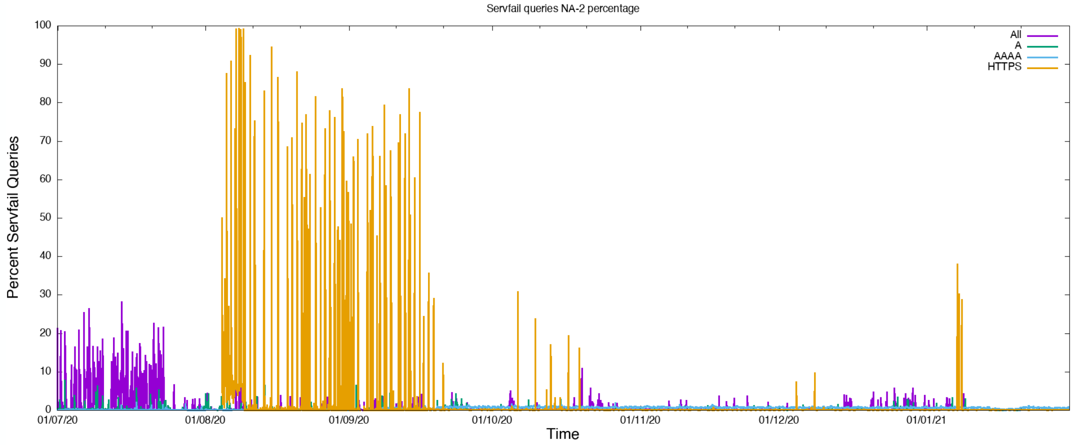

Asia Pacific

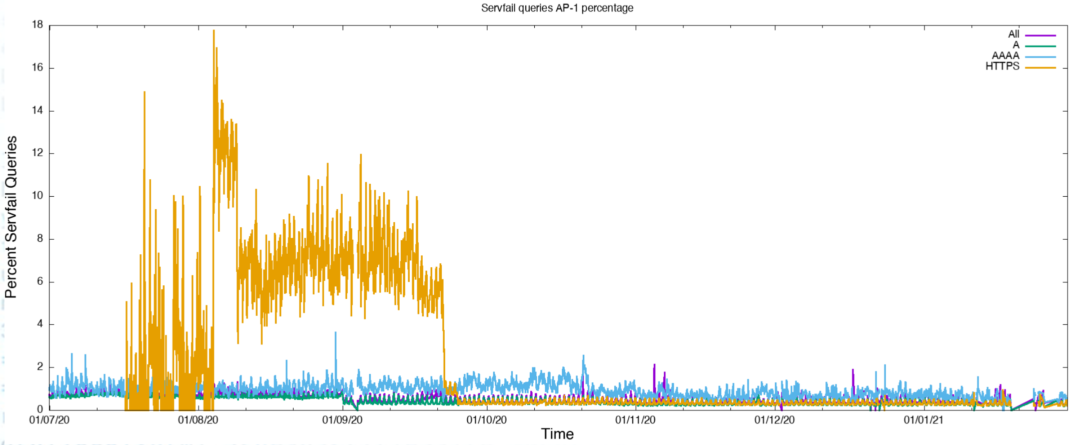

Europe

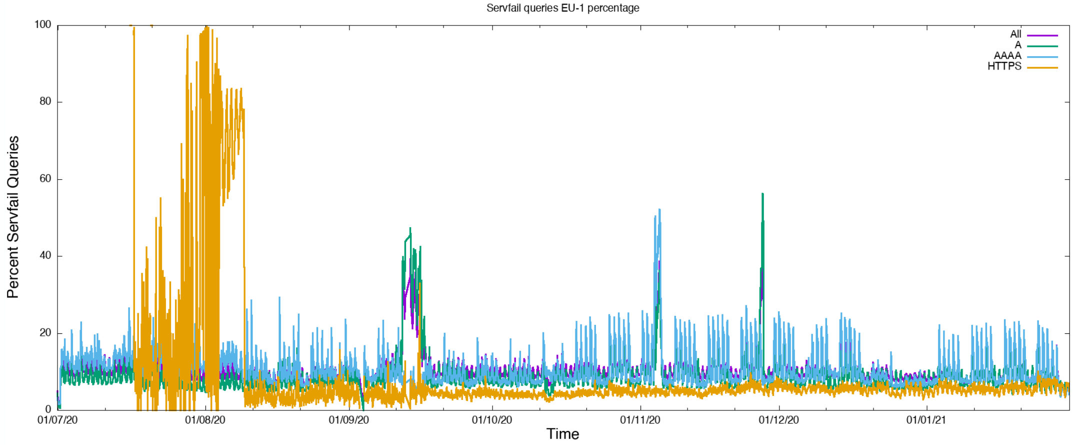

Akamai Experience the Edge

What’s the Correct Answer For a New Query Type?

- NXDOMAIN if the label does not exist for any type

```
;; ->>HEADER<<- opcode: QUERY; status: NXDOMAIN; id: 1790
;; Flags: qr rd ra; QUERY: 1; ANSWER: 0; AUTHORITY: 1; ADDITIONAL: 0

;; QUESTION SECTION:
;; bla.dns-oarc.net.              IN      HTTPS

;; AUTHORITY SECTION:
dns-oarc.net. 120 IN SOA ns1.dns-oarc.net. hostmaster.dns-oarc.net. 2021011413 300 60 604800 3600
```

- NOERROR and an empty or CNAME only answer section

```
;; ->>HEADER<<- opcode: QUERY; status: NOERROR; id: 26665
;; Flags: qr rd ra; QUERY: 1; ANSWER: 0; AUTHORITY: 1; ADDITIONAL: 0

;; QUESTION SECTION:
;; www.dns-oarc.net.              IN      HTTPS

;; AUTHORITY SECTION:
dns-oarc.net. 120 IN SOA ns1.dns-oarc.net. hostmaster.dns-oarc.net. 2021011413 300 60 604800 3600
```

- Or better answer with HTTPS resource record data

Domains Using HTTPS

- 403 million unique FQDN
- 12 million asked HTTPS RRs
- 126430 have actual HTTPS records
- Sample from February 1
  - Earlier days look similar

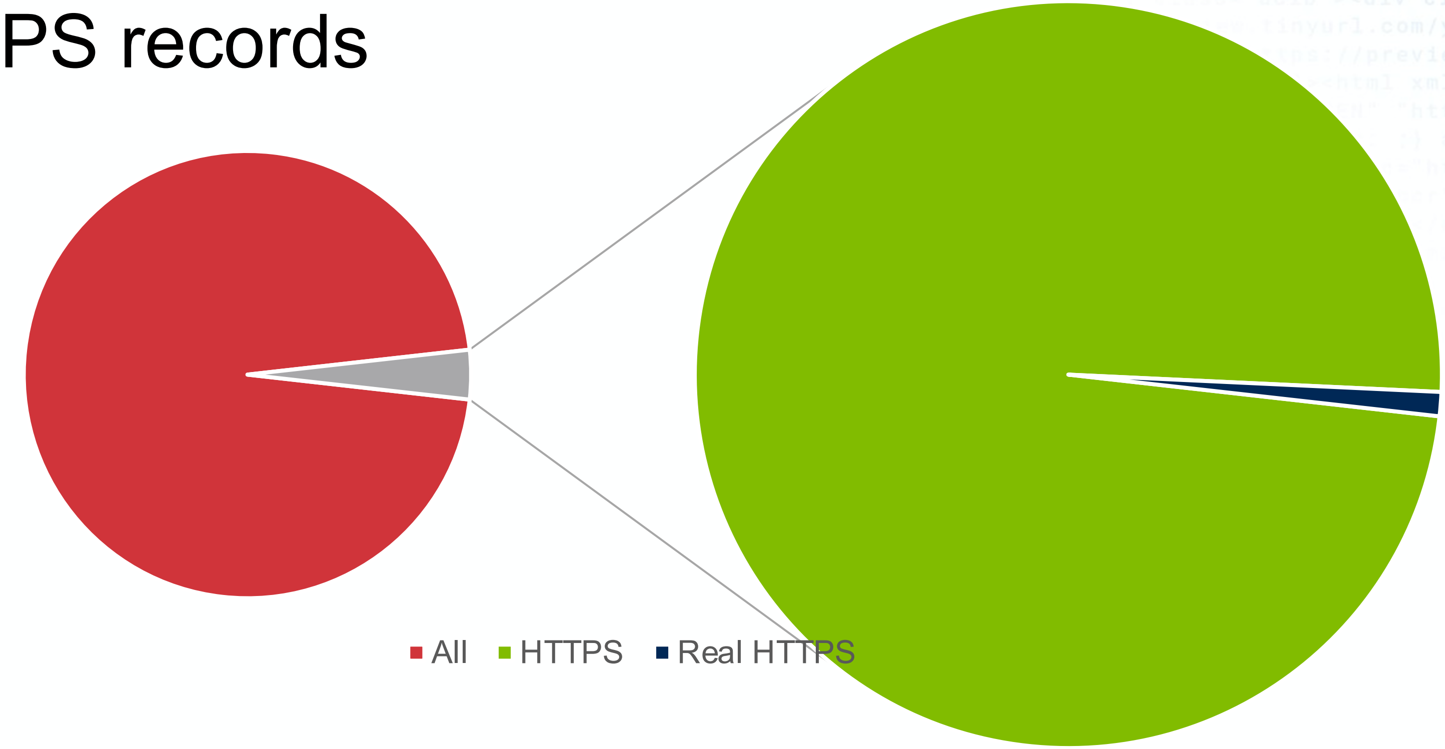

What HTTPs Records Look Like

- No domains use the alias feature ☹
- Only one domain uses more than one RR
  - Nextdns.io
- 912 have no ALPN
  - Most of them are reachable via HTTP
  - 67 are HTTP only
  - Apples implementation reaches the HTTP websites
- Nearly all have hints (99.94%)
  - Only 17% have IPv4 only hints
- No use of other keys
  - ECH
  - PORT
  - KEYXX

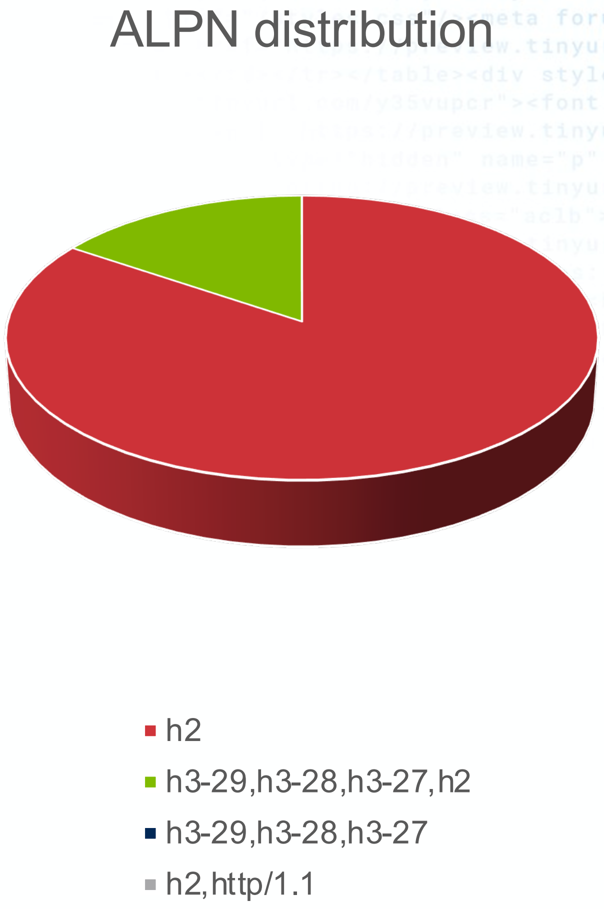

What Can Go Wrong?

- Out of the 12 million domains only 168k sometimes answer wrong
  - Means there is no coherent result code for all types
    - E.g NOERROR when asking A, SERVFAIL for HTTPS
  - Most of these are intermediate
- Only 4652 proved really wrong
- Not answering when asked for an HTTPS record
  - Some NS only answer when asked for A/AAAA
  - Some have an abritray limit at TYPE60
- Answering NXDOMAIN for HTTP, but not for A/AAAA

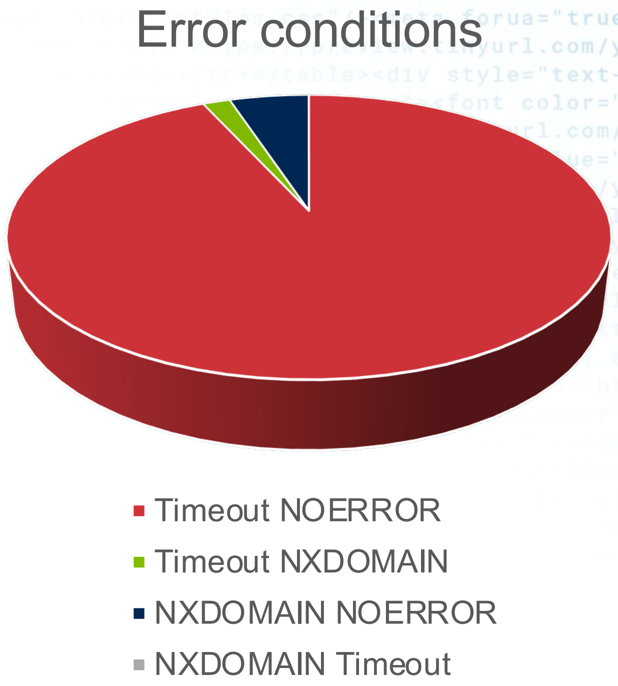

Summary

- Adoption of HTTPS record type has been relatively smooth
  - Initially there were problems
- Currently only one implementation and limited feature set usage
  - Apple iOS/macOS
- DNS is hard and their always will be problematic implementations
- To all vendors:
  - Answering is always better then dropping
  - Query types other then A/AAAA are NOT attacks
  - Answering NXDomain when you just don’t have data for the type is wrong
- Will SVCB ever be used ;-)

Thank you!
Questions?

Akamai

Experience the Edge