IT’S DNS, JIM…

… BUT NOT AS WE KNOW IT

Jim Hague
Sinodun Internet Technologies
jim.hague@acm.org
@banbury_bill
https://github.com/banburybill

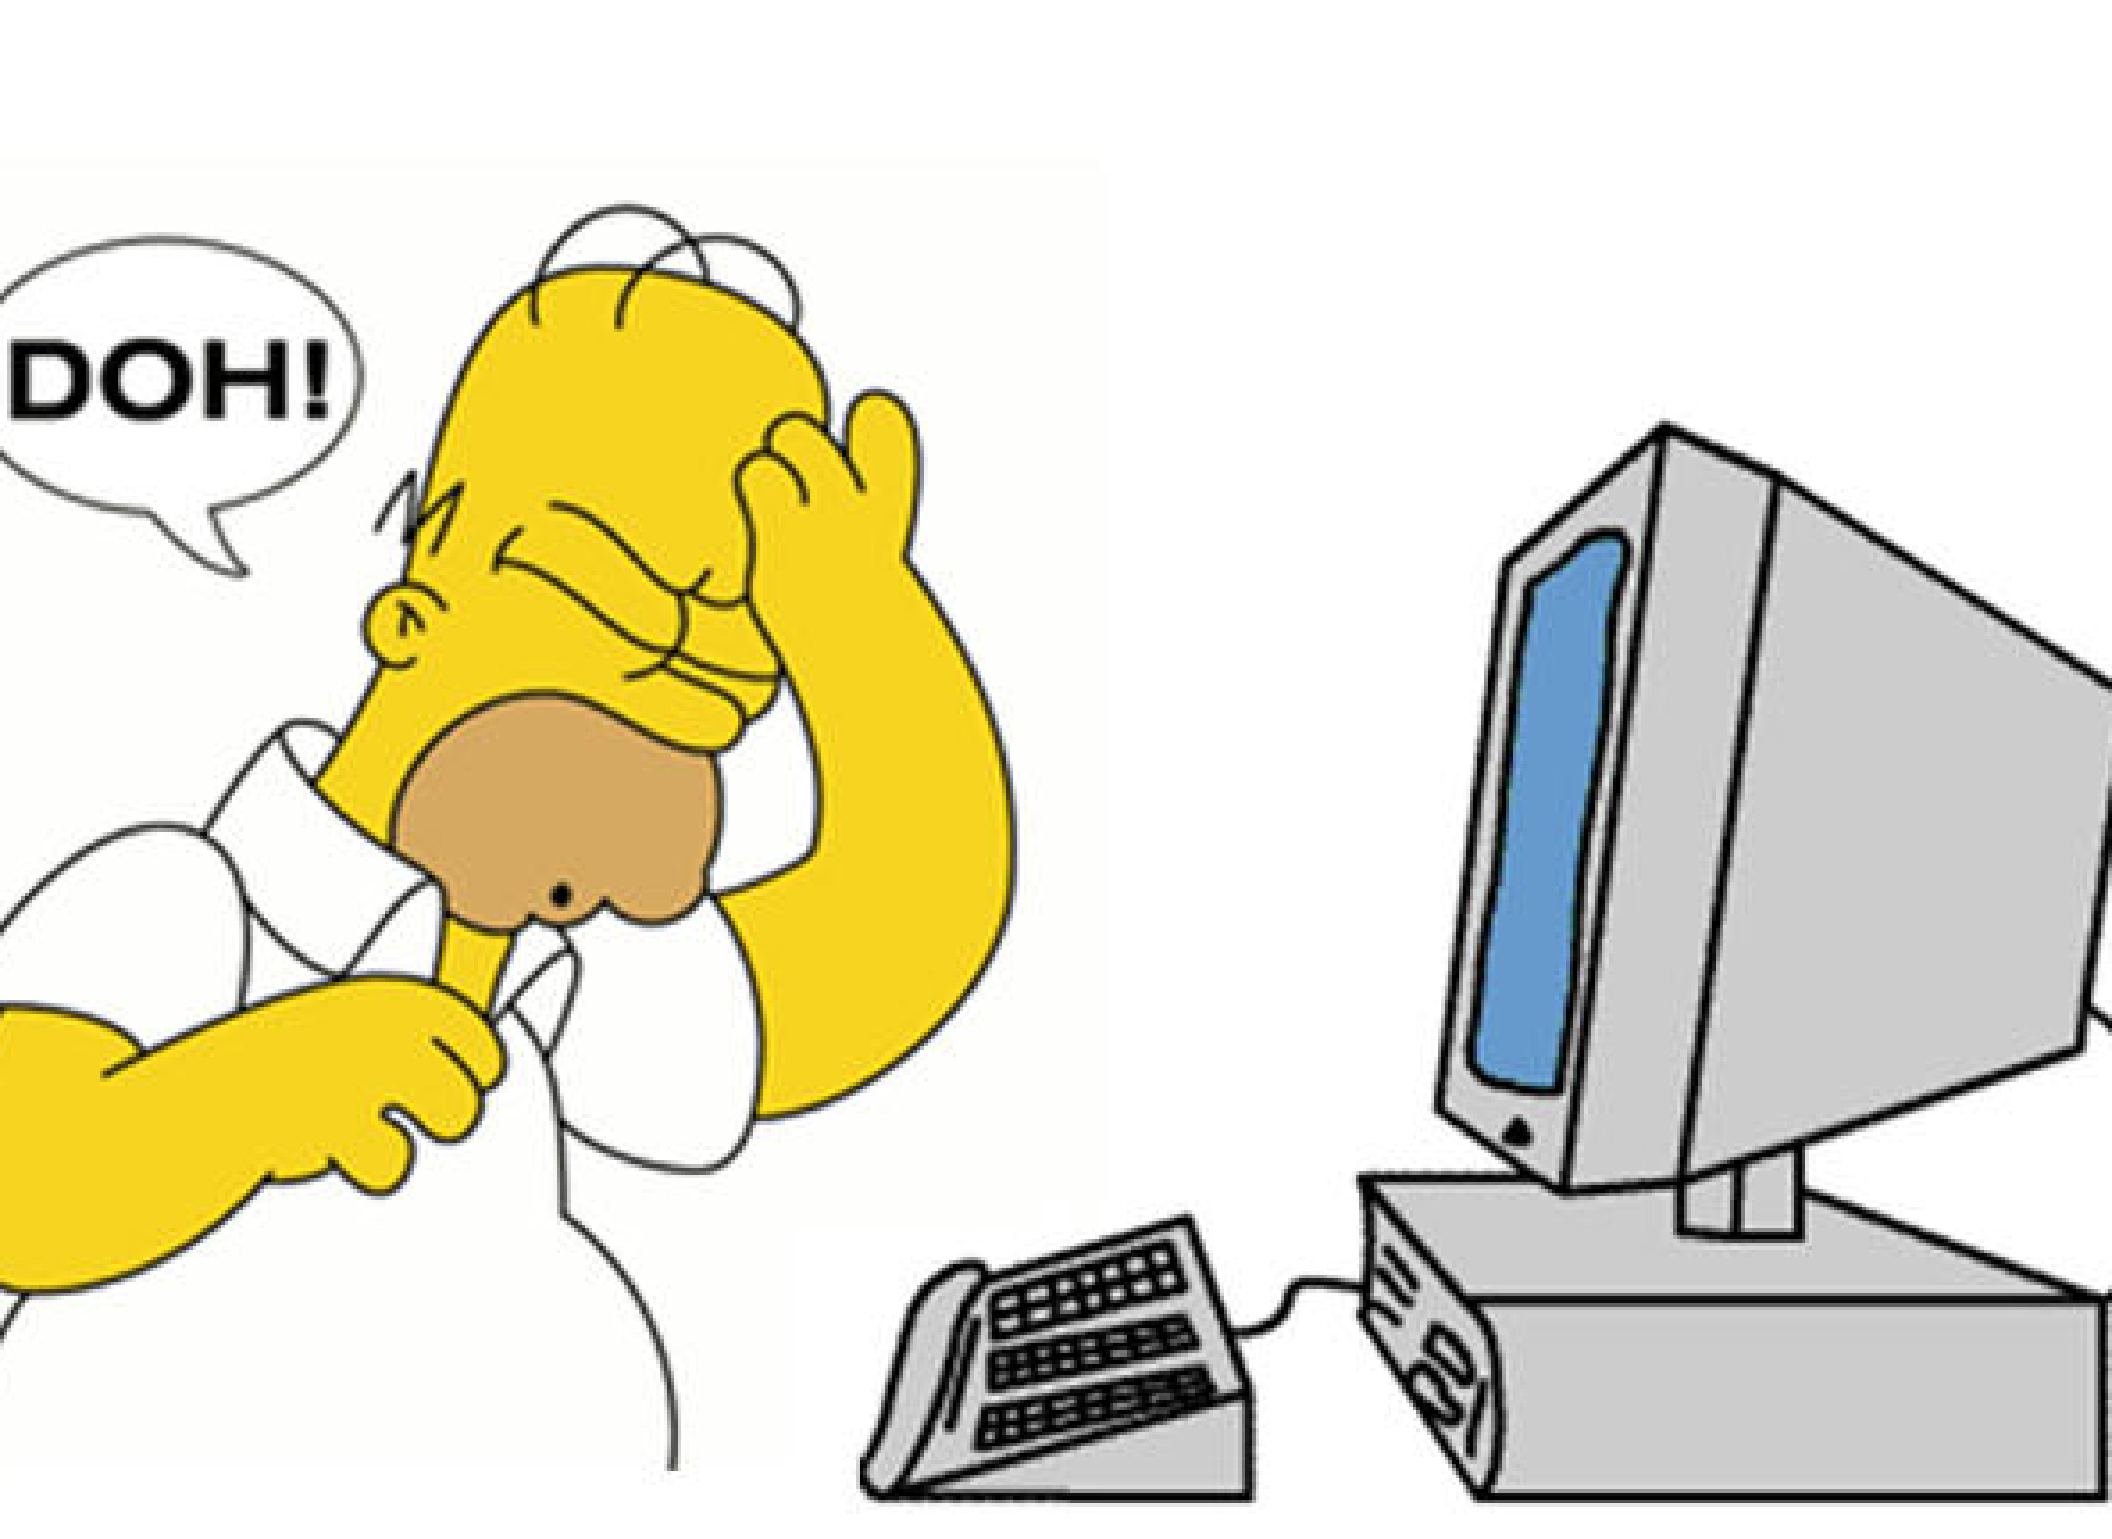
DOH!

Internet Engineering Task Force (IETF)
Request for Comments: 8484
Category: Standards Track
ISSN: 2070-1721

P. Hoff
IC
P. McMa
Mozi

October 2

DNS Queries over HTTPS (DoH)

Abstract

This document defines a protocol for sending DNS queries and getti
DNS responses over HTTPS. Each DNS query-response pair is mapped
into an HTTP exchange.


The Beginning

Network Working Group P. Mockapet
Request for Comments: 1034
Obsoletes: RFCs 882, 883, 973 November 1

DOMAIN NAMES - CONCEPTS AND FACILITIES

1. STATUS OF THIS MEMO

This RFC is an introduction to the Domain Name System (DNS), and omit many details which can be found in a companion RFC, "Domain Names - Implementation and Specification" [RFC-1035]. That RFC assumes that reader is familiar with the concepts discussed in this memo.

A subset of DNS functions and data types constitute an official protocol. The official protocol includes standard queries and their responses and most of the Internet class data formats (e.g., host addresses).

DNS

- A consistent namespace used for referring to resources.
- Maintained in a distributed manner.
- Local caching to improve performance.

THE DOMAIN NAMESPACE

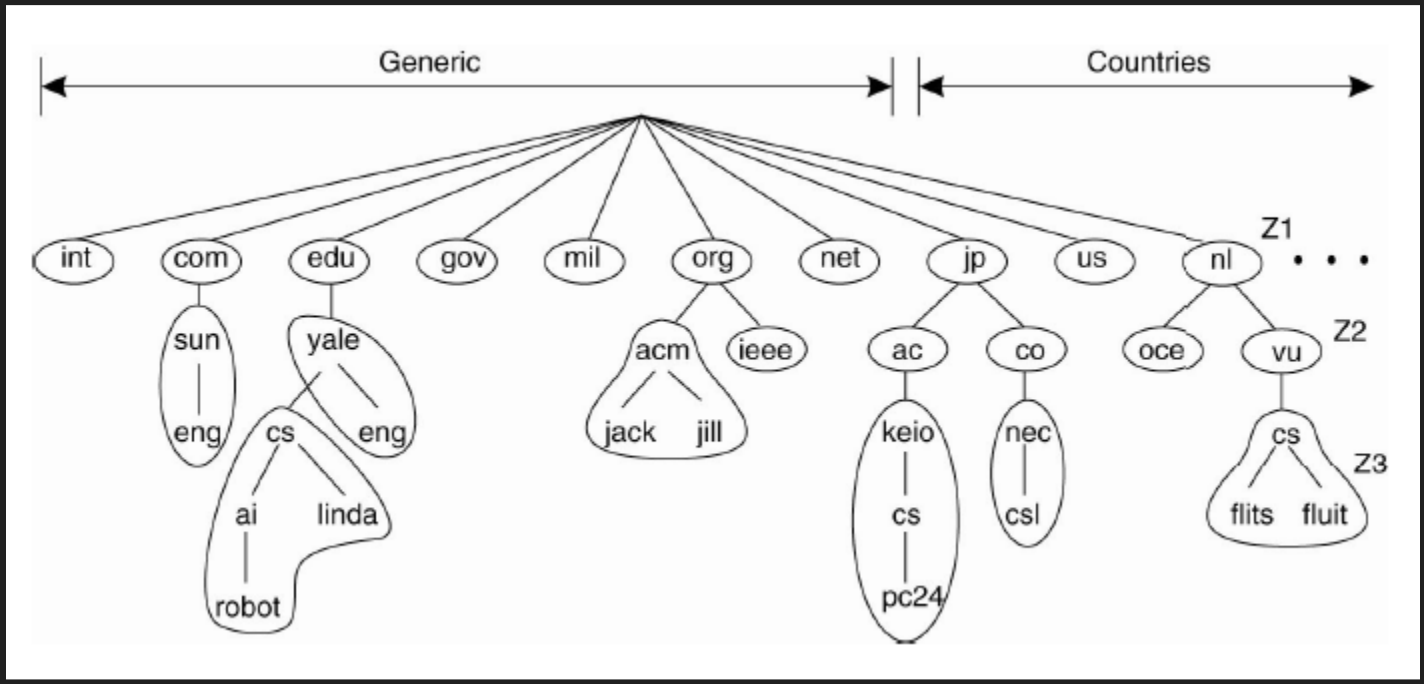

DELEGATION OF AUTHORITY

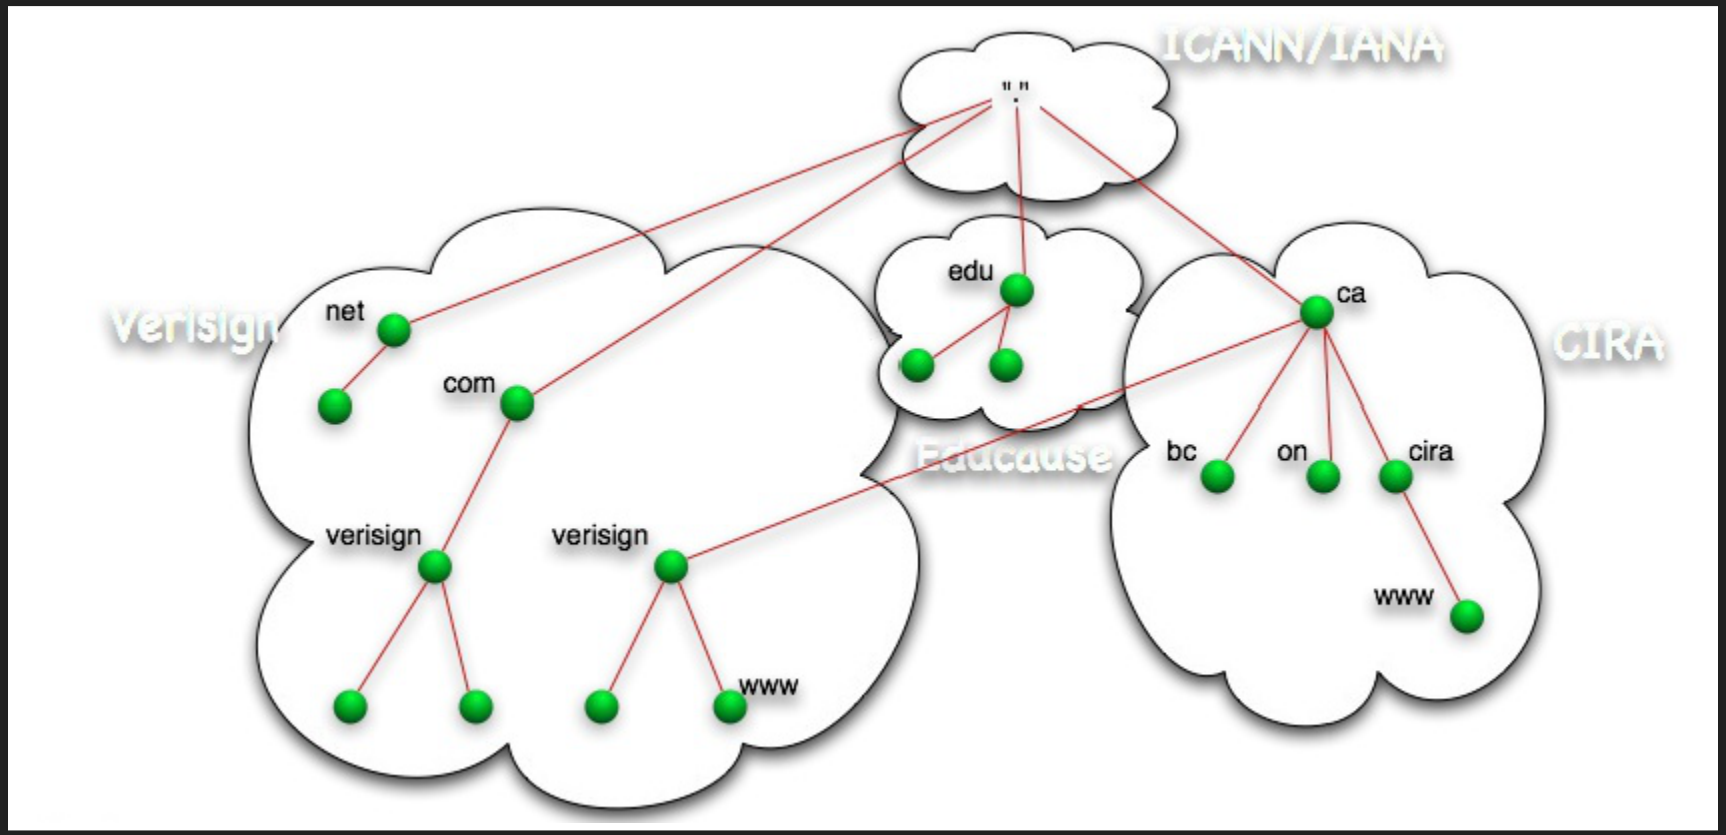

ROOT SERVERS

- A fixed list of IPv4 and IPv6 addresses for 13 servers.
  - `a.root-servers.net` .. `m.root-servers.net`

ROOT SERVERS OPERATORS

- VeriSign, Inc., University of Southern California, Cogent Communications, University of Maryland, NASA Ames Research Centre, Internet Systems Consortium, Inc., US Department of Defence, US Army Research Lab, Netnod (Sweden), RIPE, ICANN, WIDE Project (Japan)

IANA has the details

AUTHORITATIVE SERVERS

- Contain the data for a zone
- Run by the zone owner

RECURSIVE SERVERS

- Search the hierarchy to resolve queries
- Cache results and reuse them in future queries
- Typically run by ISP…
- … or 3rd party, e.g. Google, OpenDNS

STUB RESOLVER

- Your local name resolution
- Typically using recursive server(s) supplied by DHCP

A LOOK AT THE WIRE

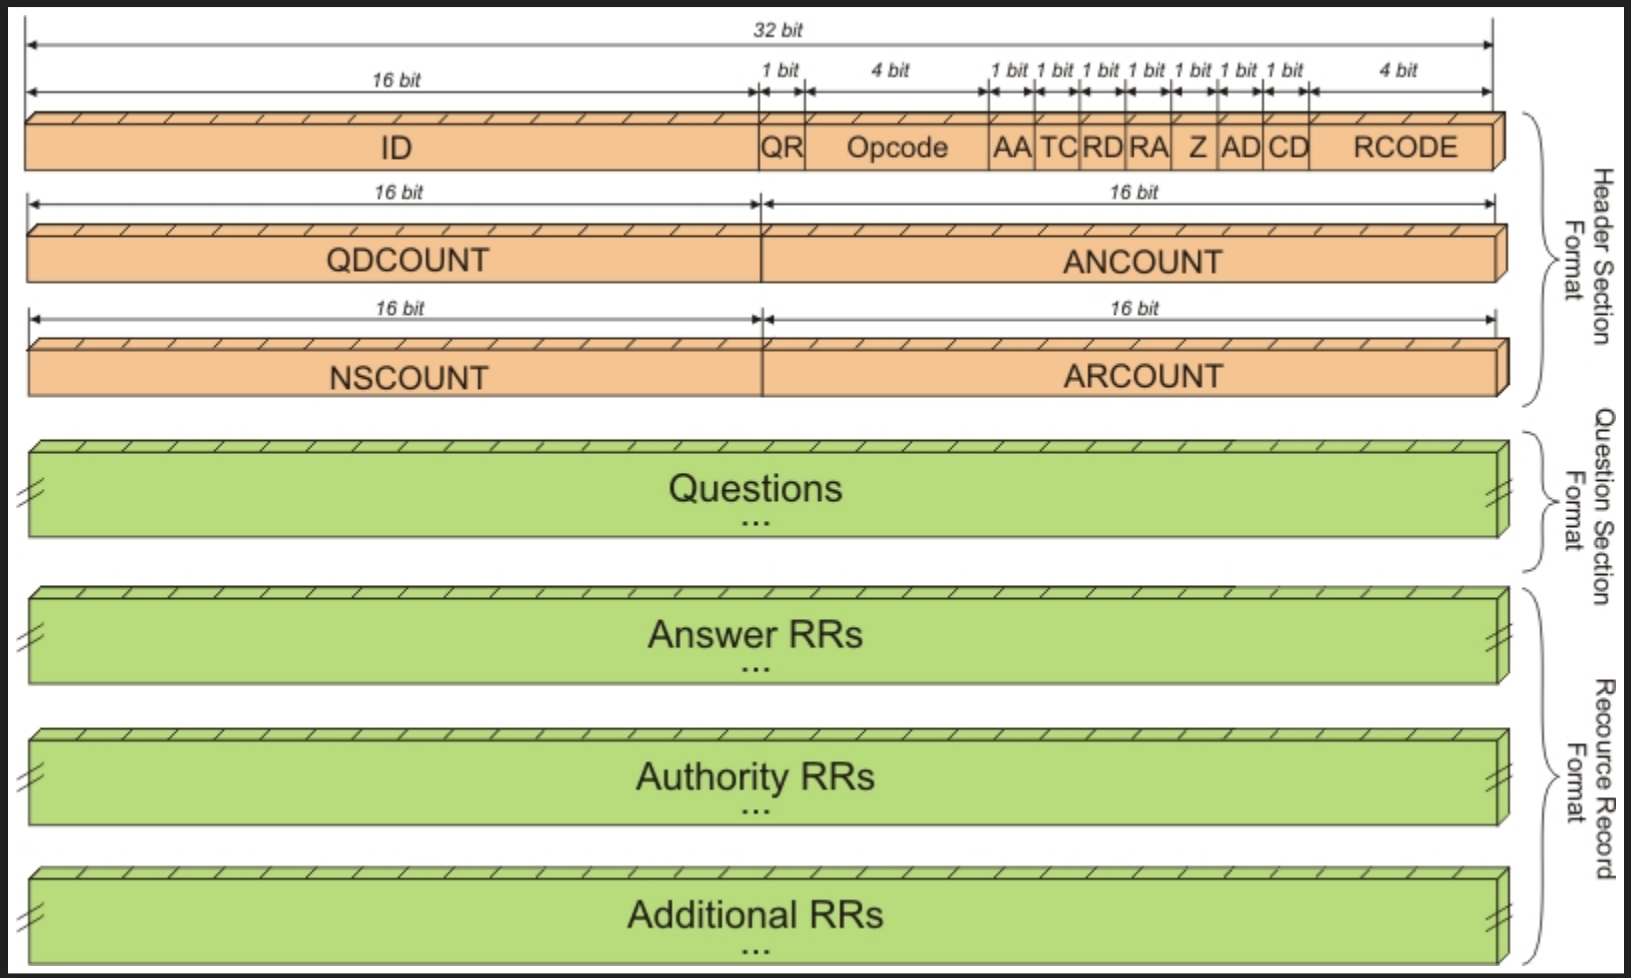

Format of a DNS message

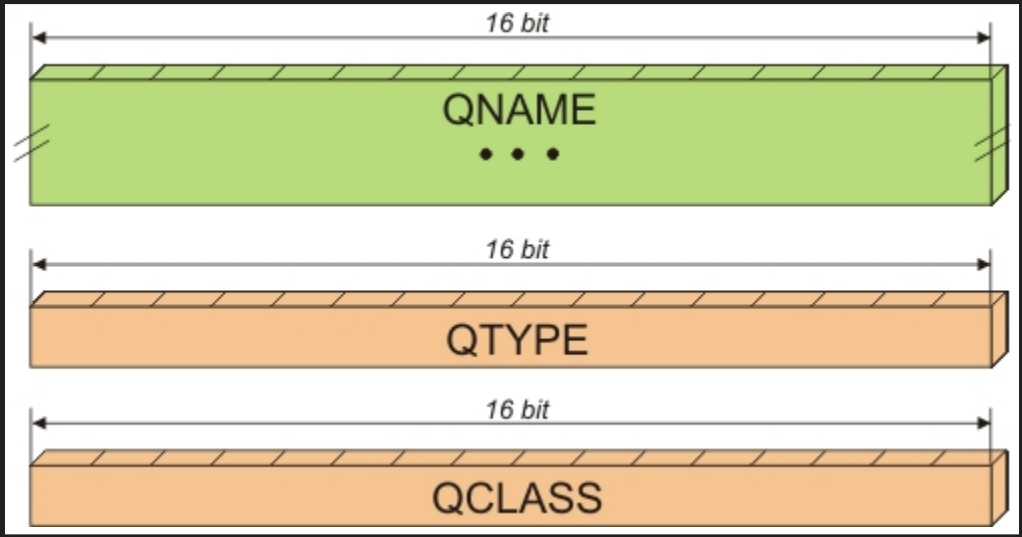

Format of a Question section

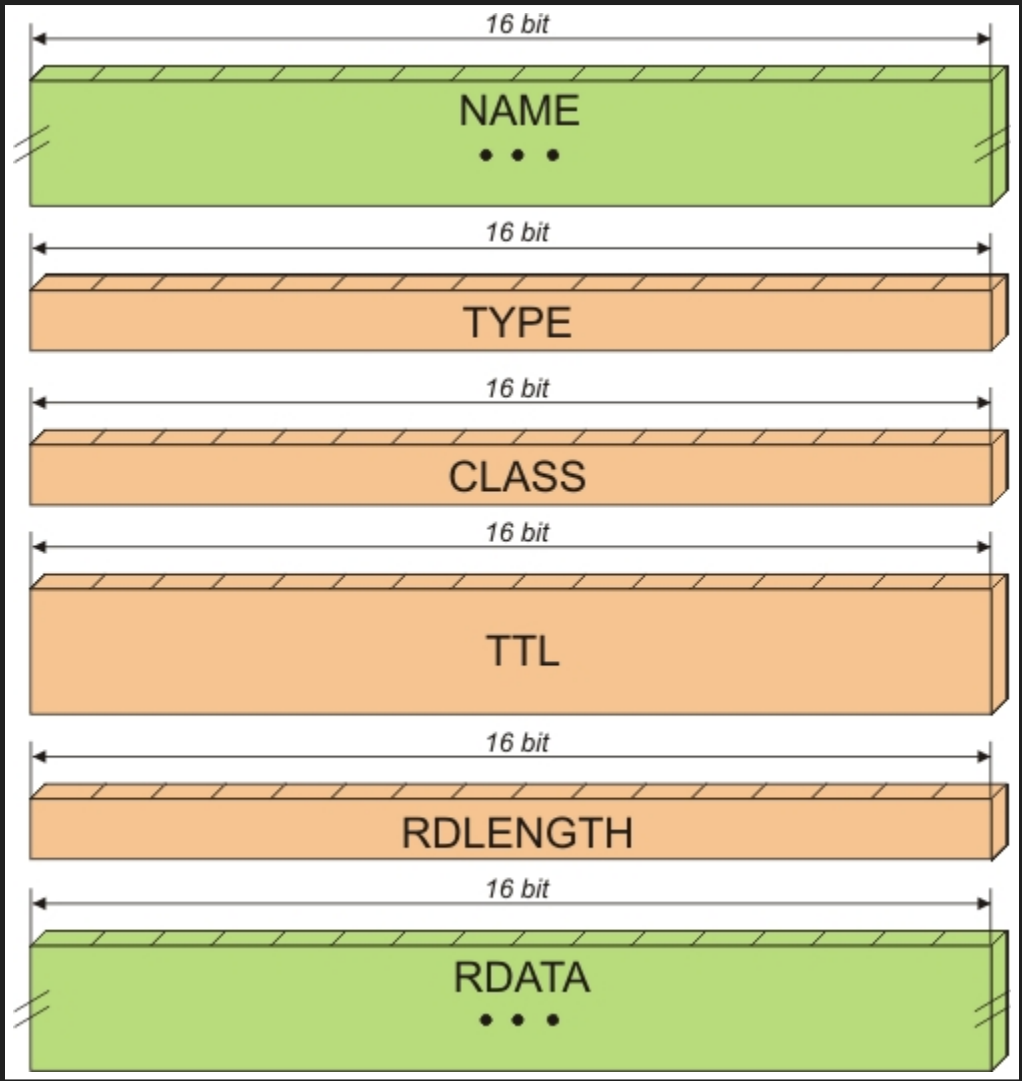

Format of a RR section

A	IPv4 address
AAAA	IPv6 address
MX	SMTP servers for domain
NS	Name servers for domain
PTR	Pointer to canonical name (for address)
SRV	Location of servers providing given service
TXT	General textual information
SOA	Start of Authority record for zone

TRANSMISSION

TRANSMISSION

- DNS uses UDP

TRANSMISSION

- DNS uses UDP
- Except when it uses TCP

Network Working Group P. V
Request for Comments: 2671
Category: Standards Track August

Extension Mechanisms for DNS (EDNS0)

...

Abstract

The Domain Name System's wire protocol includes a number of fixed fields whose range has been or soon will be exhausted and does not allow clients to advertise their capabilities to servers. This document describes backward compatible mechanisms for allowing the protocol to grow.

EDNS0

- Extends RCODE range and number of flags.
- Mechanism to allow larger UDP messages. This is necessary because of an increase in DNS RR sizes:
  - AAAA records
  - Large TXT records
  - DNSSEC

NAME		Always 00
TYPE	16 bits	OPT (41)
CLASS	16 bits	Sender UDP payload size
TTL	32 bits	uint8 extended RCODE
		uint8 version (0)
		uint16 flags
RDLEN	16 bits	Length of RDATA
RDATA		Options. Any number of:
		uint16 Option Code
		uint16 Option length
		Option data

Internet Engineering Task Force (IETF)
Request for Comments: 7871
Category: Informational
ISSN: 2070-1721

C. Contava
W. van der Ga
Goo
D. Lawre
Akamai Technolog
W. Kum
Goo

May 2

Client Subnet in DNS Queries

Abstract

This document describes an Extension Mechanisms for DNS (EDNS0) option that is in active use to carry information about the networ that originated a DNS query and the network for which the subseque response can be cached. Since it has some known operational and privacy shortcomings, a revision will be worked through the IETF f improvement.

EDNS0 ECS CLIENT SUBNET

- An unusual RFC
- Encodes client subnet into the query
  - So CDN knows geographic location of client

Network Working Group
Request for Comments: 4033
Obsoletes: 2535, 3008, 3090, 3445, 3655, 3658,
3755, 3757, 3845
Updates: 1034, 1035, 2136, 2181, 2308, 3225,
3007, 3597, 3226
Category: Standards Track

R. Are
Telematica Instit
R. Aust
M. Lar
VeriS
D. Mas
Colorado State Univers
S. R
N

March 2

DNS Security Introduction and Requirements

[omitted]

The Domain Name System Security Extensions (DNSSEC) add data origi
authentication and data integrity to the Domain Name System. This
document introduces these extensions and describes their capabilit
and limitations. This document also discusses the services that t
DNS security extensions do and do not provide.

DNSSEC

- Assures authenticity of DNS data
- Assures integrity of DNS data
  - Note it authenticates DNS data, NOT DNS servers
- Does NOT ensure confidentiality

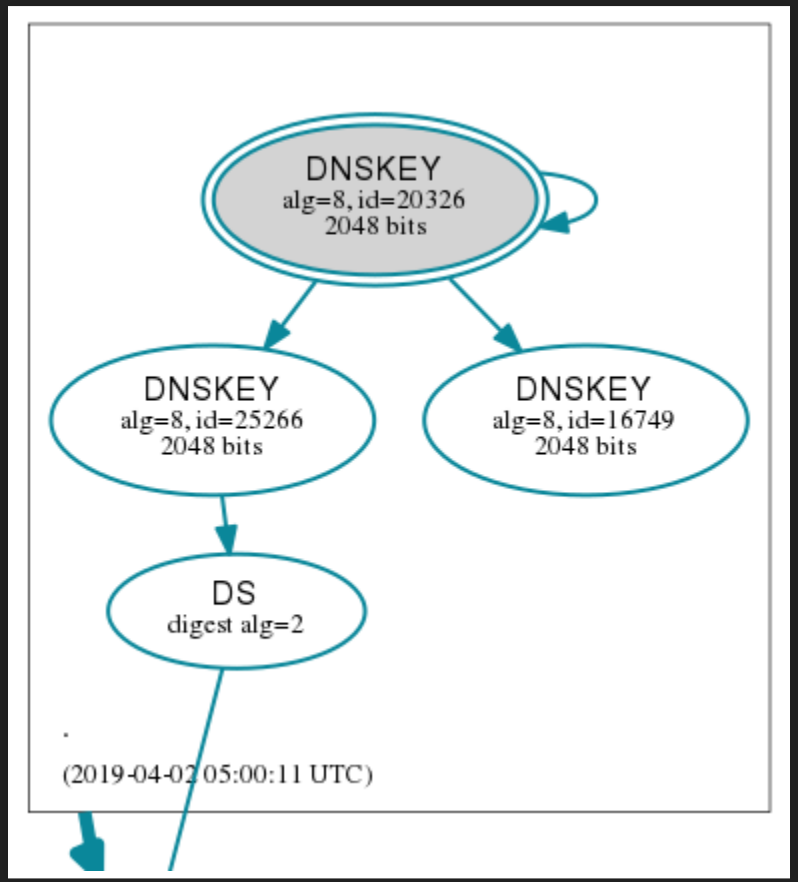
DNSKEY
alg=8, id=20326
2048 bits
DNSKEY
alg=8, id=25266
2048 bits
DNSKEY
alg=8, id=16749
2048 bits
DS
digest alg=2
.
(2019-04-02 05:00:11 UTC)

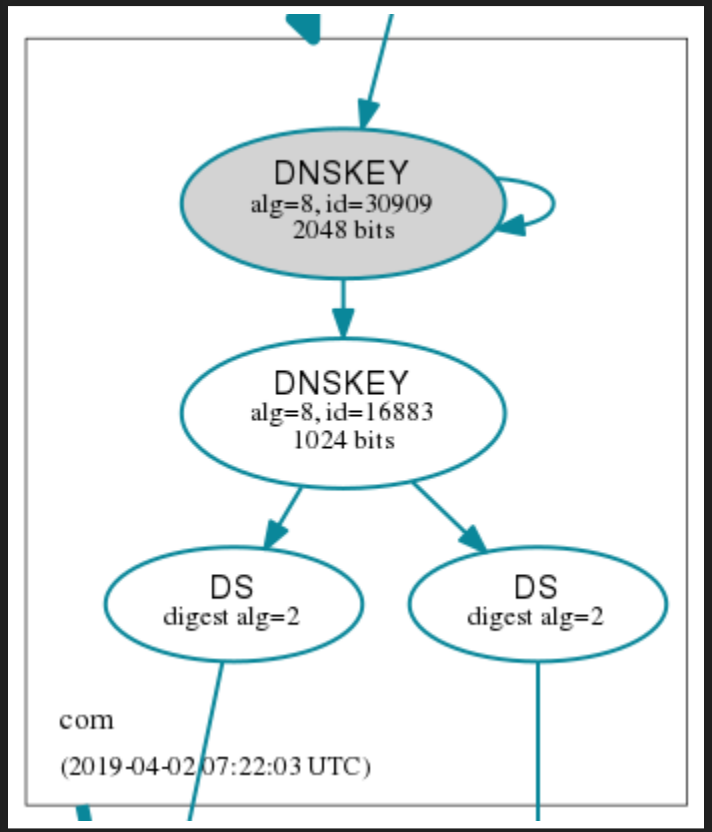
DNSKEY
alg=8, id=30909
2048 bits
DNSKEY
alg=8, id=16883
1024 bits
DS
digest alg=2
DS
digest alg=2
com
(2019-04-02 07:22:03 UTC)

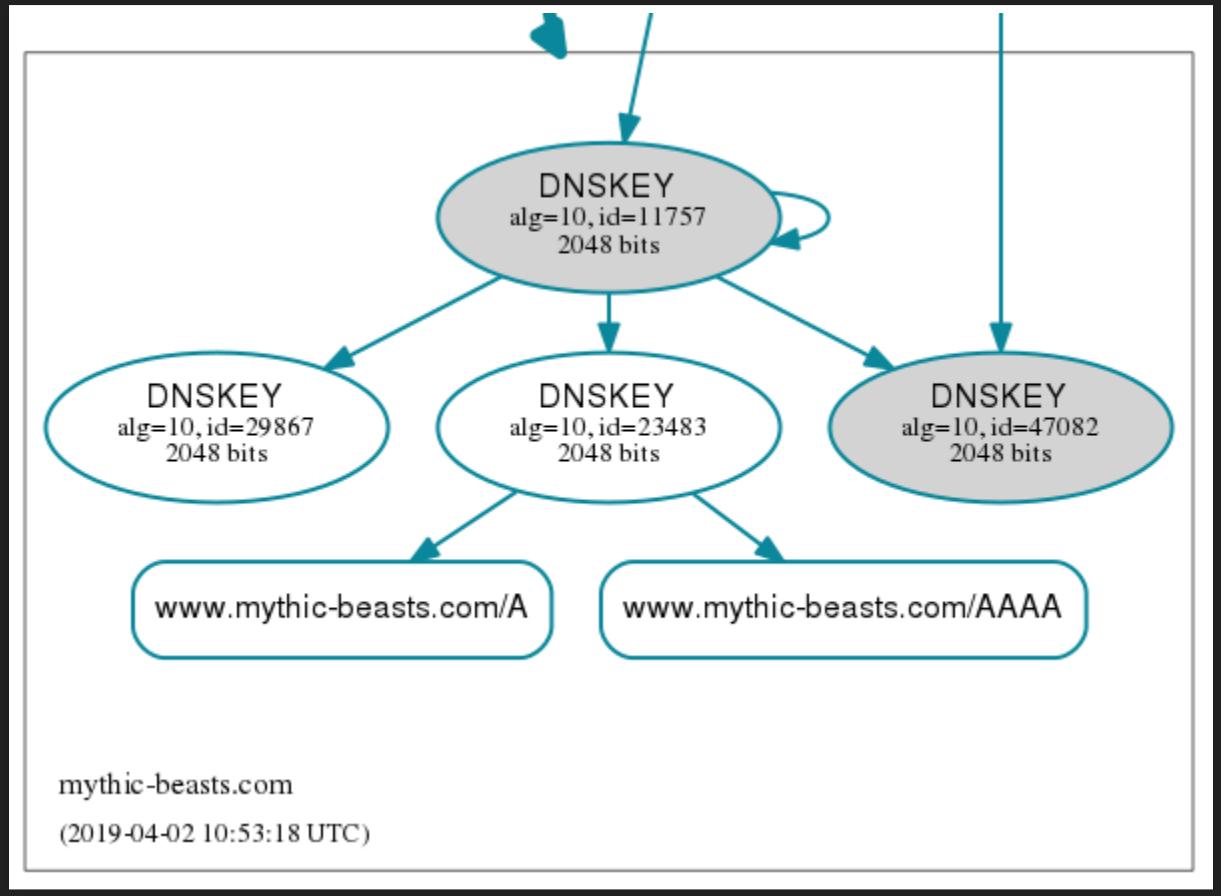
DNSKEY
alg=10, id=11757
2048 bits
DNSKEY
alg=10, id=29867
2048 bits
DNSKEY
alg=10, id=23483
2048 bits
DNSKEY
alg=10, id=47082
2048 bits
www.mythic-beasts.com/A
www.mythic-beasts.com/AAAA
mythic-beasts.com
(2019-04-02 10:53:18 UTC)

NEW DNSSEC RRS

- DNSKEY: A public key
- RRSIG: Signature of RR sets
- NSEC/NSEC3: Name existance
- DS: Digest of DNSKEY record on parent side of delegation

DNSSEC - BACK TO THE WIRE

- EDNS0 flag DO: Client groks DNSSEC.
- New main flags:
  - Authenticated Data (AD): Data is authenticated
  - Checking Disabled (CD): Client is OK to receive non-authenticated data

USING DNSSEC

- If your resolver does DNSSEC:
  - AD indicates data is authenticated
  - SERVFAIL if authentication fails

LAST MILE PROBLEM

LAST MILE PROBLEM

- Can your stub resolver validate?

LAST MILE PROBLEM

- Can your stub resolver validate?
- Can your resolving server validate?

LAST MILE PROBLEM

- Can your stub resolver validate?
- Can your resolving server validate?
- … and even if it can, can you trust the link between you and the resolving server?

LOCAL VALIDATION

- DNSSEC-trigger: https://www.nlnetlabs.nl/projects/dnssec-trigger/
- Stubby: https://getdnsapi.net/blog/dns-privacy-daemon-stubby/

DNSSEC AS PUBLIC KEY INFRASTRUCTURE

- IPSec keys (RFC4025)
- SSH host keys (RFC4255)
- Storing Certificates, CERT RR (RFC4398)
- DKIM keys (RFC4871)
- CA Authorisation (RFC6844)
- DNS Authentication of Named Entities (DANE), X.509 for TLS (RFC6698,7671)
- OpenPGP key (RFC7929)

IETF RESPONSE - TIMELINE

- 2013:
  - Snowden
- 2014:
  - RFC7285 Pervasive Monitoring is an Attack
  - DPRIVE Working Group formed - goals:
    - Encrypt Stub-Resolver DNS
    - Think about encrypting Resolver-Authoritative

DPRIVE

- 2015:
  - RFC7626 DNS Privacy Considerations
- 2016:
  - RFC7766 DNS over TCP
  - RFC7858 DNS over TLS

Internet Engineering Task Force (IETF)
Request for Comments: 7858
Category: Standards Track
ISSN: 2070-1721

Z.
L.
J. Heidem
USC/
A. Man
Independ
D. Wess
Verisign L
P. Hoff
IC

May 2

Specification for DNS over Transport Layer Security (TLS)

Abstract

This document describes the use of Transport Layer Security (TLS)
provide privacy for DNS. Encryption provided by TLS eliminates
opportunities for eavesdropping and on-path tampering with DNS
queries in the network, such as discussed in RFC 7626. In additio
this document specifies two usage profiles for DNS over TLS and
provides advice on performance considerations to minimize overhead
from using TCP and TLS with DNS.

DNS OVER TLS (DOT)

DNS over TCP, but using TLS and to port 853

DOT MODES

DOT MODES

DOT MODES

DOT MODES

DOT SUPPORT

- Clients: Android Pie, systemd, Stubby
  - Native Windows/macOS/iOS support still needed
- Servers: Unbound, Knot resolver, dnsdist, Bind via proxy
- November 2017: Quad9 public DNS (9.9.9.9)
- March 2018: Cloudflare public DNS (1.1.1.1)
- January 2019: Google public DNS (8.8.8.8)

Internet Engineering Task Force (IETF)
Request for Comments: 7816
Category: Experimental
ISSN: 2070-1721

S. Bortzme
AF

March 2

DNS Query Name Minimisation to Improve Privacy

Abstract

This document describes a technique to improve DNS privacy, a technique called "QNAME minimisation", where the DNS resolver no longer sends the full original QNAME to the upstream name server.

QNAME MINIMISATION

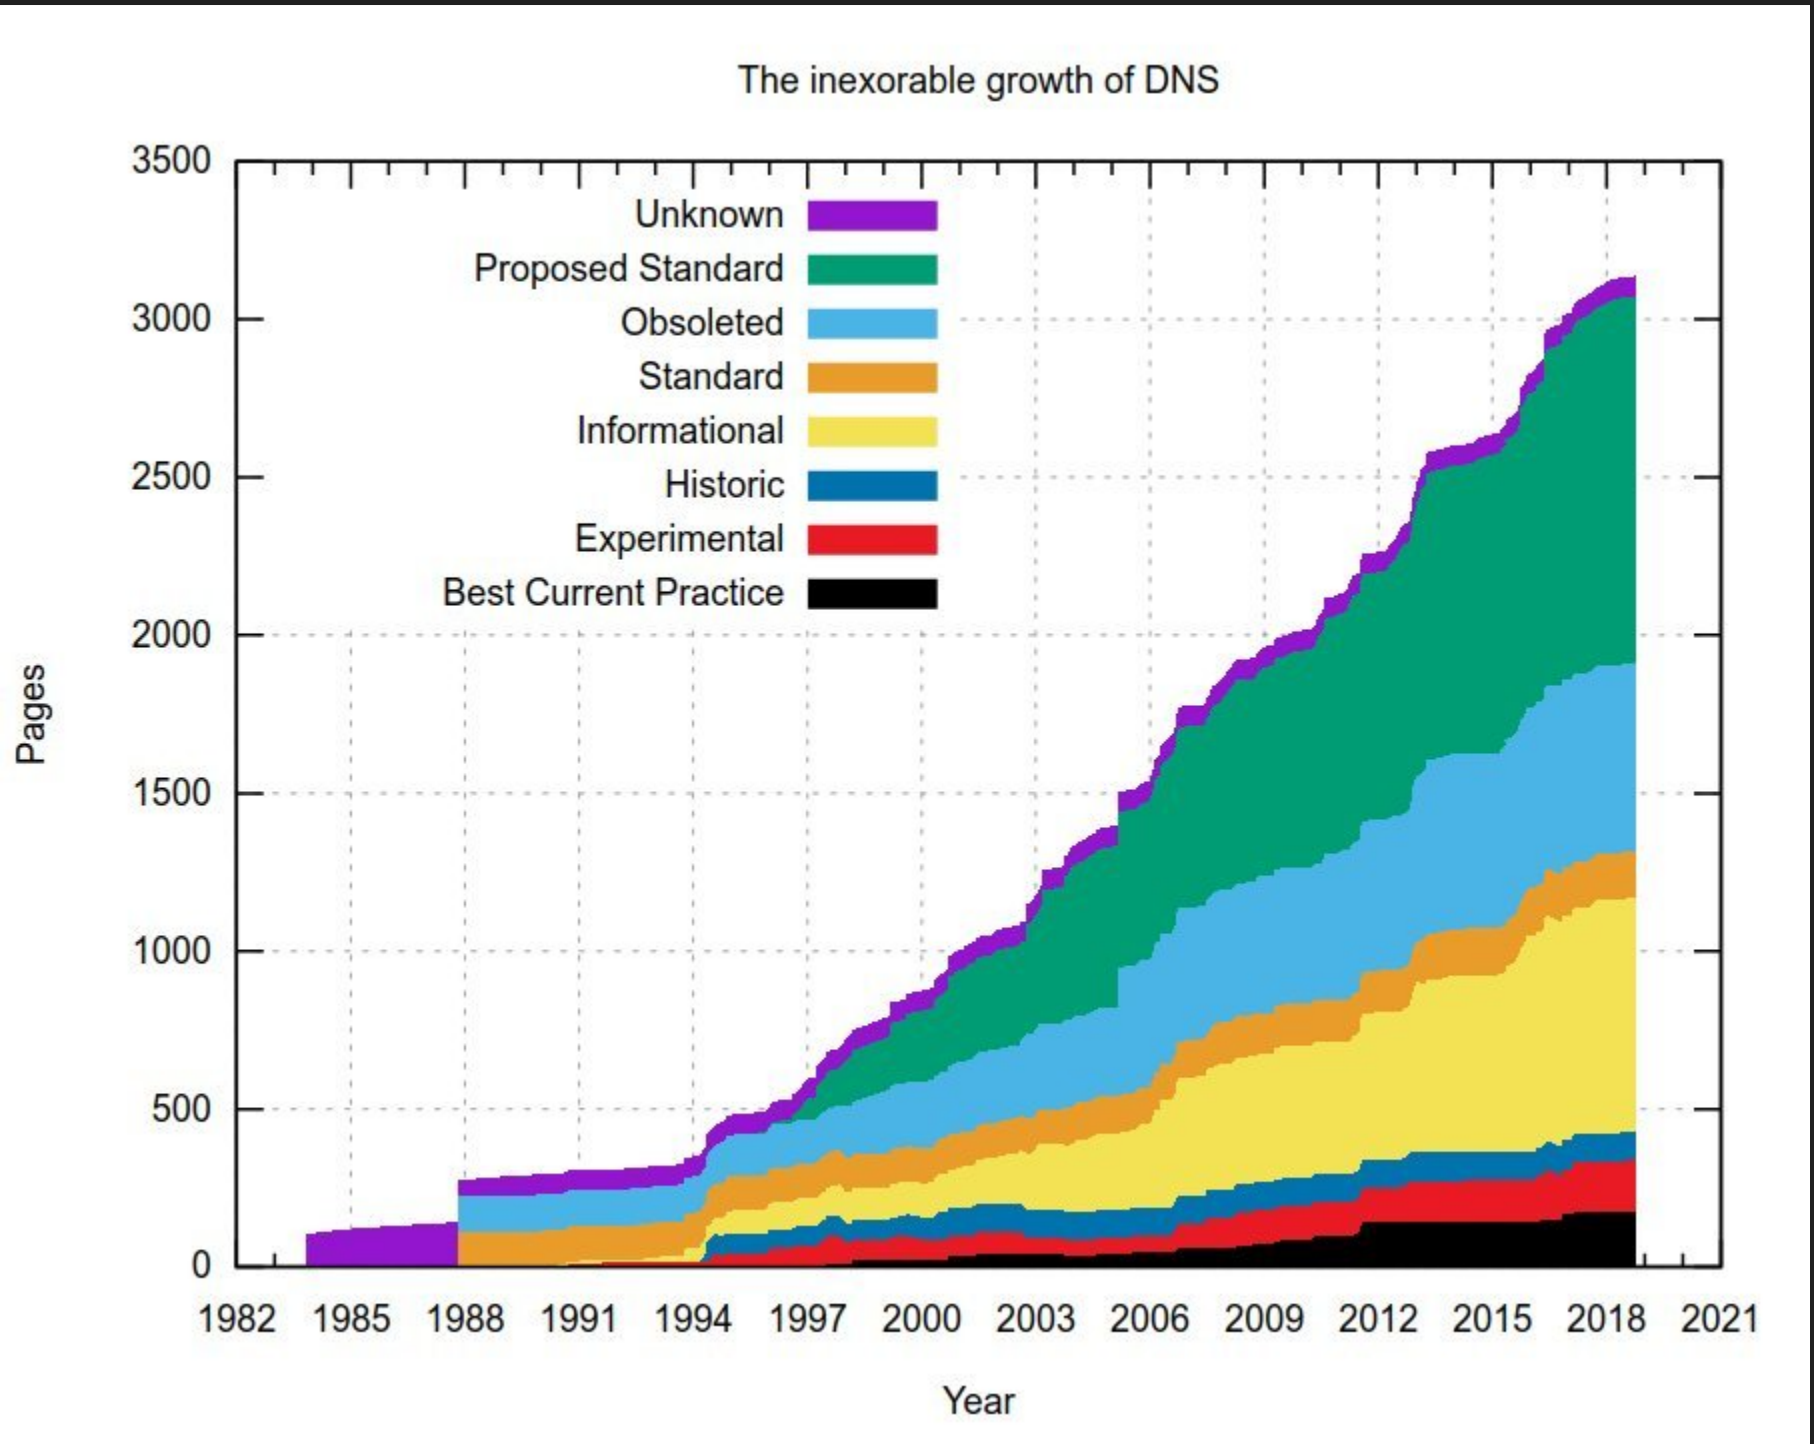
The inexorable growth of DNS
Unknown
Proposed Standard
Obsoleted
Standard
Informational
Historic
Experimental
Best Current Practice
Pages
Year
0
500
1000
1500
2000
2500
3000
3500
1982
1985
1988
1991
1994
1997
2000
2003
2006
2009
2012
2015
2018
2021

TACKLING THE CAMEL

https://powerdns.org/hello-dns/
https://powerdns.org/dns-camel/

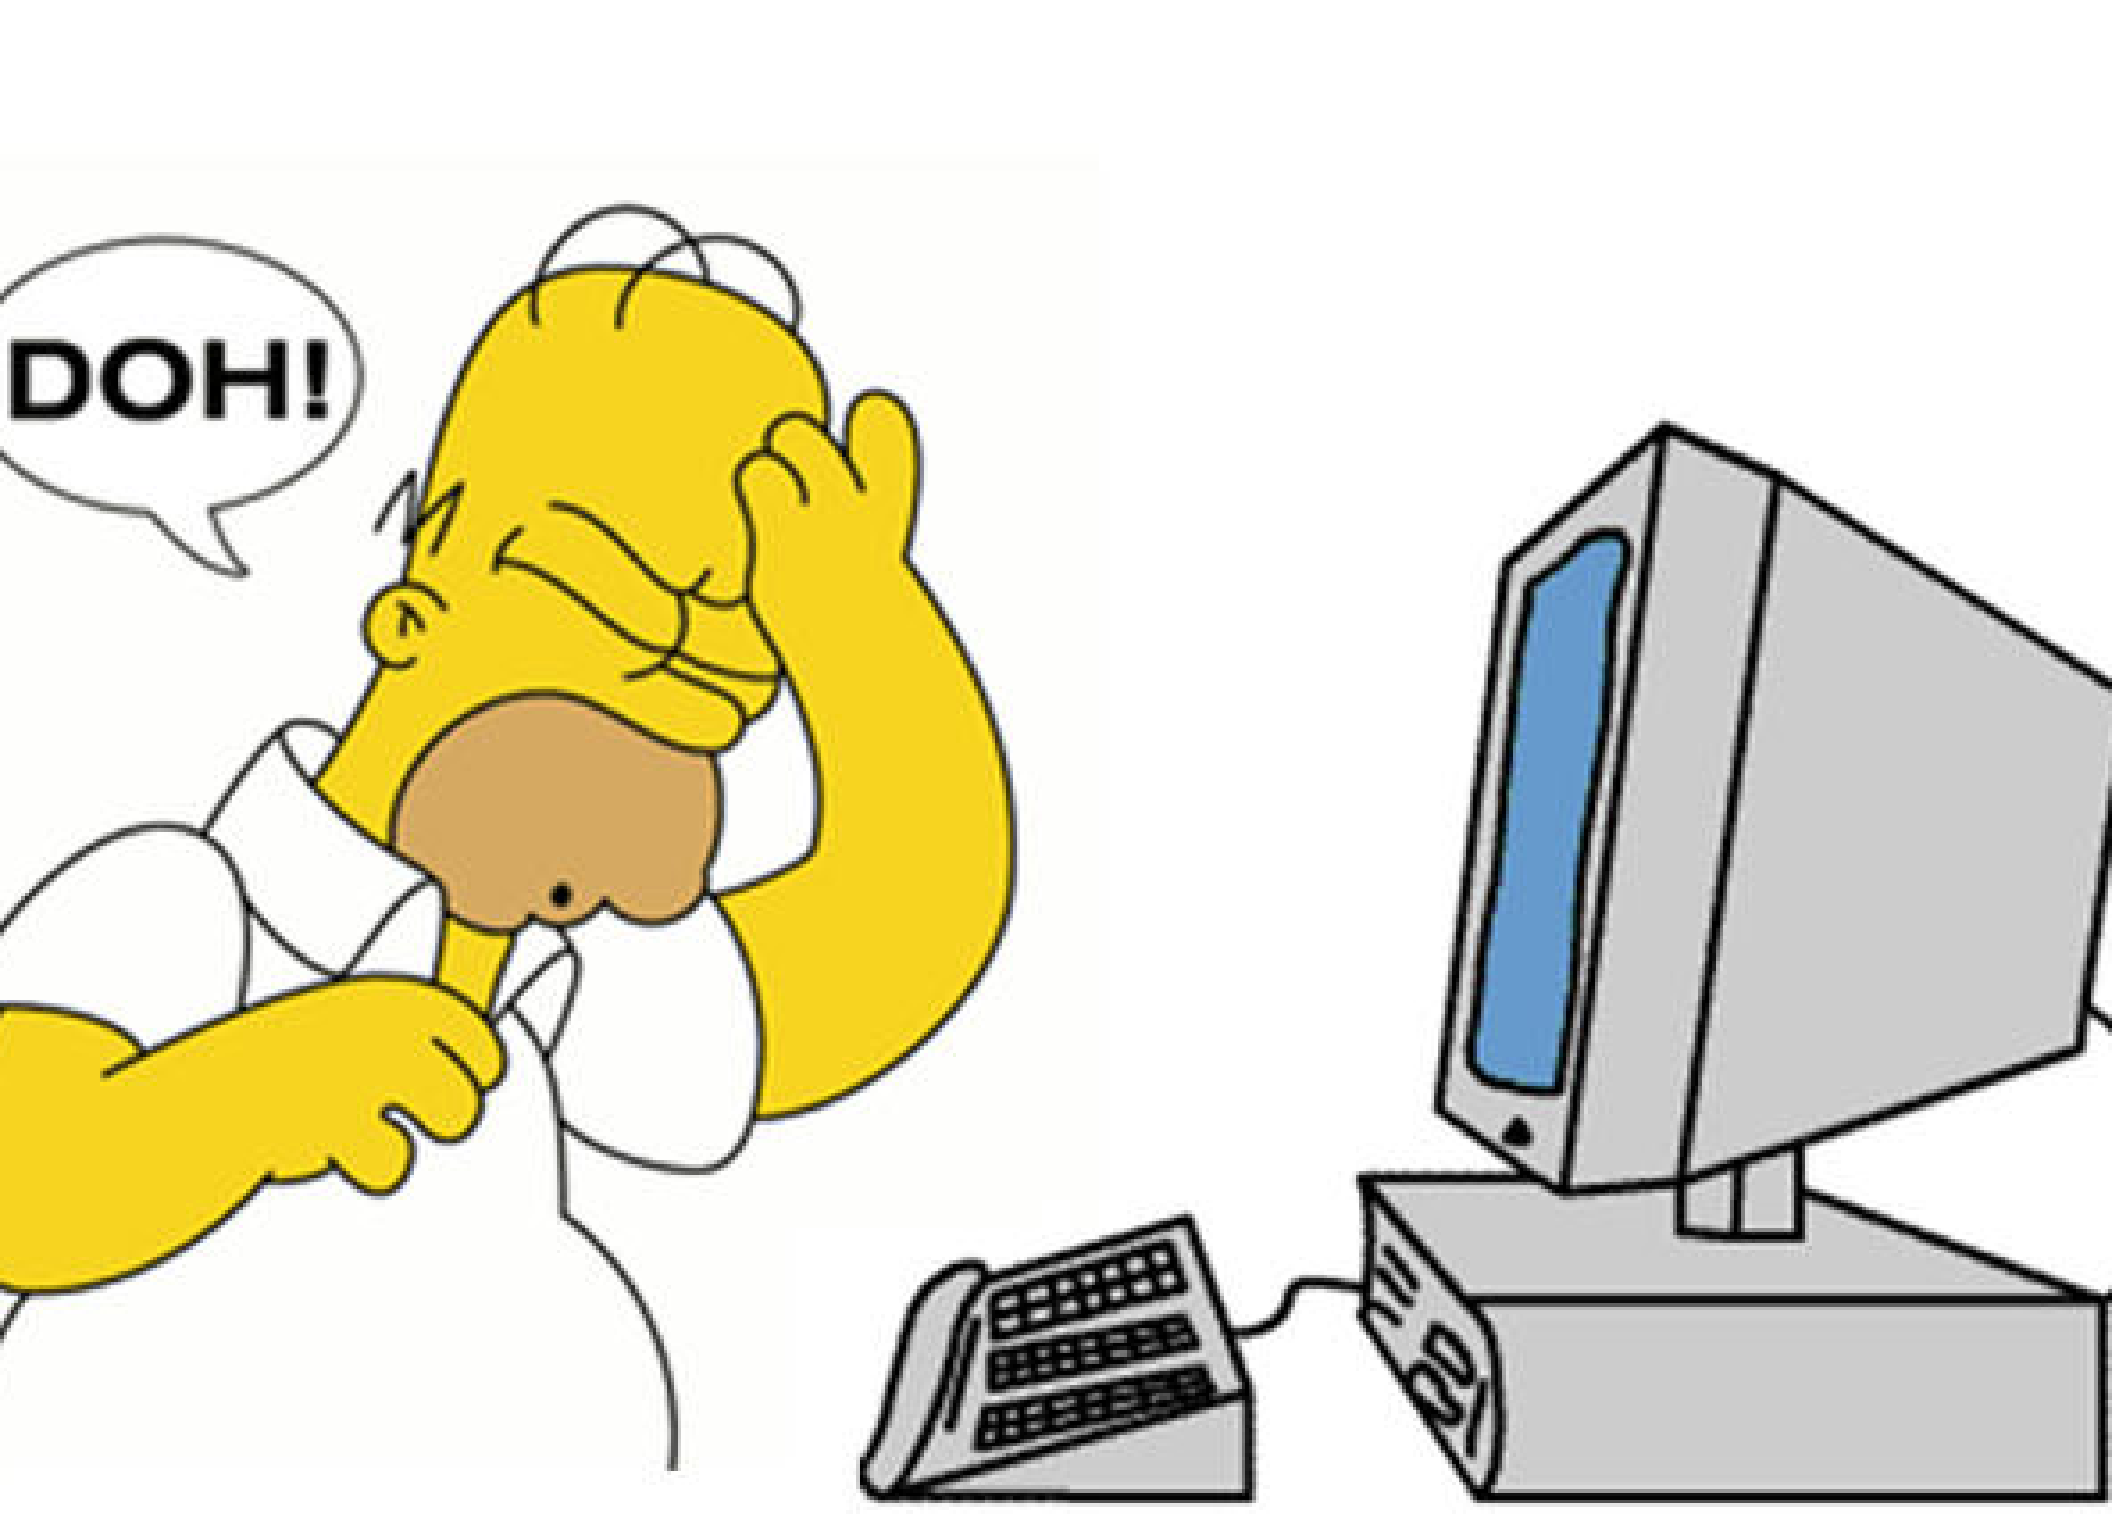
DOH!

DOH TIMELINE

- March 2017: Discussed at IETF 98
- May 2017: First draft published
- September 2017: DoH Working Group formed - goals:
  - Standardise encodings for DNS queries and responses that are suitable for use in HTTPS

DOH TIMELINE

- October 2017: DoH draft adopted by WG
- July 2018: Submitted to IESG
- August 2018: Approved
- October 2018: RFC8484 published

DNS OVER HTTPS

- Each DNS query/response is a HTTP exchange
- Must use `https` URI scheme
  - HTTP/2 is minimum recommended HTTP version
  - SHOULD use 0 in DNS ID
- Client configured via URI template (RFC6570)
  - `https://dnsserver.example.net/dns-query{?dns}`

DNS OVER HTTPS

- Defined `application/dns-message` media type
  - Same as the payload of a DNS UDP packet
  - Maximum size 65535
  - Door open to future definitions of alternate media types: DNS/JSON perhaps?
- HTTP cache control and DNS TTL need to be coordinated

DOH: HTTP GET QUERY

```
:method = GET
:scheme = https
:authority = dnsserver.example.net
:path = /dns-query?dns=AAABAAABAAAAAAAAA3d3dwdleGFtcGxlA2NvbQAAAQAB
accept = application/dns-message
```

- Query data is encoded in `base64url`.

DOH: HTTP POST QUERY

```
:method = POST
:scheme = https
:authority = dnsserver.example.net
:path = /dns-query
accept = application/dns-message
content-type = application/dns-message
content-length = 33

<33 bytes represented by the following hex encoding>
00 00 01 00 00 01 00 00  00 00 00 00 03 77 77 77
07 65 78 61 6d 70 6c 65  03 63 6f 6d 00 00 01 00
01
```

DOH: HTTP RESPONSE

```
:status = 200
content-type = application/dns-message
content-length = 61
cache-control = max-age=3709

<61 bytes represented by the following hex encoding>
00 00 81 80 00 01 00 01  00 00 00 00 03 77 77 77
07 65 78 61 6d 70 6c 65  03 63 6f 6d 00 00 1c 00
01 c0 0c 00 1c 00 01 00  00 0e 7d 00 10 20 01 0d
b8 ab cd 00 12 00 01 00  02 00 03 00 04
```

- Query: `IN` AAAA records for `www.example.com`
- Response: 1 answer record
  - Address of `2001:db8:abcd:12:1:2:3:4`
  - TTL of 3709s (0xe7d)

DOH: COMPARISON WITH DOT

DOH: COMPARISON WITH DOT

- One use case: "Allow web applications to to access DNS information via existing browser APIs"

DOH: COMPARISON WITH DOT

- One use case: "Allow web applications to to access DNS information via existing browser APIs"
- Discovery: MUST use URI template
  - So no **Opportunistic**

DOH: COMPARISON WITH DOT

- One use case: "Allow web applications to to access DNS information via existing browser APIs"
- Discovery: MUST use URI template
  - So no **Opportunistic**
- *Increased* tracking potential via HTTP headers (`User-Agent`, language, etc.)?

DOH: COMPARISON WITH DOT

- One use case: "Allow web applications to to access DNS information via existing browser APIs"
- Discovery: MUST use URI template
  - So no **Opportunistic**
- *Increased* tracking potential via HTTP headers (`User-Agent`, language, etc.)?
- New privacy concerns

DOH: CONNECTION MODELS

DOH: CONNECTION MODELS

- *Dedicated*: DoH traffic only

DOH: CONNECTION MODELS

- *Dedicated*: DoH traffic only
- *Mixed*: DoH traffic mixed with other HTTPS traffic

DOH: CONNECTION MODELS

- *Dedicated*: DoH traffic only
- *Mixed*: DoH traffic mixed with other HTTPS traffic
  - Better privacy?

DOH: CONNECTION MODELS

- *Dedicated*: DoH traffic only
- *Mixed*: DoH traffic mixed with other HTTPS traffic
  - Better privacy?
  - Impossible to block just DNS traffic

DOH: CONNECTION MODELS

- *Dedicated*: DoH traffic only
- *Mixed*: DoH traffic mixed with other HTTPS traffic
  - Better privacy?
  - Impossible to block just DNS traffic
  - **THE** big differentiator with DoT

DOH: SERVER DEPLOYMENT STATUS

DOH: SERVER DEPLOYMENT STATUS

- Large scale:
  - Cloudflare https://cloudflare-dns.com/dns-query
  - Google https://dns.google.com/experimental
  - Quad9 https://dns*.quad9.net/dns-query (3 flavours of service)

DOH: SERVER DEPLOYMENT STATUS

- Large scale:
  - Cloudflare https://cloudflare-dns.com/dns-query
  - Google https://dns.google.com/experimental
  - Quad9 https://dns*.quad9.net/dns-query (3 flavours of service)
- ~12 other test servers

DOH: CLIENT STATUS

- Firefox
- Chrome (Chromium, Bromite)
- curl
- *Intra* Android app
- `cloudflared`
- Various experimental
- GetDNS/Stubby in progress

DOH: SERVER IMPLEMENTATIONS

- `dnsdist` load balancer
- Knot resolver (branch)

DOH IN BROWSERS

- OSs are slow to offer new DNS features
- “We care about the privacy of our users”
- “Reduced latency within the browser”

WHY DOH NOT DOT - MOZILLA

WHY DOH NOT DOT - MOZILLA

- Integration: "leverage the HTTP ecosystem"

WHY DOH NOT DOT - MOZILLA

- Integration: "leverage the HTTP ecosystem"
- HTTPS everywhere: “it works … just use port 443, mix traffic”

WHY DOH NOT DOT - MOZILLA

- Integration: "leverage the HTTP ecosystem"
- HTTPS everywhere: “it works … just use port 443, mix traffic”
- Cool stuff:
  - JSON
  - Server push
  - Get DNS from location other than configured resolver

‘MOZIFLARE’

‘MOZIFLARE’

- “We’d like to turn this on for all our users”

‘MOZIFLARE’

- “We’d like to turn this on for all our users”
- “Cloudflare is our Trusted Recursive Resolver (TRR)”

‘MOZIFLARE’

- “We’d like to turn this on for all our users”
- “Cloudflare is our Trusted Recursive Resolver (TRR)”
- “.., we have a resolver we can trust to protect our users' privacy. This means **Firefox can ignore the resolver that the network provides** and just go straight to CloudFlare”

‘MOZIFLARE’ CONT.

‘MOZIFLARE’ CONT.

‘MOZIFLARE’ CONT.

‘MOZIFLARE’ CONT.

7 Layers of the OSI Model

Application	• End User layer • HTTP, FTP, IRC, SSH, DNS
Presentation	• Syntax layer • SSL, SSH, IMAP, FTP, MPEG, JPEG
Session	• Synch & send to port • API's, Sockets, WinSock
Transport	• End-to-end connections • TCP, UDP
Network	• Packets • IP, ICMP, IPSec, IGMP
Data Link	• Frames • Ethernet, PPP, Switch, Bridge
Physical	• Physical structure • Coax, Fiber, Wireless, Hubs, Repeaters

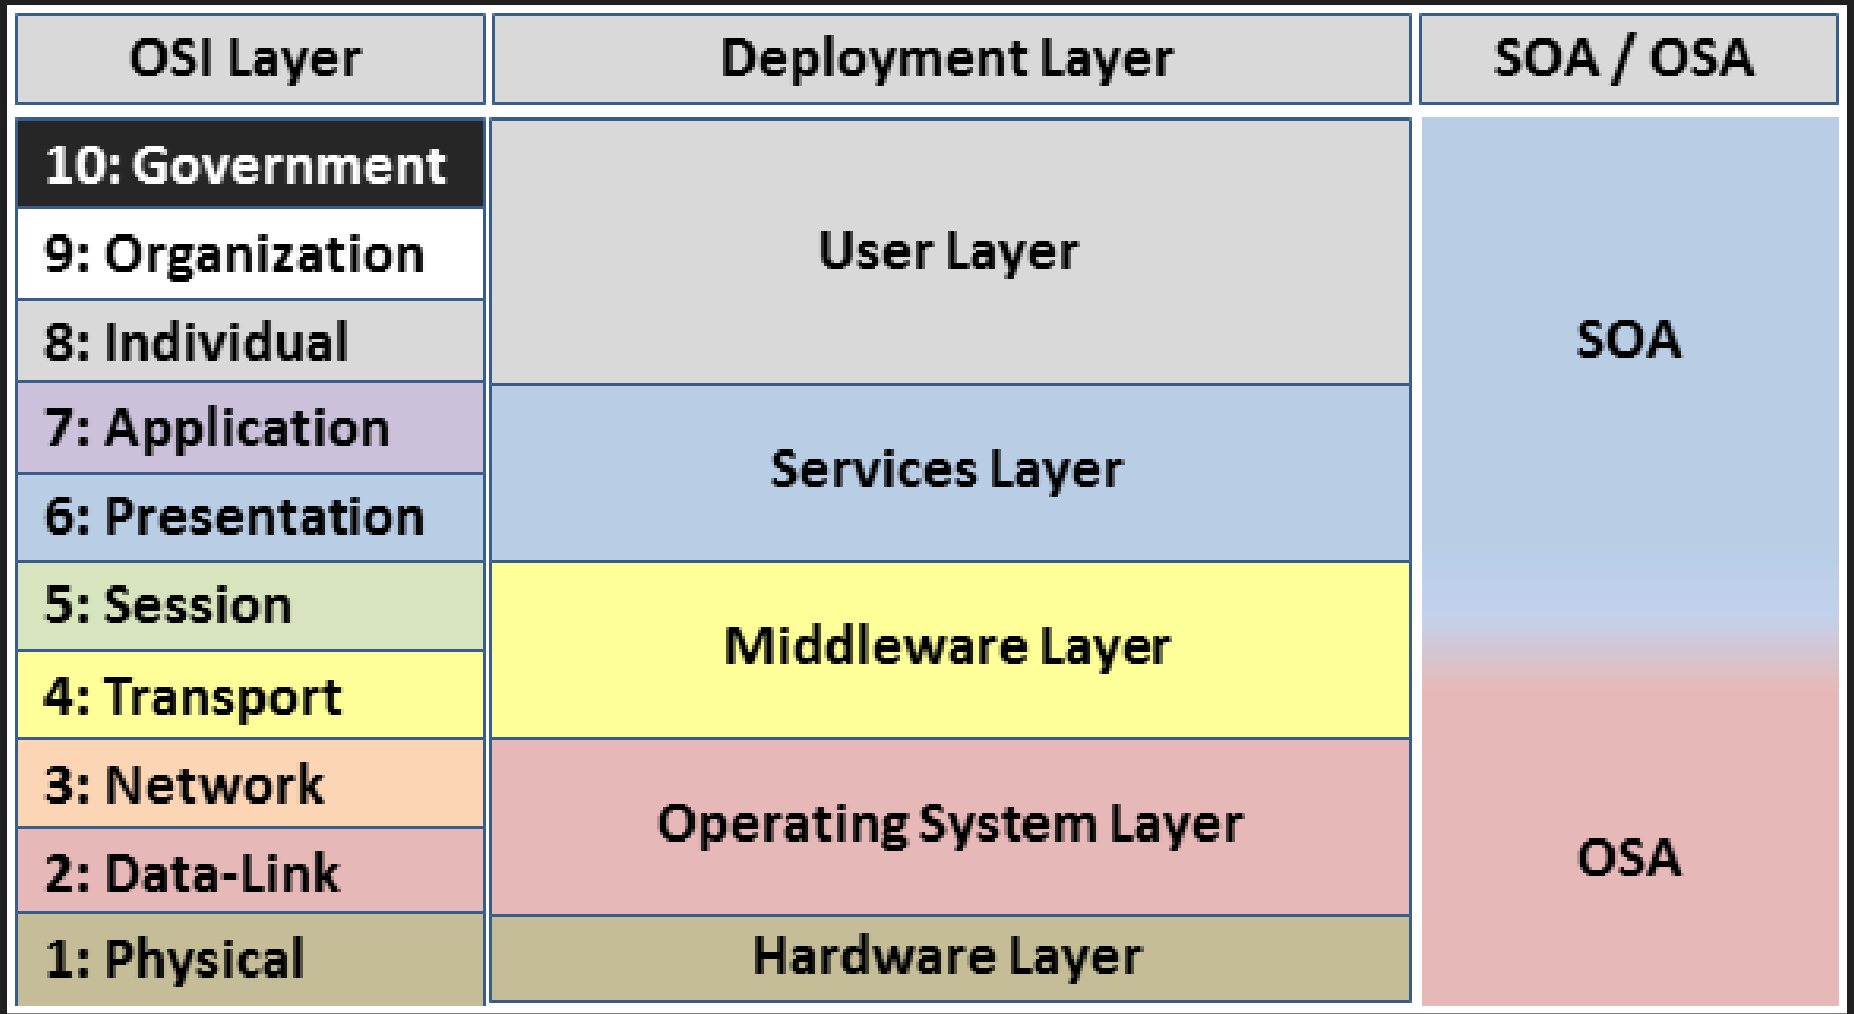

WILL THIS BE THE 5 MINUTE ARGUMENT?

INDIVIDUAL LAYER

INDIVIDUAL LAYER

- Split between system and browser resolving
  - Home router naming
  - VPN naming

INDIVIDUAL LAYER

- Split between system and browser resolving
  - Home router naming
  - VPN naming
- Configure DNS for each application?

INDIVIDUAL LAYER

- Split between system and browser resolving
  - Home router naming
  - VPN naming
- Configure DNS for each application?
- Breaks parental control service

INDIVIDUAL LAYER

- Split between system and browser resolving
  - Home router naming
  - VPN naming
- Configure DNS for each application?
- Breaks parental control service
- Informed consent

INDIVIDUAL LAYER

- Split between system and browser resolving
  - Home router naming
  - VPN naming
- Configure DNS for each application?
- Breaks parental control service
- Informed consent
- What is best choice for user?

Contract with TalkTalk is based in the same legal juristiction, and TalkTalk are subject to GDPR. Regulatory environment for handling of privacy data is understood. Cloudflare's privacy policy appears satisfactory, but Cloudflare is a US corporation, so subject to different regulatory regime, with laxer requirements.

TalkTalk 2015 data breach is compelling evidence that TalkTalk isn't a safe host for privacy-related data. Cloudflare's record is not spotless, but on balance they are more trustworthy than TalkTalk.

… OR THE FULL HALF HOUR?

ORGANISATION LAYER

ORGANISATION LAYER

- Split-horizon DNS

ORGANISATION LAYER

- Split-horizon DNS
- Local content caches

ORGANISATION LAYER

- Split-horizon DNS
- Local content caches
- Service support

ORGANISATION LAYER

- Split-horizon DNS
- Local content caches
- Service support
- Organisation does not regard its own network as belonging to attacker

ORGANISATION LAYER

- Split-horizon DNS
- Local content caches
- Service support
- Organisation does not regard its own network as belonging to attacker
- “My network, my rules”

ORGANISATION LAYER

- Split-horizon DNS
- Local content caches
- Service support
- Organisation does not regard its own network as belonging to attacker
- “My network, my rules”
  - Though if org is an ISP, do customers have a choice of ISP?

GOVERNMENT LAYER

GOVERNMENT LAYER

- Filtering banned content using DNS

GOVERNMENT LAYER

- Filtering banned content using DNS
- Malware detection and mitigation

GOVERNMENT LAYER

- Filtering banned content using DNS
- Malware detection and mitigation
- There are valid reasons organisations need **some** visibility on their DNS lookups

RELIGIOUS LAYER

RELIGIOUS LAYER

- Will DNS resolving go the way of email?

RELIGIOUS LAYER

- Will DNS resolving go the way of email?
- Internet future:

RELIGIOUS LAYER

- Will DNS resolving go the way of email?
- Internet future:
  - Are we moving inexorably towards an internet totally reliant on a few big corporations?

RELIGIOUS LAYER

- Will DNS resolving go the way of email?
- Internet future:
  - Are we moving inexorably towards an internet totally reliant on a few big corporations?
  - Are we heading for an internet where **everything** runs on HTTPS to port 443?

DNS PRIVACY NOW

- DoT via system
- Opportunistic or Strict to organisation’s chosen resolver

FIN

https://dnsprivacy.org/