Advanced Computer Networking

CYBR 230 – Jeff Shafer – University of the Pacific

DNS

DNSSEC, DNS over TLS, DNS over HTTPS

Motivation

- IP addresses are hard to remember
  - `138.9.110.12`? Or was it `.21`?
- Human-friendly names are much better
  - `engineering.pacific.edu`
- How can we translate between the two?

Early Days (prior to 1984)

- Each computer on the ARPAnet (early Internet) had a single file
  - `hosts.txt` maps all known host names to IP address
- Master list maintained by SRI Network Information Center
  - Email them if your mapping changes
  - New list produced 1-2 times a week
  - All hosts download the new list
- **Problems with this approach?**

Domain Name System (DNS)

- **Distributed database** implemented in hierarchy of many **name servers**
- **Application-layer protocol**
  - Hosts, routers, and name servers communicate to resolve names (address/name translation)
  - Core Internet function, implemented as application-layer protocol
  - Complexity at network's "edge"

DNS is Decentralized

- No single point of failure
- No distant centralized database
- Easier maintenance
  - Take one or a dozen servers offline without issue
- Support high traffic volume
- *** **Scalability** ***

How many DNS requests/second globally?

DNS: Scalability

- **Challenging to find data on global DNS requests/sec**
  - No global internet "dashboard"
  - Internet is a "network of networks"
- Would have to inquire with AT&T, Comcast, TimeWarner, Pacific, etc
  - They would have to check stats on all of their local servers
- **Google Public DNS**
  - 1+ trillion requests/day as of August 2018
  - https://security.googleblog.com/2018/08/google-public-dns-turns-8888-years-old.html
- **OpenDNS**
  - 160 billion requests/day as of October 2018
  - http://system.opendns.com/

What's in a Name?

- `engineering.pacific.edu`
  - .edu is top-level domain
  - "pacific" belongs to .edu
  - "engineering" belongs to "pacific"
  - Hierarchical! Read from right to left
- Limits?
  - Up to 127 levels of hierarchy
  - Each label can have up to 63 characters
  - Full domain name cannot exceed 253 characters

DNS: Services

- Hostname to IP address translation
  - *"www.pacific.edu" is 138.9.110.12*
- Hostname aliasing
  - Canonical, alias names
- Hostname load distribution
  - Replicated servers – Multiple IP addresses available for one name
  - *"google.com" is 74.125.239.128 or 74.125.239.135 or … or …. or … or ….*

DNS: Services

- Mail server aliasing
  - What are the **multiple** host names that receive mail for this domain?
  - 1st priority, then 2nd backup, then 3rd backup, etc…
  - Allows you to use 3rd party email services (e.g. Google Apps)
  - *Mail to "pacific.edu" is directed to "d73442a.ess.barracudanetworks.com" (SPAM filtering)*
- Other / Misc
  - SPF entries for email (Anti-spam)
  - DNSSEC (security/encryption)
  - Many other attributes…

DNS: Record Types (Distributed Database)

Resource Record (RR) format: **`(name, value, type, ttl)`**

- Type=**A**
  - *name* is **hostname**
  - *value* is **IP address**
- Type=**NS**
  - *name* is domain (e.g. foo.com)
  - *value* is **hostname** of **authoritative name server** for this domain
- Type=**CNAME**
  - *name* is alias name for some "canonical" (real) name
  - value is canonical name
- Type=**MX**
  - *value* is name of **mailserver** associated with name
- Type=**TXT**
  - *value* is machine readable text (arbitrary)

DNS: Example

```
$ dig pacific.edu any

; <<>> DiG 9.8.3-P1 <<>> pacific.edu any
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 5270
;; flags: qr rd ra; QUERY: 1, ANSWER: 9, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;pacific.edu.                   IN   ANY

;; ANSWER SECTION:
pacific.edu.        59   IN   A    138.9.110.12
pacific.edu.        21599    IN   NS   ns-110.awsdns-13.com.
pacific.edu.        21599    IN   NS   ns-1289.awsdns-33.org.
pacific.edu.        21599    IN   NS   ns-2044.awsdns-63.co.uk.
pacific.edu.        21599    IN   NS   ns-705.awsdns-24.net.
pacific.edu.        899  IN   SOA  ns-110.awsdns-13.com. awsdns-
hostmaster.amazon.com. 1 7200 900 1209600 86400
pacific.edu.        299  IN   MX   10 d73442a.ess.barracudanetworks.com.
pacific.edu.        299  IN   MX   10 d73442b.ess.barracudanetworks.com.
pacific.edu.        299  IN   TXT  "v=spf1 ip4:138.9.240.95 ip4:138.9.110.64
ip4:138.9.110.74 include:_spf.google.com
include:spf.protection.outlook.com include:_spf.qualtrics.com ~all"
```

Resource Record Type

Resource Record Value

DNS: Name Resolution

- Two types
- **Recursive**
  - The server you contact provides the final answer
  - *Behind the scenes, it may make several consecutive requests*
- **Iterative**
  - The server you contact directs you to a different server to get (closer to) the final answer

en.wikipedia.org = 208.80.154.224

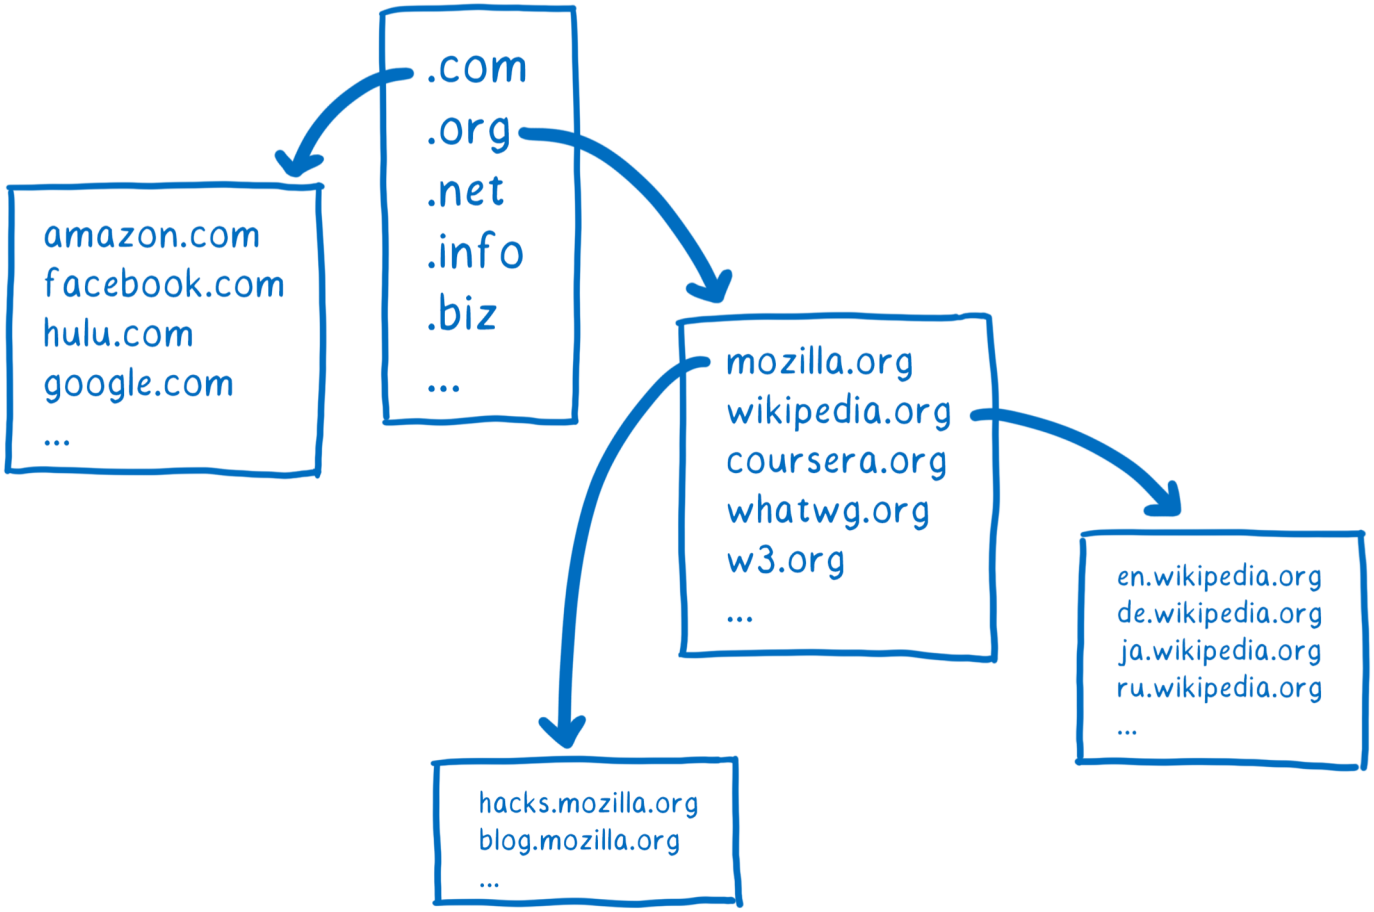

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

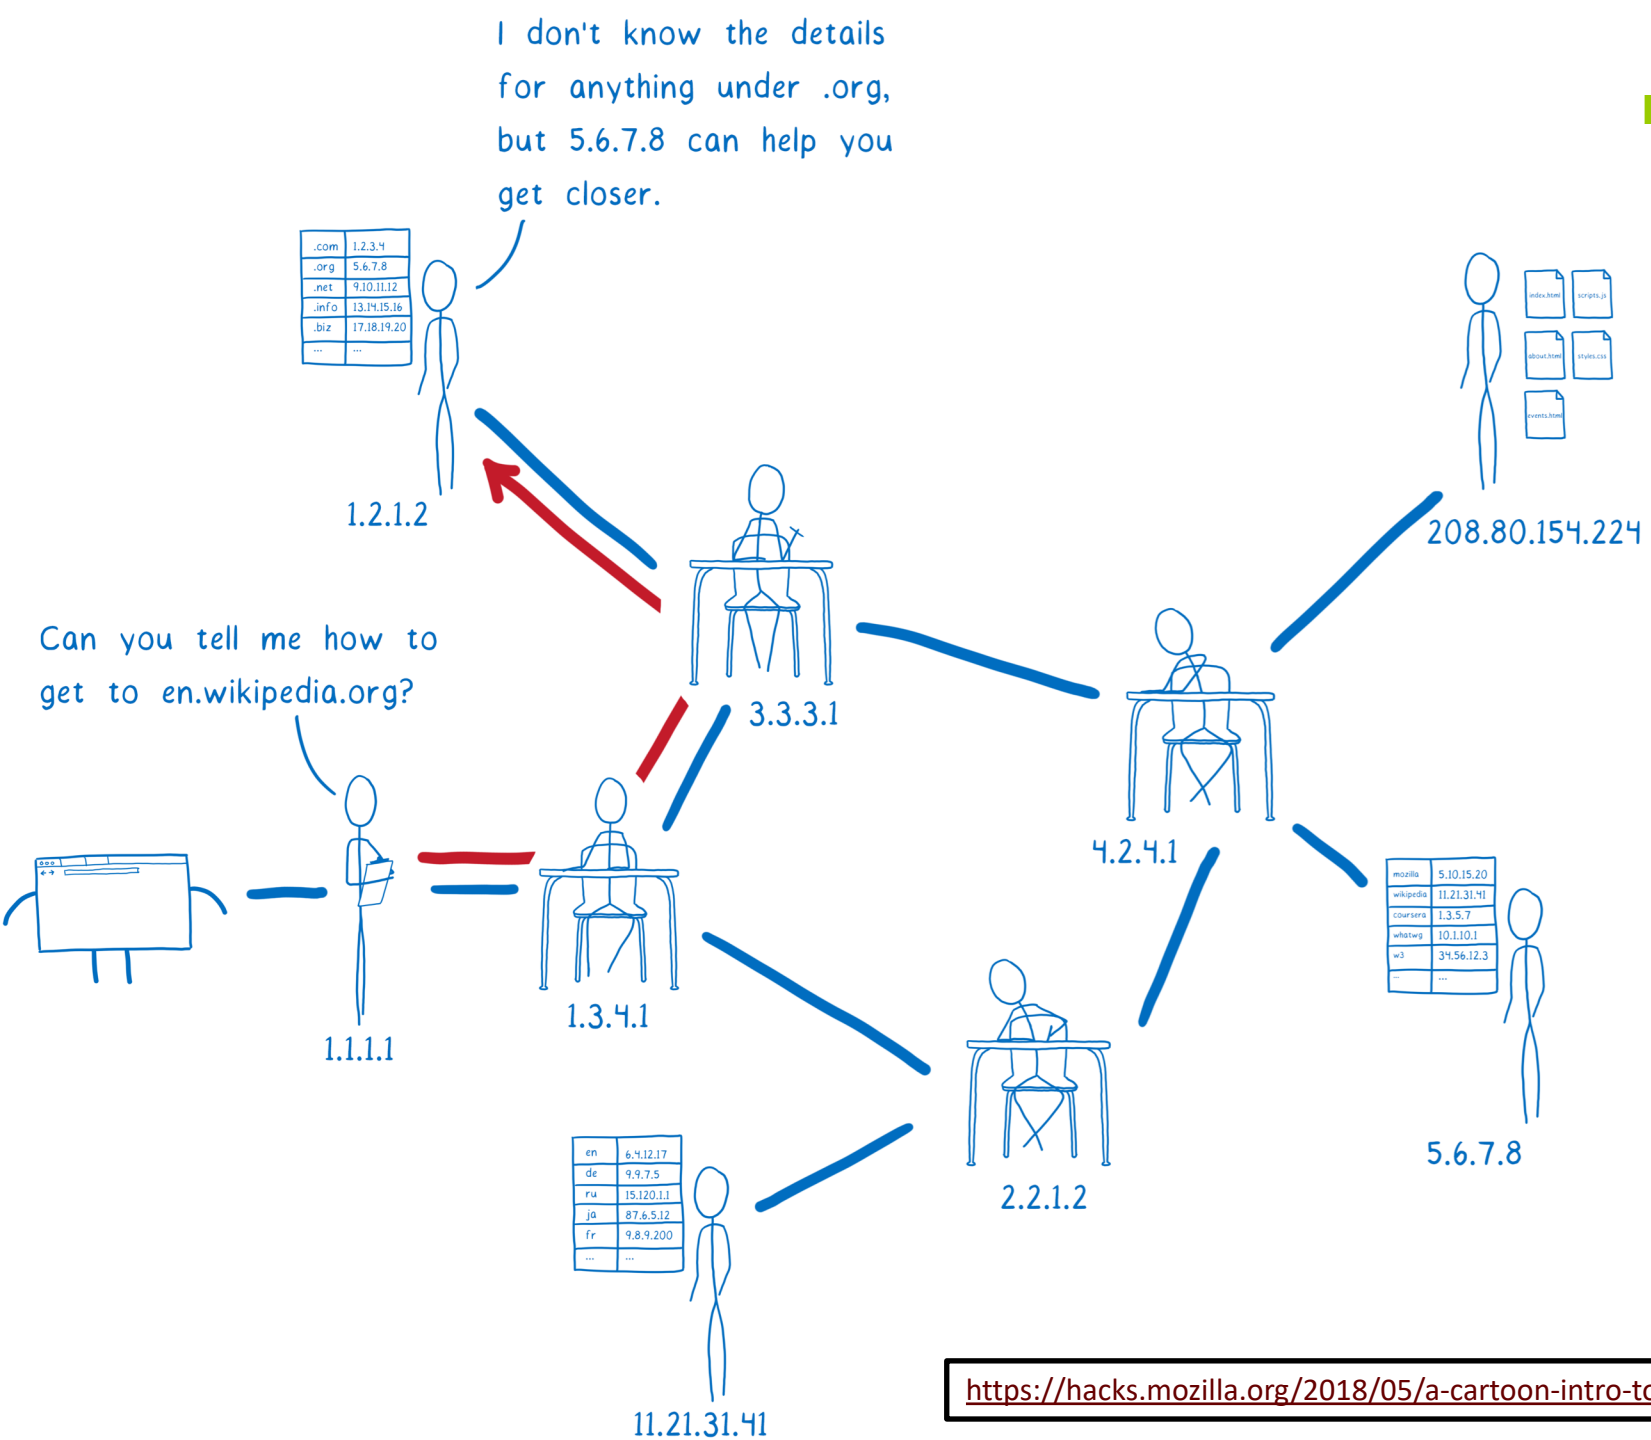

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

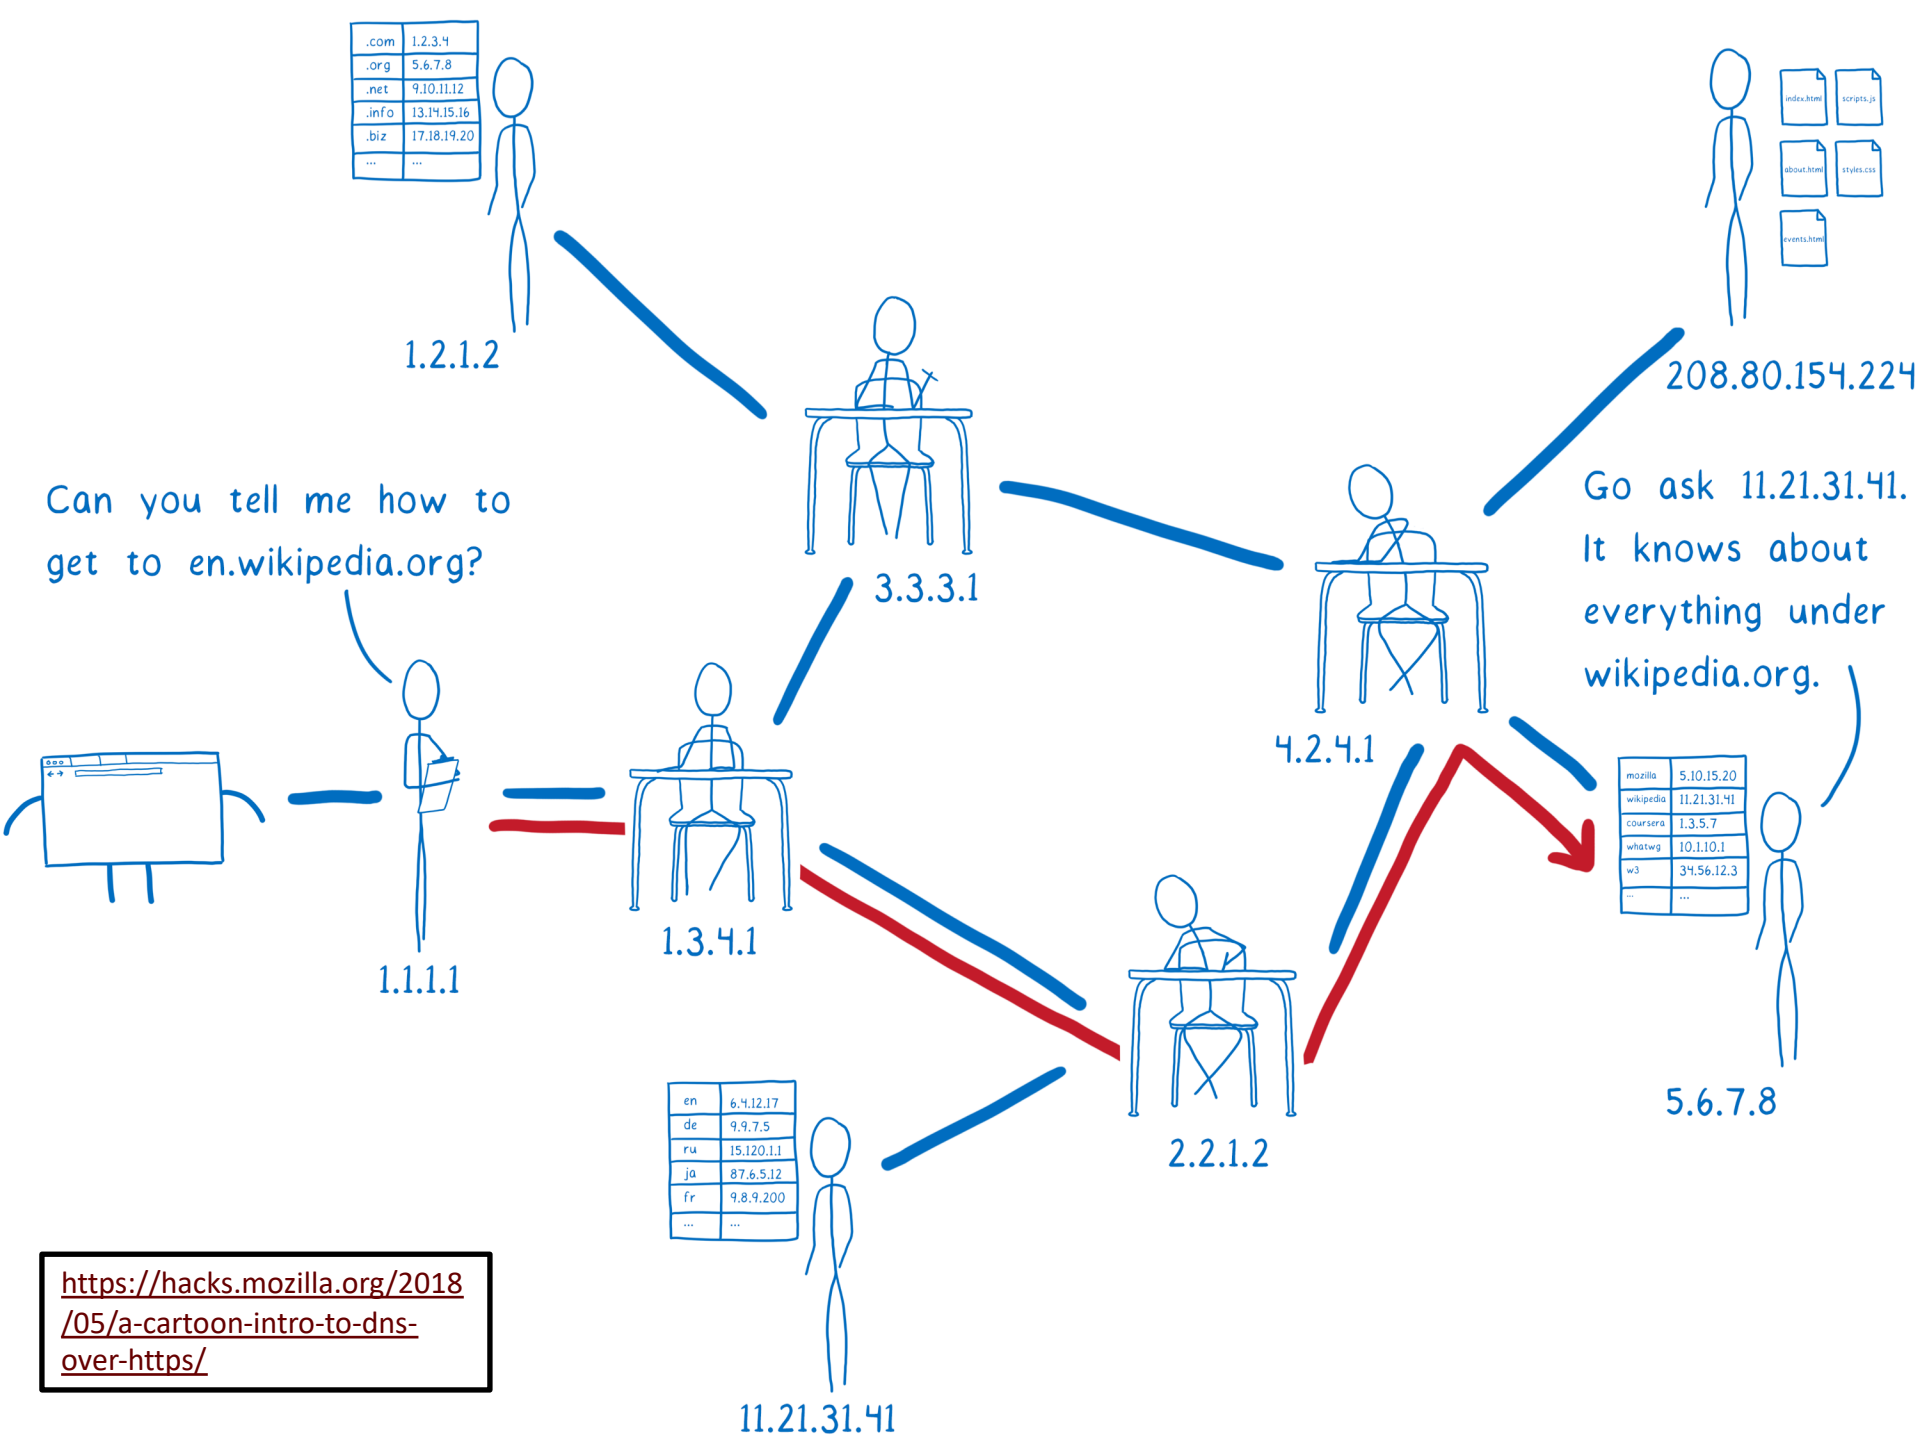

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

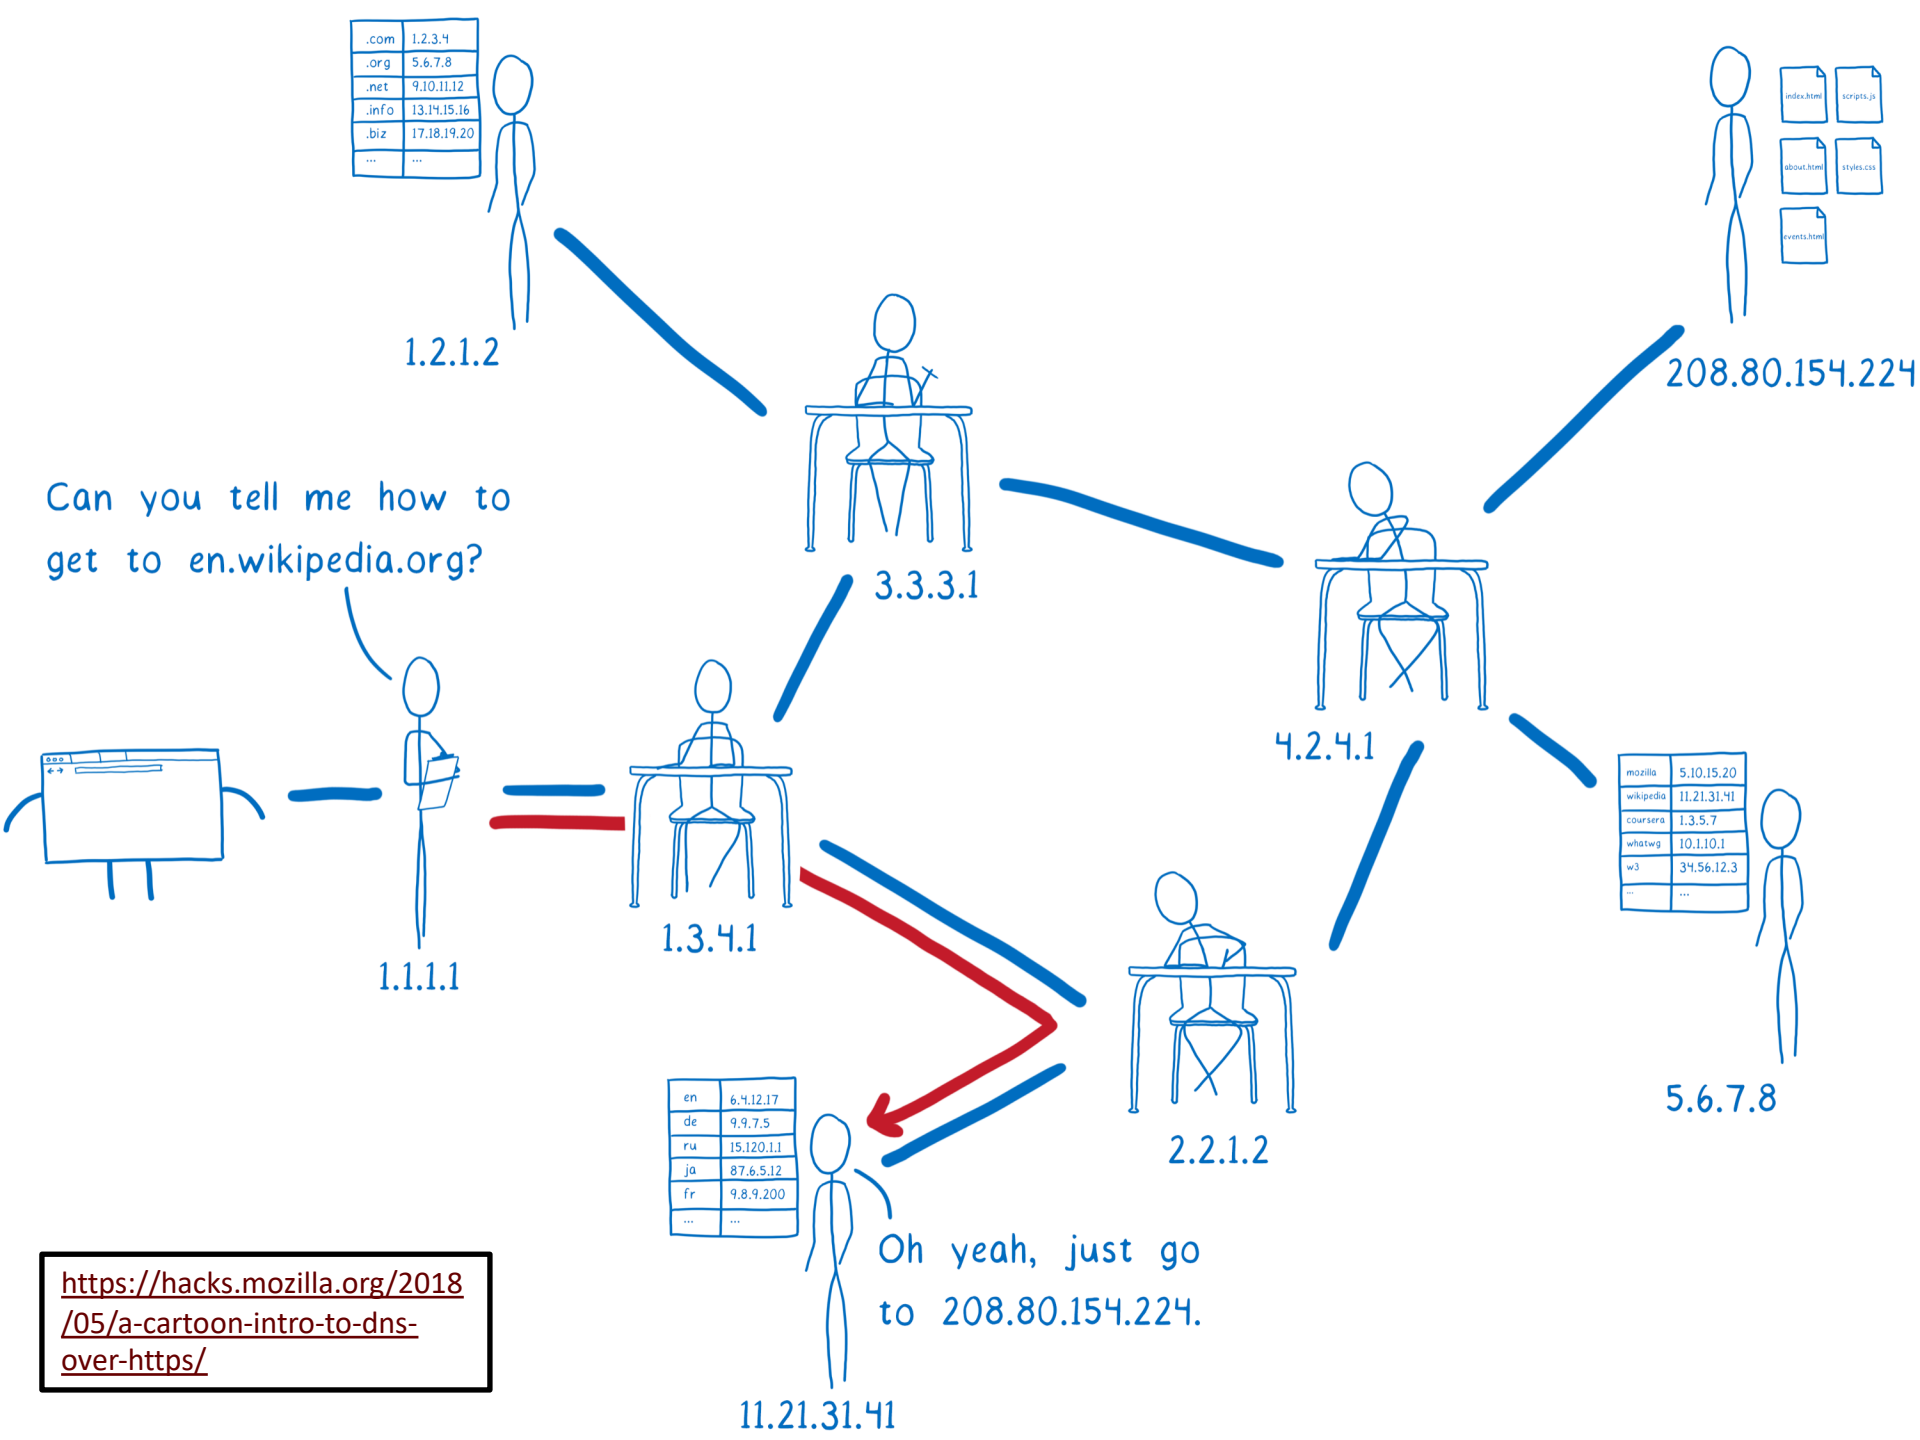

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

DNS: Root Name Servers

- Contacted by local name server that can not resolve top-level domain
- Root name server:
  - Contacts authoritative name server for TLD if name mapping not known
  - Gets mapping
  - Returns mapping to local name server

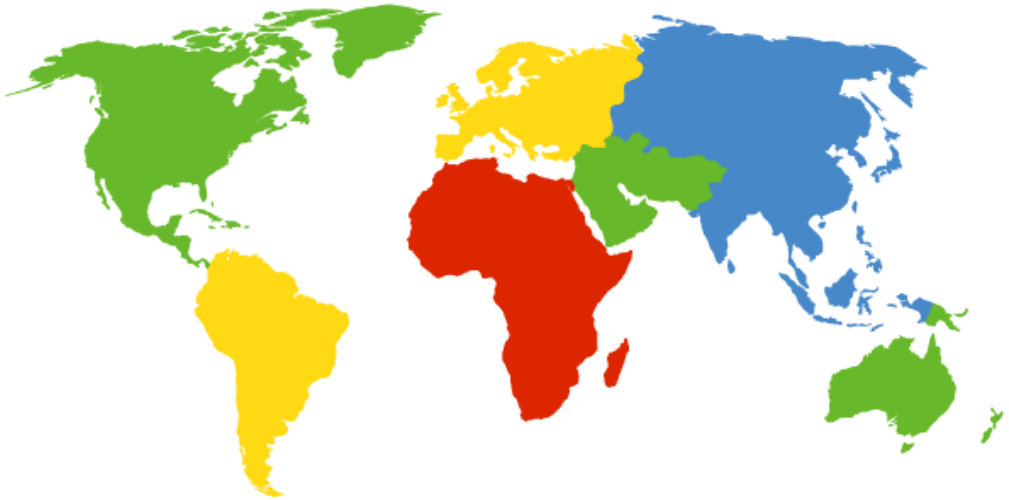

13 root name "servers" worldwide labeled `a` - `m`

- Each "server" is really a cluster
- Some clusters are geographically distributed
- 504 total in Fall 2014

DNS: Root Name Servers

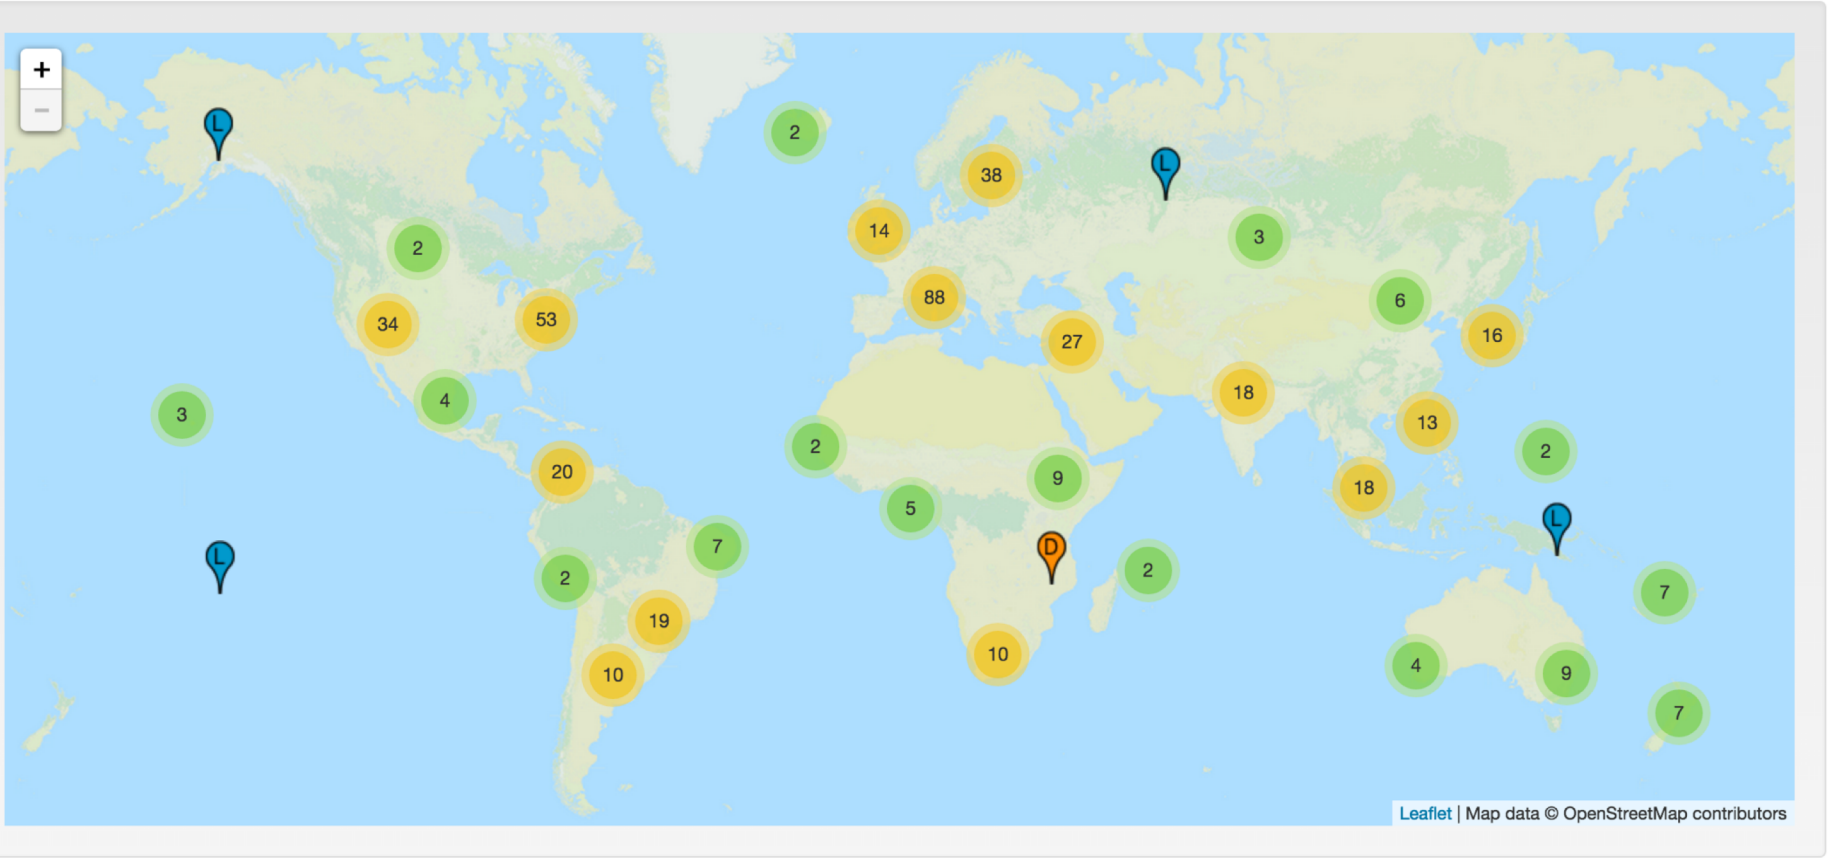

http://www.root-servers.org/

TLD and Authoritative Servers

- **Top-level domain** (TLD) servers
  - Responsible for com, org, net, edu,… and all top-level country domains (uk, fr, ca, jp, …)
  - Server maintainers
    - VeriSign for `.com`, `.net` TLDs
    - Educause for `.edu` TLD
- **Authoritative** DNS servers:
  - Organization's DNS servers, providing authoritative hostname to IP mappings for organization's servers
  - Can be maintained by organization or service provider

Local Name Server (Cache)

- **Aka "Stub Resolver"**
- Not part of previous hierarchy
- Each ISP (residential ISP, company, university) has one or more
- When host makes DNS query, query is sent to its local DNS server
  - Maintains local cache of common DNS records
    - *www.facebook.com?*
  - Acts as proxy, forwards query into hierarchy and provides eventual reply
- **You typically know this server's IP address from DHCP (upon connecting to the network)**

DNS and UDP

- DNS uses UDP by default
  - It *can* use TCP, but it's rare
  - **Isn't this unreliable?**
- Why use UDP
  - Faster (in three ways!)
    - No need to establish a connection (RTT/latency overhead)
    - Lower per-packet byte overhead in UDP header
    - Less packet processing by hosts
  - Reliability not needed
    - DNS will just re-request if no response received (2-5 seconds)

DNS and Security

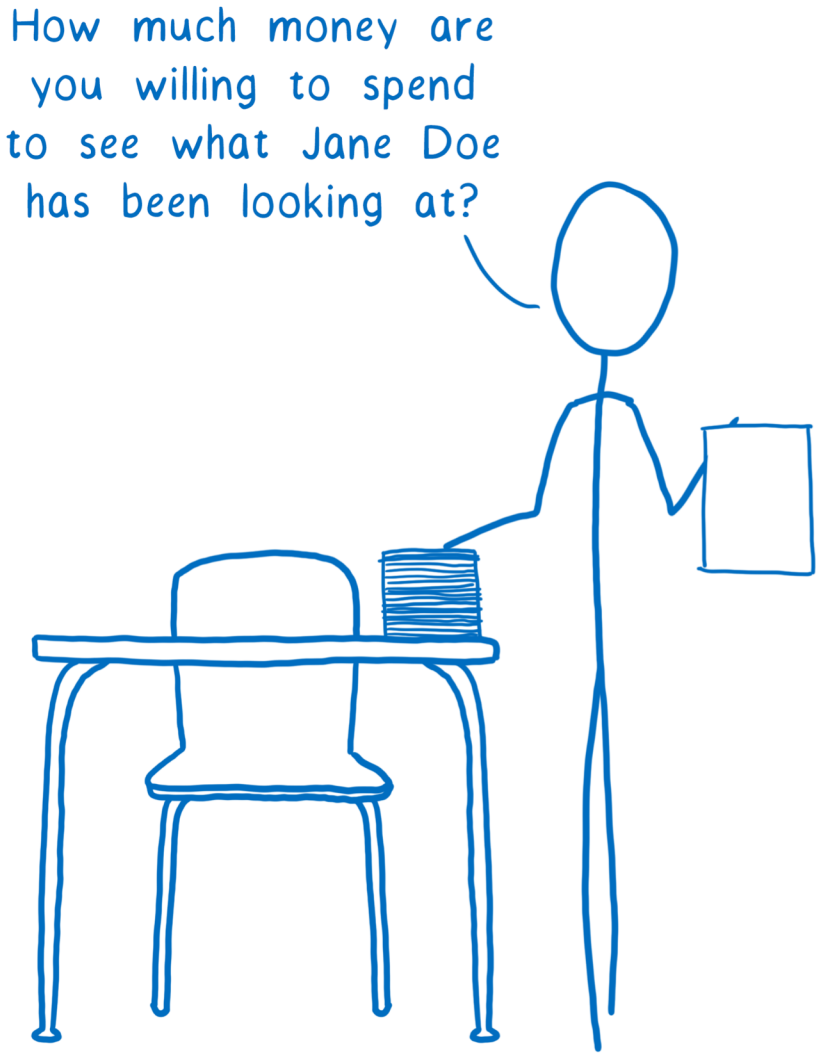

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

DNS and Security

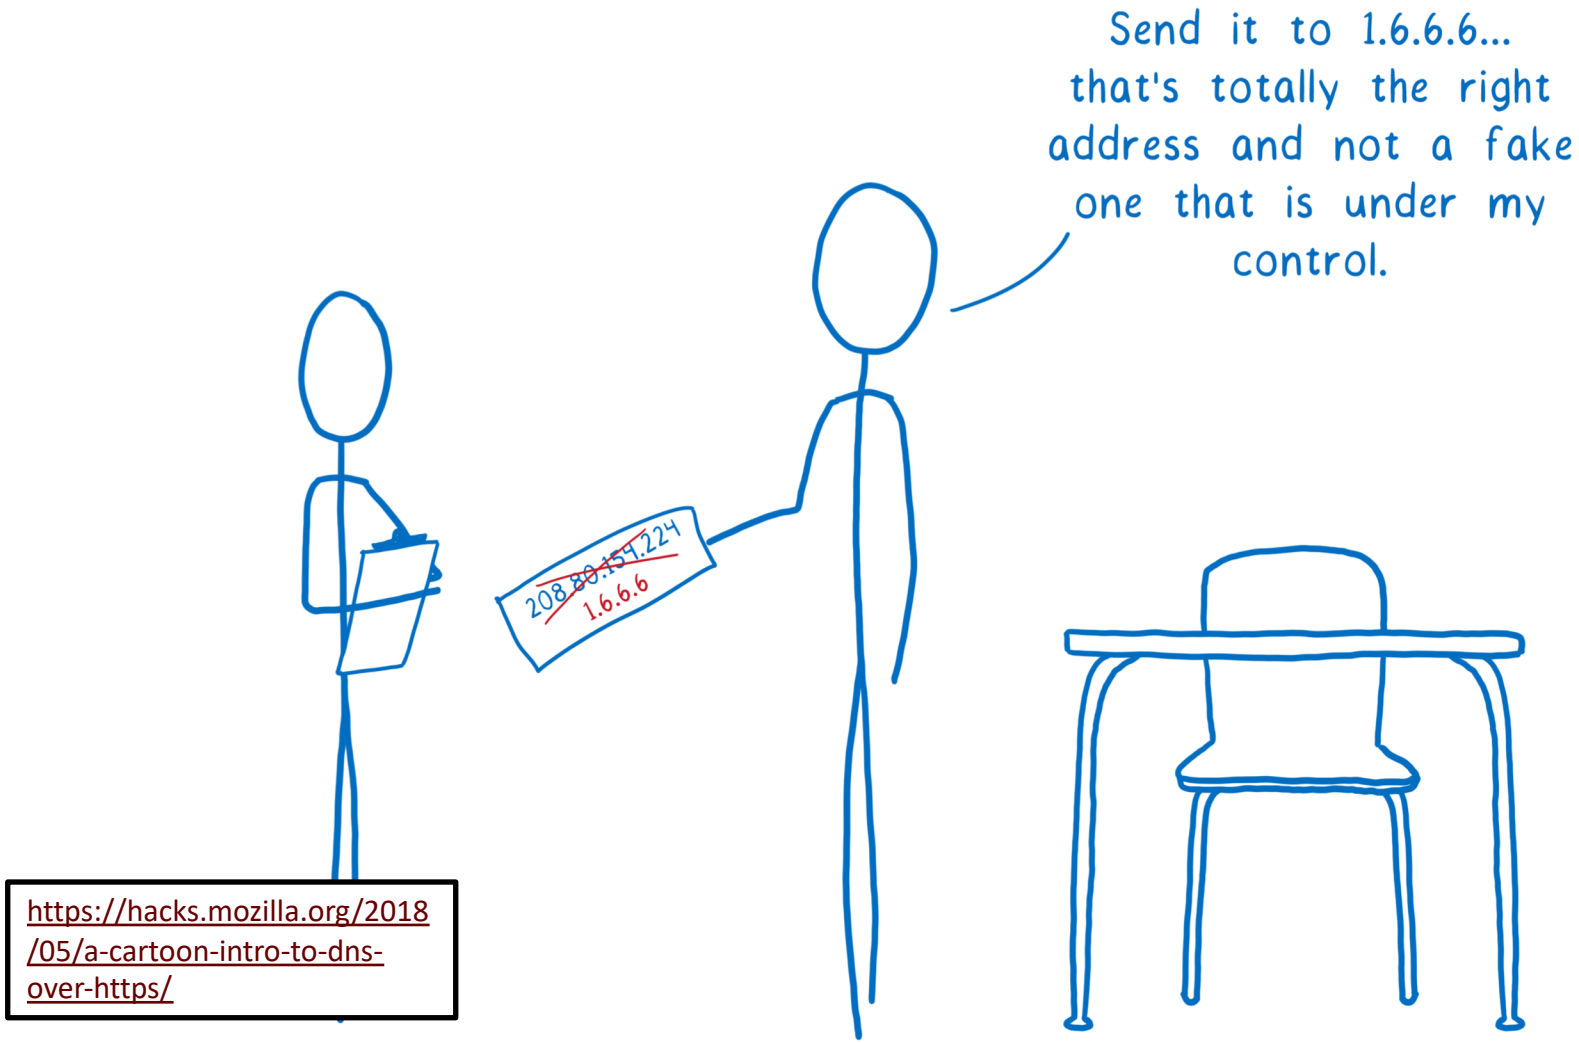

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

DNS and Security

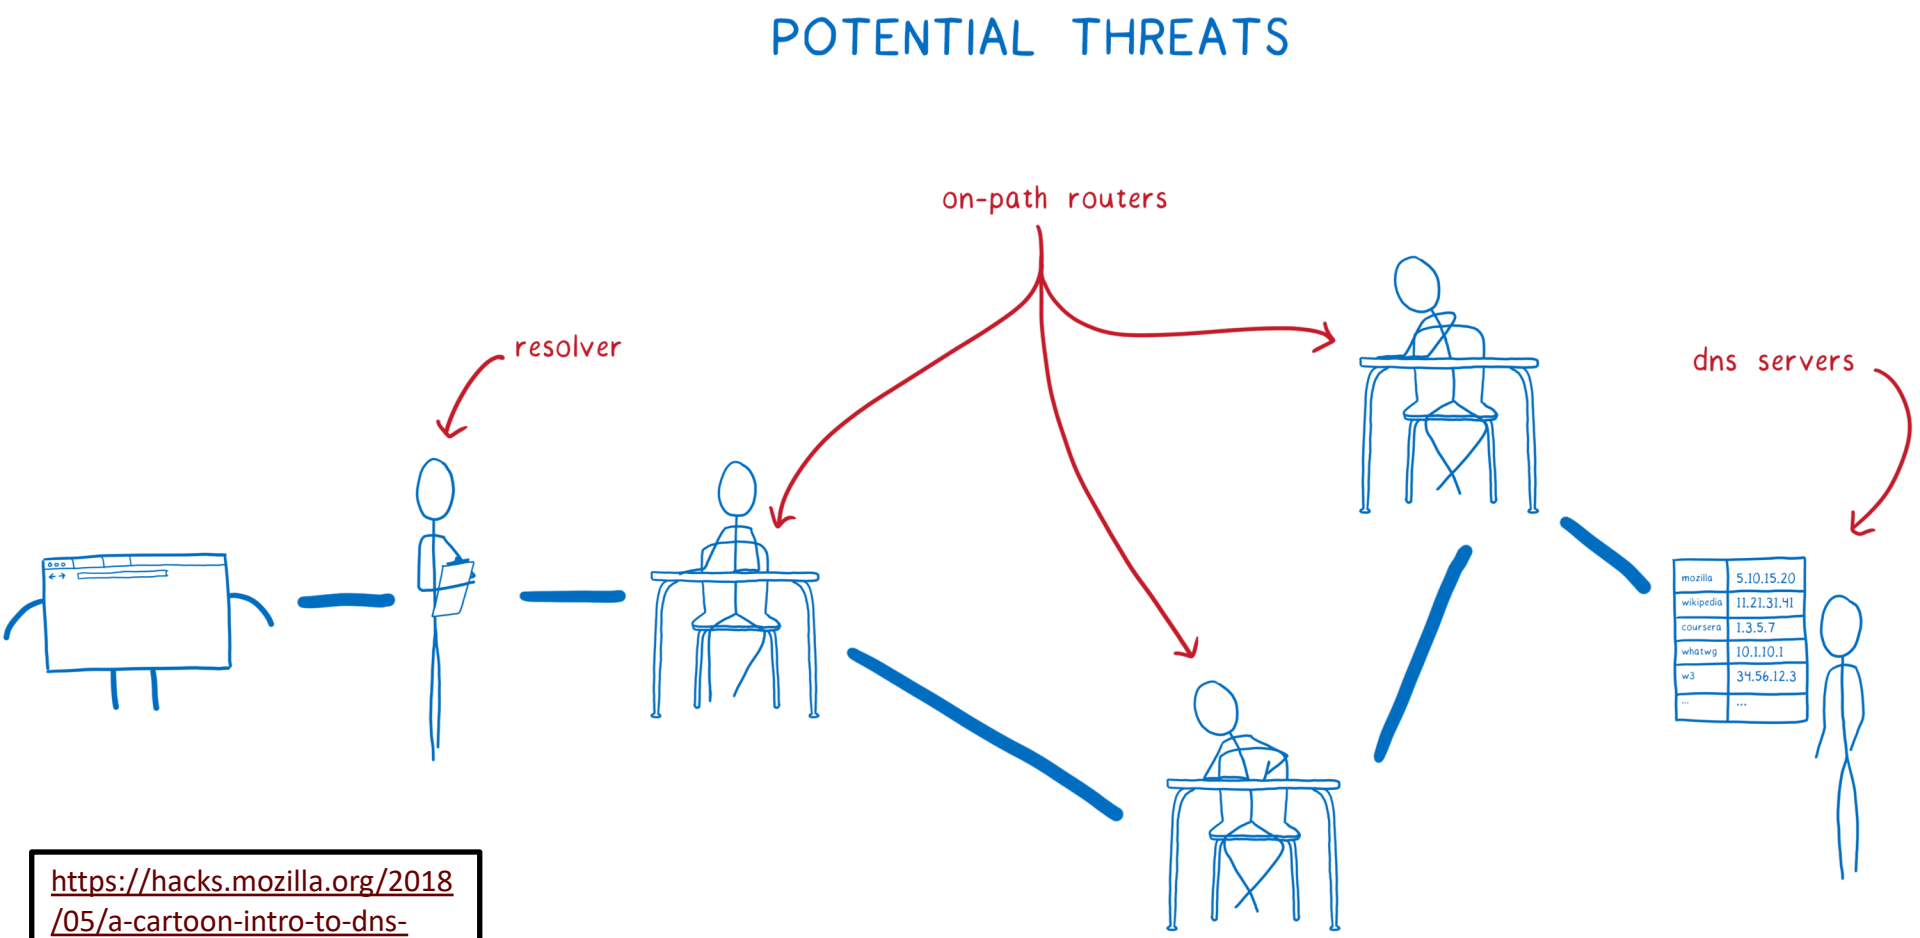

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

DNS and Security

- **Confidentiality**
  - Traditional DNS request and reply (over UDP) is plaintext
    - ISP spies on your Internet usage for profit?
    - NSA spies on your Internet usage for control?
    - DNS is *not just for names*
  - Solutions: **DNS over HTTPS**, **DNS over TLS**
- **Integrity**
  - Traditional DNS request and reply (over UDP) is unsigned
  - ISP tampers with reply message? (NXDOMAIN replaced with ad-laden site)
  - Governments tamper with reply message? (Domain blocked by court order)
  - Hackers tamper with reply message? (Redirect to malware site)
  - Solutions: **DNSSEC** (and DNS over HTTP/TLS)
- **Availability**
  - *Addressed by DNS distributed database design*

https://dnsprivacy.org

DNSSEC

DNSSEC

- **Domain Name System Security Extensions** (**DNSSEC**)
- Validate that a DNS response has not been tampered with
  - IP addresses, TXT, MX, etc… (all data protected)
- Uses public/private keys and signatures
- Prevents some attacks against clients (e.g. DNS cache poisoning)
- Does **not** provide confidentiality
  - Communication between client and server is in plaintext

DNSSEC Basics

- All records of same type (AAAA, …) grouped into resource record set (RRSet)
  - The RRSet is digitally signed, not individual record

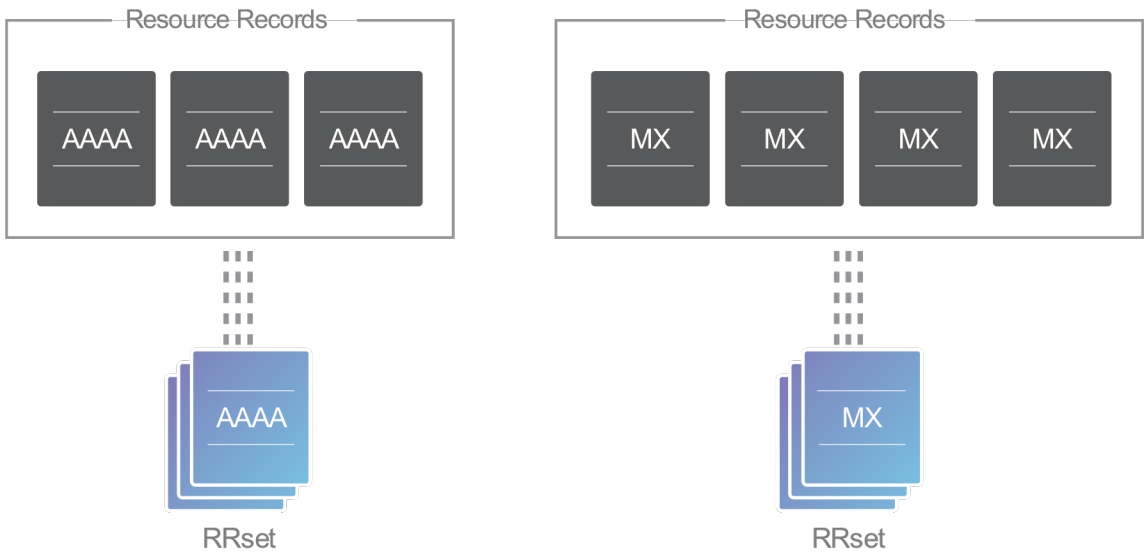

https://www.cloudflare.com/dns/dnssec/how-dnssec-works/

DNSSEC Basics

- Each zone has Zone-Signing Key (ZSK)
  - Private key signs entire RRset
    - Signature saved in RRSig record (stored in DNS)
  - Public key verifies entire RRSet
    - Key saved in DNSKey record (stored in DNS)

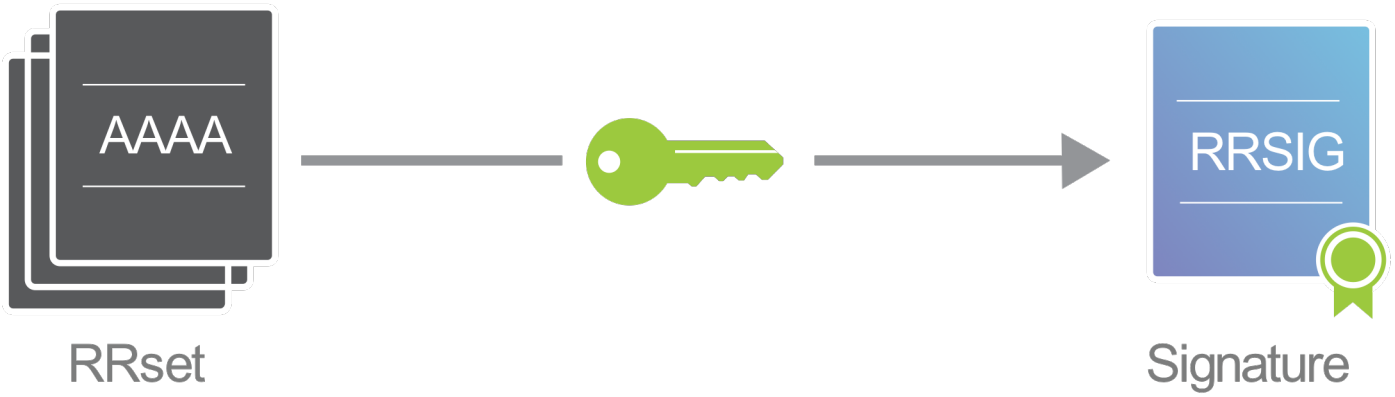

https://www.cloudflare.com/dns/dnssec/how-dnssec-works/

DNSSEC Basics

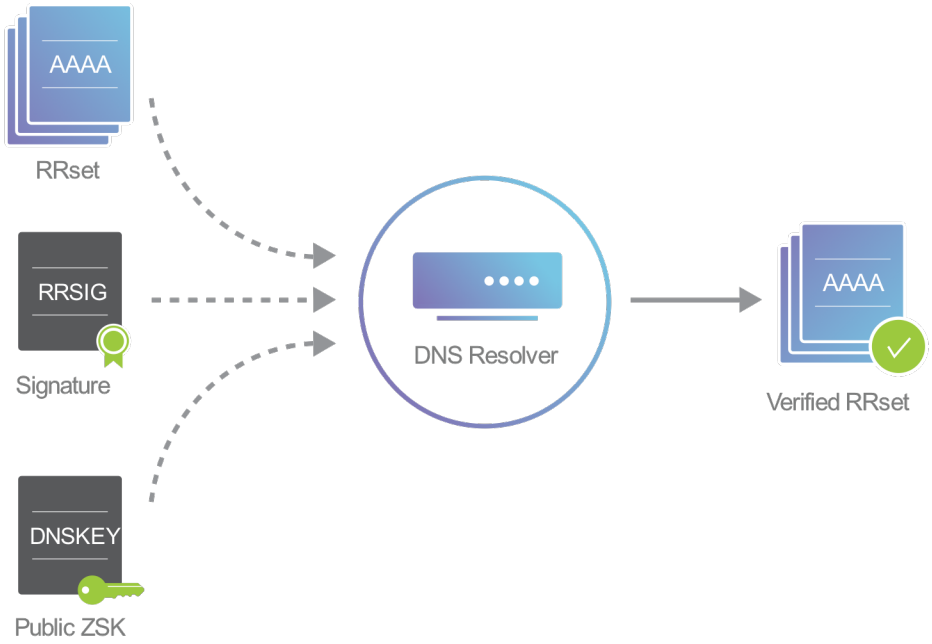

- Resolver pulls a particular record (AAAA) along with the RRSig *(which signs the record set)* and public Zone Signing Key *(which verifies the RRSig)*
  - Resolver verifies signature
  - **But how does it trust the public ZSK?**

https://www.cloudflare.com/dns/dnssec/how-dnssec-works/

DNSSEC Basics

- DNSSEC name servers have Key Signing Keys (KSK)
  - KSK is used to sign public Zone Signing Key (ZSK)
  - Name server publishes public KSK in a DNSKey record
    - **But how do we trust these KSKs?**

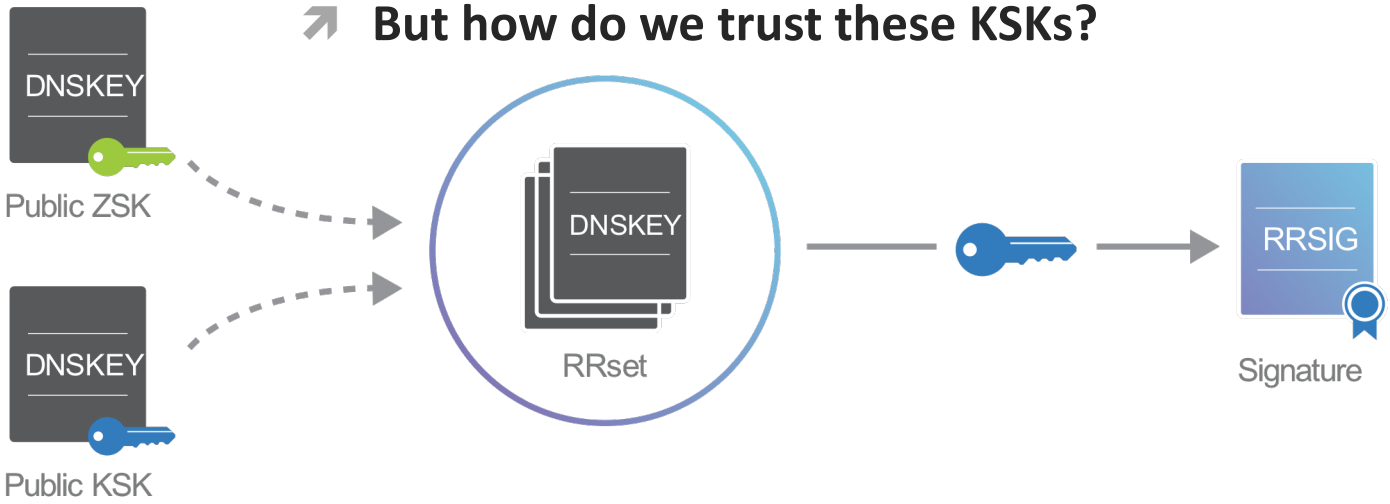

https://www.cloudflare.com/dns/dnssec/how-dnssec-works/

DNSSEC Basics

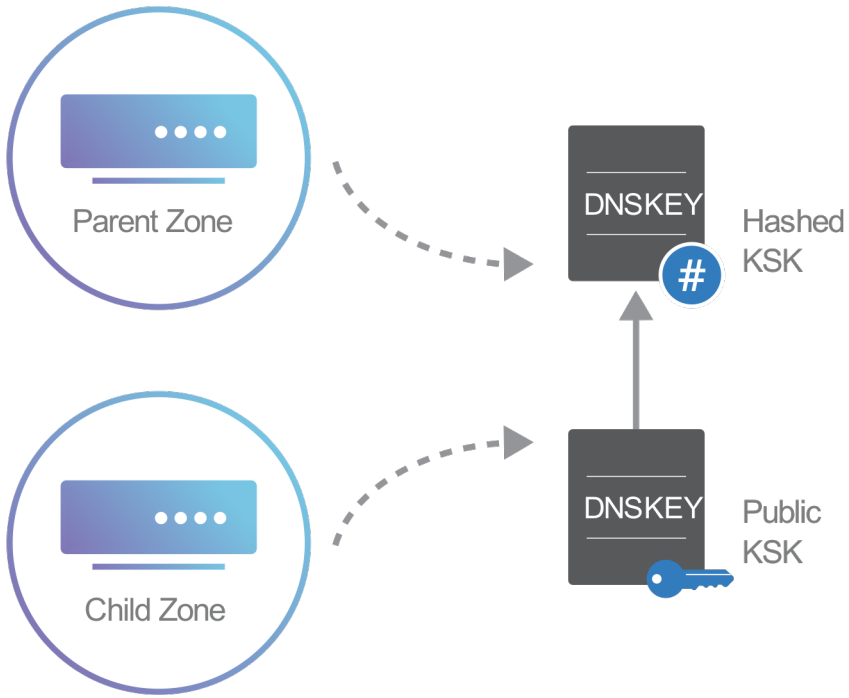

- Delegation Signer (DS) records allow trust to be transferred from Parent zone to Child Zone
- Hash of DNSKey record (containing KSK) is produced by zone operator (e.g. `example.com`) and given to parent zone (e.g. `.com`)

https://www.cloudflare.com/dns/dnssec/how-dnssec-works/

DNSSEC Basics

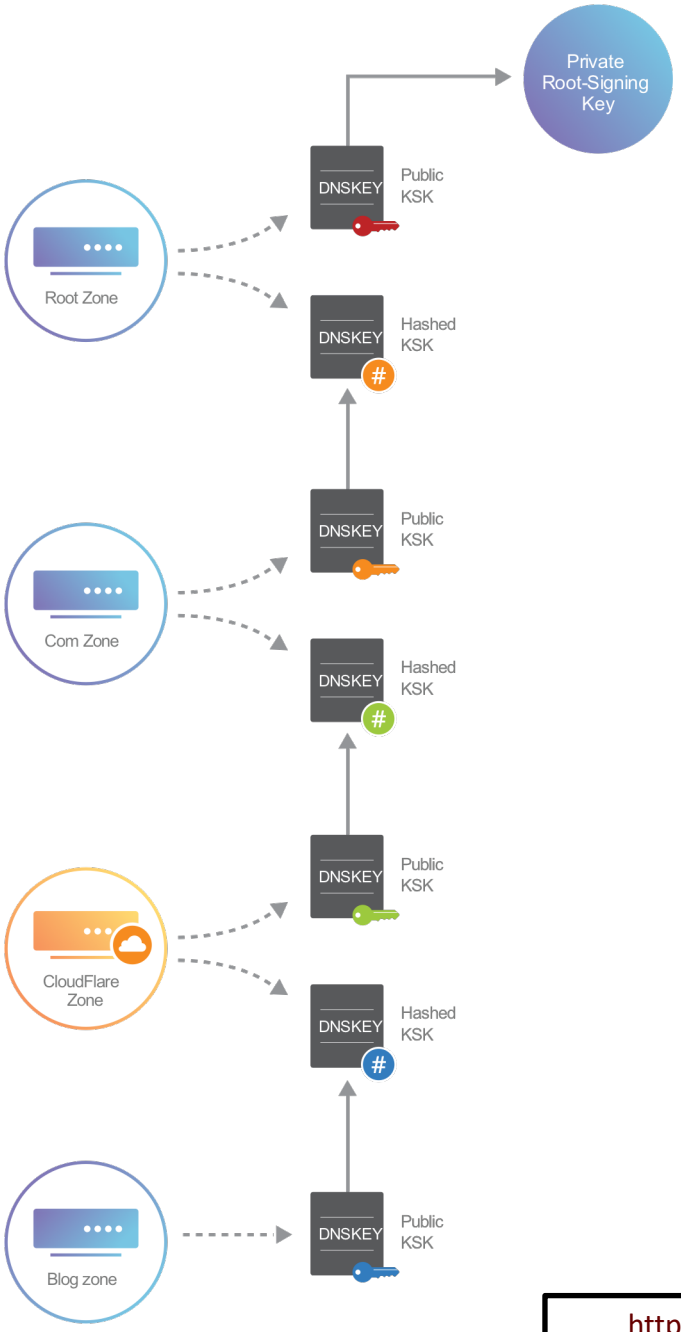

- Chain of Trust
  - Root is **self-signed** – societal engineering challenge to rotate root KSKs on periodic basis
- DNS Root zone KSK last rotated October 11 2018
  - Previous key was from 2010

https://www.cloudflare.com/dns/dnssec/how-dnssec-works/

Use of DNSSEC Validation for World (XA)

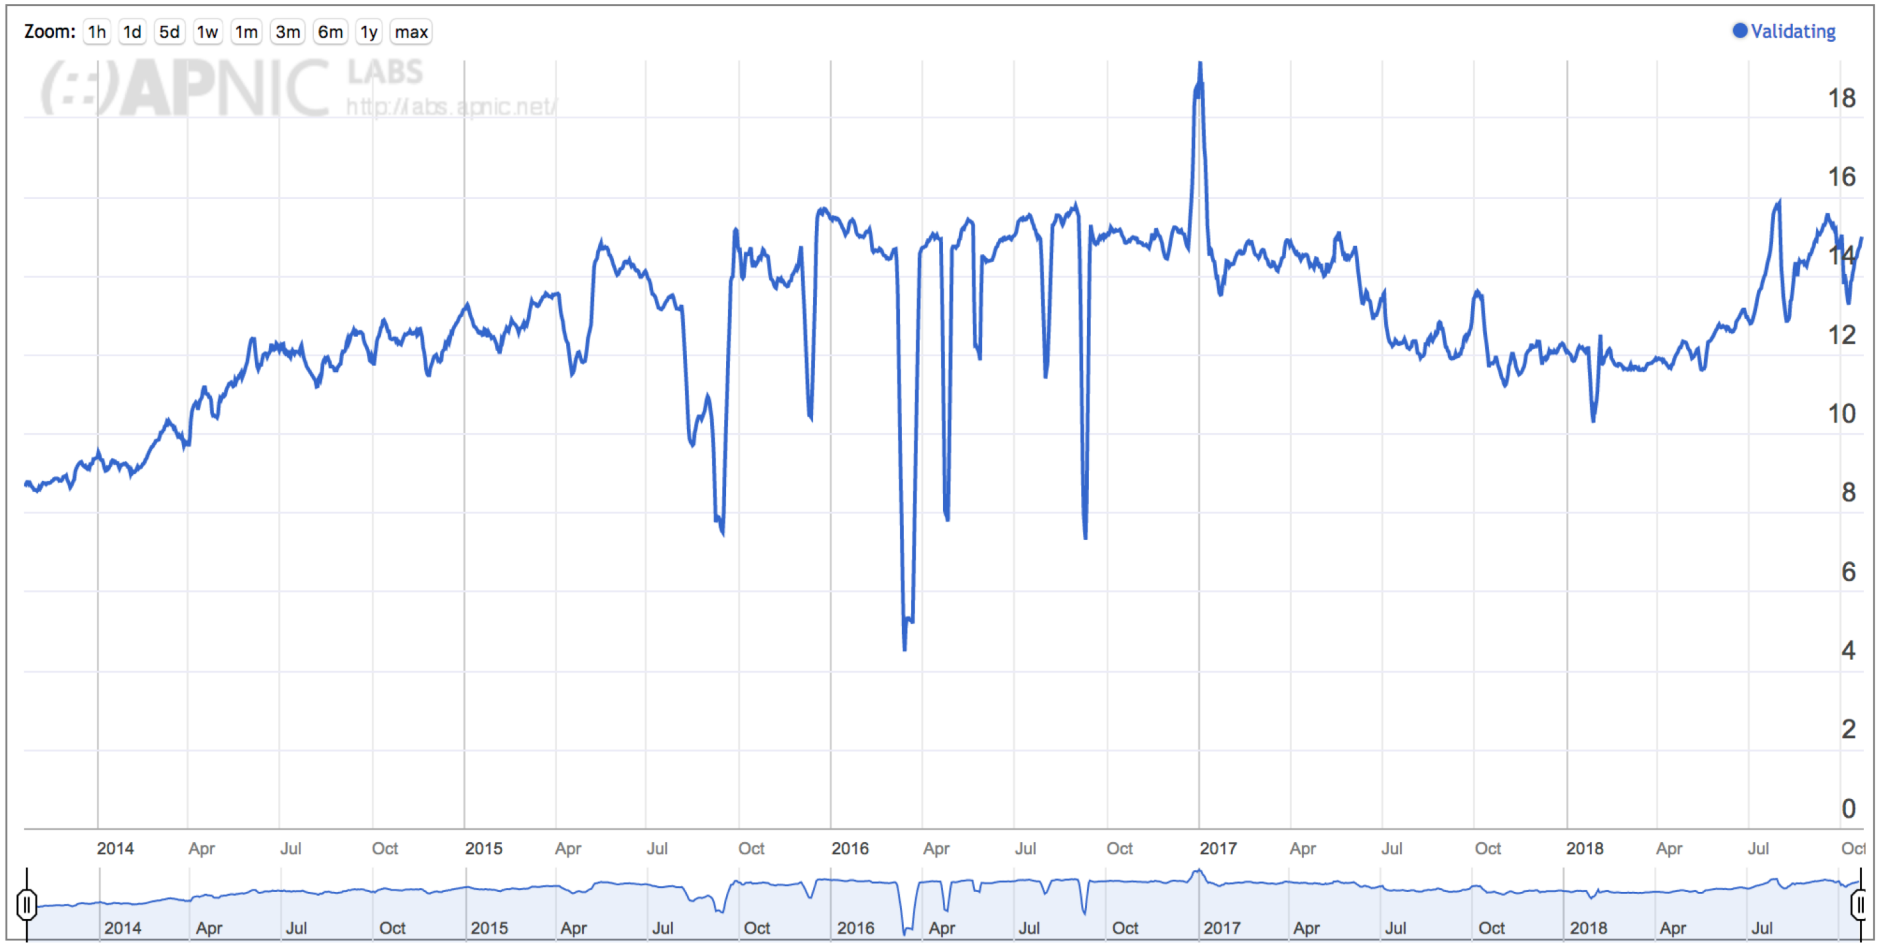

- Very slow adoption of DNSSEC – Is trend even increasing?
  - RFC 4033 released in 2005

http://stats.labs.apnic.net/dnssec/XA?c=XA&x=1&g=1&r=1&w=7&g=0

https://blog.apnic.net/2018/02/26/peak-dnssec/

DNSSEC Validation Rate by country (%)

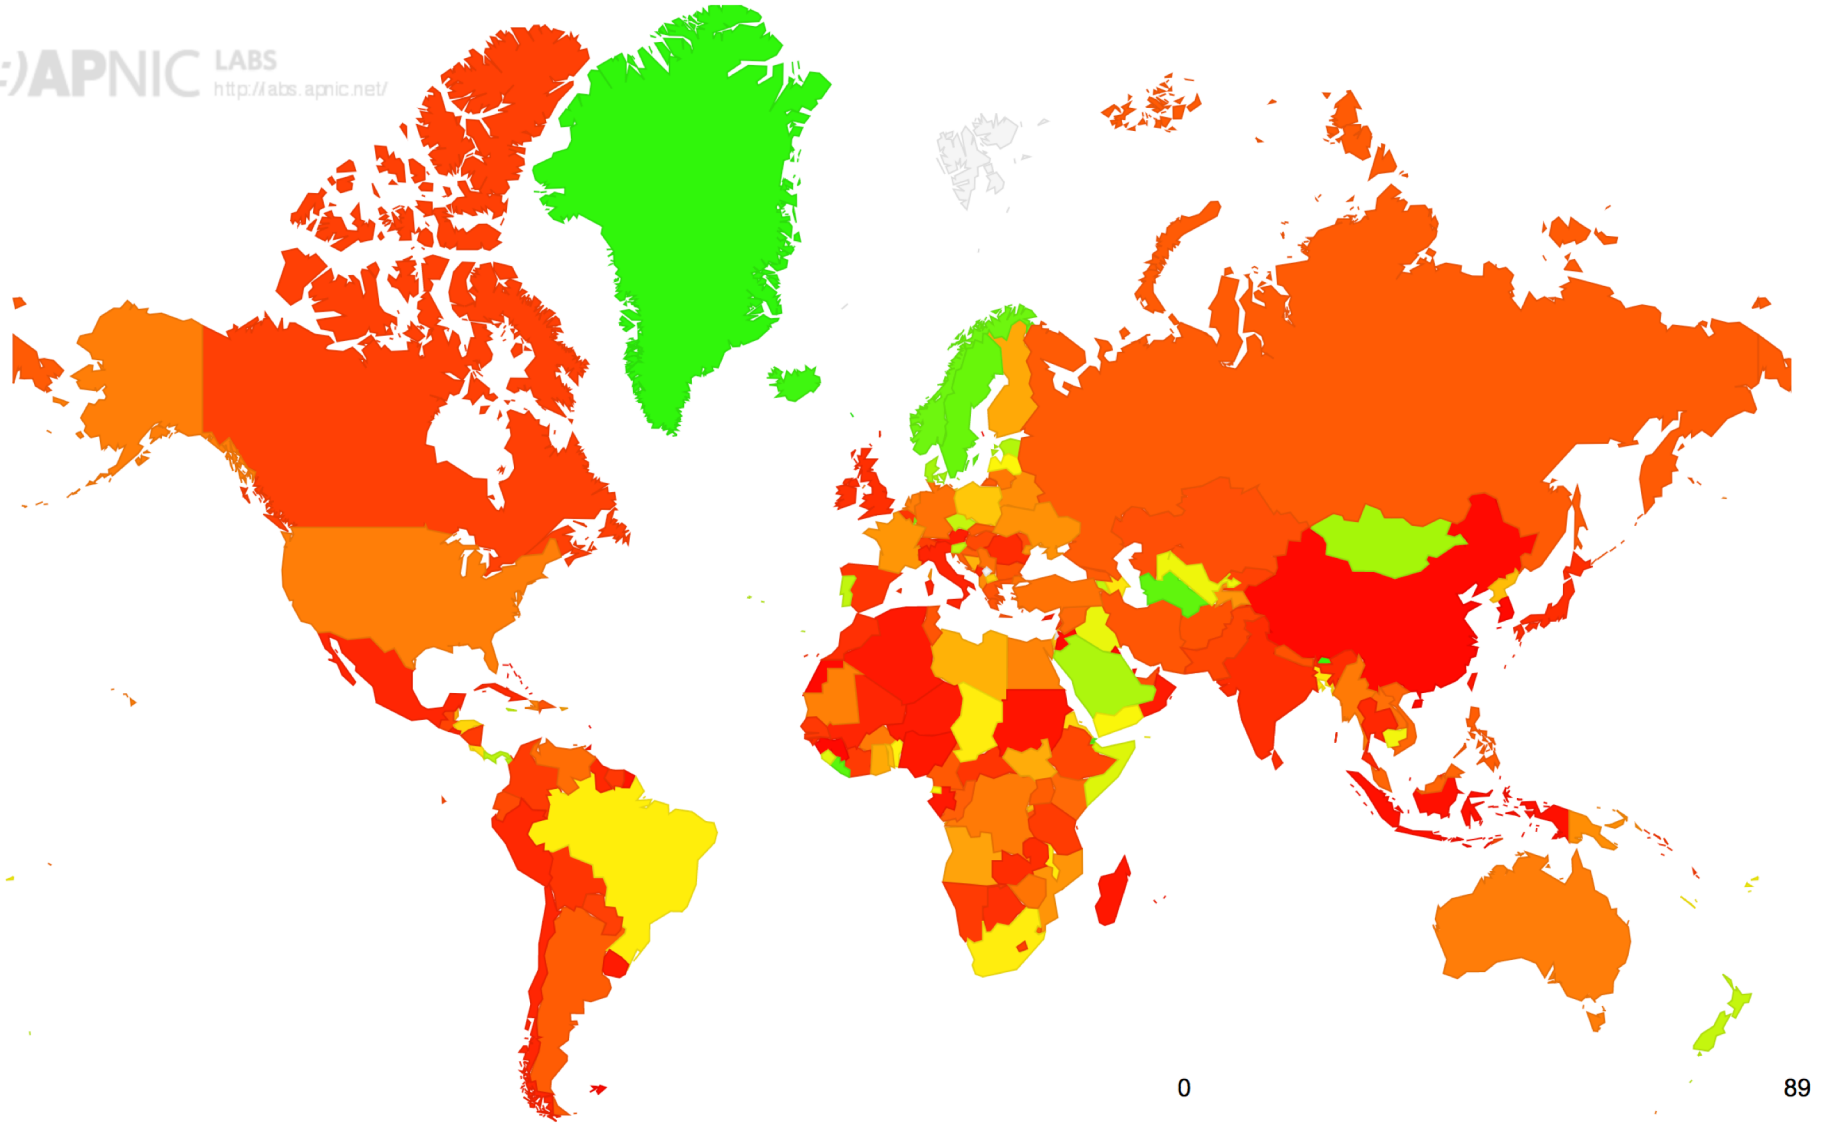

DNS over TLS

Motivation

- DNS sent over plaintext is vulnerable to snooping and manipulation
  - Encrypt it!

DNS Over TLS

- Encrypt DNS queries/responses over TLS connection
  - Uses existing DNS functionality to send queries/responses over TCP (infrequently used)
    - Now just encrypted via TLS
  - TCP port 853
  - [RFC 7858]
- Provides Confidentiality + Integrity for MITM attacks (no eavesdropping / no tampering)
  - **Key caveat** – The DNS server itself could provide a malicious reply. This does not replace need for DNSSEC!

Adoption

Clients

- Linux Systemd
  - Implemented as-of June 2018 (v239)
  - Off by default
- Stubby
  - Local DNS Privacy Stub Resolver
  - https://dnsprivacy.org/wiki/display/DP/DNS+Privacy+Daemon+-+Stubby

DNS Services

- Cloudflare (1.1.1.1)
- Google Public DNS (8.8.8.8)
- Quad9 (9.9.9.9)

DNS over HTTPS (DOH)

Motivation

- Network operators use DNS as means to enforce policy
  - *Though Shall NOT Access That Website*
  - Oppressive government?
  - Oppressive network operator?
  - Responsible network operator trying to save you from visiting malicious websites?

DNS Over HTTPS

- Encode DNS queries and responses over HTTPS
  - [RFC 8484, *draft standard as of Oct 2018*]
- Advantages *(for web browsers)*
  - Privacy (DNS request/response encrypted)
  - Tamper resistance: Network operators can't block DNS without also blocking HTTPS
    - *Which is very obvious to the end user*
    - *Network operators can't even tell there's DNS data being sent*
  - Reduced latency
    - HTTP/2 server push
    - Browser can do DNS directly, no need to invoke system resolve
  - Proxying and caching will work for DNS too

https://datatracker.ietf.org/wg/doh/about/

DNS Over HTTPS

- Can work independently from existing DNS methods
  - UDP
  - TLS [RFC 7857]
  - DTLS [RFC 8094]
- DNS response data (identical bytes as UDP response) is placed in HTTPS playload
  - MIME type: `application/dns-message`
  - HTTP/2 server push can even send values to client in advance of request
- **Key caveat** – The DNS server itself could provide a malicious reply. This does not replace need for DNSSEC!

Adoption

Clients

- Firefox (v62+)
  - Not enabled by default
  - [2018] Firefox Nightly sending requests via traditional resolver and DOH and measuring performance/accuracy
  - Average 6ms slowdown
- Chrome
  - Development in Chromium ongoing

DNS Services

- Cloudflare (1.1.1.1)
- Google Public DNS (8.8.8.8)
- Quad9 (9.9.9.9)

DOHysteria

/dəʊ hɪˈstɪərɪə/

noun

exaggerated or uncontrollable emotion or excitement surrounding DNS over HTTPs

Origin: Geoff Huston

DOH Challenges

- Browser can use a different DNS namespace than the rest of your computer (email, chat, etc.) that still uses the system resolver
  - Browser can "punch-through" local infrastructure
- Implications on naming consistency across applications?
  - *DNS Split-Horizon hacks: PacificNet example with students using Google Public DNS while on-campus*

https://blog.apnic.net/2018/10/23/dns-oarc-29-diving-into-the-dns/

DNS Trivia

- Web browsers (and other applications) use the resolver provided by the operating system, which is (typically) provided via DHCP
  - ***But they don't <u>have</u> to….***

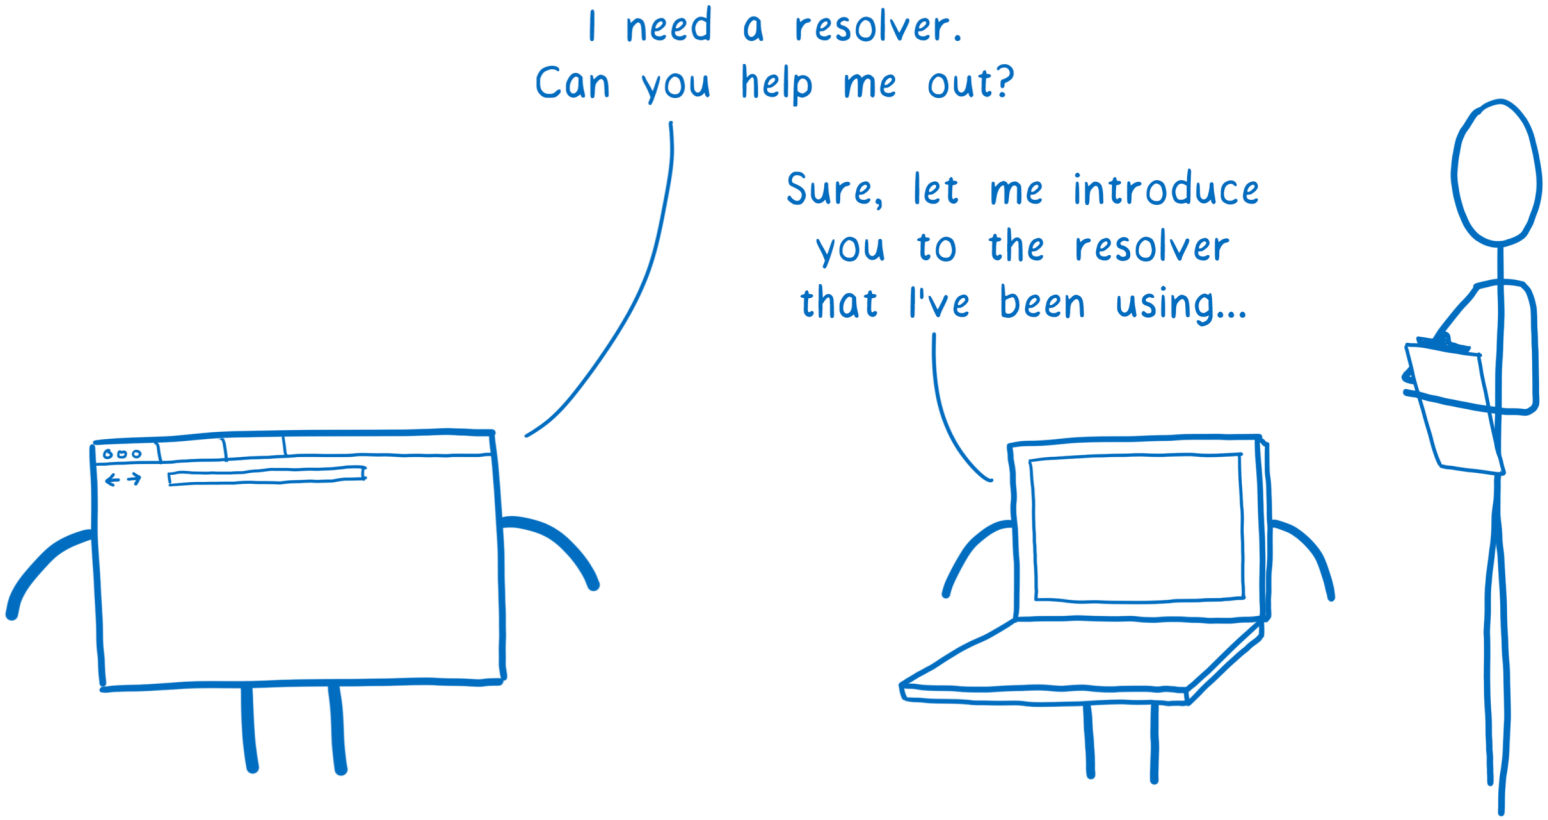

https://hacks.mozilla.org/2018/05/a-cartoon-intro-to-dns-over-https/

Cloudflare (1.1.1.1)

- DNSSEC: **YES** DNSoTLS: **YES** DOH: **YES**
- Privacy policy
  - Discard all personally identified information after 24 hours
  - Never sold to third parties

https://blog.cloudflare.com/announcing-1111/

Google Public DNS (8.8.8.8)

- DNSSEC: **YES** DNSoTLS: **NO** DOH: **YES**

https://blog.cloudflare.com/announcing-1111/

IBM Quad 9 (9.9.9.9)

- DNSSEC: **YES** DNSoTLS: **YES** DOH: **YES**
- Blocks access to domains considered threat to security by “threat intelligence partners”

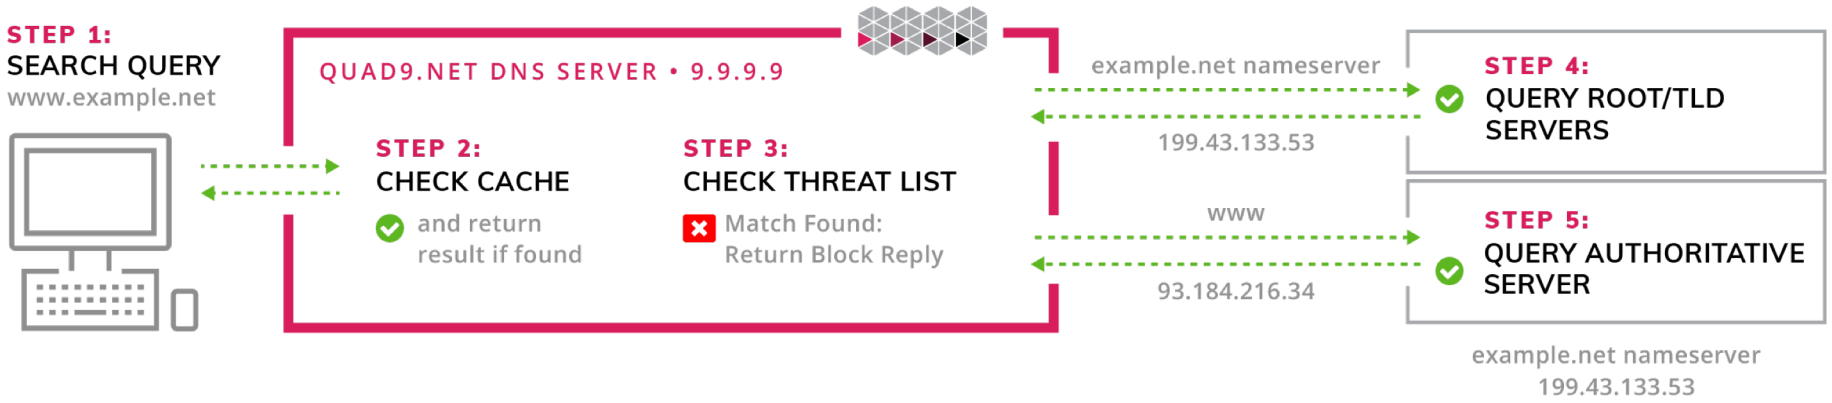

IBM Quad 9 (9.9.9.9)

Quad9 Threat Intelligence Partners (https://www.quad9.net/about/)

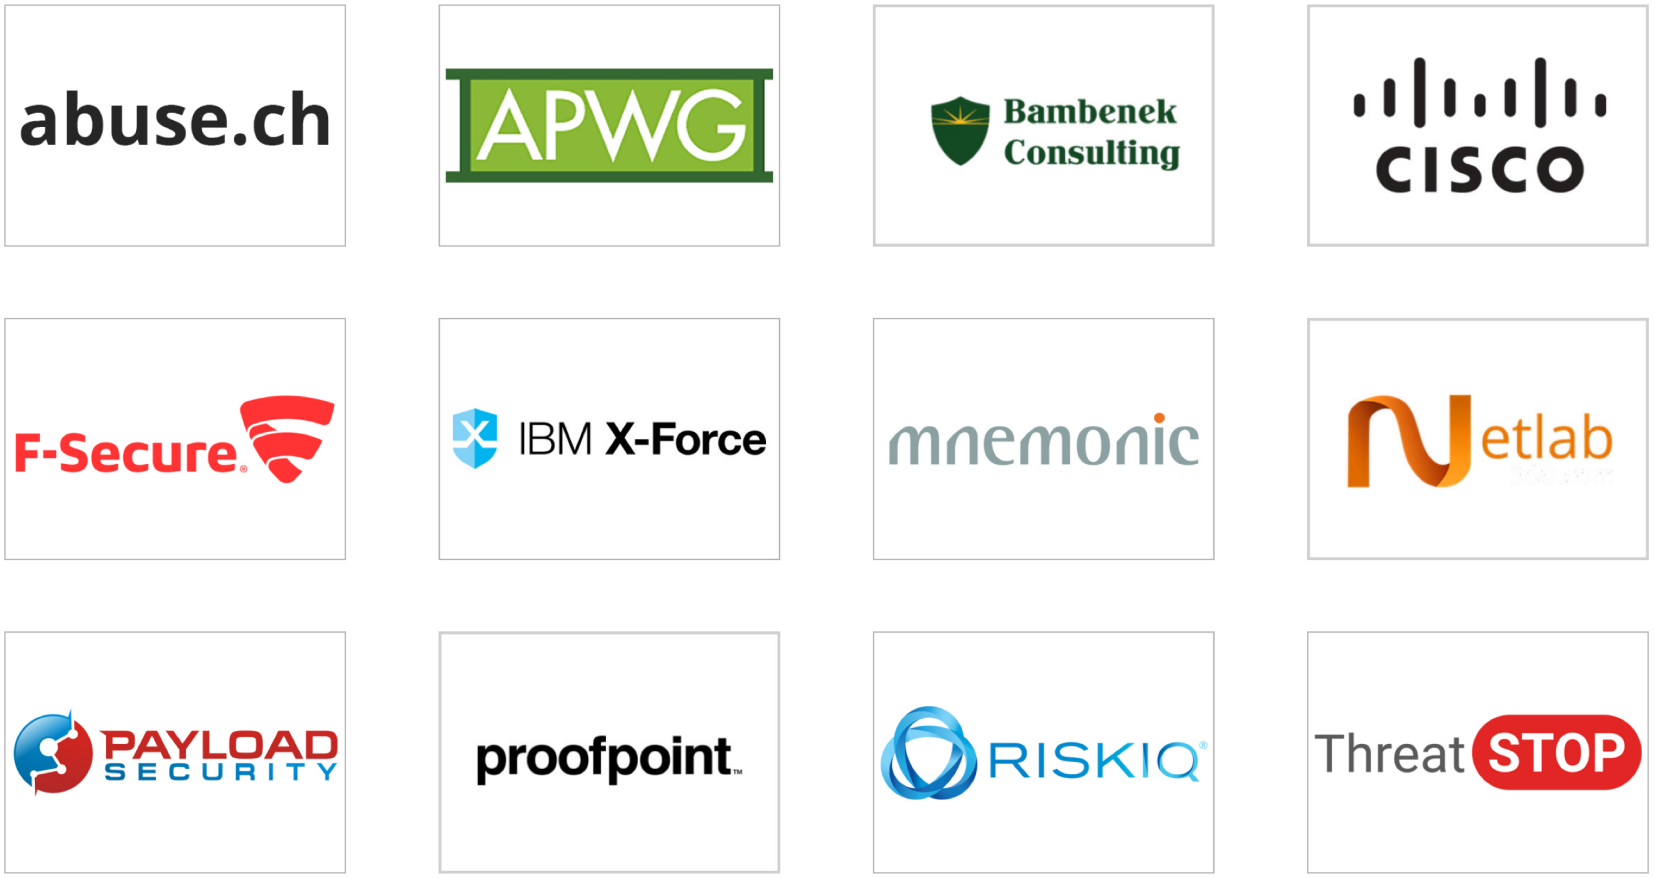

DNS and Security

- https://dnsprivacy.org
  - Good reference for ongoing work in this area!