Past, Present, and Future of DNS Resolution

Paul Vixie, CEO
Farsight Security, Inc.

2019-06 #firstcon2019

Abstract

- The Domain Name System has been a critical enabler of Internet growth since its inception in 1987. In the decades since then, the DNS *resolution* process has evolved from the LAN to the WAN, and to Anycast; it now includes DNSSEC *validation*, Extended DNS (EDNS) Client Subnet, larger message sizes, and I18N. The resolution process has also been abused for surveillance, advertising insertion, and exfiltration. Today the DNS resolution process is poorly understood, and yet under forced revision. The trend is for DNS to be carried inside HTTPS where it cannot be monitored or controlled except by servers and clients themselves, and the dangers this will yield must be studied and discussed while the future remains flexible. Dr. Vixie (Keio, 2012) will describe the past and present of DNS, and discuss its likely near term future.

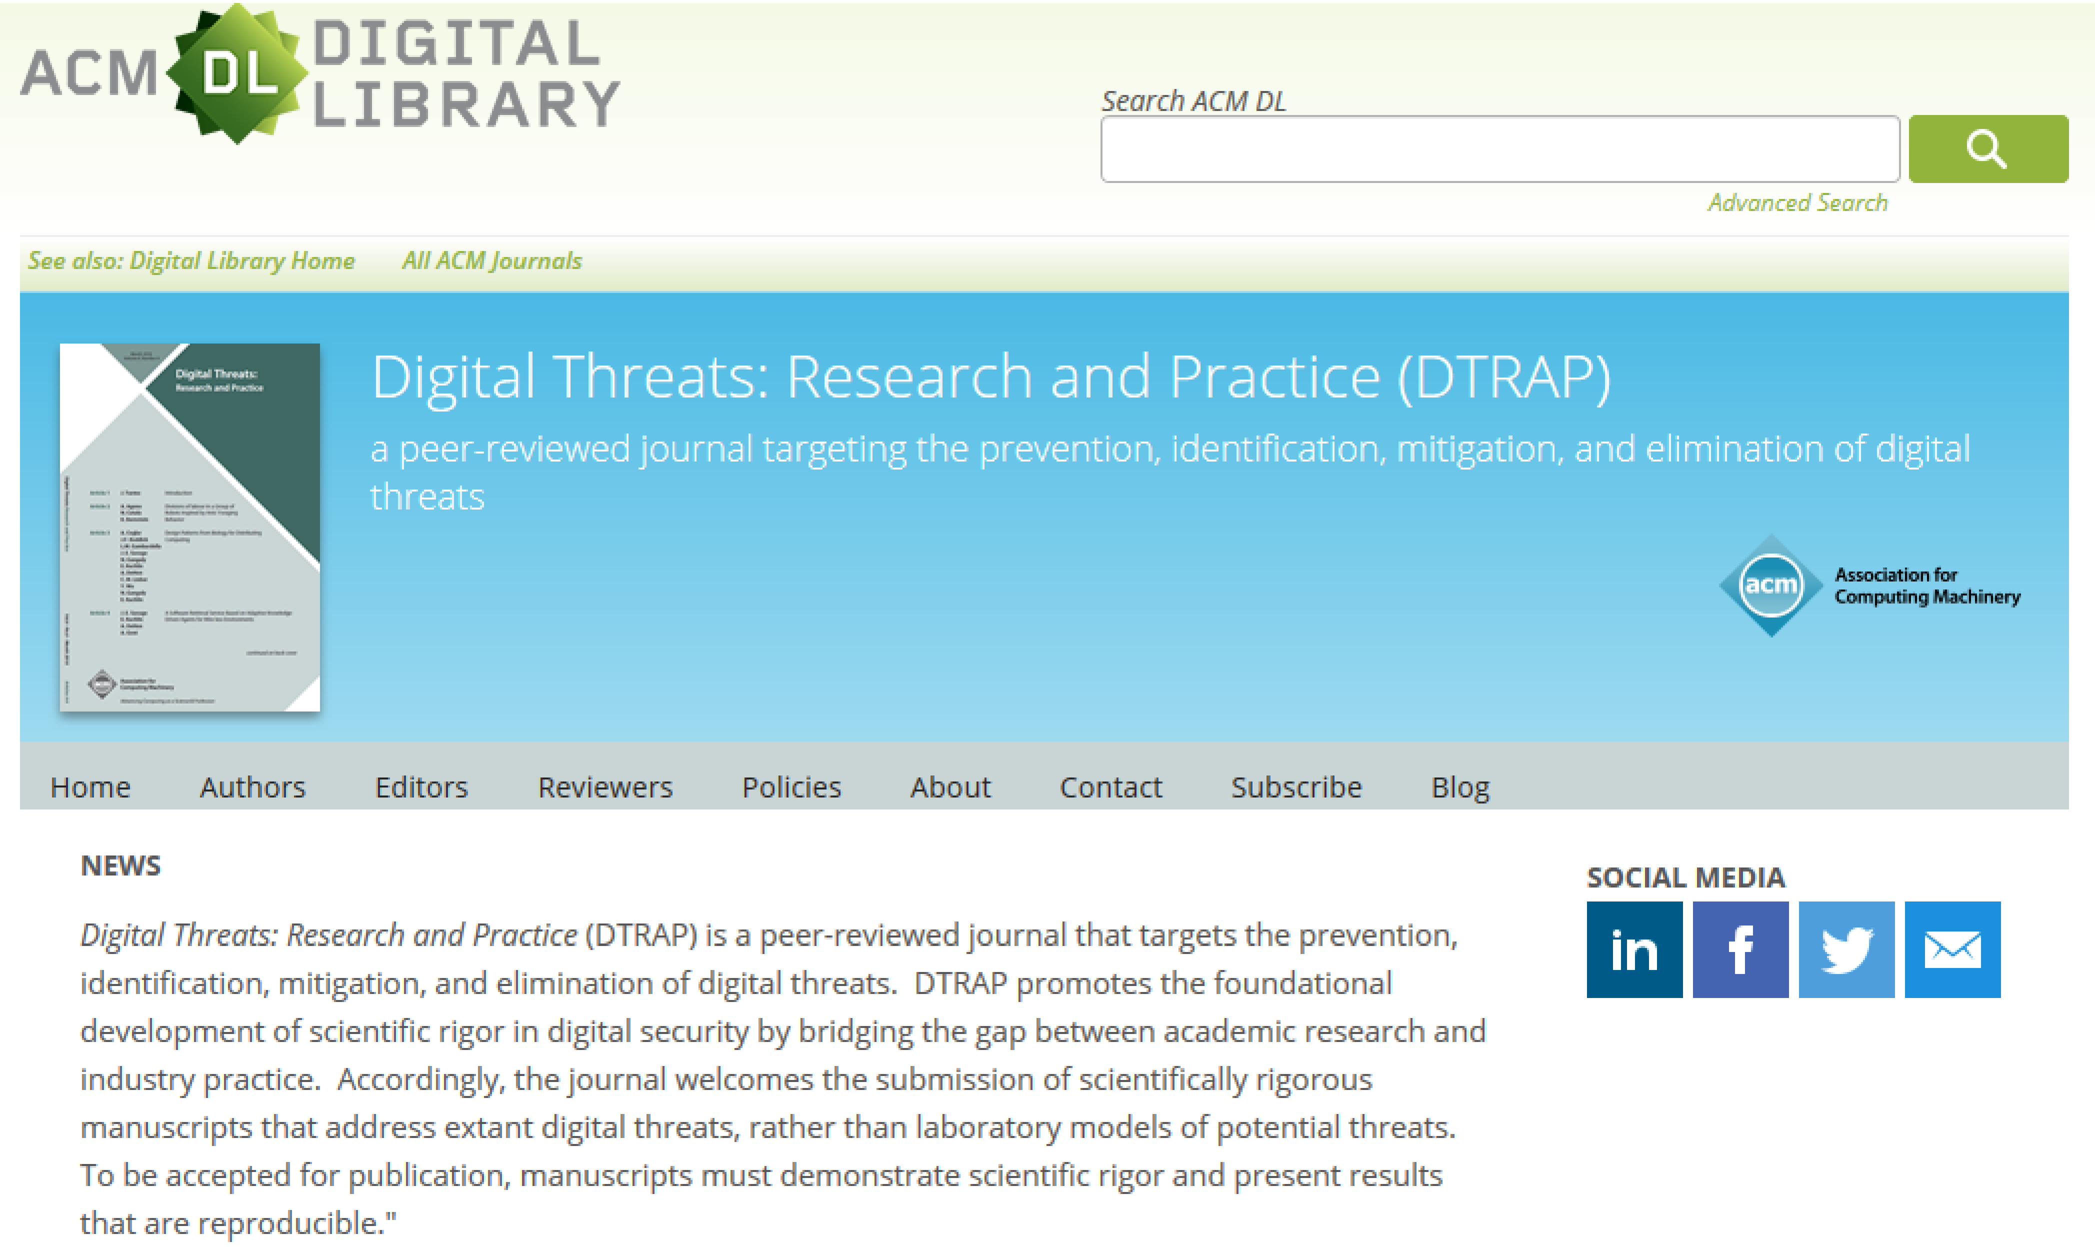

Digital Threats: Research and Practice (DTRAP)

a peer-reviewed journal targeting the prevention, identification, mitigation, and elimination of digital threats

NEWS

Digital Threats: Research and Practice (DTRAP) is a peer-reviewed journal that targets the prevention, identification, mitigation, and elimination of digital threats. DTRAP promotes the foundational development of scientific rigor in digital security by bridging the gap between academic research and industry practice. Accordingly, the journal welcomes the submission of scientifically rigorous manuscripts that address extant digital threats, rather than laboratory models of potential threats. To be accepted for publication, manuscripts must demonstrate scientific rigor and present results that are reproducible."

Pre-History

- Before DNS, other methods were used:
  - HOSTS.TXT file, fetched nightly by FTP
  - /etc/hosts file, edited by local sysadmins
  - Sun YP (NIS), to serve a LAN or campus
  - NeXT NetInfo, less portable than YP/NIS
  - SRI Hostname, proposed but never deployed
- Some of these had broader goals than IP address lookup
  - But to grow the Internet, IP address lookup had to scale by 10^8
  - Thus the driver for DNS was simple name → address translation

DNS and 4.3 BSD UNIX (~1986)

- BIND = **B**erkeley **I**nternet **N**ame **D**omain
  - Stub resolver (`libresolv`, `gethostbyname`(), `gethostbyaddr`())
  - Name server (named, named-xfer)
  - CLI utilities (nslookup)
- Funded by a USG contract, but was too late for 4.3 BSD UNIX release
  - Thus, first formal release was a .TAR file published via Anonymous FTP
  - Deployment was rapid and broad by the standards of that era
  - A patched system simply got a new `gethostbyname`(), `gethostbyaddr`()
- `libresolv`, a stub resolver, could operate even without a config file
  - Default recursive server was 0.0.0.0 (local host) and was used if it answered

DNS System Architecture

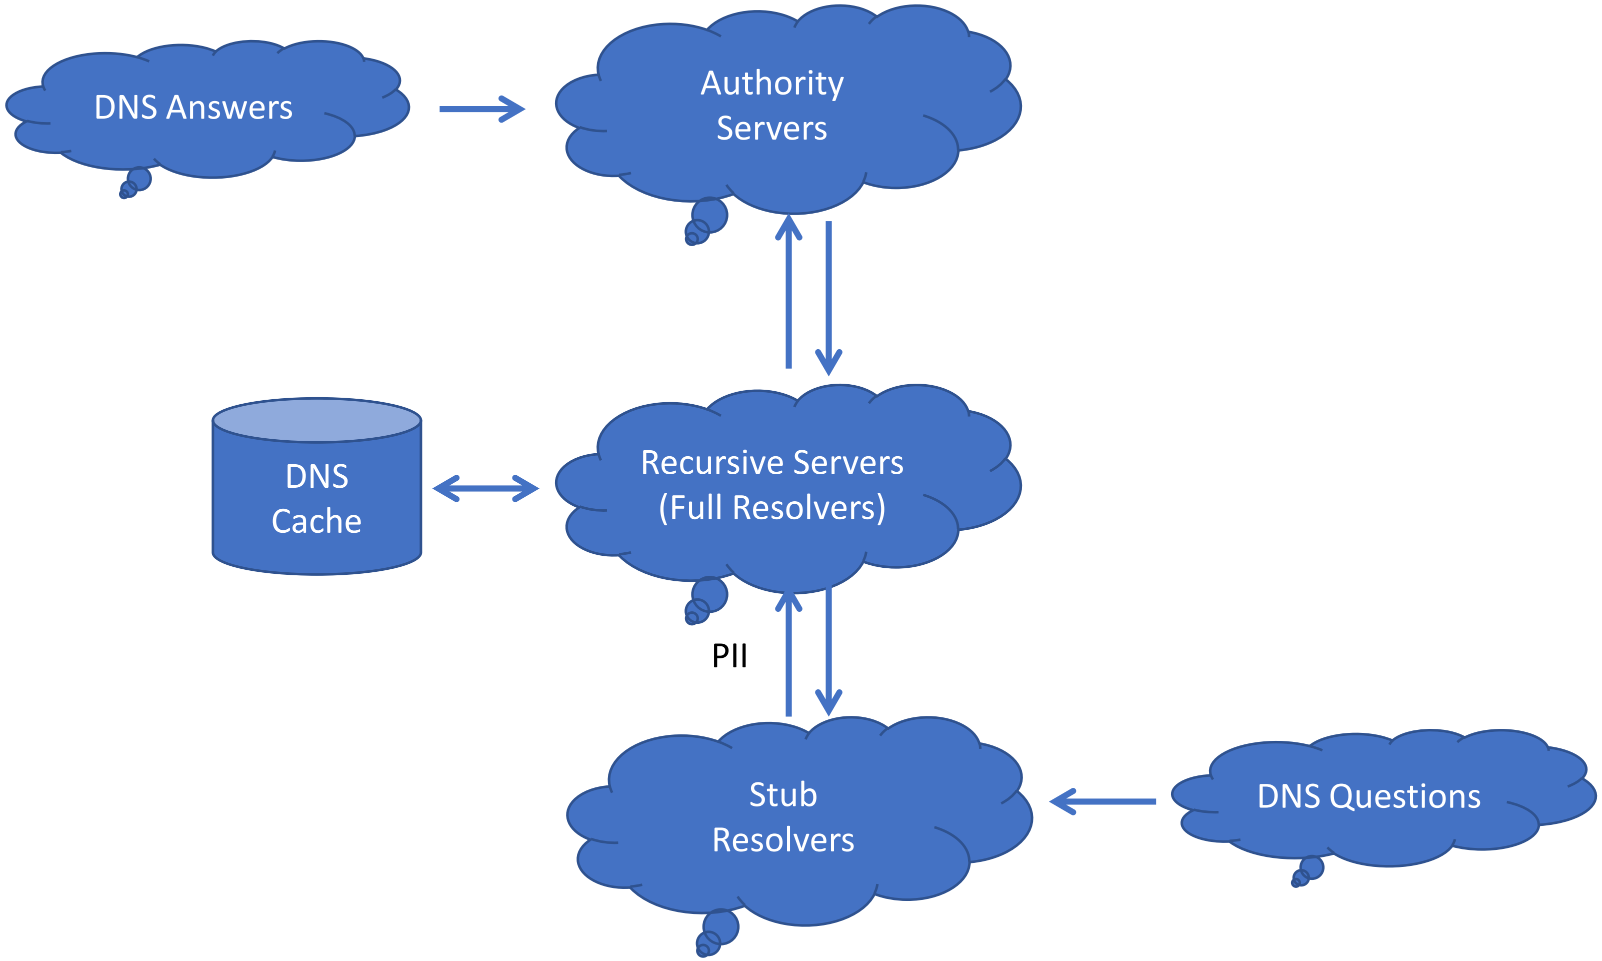

Internet System Topology

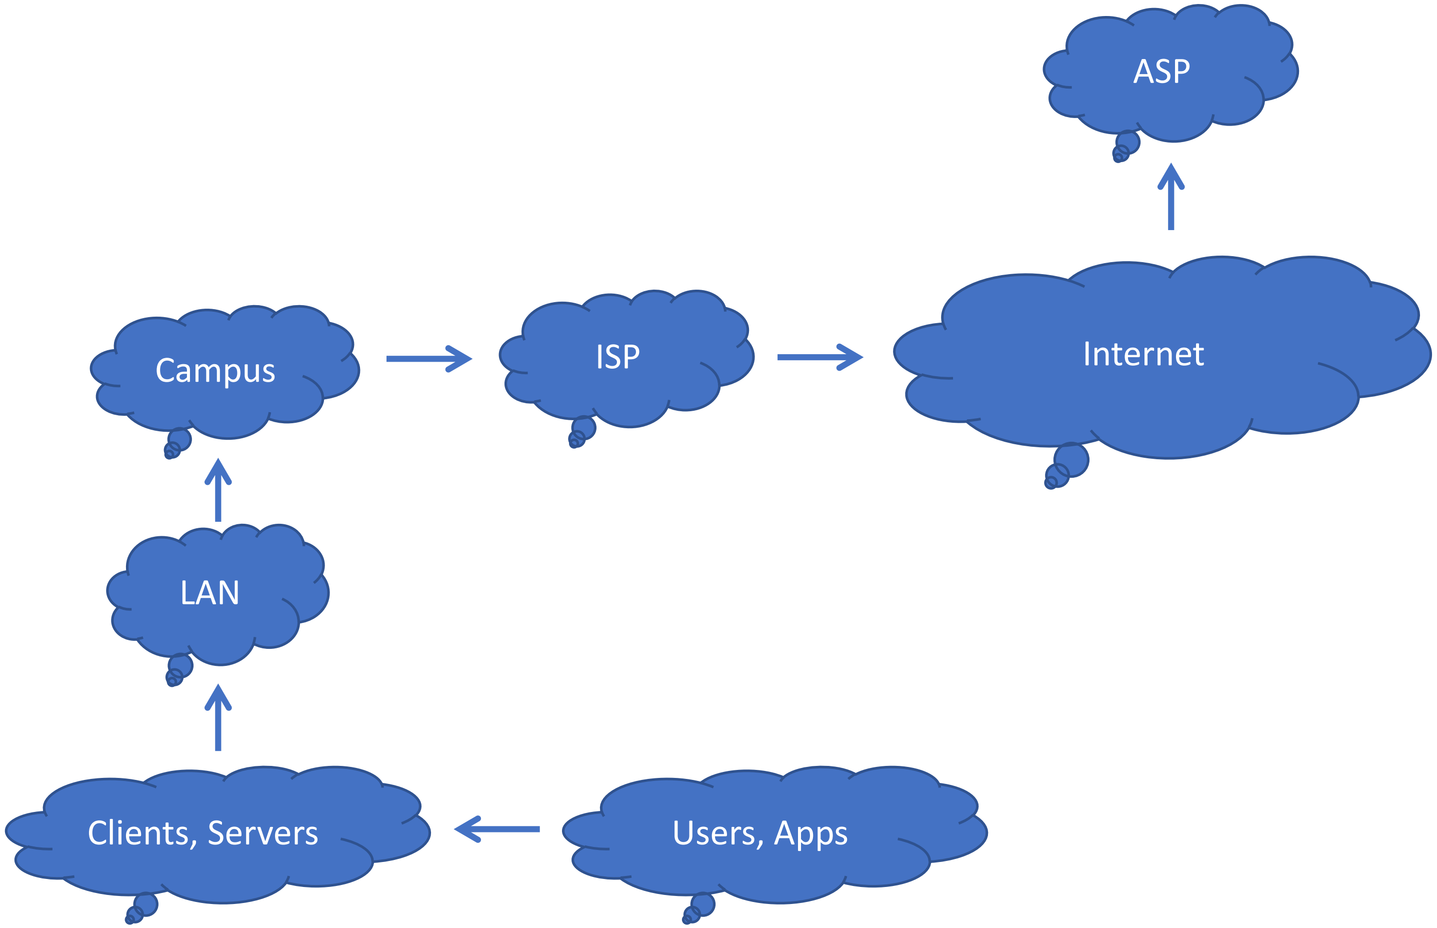

About IP Anycast

- BGP is a link state distance vector multi-path routing protocol
- Reachability is in terms of prefixes, such as 192.5.5.0/24
- Intermediate systems can use the best, or several, or all paths
- Some networks are multi-homed (have more than one path outward)
- Anycast looks exactly like a multi-homed network, as long as:
  - All reachable address+port endpoints are present at every anycast instance
- M-Root was the first root name server to practice IP anycast
  - Two identical servers, each connected to an Internet Exchange Point
- IP anycast is now *de rigeur* – at least, for DNS

Commercialization and Privatization (~1995)

- As the Internet began to outgrow its academic/government origins:
  - The number of connected networks doubled every month for quite a while
  - Most of these new networks did not speak BGP or connect to an IXP
  - Businesses such as UUNET and IIJ were created to service this new market
- Non-technical businesses were never told to run their own RDNS
  - RDNS thus moved away from the LAN/Campus and into the local ISP
  - ISP’s interests were well aligned: caching meant less upstream traffic
  - Nevertheless, running a LAN or Campus RDNS was still common

Economic Products of Centralization

- End systems (hosts, routers, gateways) are rarely upgraded
  - Creates a “long tail” problem which limits systemic innovation
  - Makes deliberate first- or last-mover policies practical
- ISP and ASP systems are rapidly and often upgraded
  - Makes protocols like IPv6, EDNS and DNSSEC more deployable
  - Creates opportunities for abuse of power (surveillance; ad insertion)
- There isn’t a simple, timeless, or universal winning position
  - Like all build vs. buy decisions, centralization is a case by case matter
  - Mistakes will be made; tension will exist; powers will be abused

Anycast RDNS (~2005)

- OpenDNS was created to provide RDNS services to the whole Internet
  - This was seen as innovative and/or controversial at the time
- Early business model included NXDOMAIN redirection
  - So a typographic error in a web browser led to an advertising page
- Another business model was to intercept www.google.com
  - Each search was redirected to Google after keywords were extracted
  - This led directly to Google's investment in RDNS which became 8.8.8.8

Internet System Topology, Revisited

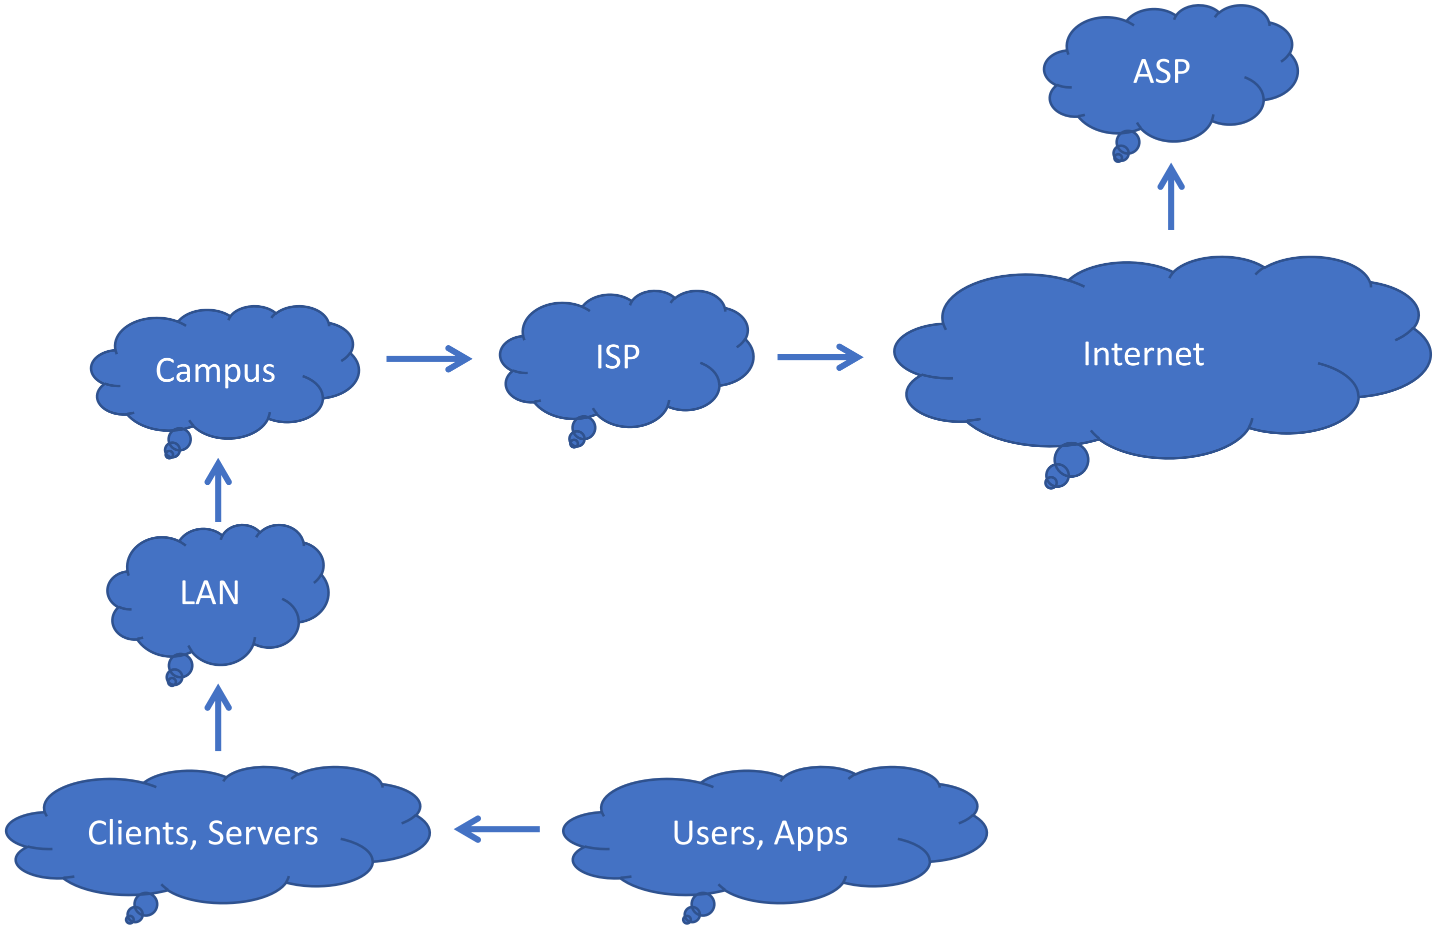

DNSSEC Validation

- DNSSEC is a public-key cryptography system to protect DNS data
  - Publishers generate signatures; validators check those signatures
- It was developed as a way to prevent RDNS cache pollution
  - Because otherwise, UDP/53 can be fooled into accepting off-path answers
- Typical RDNS servers now validate
  - So, a stub has to trust its RDNS to introduce it to sensitive secure services
- DANE is an example of a DNSSEC application
  - It avoids the need for X.509 CA’s, by allowing a DNS data owner to self-sign
- Can a non-contracted party such as a “public DNS” be trusted?
  - Not if they don’t know my local naming, or to introduce me to my own bank

Abuse of Side Effects → Loss of Privacy

- Meanwhile back at the authority servers, enter the CDN
  - Content Delivery Networks wanted to optimize web server selection
- They did this by estimating a browser’s location from its DNS queries
  - However, the DNS queries they received were from RDNS, not stub resolvers
- Anycast RDNS blurred the inputs to this topologic estimation
  - CDN’s therefore pushed for a way to learn the stub resolver’s IP address
- Thus: EDNS Client Subnet (ECS)
  - Increased RDNS implementation and diagnostic complexity
  - Reduced end-user privacy since the “blender effect” was no longer present

Internet System Topology, Extended

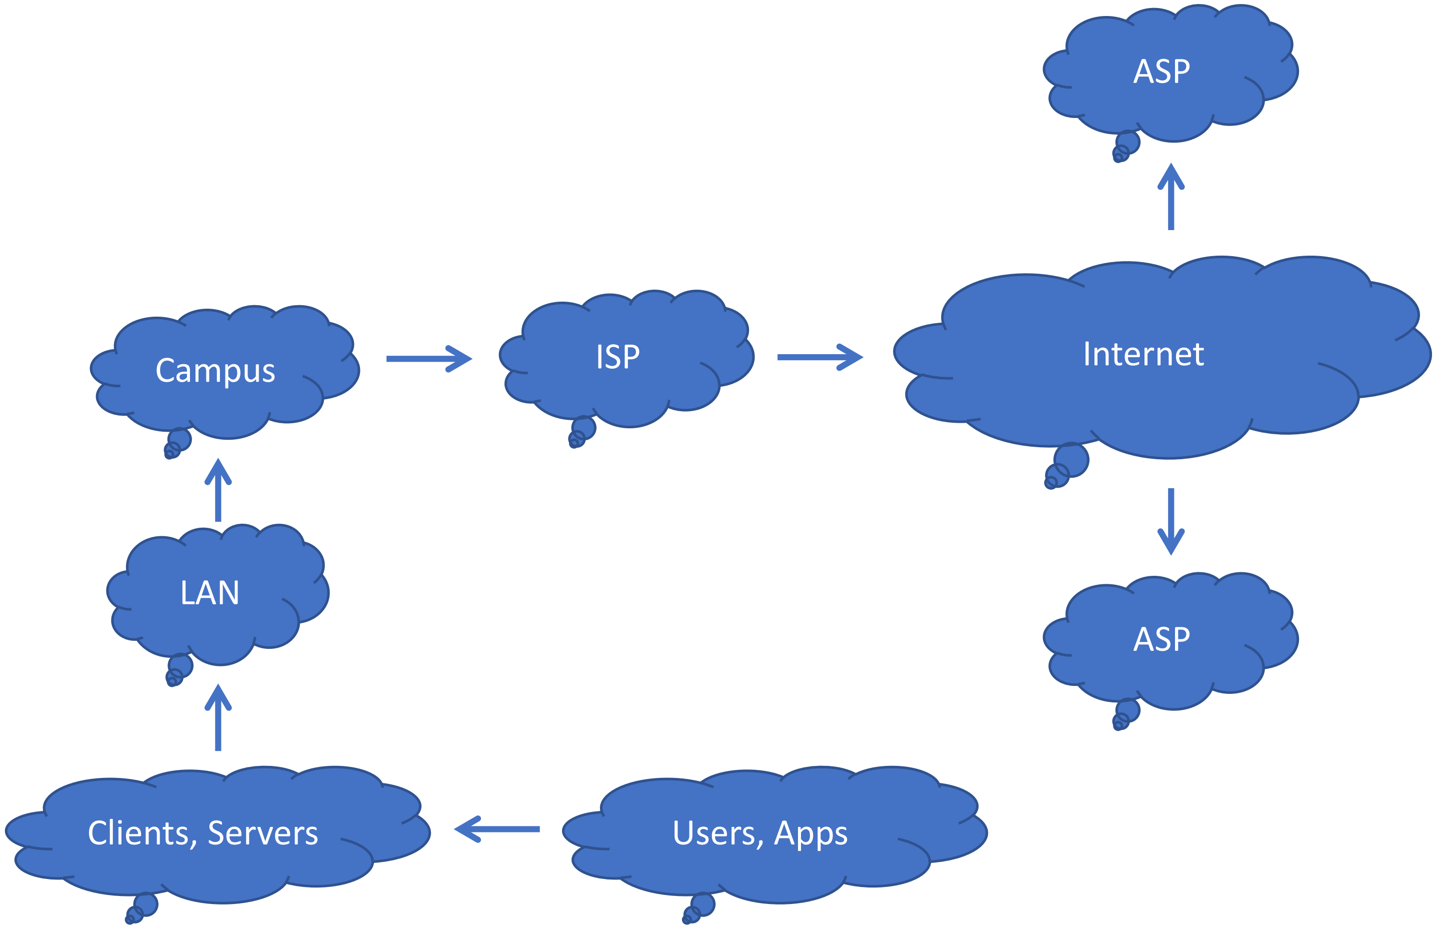

Mobile IP (~2010)

- In this decade, almost all new IP growth is in mobile devices
  - When they use WiFi, their ISP is the home or enterprise network
  - When they use 3G++, their ISP is the phone company
- Regulations for phone companies are still catching up to the Internet
  - Meanwhile, surveillance and interception inside RDNS is often legal
  - The log of stub DNS transactions is very useful for commercial, law enforcement, and both domestic and foreign intelligence purposes
  - The ability to insert false answers is very useful for ad insertion, censorship
  - These powers have been widely abused (see Munk School reports)

Enter Several Kinds of DNS Privacy

- First there was DNS Crypt, which is still supported by OpenDNS/Cisco
  - This protects the stub-to-RDNS data path, but was never broadly adopted
- Then there was DNS Over TLS (DoT), which is being deployed now
  - This is a new transport for any/all DNS transactions, above or below RDNS
  - This is TCP/853, is better than TCP/53, and probably better than UDP/53
  - Network operators can forbid, but cannot surveil or intercept, DoT
- Finally there is DNS Over HTTPS (DoH), also being deployed now
  - This is a new transport for stub-to-RDNS, so, a lot like DNS Crypt
  - Since it uses TCP/443, a network operator will “think twice before blocking it”
  - DoH disintermediates parental controls at home, and company policy at work

DNS System Architecture, Extended

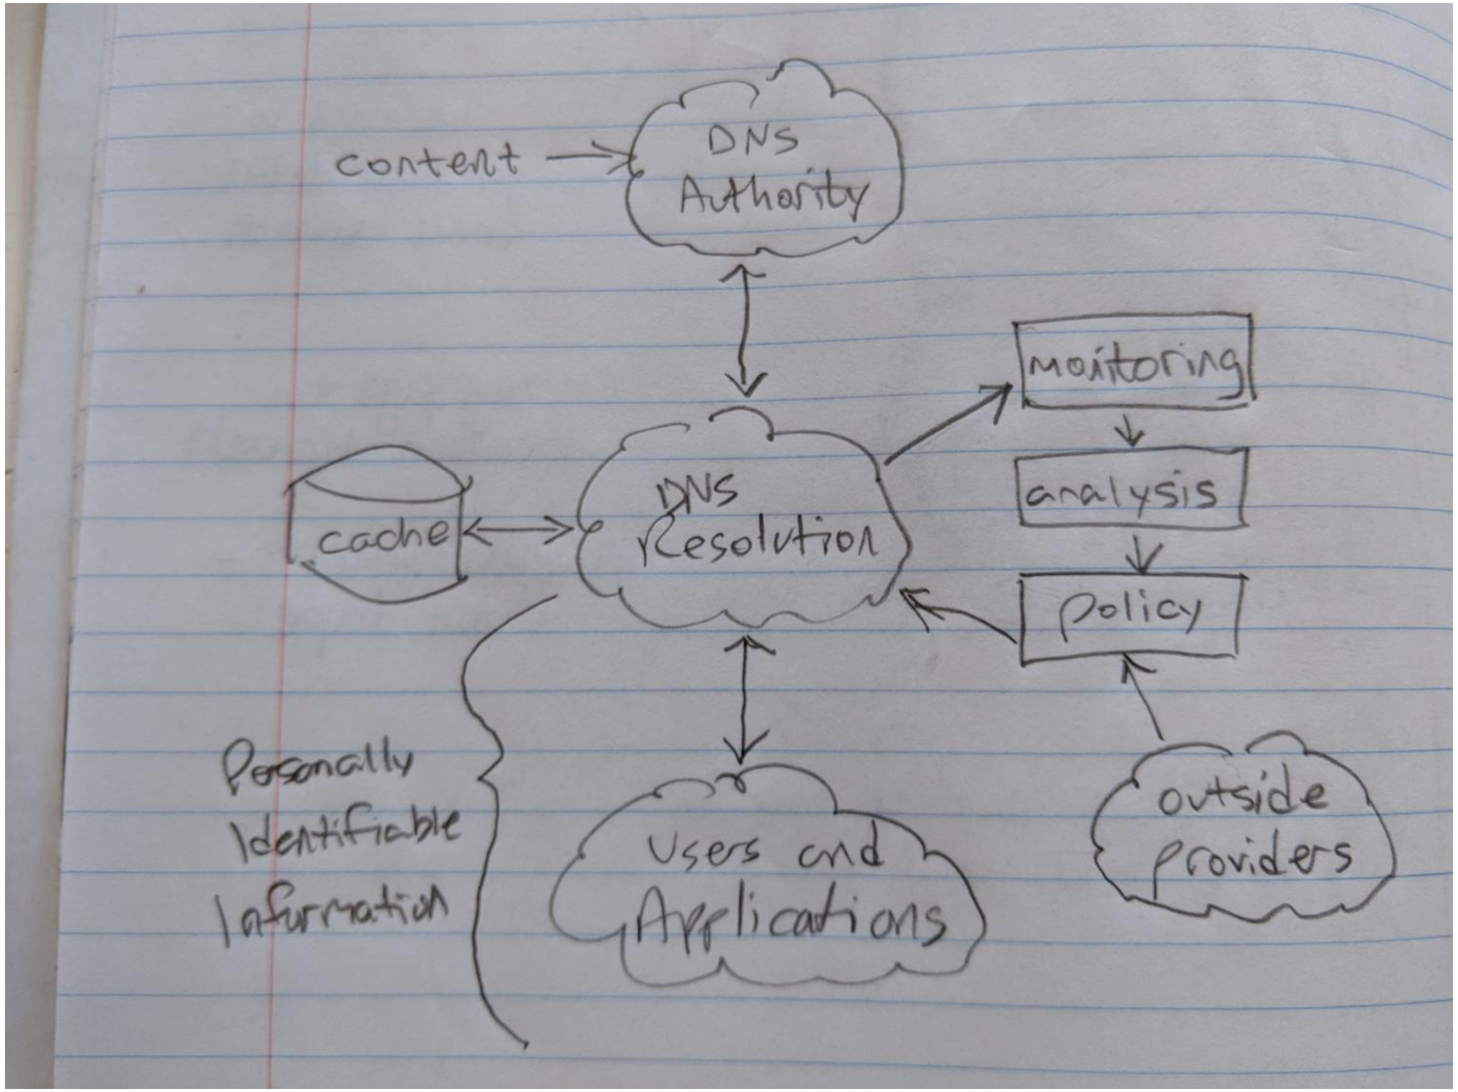

Problems with DoH, part 1

- It’s a political act not a technical one
    - Encrypting stub-to-RDNS but not subsequent flows adds no actual privacy
    - An eavesdropper can guess answers based on what happens afterward
    - Guessing the questions once you know the answers is trivial data science
- To stay out of jail in an authoritarian regime, you need a VPN
    - And once you have a VPN, what value would DoH add?
- Also note, many names are resolvable locally but not remotely
    - Most companies have their own internal-only TLD’s like .CORP or .GOOG
- The web is not the whole Internet; browsers can launch helper apps
    - Helper apps will use the normal stub resolver, getting different DNS answers

Problems with DoH, part 2

- DoH cannot differentiate between three network operators:
  - Parents, who use RDNS filtering as part of their family Internet controls
  - Sysadmins, who use RDNS filtering to block spam and malware
  - Security teams, who use RDNS monitoring to detect new malware infections
- It’s going to become broadly necessary to control TCP/443 (HTTPS):
  - Service networks will proxy or whitelist known-safe external API servers
  - Access networks will add HTTPS MITM, or simply require SOCKS for outbound
  - Any CDN who offers DoH will have to be blacklisted, because of malware

Problems with DoH, Summary

- DoH’s costs would be tolerable if there was an accompanying benefit
  - However, DoH is a political act, adding no actual or effective privacy
- It’s not likely that authoritarian governments will allow DoH
  - Therefore the result of DoH will be more aggressive filtering, not privacy
- Possession is said to be 90% of the law
  - On the Internet that has meant: “my network; my rules”
  - On the Web that appears to mean: “my network; DoH’s rules”
- As a form of Internet governance, DoH shows the worst of all worlds

Now Under Consideration: Resolverless DNS

- Web content providers and their CDN’s want better performance
  - Which means, faster time-to-first-impression
- Most content includes many object references (images, scripts)
  - The time taken for a browser to look up these DNS names is measurable
- Therefore a new IRTF WG is studying “Resolverless DNS”
  - So, DNS data would be “pushed” as part of a normal web content fetch
- No plan so far indicates that DNSSEC signatures would be included
  - Apparently, DNSSEC wasn’t deployed fast enough to seem relevant?

Mozilla’s Plans For Firefox (as of March 2019)

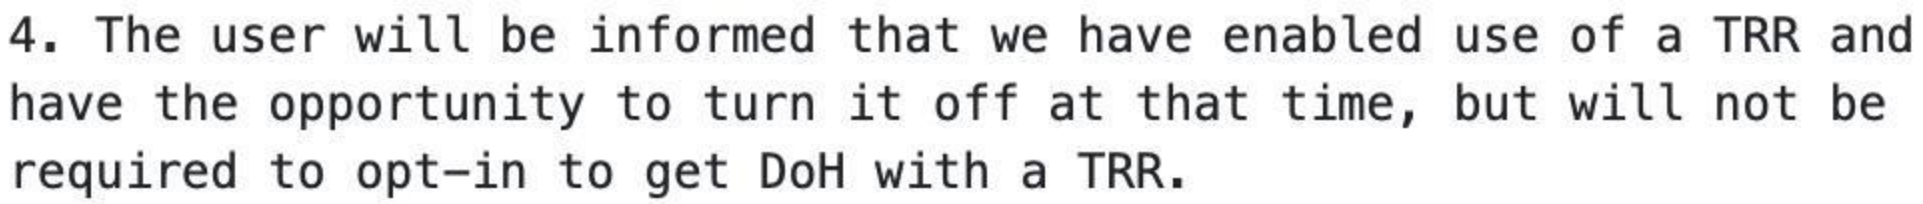

4. The user will be informed that we have enabled use of a TRR and have the opportunity to turn it off at that time, but will not be required to opt-in to get DoH with a TRR.

Cooperation Is Alignment

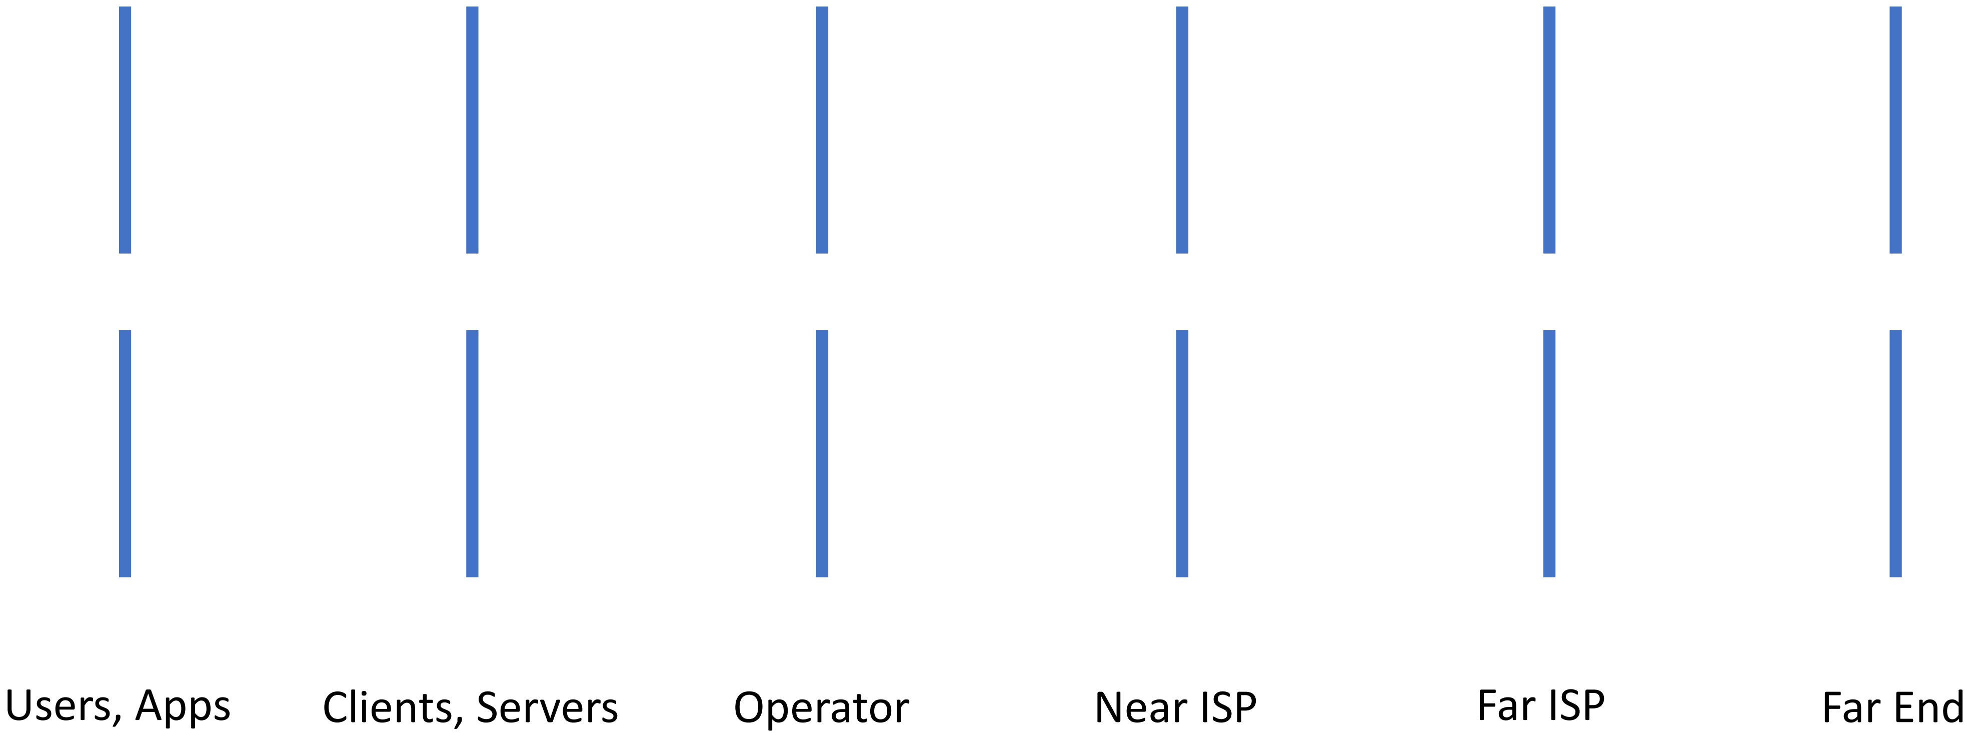

End Notes

- Every innovator solves the problems their customers have
  - Not every innovator knows or cares about systemic costs
- Time to market, not quality, is the primary success trigger
  - Total revenue before obsolescence is the primary success metric
- Adding complexity to a system will externalize complexity’s costs
  - One estimate is that the total lifetime cost of an “IF” statement is USD 10K
- DNS is the first and only system of its kind that has scaled to 10^9
  - Distributed, coherent, reliable, autonomous, and hierarchical
- Keeping DNS working is not a simple task on the easiest day
  - The war for control over the DNS resolution path is costly and damaging

Current Events: SIE Europe

- Non-profit limited liability company (UG), incorporated in Germany
  - Founders/owners: Paul Vixie (FSI), Christoph Fischer (BFK), Peter Kruse (CSIS)
  - BFK (hosting) and FSI (devops) will operate the infrastructure (cost recovery)
- Collect data from European participants (academic and commercial)
  - Raw and filtered real time data will remain inside GDPR-governed territory
- Give a little; get a lot: participants will get access to combined data
  - Filtered or deduped (LAN/UDP or AXA/TCP); or stored (DNSDB API); no “raw”
- Constraints on redistribution and derivative works
  - No redistribution of raw, filtered, deduped, or stored forms except by FSI
  - Derivative works must remain in Europe or be artificially delayed
- Pre-launched March 2018; official launch October 2018; now live
  - Learned so far: we’ll need a separate raw channel for sandbox (inhuman) traffic