DNS Wars: Episode IV

A New Bypass

Dr. Paul Vixie, CEO
Farsight Security, Inc.

2019-09 EuroBSDCon

Abstract

- Due to pervasive unpreparedness of users, applications, operating systems, and protocols, DNS has become an essential control point for “cyber” security. Most networks have a mix of legacy, modern, safe, and unsafe devices attached to them, and this condition won’t change as quickly as the Beyondcorp initiative might suggest. However, DNS is also an important control point for authoritatian regimes, and so “bypass” innovation is continuous, rapid, and ambitious. Special attention is deserved by the “DNS over HTTP” or “DoH” protocol now being strongly pushed by Mozilla, CloudFlare, and others. A brief mention will be made of IRTF Resolverless DNS.

```
[fbsd.local:amd64] telnet www.fsi.io 444
Trying 104.244.14.108...
^C
```

```
[10.0.2.15].50039 => [104.244.14.108].444 (TCP SYN)
pdns: "archive.farsightsecurity.com"
pdns: "web1.iad1.fsi.io"
pdns: "farsightsecurity.com"
pdns: "fsi.io"
pdns: "www.farsightsecurity.com"
pdns: "www.fsi.io"
pdns: "fastrpz.com"
[10.0.2.15].50039 => [104.244.14.108].444 (TCP SYN)
pdns: "archive.farsightsecurity.com"
pdns: "web1.iad1.fsi.io"
pdns: "farsightsecurity.com"
pdns: "fsi.io"
pdns: "www.farsightsecurity.com"
pdns: "www.fsi.io"
pdns: "fastrpz.com"
```

```
/* Pseudo code:
 *
 * get recently observed domain names pointing to this IP.
 * for each such domain name:
 *      try DoH.
 *      if successful:
 *              send RST back to initiator.
 *              add this address to the "no go" IPFW table.
 *      else:
 *              resubmit SYN packet in outbound direction.
 *              add this address to the "ok, go" IPFW table.
 */
```

```
flush

table   no-doh  create type addr
table   no-doh  unlock

table   go-doh  create type addr
table   go-doh  unlock

add     pass            tcp     from any to table(go) 444 setup out
add     deny            tcp     from any to table(no) 444 setup out
add     divert 444      tcp     from any to any 444 setup out

add     pass            all     from any to any
```

```
/* Read - eval - print. */
while ((x = recvfrom(divert_socket, inpkt, sizeof inpkt, 0,
                     (struct sockaddr *)&from, &fromlen)) > 0)
{
        const char *msg;

        if ((msg = input(inpkt, x, from, fromlen)) != NULL)
                fprintf(stderr, "input: %s\n", msg);
}
```

```
char *name;

pdns_namebyaddr_ask(ip->ip_dst);
while ((name = pdns_namebyaddr_next()) != NULL) {
        if (want_hdrdump)
                printf("pdns: \"%s\"\n", name);
        free(name);
}
```

DNS System Architecture (Traditional)

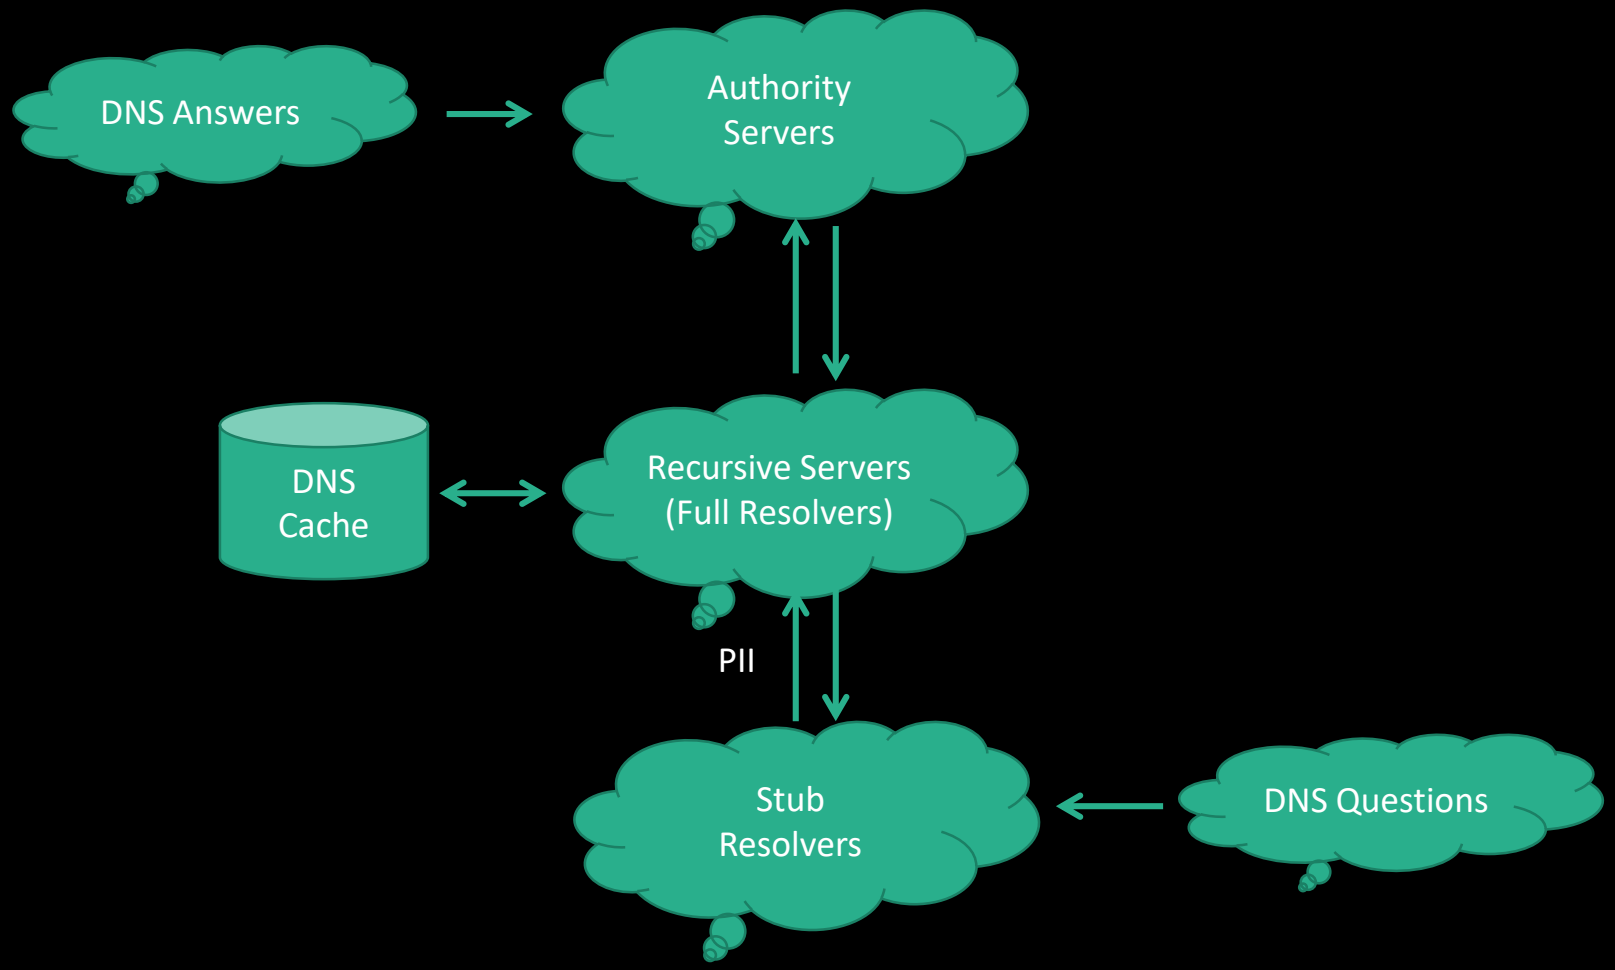

Internet System Topology, ~1999

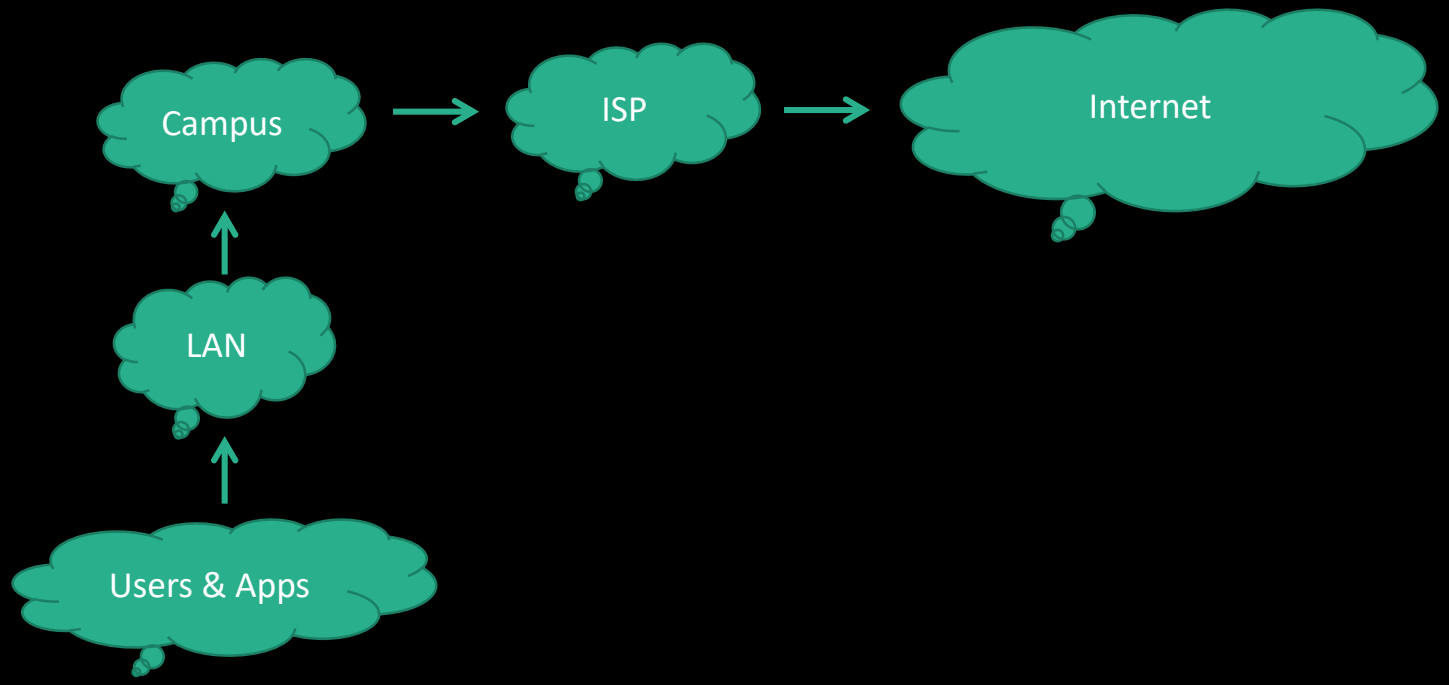

VeriSign™ SiteFinder™ (Episode I)

- SiteFinder is in the advertising business
- VeriSign is in the DNS business
  - So, *.COM happened
  - Then delegation-only happened
  - Then delegation-only-except happened
- All ICANN SSAC members were sued, along with ICANN itself
  - Cool t-shirts created by Olaf Kolkman (then at RIPE)
  - Lawsuit resolved – ploy for .COM contract negotiations?

ISP and Anycast RDNS (Episode II)

- 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
  - Noting, ISP's were doing this on-path long before OpenDNS did it
- OpenDNS also intercepted lookups for 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
- There are roughly 200 more quad-N's available
  - Get yours while they last!

Abuse of Side Effects (Episode III)

- 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
  - So privacy-abrasive that not even CloudFlare (1.1.1.1) supports it

Internet System Topology, ~2019

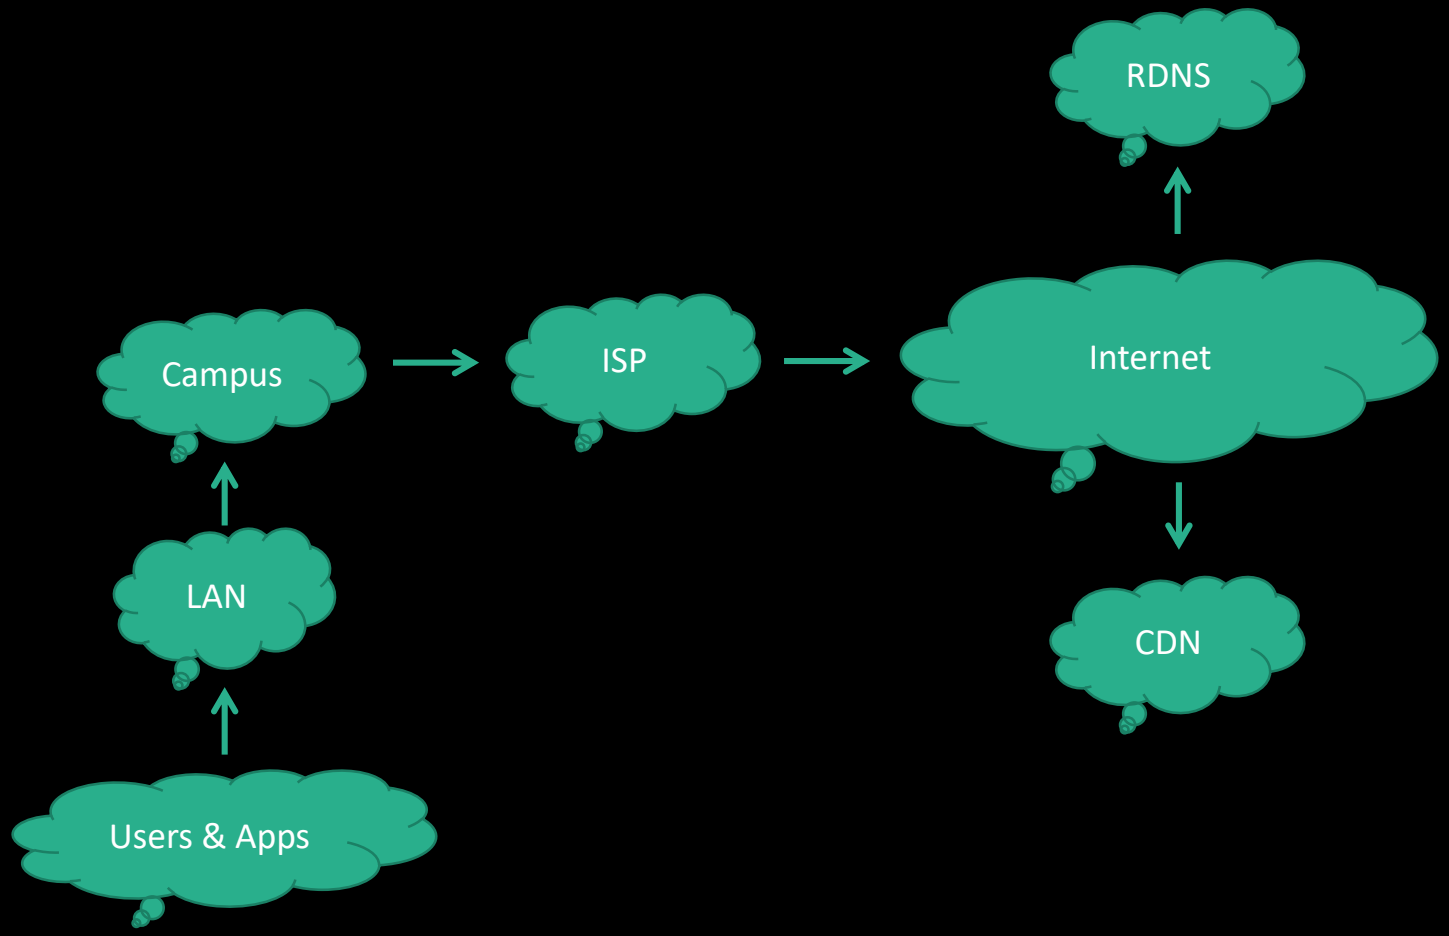

Several Kinds of DNS Privacy

- 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, better than TCP/53, and w/ TCPFO often better than UDP/53
  - Network operators can forbid, but cannot surveil or intercept, DoT
- Next came DNS Over HTTPS (DoH), because, why not?
  - This is a new transport for stub-to-RDNS, so, a lot like DNS Crypt
  - Since it uses TCP/443 a network operator may “think twice before blocking it”
  - DoH disintermediates parental controls at home, and company policy at work

DNS System Architecture, As Amended

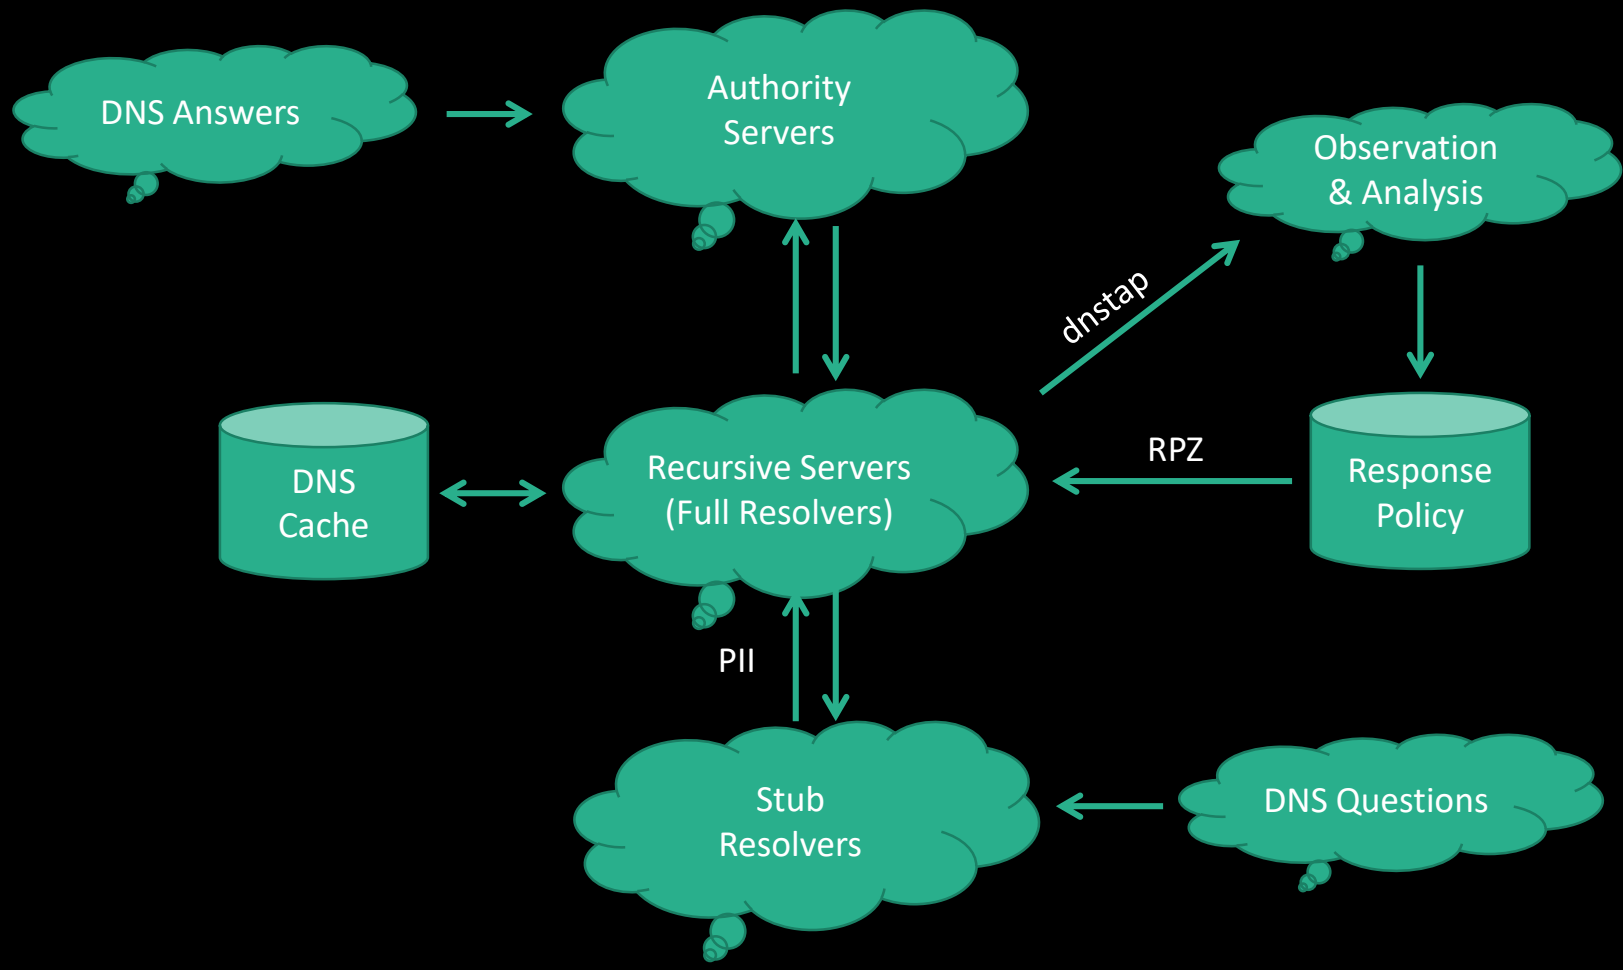

Problems with DoH, part 1

- It’s a political project, 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 these 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
  - Authoritarian government, who uses RDNS for “thought control”
- It’s going to become broadly necessary to control TCP/443 (HTTPS):
  - Service networks will proxy or whitelist known-safe external API servers
  - Networks can no longer HTTPS MITM, so, require proxy for all outbound?
  - Any IP offering DoH will be widely blacklisted, because of malware
  - This increases complexity, cost, and vulnerability for almost every network

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
- Some people think CCP has theoretical resource limits for GFW
  - Some people don't
- 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-next-eyeball
- 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?
- Remembering the DNS rebinding attacks makes this controversial
  - Any browser that looks at DNS <META> headers won't reach the enterprise

In 2019, Polite Behaviour Has Been Redefined

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

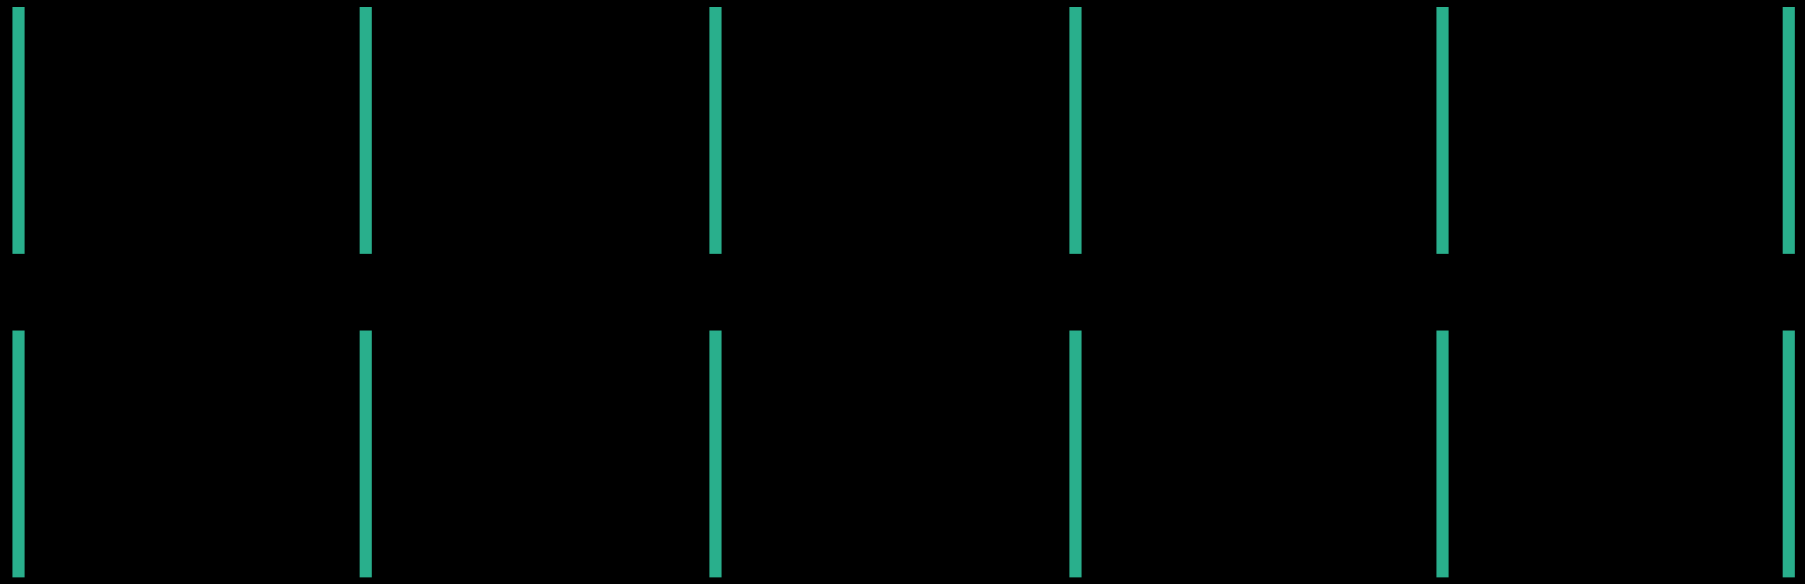

Expensive (Imposed) Choices

- Faced with Internet Standards for RDNS bypass, a NetOp can either:
  - Allow malware, intruders, supply chain poison, BYOD to bypass DNS controls
  - Stop thinking any network can ever be secure, move beyond “perimeters”
  - Create smaller networks having explicit whitelists for must-be-reached
  - Allow Chromecast, Chrome, IoT unlimited access to their motherships
  - Proxy *everything*, effectively strip-searching outbound at the perimeter
- …Or:
  - Follow the tradition that possession is 9/10th of the law
  - Establish an AUP and enforce it for all outbound communications
  - Get creative about what (few) requires a proxy and what (many) does not

Consent Matters

- In ~1995, MAPS said “all communications should be consensual”
  - Many spammers disagreed, and the culture war was lost due to legal fees
- Most people just want to extract some value and punch out
  - The Internet has evolved to be the perfect accountability launderer
- CDNs sometimes claim that they aren’t hosting, only caching
  - This is a grave, ugly abuse of the DMCA’s safe harbor provisions
- Whois was already 90% dead due to ineffective governance
  - So I’m not ready to blame GDPR for finishing off the remaining 10%
- I am not yet ready to build and fund infrastructure for the opposition
  - *My network, my rules* – or else, whose network is it, actually?

Demo, using FreeBSD, IPFW, and IPDIVERT

- I am exploring ways to use minimal force to defend networks
  - In particular, I won't be provoked into overreaction (proxy by default)
- Whistle Communications gave FreeBSD a lot of terrific technology
  - IPDIVERT, originally to support NAT, turns out to have other uses also
- Passive DNS means that gethostbyaddr() isn't as dead as we thought
  - PTR records are waning, but Name ⇒ IP ⇒ Name is still somewhat possible
- TLS 1.3 and Encrypted SNI prevent traditional enterprise MITM
  - But, enumeration attacks may remain practical with some crowd sourcing

Real Persons

- www.axios.com, 2019-09-19:
  - «DoH advocates argue that their preferred protocol has a key advantage over DoT. DoH uses the same pathways as web browsing, making it impossible to block without blocking all web browsing. DoT doesn't disguise itself that way.

    But Vixie believes that puts the security of the few over that of the many. “With DoH, they are solving a problem that most of the world doesn't have by creating a problem that everyone in the world will have,” he said.»

Corporate Persons

- Ibid:
  - «Mozilla says that many concerns are already being addressed on its end. "Our deployment plan will disable DoH if parental controls are in place," said Selena Deckelmann, senior director of engineering, adding the same will be true when Firefox detects certain security products.»
  - «And Cloudflare notes that parental filters that operate before starting to connect to a website will still work. "Someone looking to use DoH to keep their web browsing data private can apply parental filters or security products on their DoH endpoint," said Alissa Starzak, Cloudflare head of policy.»

Unreal Persons

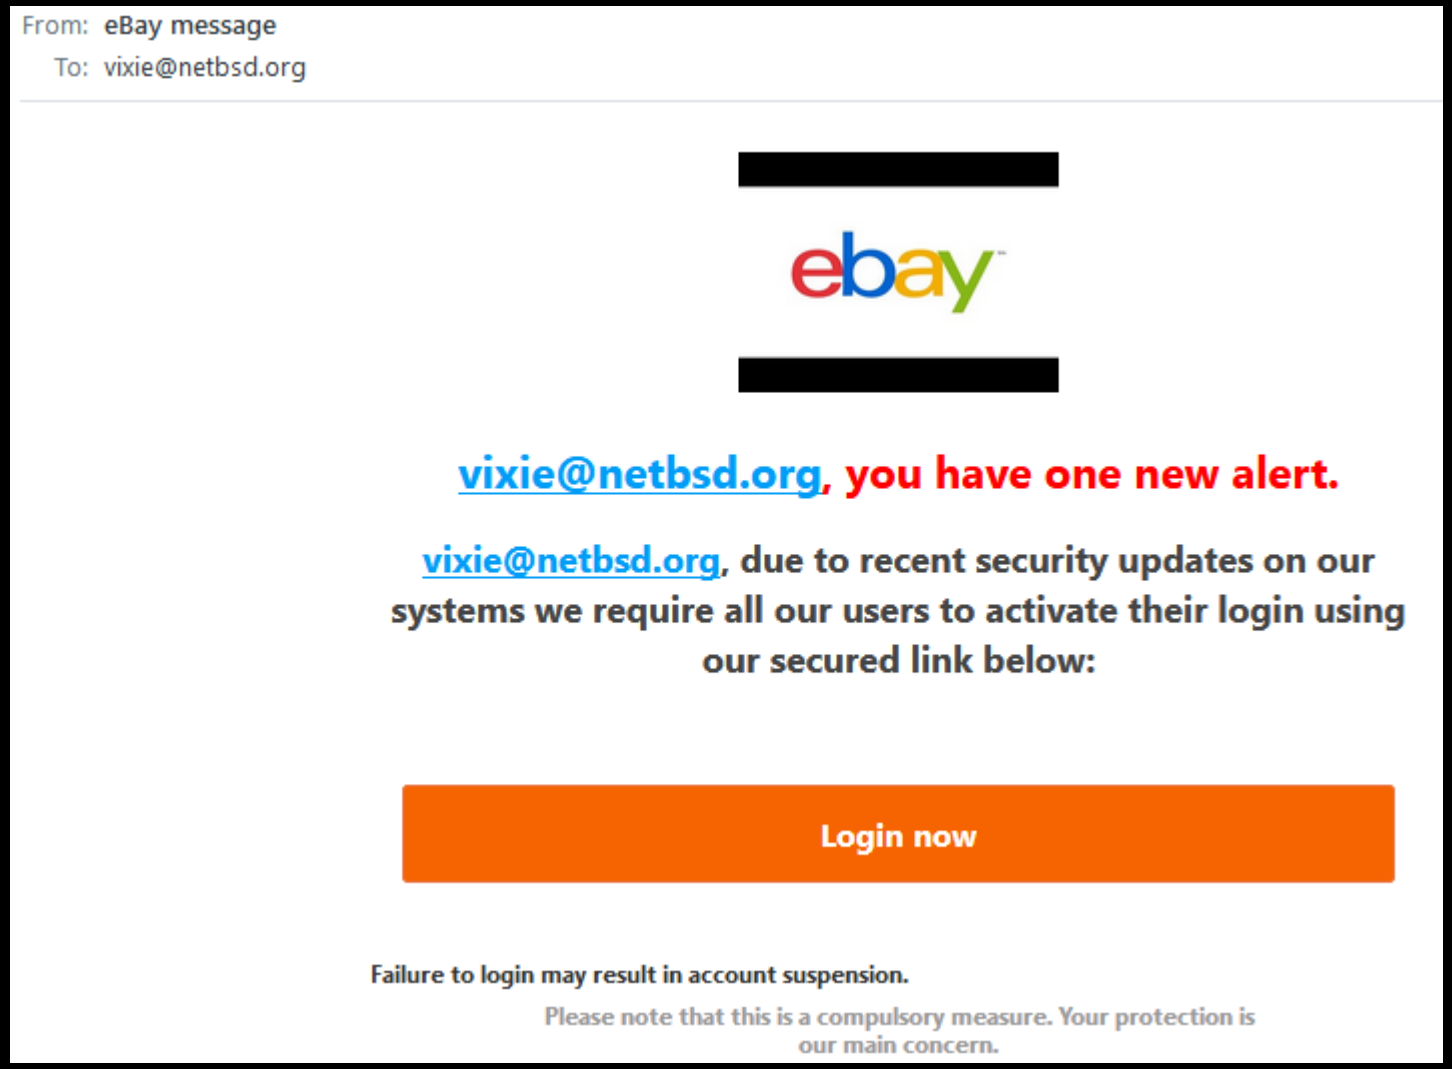

End Notes

- Every innovator solves the problems their/they customers have
  - Not every innovator knows or cares about systemic costs
- DNS is the first and only system of its kind that has scaled by 10^9
  - Distributed, coherent, reliable, autonomous, and hierarchical – unique!
- 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
- As in politics, economics, and climate change, this future is brutal
  - Our consent is no longer sought, and can only be withheld at notable cost