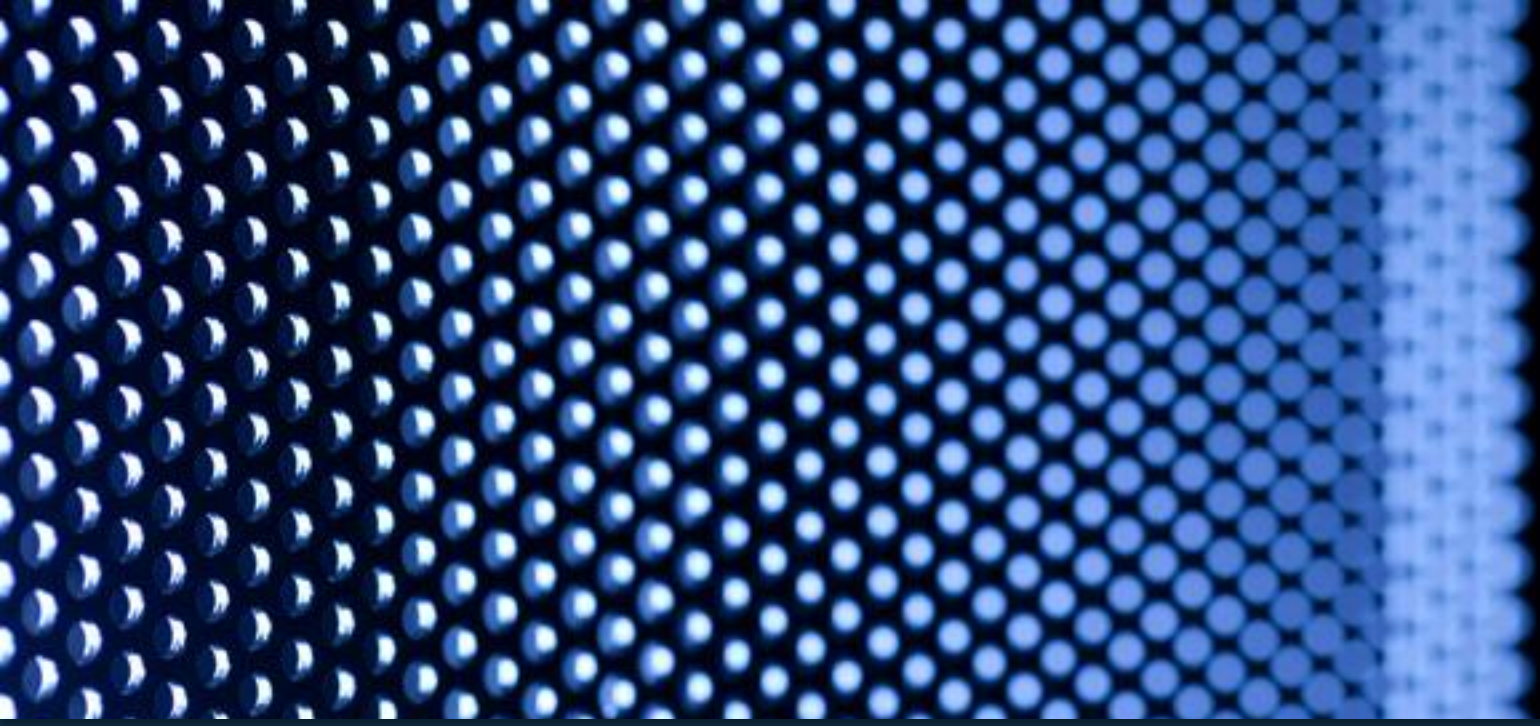

VERISIGN

DNS Privacy Overview

Allison Mankin & Shumon Huque, Verisign Labs

DNS-OARC Fall Workshop

October 3, 2015

Background

- DNSSEC (RFC 4033) specifically has no confidentiality requirement
- DNSSEC did consider a privacy requirement (avoidance of zone enumeration) in adding NSEC3 to the extensions
  - Consistent with guidance and protocols for confidentiality for zone transfers
- Outside IETF, services such as dnscurve and dnscrypt offered confidentiality
  - Did not get on standards radar
- This changed with PERPASS effort and its output, RFC 7258
- IETF formed DNS Private Exchange (DPRIVE) WG in 2014
- DPRIVE has just issued its first RFC, DNS Privacy Considerations (RFC 7626)

RFC 7258 - Pervasive Monitoring is an Attack

- Essential message conveyed by its abstract (entirety):
  - Pervasive monitoring is a technical attack that should be mitigated in the design of IETF protocols, where possible.
- Focus on meta-data in addition to data plane
  - Some attention to this previously, such as IPv6 privacy addresses
  - Renewal of focus and effort
- Consider broad range of risks
  - Protocol design issues
  - Interactions/intersections between protocols
  - Side channels – for example, size- and timing-based information leakage

RFC 7624 – Confidentiality Threat Model

- Follows on from RFC 7258
- More detail and terminology
- More linkage to the Privacy Considerations Best Current Practices (BCP) (RFC 6973)
- Background and bibliography on in-the-wild Pervasive Monitoring
- Places DNS privacy in broad context (3.1, 3.2, 3.3.2, 5.2)

RFC 7626 – DNS Privacy Considerations

- Expert coverage of risks throughout DNS ecosystem
- Linkage to RFC 6973 (Privacy Considerations for Internet Protocols)
- Rebuts “alleged public nature of DNS data”
- Covers:
  - Targets in the DNS data
  - Places in the DNS ecosystem where data may be tapped
  - Places in the DNS ecosystem where data is collected, that may be misused or compromised
  - Indirect sources of privacy disclosure such as cache snooping (timing probes)

Privacy Evaluation

- An individual draft
- Presentations in DPRIVE at IETF-91, IETF-92, and IETF-93
- Attempt to connect IETF efforts with privacy formalisms
- Supports quantitative evaluation of privacy methods (on their own or combined)

- draft-am-dprive-eval-01.txt

Overview of DNS Privacy Risks

DNS Privacy Risks

- DNS data may be at risk of disclosure:
  - Between client and recursive
  - At recursive name server
  - Between recursive and authoritative
  - At authoritative name server
- Data may also be at risk of modification: privacy risk if client misdirected
- Important to consider such risks as part of overall privacy strategy
- Presentation will be light on modification/DNSSEC angle

Risk 1: Between Client and Recursive

- Client effectively reveals browsing history via DNS traffic to recursive name server
- Adversary must be “on path” to see it, but it’s all in one place
- Risk increases when recursive name server deployed outside organization
- How to protect against eavesdropping?

Risk 1: Between Client and Recursive

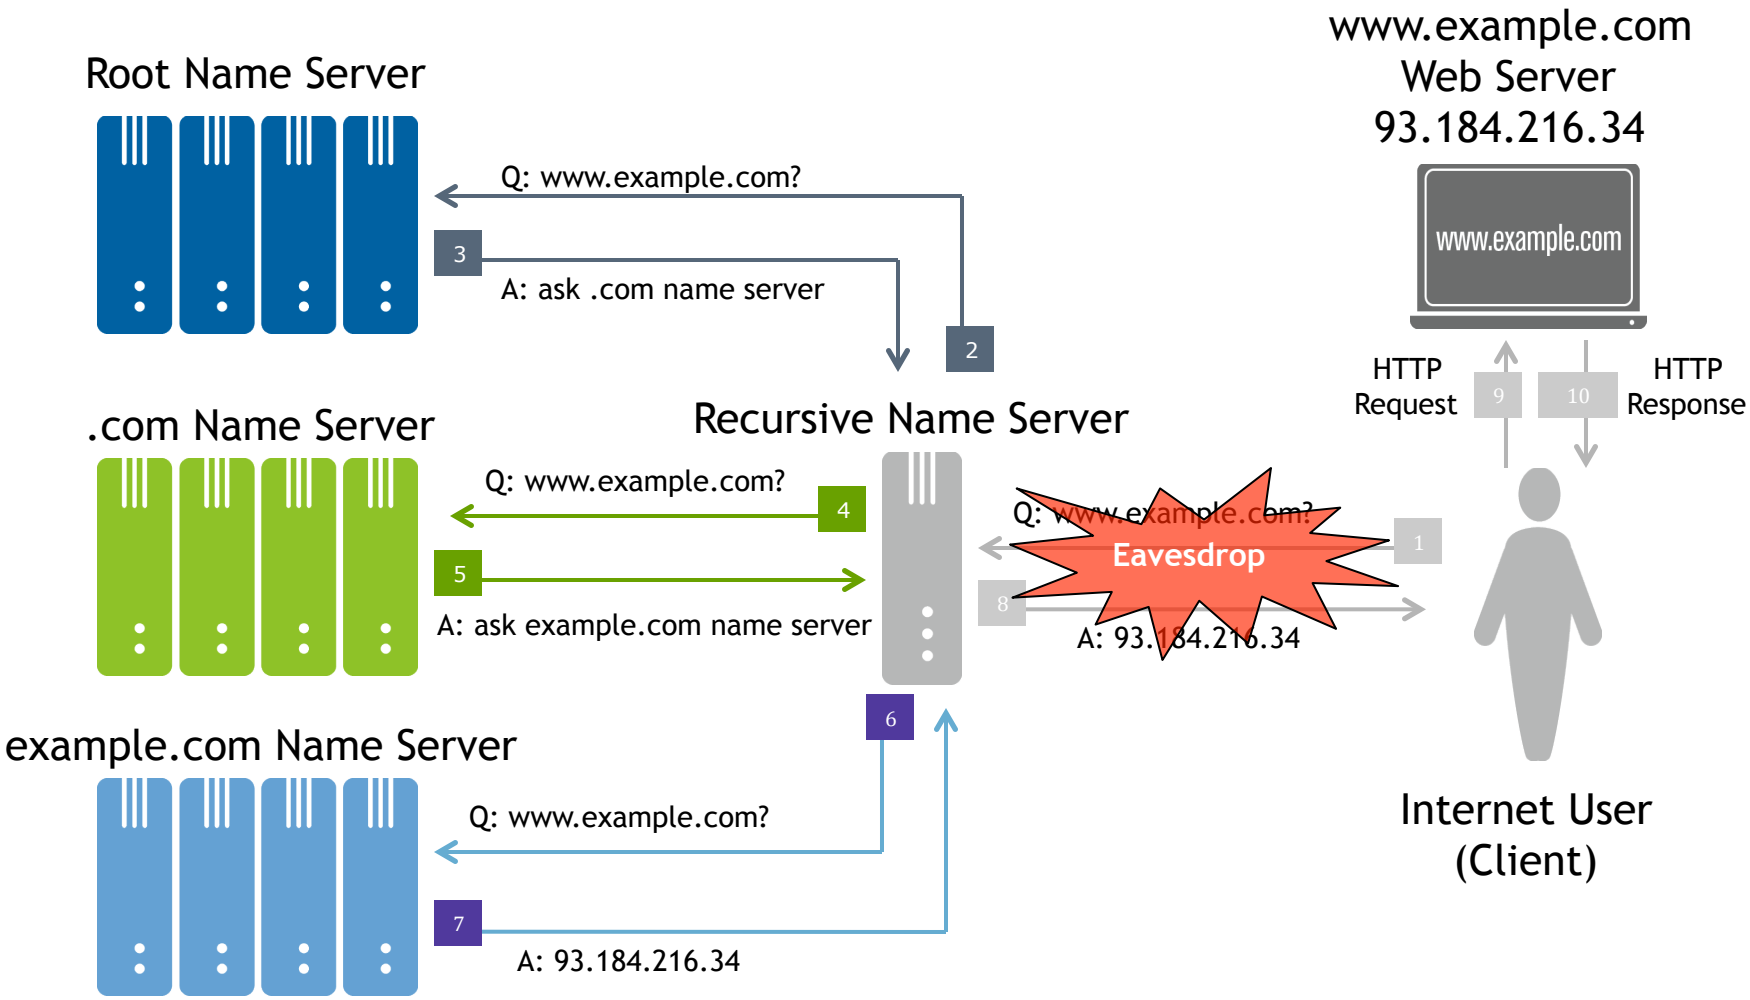

Risk 2: at Recursive Name Server

- Recursive name server learns client's browsing (and other) history through its DNS traffic
- Adversary may compromise server systems to get this data
- Server itself may be "adversary," misusing data …
- How to protect against compromise, misuse?

Risk 2: at Recursive Name Server

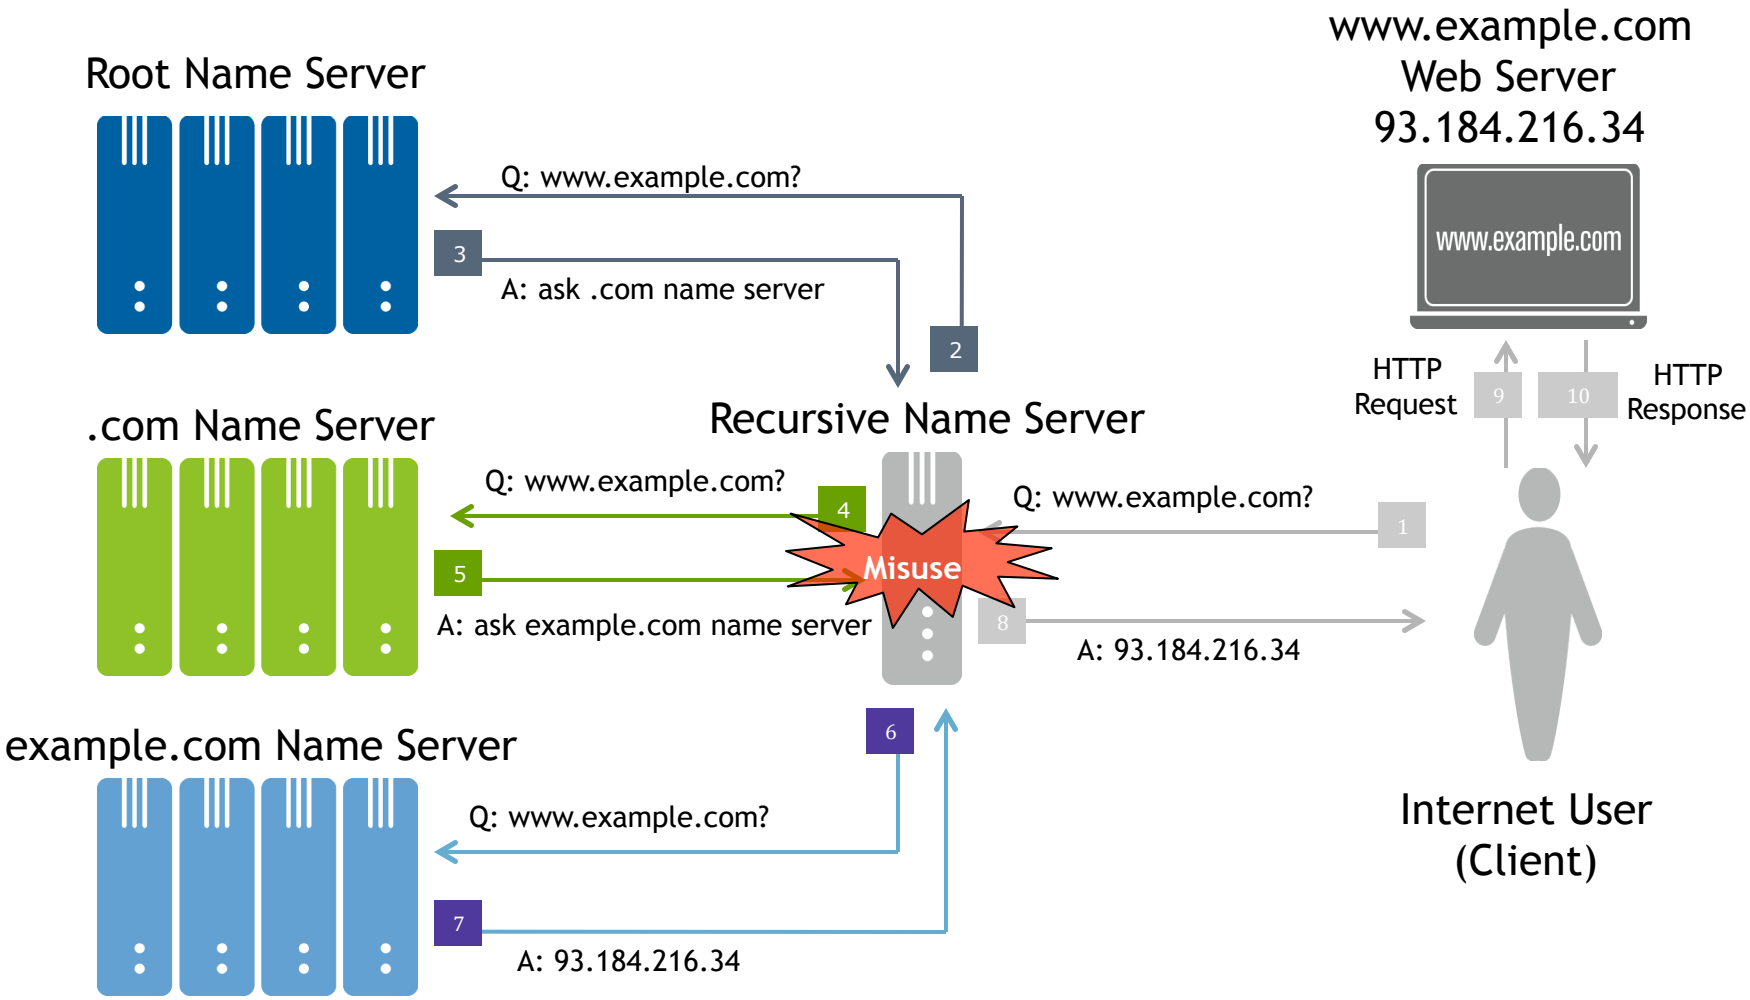

Risk 3: Between Recursive and Authoritative

- Recursive name server reveals samples of *community's* lookup history via DNS traffic to authoritative name servers
- Adversary again must be "on path" to see traffic, but all in one place
- Authoritative name servers by definition deployed *outside* organization
- How to protect against eavesdropping?

Risk 3: Between Recursive and Authoritative

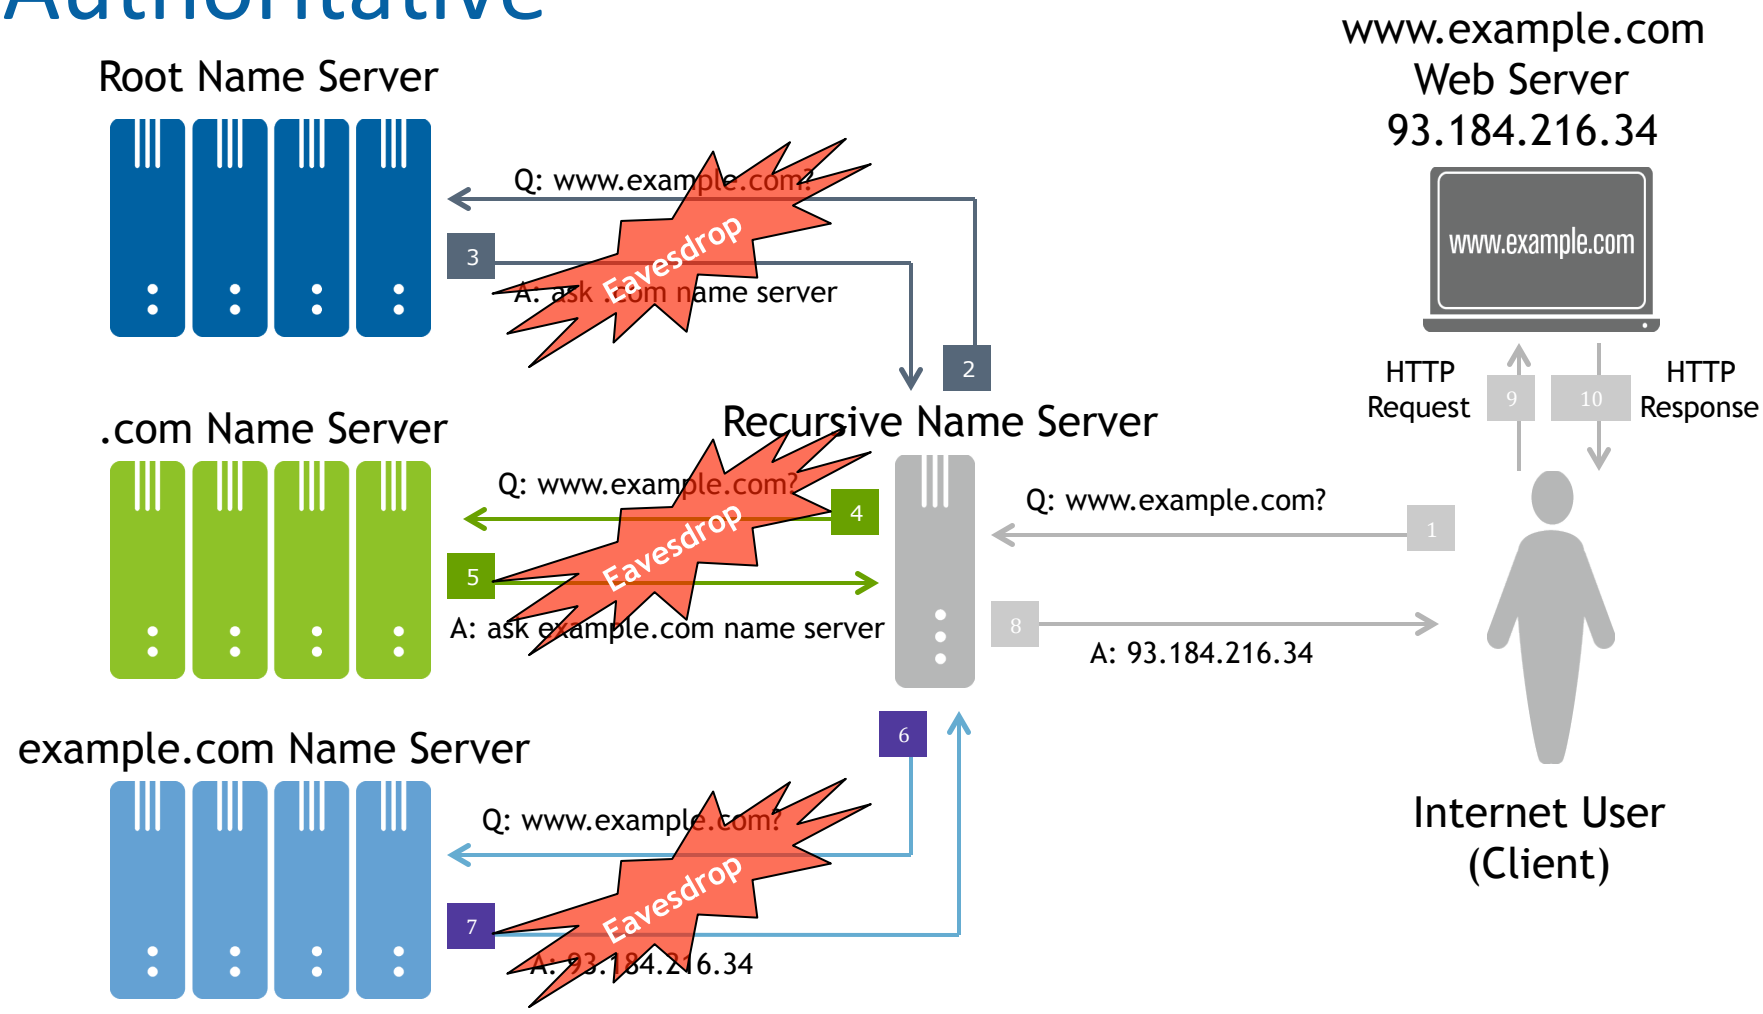

Risk 4: at Authoritative Name Server

- Authoritative name server learns samples of recursive's community's browsing history
- Adversary may again try to compromise server systems to get this data
- Server itself may again be "adversary"
- How to protect against compromise, misuse?

- A hybrid risk: authoritative server learns recursive client's identity via the use of *edns-client-subnet* option by the intervening recursive server. This is done normally for service optimization purposes, but nonetheless represents a privacy leakage.

Risk 4: At Authoritative Name Server

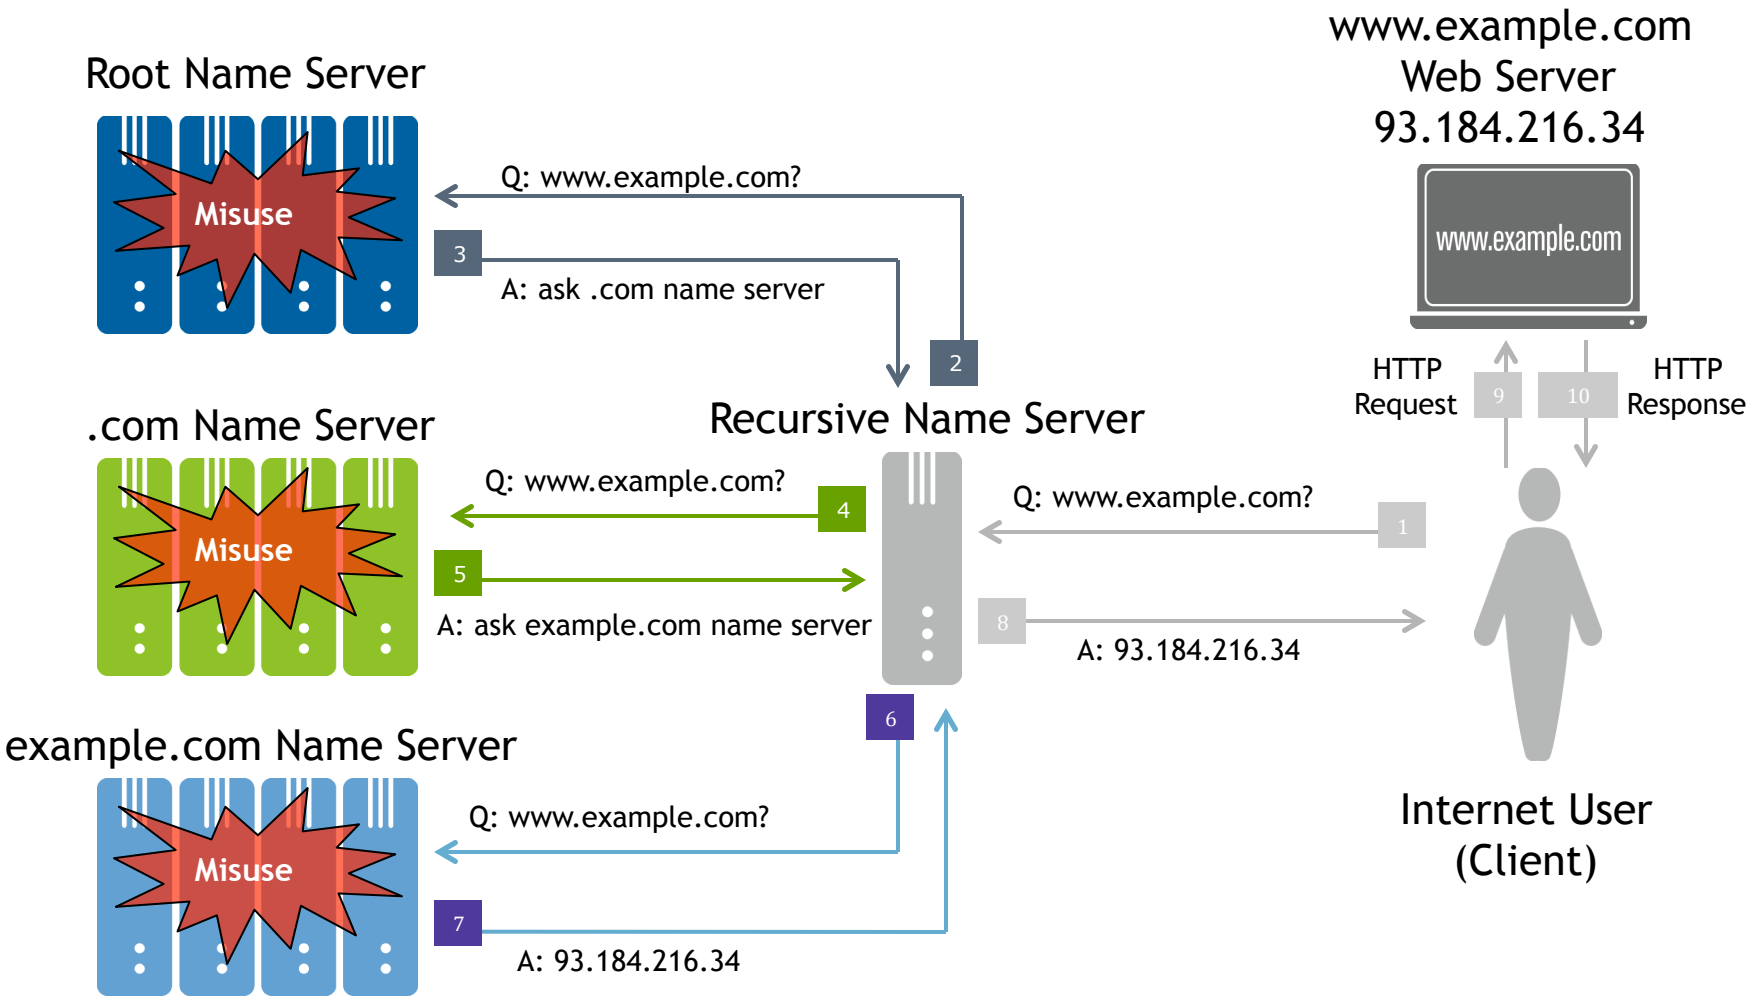

Summary of DNS system risks

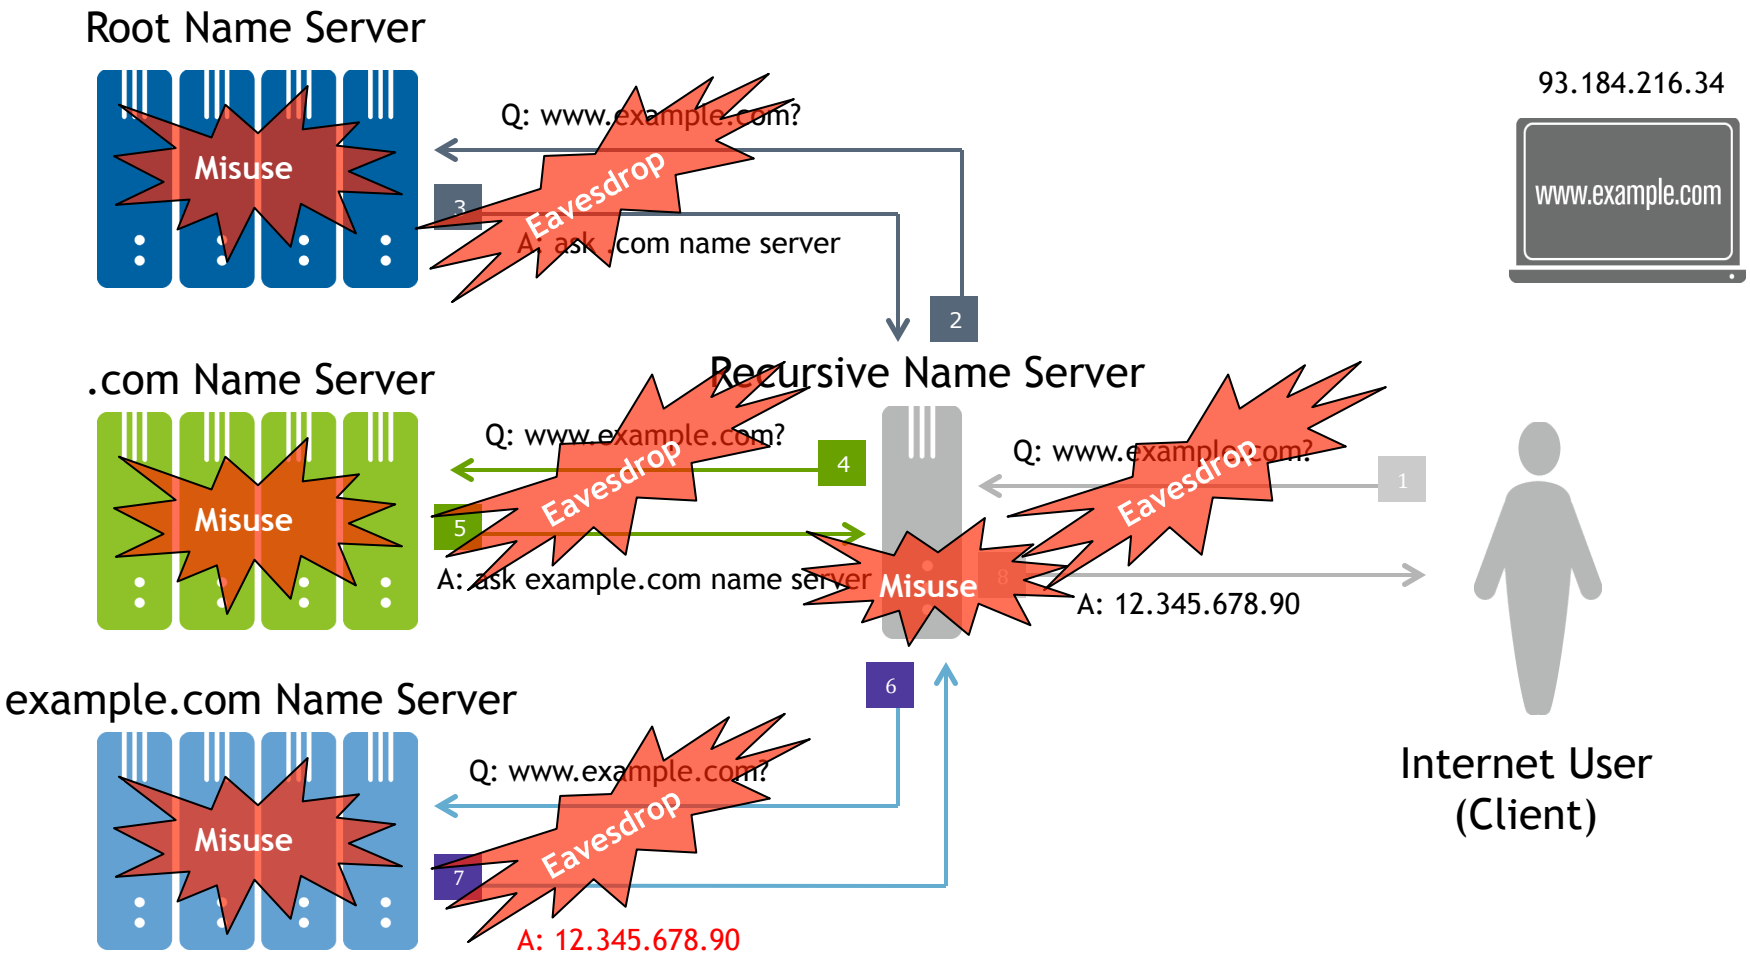

Overview of Mitigations

Mitigating DNS Privacy Risks

- Data handling policies can help mitigate the risks
- Technical enhancements to DNS have also been introduced & proposed in recent years to mitigate these risks:
  - DNS-over-TLS
  - qname-minimization
  - DANE and DNSSEC*

- (*DNSSEC might help in the sense that unauthorized modification of DNS traffic can present a privacy risk if a client is misdirected to a resource in the control of an adversary.)

Mitigation 1: DATA HANDLING

- Data handling policies, technologies and audits can mitigate risk of compromise, misuse of data at recursive, authoritative servers
- Root, top-level domain servers generally operate under established agreements
- Other authoritative name servers, recursive name servers may not

Risks 2 & 4: Misuse

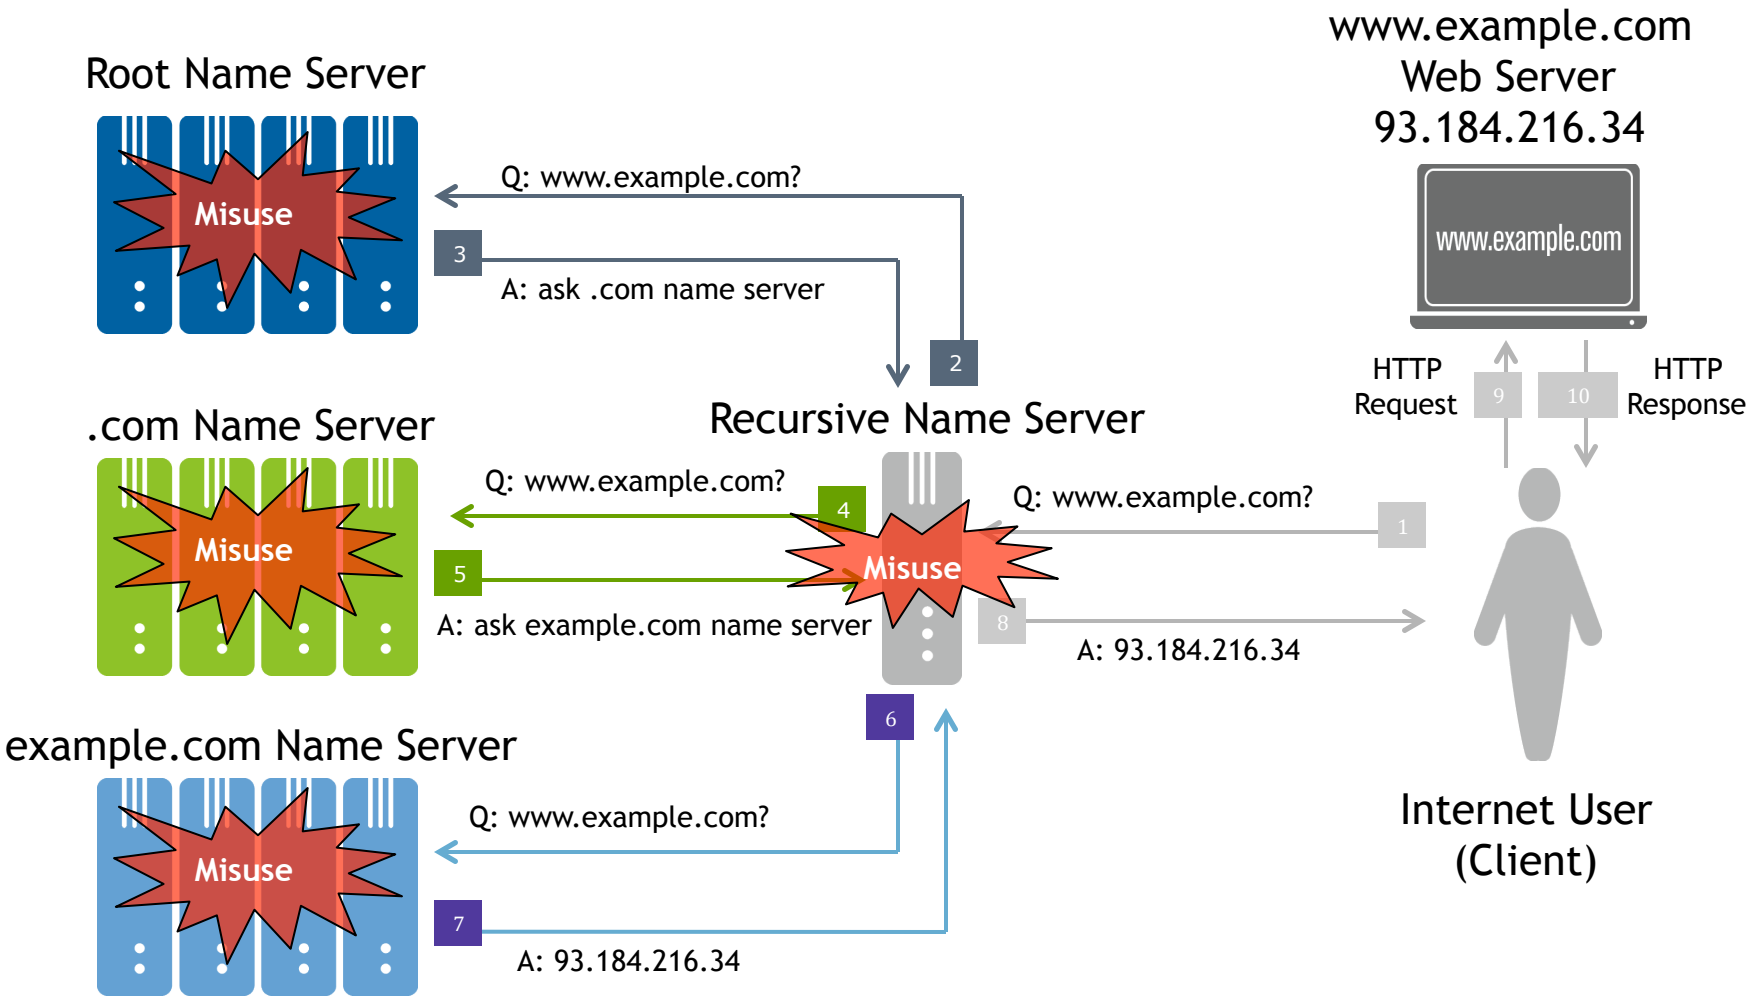

Mitigation 1: Data handling

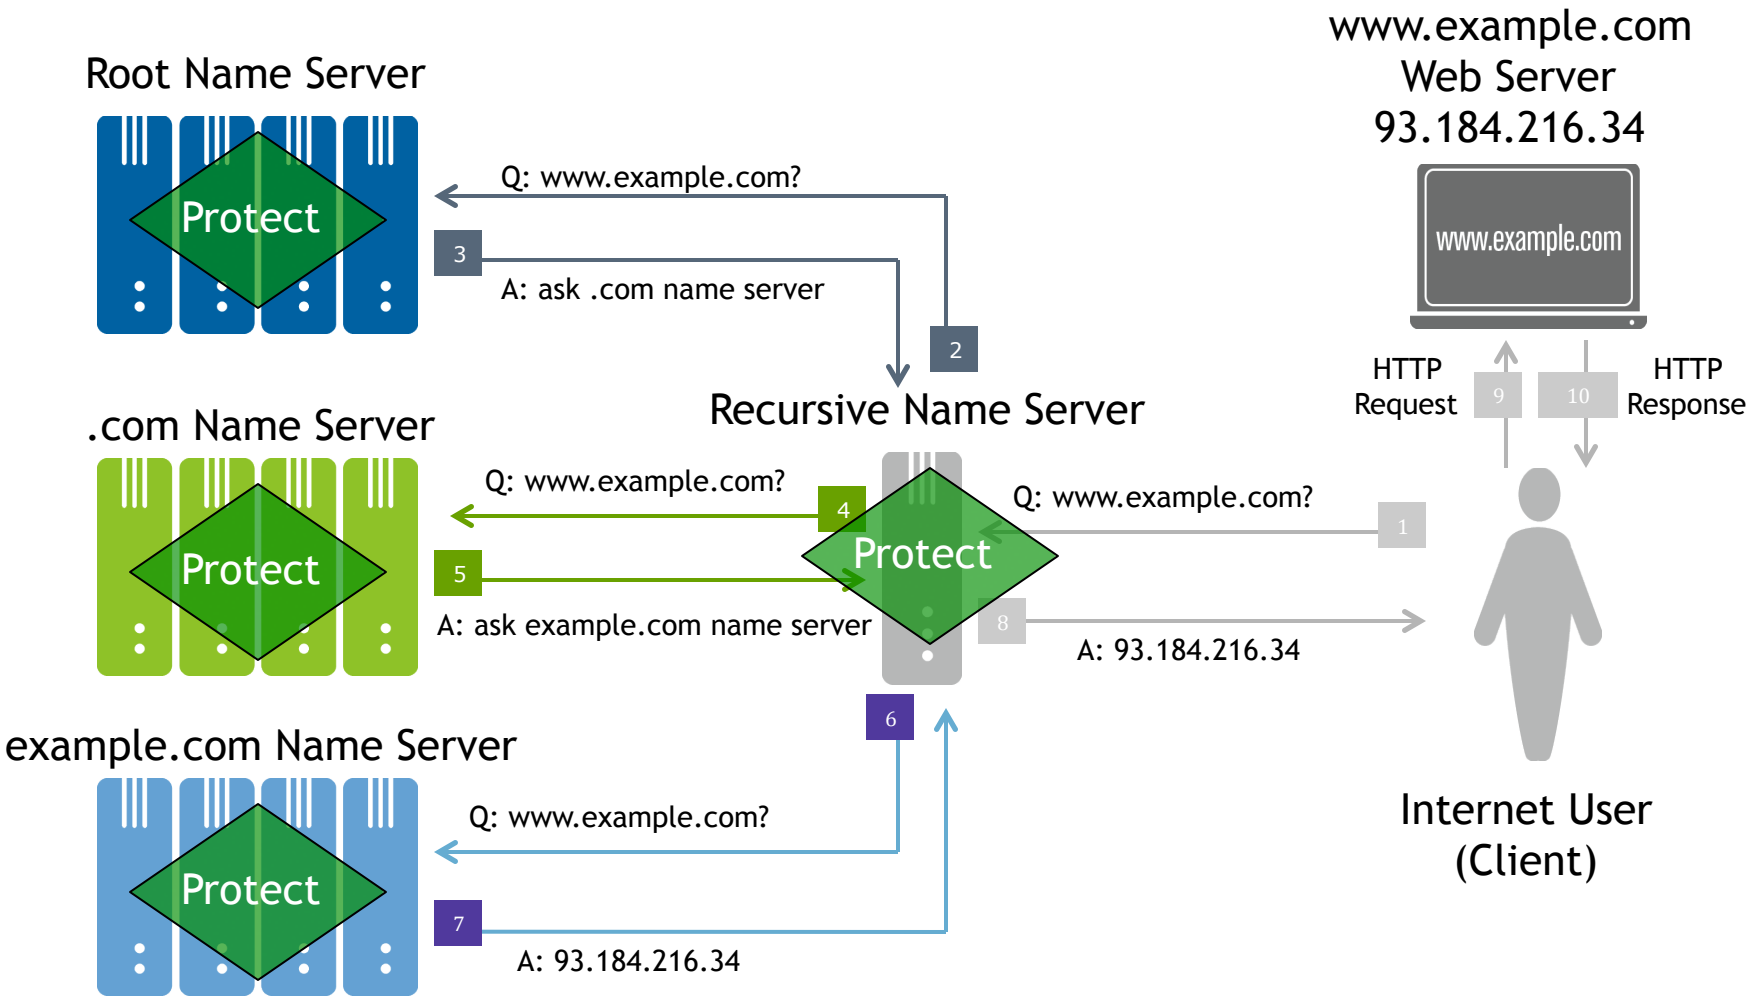

Mitigation 2: Encryption (DNS-over-TLS etc.)

- Like other Internet protocols, DNS can be made more secure and information disclosure can be reduced by running over Transport Layer Security (TLS)
- IETF DPRIVE working group currently developing DNS-over-TLS specification and others
- Mitigates eavesdropping (risks 1 & 3)
  - Also mitigates modification in transit

Mitigation 2: Encryption (DNS-over-TLS)

- DNS Over TLS: Initiation and Performance Considerations
  - https://tools.ietf.org/html/draft-ietf-dprive-dns-over-tls
  - New well known port (TBD) for DNS over TLS
  - TLS: follow best practices of RFC 7525
  - Two profiles defined: an opportunistic profile (no server authentication), and a pre-configured profile.
  - Details of performance considerations and recommendations:
    - Connection reuse, pipelining, out-of-order response processing, use of TCP Fast Open if available, use of TLS session resumption, and other optimizations.
  - Implementations already emerging (see next talk!)

Risks 1 & 3: Eavesdropping

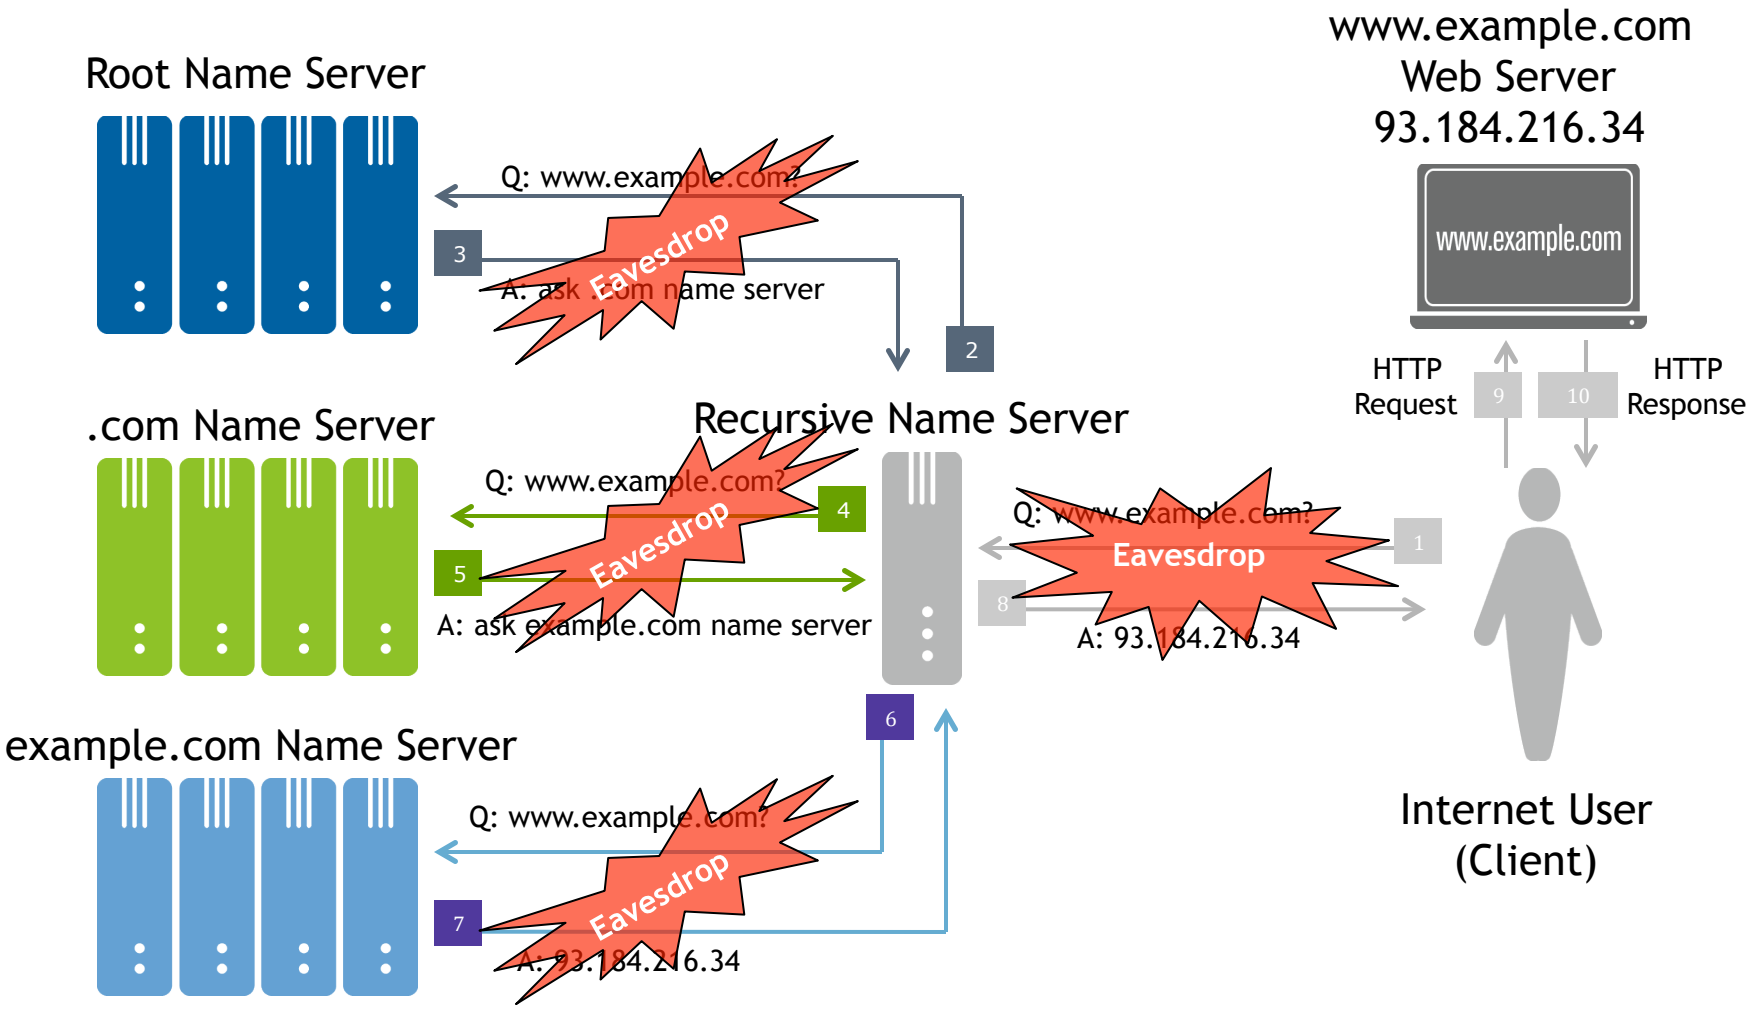

Mitigation 2: Encryption (DNS-over-TLS etc.)

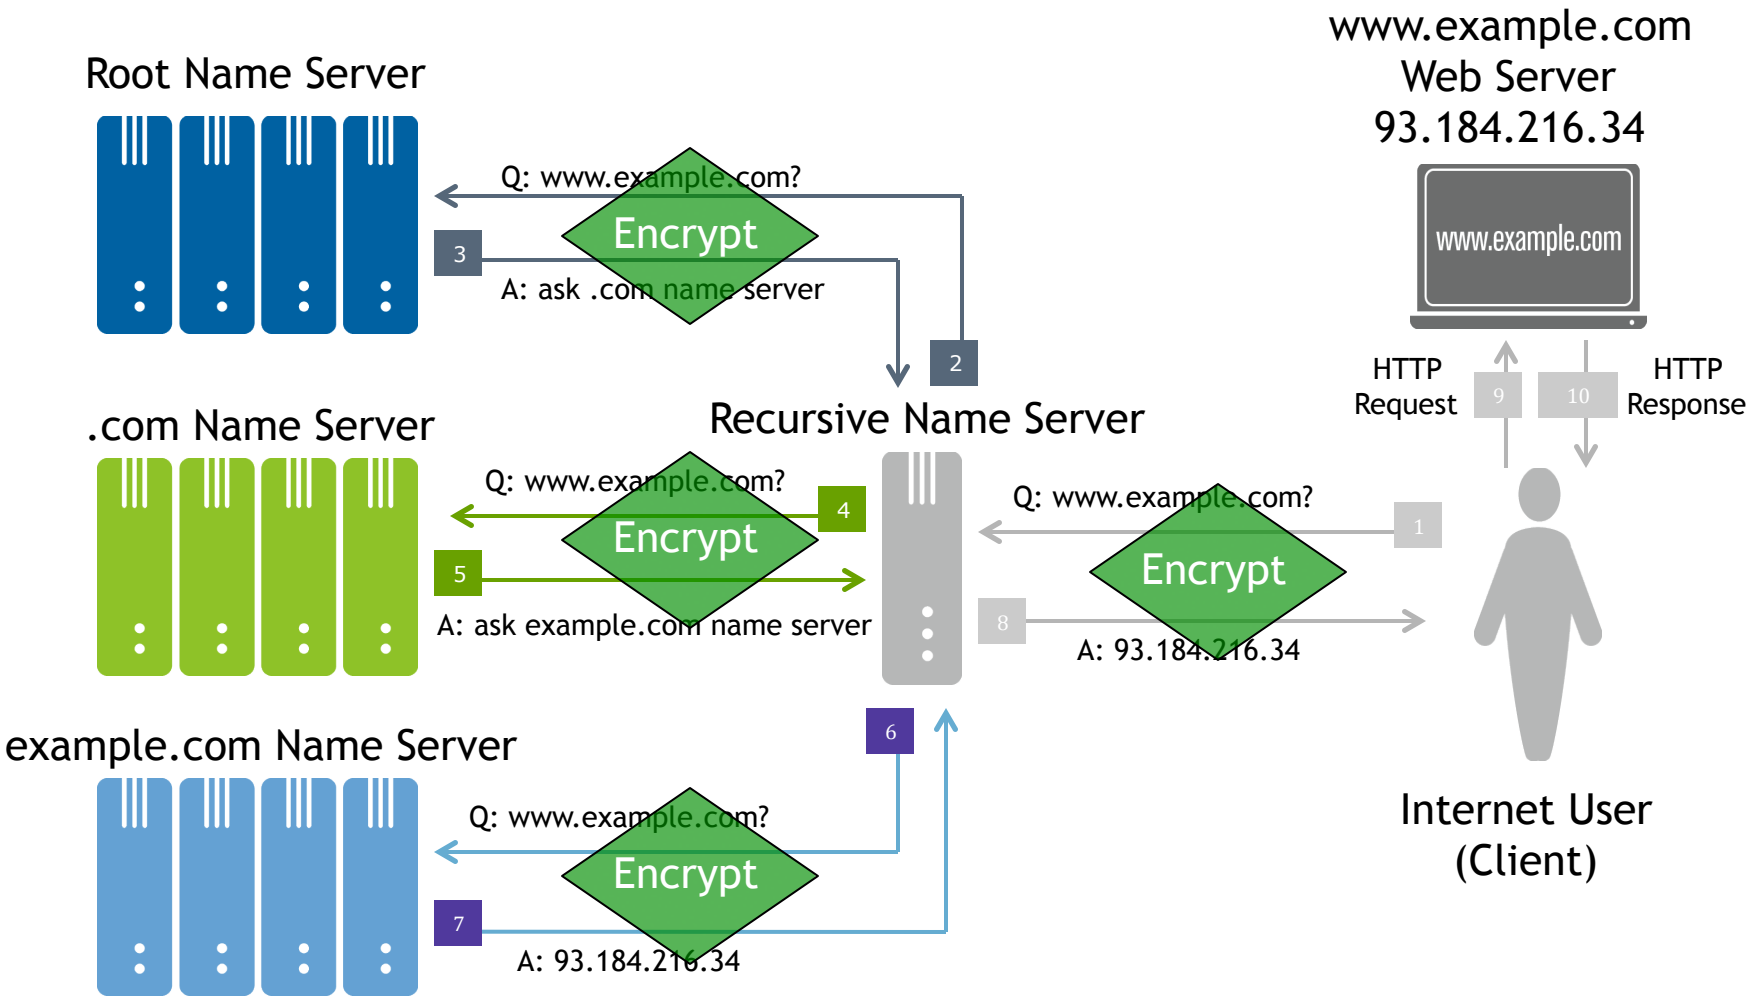

Mitigation 3: Qname Minimization

- DNS information disclosure can be reduced by asking authoritative only enough for referral to next server – not full query name ("qname") each time
- IETF DNSOP working group currently developing qname minimization spec
  - Completed DNSOP WGLC and soon will go to IETF Last Call
- Partially mitigates eavesdropping (risk 3) w/o encryption or changing authoritative

- For a more detailed treatment, see "Query-name Minimization and Authoritative Server Behavior – S. Huque", Spring 2015 DNS-OARC workshop: https://indico.dns-oarc.net/event/21/contribution/9

Risk 1 & 3: Eavesdropping

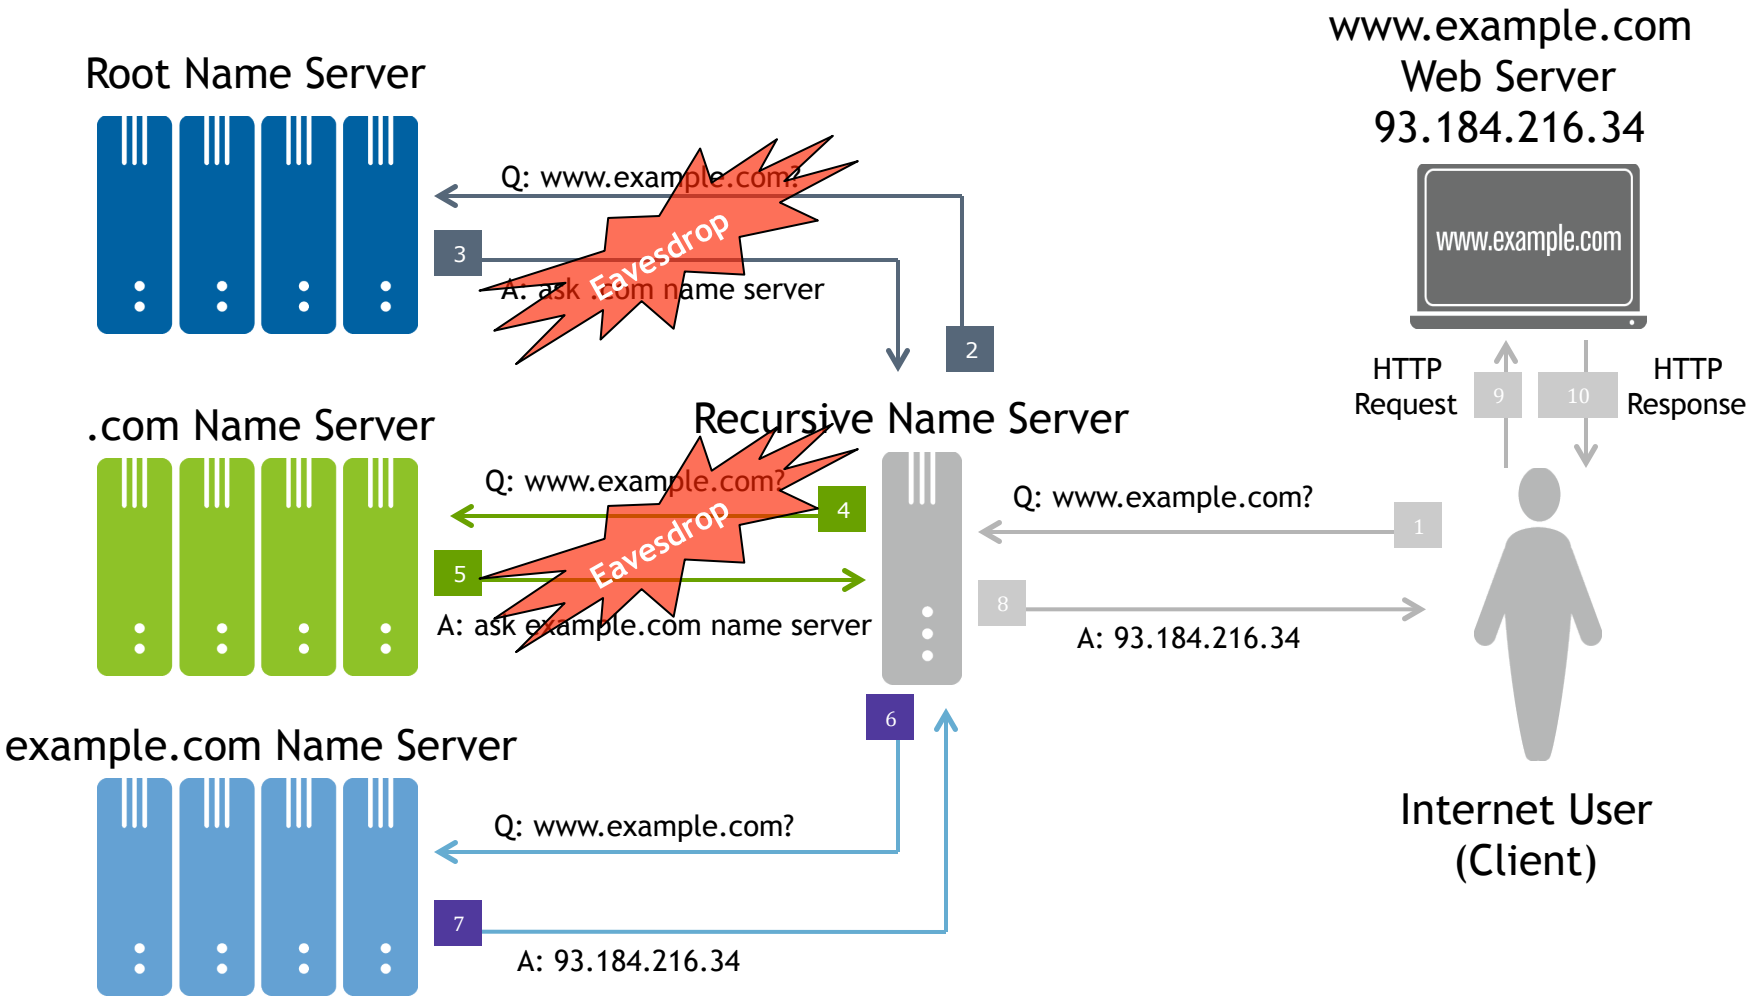

Mitigation 3: Qname Minimization

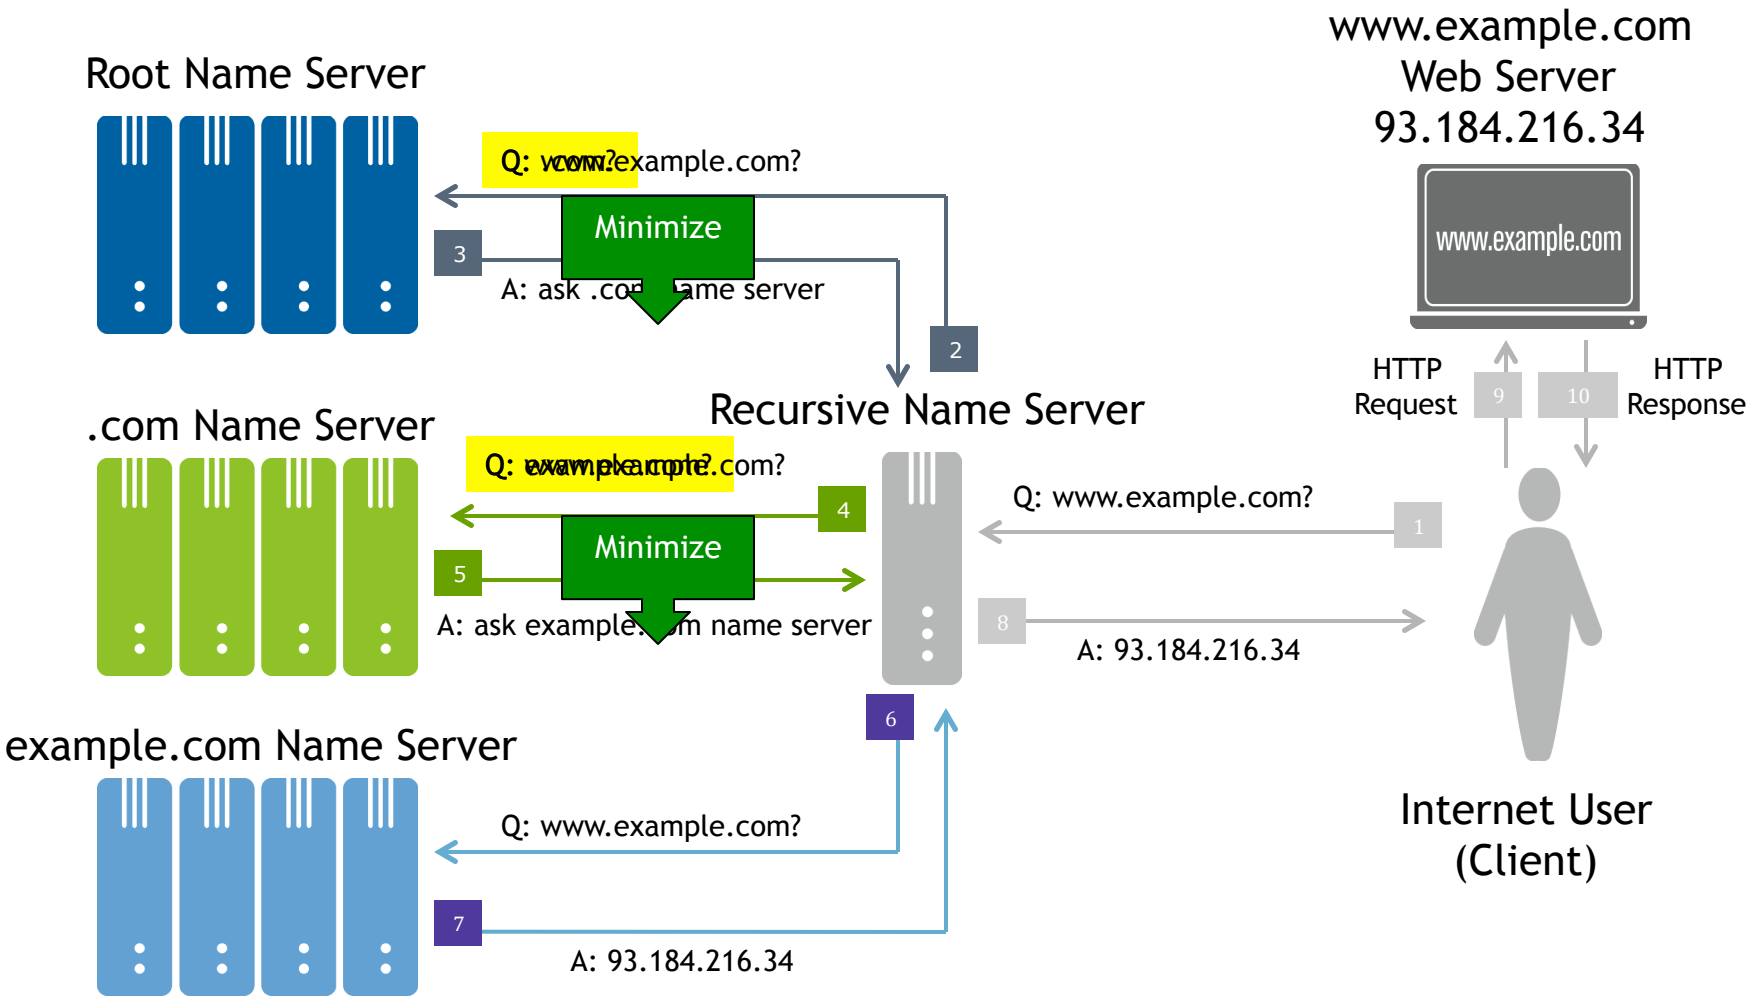

Summary: Risk Mitigation Matrix

Mitigations	DNS System Level Risks			
	Client to Recursive	At Recursive	Recursive to Authoritative	At Authoritative
Data Handling (Policies)		Mitigate Misuse		Mitigate Misuse
Encryption (DNS-over-TLS etc.)	Mitigate Monitoring		Mitigate Monitoring	
qname minimization			Mitigate Monitoring	Mitigate Monitoring

Some Additional Risks and Mitigations

Zone Enumeration

- Consider zones with policy limiting access to data as a whole
  - Access control for AXFR and IXFR, and channel encryption
- DNSSEC proof of non-existence, NSEC, re-opens this risk
  - Enclosing proof has plaintext names, and adversary can zone-walk through random queries (RFC 4033-4035)
- NSEC3 (RFC 5155) mitigates zone-walking through hashing, but now can be compromised by well-resourced adversary
- Research proposal, NSEC5 (i-d ref) mitigates this attack
  - Ongoing discussion in DNSOP WG – tradeoffs of risk versus cost (due to online signing)
  - Tradeoff may be in favor for DANE zones where enumeration would produce catalog of public keys

Side Channels

- Even when a data flow is encrypted, private information may be inferred by various means
- Side-channel attacks – well known ones include:
  - Size-based
  - Timing-based
- Cache snooping is an example of a timing-based attack
  - In some cases, in-cache responses (faster than not in-cache ones) can reveal what names are queried by the target individual
  - Adversary needs to identify recursive used by target and gain access
- Another form of cache snooping: targeted RD=0 queries:
  - DNS Cache Snooping, Feb 2004 (L Grangeia)
  - http://cs.unc.edu/~fabian/course_papers/cache_snooping.pdf

Size-based Side Channels

- Size-based attacks have been practiced on TLS, Skype and other encrypted traffic
- DNS once encrypted still has some predictable query/ response patterns
- Another advantage for practicality of this attack is that adversary may have access to known plaintext (by making its own queries)
  - Shulman IRTF ANRP award paper at IETF-93 stimulated discussions in DPRIVE and TLS WGs
- Known mitigation is to pad requests & responses so that they have uniform length

Size-based Side Channels (cont.)

- DPRIVE: draft for an EDNS(0) padding option:
  - https://tools.ietf.org/html/draft-mayrhofer-edns0-padding
- TLS: multiple choices
  - Existing padding options, but they have been impacted in TLS by some attacks (Poodle, …)
  - Create new application padding option that TLS stacks could use for DNS (in our case)
  - Wait a bit longer for TLS 1.3, which has been addressing a requirement for cryptographically analyzed padding and is a green field

Leakage of DNS Names by Other Protocols

- Impact of developing privacy enhancements for DNS
- Before, with no DNS privacy, pressure was low to avoid DNS name disclosures in plaintext in other protocols
- That may be changing:
  - TLS use of cleartext domain names in handshake, now recognized as a risk
  - DHCP – an Anonymity Profile document that is currently in WGLC provides options that allow an end-system not to expose its FQDN (this was a PERPASS outcome)
    - https://tools.ietf.org/html/draft-ietf-dhc-anonymity-profile

DNS name leakage in TLS

- Server Name Indication extension (SNI) exposes the domain name of the intended server
  - An issue where many named services are hosted on common platforms like large CDNs.
  - Tricks to obfuscate the server name have already emerged. See "domain fronting" www.icir.org/vern/papers/meek-PETS-2015.pdf
- DNS names also exposed in TLS Certificate messages.
- TLS1.3 protocol designers are discussing ways to encrypt and prevent these exposures.
- (Note: SNI encryption is at best a partial solution to hiding a service name. A more complete solution involves mechanisms well beyond just the DNS, such moving servers into anonymity networks. See Facebook's Tor hidden servicefor example at "facebookcorewwwi.onion".)

Summing Up

- Background and Risk Overview
  - RFCs 7258, 7624, 7626
  - Privacy evaluation
- DNS Privacy Risks – System View
  - Between client and recursive
  - At recursive
  - Between recursive and authoritative
  - At authoritative
- Mitigations – System View
  - Data Handling (Policies)
  - Query Confidentiality
  - Qname Minimization
- Additional Risks and Mitigations
  - Enumeration
  - User Identifier LHS
  - Side-channels
  - Size-based side channel
  - Research (no slide)
    - Transitivity networks
    - DNS Ecosystem variants
    - Unlucky Few
- Domain Name Leakage in Other Protocols
  - TLS server name extension
  - DHCP FQDN option
  - More work to be done!

Questions/Comments?