Domain Name System

Part One

CS249i - Modern Internet
Stanford
Fall 2021

Gautam Akiwate

UC San Diego

Goal

Understand
the rationale, and operational workings of the
Domain Name System (DNS).

Problem

- Communication on the Internet via IP
- Hard to remember IP addresses
- Easier to remember names.
    - Slightly harder to type them in correctly!
- Map Names to IP addresses
    - theo → 35.186.238.101

Centralized Solution: Historical Solution

Centralized Solution: Historical Solution

- hosts.txt file that has mappings for all hosts
  - organization : host → IP address
  - /etc/hosts
- Stanford Research Institute (SRI) kept main copy
  - Single place to update records
  - Download hosts.txt file periodically

Centralized Solution: Historical Solution

- hosts.txt file that has mappings for all hosts
  - organization : host → IP address
  - /etc/hosts
- Stanford Research Institute (SRI) kept main copy
  - Single place to update records
  - Download hosts.txt file periodically
- Problems
  - Fragile
  - Hard to scale
  - Hard to keep in sync

Decentralized Solution: Intuition

hosts.txt

organization : host → IP address

Decentralized Solution: Intuition

hosts.txt

organization : host → IP address

hosts.txt

organization→ IP address of organization.txt

organization.txt

host → IP address

Decentralized Solution: Intuition

hosts.txt

organization⟶ IP address of organization.txt

organization.txt

host ⟶ IP address

Decentralized Solution: Intuition

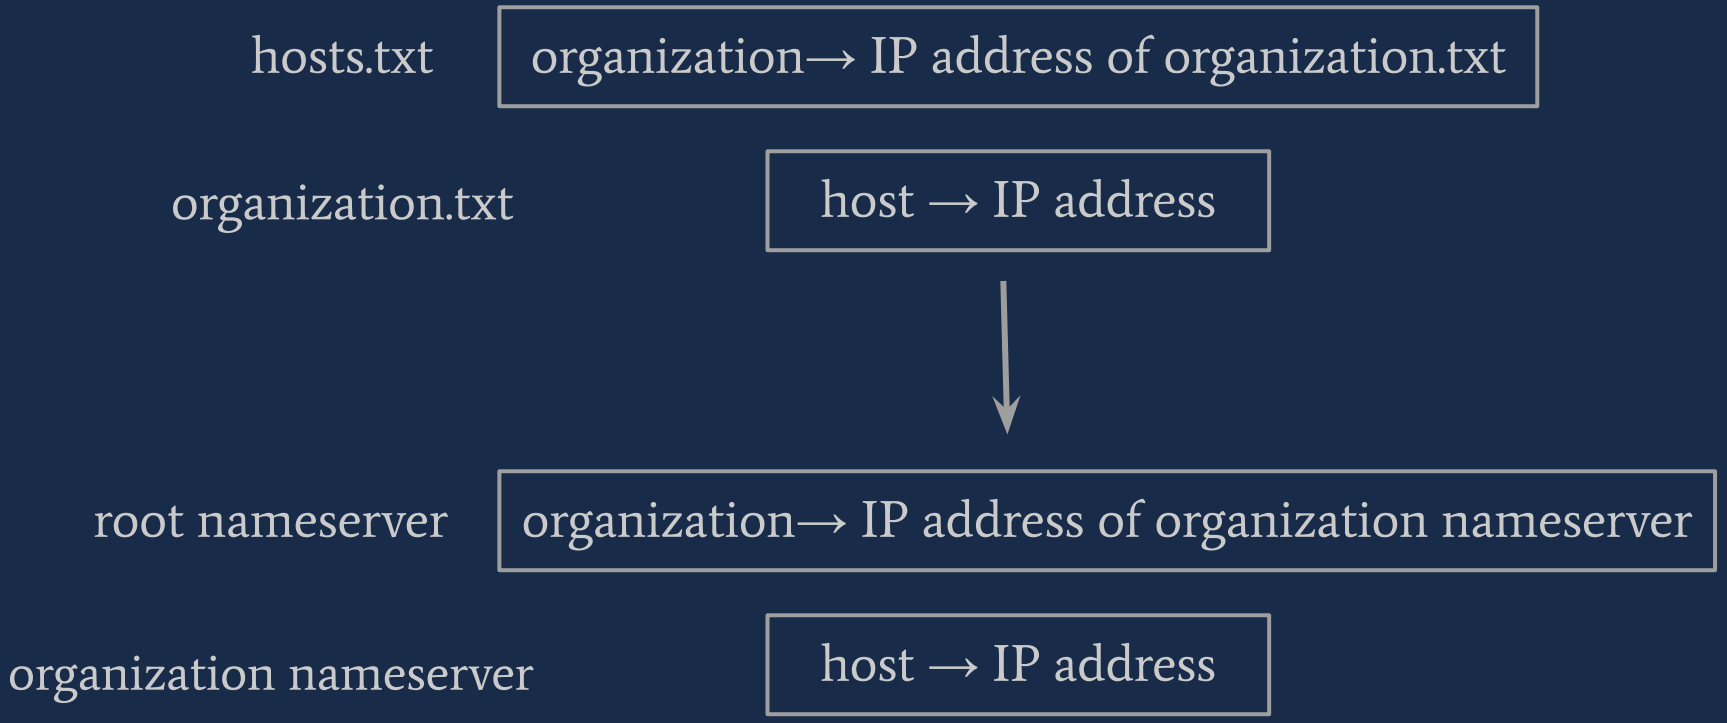

Hierarchical Namespace

DNS Root

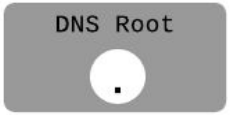

root
authoritative
nameservers

Hierarchical Namespace

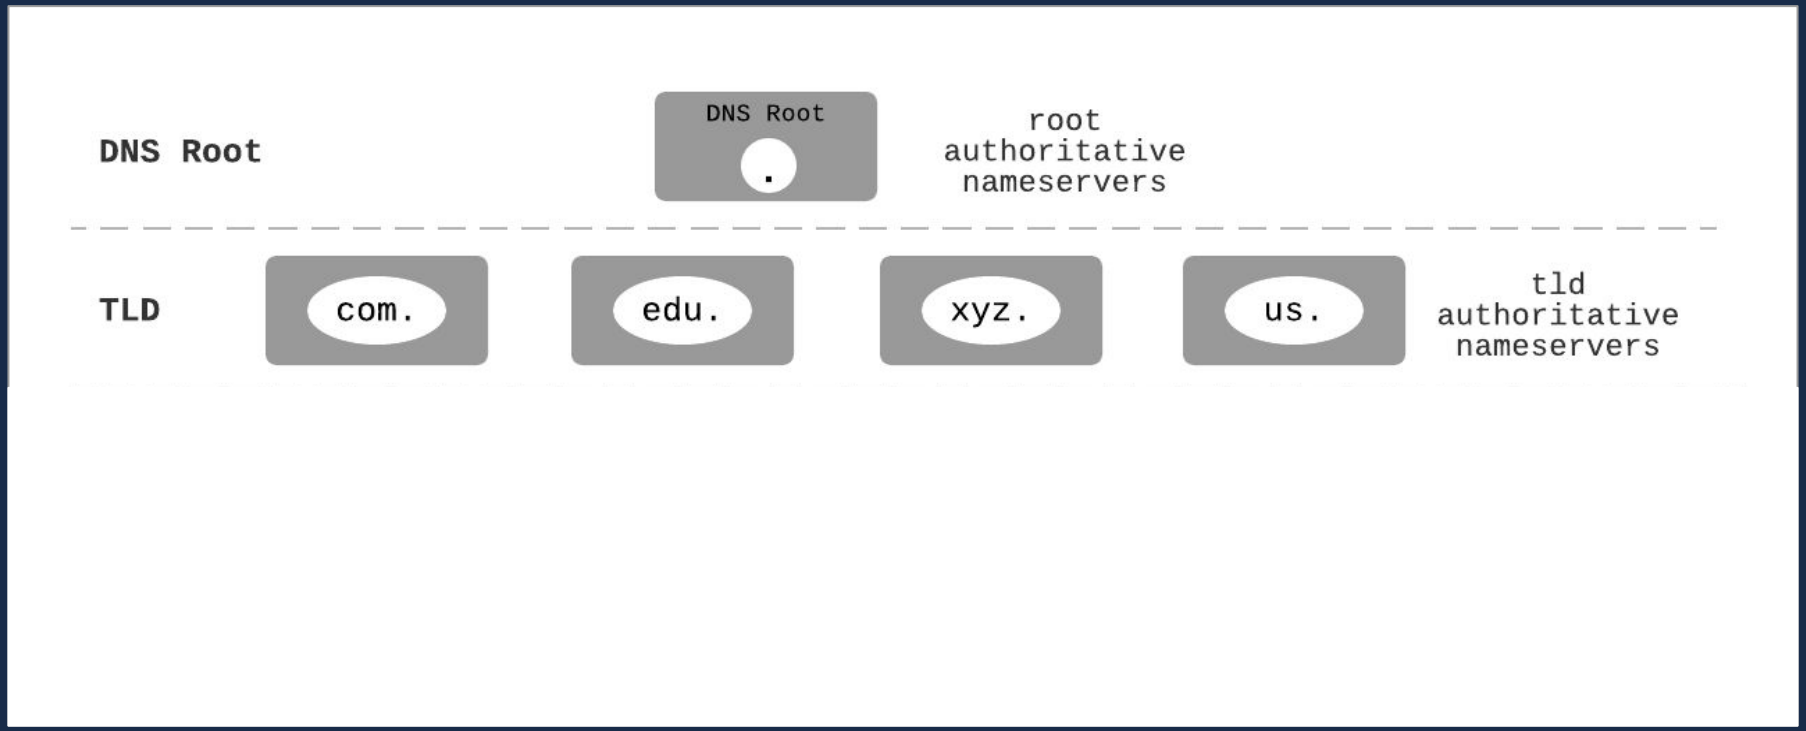

Hierarchical Namespace

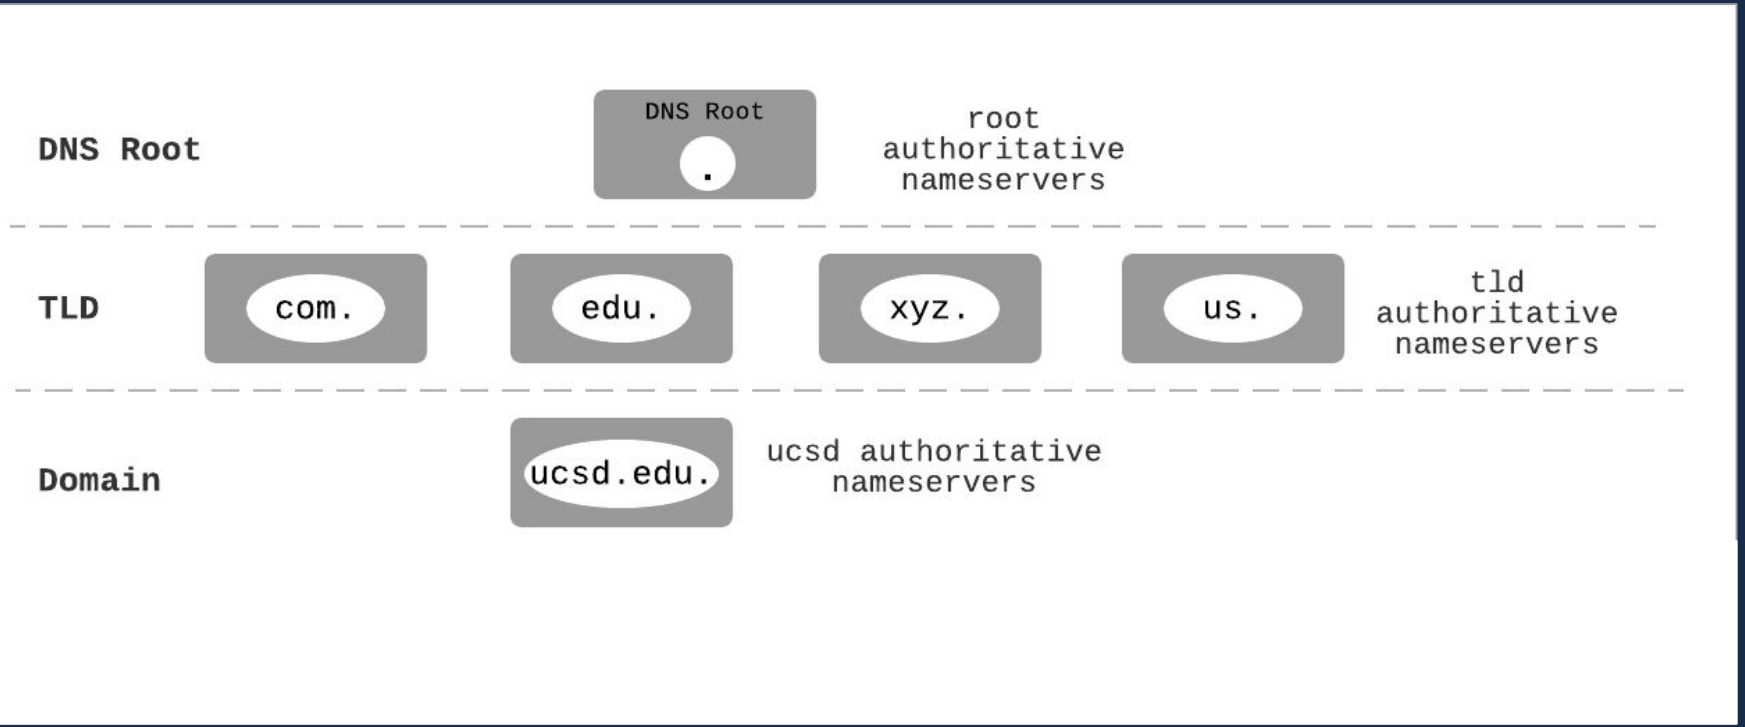

Hierarchical Namespace

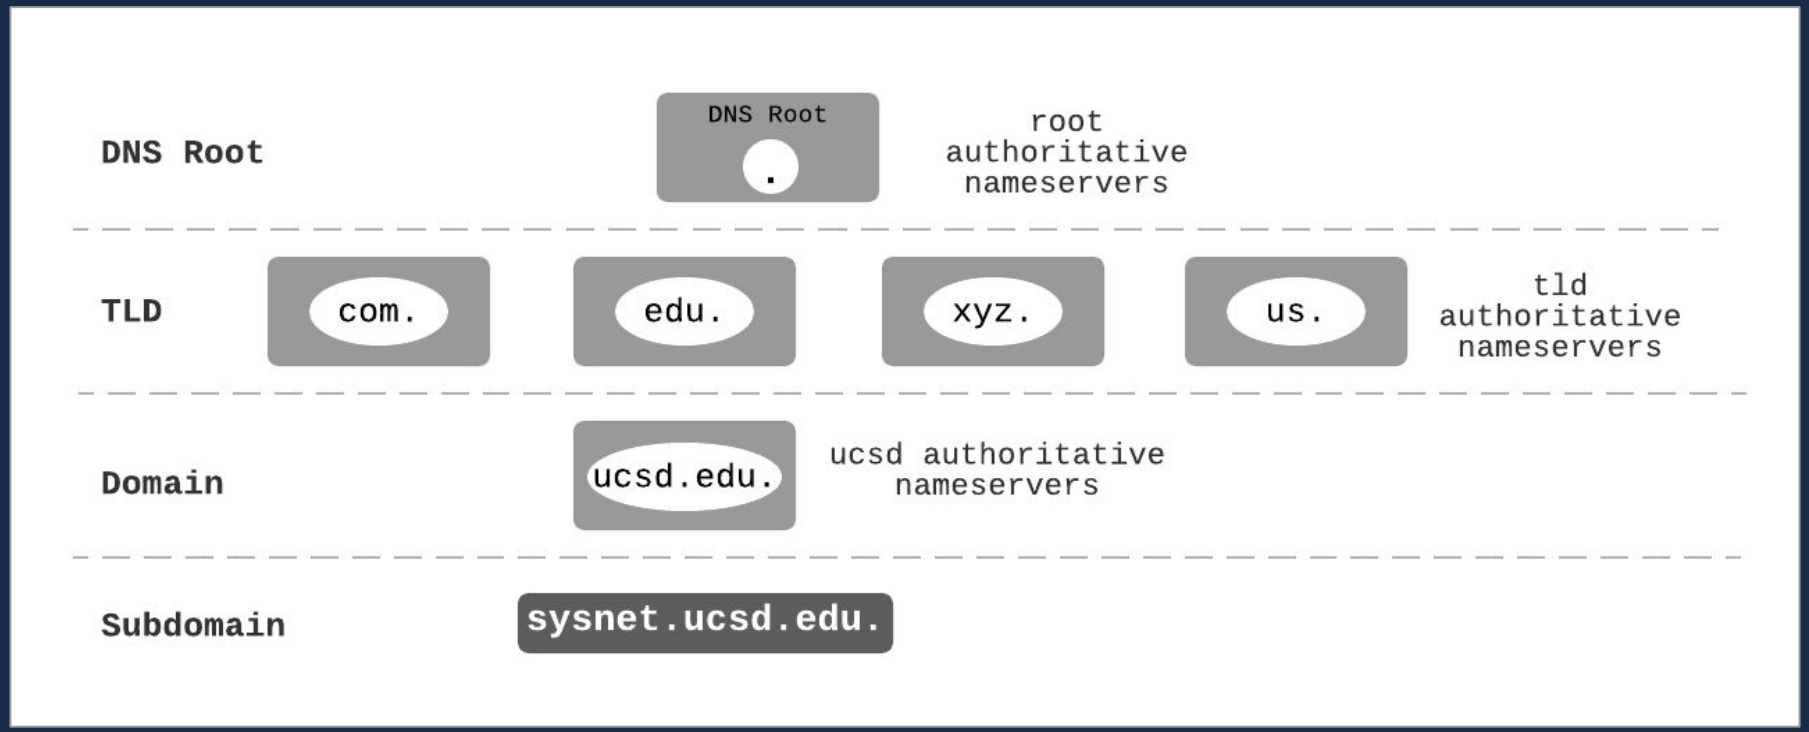

Life of a DNS Query

Life of a DNS Query

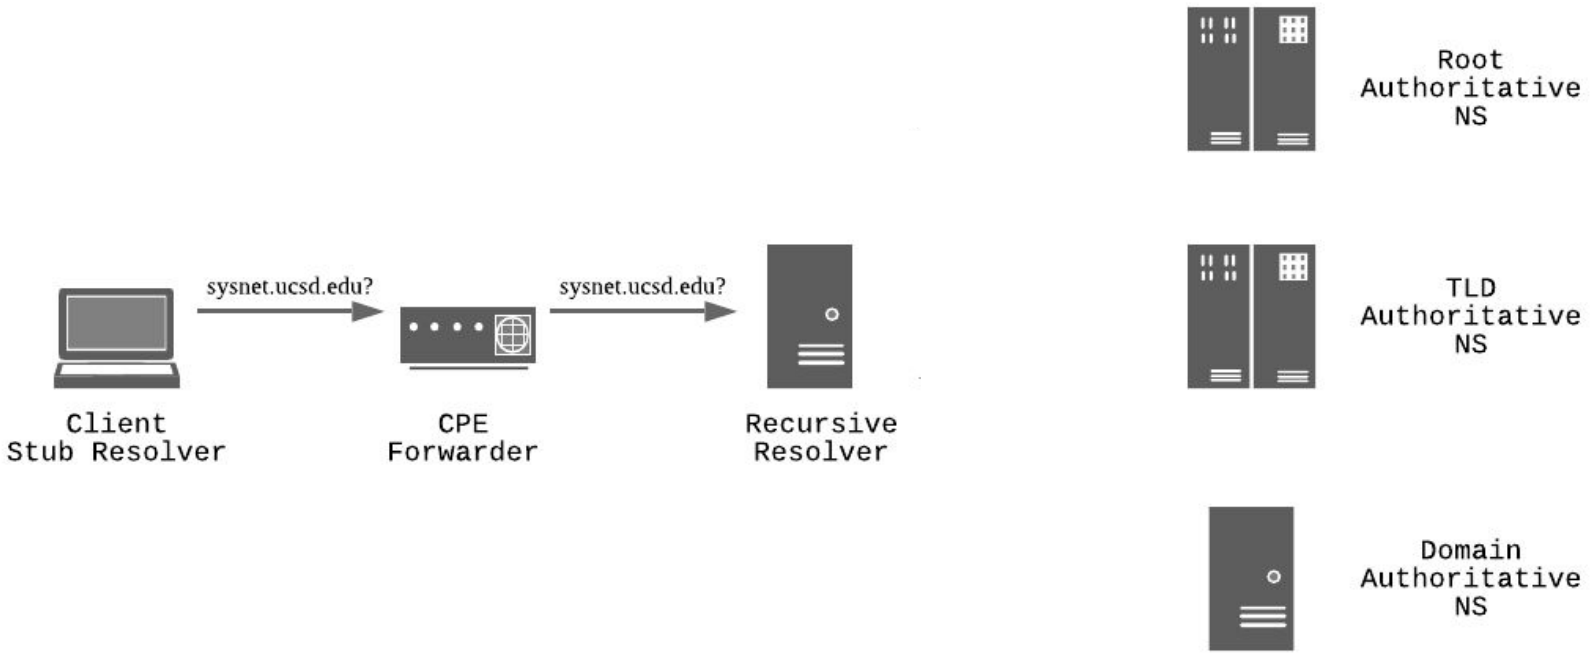

Life of a DNS Query

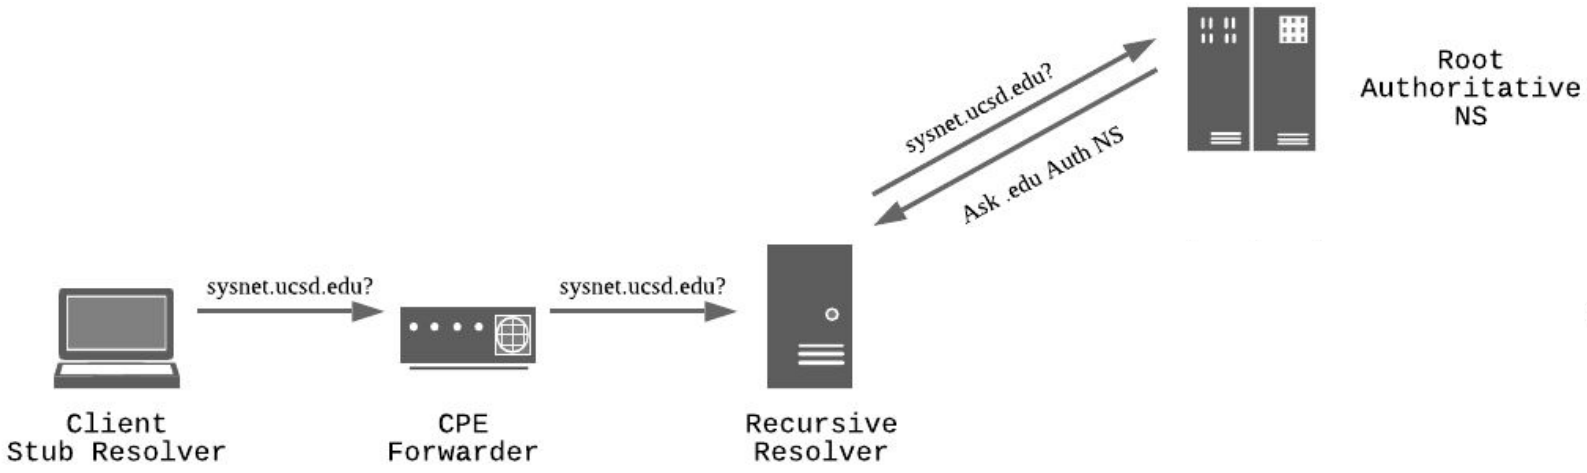

Life of a DNS Query

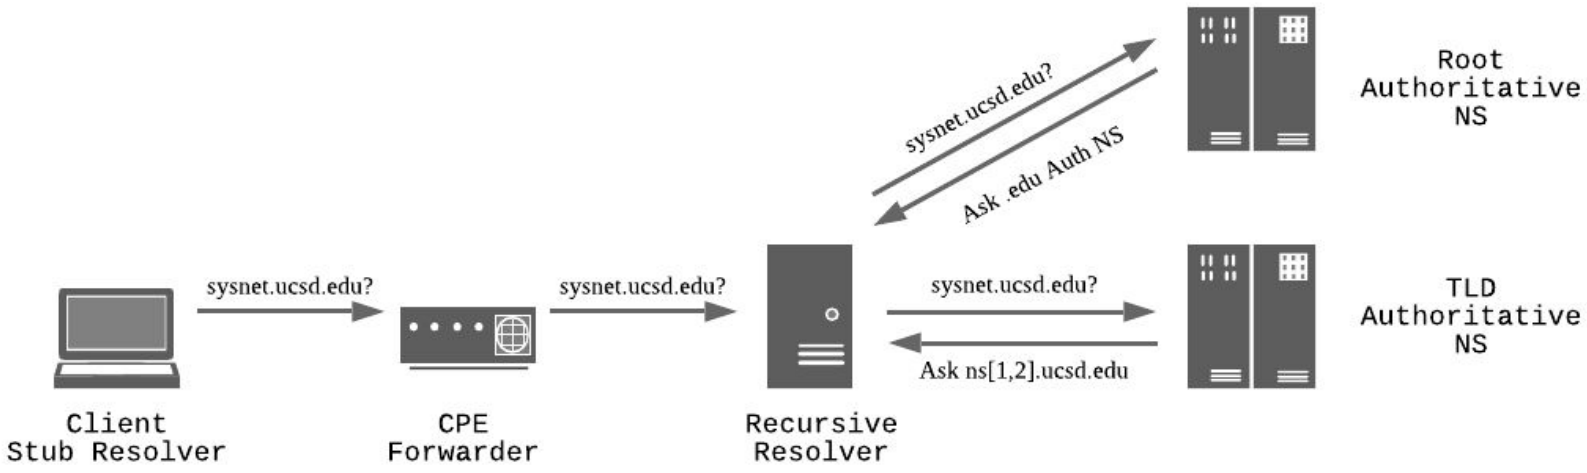

Life of a DNS Query

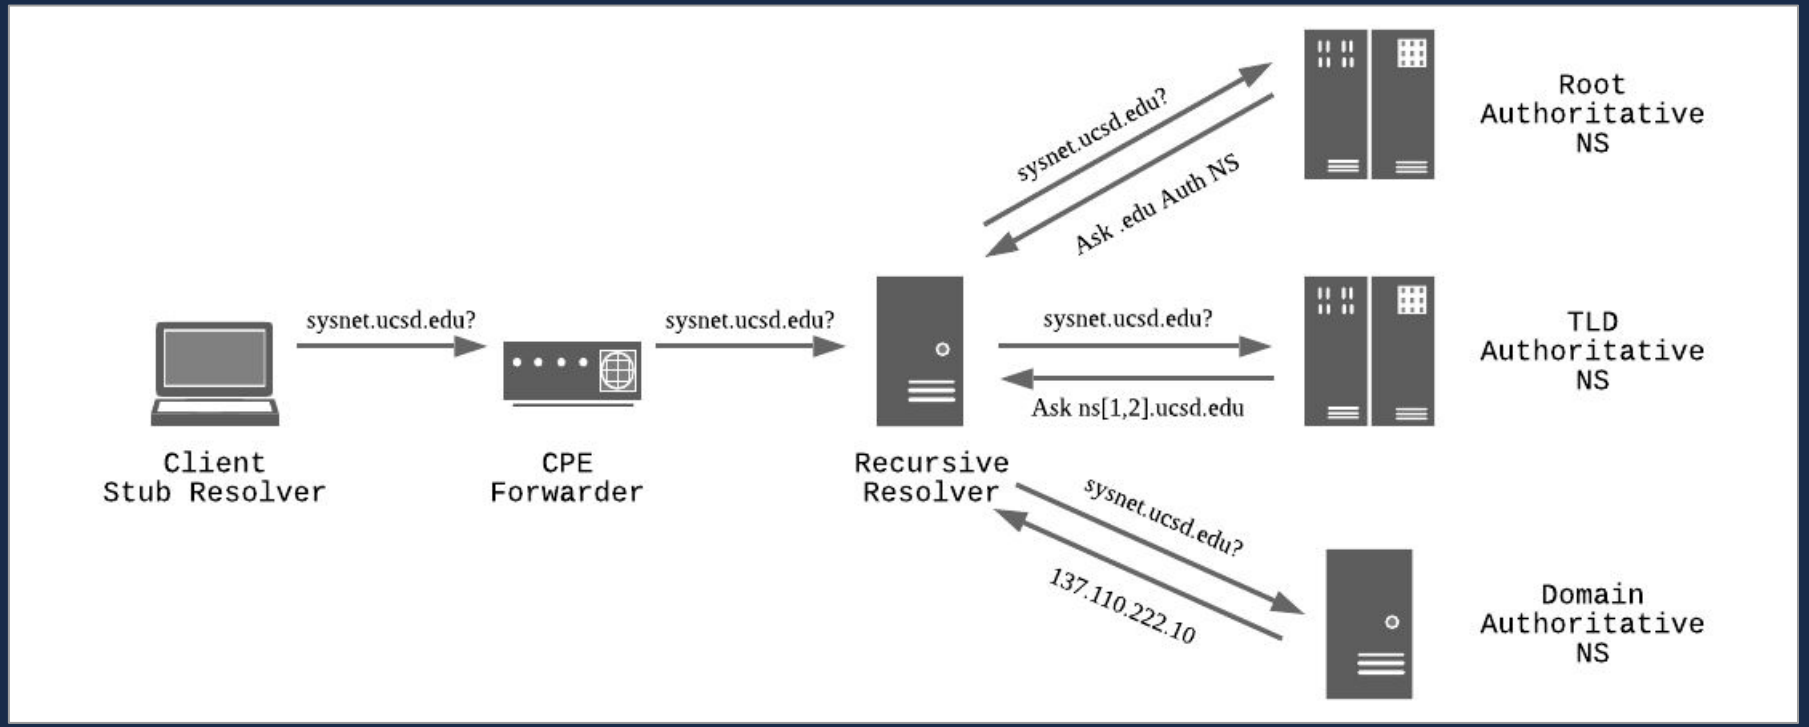

Life of a DNS Query

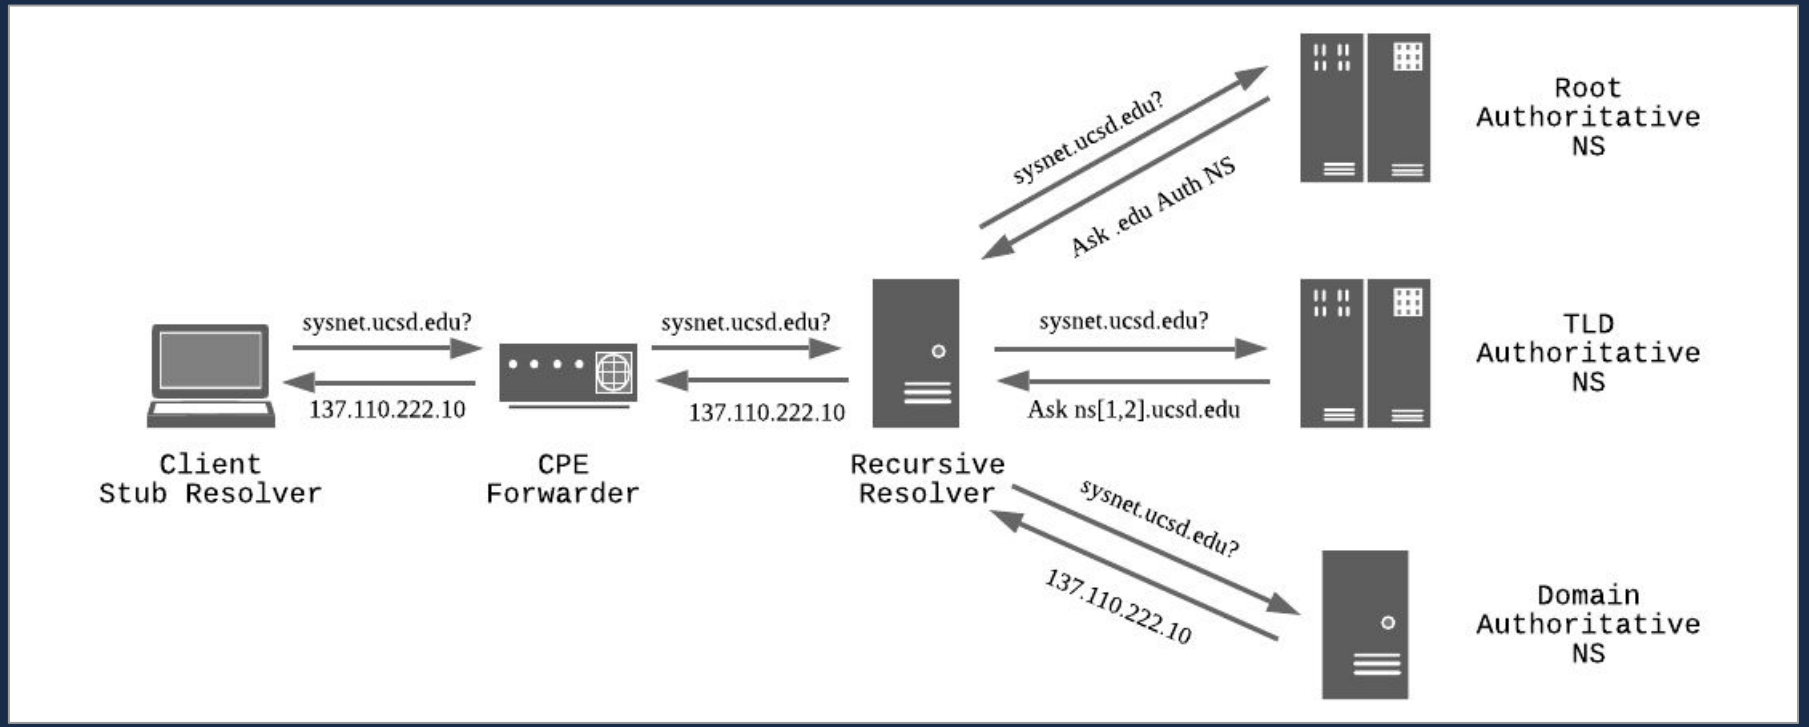

DNS Query Anatomy

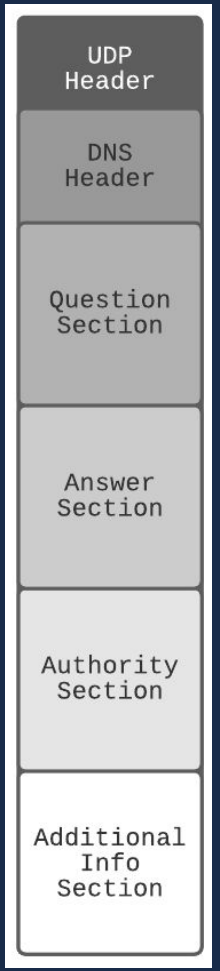

DNS Query Anatomy

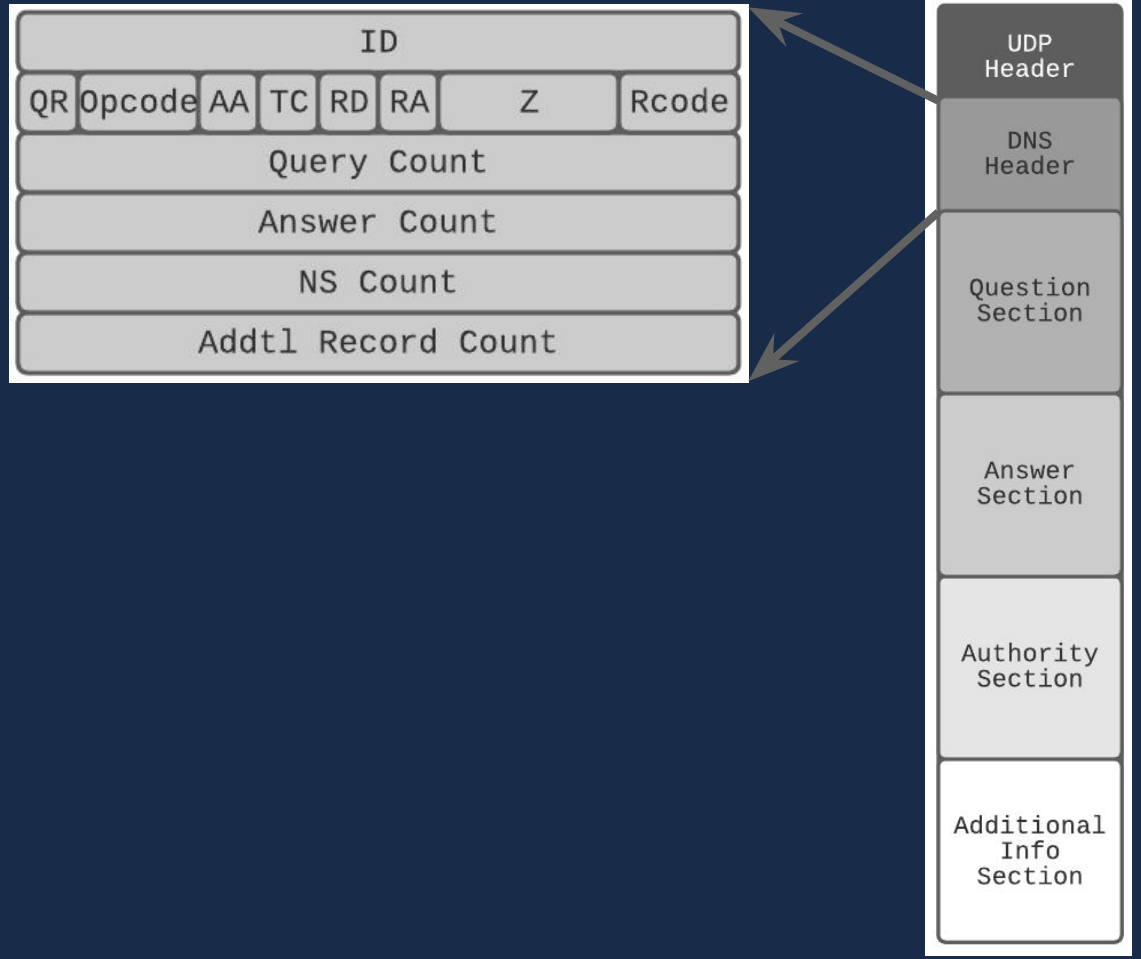

DNS Query Anatomy

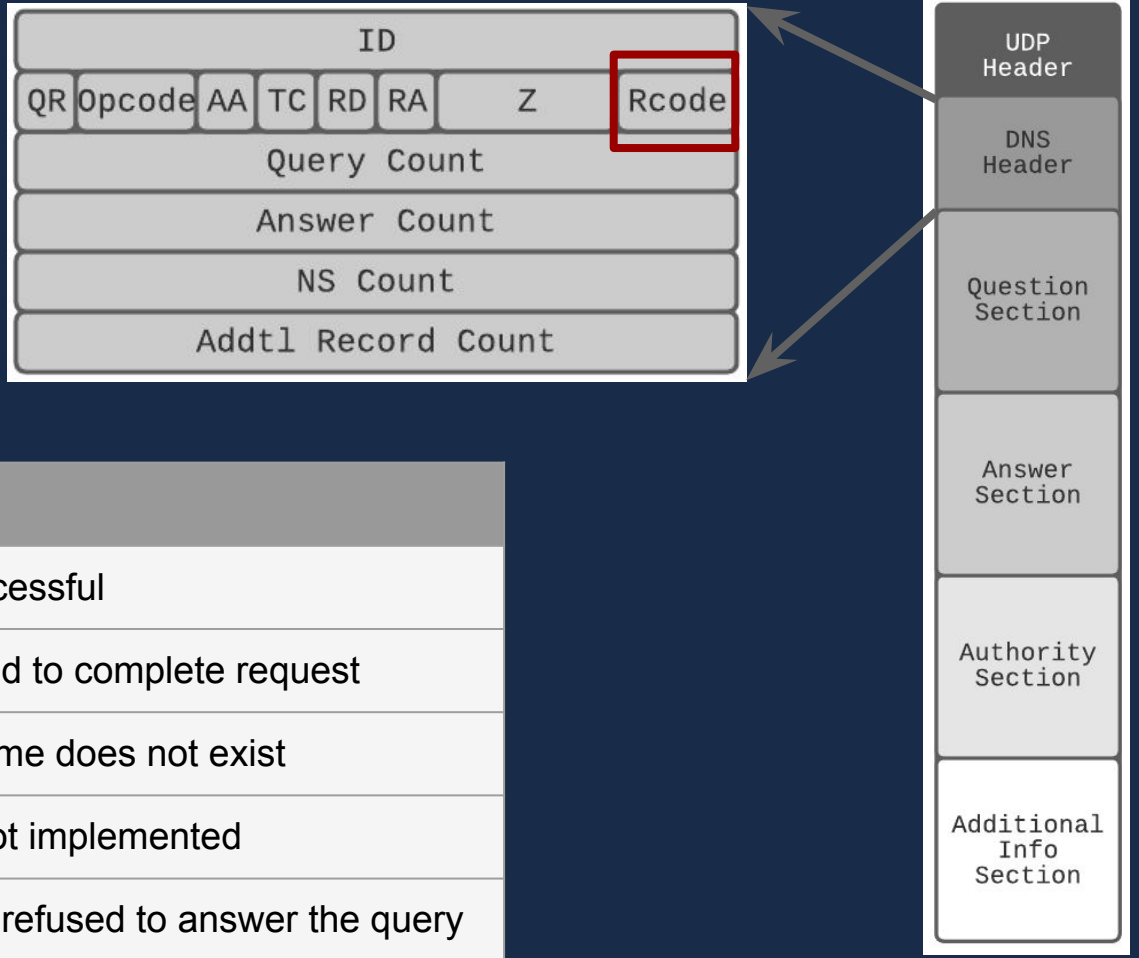

Rcode	Message	Function
0	NOERROR	Query Successful
2	SERVFAIL	Server failed to complete request
3	NXDOMAIN	Domain name does not exist
4	NOTIMP	Function not implemented
5	REFUSED	The server refused to answer the query

DNS Query Anatomy

Query Name (QNAME) -- Domain to resolve!

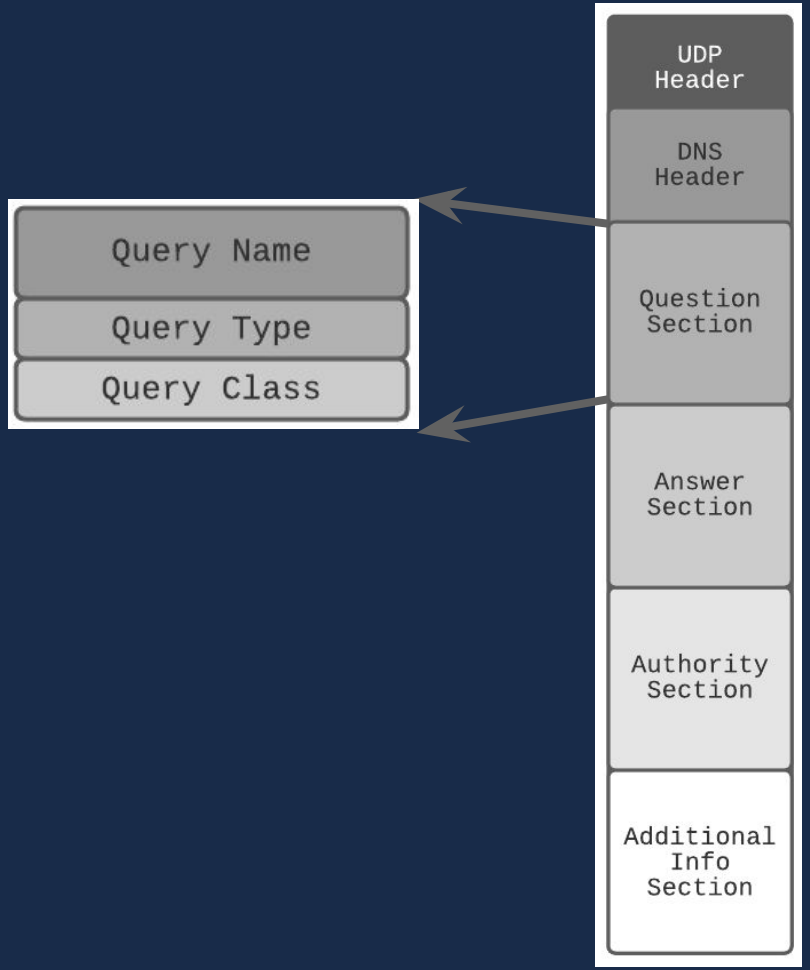

DNS Query Anatomy

Query Name (QNAME) -- Domain to resolve!

Query Class (QCLASS)

- CHAOS -- Used for debugging
- IN -- Internet

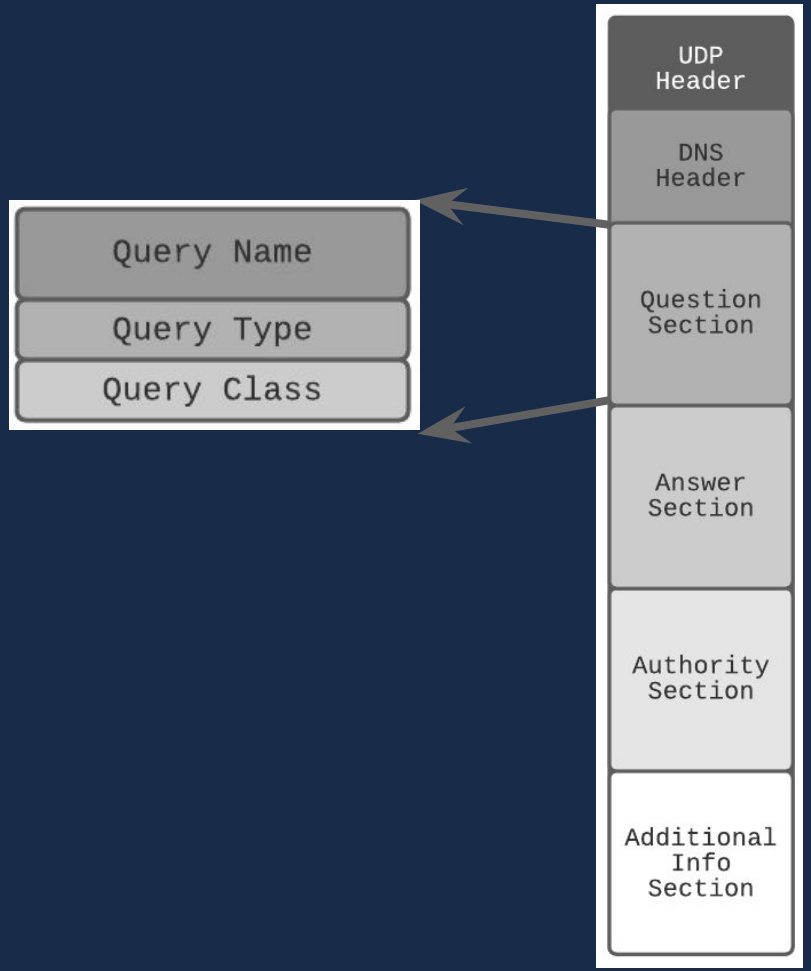

DNS Query Anatomy

Query Name (QNAME) -- Domain to resolve!

Query Class (QCLASS)

- CHAOS -- Used for debugging
- IN -- Internet

Query Type (QTYPE)

- NS -- Authoritative nameserver for domain
- A -- IPv4 Address
- AAAA -- IPv6 Address
- MX -- Mail Exchange Records

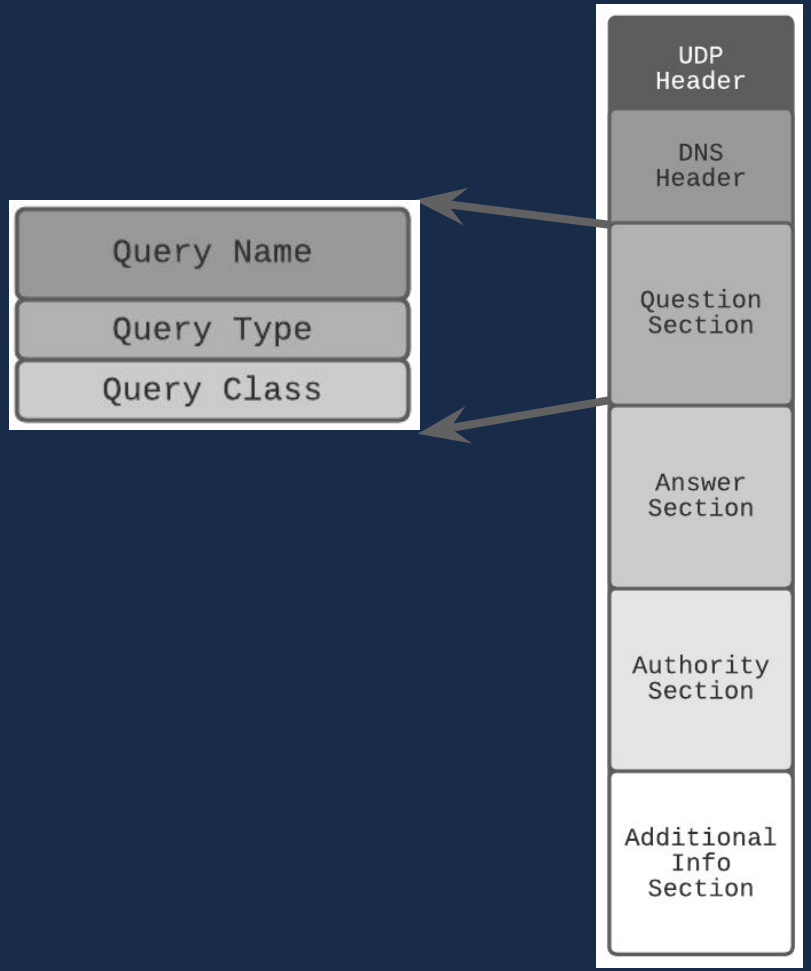

DNS Query Anatomy

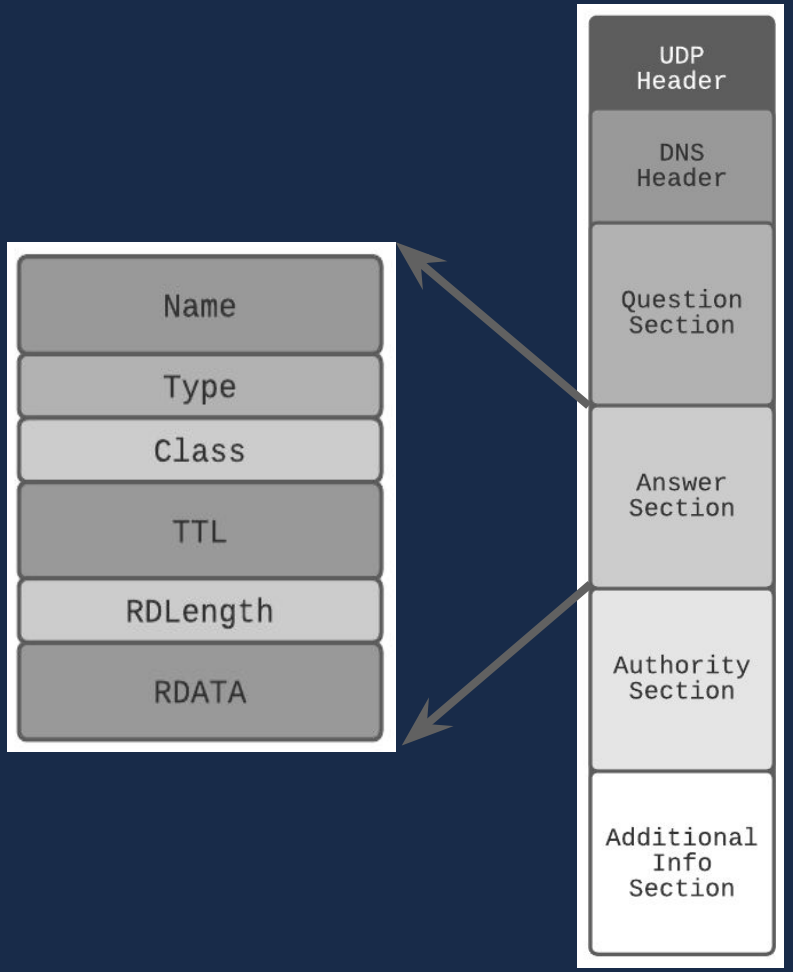

DNS Query Anatomy

TTL -- How long to cache answer

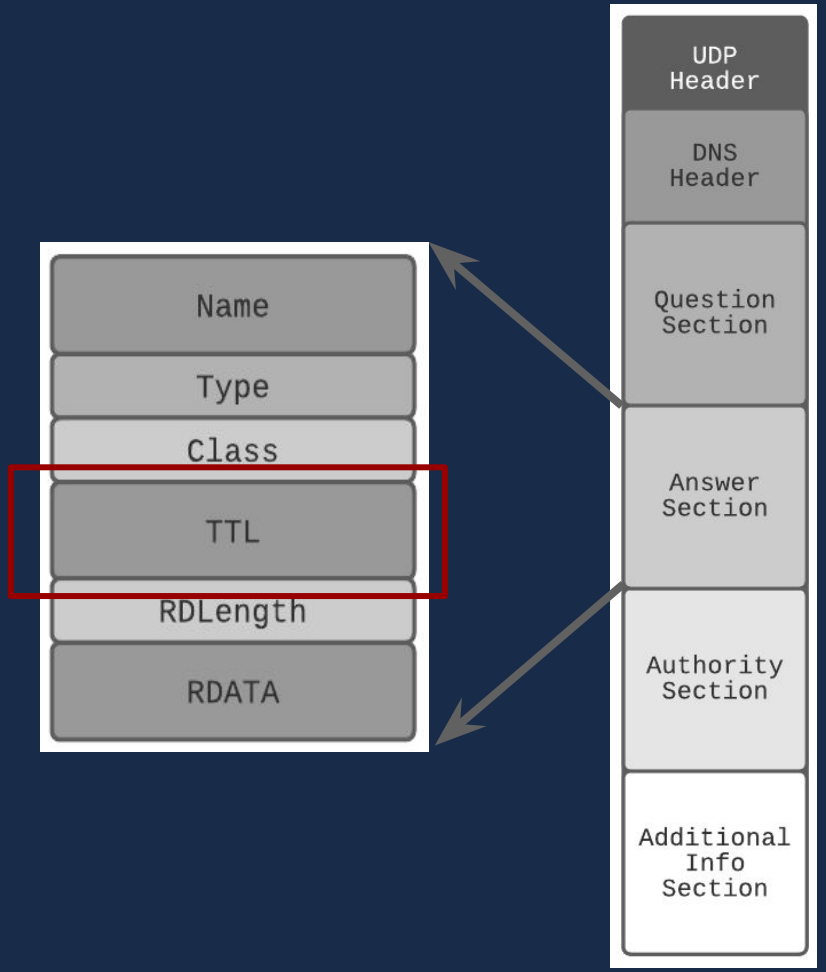

Caching

- Cache DNS Responses
  - Reduces load
  - Improves latency
  - Reuse of previous queries
- Negative Caching
- How long to cache?
  - Time To Live (TTL)

> “The caching discipline of the DNS works well, and given the unexpectedly bad performance of the Internet, was essential to the success of the system.”

Live Demo

Discussion

Failures

Reliability

Integrity

Confidentiality

Discussion: Failures

- Why can a DNS query fail?

Discussion: Failures

- Misconfiguration?
  - Typos
  - Misconfigured authoritative nameserver
- Hardware/Network Failures
  - Unreachable Nameserver
- Large Traffic Volume
  - DoS Attacks

Discussion: Reliability

- Why use UDP? Why not TCP?
- Reliability through replication
  - Two authoritative nameserver per domain
  - What about root servers? TLD authoritative NS?
- Reliability across the entire life cycle?

Discussion: Reliability

Root Servers

- Only 13 root servers?

https://blog.verisign.com/security/verisign-perspective-root-server-attacks/

Root Servers

- Only 13 root servers?

https://blog.verisign.com/security/verisign-perspective-root-server-attacks/

Anycast Primer

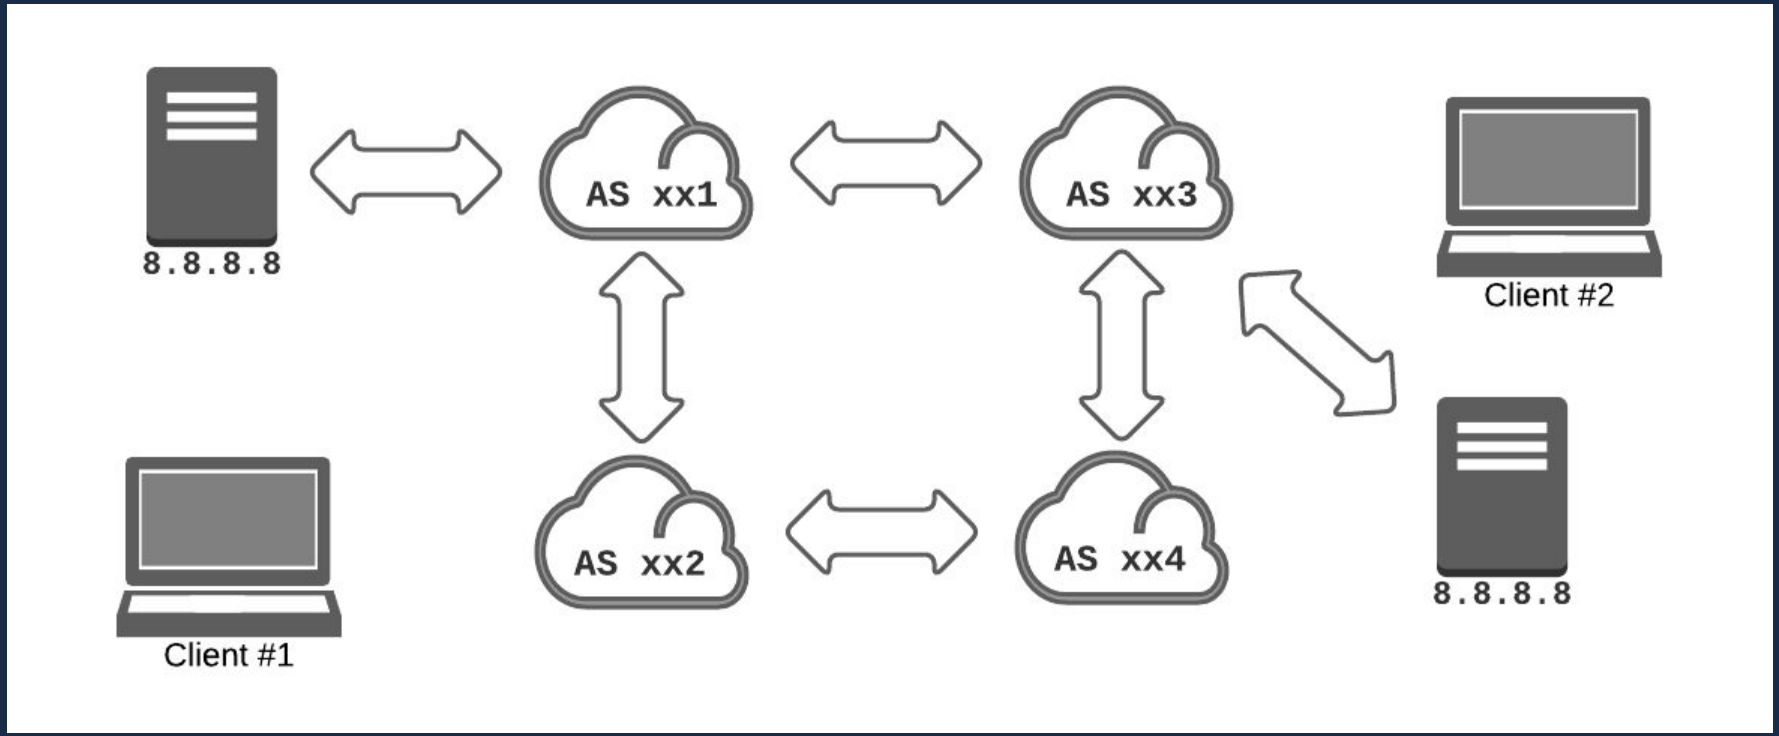

Anycast Primer

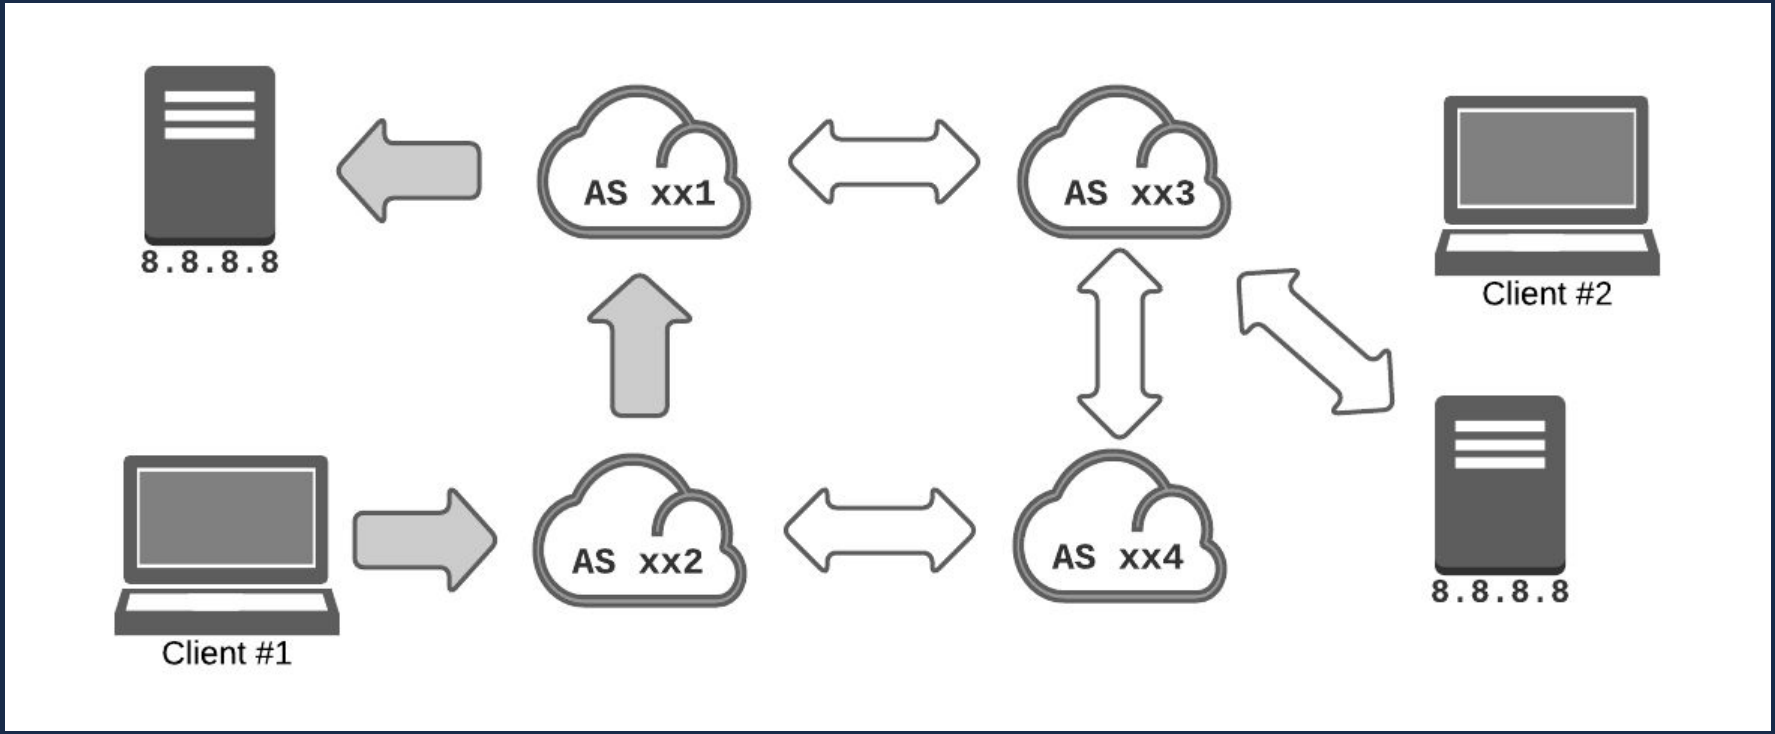

Anycast Primer

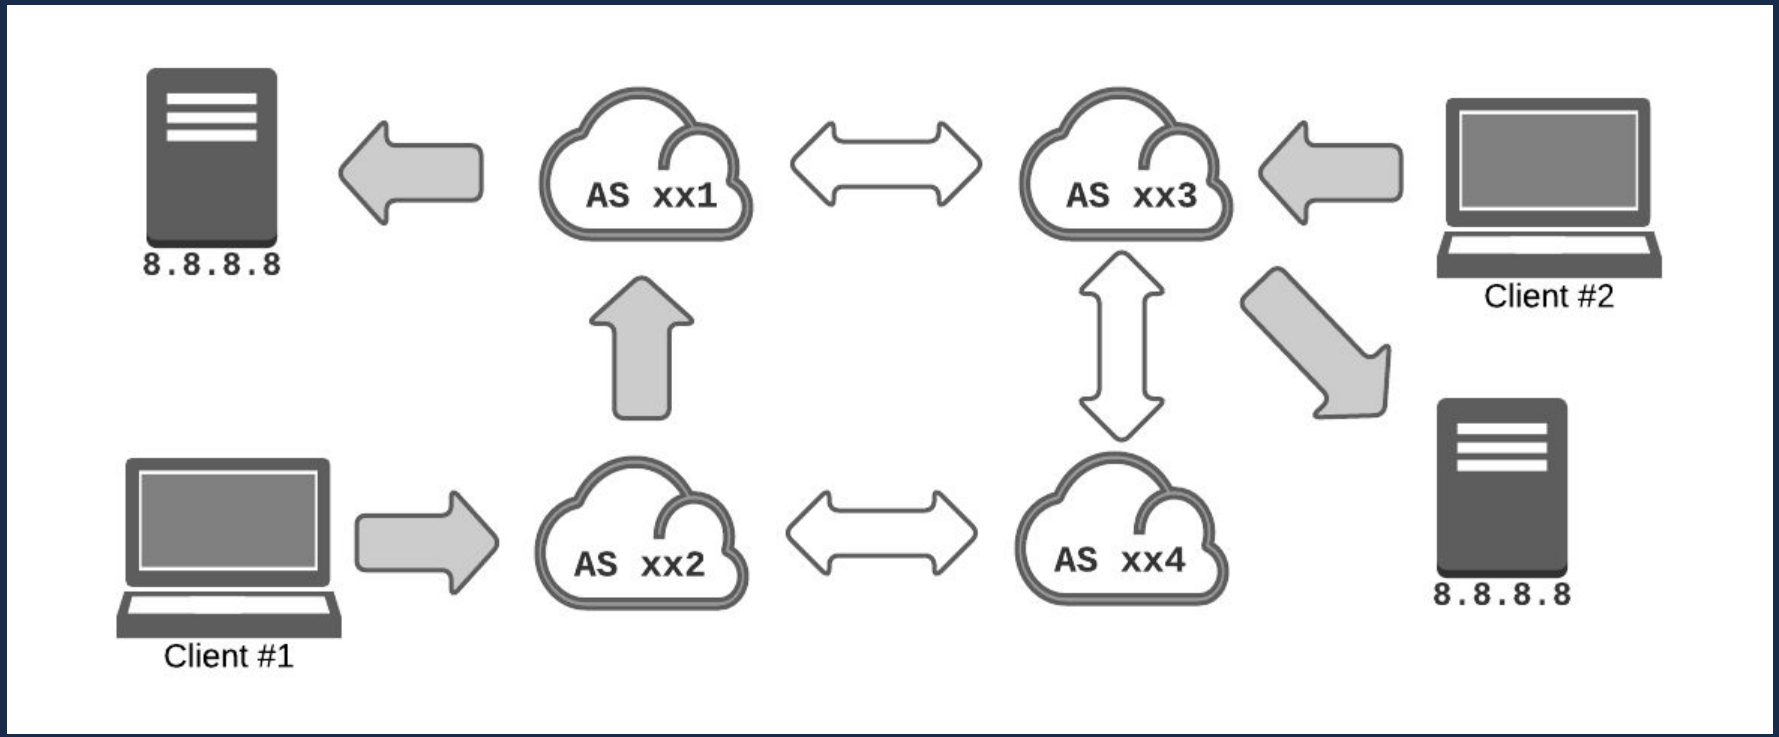

Discussion: Reliability

Discussion: Reliability

Discussion: Reliability

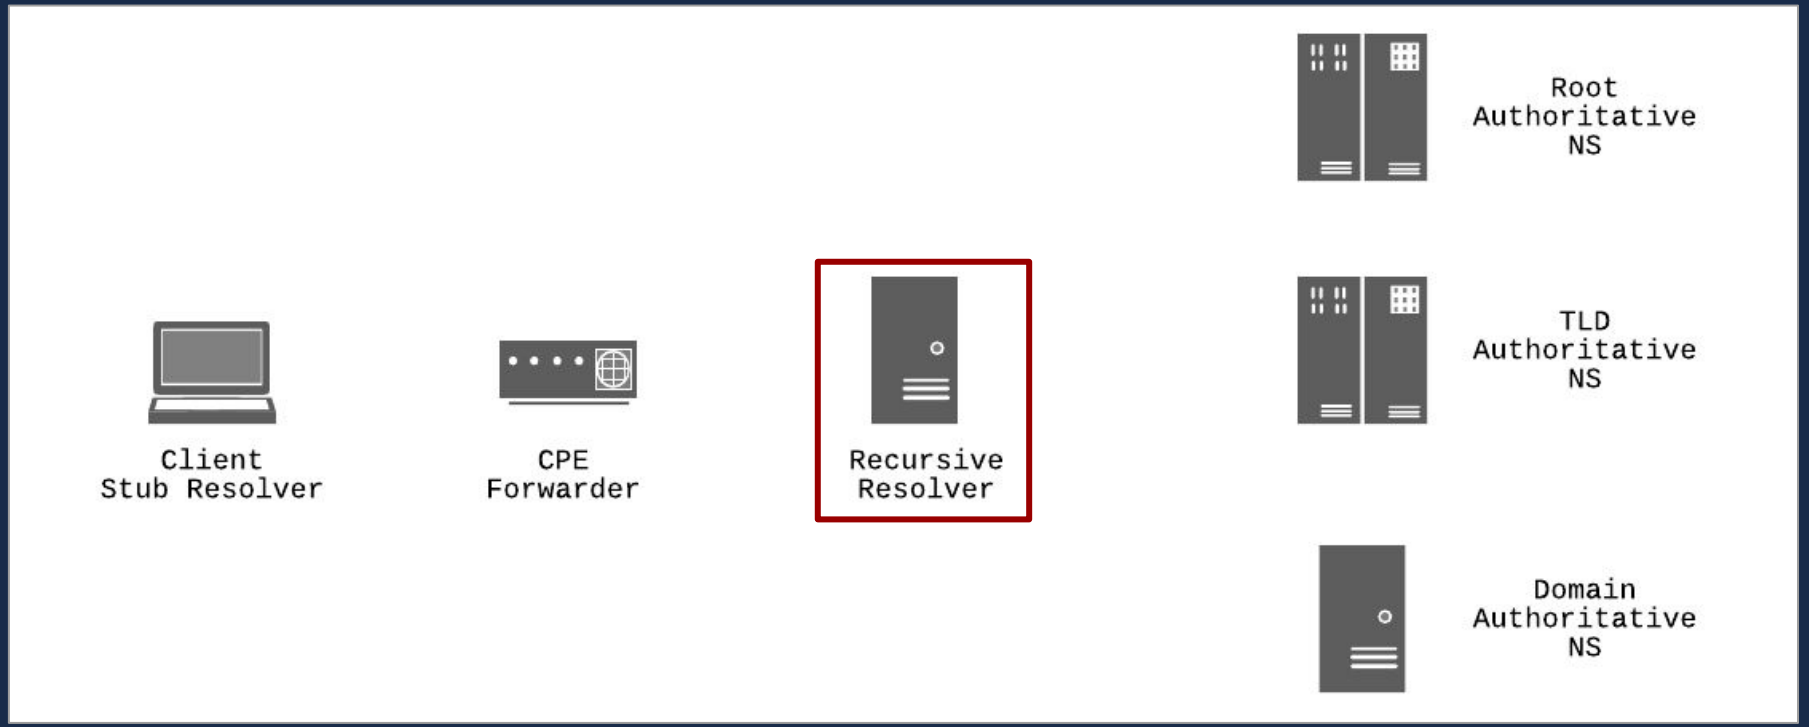

Recursive Resolvers

- Typically, assigned by DHCP. Defaults to ISP Nameservers.
- Recent shift to large public public resolvers.
    - Smaller ISPs default to using Google Public DNS.
- Google, Cloudflare, OpenDNS, Quad9 -- Also use anycast.
- Why use large public resolvers?
    - Can queries still be intercepted? Modified?

Discussion: Reliability

- Why use UDP? Why not TCP?
- Reliability through replication
  - Two authoritative nameserver per domain
  - What about root servers? TLD authoritative NS?
- Reliability across the entire life cycle?

- **Anycast adds another layer of reliability across the query life cycle!**

Discussion: Integrity

- Minimal security considerations in original DNS design.
- How to guarantee integrity of response?
  - Guarantee response has not been modified.
- But in order to do that, how to extend DNS?

Extension Mechanisms for DNS (EDNS)

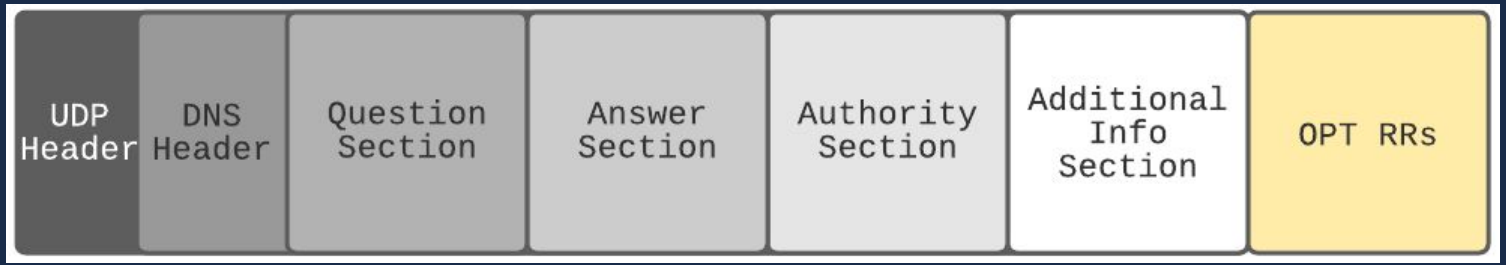

Add additional section to the end of a DNS packet.

EDNS sections skipped in old resolvers, and nameservers.

Used to implement DNSSEC, and ECS.

DNS Security Extension (DNSSEC)

- Add signature to DNS Records
  - Validate signature to ensure integrity of response
- Low adoption rate
  - Complicated to deploy
- Not all resolvers support DNSSEC.
  - Public DNS Resolvers support DNSSEC

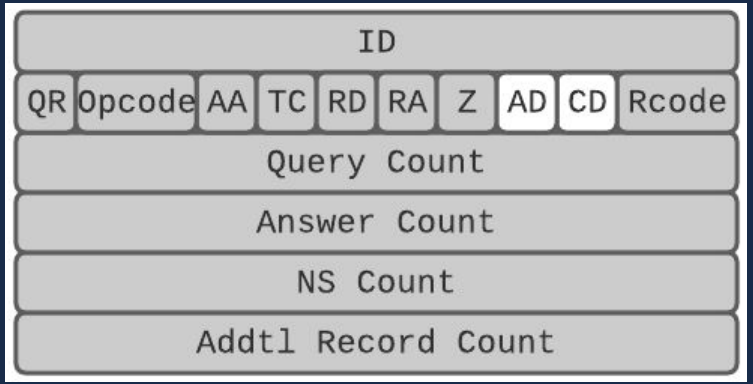

EDNS Client Subnet (ECS)

CDNs with large number of PoPs.

How to ensure response is mapped to closest PoP for client?

ECS allows recursive resolvers to supply the prefix of client IP

`137.110.222.10 → 137.110.222.0/24`

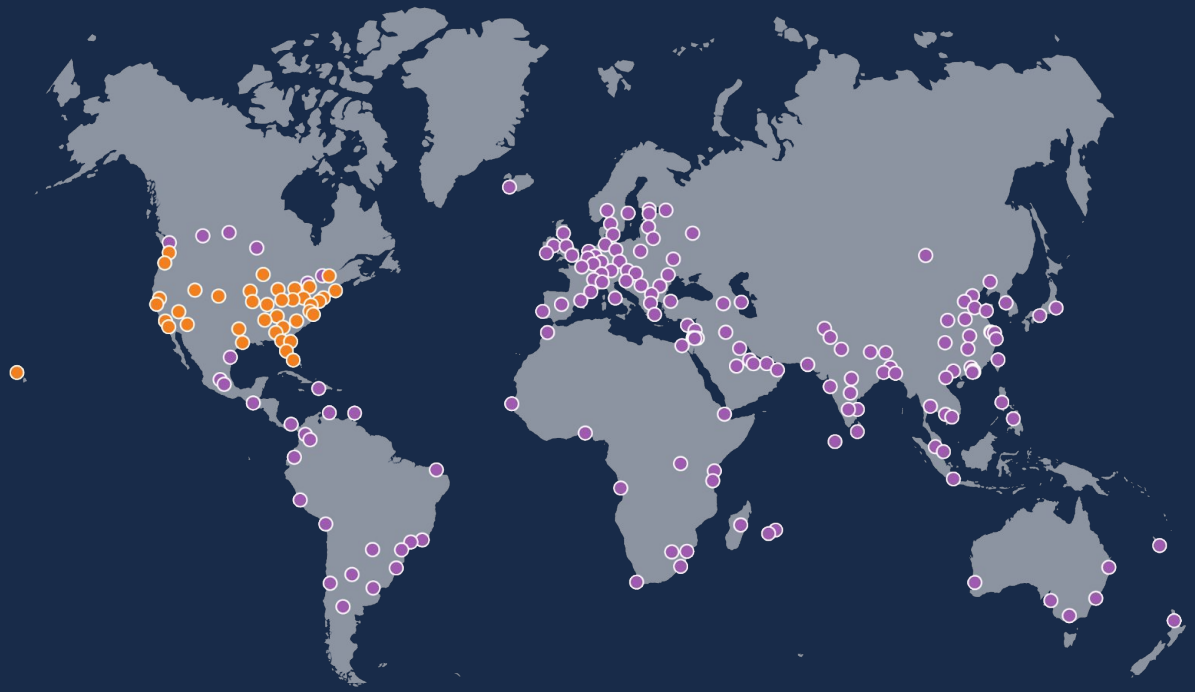

https://blog.cloudflare.com/introducing-regional-services/

Discussion: Confidentiality

Are DNS queries confidential?

Can we do better?

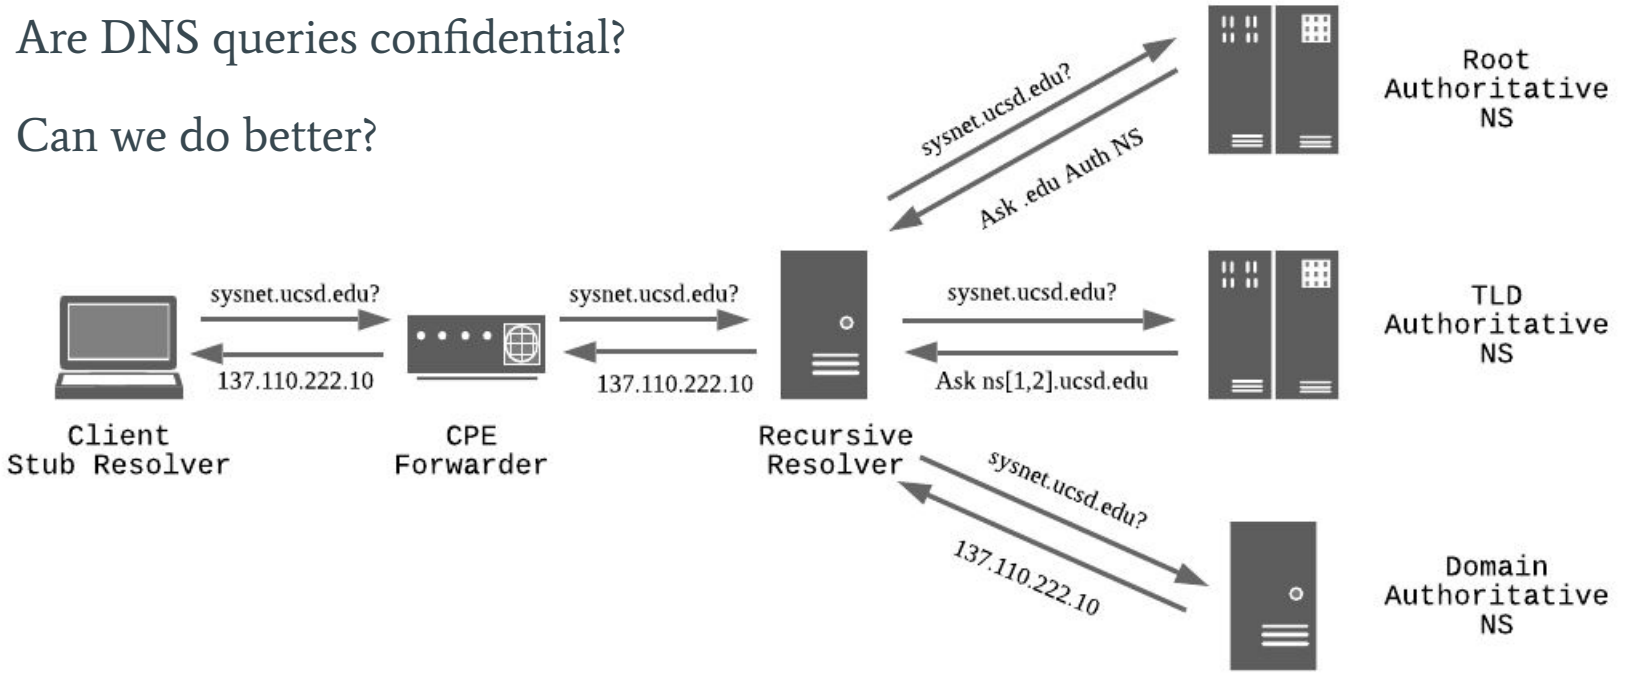

QName Minimization

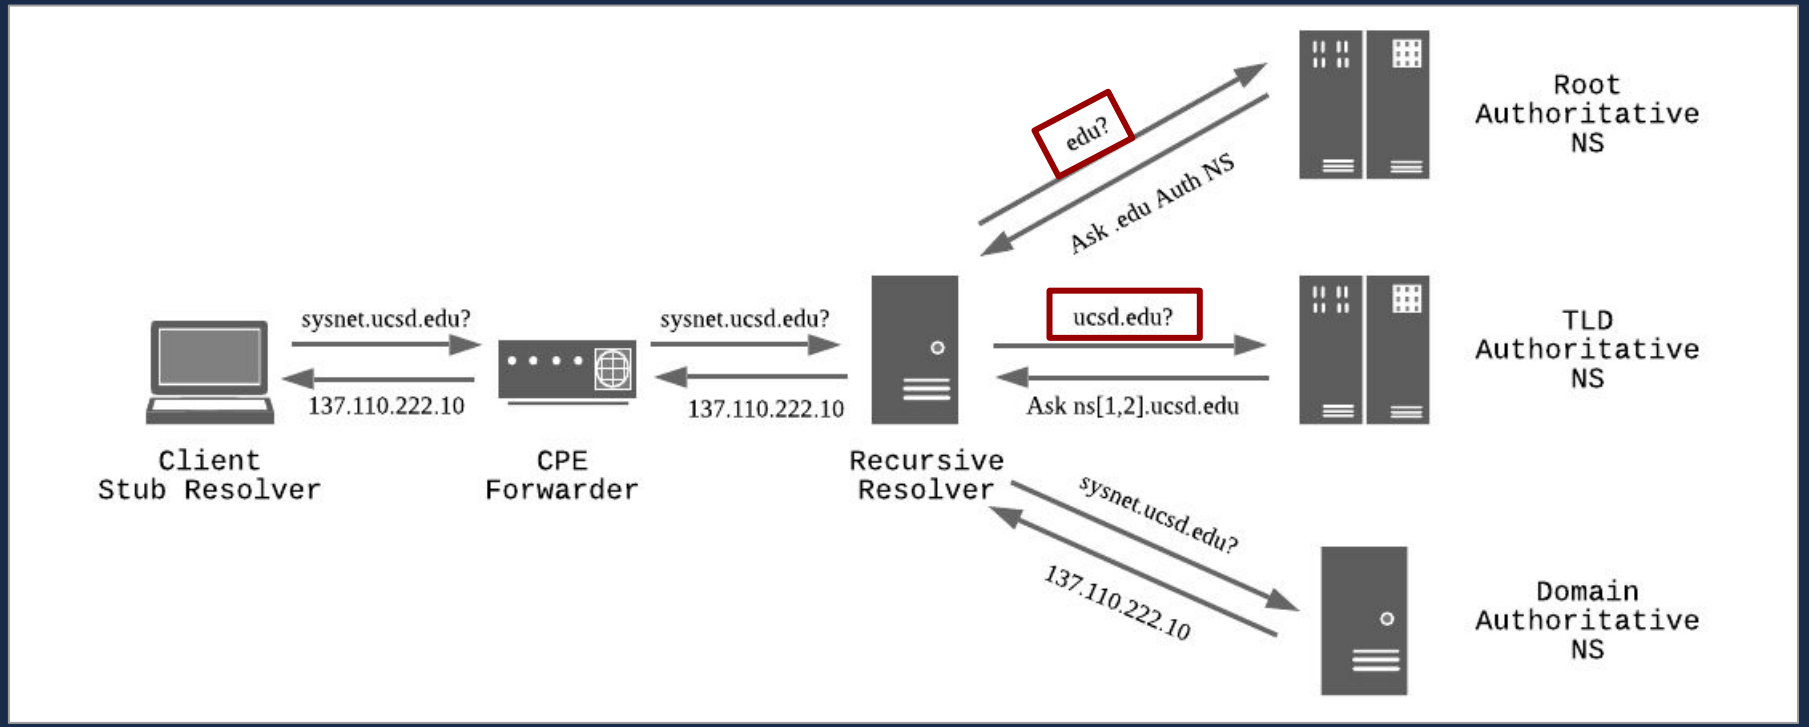

Encryption: DoH/DoT

DNS over HTTPs

DNS over TLS

Encrypted queries to recursive resolver?

Confidentiality? From whom?

What about ECS?

Questions?