CS 43: Computer Networks

Naming and DNS

September 26, 2024

Where we are

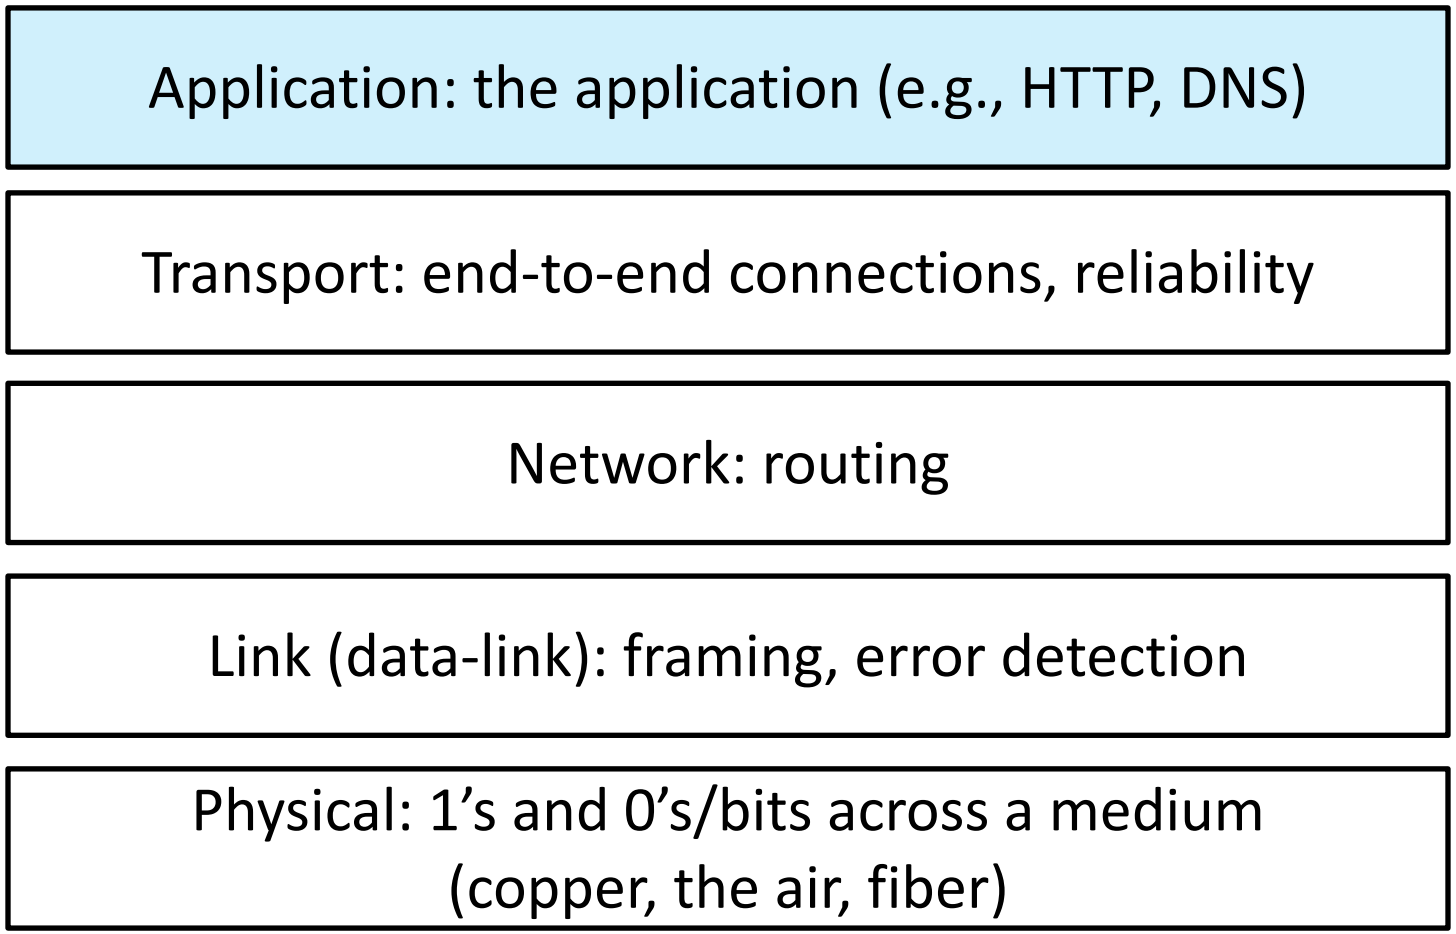
Application: the application (e.g., HTTP, DNS)
Transport: end-to-end connections, reliability
Network: routing
Link (data-link): framing, error detection
Physical: 1’s and 0’s/bits across a medium
(copper, the air, fiber)

Today

- Identifiers and addressing

- Domain Name System
  - Telephone directory of the Internet
  - Protocol format
  - Caching: Load balancing
  - Security Challenges

Goals of DNS

A wide-area distributed database

Possibly biggest such database in the world!

Goals

- Scalability; decentralized maintenance
- Robustness
- Global scope
- Names mean the same thing everywhere
- Distributed updates/queries
- Good performance

DNS Root Server and Politics

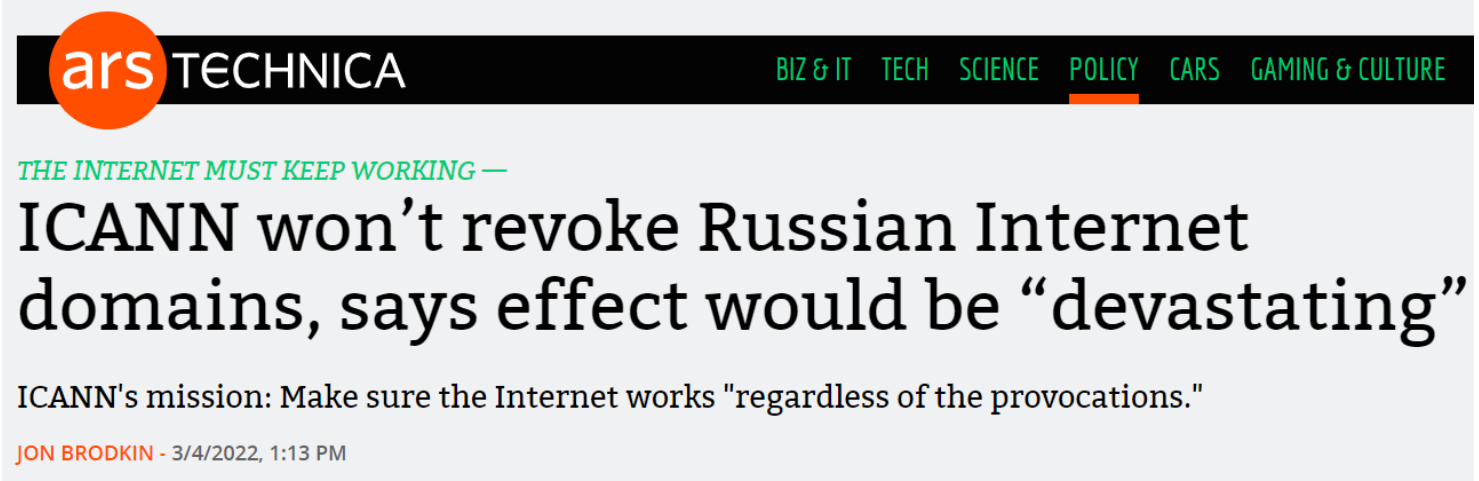

DNS: Application Layer Protocol

- distributed database
  - implemented in hierarchy of many name servers.
- application-layer protocol:
  - hosts communicate to name servers
  - resolve names → addresses
- *Core Internet function, implemented as application-layer protocol*

DNS: Domain Name System

People: many identifiers:

- name, swat ID, SSN, passport #

Internet hosts (endpoints), routers (devices inside a n/w):

- “name”, e.g., www.google.com - used by humans
- IP address (32 bit) - used for addressing packets

How do we map between IP address and name, and vice versa ?

Where

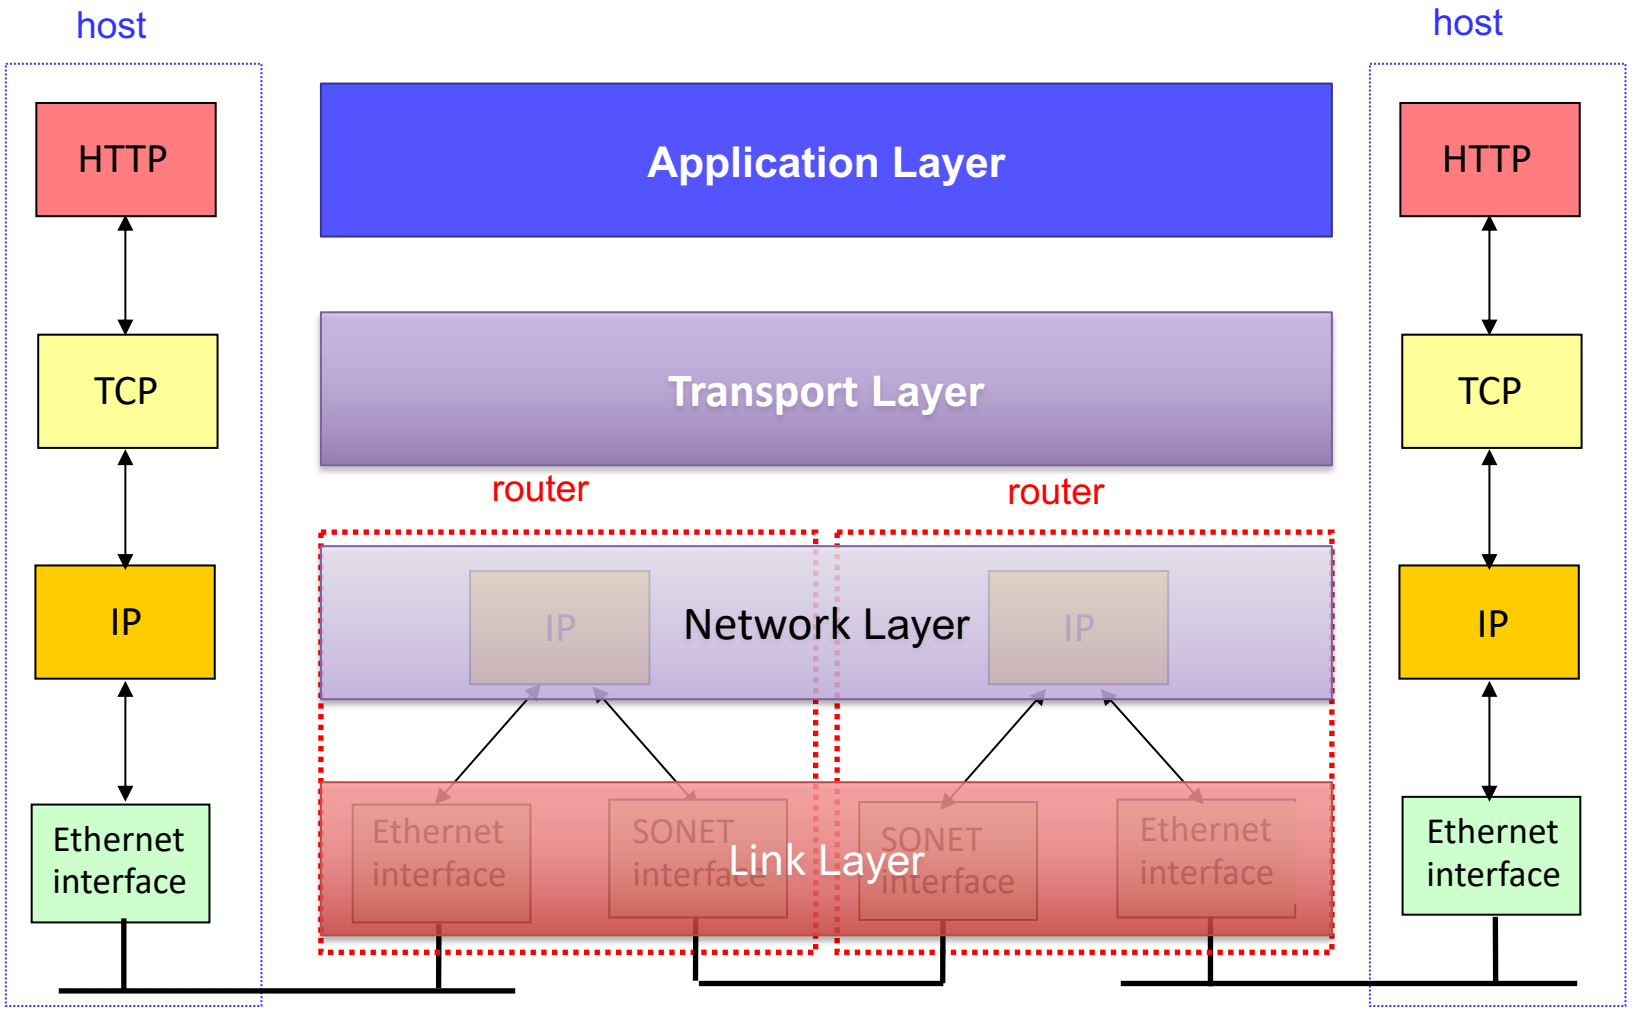
host
host
HTTP
TCP
IP
Ethernet interface
Application Layer
Transport Layer
router
router
IP
Network Layer
IP
Ethernet interface
SONET interface
Link Layer
SONET interface
Ethernet interface
HTTP
TCP
IP
Ethernet interface

Recall: TCP/IP Protocol Stack

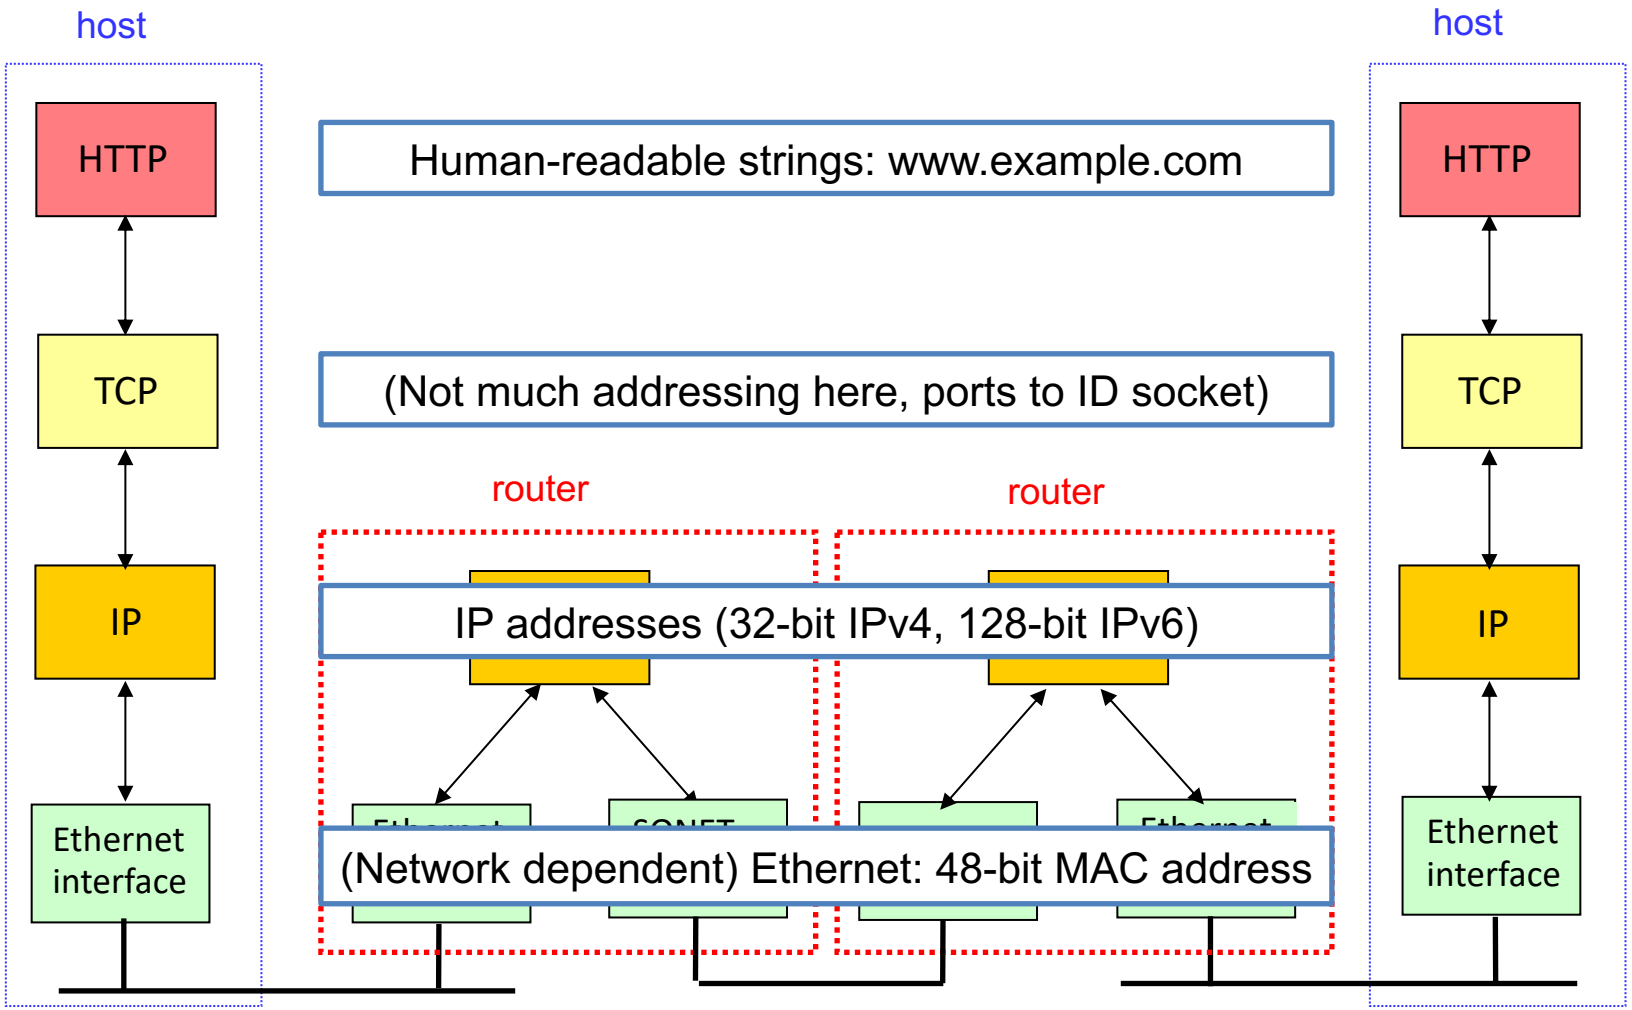

DNS: domain name system

- distributed database implemented in hierarchy of many name servers.
- application-layer protocol: hosts, name servers communicate to resolve names → addresses
  - *note: core Internet function, implemented as application-layer protocol*
  - *complexity at network's "edge"*

Why do we need to map names to IP addresses? Why not route on names at the network layer?

A. Domain names are hierarchical, so we can route on domain names too.
B. Domain names are variable length, vs IP are fixed length, some changes will be required to switch.
C. With domain names we wouldn’t know where to route to geographically.
D. Some other reason.`

Identifiers

- Host name (e.g., www.swarthmore.edu)
  - Used by humans to specify host of interest
  - Unique, selected by host administrator
  - Hierarchical, variable-length string of alphanumeric characters
- IP address (e.g., 130.58.68.164)
  - Used by routers to forward packets
  - Unique, topologically meaningful locator
  - Hierarchical namespace of 32 bits

Identifiers

- IP address (e.g., 130.58.68.164)
  - Used by routers to forward packets
  - Unique, topologically meaningful locator
  - Hierarchical namespace of 32 bits
- MAC address (e.g., D8:D3:85:94:5F:1E)
  - Used by network adaptors to identify frames
  - Unique, hard-coded identifier burned into network adaptor
  - Flat name space (of 48 bits in Ethernet)

What's in a name?

- Host name: web.cs.swarthmore.edu
  - Domain: registrar for each top-level domain (e.g., .edu)
  - Host name: local administrator assigns to each host

- IP addresses: 130.58.68.164
  - Prefixes: ICANN, regional Internet registries, and ISPs
  - Hosts: static configuration, or dynamic using DHCP

- MAC addresses: D8:D3:85:94:5F:1E
  - OIDs: assigned to vendors by the IEEE
  - Adapters: assigned by the vendor from its block

Mapping Between Identifiers

- Domain Name System (DNS)
  - Given a host name, provide the IP address
  - Given an IP address, provide the host name
- Address Resolution Protocol (ARP)
  - Given an IP address, provide the MAC address
  - To enable communication within the Local Area Network
- Dynamic Host Configuration Protocol (DHCP)
  - Automates host boot-up process
  - Given a MAC address, assign a unique IP address
  - … and tell host other stuff about the Local Area Network

What's the biggest challenge for DNS?

A. It's old.

B. The fact that the Internet is global.

C. The fact that DNS is now critical infrastructure.

D. The sheer number of name lookups happening at any given time.

E. How and when the name -> IP address mapping should change.

In the old days…

- Pre-1982, everyone downloads a "hosts.txt" file from SRI
- Pre-1998, Jon Postel, researcher at USC, runs the Internet Assigned Numbers Authority (IANA)
  - RFCs 882 & 883 in 1983
  - RFCs 1034 & 1035 in 1987

Emailed 8/12 root DNS servers, asked change to his authority. They did.

http://www.wired.com/wiredenterprise/2012/10/joe-postel/

Since 1998…

- Control of Internet Assigned Numbers Authority (IANA) transferred to Internet Corporation for Assigned Names and Numbers (ICANN)
  - ICANN is a private non-profit (formerly) blessed by US DOC
  - Global advisory committee for dealing with international issues
  - 2000's: Many efforts for UN control, US resisted
  - 2016: ICANN no longer partnered with DOC

Who should control DNS?

A. US government

B. UN / International government

C. Private corporation

D. Someone else

Recent Controversy

- Is ICANN working in the world's best interest?
- New "top level domains" added, for auction

Contributing to change

Internet standards ›

RFCs

Memos in the RFC document series contain technical and organizational notes about the Internet.

Open for Public Comment	Open Date	Close Date
Modification of Domains Protected Marks List Service	15 Aug 2018 23:59	24 Sep 2018 23:59

Division of Computer and Network Systems (CNS)

CNS invents new computing and networking technologies, while ensuring their security and privacy, and finds new ways to make use of current technologies.

Office of Science and Technology Policy

In 1976, Congress established the White House Office of Science and Technology Policy (OSTP) to provide the President and others within the Executive Office of the President with advice on the scientific, engineering, and technological aspects of the economy, national security, homeland security, health, foreign relations, the environment, and the technological recovery and use of resources, among other topics.

https://www.whitehouse.gov/ostp/ https://www.nsf.gov/cise/cns/about.jsp
https://www.icann.org/public-comments#open-public
https://www.ietf.org/standards/rfcs/

Uses of DNS

Hostname to IP address translation

- Reverse lookup: IP address to hostname translation

Host name aliasing: other DNS names for a host

- Alias hostnames point to canonical hostname

Email: look up domain's mail server by domain name

Different DNS Mappings

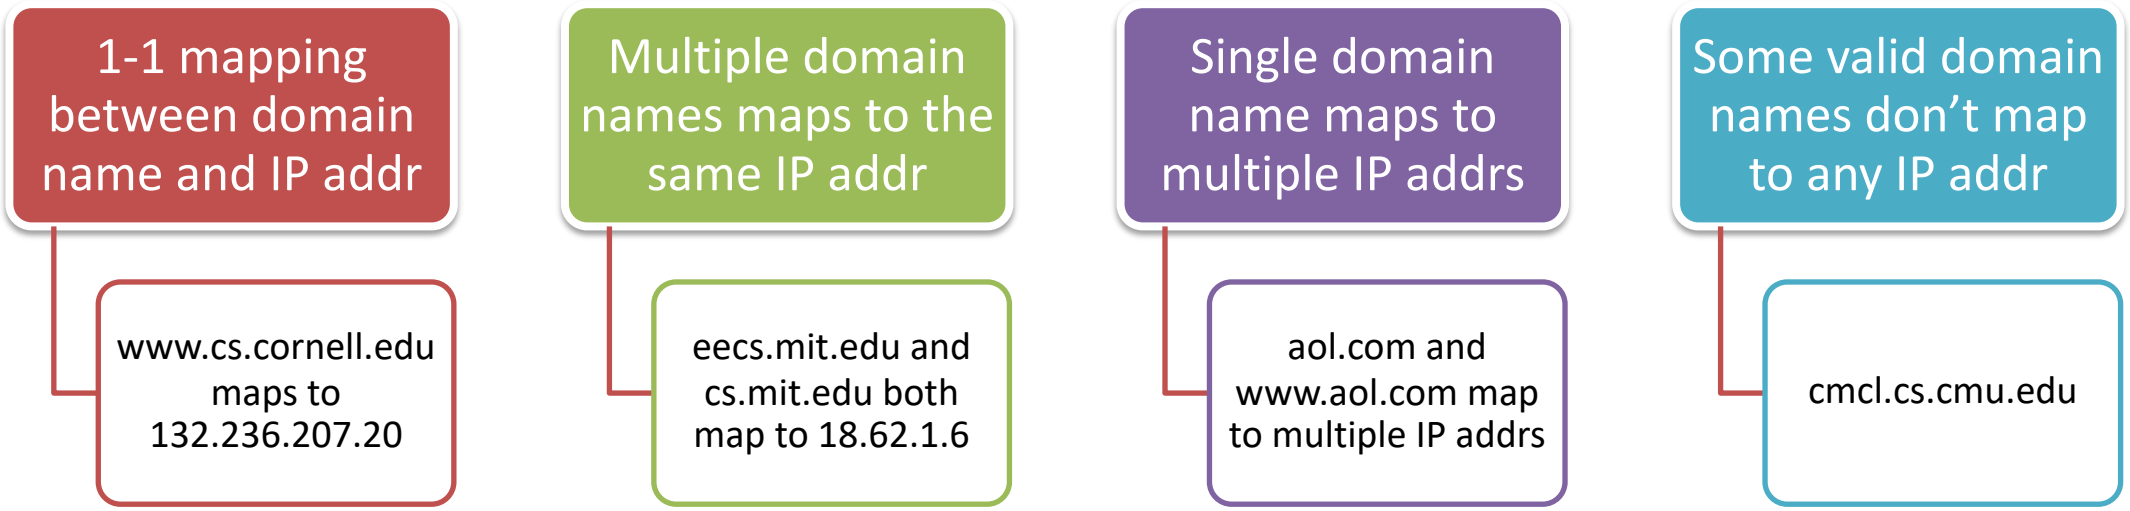

DNS Services

- DNS is an application-layer protocol. E2E design!
- It provides:
  - Hostname to IP address translation
  - Host aliasing (canonical and alias names)
  - Mail server aliasing
  - Load distribution (one name may resolve to multiple IP addresses)
  - Lots of other stuff that you might use a directory service to find. (Wikipedia: List of DNS record types)

DNS Records

DNS: distributed DB storing resource records (RR)

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" (the real) name
- **www.ibm.com** is really
- servereast.backup2.ibm.com
- **value** is canonical name

type=MX

- **value** is name of mailserver associated with name

DNS Packet

Lightweight exchange of *query* and *reply* messages, both with same message format

Primarily uses UDP for its transport protocol, which is what we'll assume

Frequently, both clients and servers use port 53

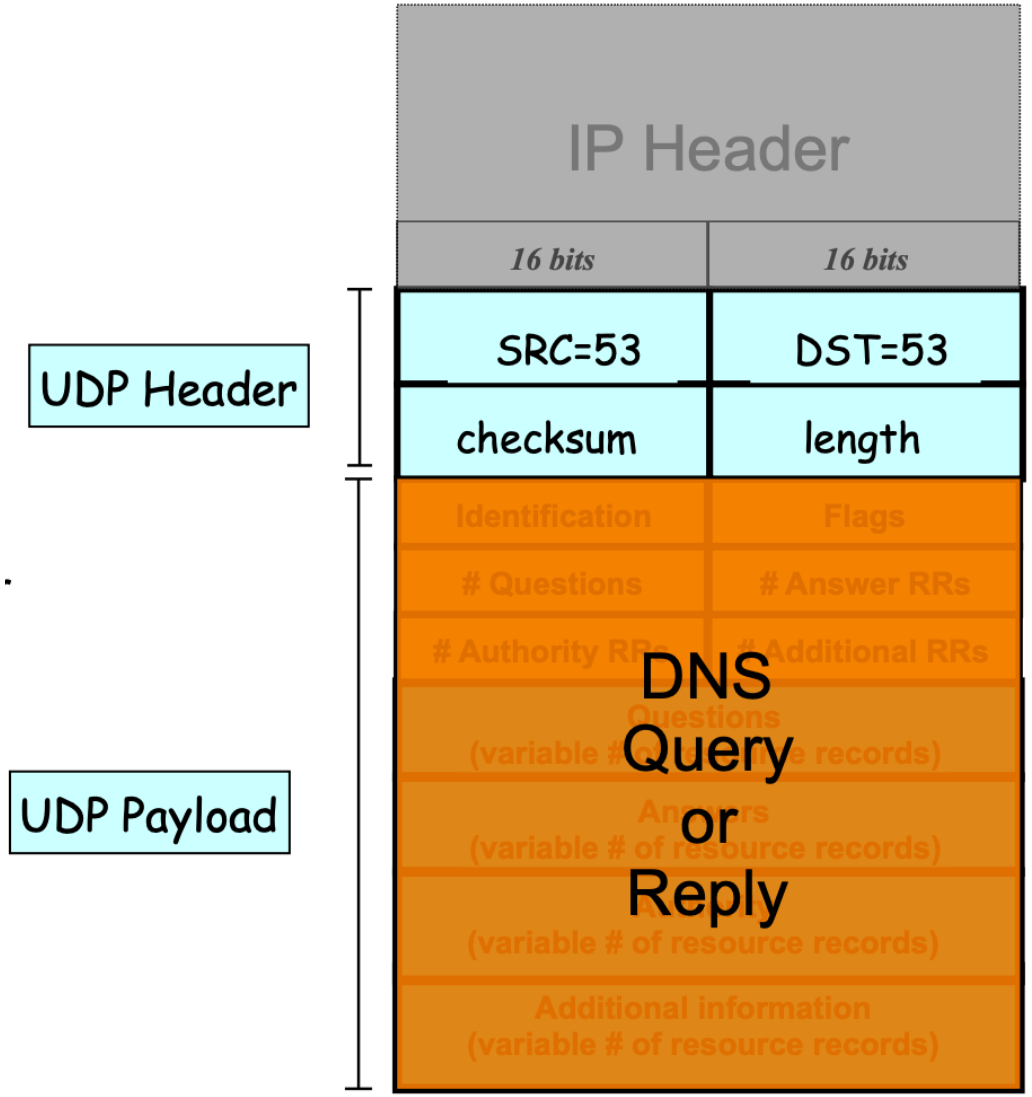

DNS Packet over UDP (no reliability!)

DNS requests sent over UDP

Four sections: questions, answers, authority, additional records

Query ID:
16 bit random value
Links response to query

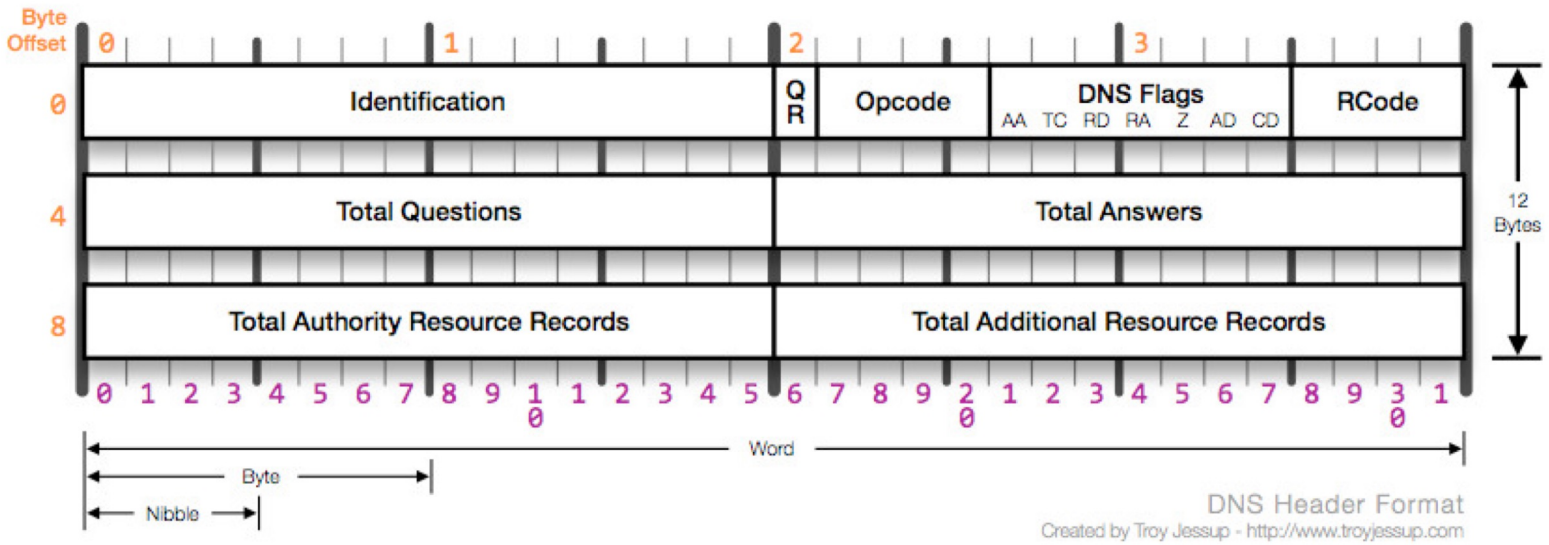

DNS protocol, messages

- query and reply messages, both with same message format

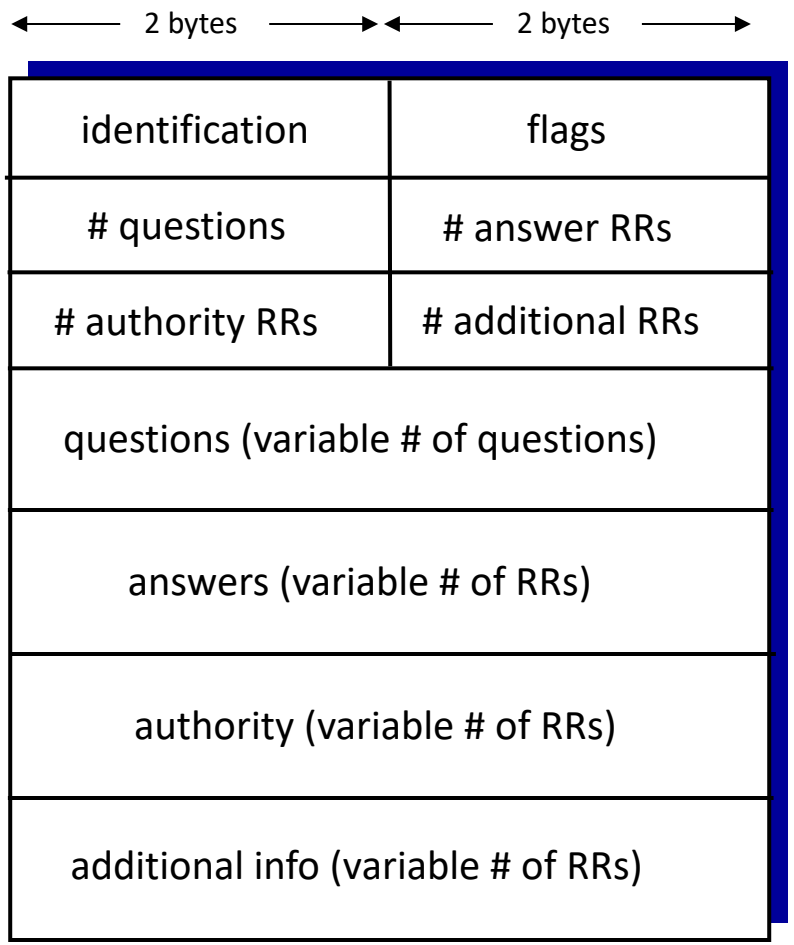

DNS protocol, messages

- query and reply messages, both with same message format

Message header

- identification: 16 bit # for query, reply to query uses same #
- flags:
  - query or reply
  - recursion desired
  - recursion available
  - reply is authoritative

Sent via UDP!

- No connection established
- Not reliable

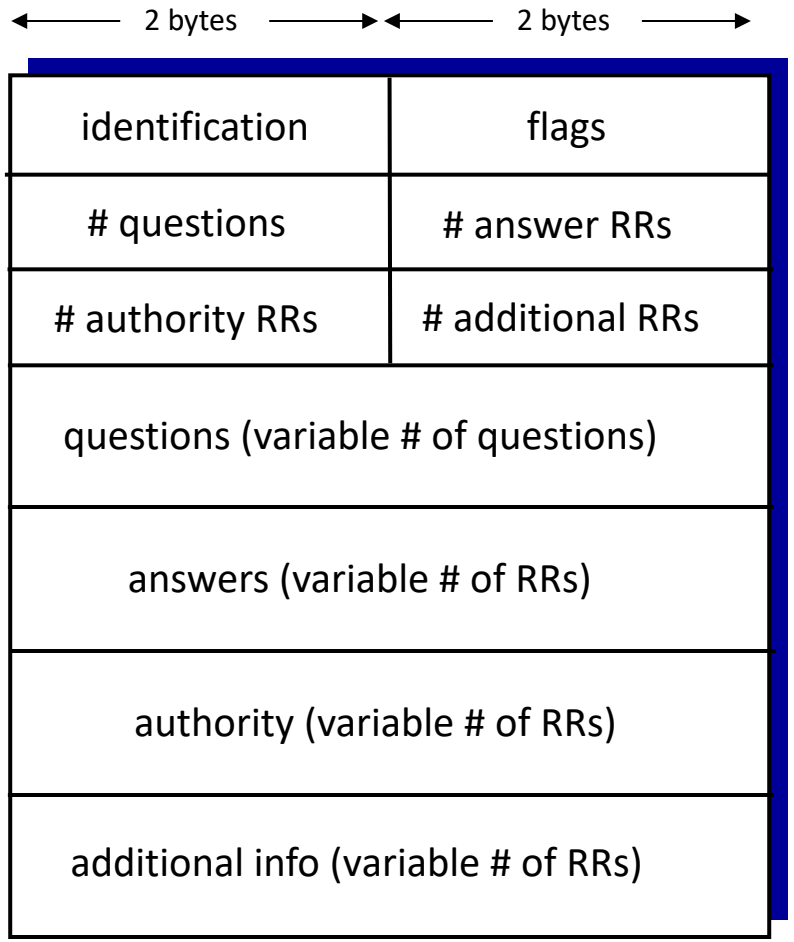

Request

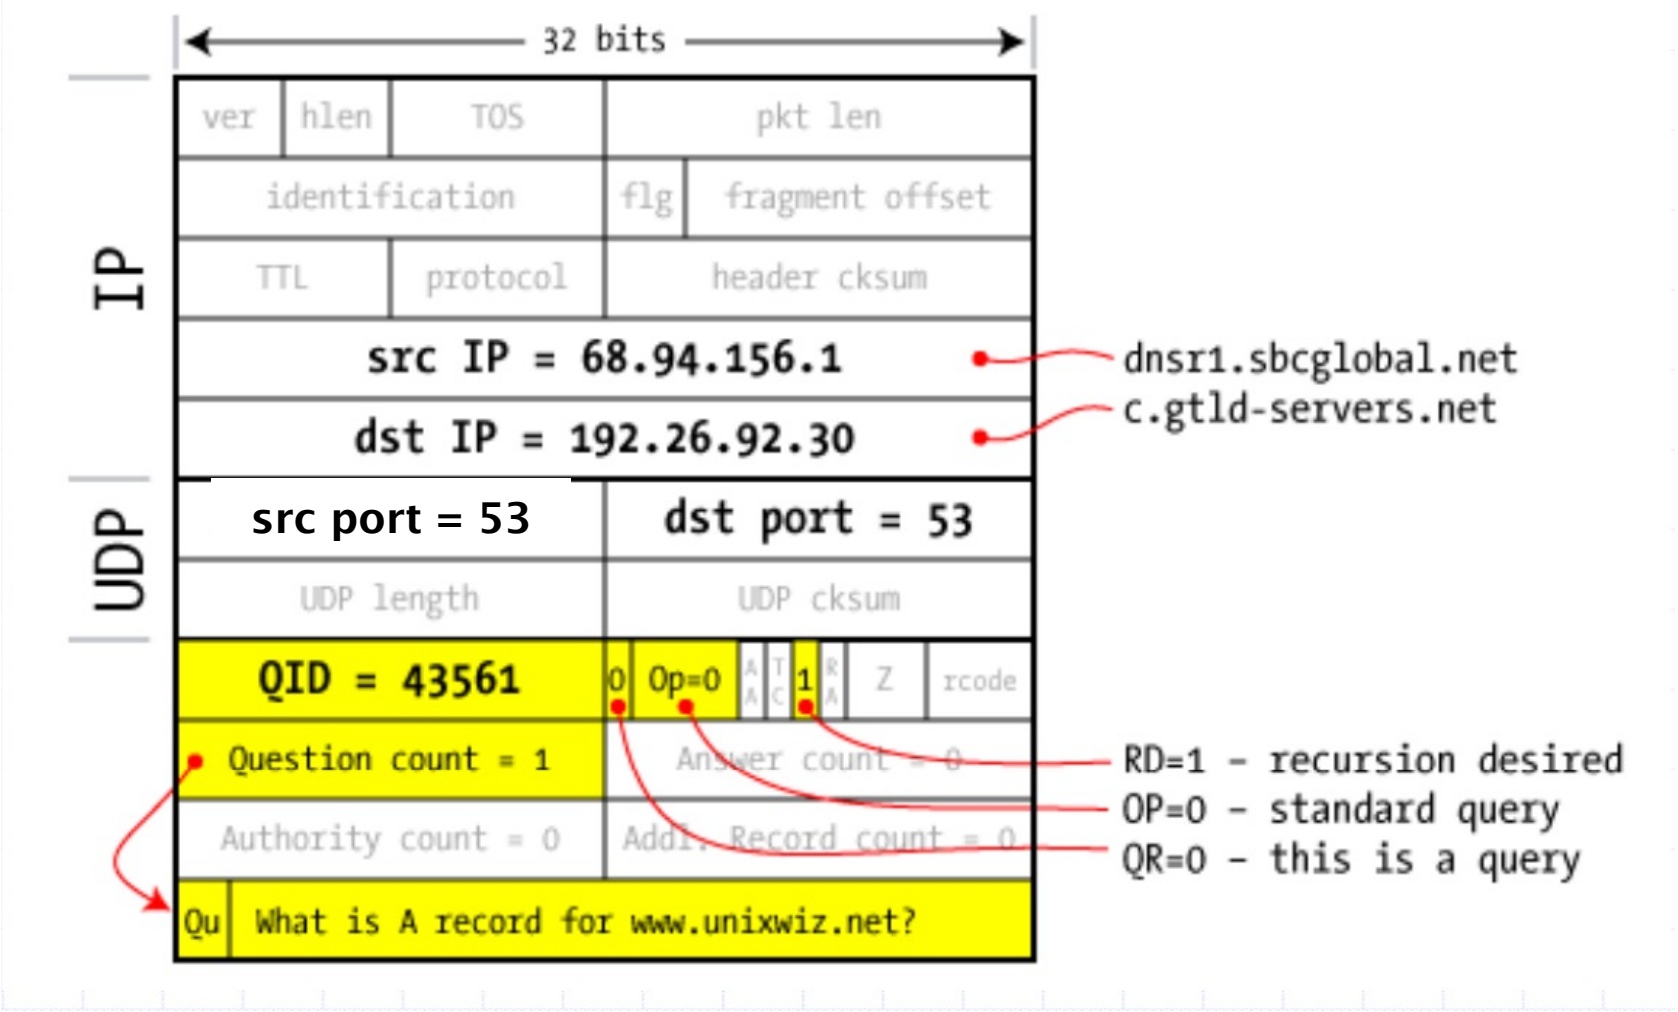
32 bits
IP
UDP
ver
hlen
TOS
pkt len
identification
flg
fragment offset
TTL
protocol
header cksum
src IP = 68.94.156.1
dst IP = 192.26.92.30
src port = 53
dst port = 53
UDP length
UDP cksum
QID = 43561
0
Op=0
1
Z
rcode
Question count = 1
Authority count = 0
Qu
What is A record for www.unixwiz.net?
dnsr1.sbcglobal.net
c.gtld-servers.net
RD=1 - recursion desired
OP=0 - standard query
QR=0 - this is a query

Response

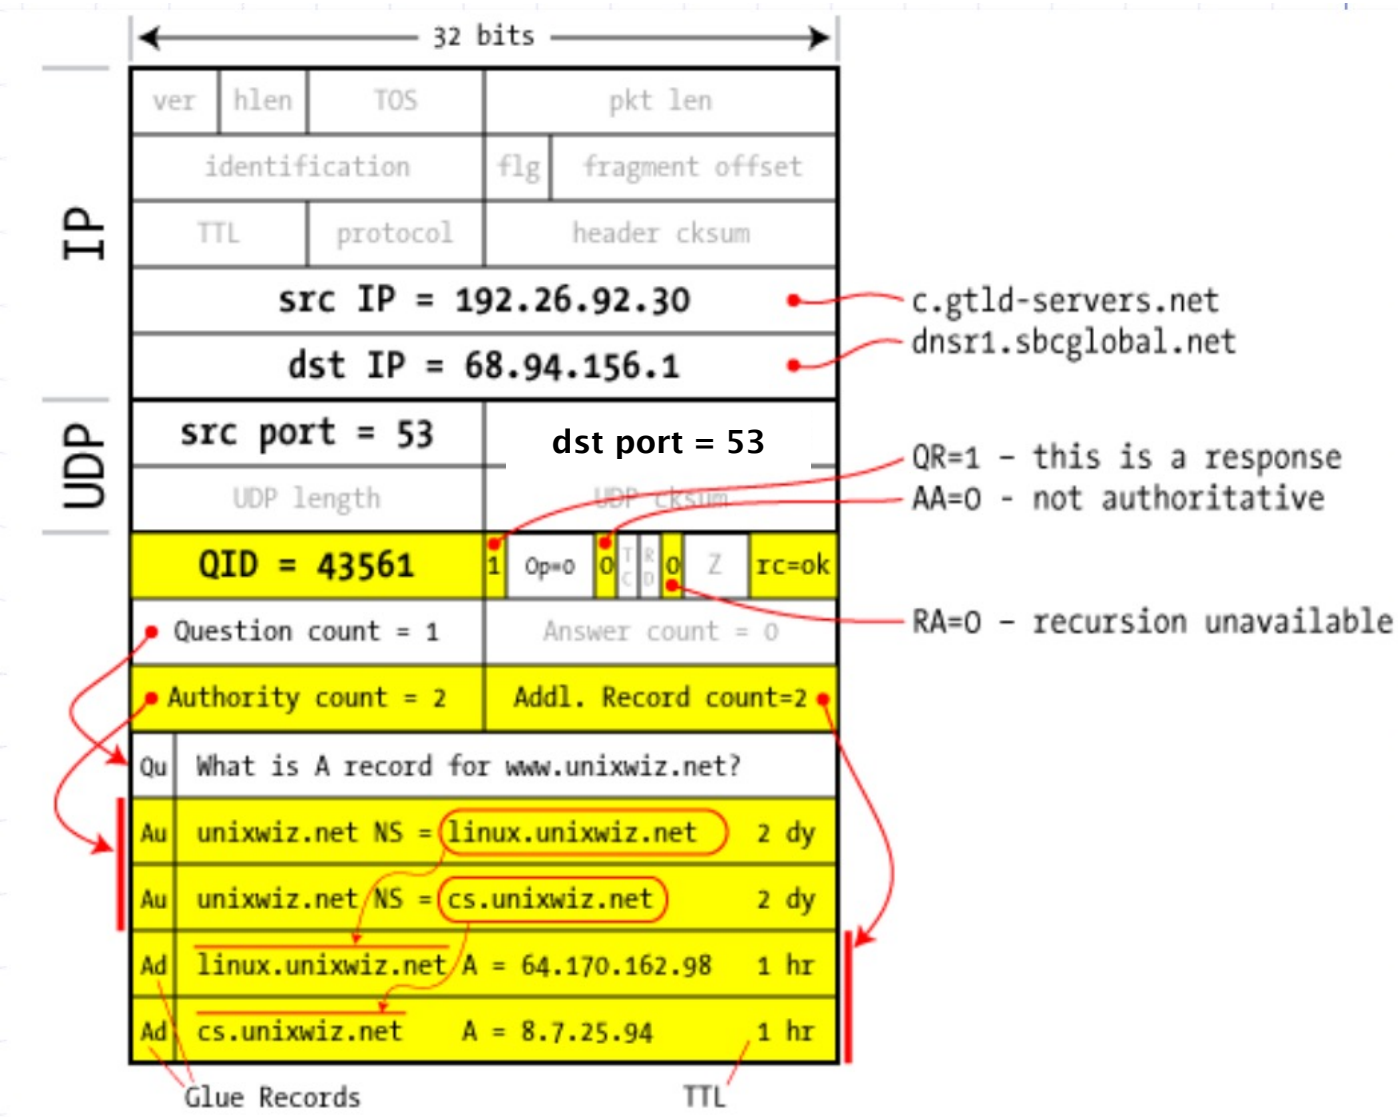
32 bits
IP
ver
hlen
TOS
pkt len
identification
flg
fragment offset
TTL
protocol
header cksum
src IP = 192.26.92.30
dst IP = 68.94.156.1
UDP
src port = 53
dst port = 53
UDP length
UDP cksum
QID = 43561
1
Op=0
0
0
Z
rc=ok
Question count = 1
Answer count = 0
Authority count = 2
Addl. Record count=2
Qu
What is A record for www.unixwiz.net?
Au
unixwiz.net NS = linux.unixwiz.net
2 dy
Au
unixwiz.net NS = cs.unixwiz.net
2 dy
Ad
linux.unixwiz.net A = 64.170.162.98
1 hr
Ad
cs.unixwiz.net A = 8.7.25.94
1 hr
c.gtld-servers.net
dnsr1.sbcglobal.net
QR=1 - this is a response
AA=0 - not authoritative
RA=0 - recursion unavailable
Glue Records
TTL

Authoritative Response

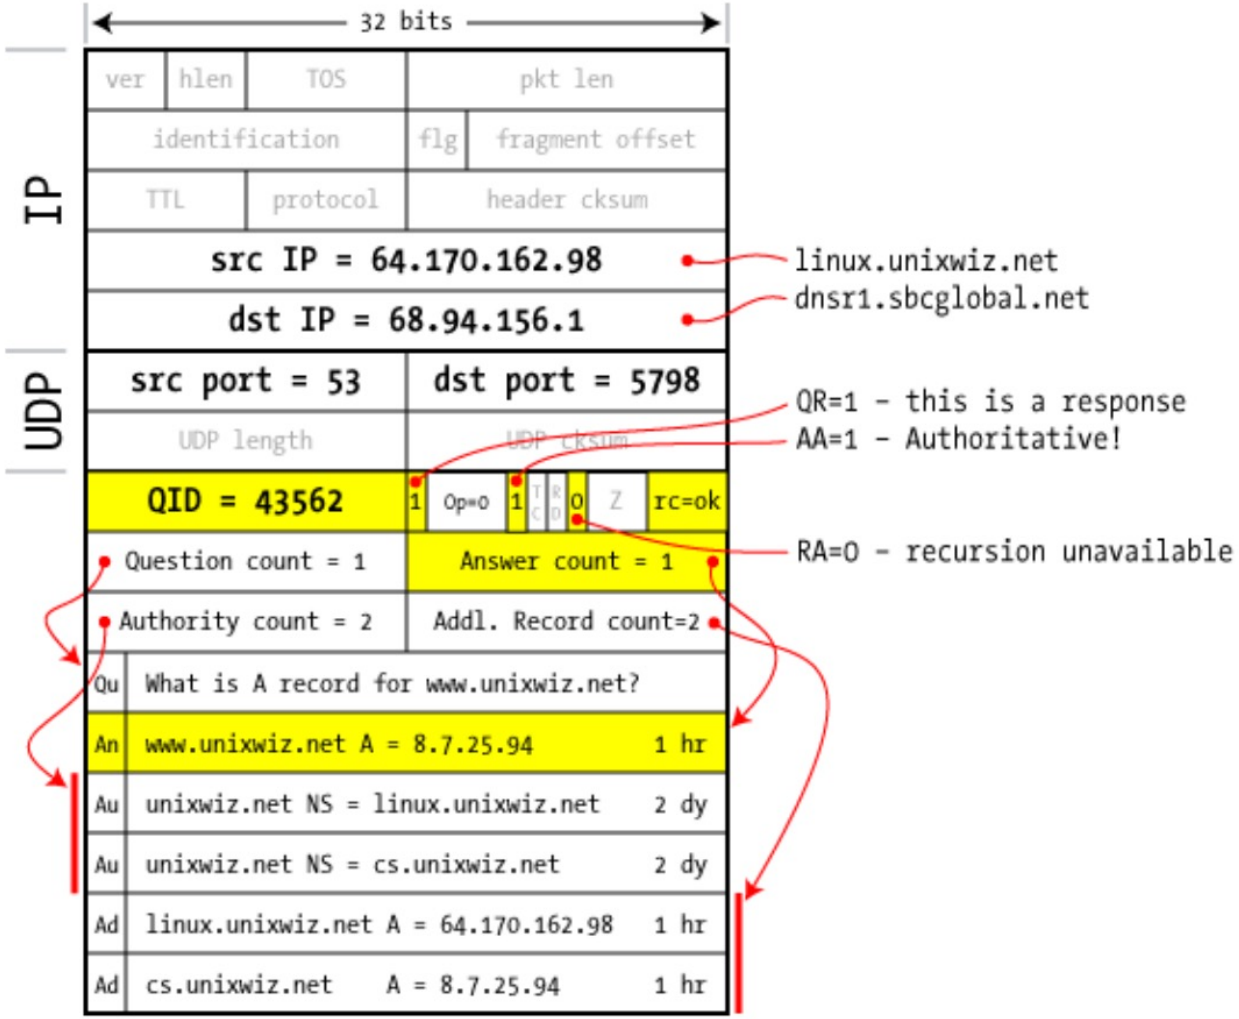

Domain Name System (DNS)

- Distributed administrative control
  - Hierarchical name space divided into zones
  - Distributed over a collection of DNS servers
- Hierarchy of DNS servers
  - Root servers
  - Top-level domain (TLD) servers
  - Authoritative DNS servers
- Performing the translations
  - Local DNS servers
  - Resolver software

DNS: a distributed, hierarchical database

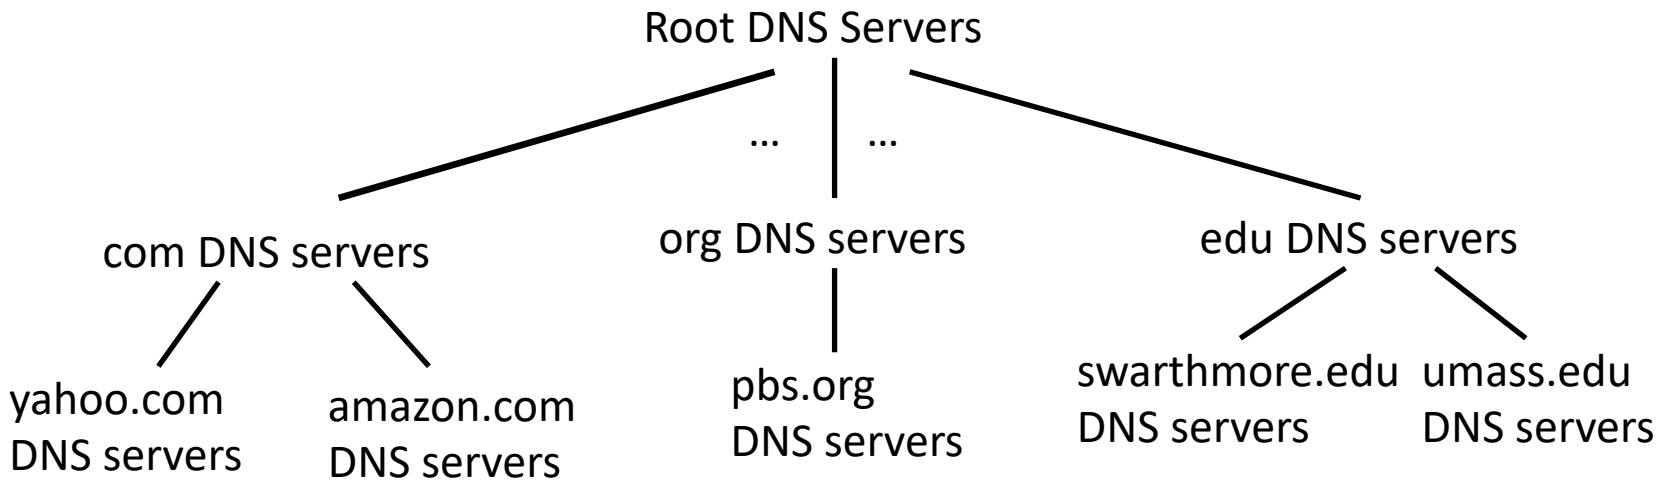

DNS: a distributed, hierarchical database

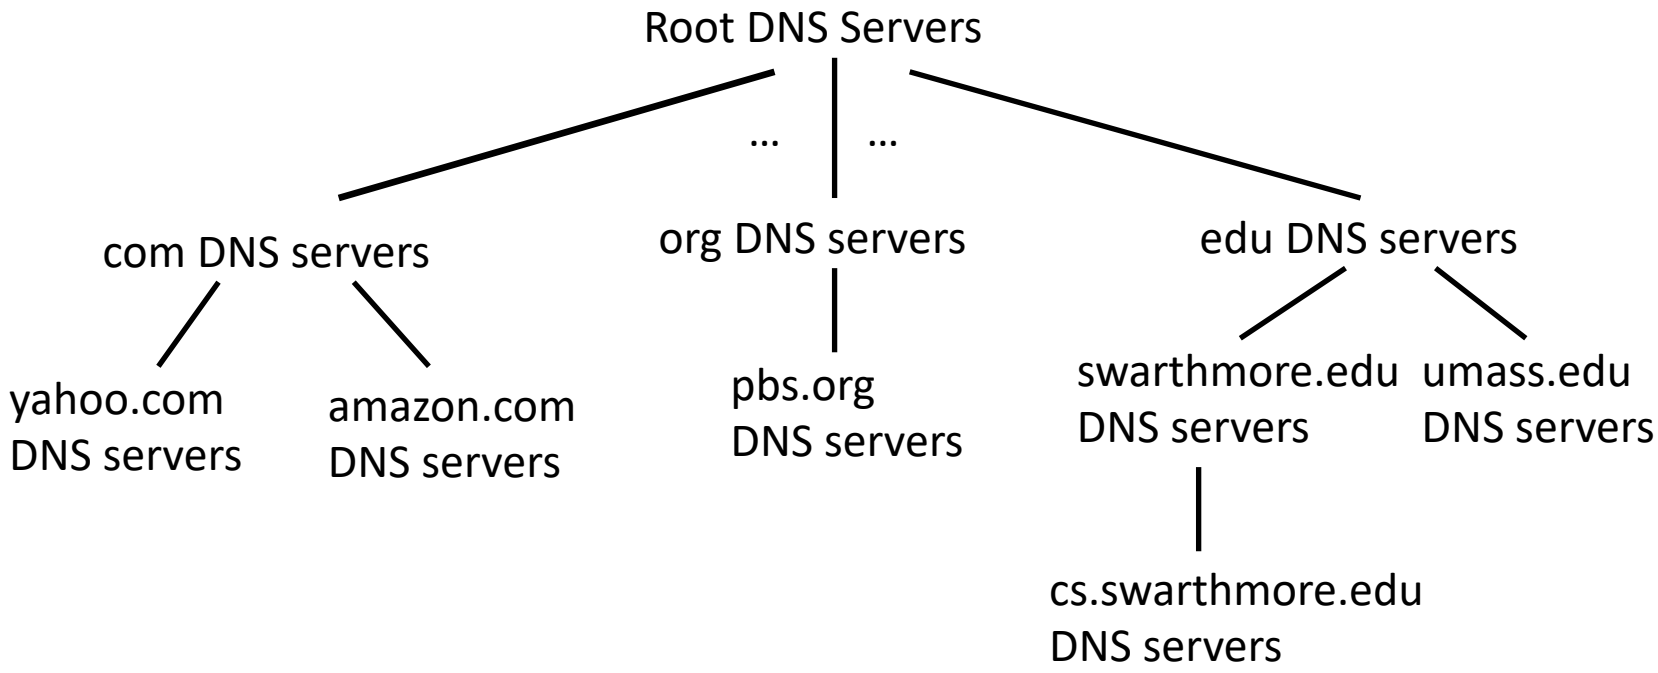

DNS: a distributed, hierarchical database

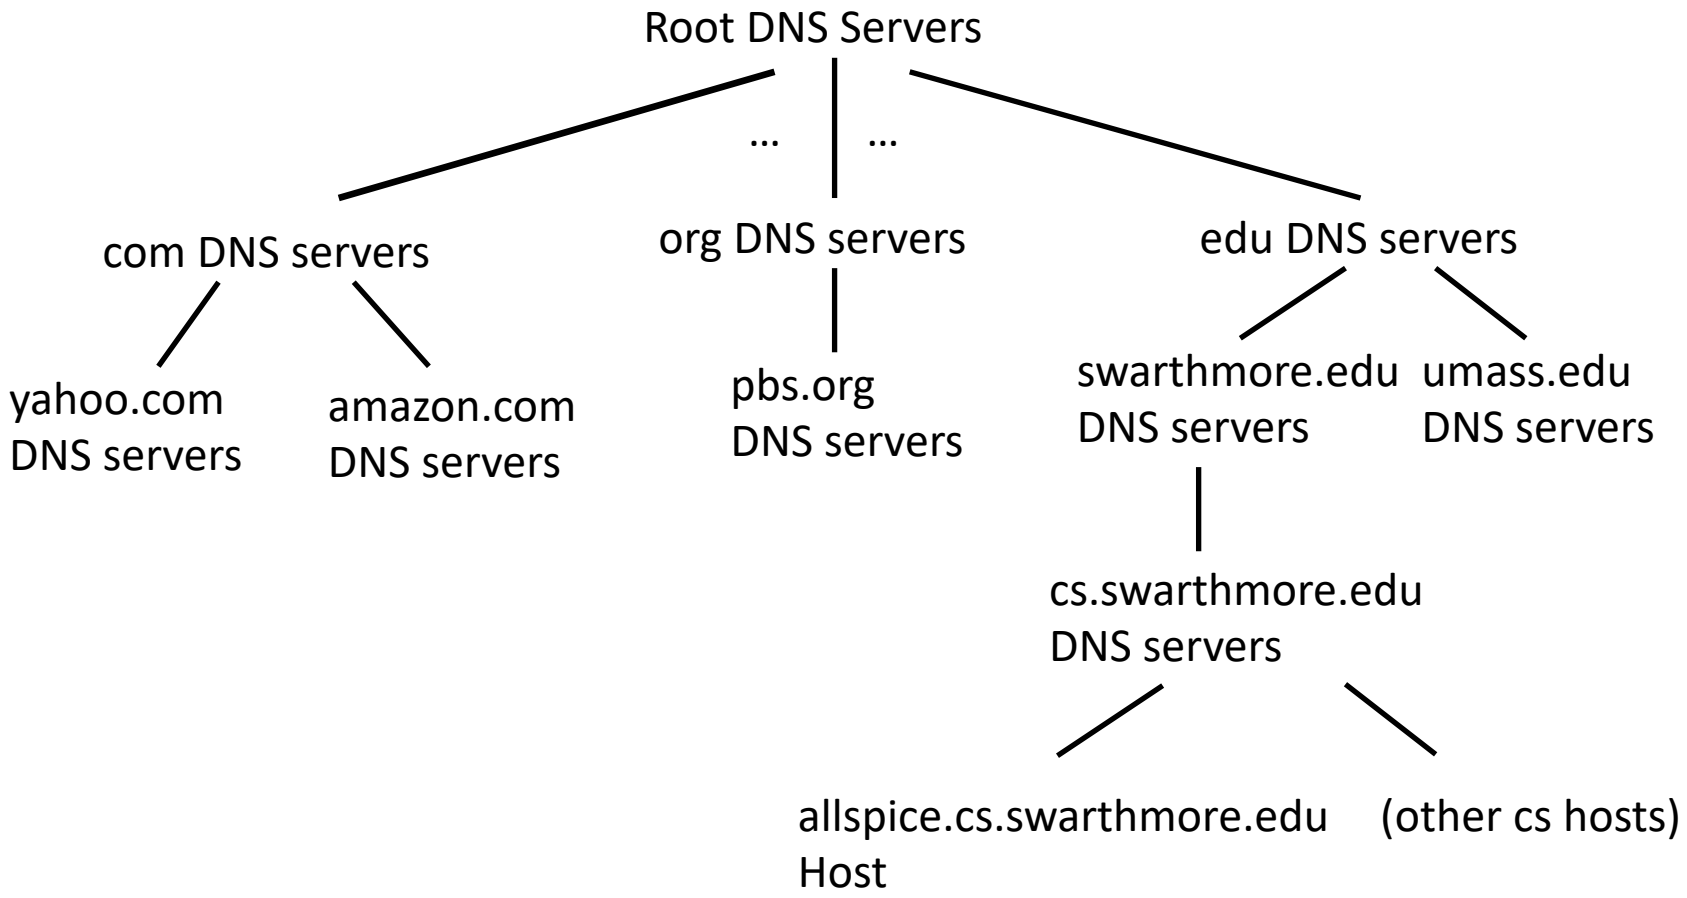

- allspice.cs.swarthmore.edu.

Nameless root,
Usually implied.

Why do we structure DNS like this? Which of these helps the most? Drawbacks?

A. It divides up responsibility among parties.

B. It improves performance of the system.

C. It reduces the size of the state that a server needs to store.

D. Some other reason.

DNS: a distributed, hierarchical database

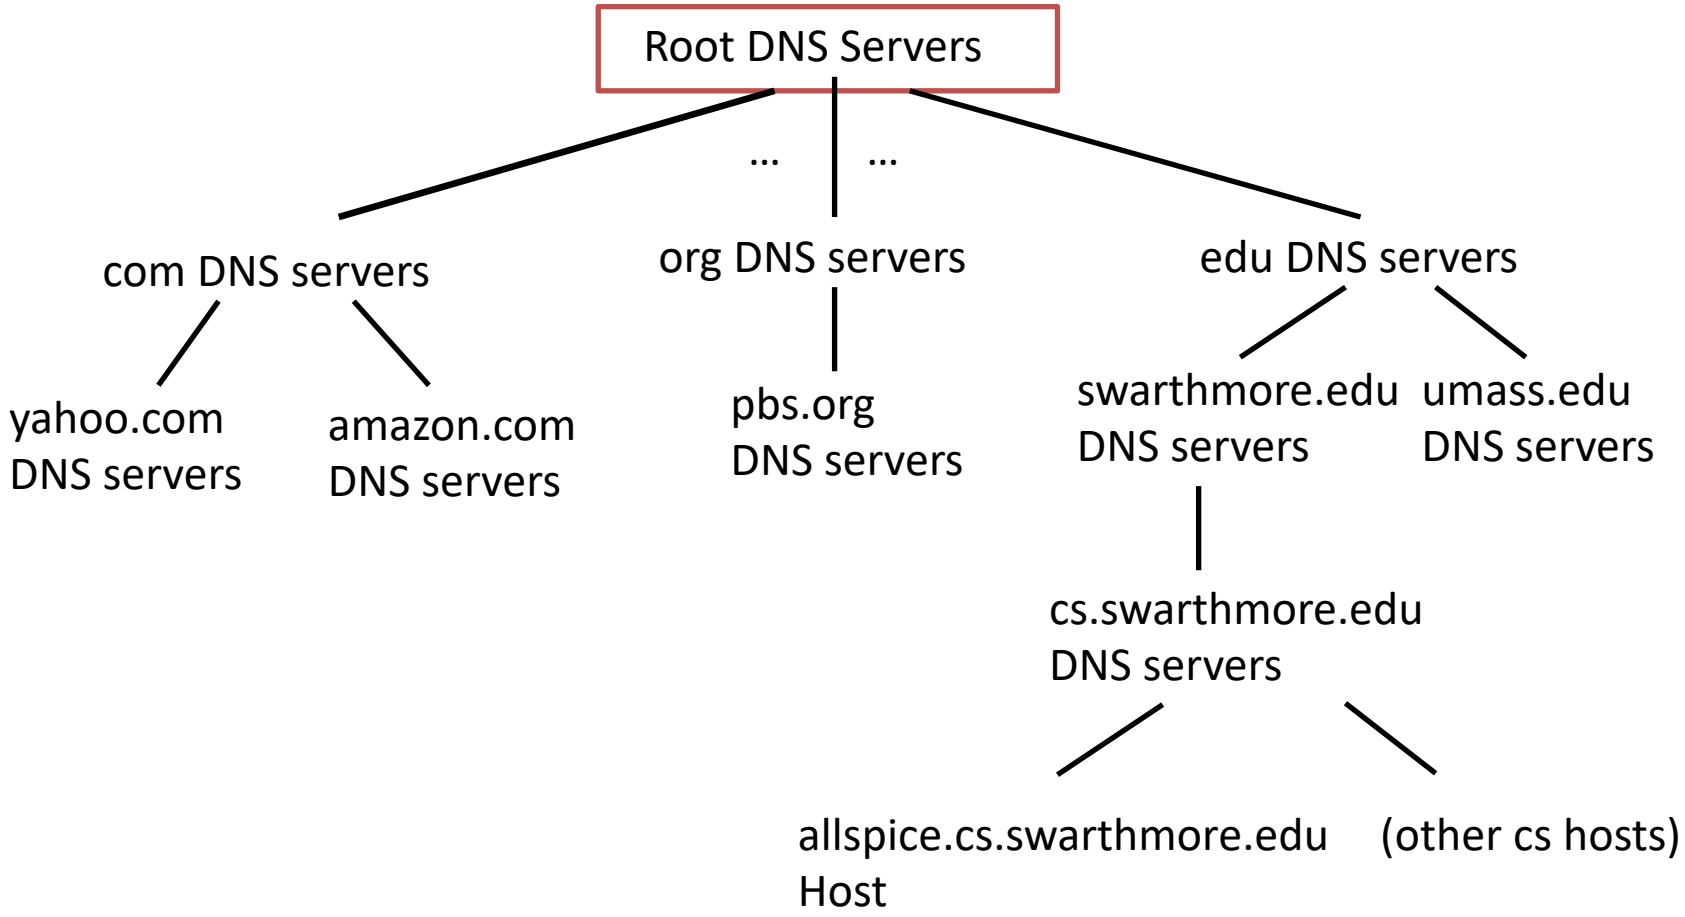

DNS: Root Name Servers

- Root name server:
  - Knows how to find top-level domains (.com, .edu, .gov, etc.)
  - How often does the location of a TLD change?

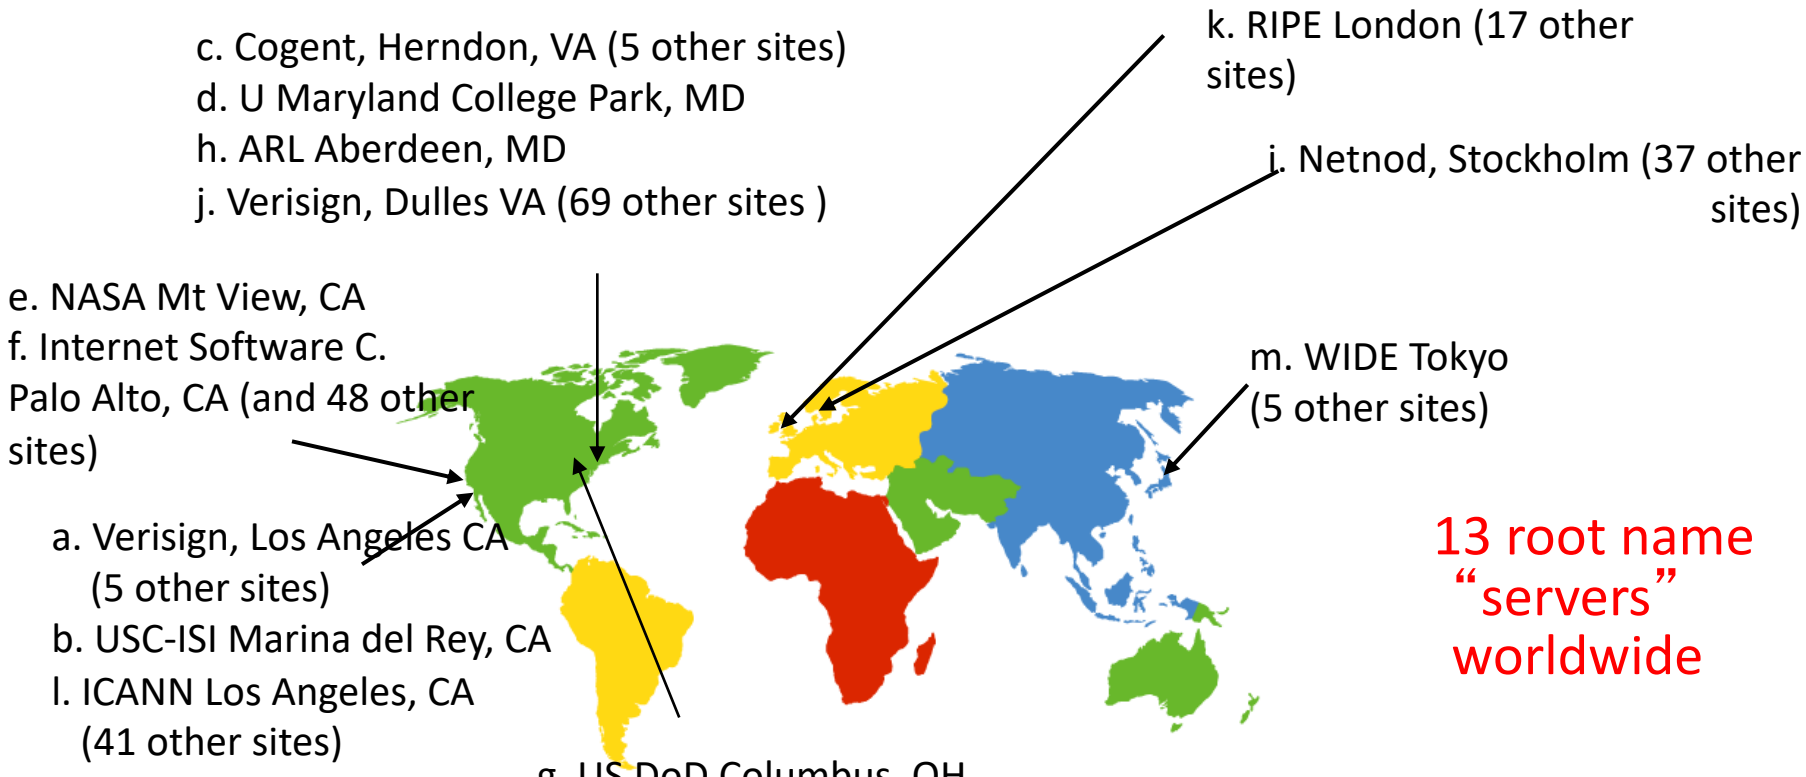

13 root name "servers" worldwide

DNS: Root Name Servers

- Root name server:
  - Knows how to find top-level domains (.com, .edu, .gov, etc.)
  - How often does the location of a TLD change?
  - approx. 400 total root servers
  - Significant amount of traffic is not legitimate

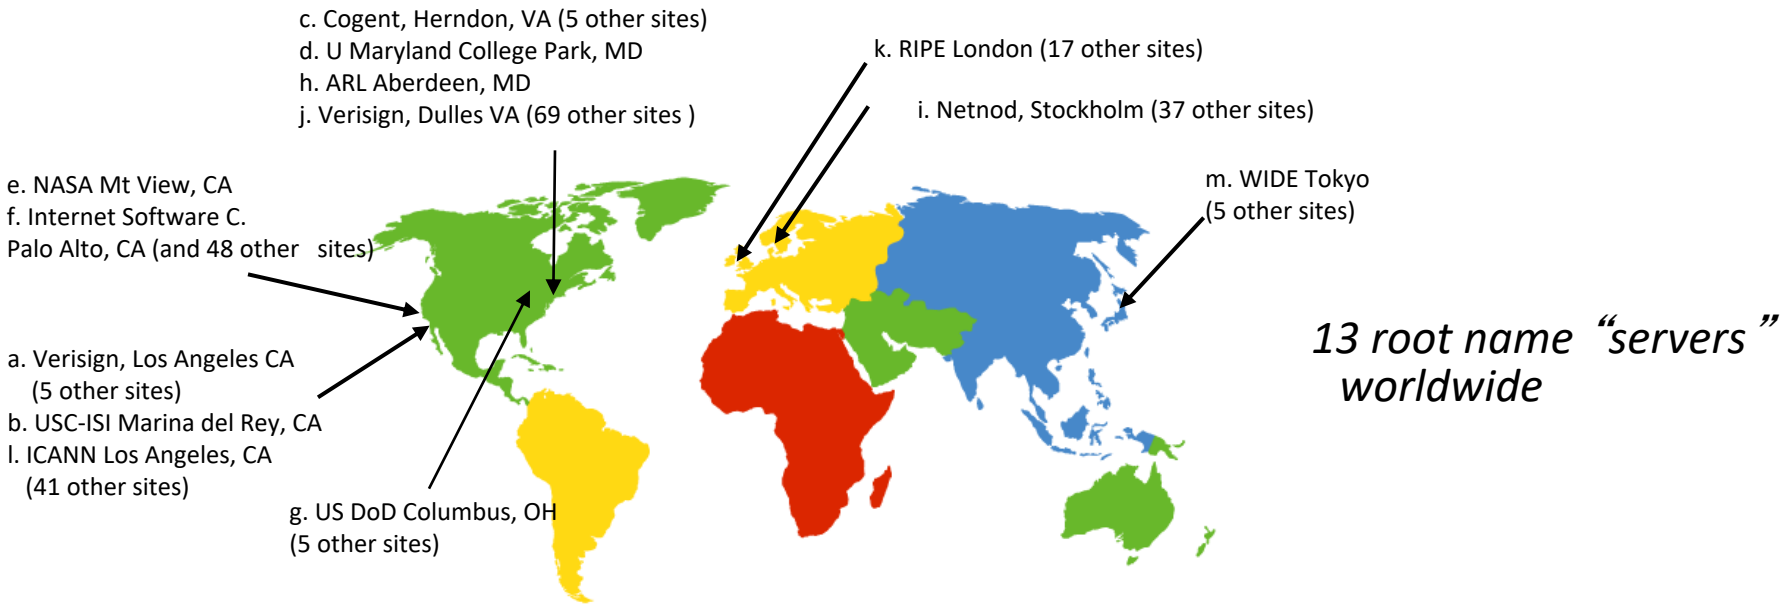

13 root name "servers" worldwide

DNS: a distributed, hierarchical database

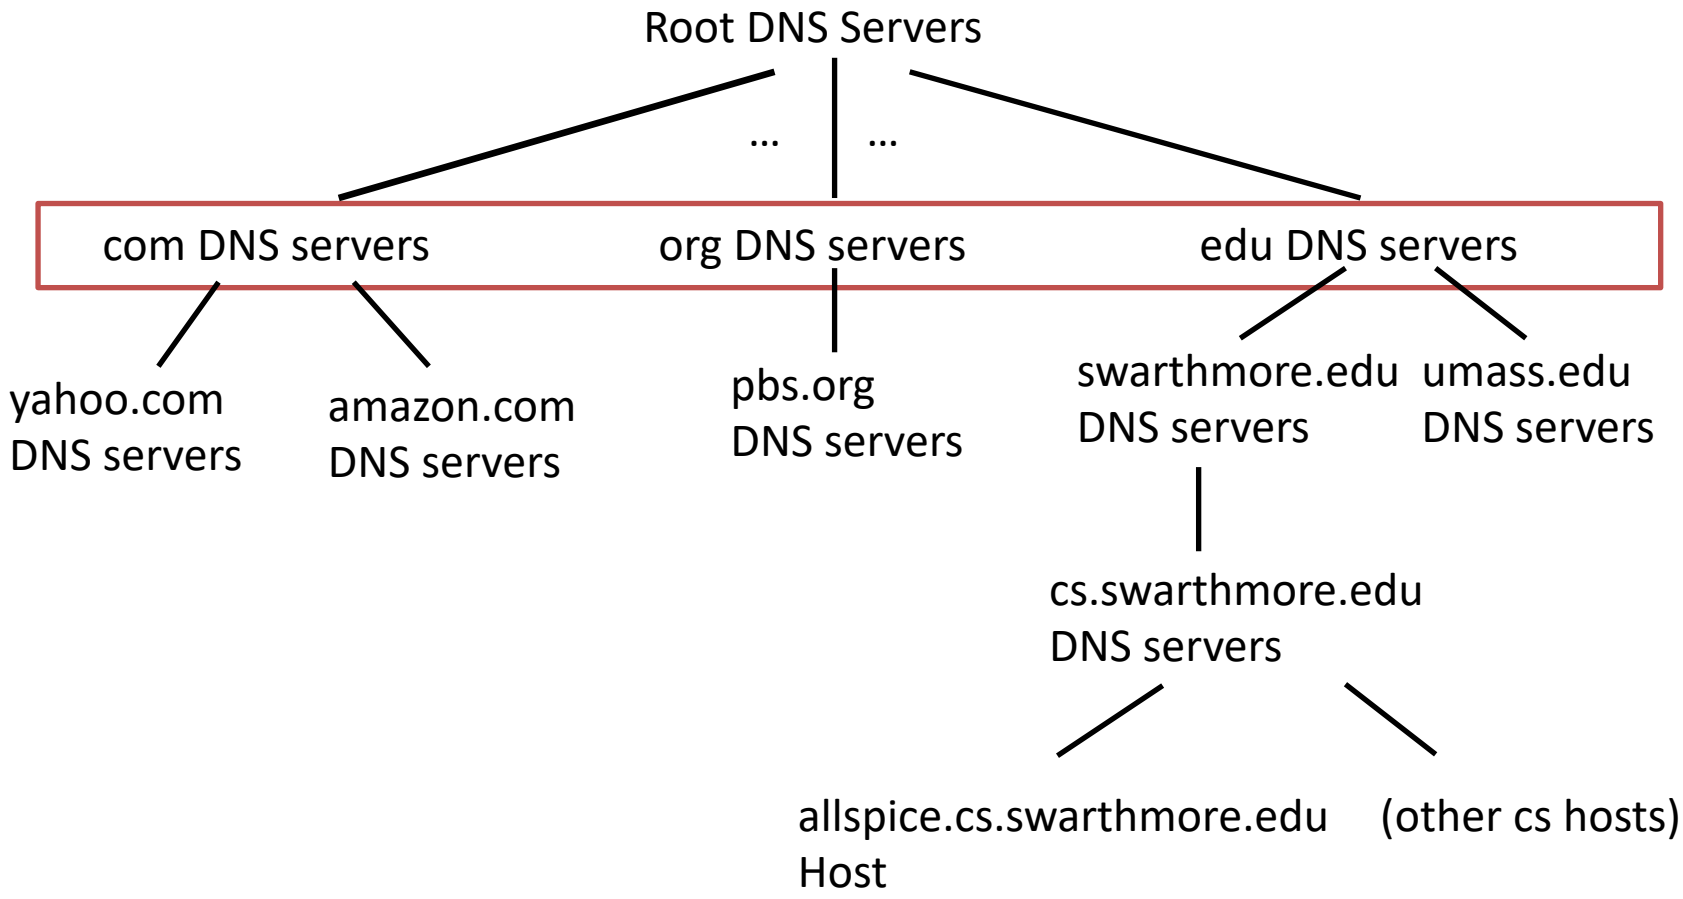

Top Level Domains

Top-level domain (TLD) servers:

- Responsible for com, org, net, edu, gov, aero, jobs, museums, and all top-level country domains, e.g.: uk, fr, de, ca, jp, etc.
- Verisign maintains servers for .com and .net TLD
- Educause for .edu TLD (Verisign actually runs backend)
- Others managed by corresponding entity (e.g., local governments or companies)

DNS: a distributed, hierarchical database

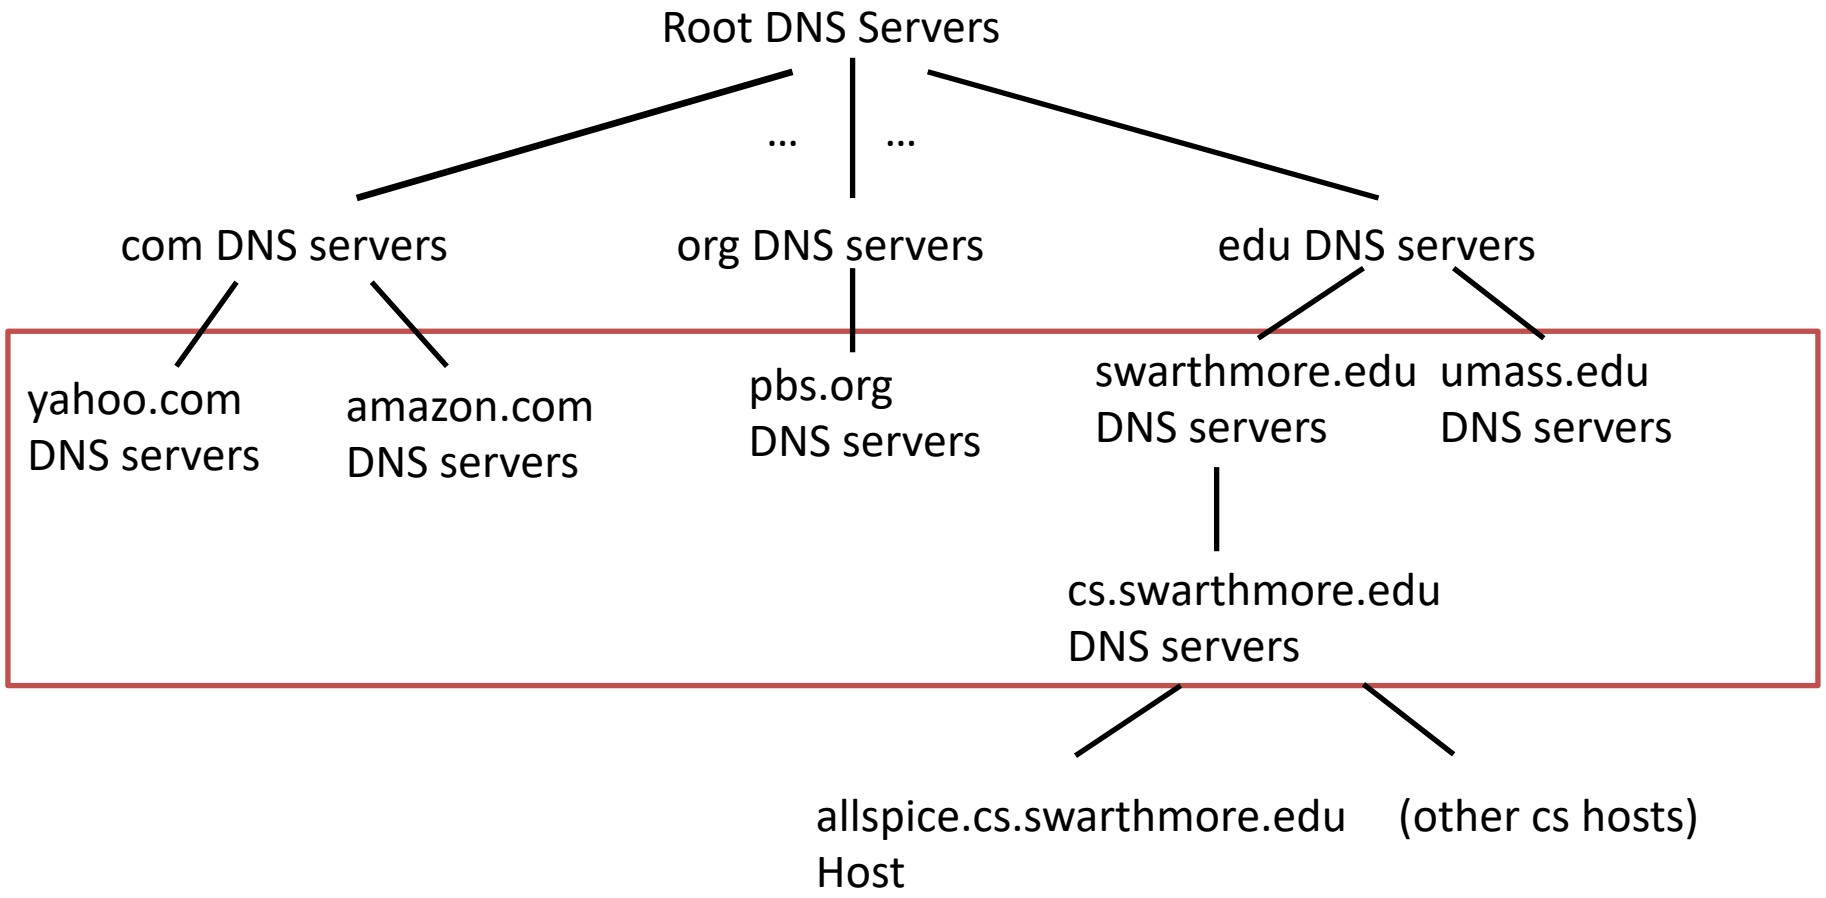

Authoritative Servers

Authoritative DNS servers:

- Organization's own DNS server(s), providing authoritative hostname to IP mappings for organization's named hosts
- Can be maintained by organization or service provider, easily changing entries
- Often, but not always, acts as organization's local name server (for responding to look-ups)

Resolution Process

- End host wants to look up a name, who should it contact?
  - It could traverse the hierarchy, starting at a root
  - More efficient for ISP to provide a local server
- ISP's local server for handling queries not necessarily a part of the pictured hierarchy

Local DNS Name Server

- Each ISP (residential ISP, company, university) has (at least) one
  - also called “default name server”
- When host makes DNS query, query is sent to its local DNS server
  - has local cache of recent name-to-address translation pairs (but may be out of date!)
  - acts as proxy, forwards query into hierarchy

DNS name resolution example #1

- allspice wants IP address for gaia.cs.umass.edu

iterative query:

- contacted server replies with name of server to contact
- "I don't know this name, but ask this server"

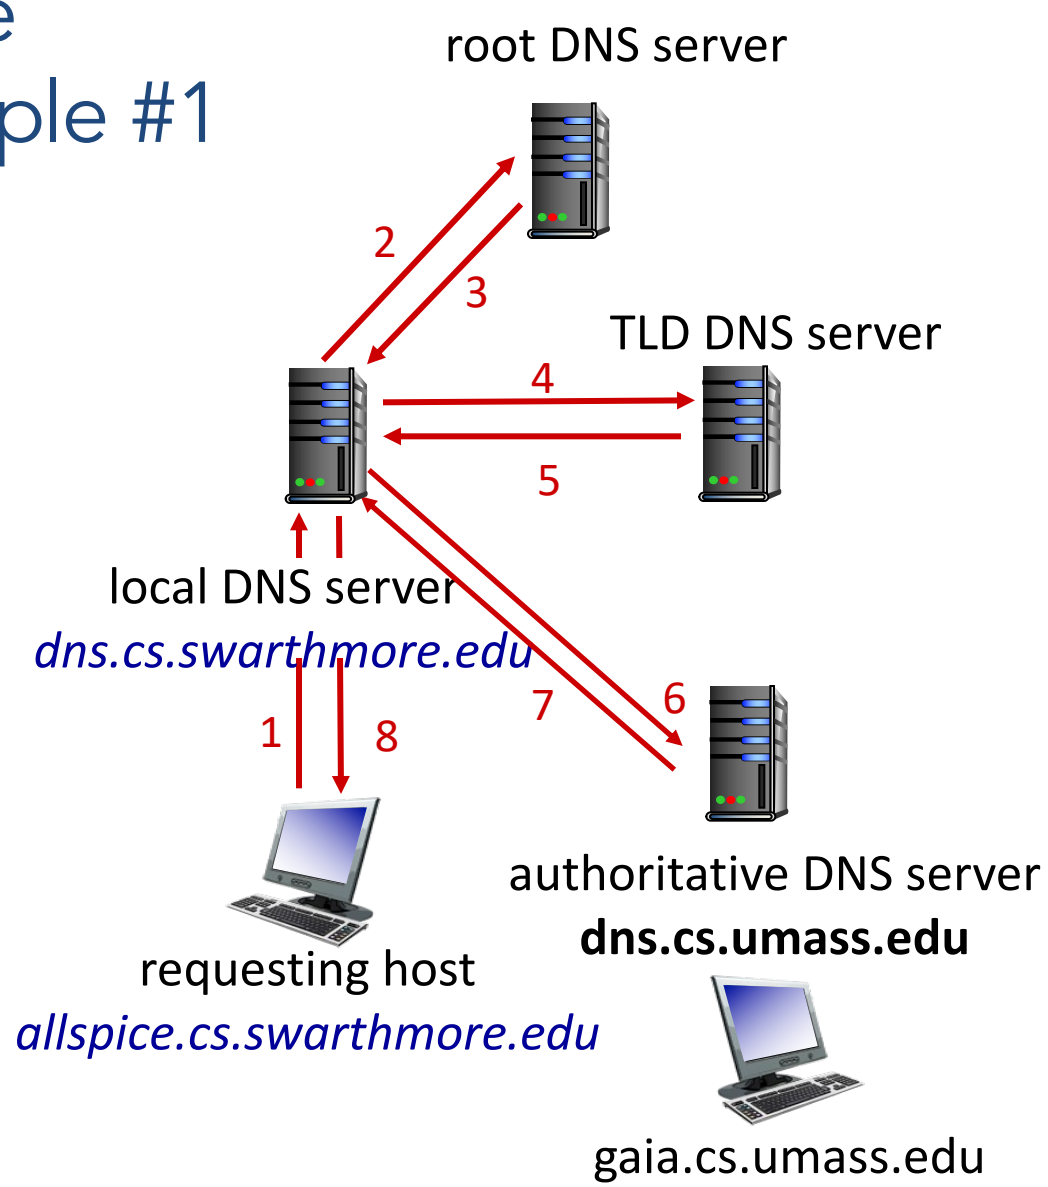

DNS name resolution example #2

recursive query:

- each server asks the next one, in a chain

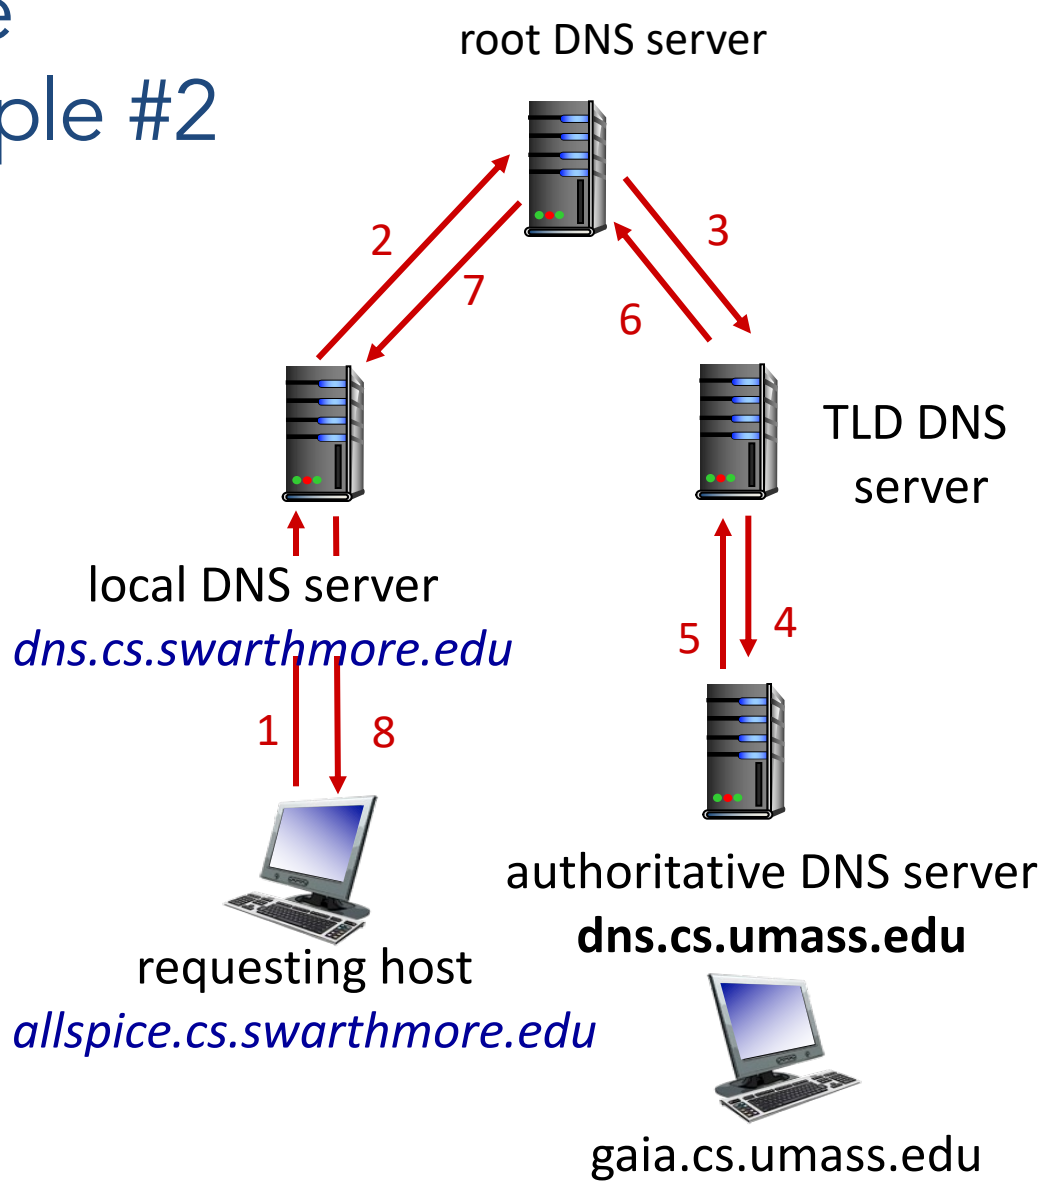

Which would you use? Why?

A. Iterative

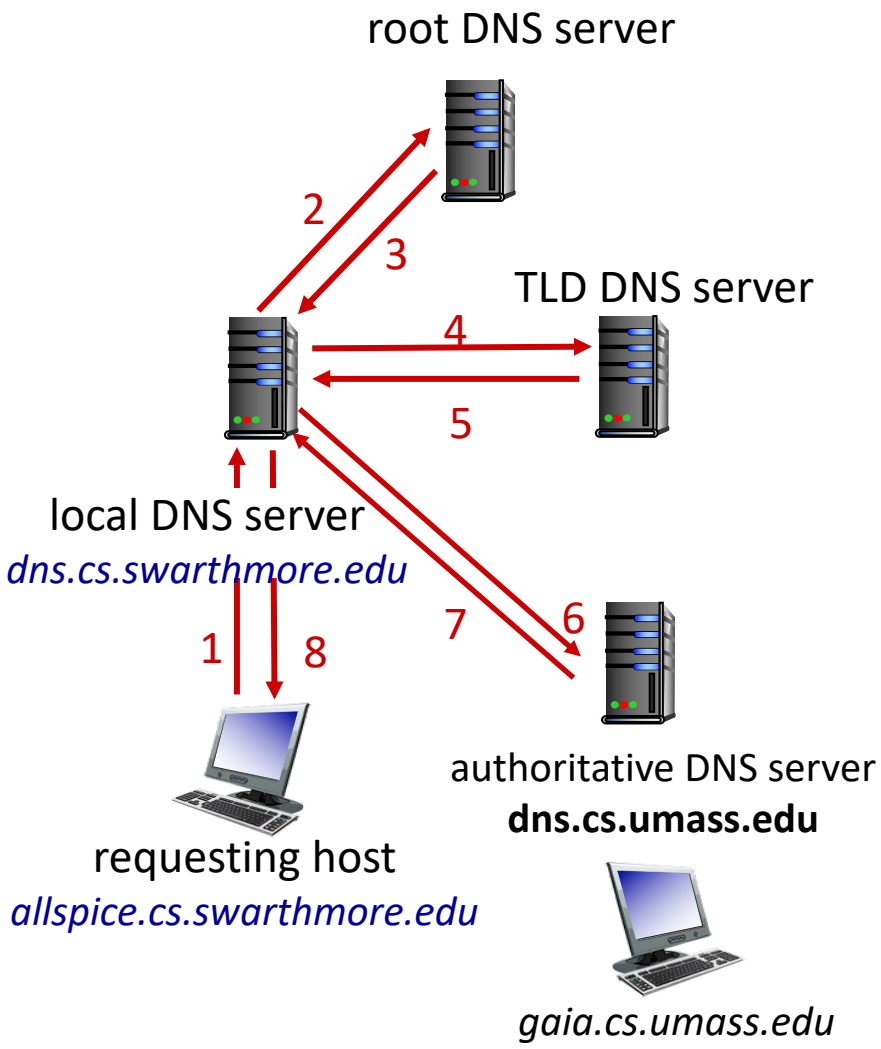

B. Recursive

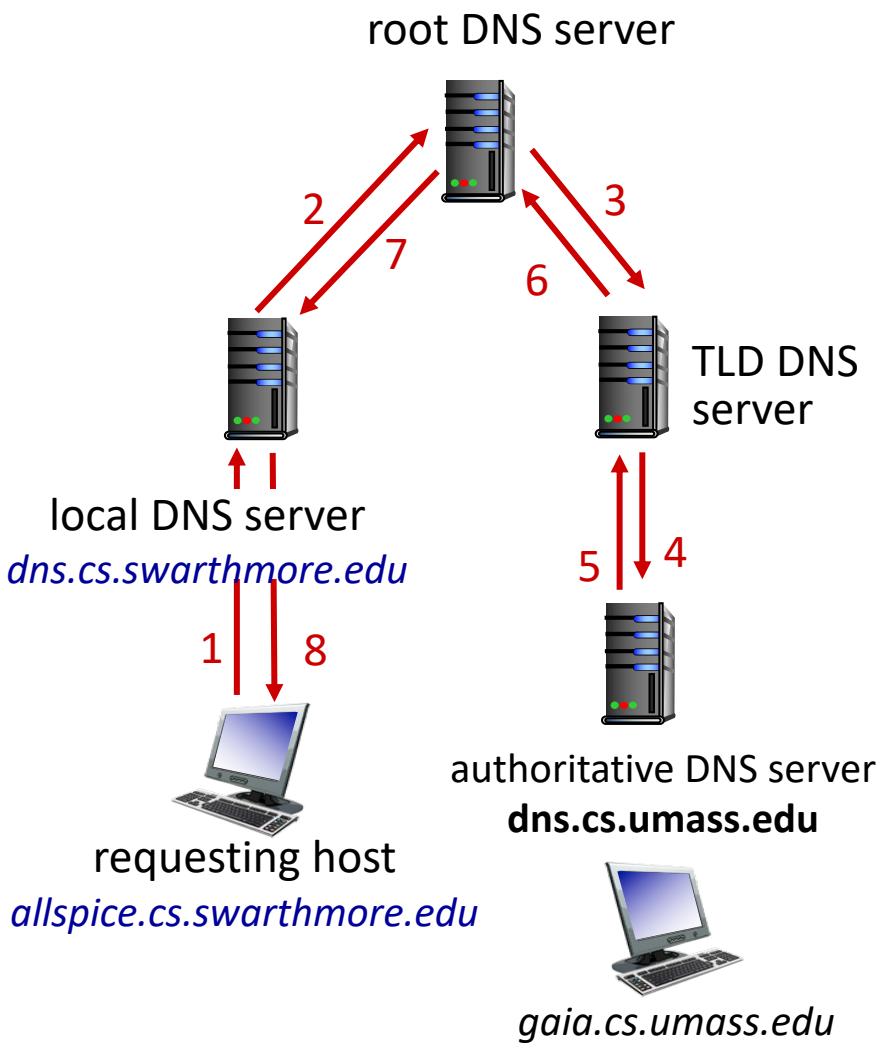

Example: iterative query using dig()

```
dig . ns

dig +norec demo.cs.swarthmore.edu @a.root-servers.net

dig +norec demo.cs.swarthmore.edu @a.edu-servers.net

dig +norec demo.cs.swarthmore.edu @ibext.its.swarthmore.edu

    demo.cs.swarthmore.edu.  259200 IN A  130.58.68.26
```

DNS as Indirection Service

- DNS gives us very powerful capabilities
  - Not only easier for humans to reference machines!

- Changing the IPs of machines becomes trivial
  - e.g. you want to move your web server to a new host
  - Just change the DNS record!

Aliasing and Load Balancing

- One machine can have many aliases

  - One domain can map to multiple machines

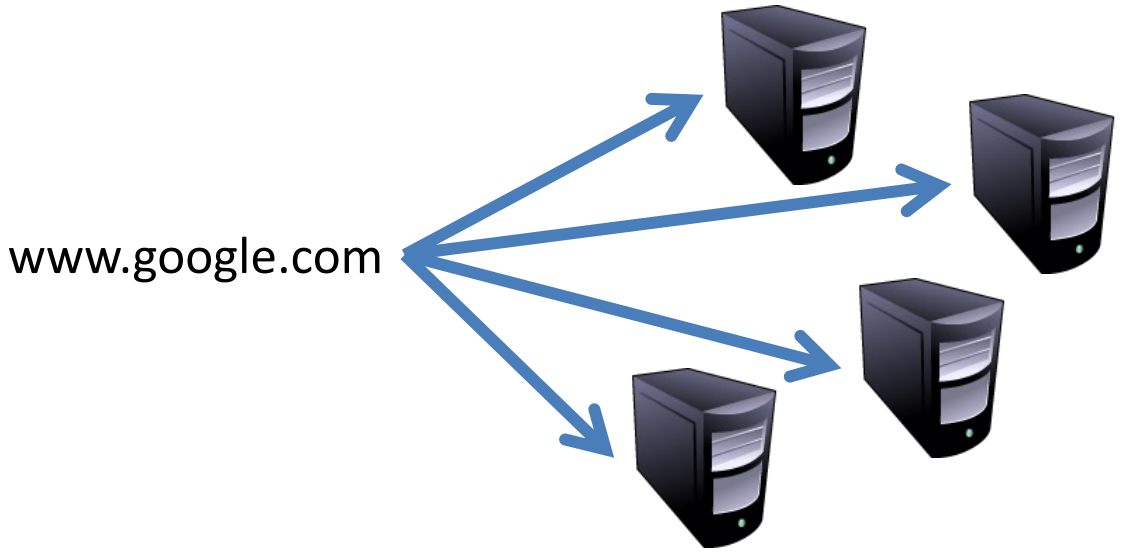

Content Delivery Networks

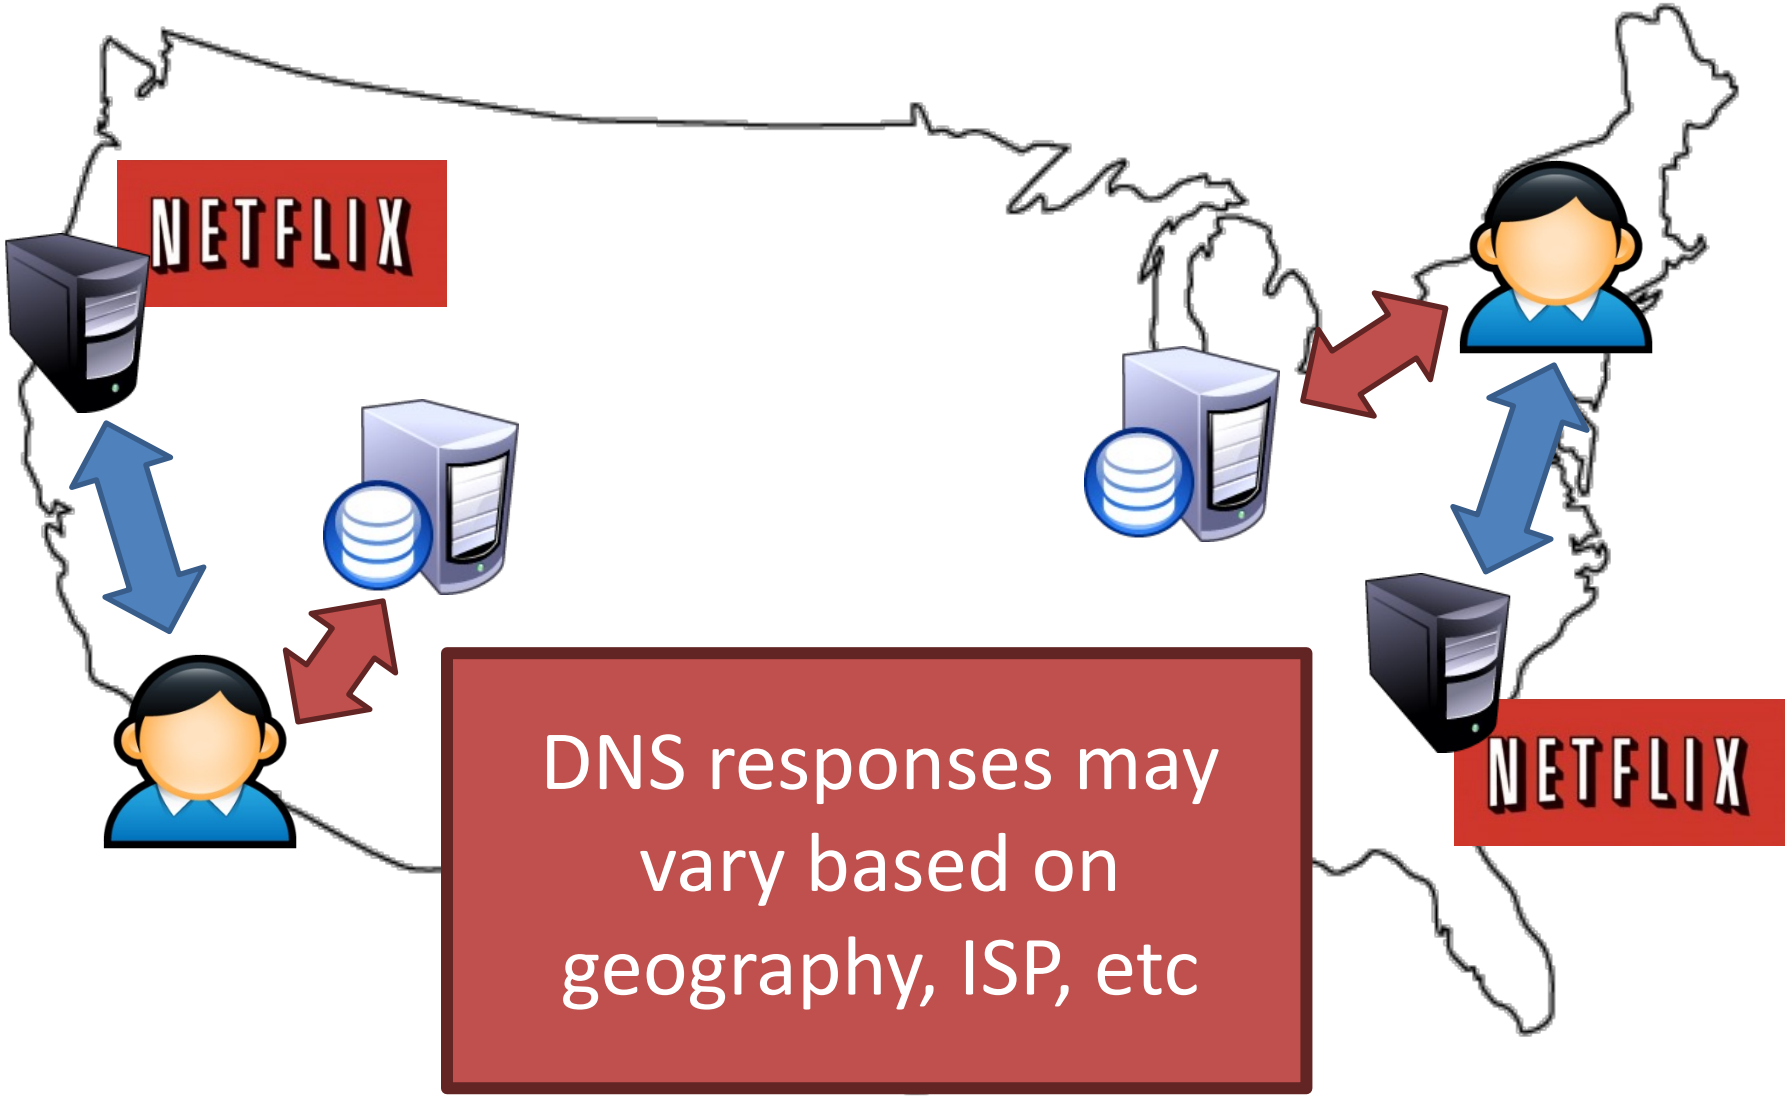

DNS Query Process and Cache

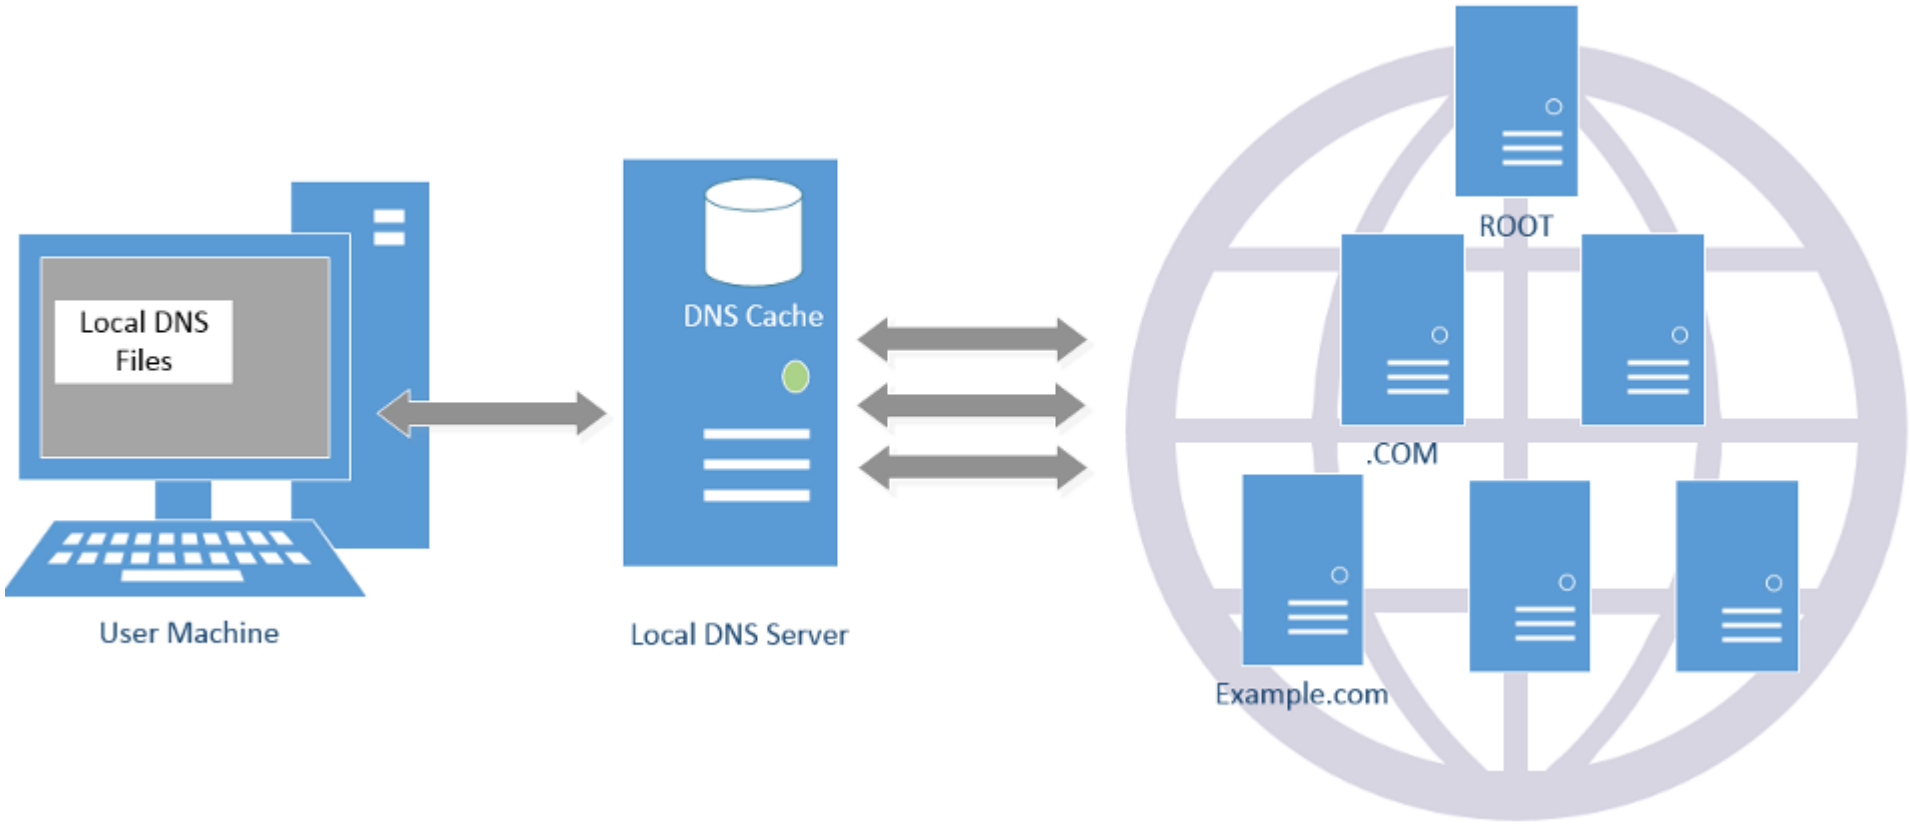

Caching

- Once (any) name server learns a mapping, it caches mapping
  - cache entries timeout (disappear) after some time (TTL: time to live)
  - TLD servers typically cached in local name servers
  - Thus root name servers not often (legitimately) visited

Caching

- Once (any) name server learns a mapping, it caches mapping
  - cache entries timeout (disappear) after some time (TTL: time to live)
  - TLD servers typically cached in local name servers.
  - Root name servers not often (legitimately) visited
- (+) Subsequent requests need not burden DNS
- (-) Cached entries may be out-of-date (best effort!)
  - If host's name or IP address changes, it may not be known Internet-wide until all TTLs expire

The TTL value should be…

A. Short, to make sure that changes are accurately reflected

B. Long, to avoid re-queries of higher-level DNS servers

C. Something else

Inserting (or changing) records

- Example: new startup “Network Utopia”
- Register networkuptopia.com at DNS registrar
  - provide names, IP addresses of authoritative name server (primary and secondary)
  - registrar inserts two RRs into .com TLD server
    - (networkutopia.com, dns1.networkutopia.com, NS)
    - (dns1.networkutopia.com, 212.212.212.1, A)
- Set up authoritative server at that name/address
  - Create records for the services:
    - type A record for www.networkuptopia.com
    - type MX record for @networkutopia.com email

Attacking DNS

DDoS attacks

- Bombard root servers with traffic
  - Not successful to date
  - Traffic Filtering
  - Local DNS servers cache IPs of TLD servers, bypassing root
- Bombard TLD servers
  - Potentially more dangerous

Redirect attacks

- Man-in-middle
  - Intercept queries
- DNS poisoning
  - Send bogus replies to DNS server that caches

Exploit DNS for DDoS

- Send queries with spoofed source address: target IP
- Requires amplification

Tools

- dig
  - $ dig cs.swarthmore.edu
  - $ dig cs.swarthmore.edu ns
  - $ dig @dns.cs.swarthmore.edu cs.swarthmore.edu mx
  - $ man dig
- host
  - $ host cs.swarthmore.edu
  - $ host -t ns cs.swarthmore.edu
  - $ host -t mx cs.swarthmore.edu dns.cs.swarthmore.edu
  - $ man host

Tools (cont)

- nslookup
  - $ nslookup cs.swarthmore.edu
  - $ nslookup cs.swarthmore.edu dns.cs.swarthmore.edu
- whois
  - $ whois google.com
  - $ whois swarthmore.edu

How many answers
Time to live in seconds
How many additional records?

```
$ dig @a.root-servers.net www.freebsd.org +norecurse
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 57494
;; QUERY: 1, ANSWER: 0, AUTHORITY: 2, ADDITIONAL: 2

;; QUESTION SECTION:
;www.freebsd.org.          IN     A

;; AUTHORITY SECTION:
org.                172800 IN     NS     b0.org.afilias-nst.org.
org.                172800 IN     NS     d0.org.afilias-nst.org.

;; ADDITIONAL SECTION:
b0.org.afilias-nst.org.   172800 IN     A      199.19.54.1
d0.org.afilias-nst.org.   172800 IN     A      199.19.57.1
```

How many answers
Time to live in seconds
How many additional records?

```
$ dig @a.root-servers.net www.freebsd.org +norecurse
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 57494
;; QUERY: 1, ANSWER: 0, AUTHORITY: 2, ADDITIONAL: 2

;; QUESTION SECTION:
;www.freebsd.org.               IN      A

;; AUTHORITY SECTION:
org.                    172800 IN      NS      b0.org.afilias-nst.org.
org.                    172800 IN      NS      d0.org.afilias-nst.org.

;; ADDITIONAL SECTION:
b0.org.afilias-nst.org.    172800 IN      A       199.19.54.1
d0.org.afilias-nst.org.    172800 IN      A       199.19.57.1
```

Glue records

How many answers?
How many additional records?

(authoritative for org.)

```
$ dig @199.19.54.1 www.freebsd.org +norecurse
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 39912
;; QUERY: 1, ANSWER: 0, AUTHORITY: 3, ADDITIONAL: 0

;; QUESTION SECTION:
;www.freebsd.org.               IN      A

;; AUTHORITY SECTION:
freebsd.org.            86400   IN      NS      ns1.isc-sns.net.
freebsd.org.            86400   IN      NS      ns2.isc-sns.com.
freebsd.org.            86400   IN      NS      ns3.isc-sns.info.
```

How many answers?
How many additional records?

(authoritative for org.)

```
$ dig @199.19.54.1 www.freebsd.org +norecurse
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 39912
;; QUERY: 1, ANSWER: 0, AUTHORITY: 3, ADDITIONAL: 0

;; QUESTION SECTION:
;www.freebsd.org.               IN      A

;; AUTHORITY SECTION:
freebsd.org.            86400   IN      NS      ns1.isc-sns.net.
freebsd.org.            86400   IN      NS      ns2.isc-sns.com.
freebsd.org.            86400   IN      NS      ns3.isc-sns.info.
```

(authoritative for freebsd.org.)

```
$ dig @ns1.isc-sns.net www.freebsd.org +norecurse
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 17037
;; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 3

;; QUESTION SECTION:
;www.freebsd.org.               IN      A

;; ANSWER SECTION:
www.freebsd.org.        3600    IN      A       69.147.83.33

;; AUTHORITY SECTION:
freebsd.org.            3600    IN      NS      ns2.isc-sns.com.
freebsd.org.            3600    IN      NS      ns1.isc-sns.net.
freebsd.org.            3600    IN      NS      ns3.isc-sns.info.

;; ADDITIONAL SECTION:
ns1.isc-sns.net.        3600    IN      A       72.52.71.1
ns2.isc-sns.com.        3600    IN      A       38.103.2.1
ns3.isc-sns.info.       3600    IN      A       63.243.194.1
```

How many answers?
How many authoritative records?
How many additional records?

(authoritative for freebsd.org.)

```
$ dig @ns1.isc-sns.net www.freebsd.org +norecurse
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 17037
;; QUERY: 1, ANSWER: 1, AUTHORITY: 3, ADDITIONAL: 3

;; QUESTION SECTION:
;www.freebsd.org.          IN     A

;; ANSWER SECTION:
www.freebsd.org.    3600   IN     A       69.147.83.33

;; AUTHORITY SECTION:
freebsd.org.        3600   IN     NS      ns2.isc-sns.com.
freebsd.org.        3600   IN     NS      ns1.isc-sns.net.
freebsd.org.        3600   IN     NS      ns3.isc-sns.info.

;; ADDITIONAL SECTION:
ns1.isc-sns.net.    3600   IN     A       72.52.71.1
ns2.isc-sns.com.    3600   IN     A       38.103.2.1
ns3.isc-sns.info.   3600   IN     A       63.243.194.1
```

How many answers?
How many authoritative records?
How many additional records?

DNS security

DNS Vulnerabilities:

- No authentication
- Connectionless transport layer protocol (UDP)

DNS Attacks:

- Amplification Attack
- Cache Poisoning
- Man-in-the-middle
- DNS Redirection
- DDoS
- DNS Injection

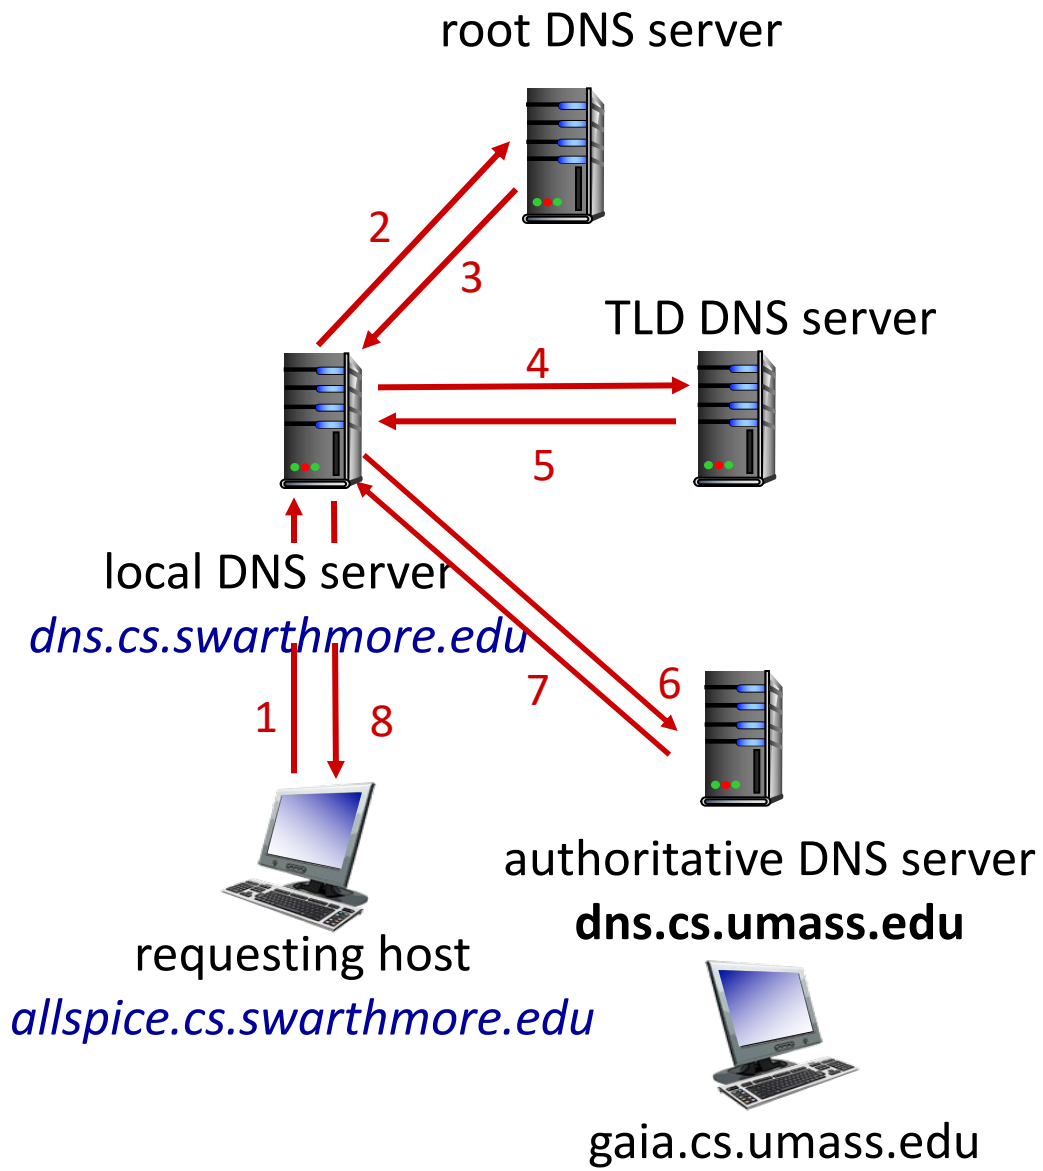

Attacking DNS

DDoS attacks

- Bombard root servers with traffic
  - Not successful to date
  - Traffic Filtering
  - Local DNS servers cache IPs of TLD servers, bypassing root
- Bombard TLD servers
  - Potentially more dangerous

Redirect attacks

- Man-in-middle
  - Intercept queries
- DNS poisoning
  - Send bogus replies to DNS server that caches

Exploit DNS for DDoS

- Send queries with spoofed source address: target IP
- Requires amplification

DNS Query Process and Cache

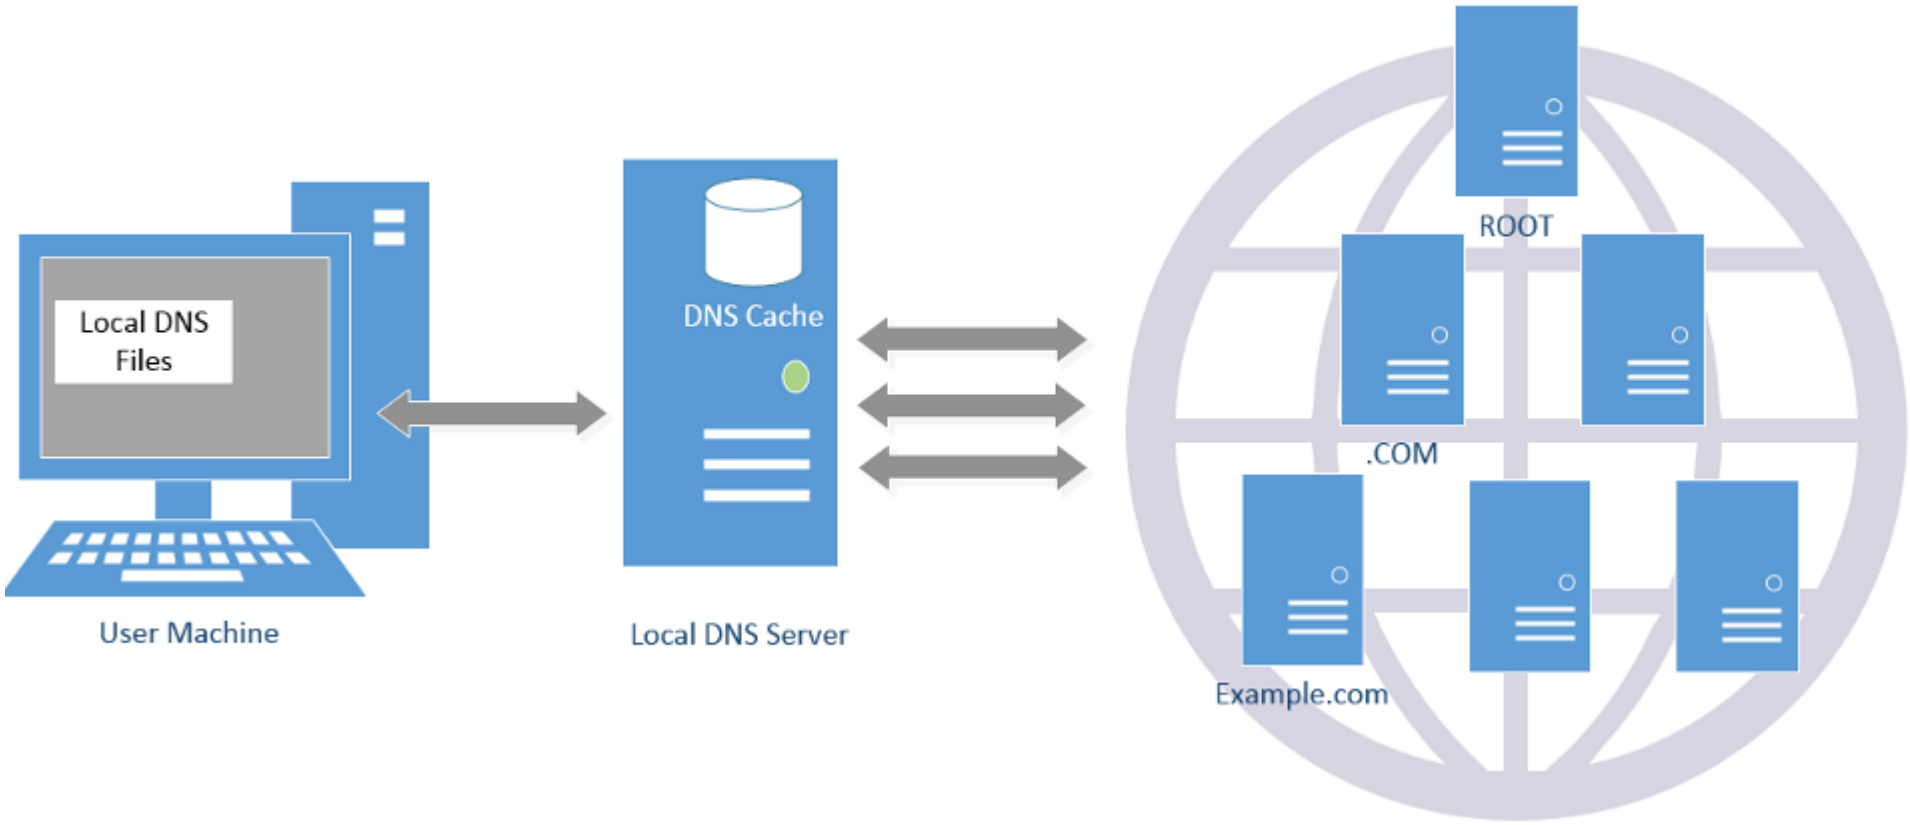

Attack Surface Overview

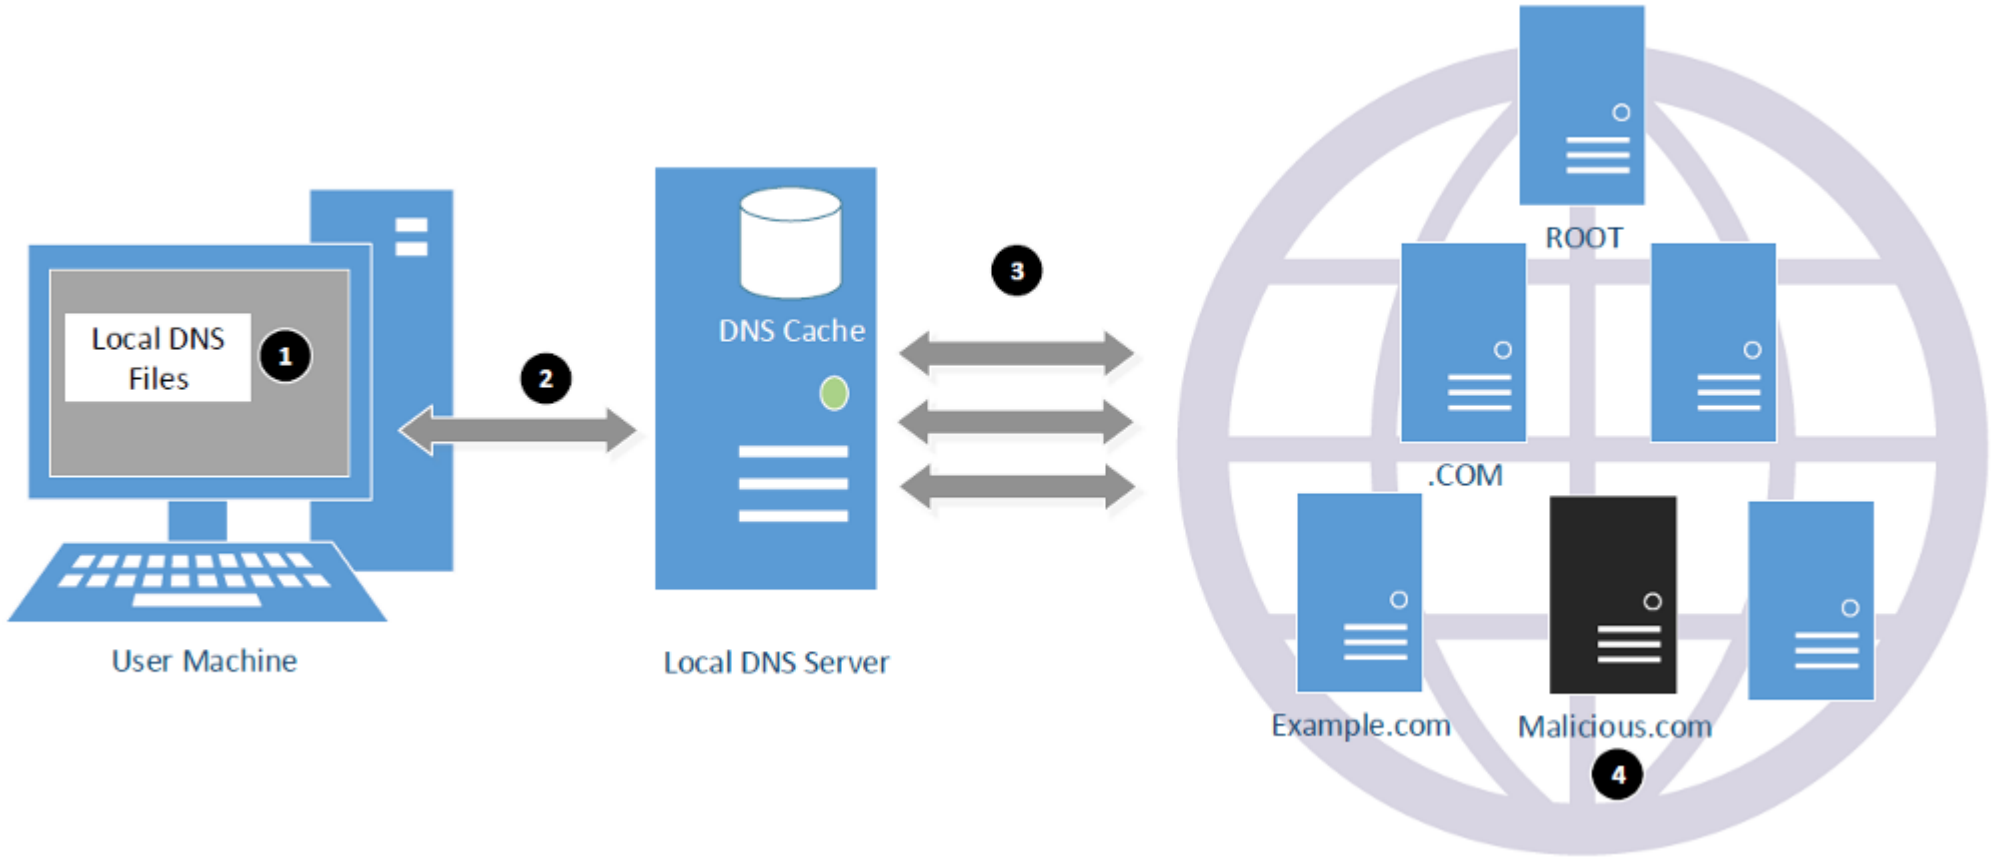

Denial Of Service

- Flood DNS servers with requests until they fail
- October 2002: massive DDoS against the root name servers
  - What was the effect?
  - … users didn't even notice
  - Root zone file is cached almost everywhere
- More targeted attacks can be effective
  - Local DNS server → cannot access DNS
  - Authoritative server → cannot access domain

DNS Hijacking

- Infect their OS or browser with a virus/trojan
  - e.g. Many trojans change entries in /etc/hosts
  - *.bankofamerica.com → evilbank.com
- Man-in-the-middle

  - Response Spoofing
    - Eavesdrop on requests
    - Outrace the servers response

DNS S

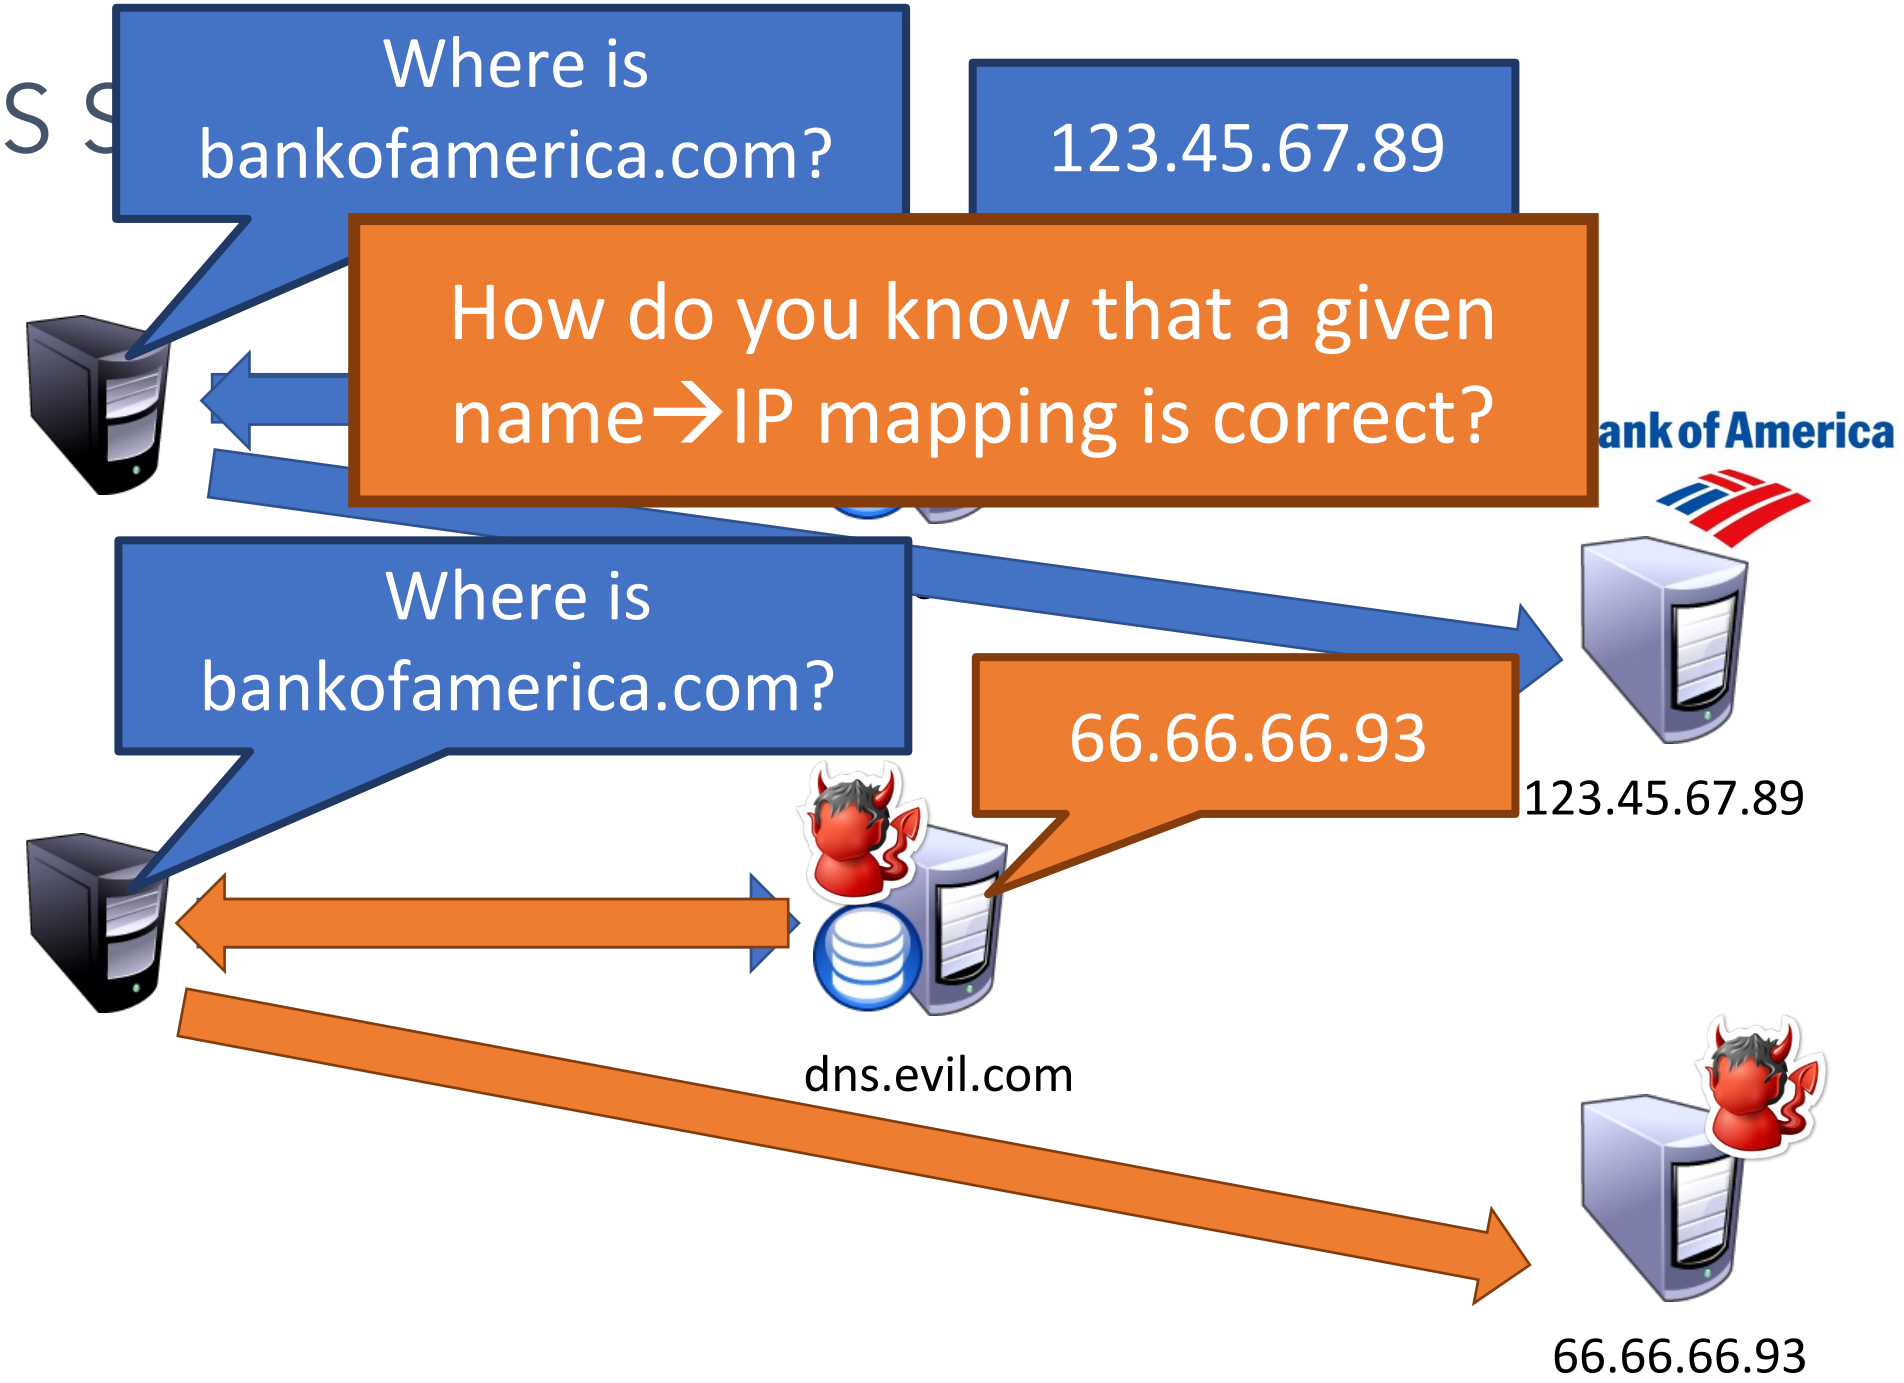
Where is bankofamerica.com?
123.45.67.89
How do you know that a given name→IP mapping is correct?
ank of America
Where is bankofamerica.com?
66.66.66.93
123.45.67.89
dns.evil.com
66.66.66.93

DNS C

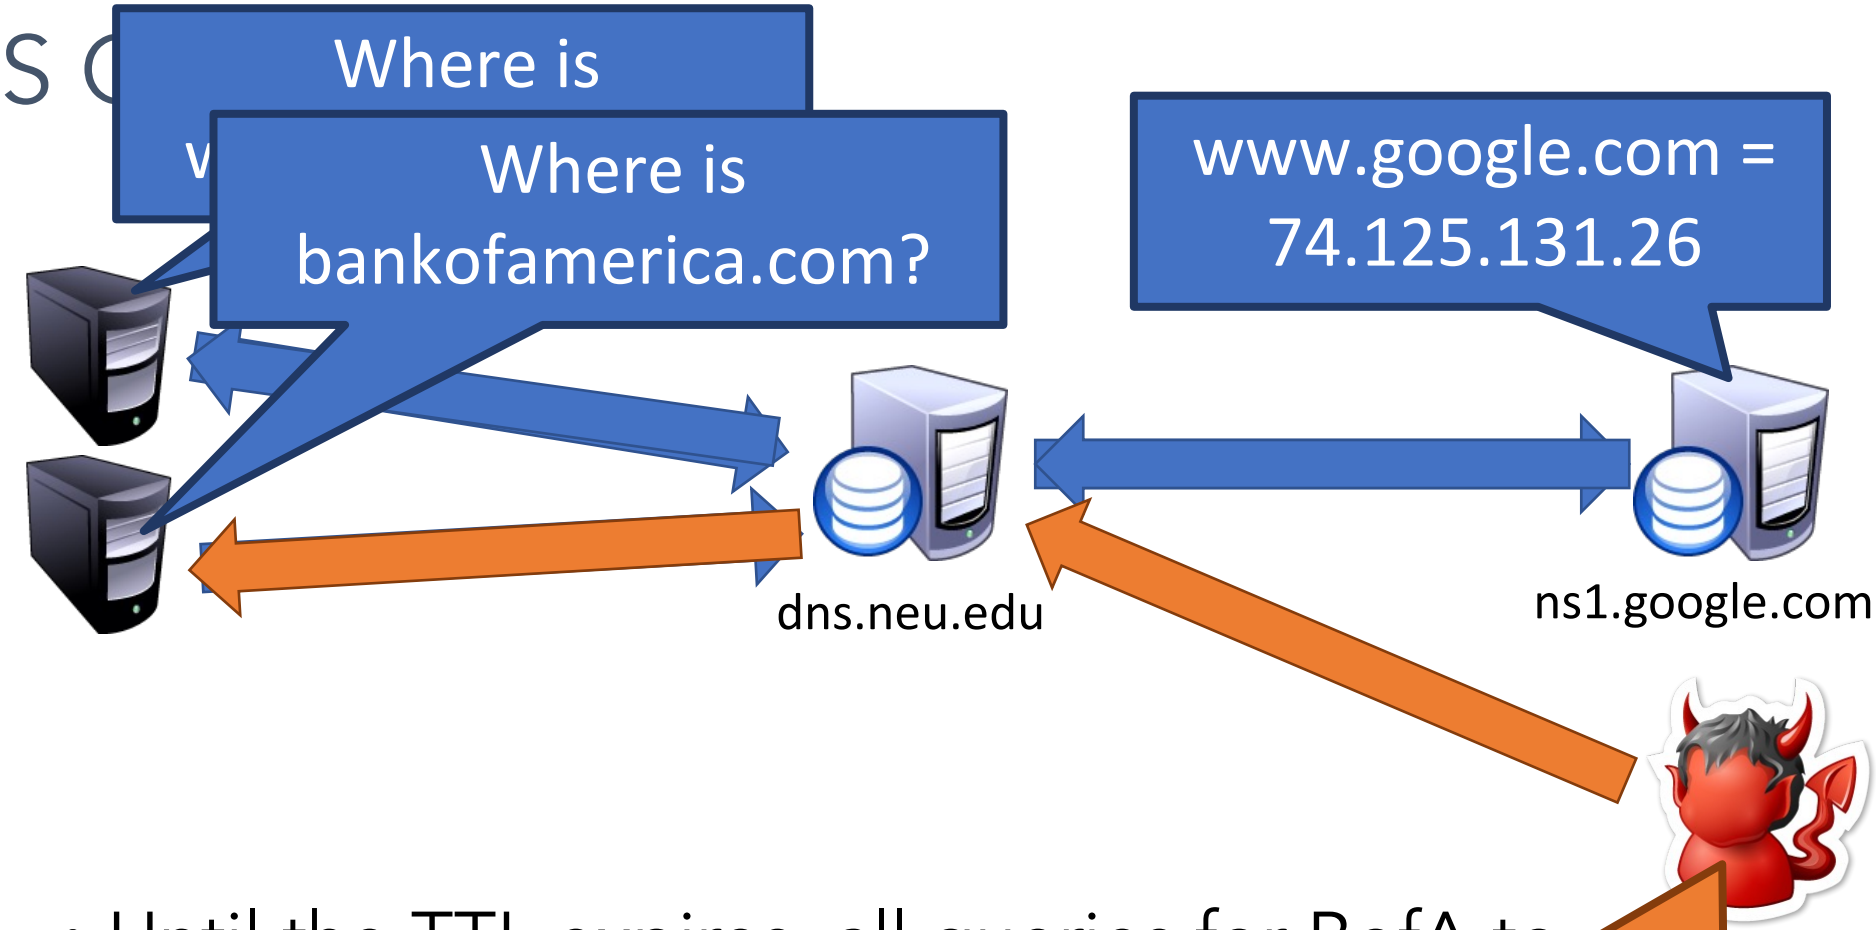

- Until the TTL expires, all queries for BofA to dns.neu.edu will return poisone
- Much worse than spoofing/mar
  - Whole ISPs can be impacted!

DNSSEC Hierarchy of Trust

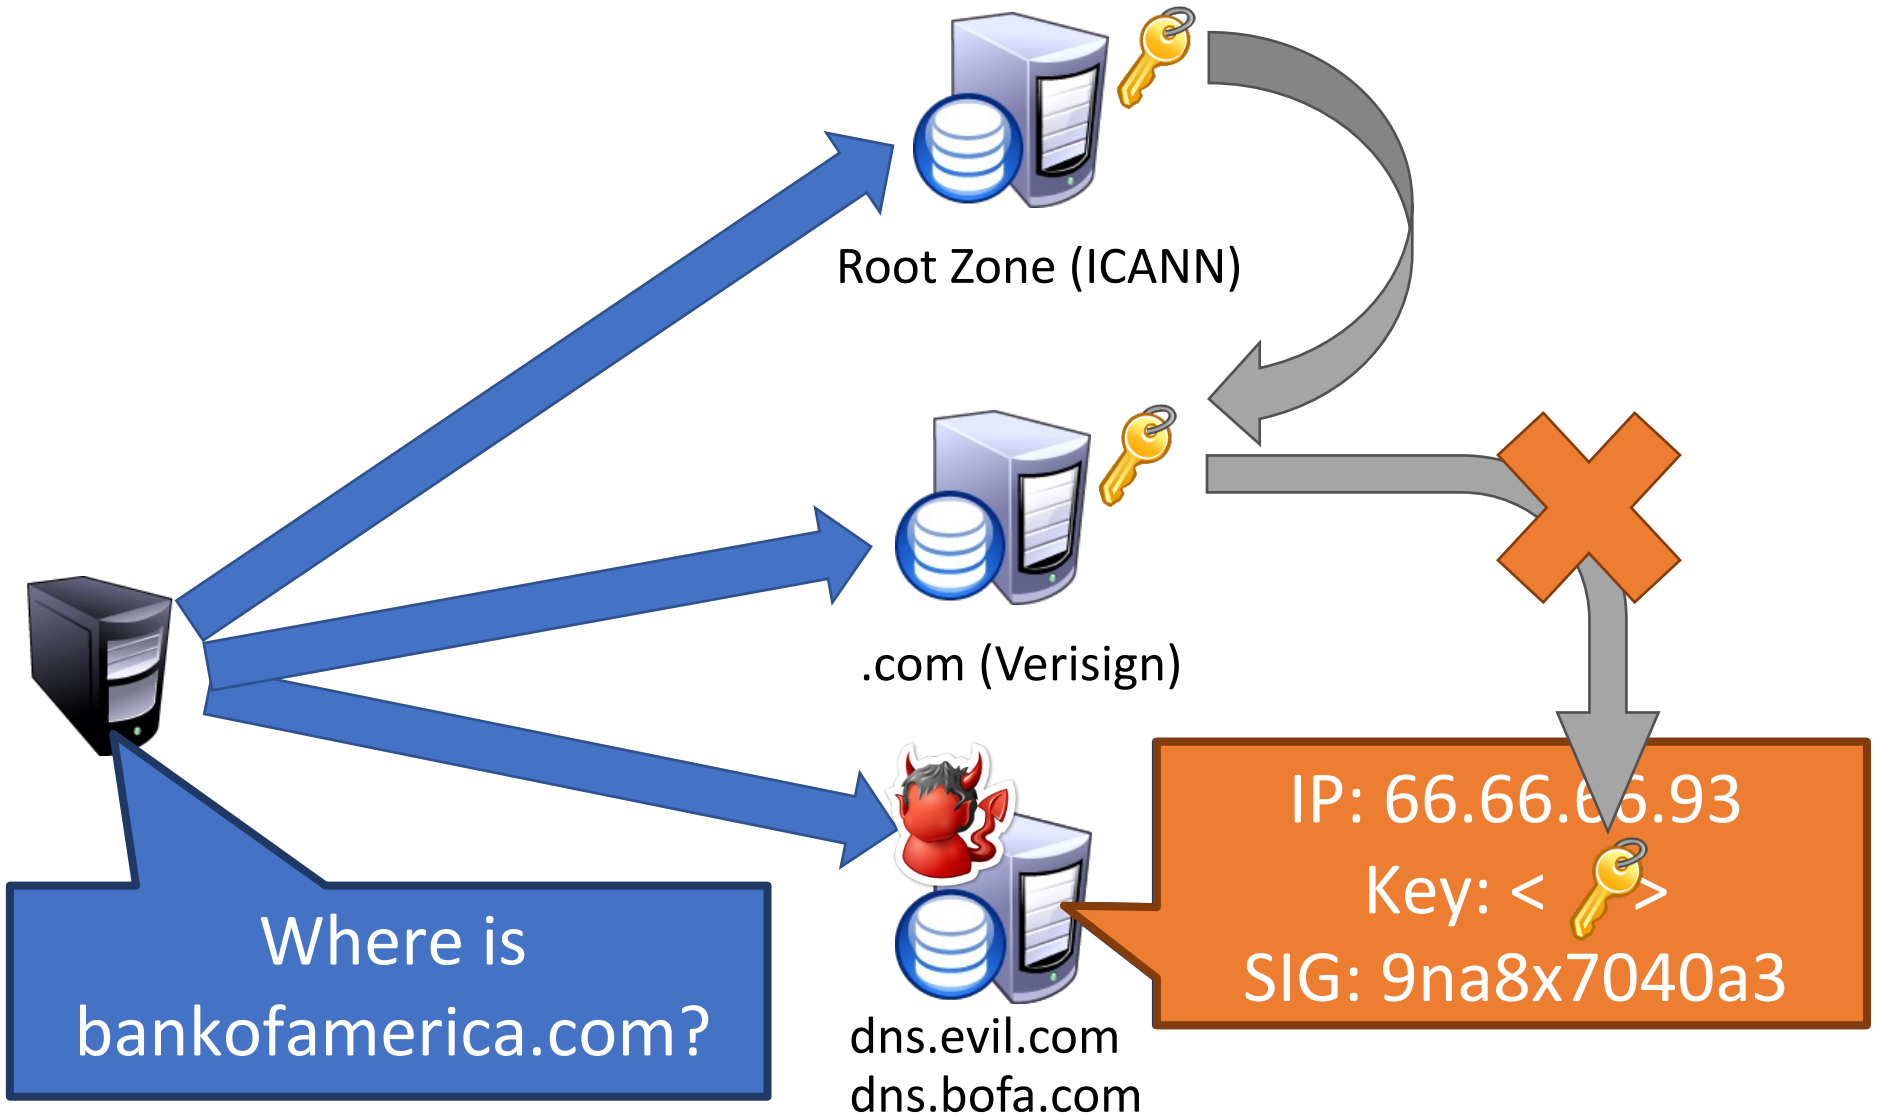

Solution: DNSSEC

- Cryptographically sign critical resource records
  - Resolver can verify the cryptographic signature
- Two new resource types
  - Type = DNSKEY
    - Name = Zone domain name
    - Value = Public key for the zone
  - Type = RRSIG
    - Name = (type, name) tuple, i.e. the query itself
    - Value = Cryptographic signature of the query results

Creates a hierarchy of trust within each zone

Prevents hijacking and spoofing

Summary

- DNS maps human readable names to IP addresses
- DNS arranged into a hierarchy
  - Scalability / distributed responsibility
  - Autonomous control of local name servers
- Caching crucial for performance