CS 43: Computer Networks

Naming and DNS

September 24, 2025

Slides adapted from Kurose & Ross, Vasanta Chaganti, Kevin Webb

Where we are

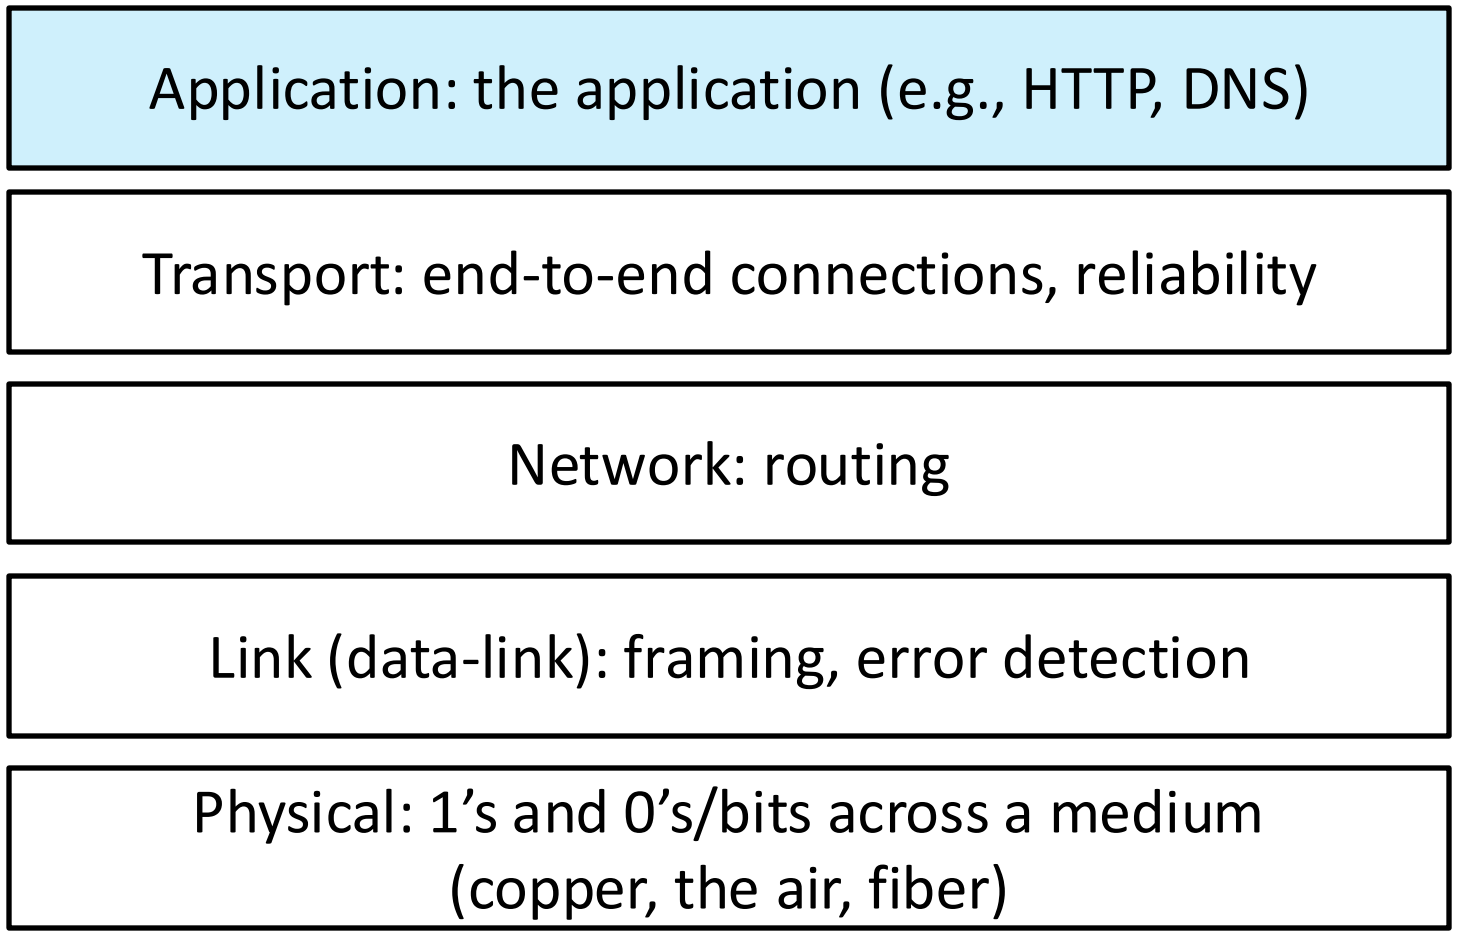
Application: the application (e.g., HTTP, DNS)
Transport: end-to-end connections, reliability
Network: routing
Link (data-link): framing, error detection
Physical: 1's and 0's/bits across a medium
(copper, the air, fiber)

Today

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

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/

Who should control DNS?

A. US government

B. UN / International government

C. Private corporation

D. Someone else

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

DNS Root Server and Politics

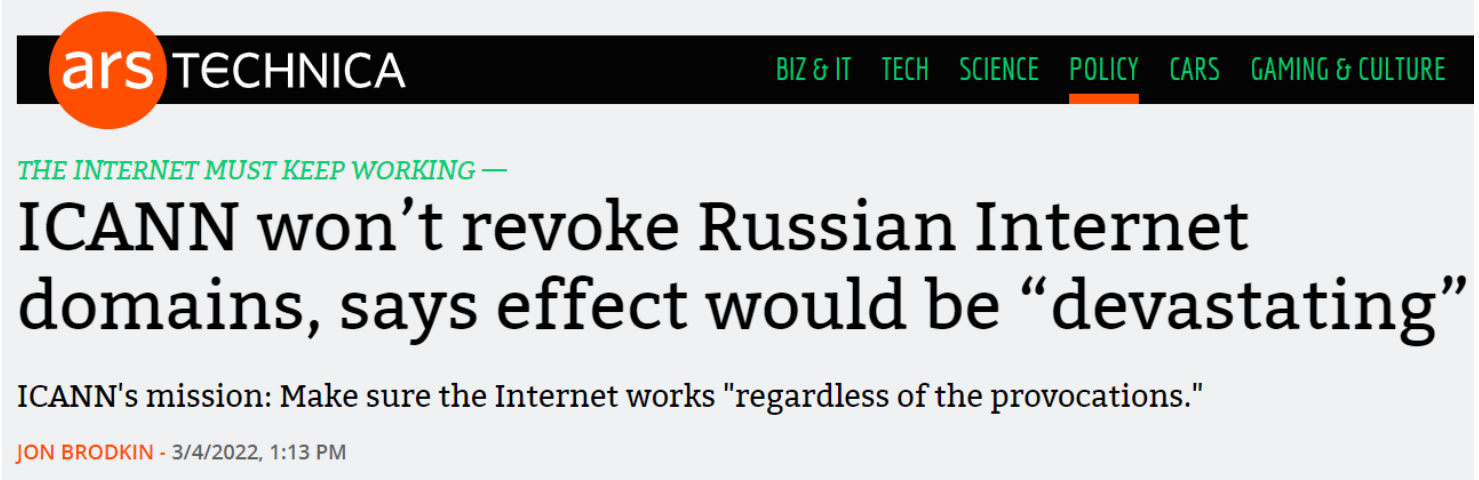

Uses of DNS

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

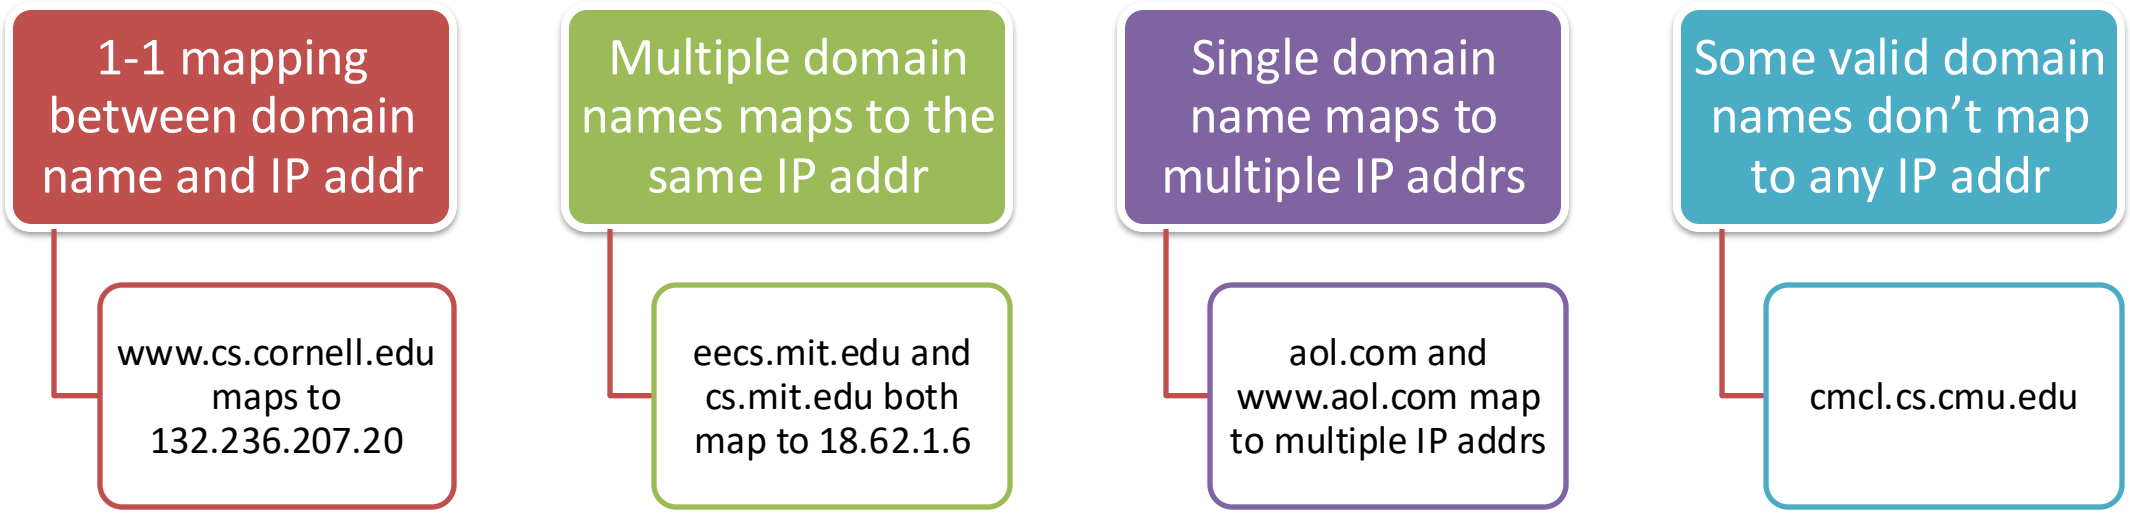

Domain Name System (DNS)

- 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

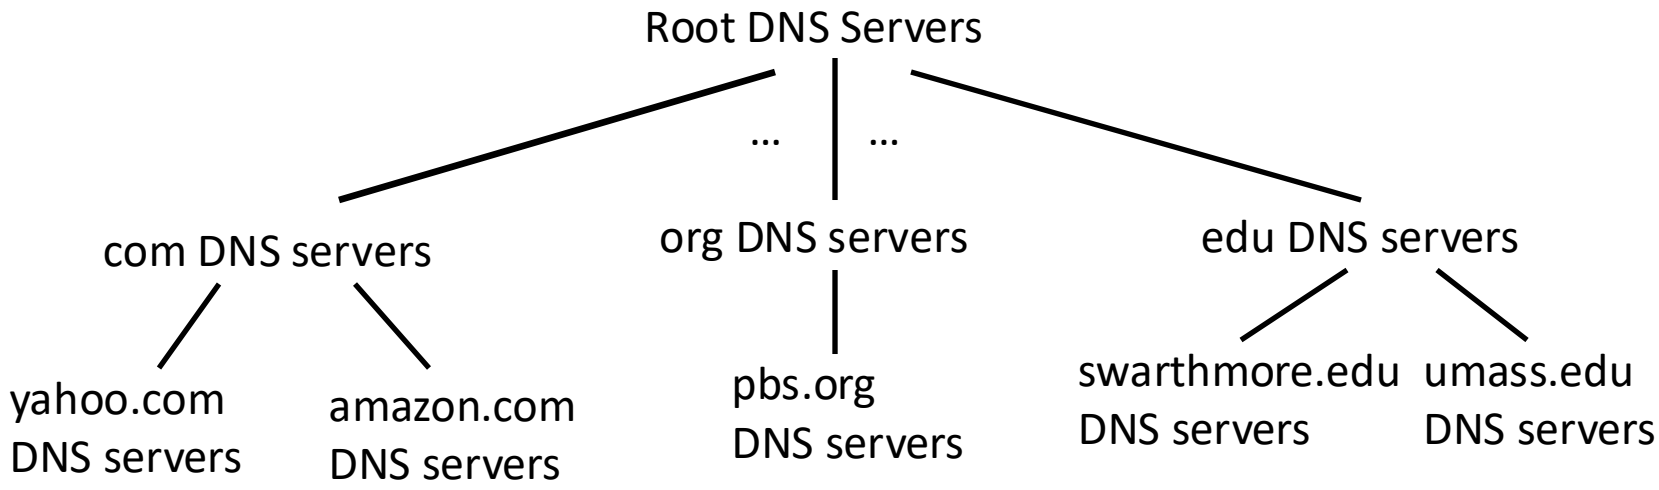

DNS: a distributed, hierarchical database

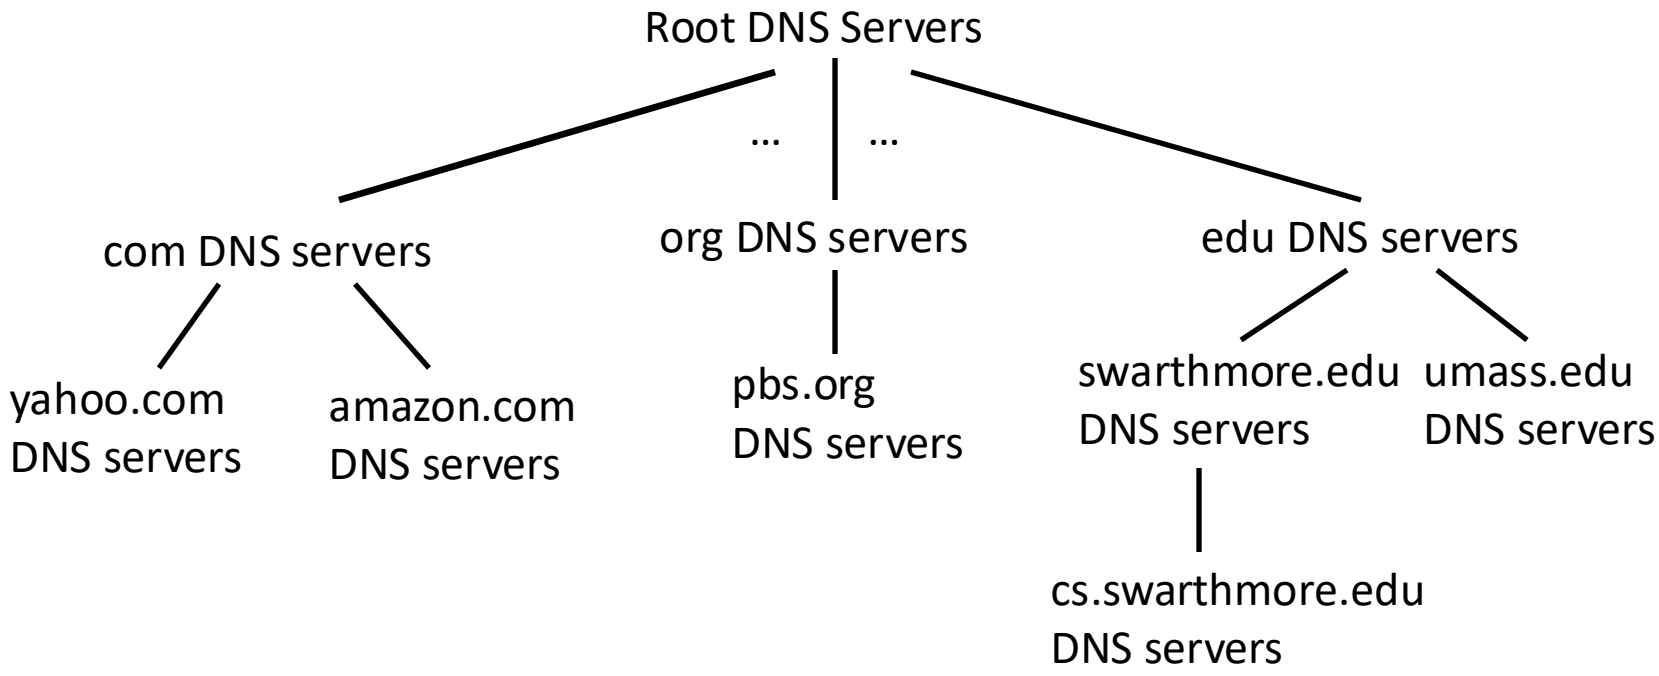

DNS: a distributed, hierarchical database

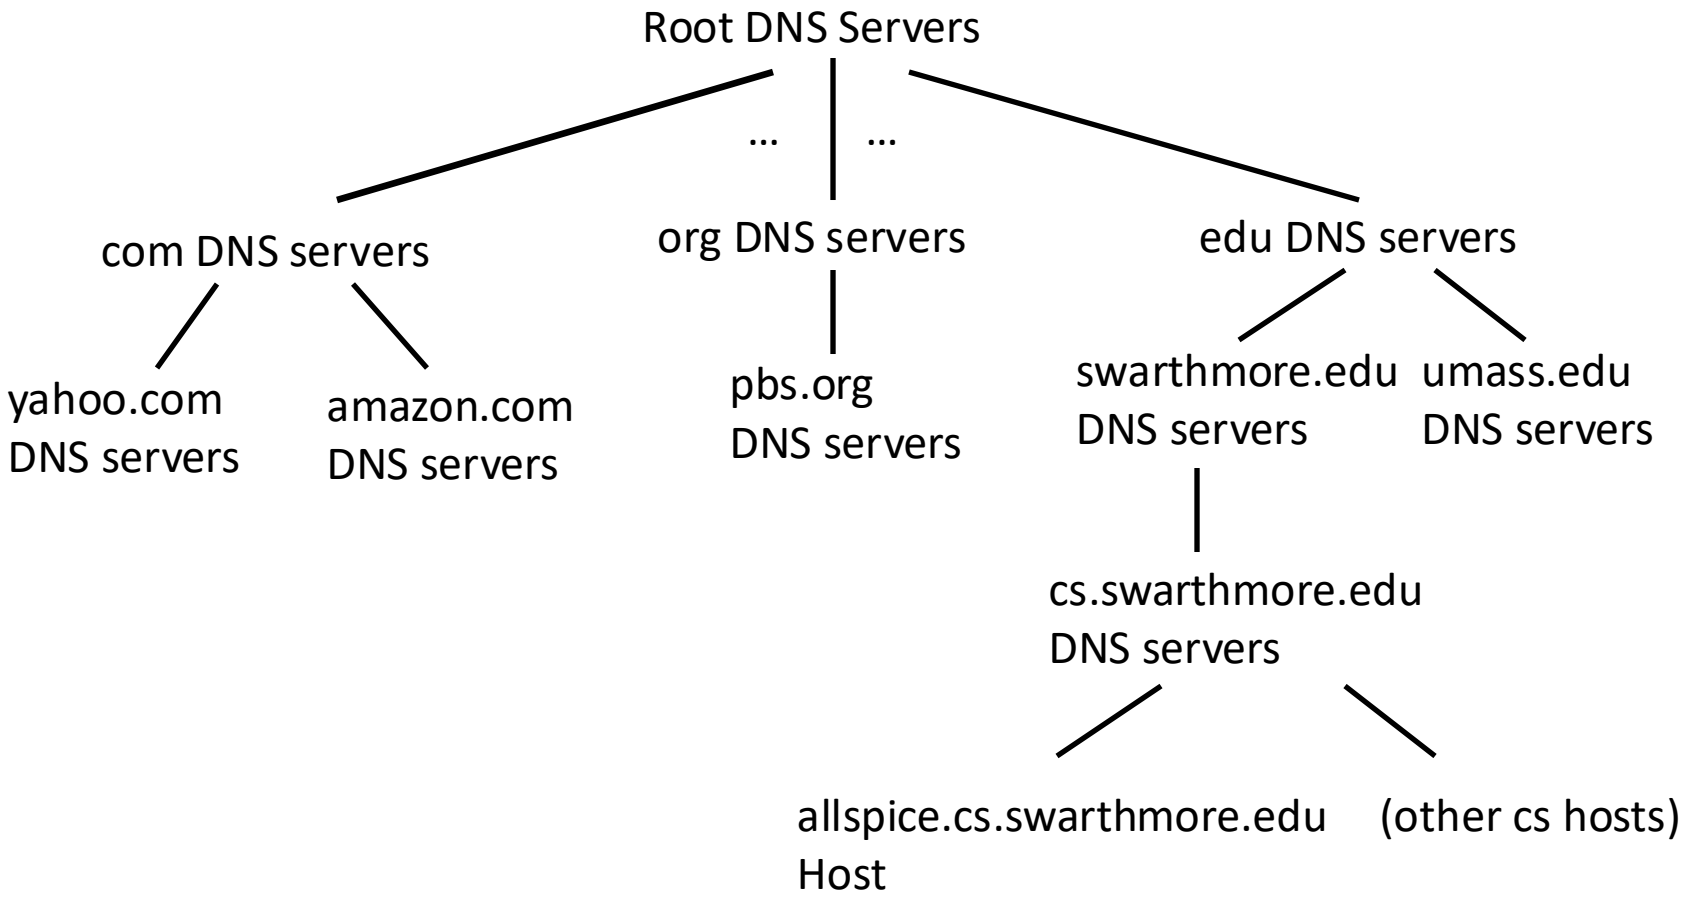

- allspice.cs.swarthmore.edu.

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.

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 overall but individual end hosts (assuming no caching) have a look-up overhead of traversing the hierarchy .

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

D. Some other reason.

DNS: a distributed, hierarchical database

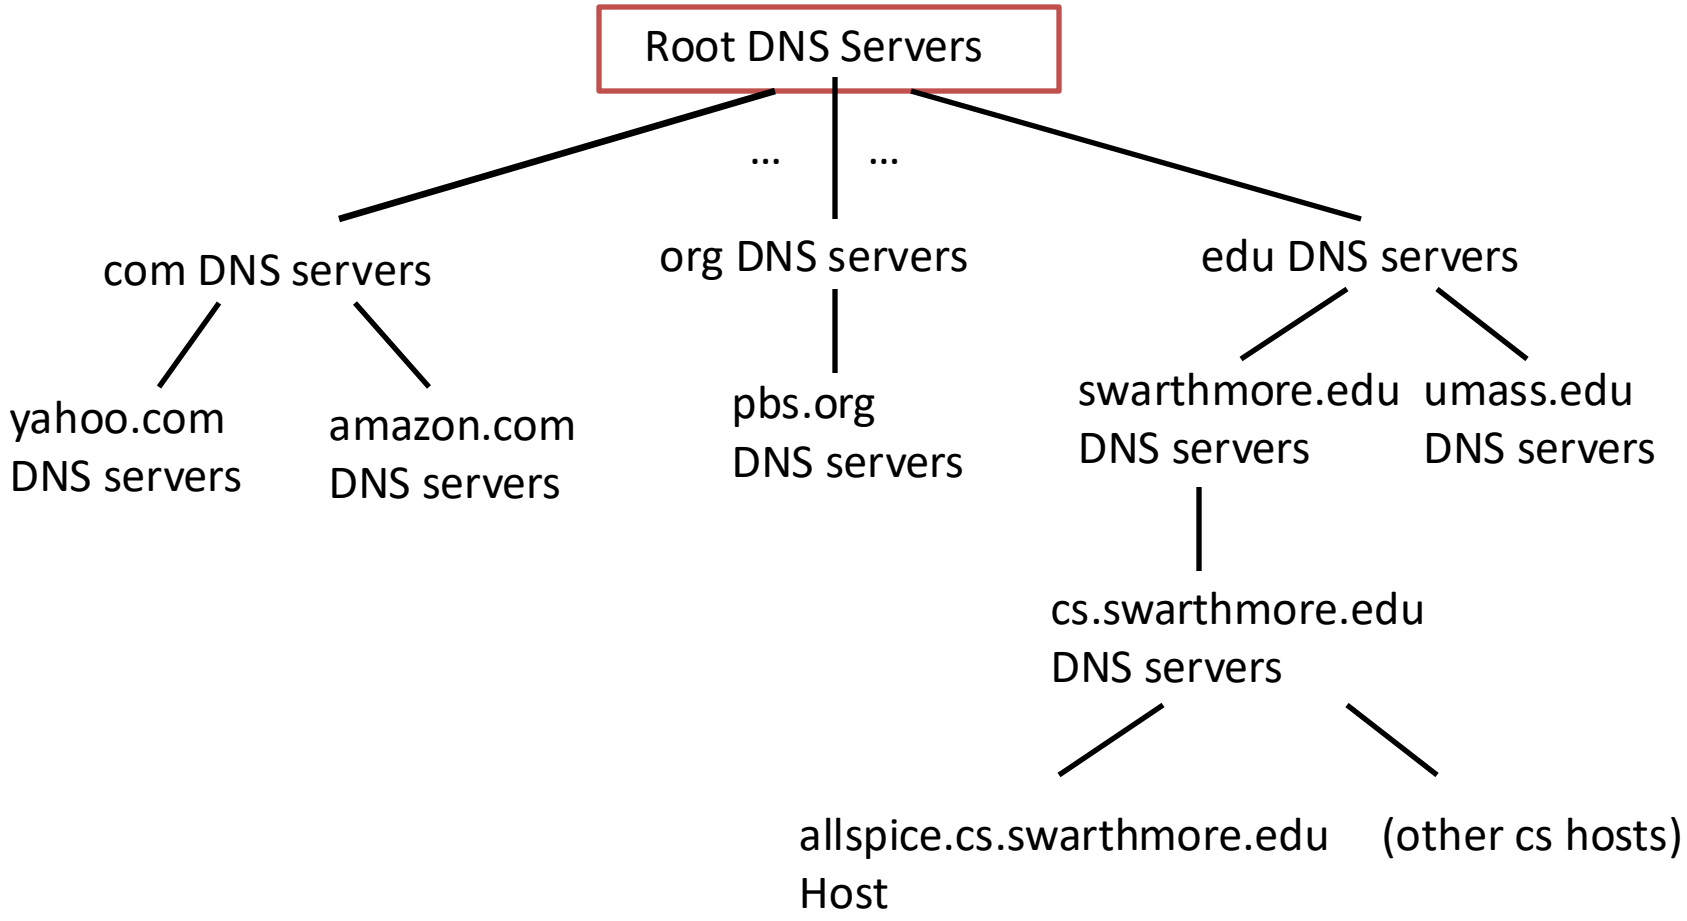

DNS: Root Name Servers

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

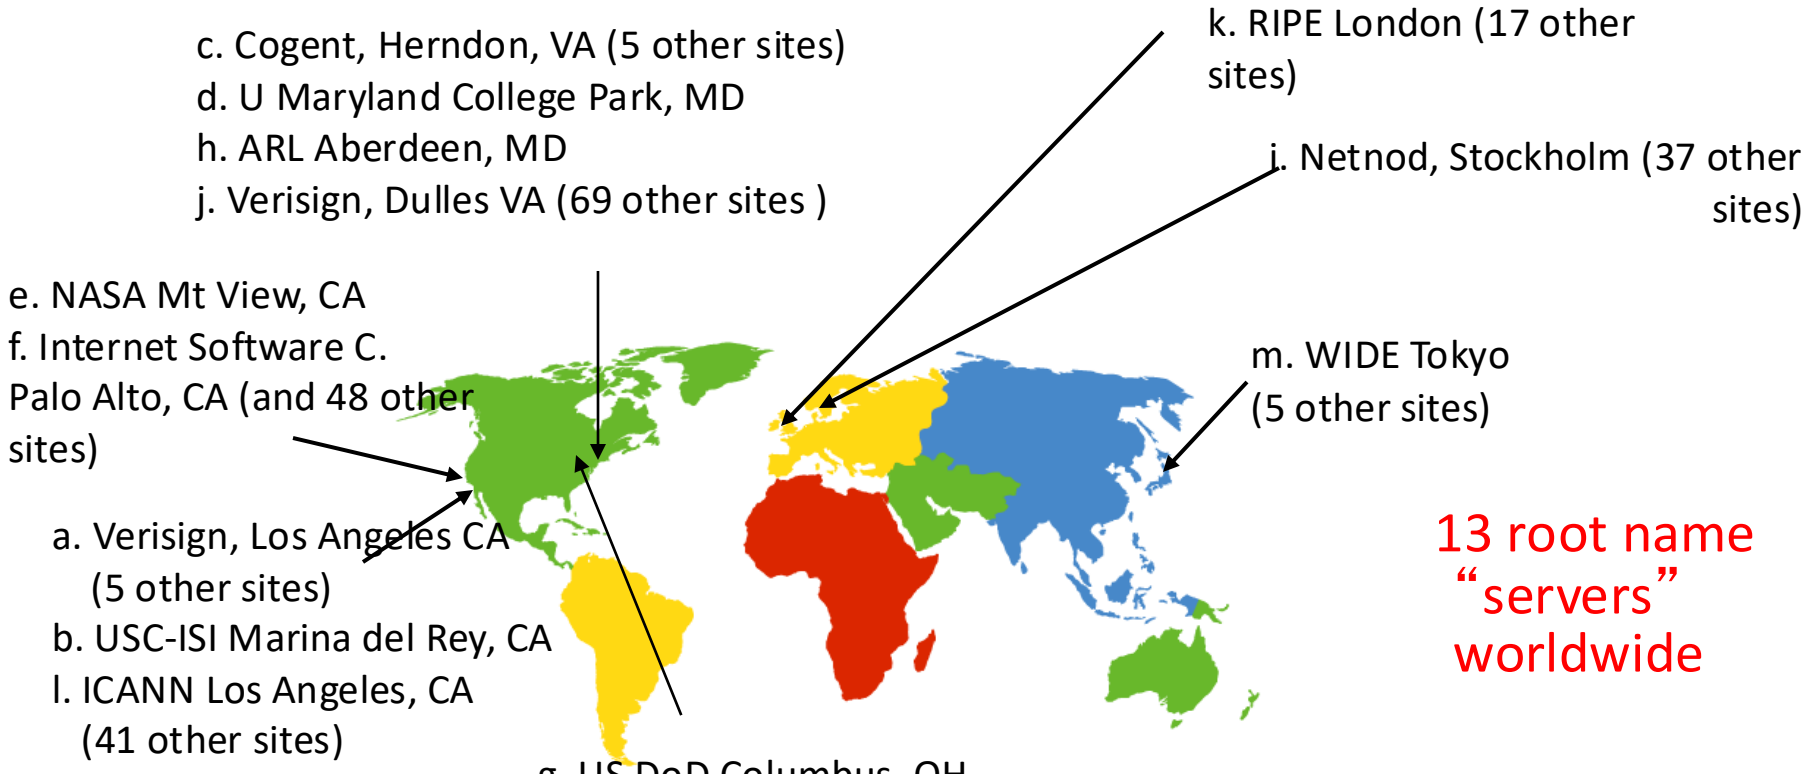

DNS: Root Name Servers

- Root name server:
  - Knows how to find top-level domains (.com, .edu, .gov, etc.)
  - How often does the location of a TLD change?
  - approx. 1700 total root servers
  - Significant amount of traffic is not legitimate

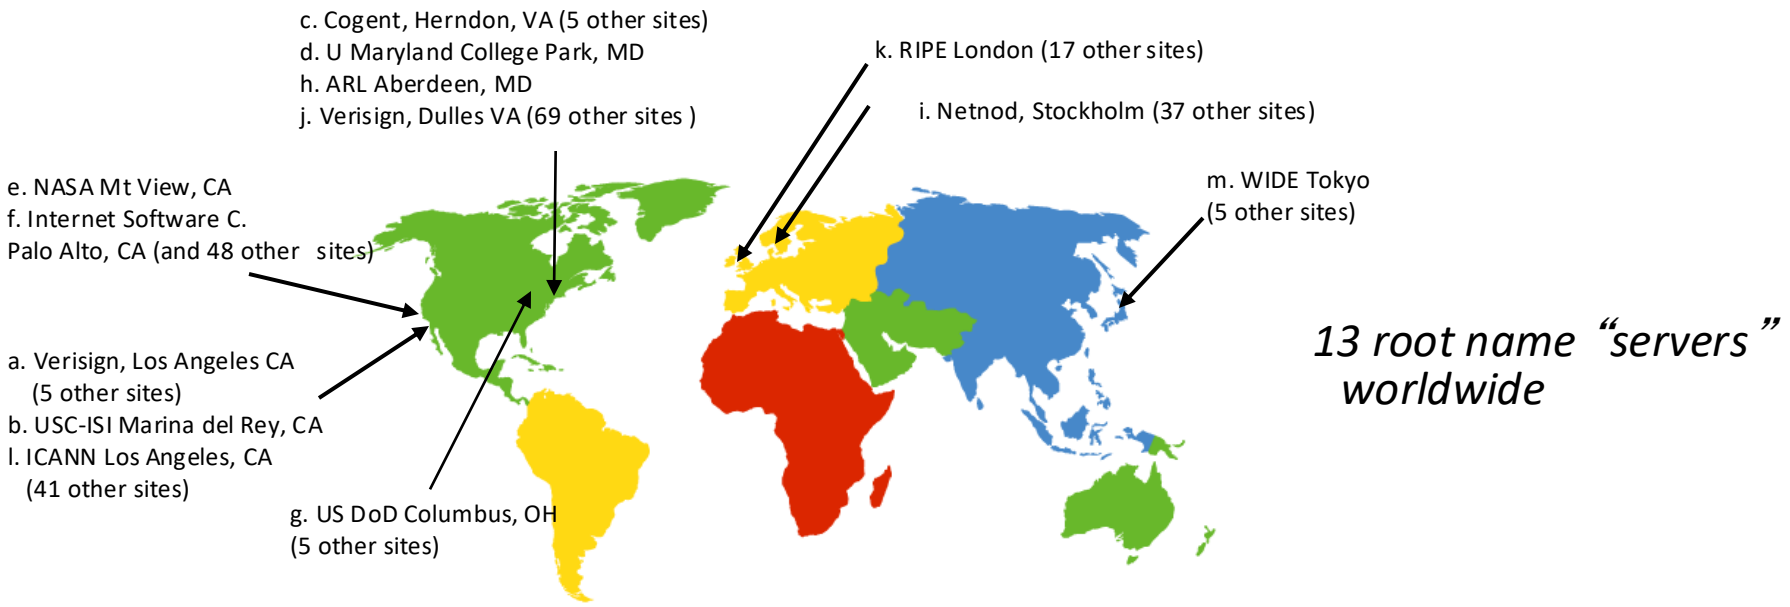

13 root name "servers" worldwide

Root Name Servers

- Responsible for the Root Zone File

```
com.        172800 IN   NS  a.gtld-servers.net.
com.        172800 IN   NS  b.gtld-servers.net.
com.        172800 IN   NS  c.gtld-servers.net.
```

  - In practice, most systems cache this information
  - Lists the TLDs and who controls them
  - ~272KB in size

Top Level Domain (TLD) servers

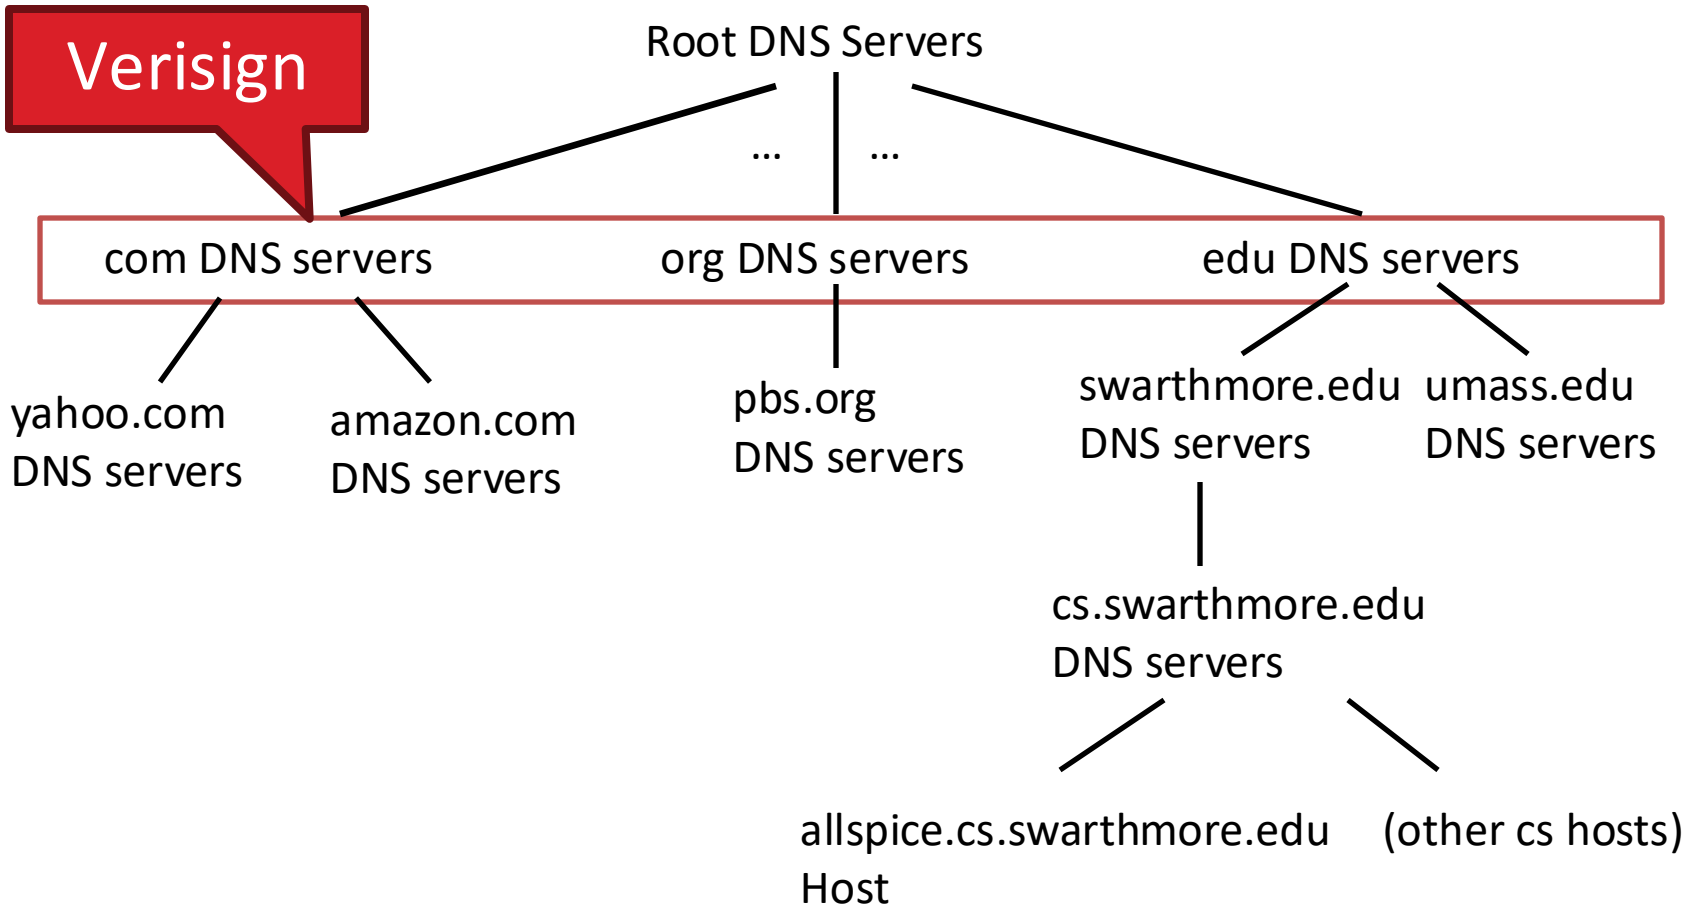

Top Level Domain (TLD) servers

- who maintains the servers?:
  - Verisign: .com, .net
  - Educause: .edu (Verisign backend)
  - local governments or companies

- Responsible for:
  - com, org, net, edu, gov, aero, jobs, museums,
  - all top-level country domains, e.g.: uk, fr, de, ca, jp, etc

Authoritative Servers

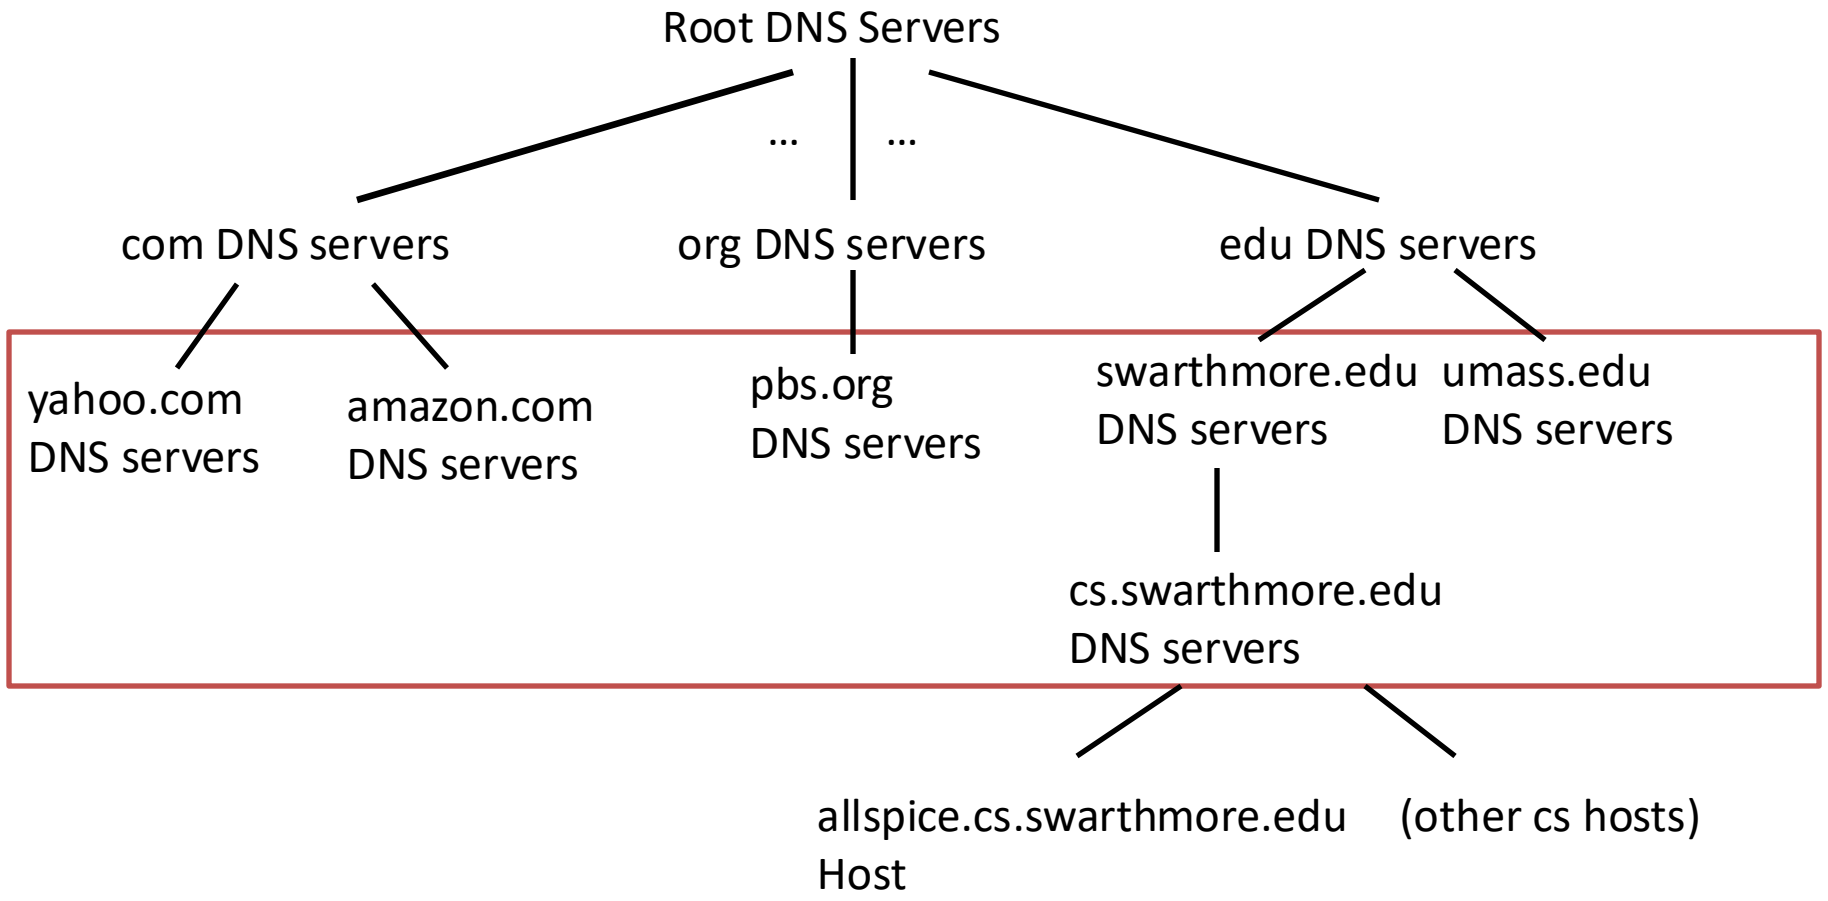

Authoritative Servers

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"

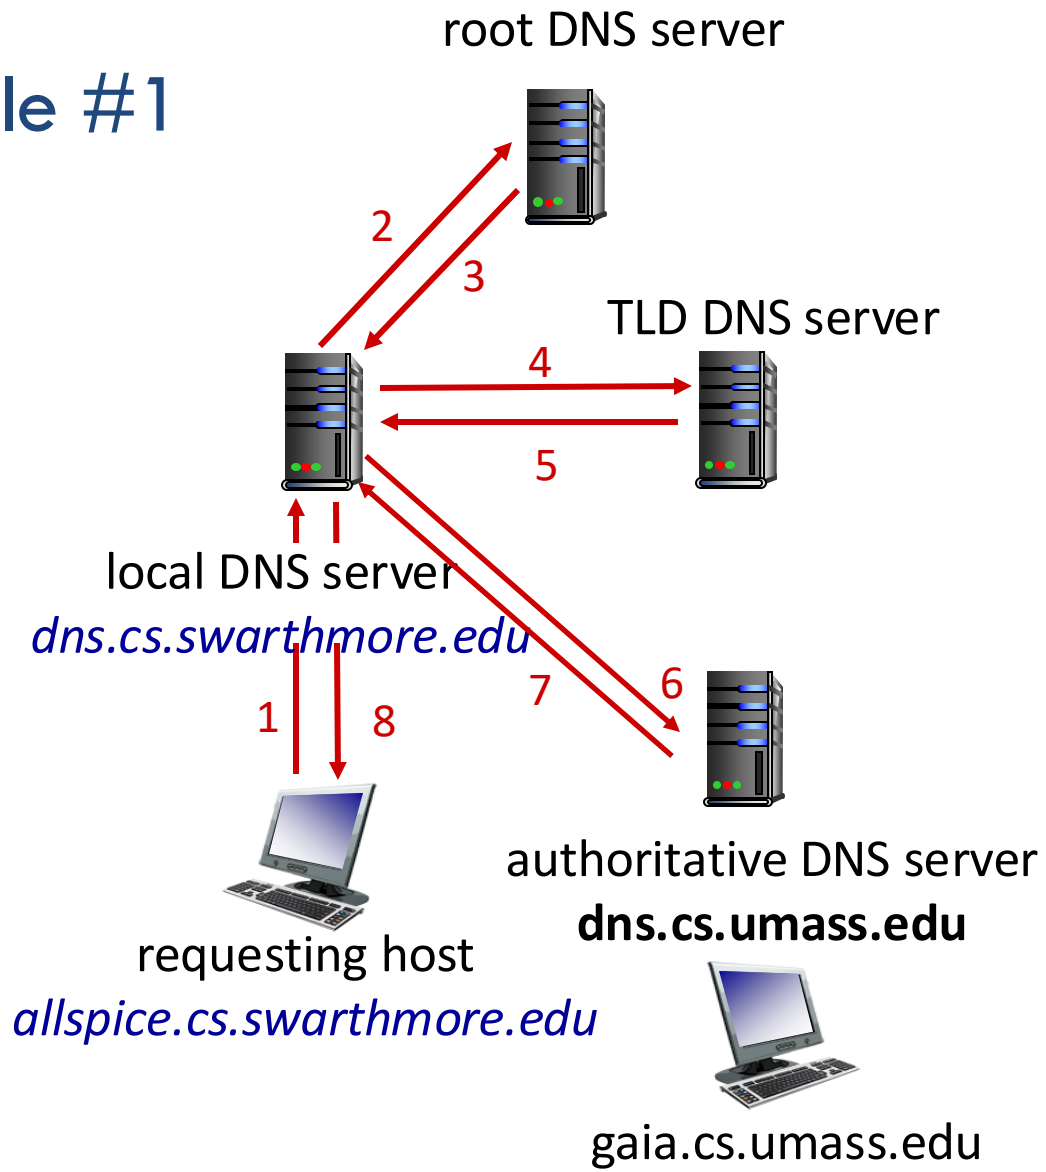

DNS name resolution example #2

recursive query:

- each server asks the next one, in a chain

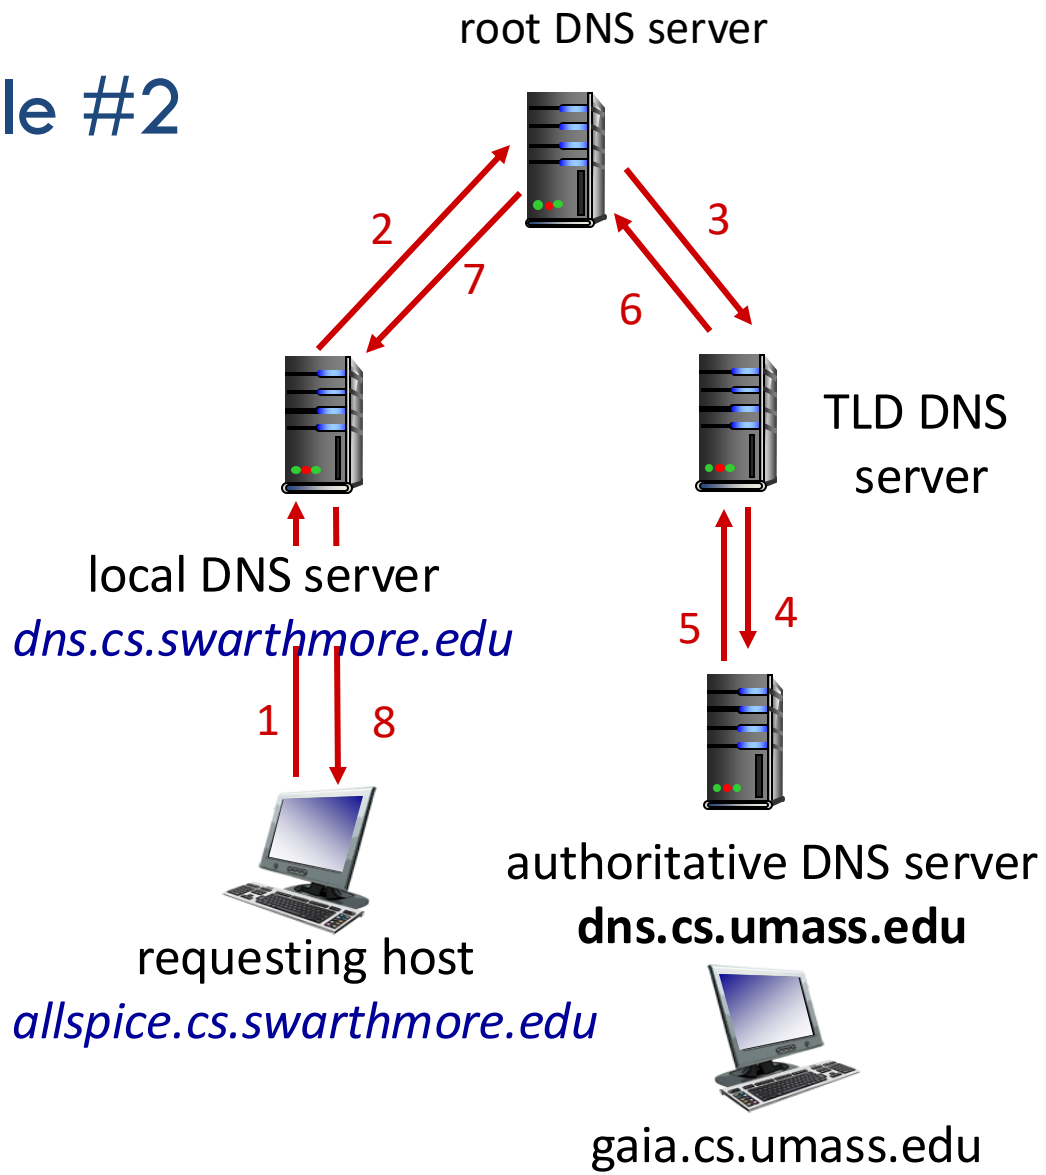

Which would you use? Why?

A. Iterative

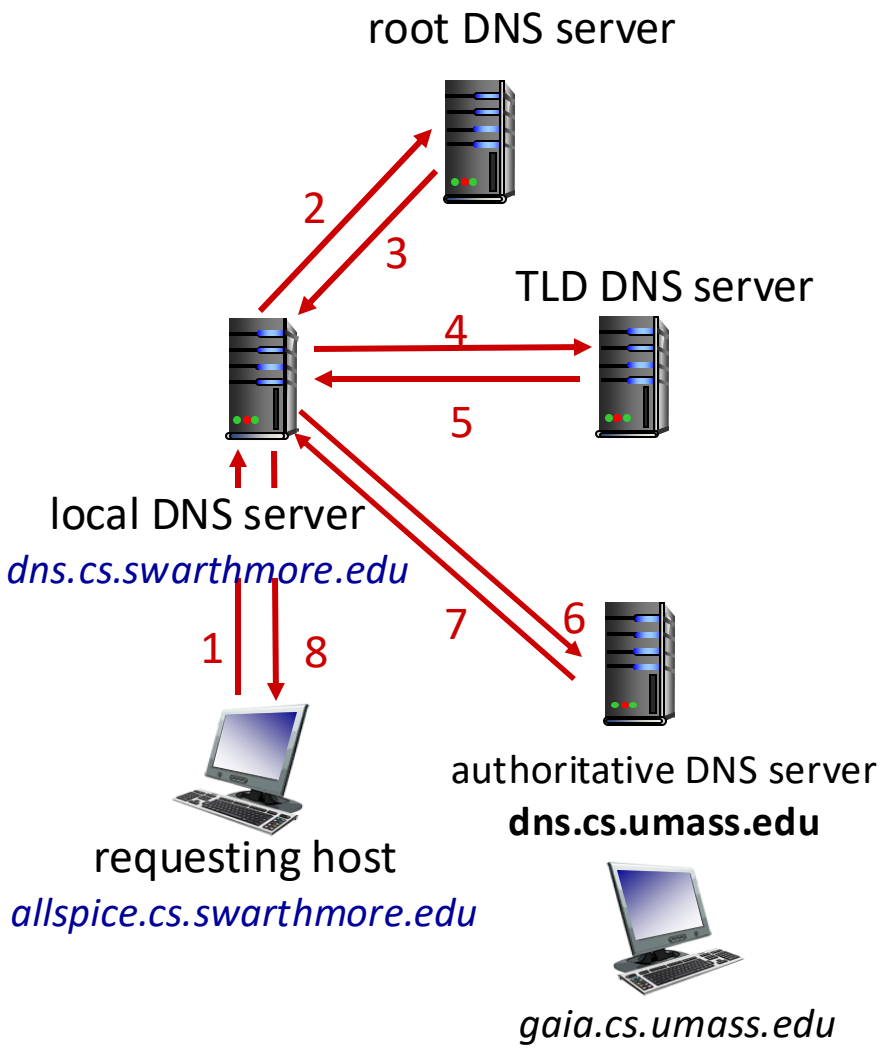

B. Recursive

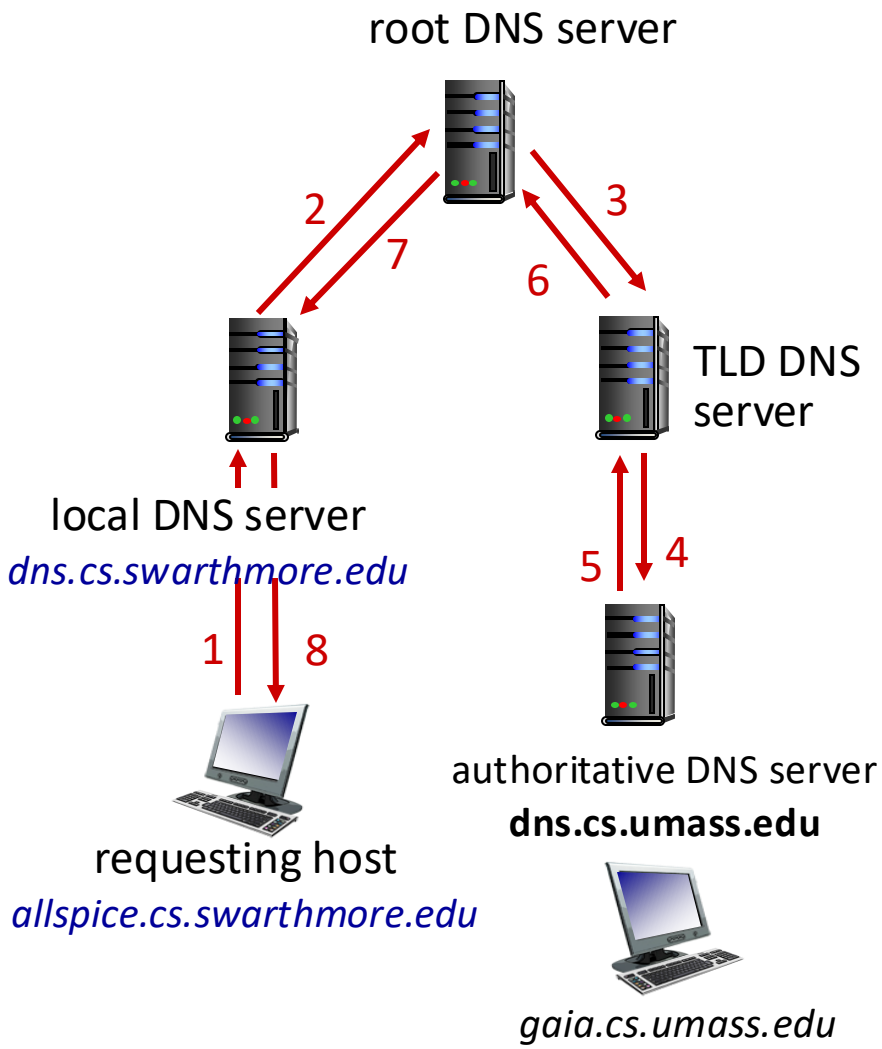

DNS protocol, messages

- query and reply messages, both with same message format

Message header

- identification: 16 bit id for query, reply to query uses same id.
- flags: recursion, query/reply
- # Resource Records to follow

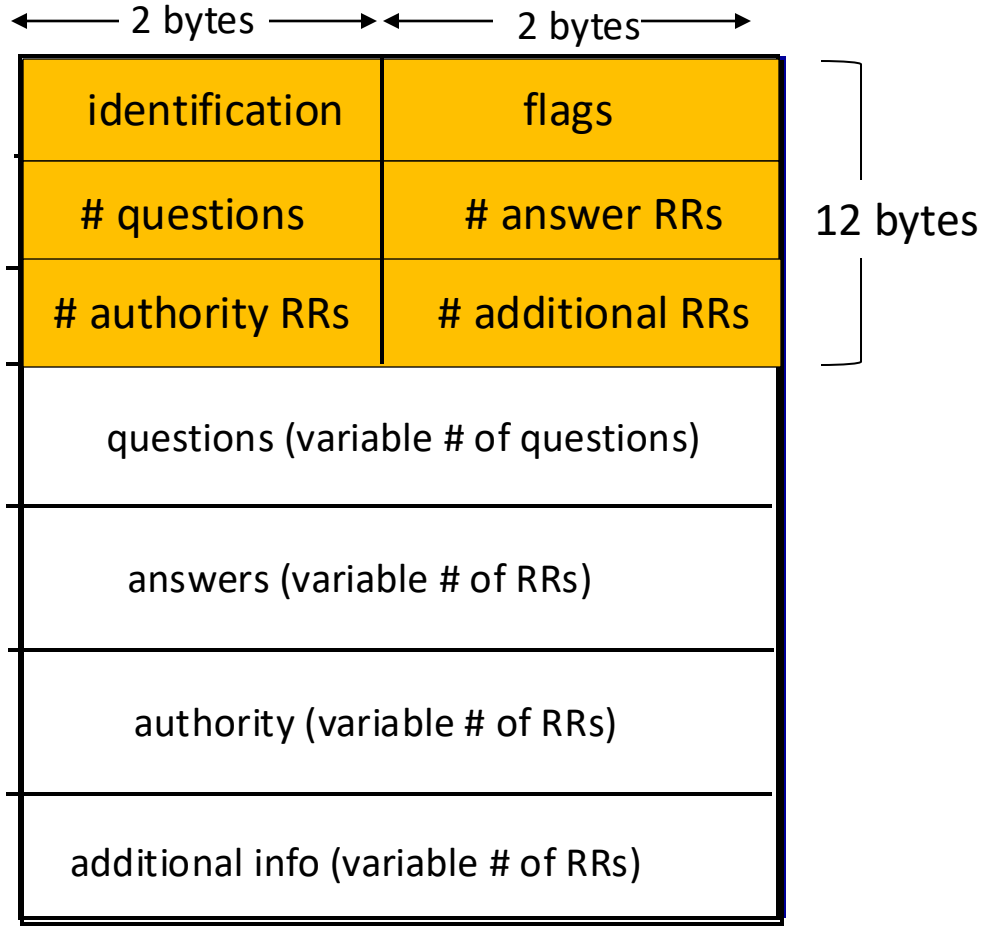

DNS protocol, messages

- query and reply messages, both with same message format!

Binary Protocol!

- Delimiters: pre-defined lengths/field
- Names: <len><name>

Sent via UDP (User Datagram Protocol)

- No connection established
- Not reliable!

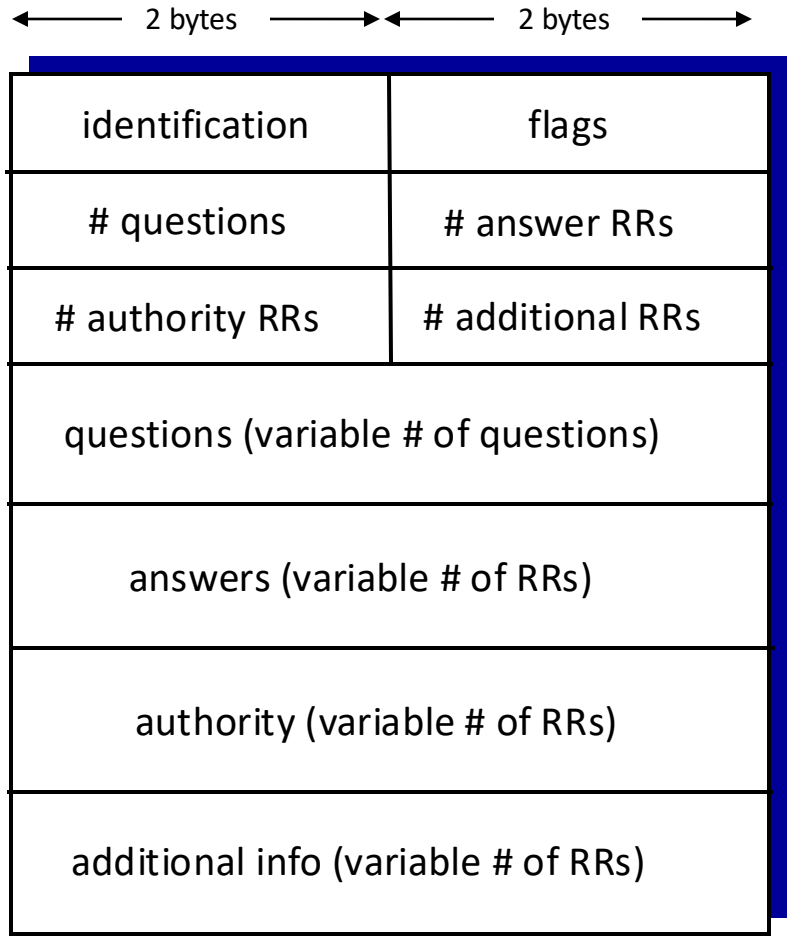

DNS Types

RR format: `(name, value, type, ttl)`

- Type = A / AAAA
  - Name = domain name
  - Value = IP address
  - A is IPv4, AAAA is IPv6

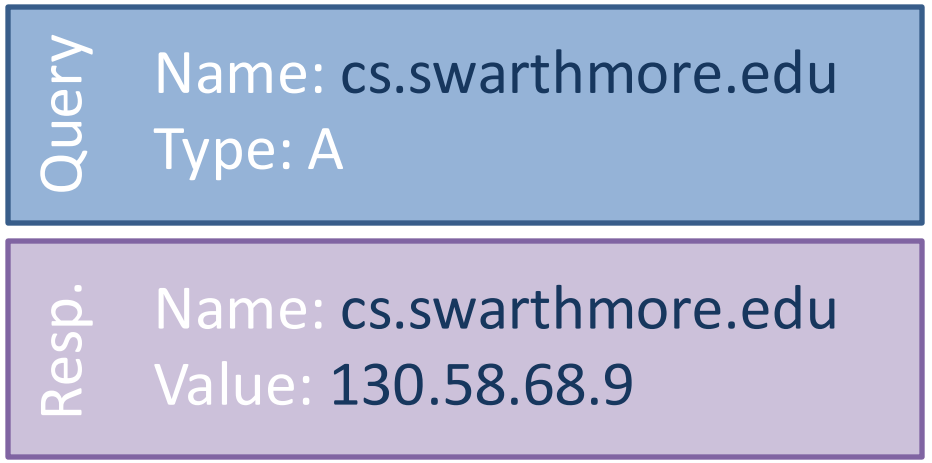

- Type = NS
  - Name = partial domain
  - Value = name of DNS server for this domain
  - "Go send your query to this other server"

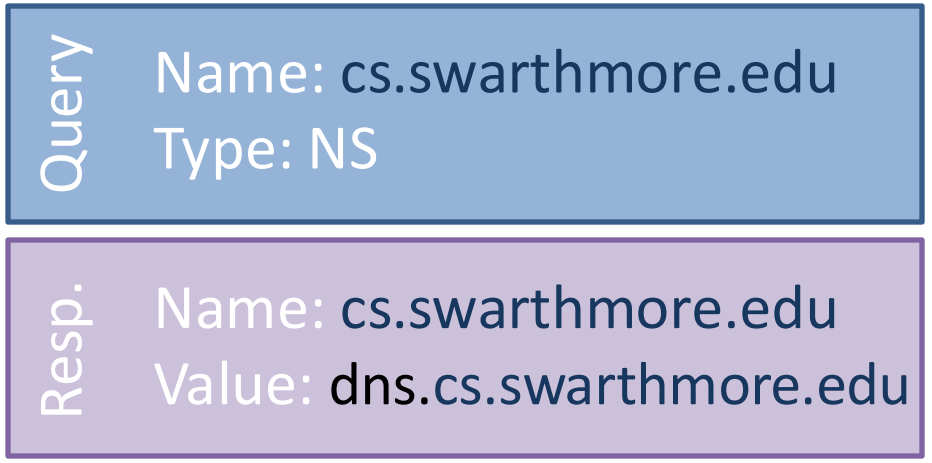

DNS Types, Continued

RR format: `(name, value, type, ttl)`

- Type = CNAME
  - Name = hostname
  - Value = canonical hostname
  - Useful for aliasing
  - CDNs use this

Query
Name: foo.mysite.com
Type: CNAME

Resp.
Name: foo.mysite.com
Value: bar.mysite.com

- Type = MX
  - Name = domain in email address
  - Value = canonical name of mail server

Query
Name: cs.umass.edu
Type: MX

Resp.
Name: cs.umass.edu
Value: barramail.cs.umass.edu.

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

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

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 Query Process and Cache

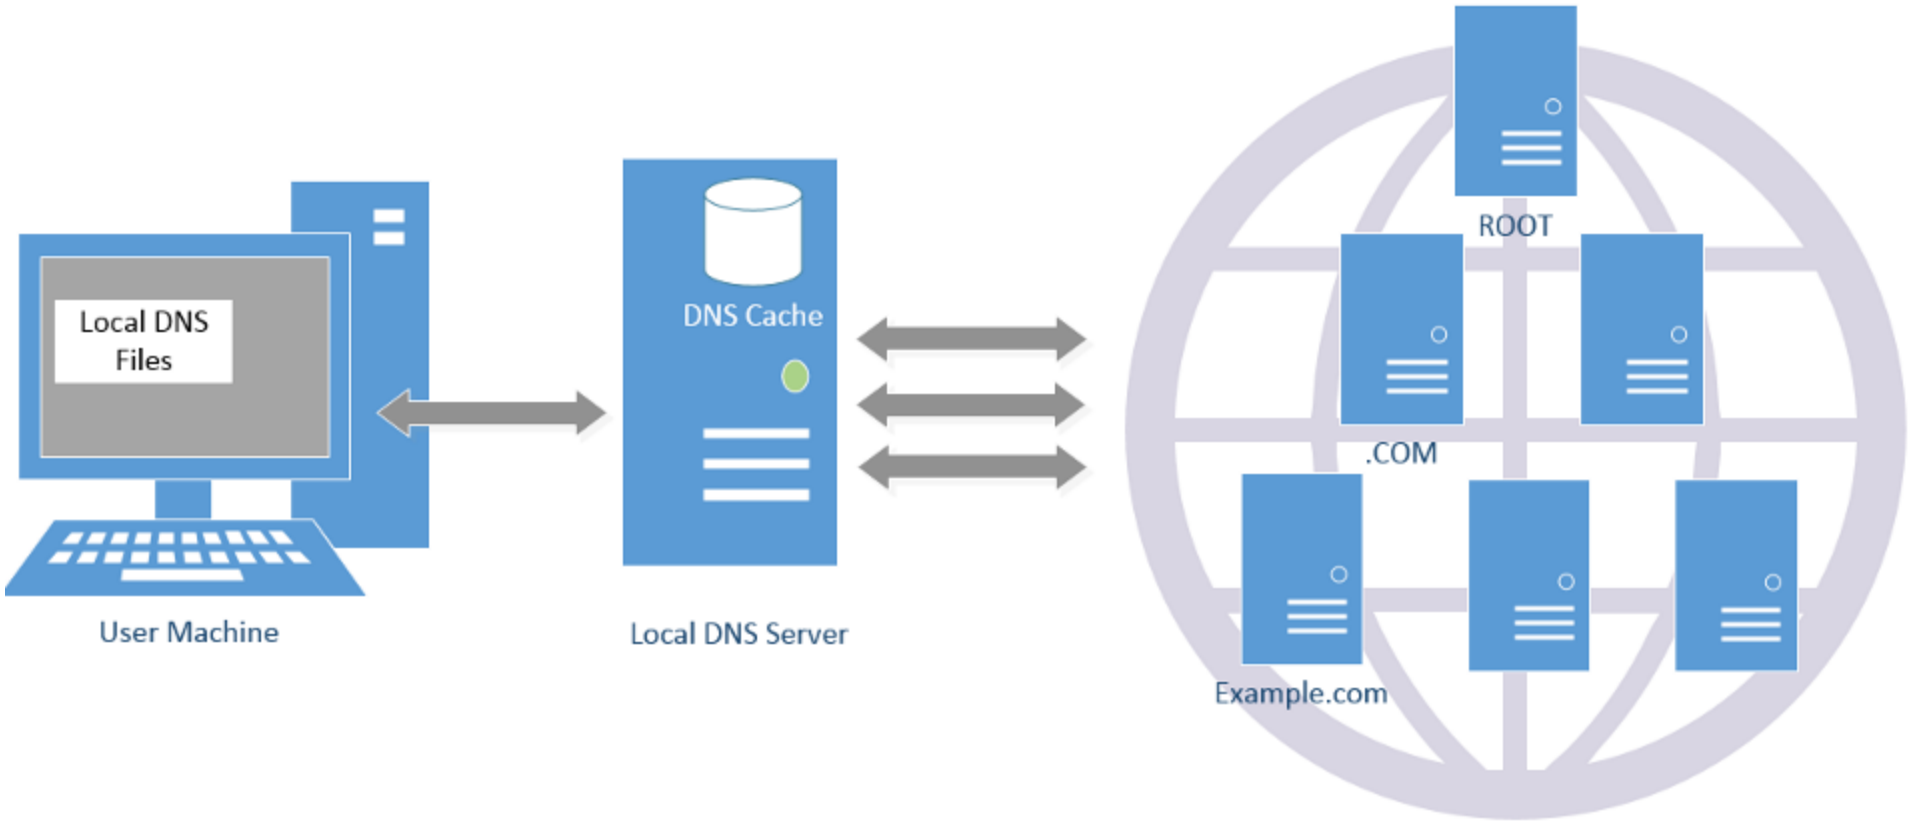

Caching

- 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

DNS Caching

- Why cache?
  - apprx. 1 sec latency before starting a download
  - Popular sites visited often
- Where to cache?
  - Local DNS server
  - Browser

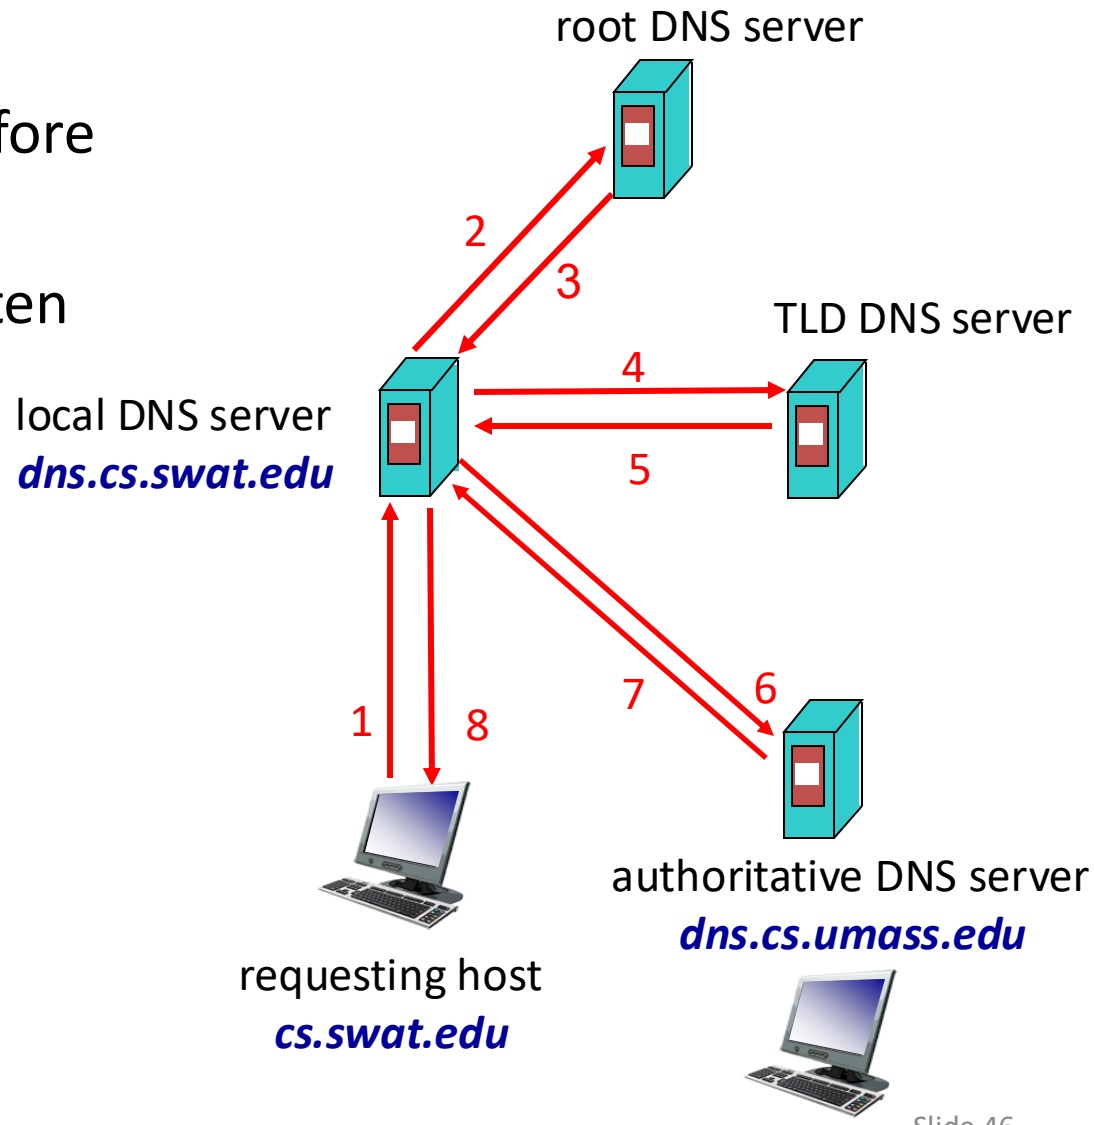

DNS Caching

- When to cache?
  - learn a mapping? cache!
  - any name server can cache
- For how long?
  - until Time To Live (expires)
- What to cache?
  - TLD servers cached – almost never change
  - Root name servers usually, not visited legitimately

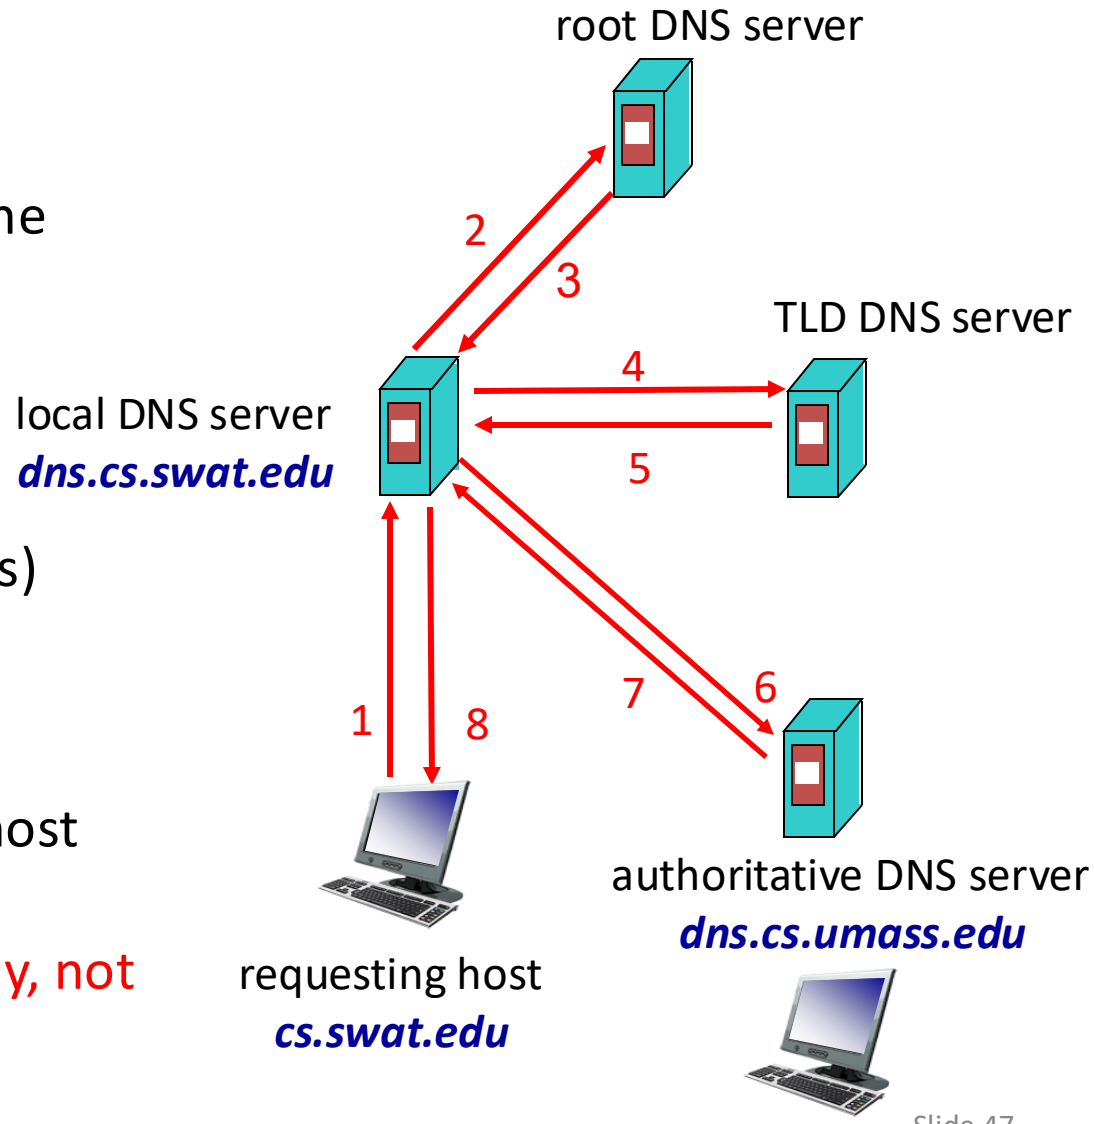

Caching

- 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

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

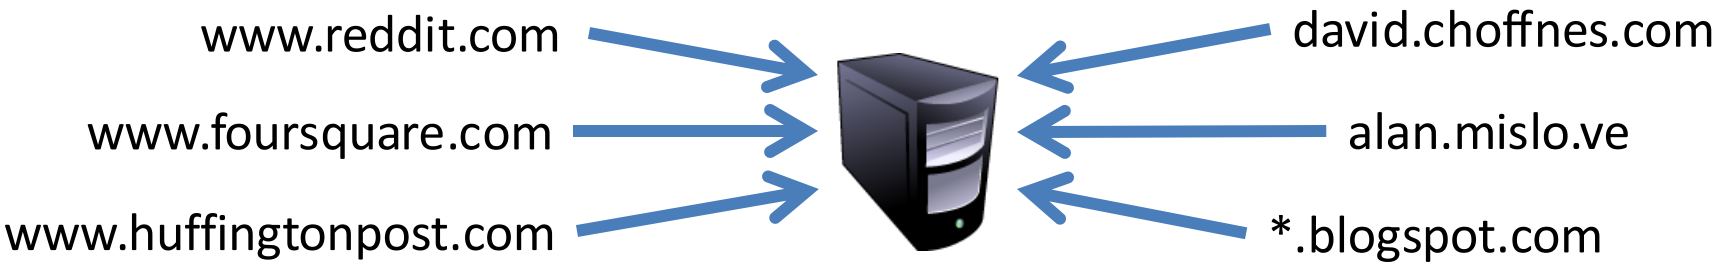

  - One domain can map to multiple machines

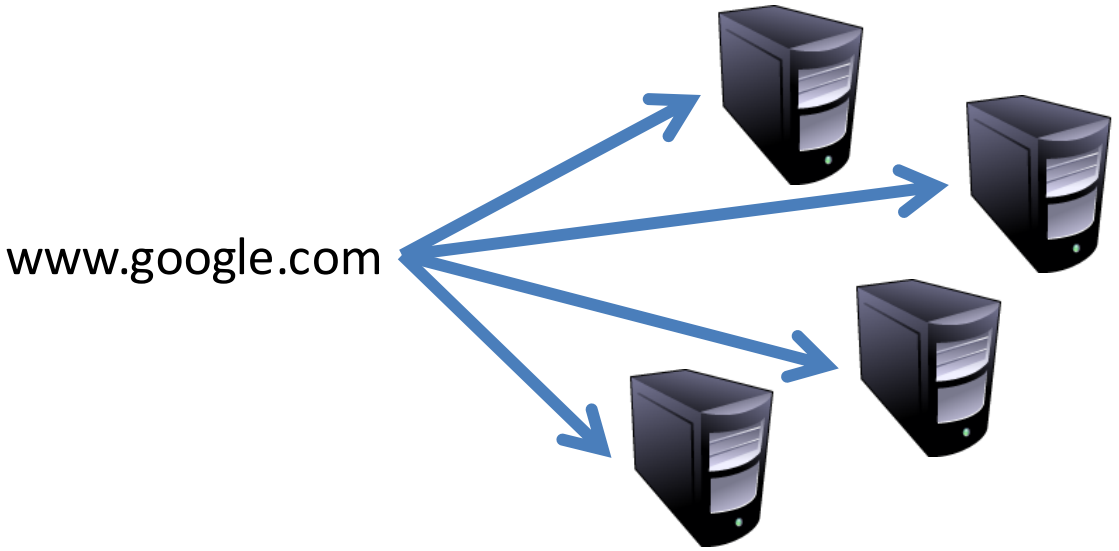

Content Delivery Networks

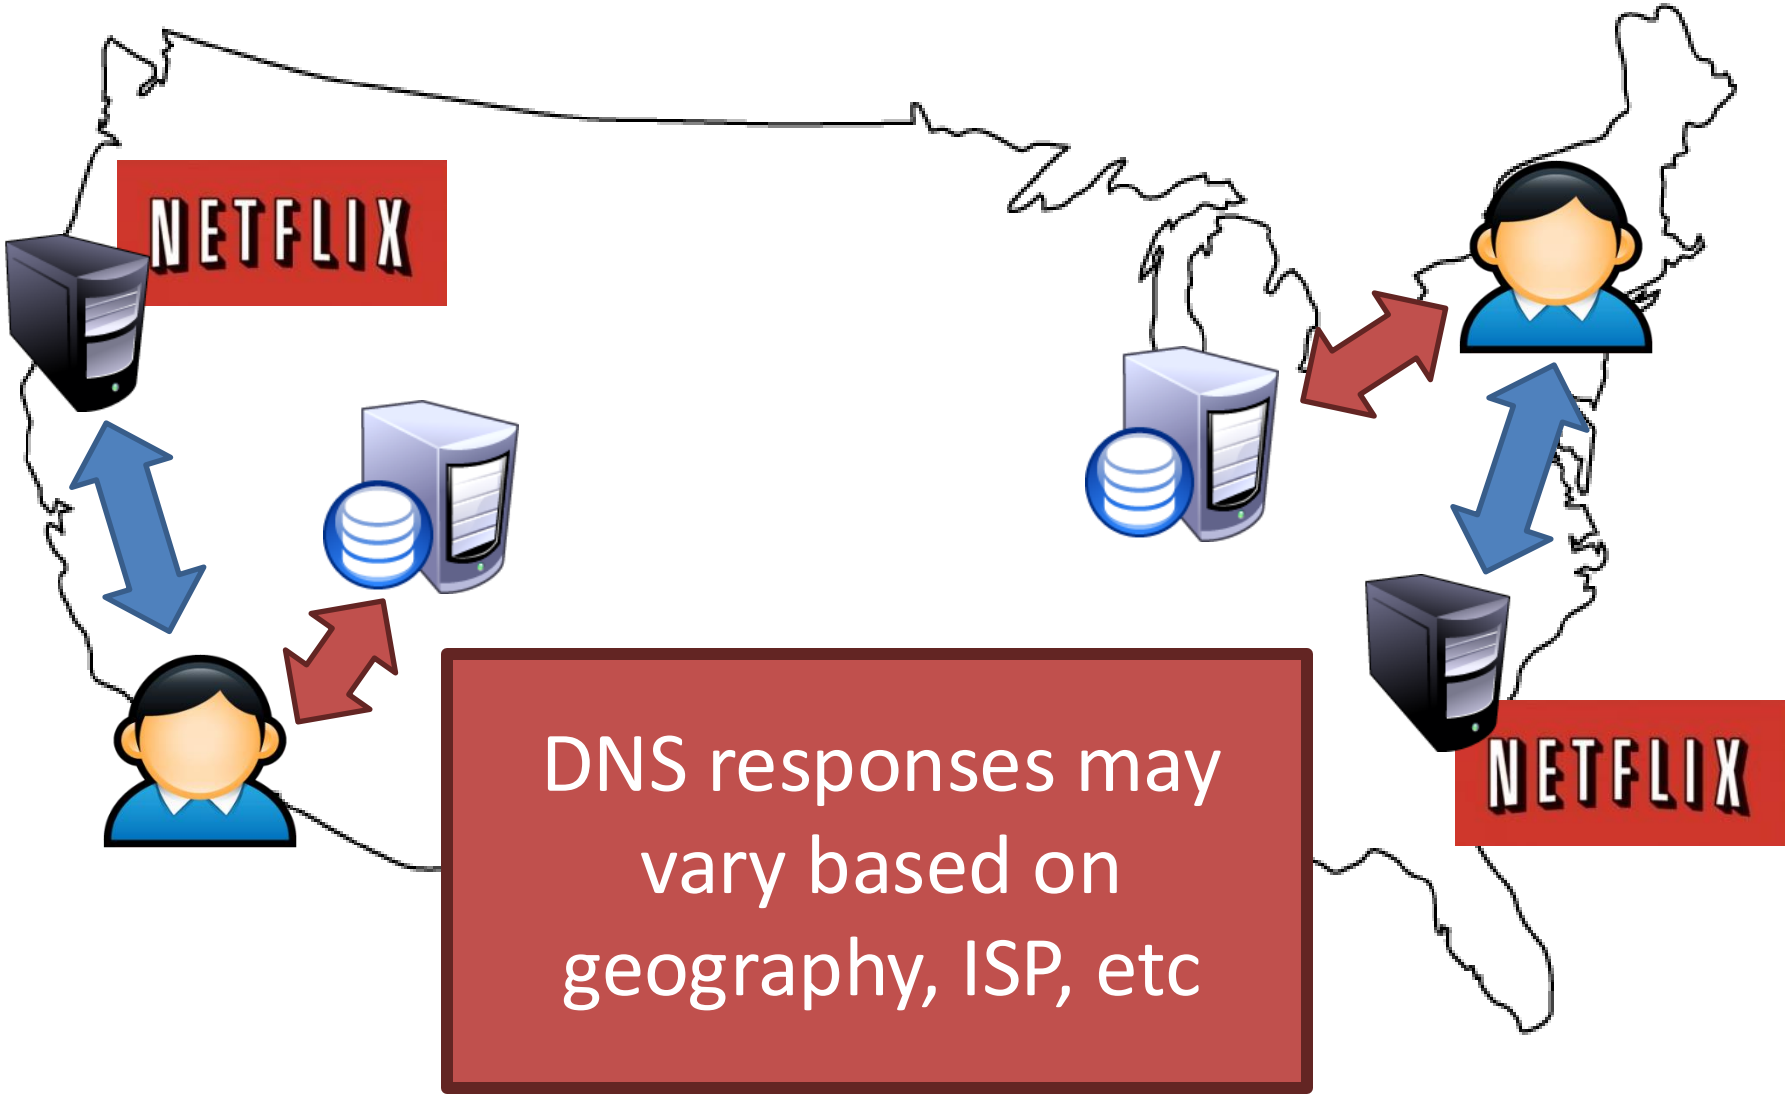

Inserting (or changing) records

Example: new startup "Network Utopia"

- Step 1: Register networkuptopia.com at DNS registrar
  - provide names, IP addresses of authoritative name server (primary and secondary)

Inserting (or changing) records

Example: new startup “Network Utopia”

- Step 2: Registrar inserts two RRs into .com TLD server
  - (networkutopia.com, dns1.networkutopia.com, NS)
  - (dns1.networkutopia.com, 212.212.212.1, A)

Inserting (or changing) records

Example: new startup "Network Utopia"

- Step 3: Set up authoritative server at that name/address
  - Create records for the services:

Inserting (or changing) records

Example: new startup "Network Utopia"

- Step 3: 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

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

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