CS 43: Computer Networks

06:Domain Name Systems

September 24, 2020

SWARTHMORE COLLEGE

Today

- Identifiers and addressing

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

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 ?

DNS: Application Layer Protocol

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

Where

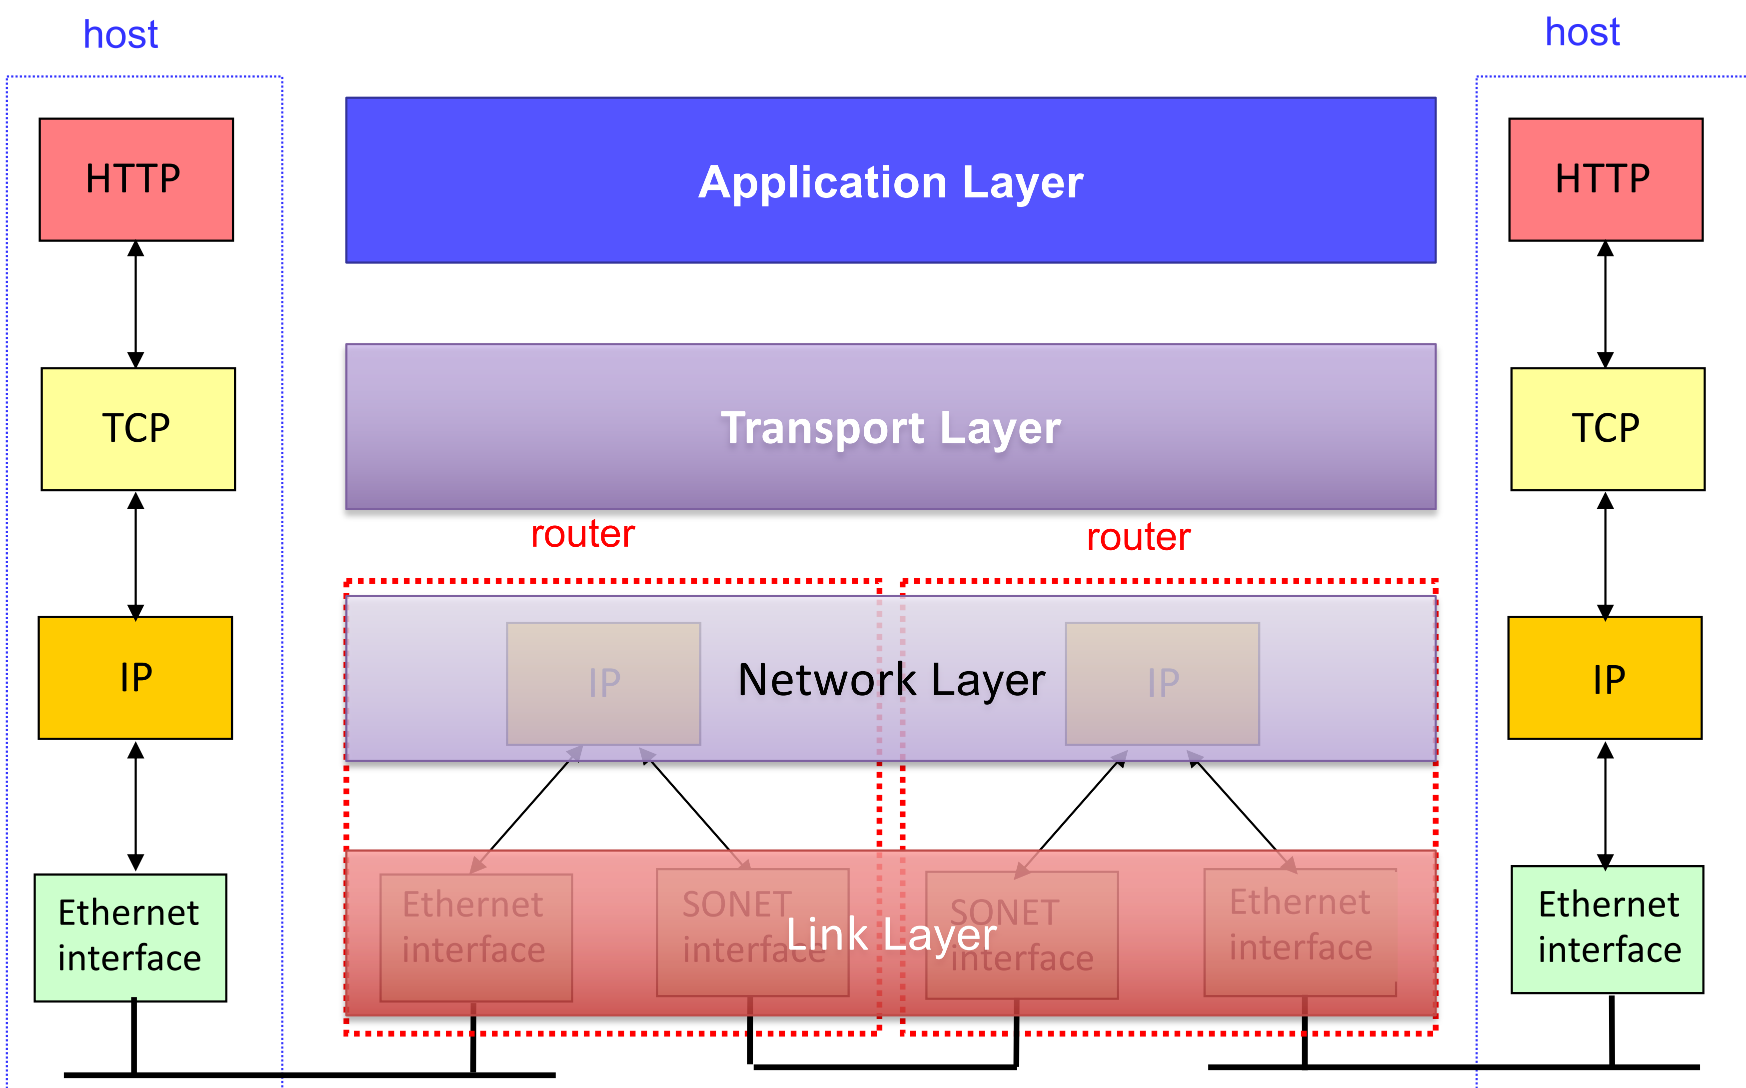

host
HTTP
TCP
IP
Ethernet interface
Application Layer
Transport Layer
router
router
IP
Network Layer
IP
Ethernet interface
SONET interface
Link Layer
SONET interface
Ethernet interface
host
HTTP
TCP
IP
Ethernet interface

Recall: TCP/IP Protocol Stack

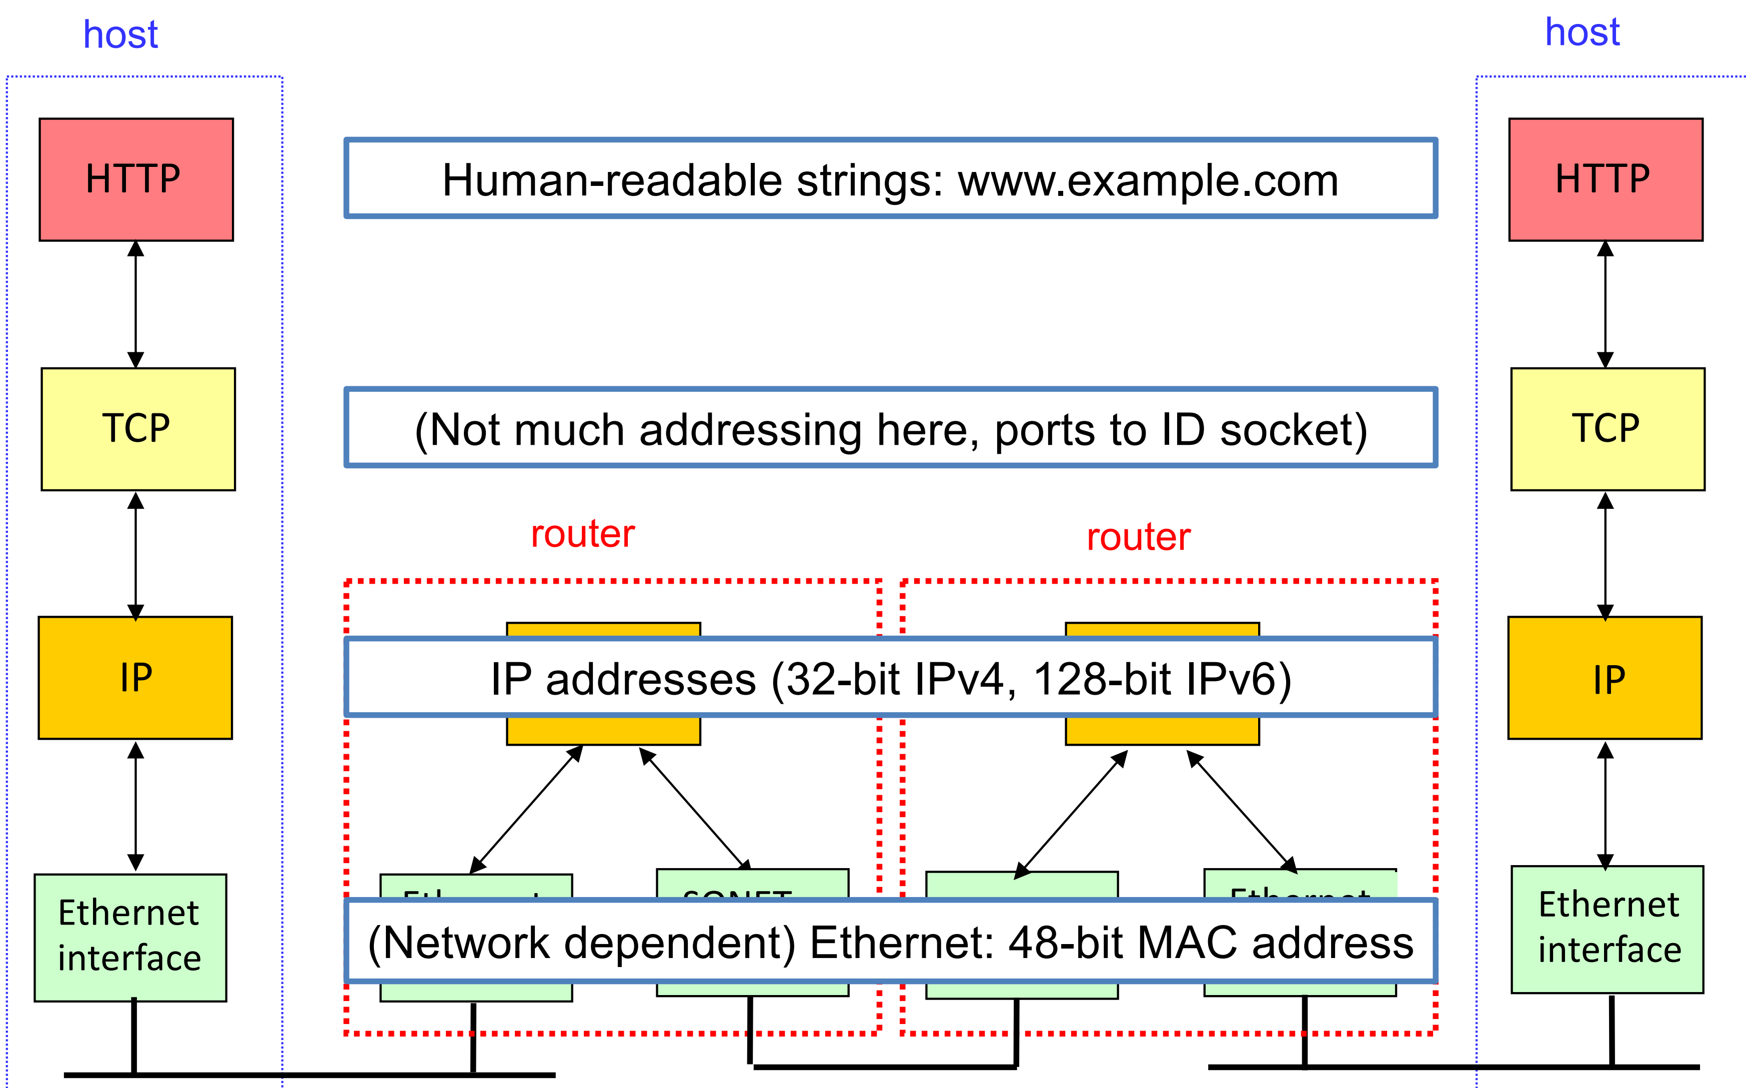

DNS: domain name system

- 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.

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 (Named Data Networking Efforts).

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 (mostly true).

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

Mapping Between Identifiers

- Domain Name System (DNS)
  - Given a host name, provide the IP address
  - Given an IP address, provide the host name

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.

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

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: a distributed, hierarchical database

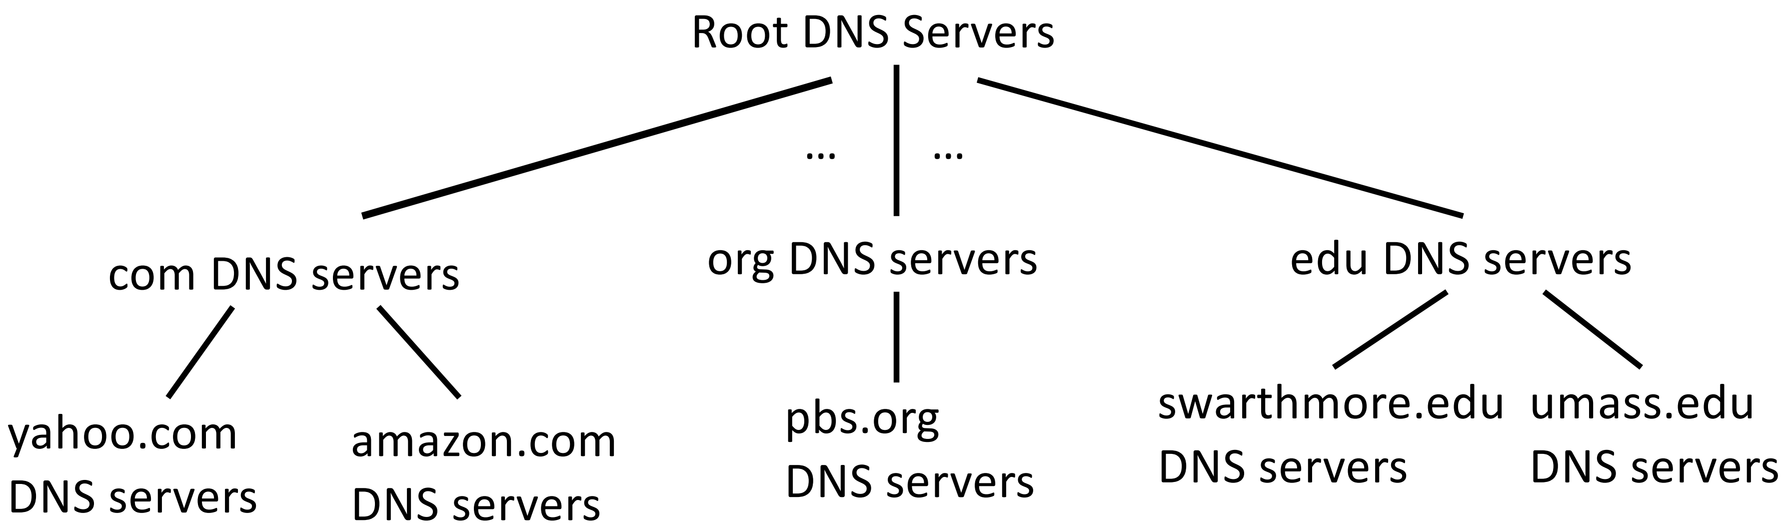

DNS: a distributed, hierarchical database

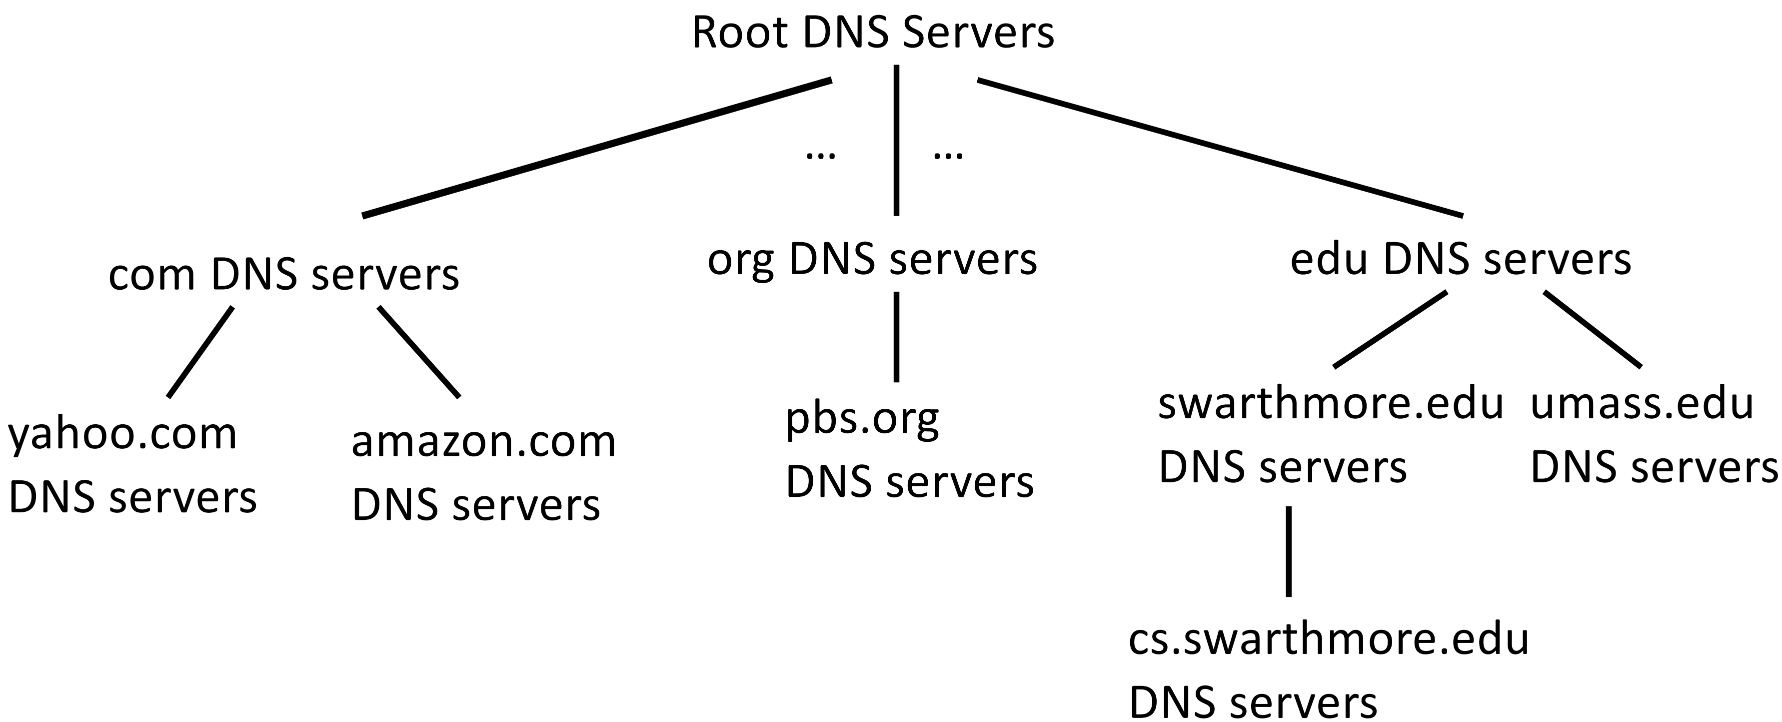

DNS: a distributed, hierarchical database

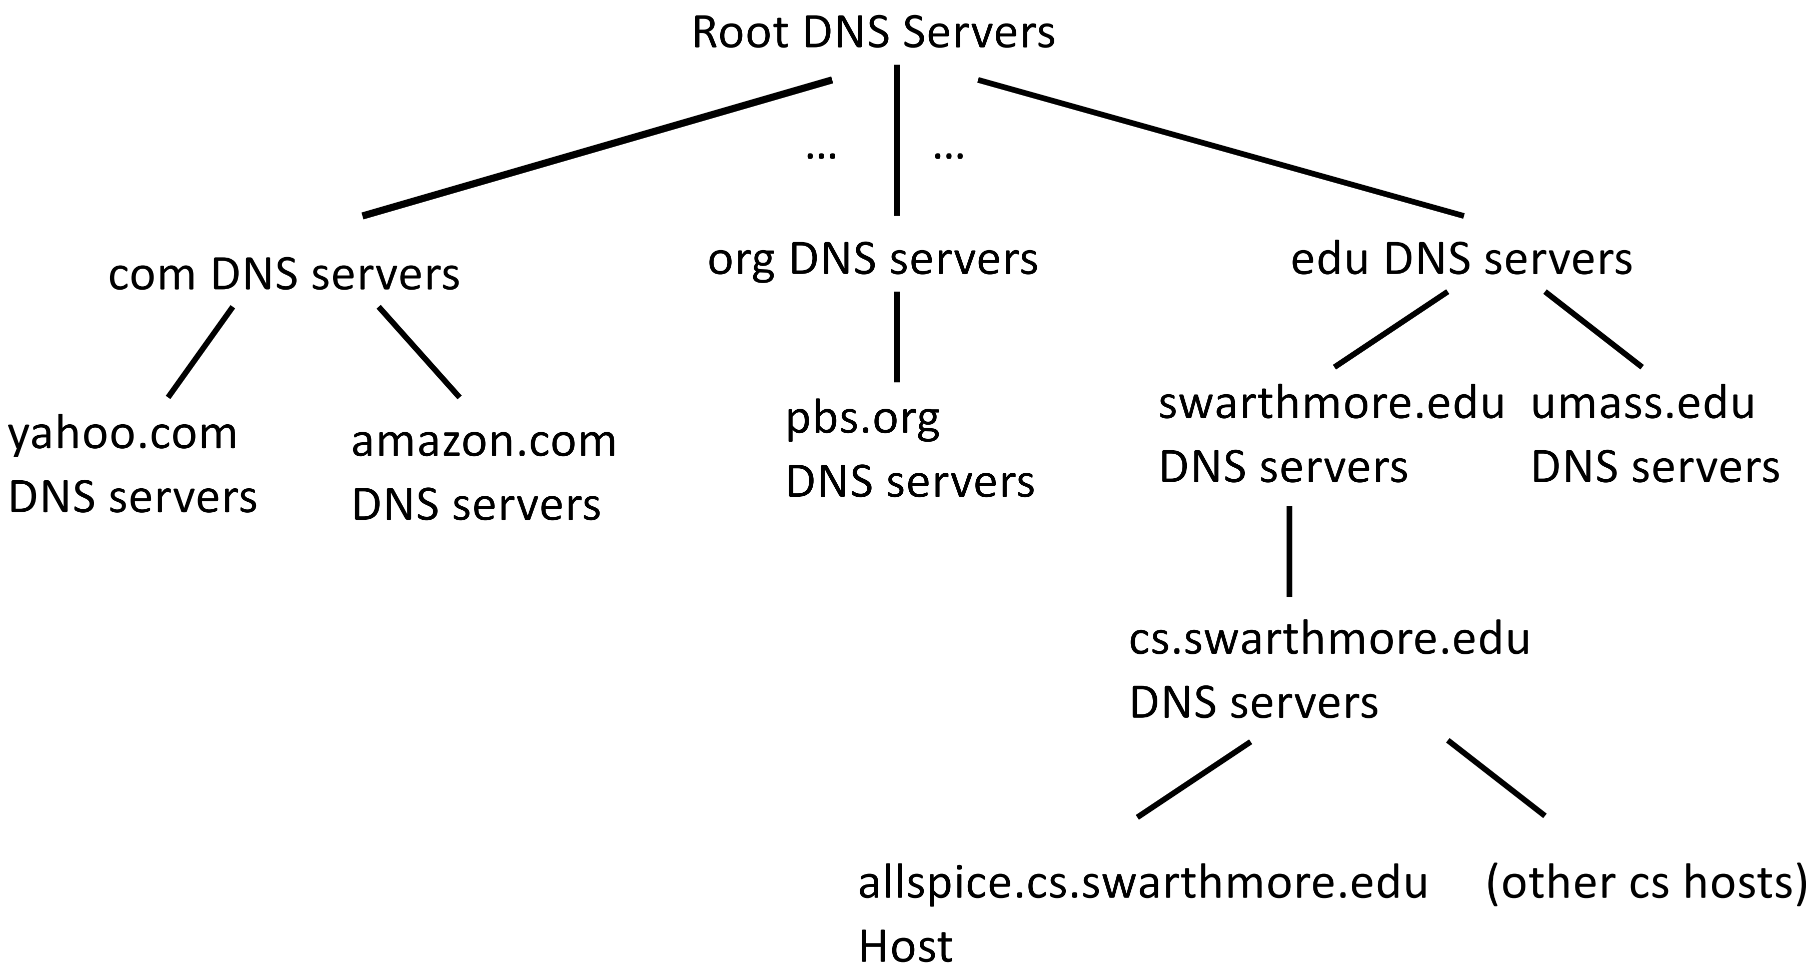

- allspice.cs.swarthmore.edu.

Nameless root,
Usually implied.

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

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 (without caching) have a look-up overhead of traversing the DNS hierarchy .

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

D. Some other reason.

DNS: a distributed, hierarchical database

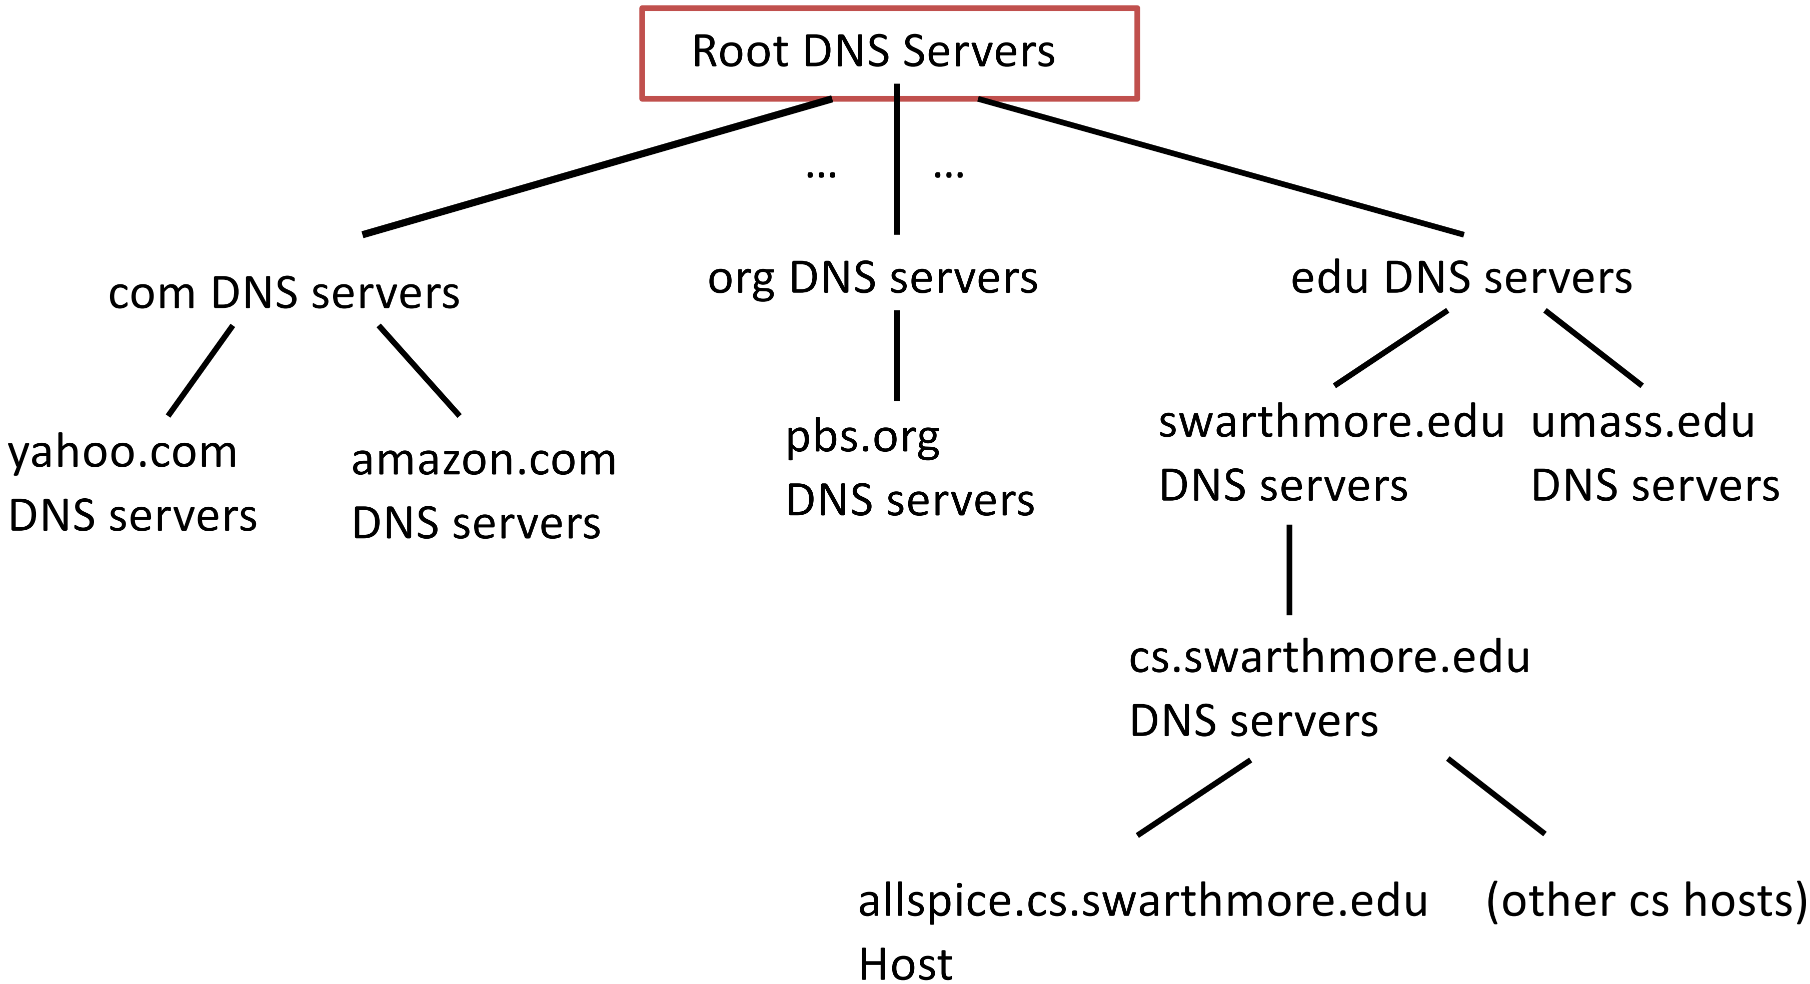

DNS: Root Name Servers

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

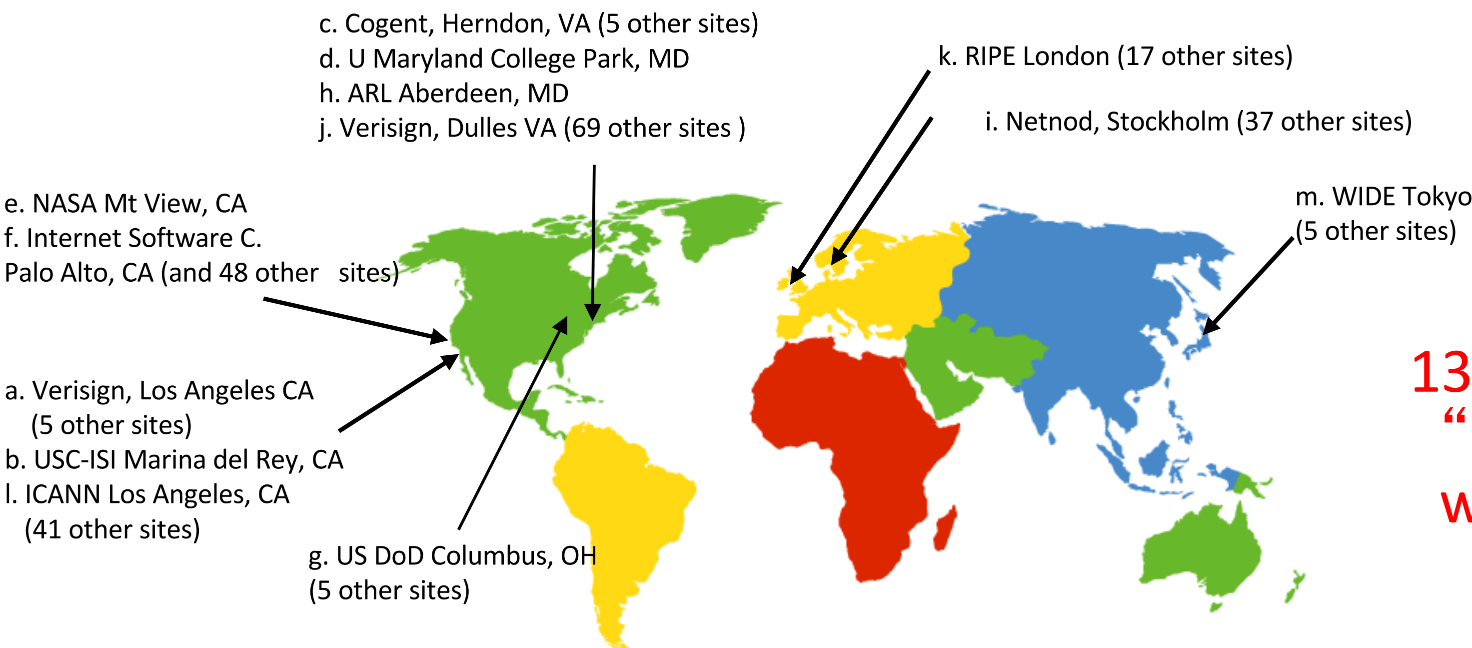

13 root name "servers" worldwide

DNS: a distributed, hierarchical database

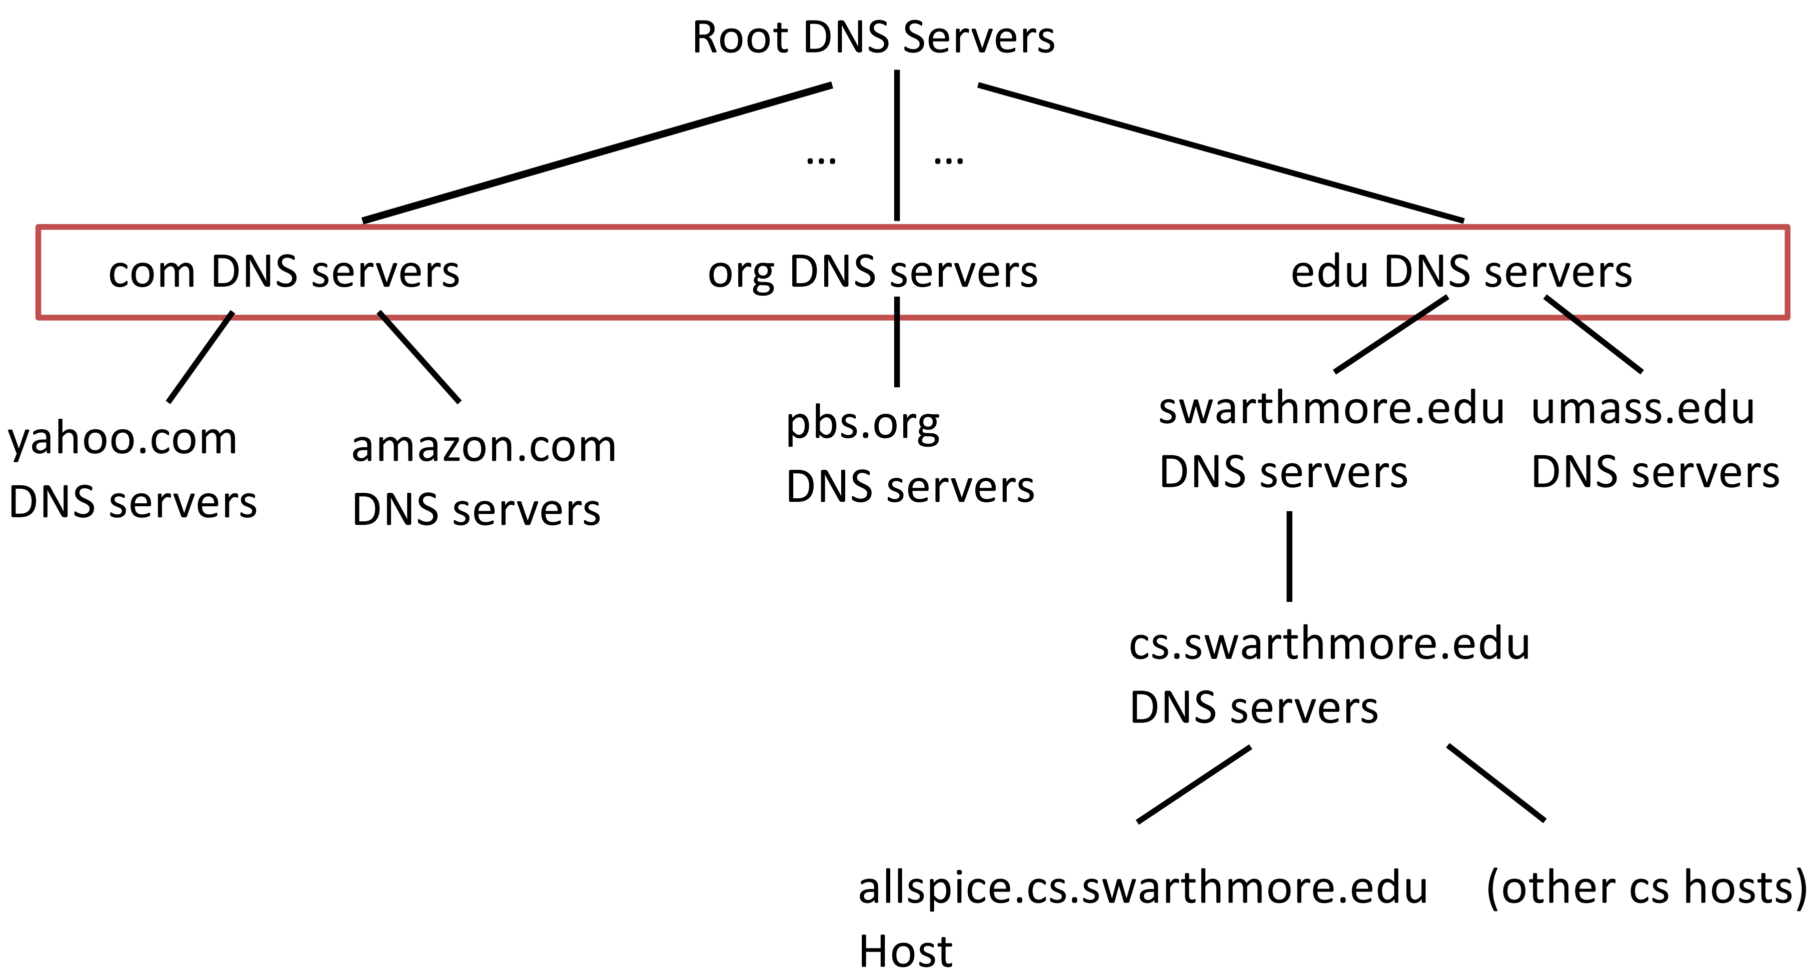

Top Level Domains

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

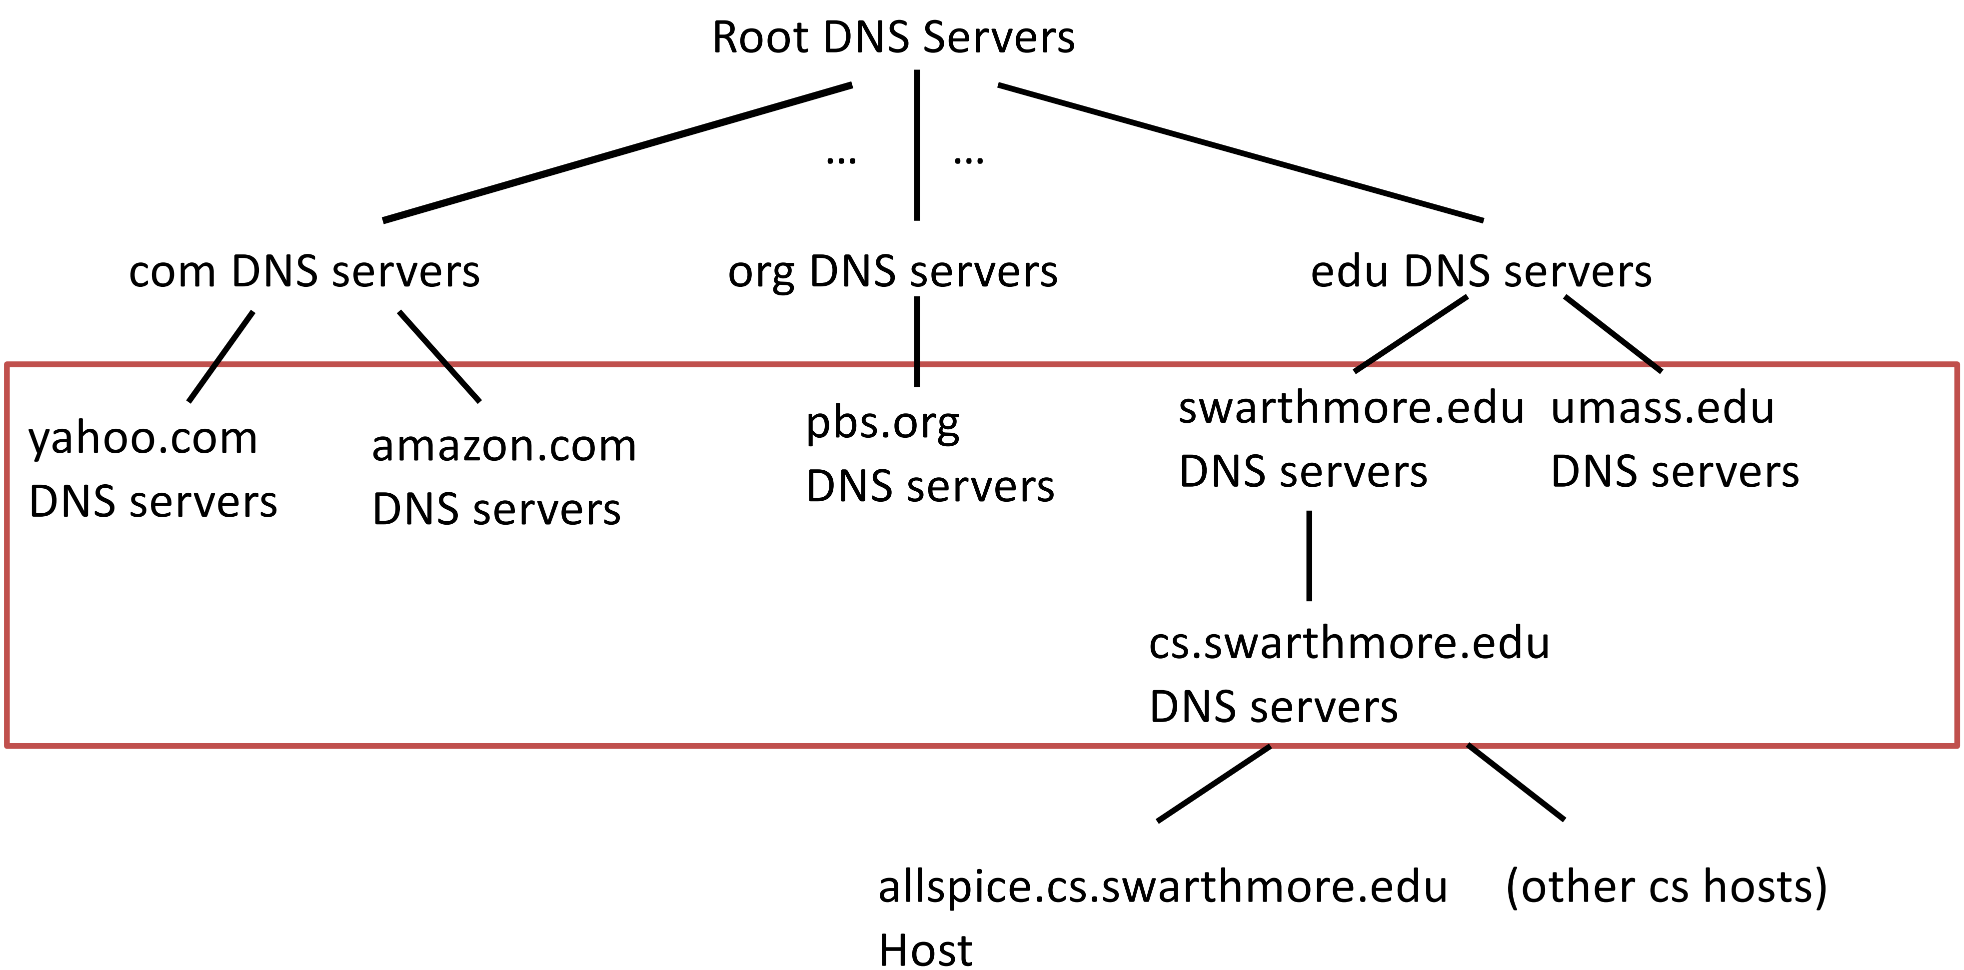

Authoritative Servers

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"

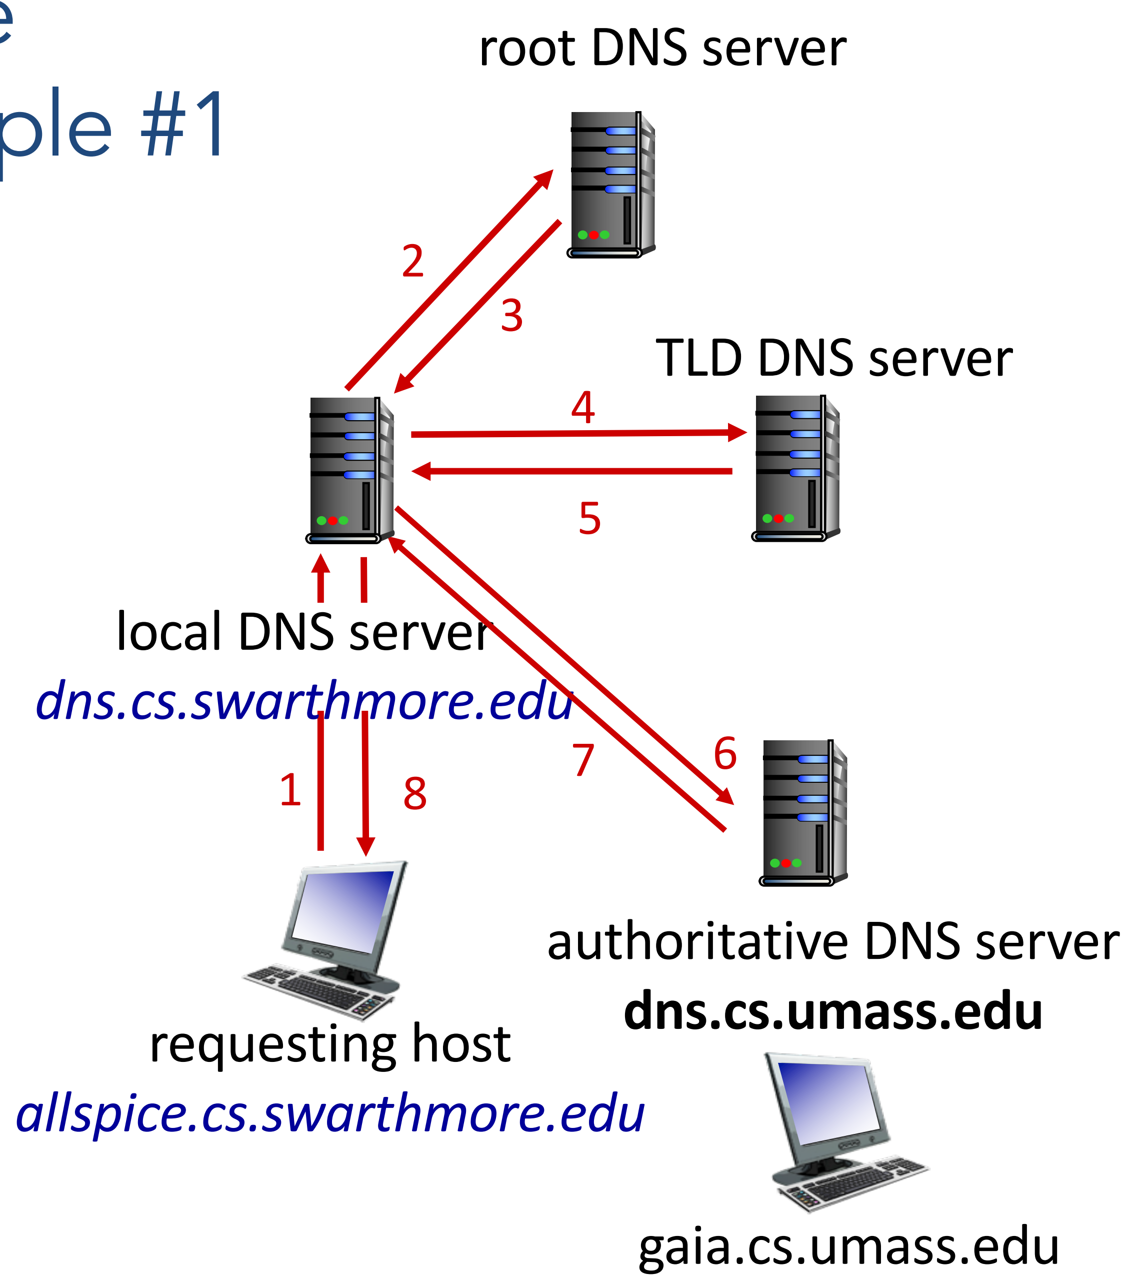

DNS name resolution example #2

recursive query:

- each server asks the next one, in a chain

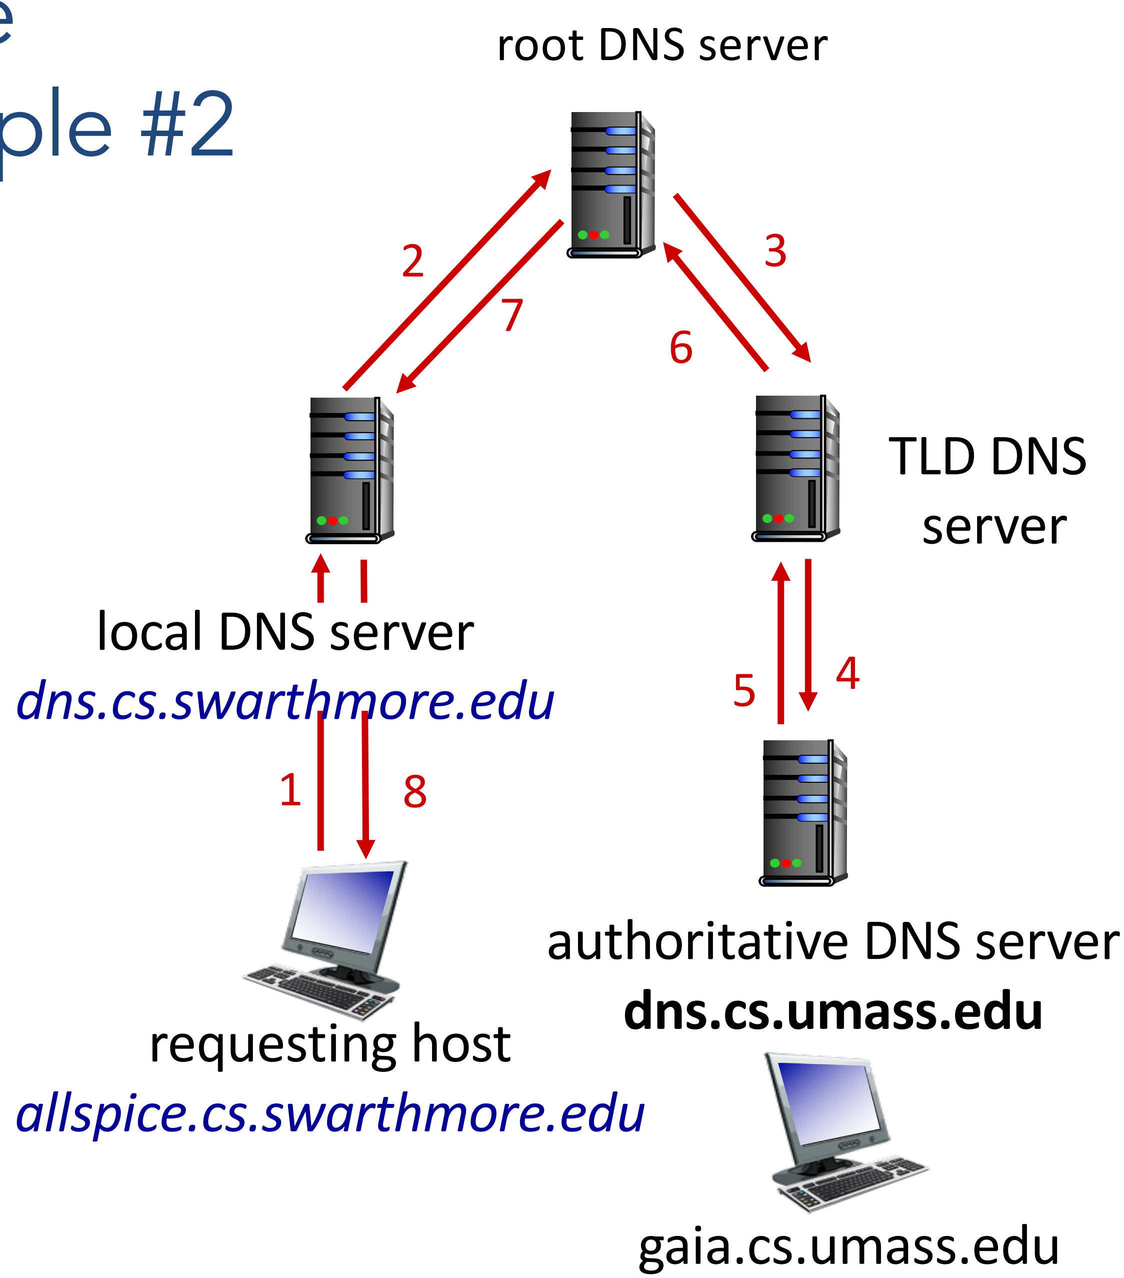

Which would you use? Why?

A. Iterative

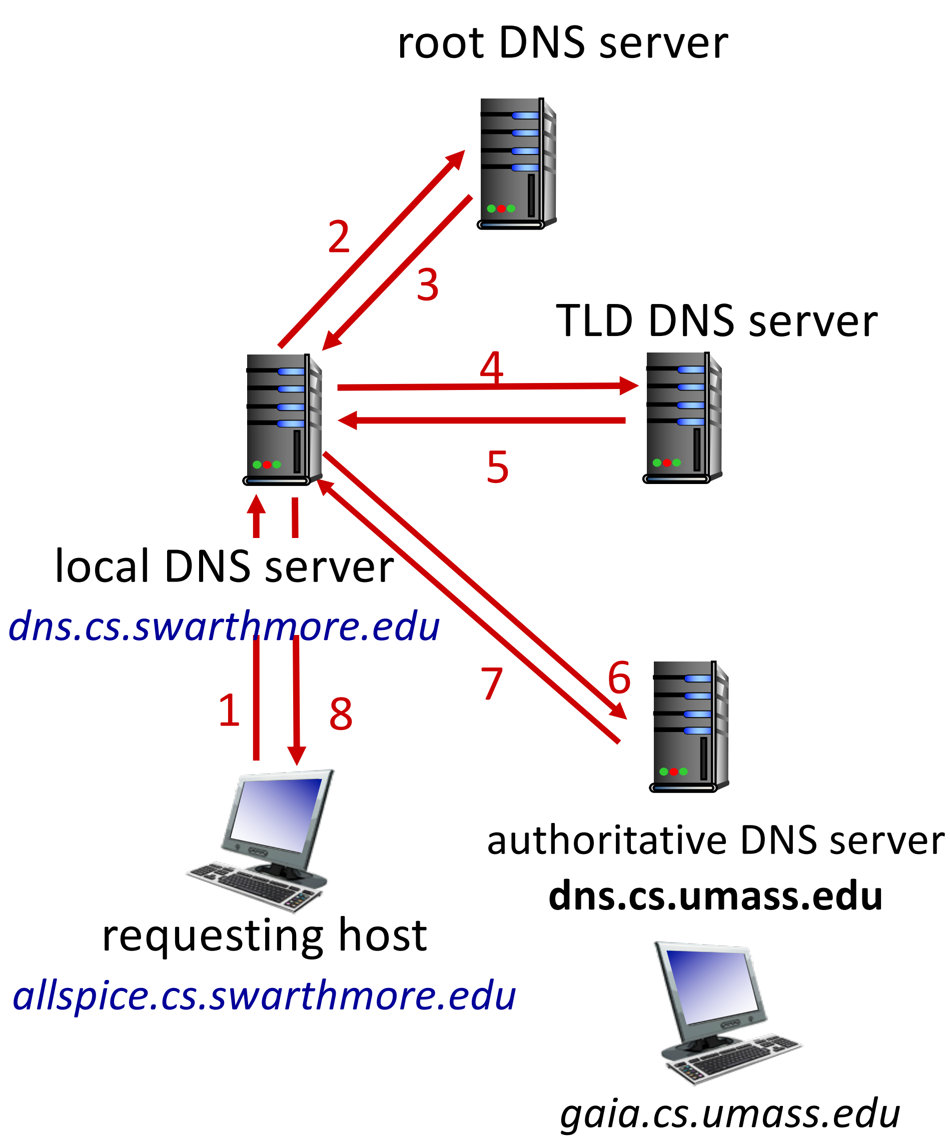

B. Recursive

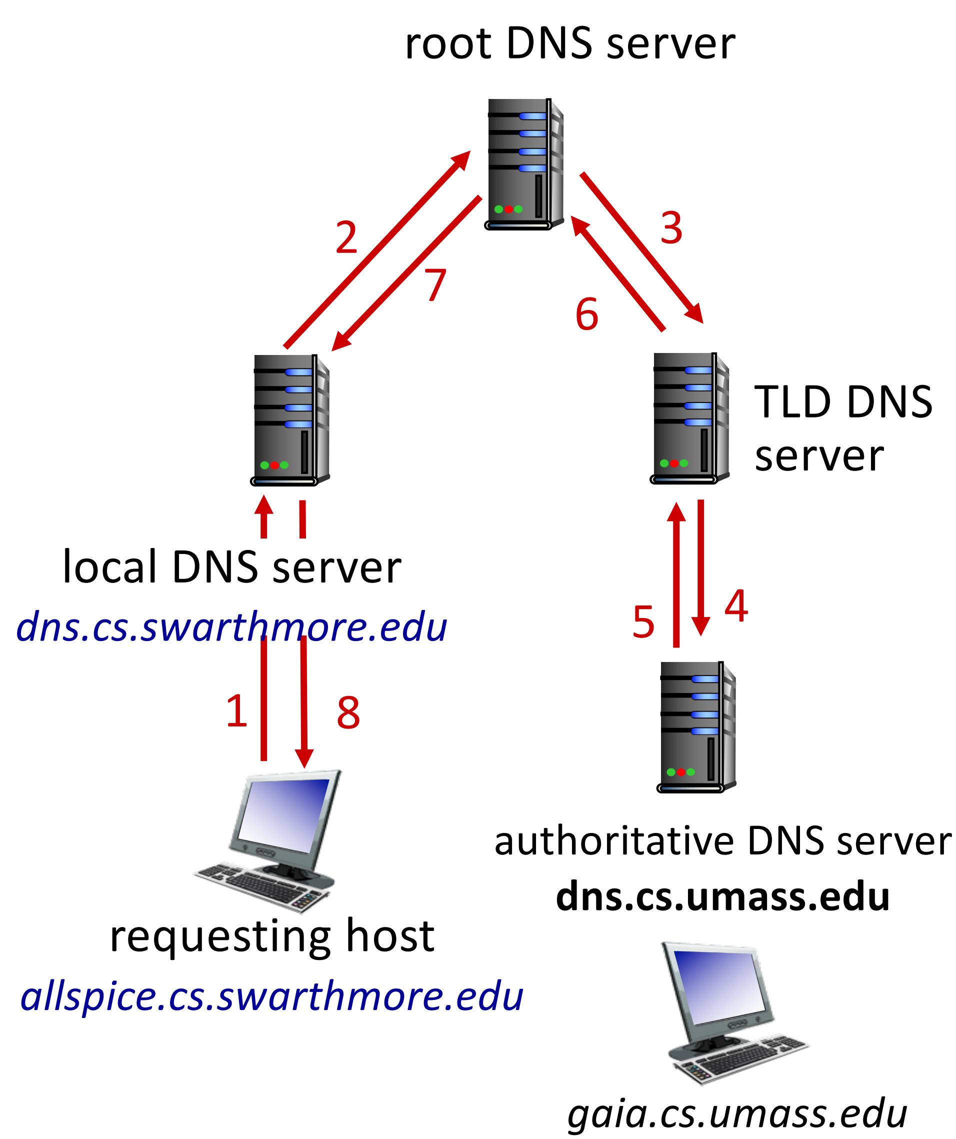

Which would you use? Why?

A. Iterative

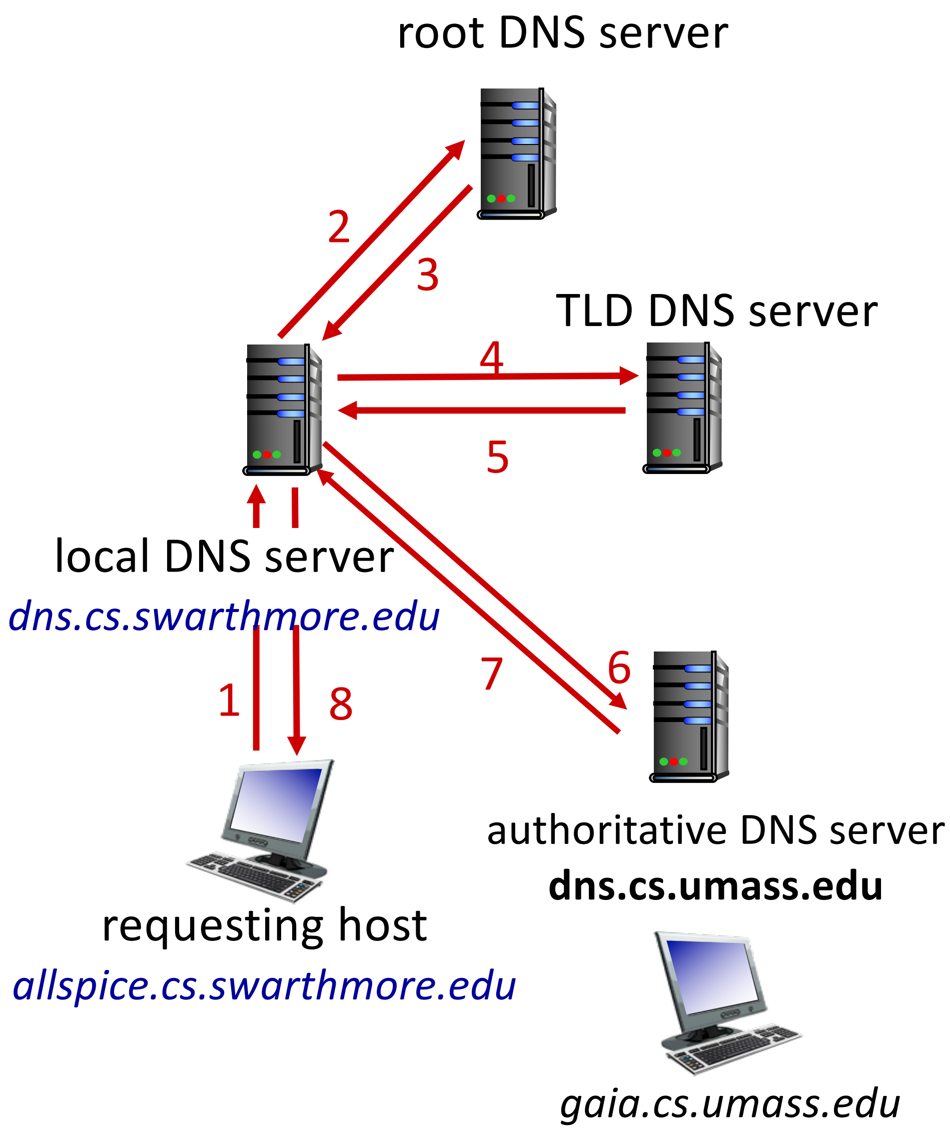

B. Recursive

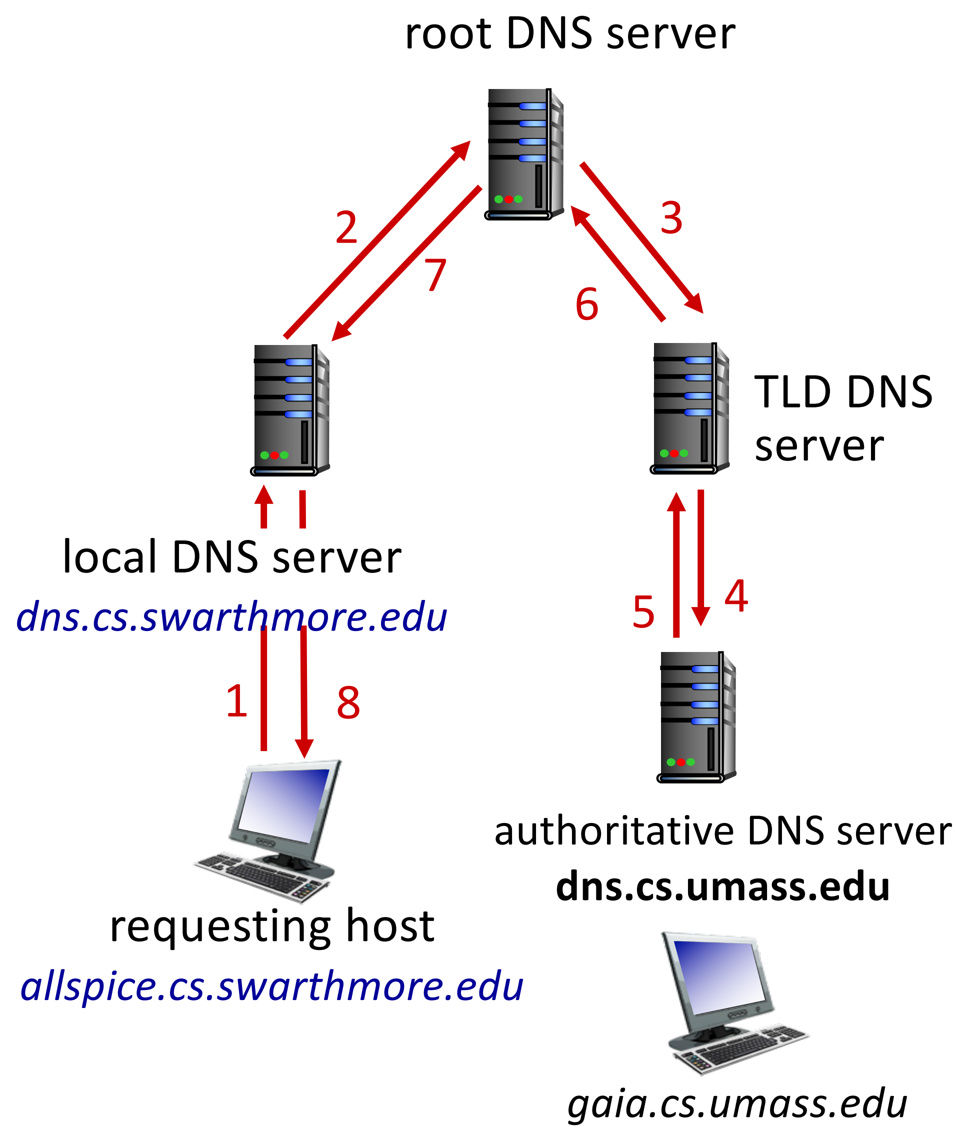

Summary

- 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