Application Layer

Raj Jain
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@wustl.edu

Audio/Video recordings of this lecture are available on-line at:

http://www.cse.wustl.edu/~jain/cse473-26/

Student Questions

©2026 Raj Jain



Overview

1. Network Application Architecture
2. HyperText Transfer Protocol (HTTP)
3. File Transfer and Email Protocols
4. Domain Name Service
5. Peer-to-Peer Applications

Note: This class lecture is based on Chapter 2 of the textbook (Kurose and Ross) and the figures provided by the authors.

Student Questions

- [Book Figure 2.1] In my own home, I hardwire ethernet devices into the same router that provides Wi-Fi for wireless devices. Why are they separate here?

Wi-Fi is a Layer 2 technology. A Wi-Fi device is usually a Layer 2 switch called an access point. A box may package any number of devices in one box.

- 1 Kb = 1000 bits or 1024 bits?

K=1024, k=1000 (lower case is lower). Computer bits are counted in K. Networking capacities are always in k. Storage is measured in K but sold in k. This is why 1TB disk is only 931.3 GB when installed.

©2026 Raj Jain



Network Application Architectures

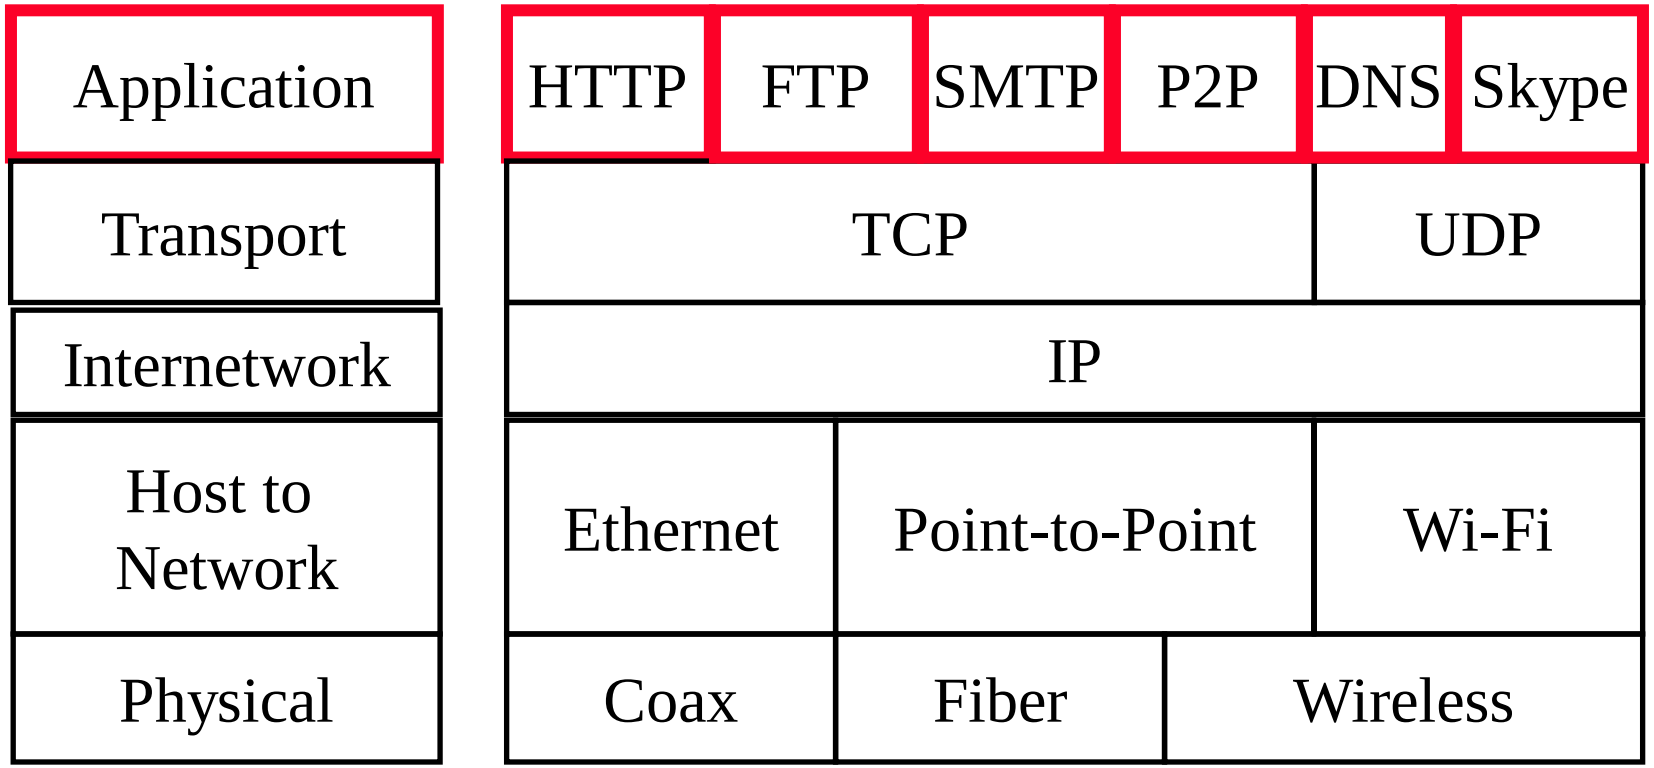
1. Protocol Layers
2. Client-Server vs. Peer-to-Peer
3. Process Communication
4. Names, Addresses, Ports
5. Transports

Student Questions

©2026 Raj Jain

Protocol Layers

- Top-Down approach



Student Questions

- Why Ethernet and point-to-point are paratactic?

All protocols of the same layer are paratactic. Paratactic=Placed next to each other

- Is it ok to assume that the transport layer uses the application layer components to run, the internetwork layer uses the transport layer, and so on?

No. The application layer uses the transport layer to send packets, and the transport layer uses the Internetwork layer.

- Do you mean the top box is the top layer, and the items at the bottom are the lower layer?

Yes.

©2026 Raj Jain

Network Application Architectures

- Client-Server
- Peer-to-Peer

Student Questions

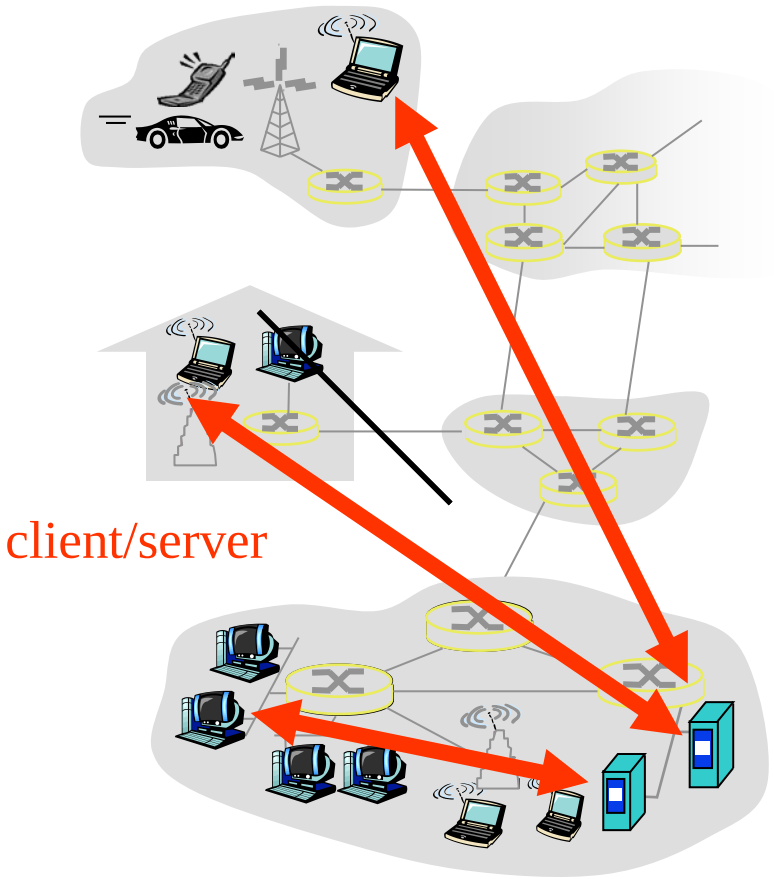
- Can one connection consist of both client-server and peer-to-peer links? As in one hop is client-server while the next is peer-to-peer.

The client is served by the server. Peer-to-peer does not imply any service. Both client-server and peer-to-peer connections can be 1-to-1, 1-to-many, many-to-1, or many-to-many.

©2026 Raj Jain

Client-Server

- Clients: Request service
- Server: Provides a service. Waits for clients
- The server is always up
- Clients do not communicate directly with each other
- Server = Data Center
- Example: Web Server, Search Engine, Social Networking



Student Questions

- There is a B/S Architecture. What is the difference between C/S and B/S? Can I consider B/S as a type of C/S?

B/S=Browser-Server
Browser is a client.
Therefore, yes,
B/S ⊆ C/S.

- Is client-server more secure than peer-to-peer?

Yes. Getting viruses from torrent websites, programs, and downloaded software is easy.

- If a client wants to communicate with another client in C/S architecture, should they send the packet to the server and let the server retransmit it?

No. In a client-to-client connection, the originator becomes the client, and the recipient is the server.

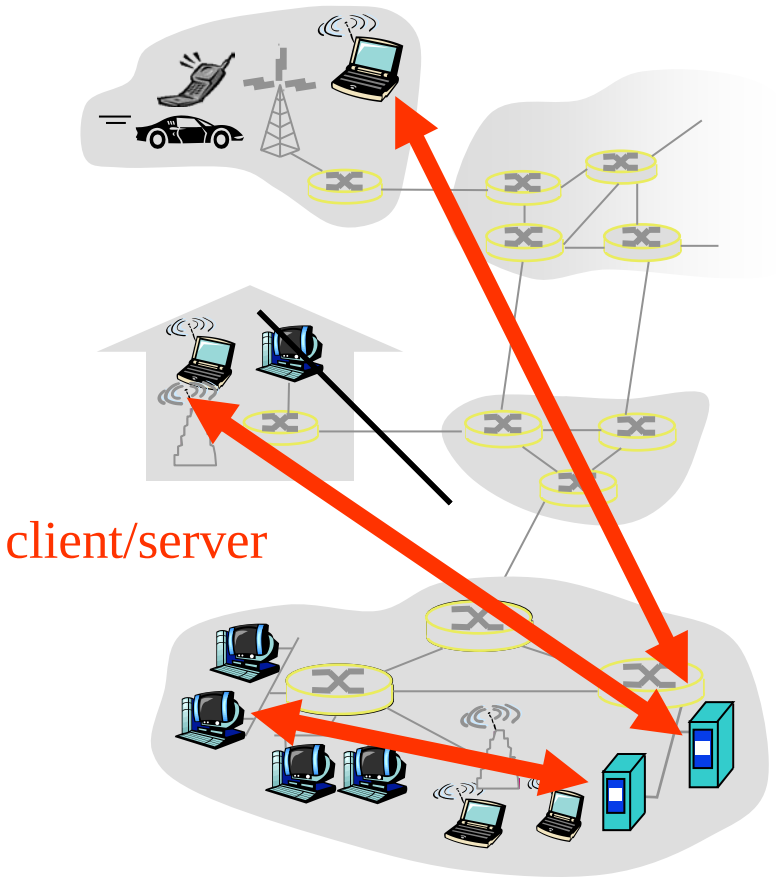
- Are there client-client or server-server connections directly?

Yes. These are called peer-to-peer, as shown in the next slide.

©2026 Raj Jain

Client-Server

- Clients: Request service
- Server: Provides a service. Waits for clients
- The server is always up
- Clients do not communicate directly with each other
- Server = Data Center
- Example: Web Server, Search Engine, Social Networking



Student Questions

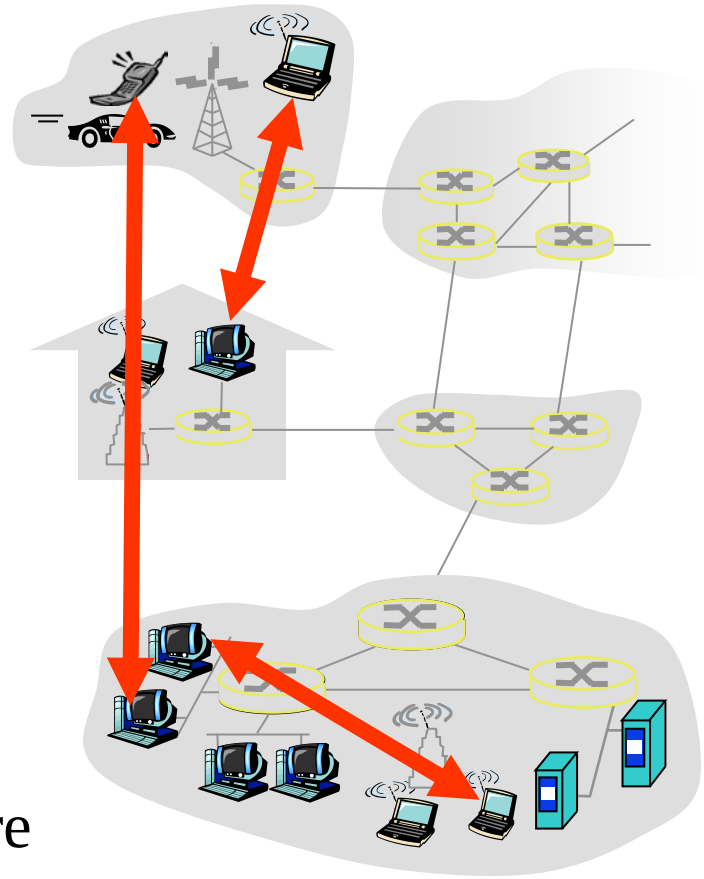
- Some streaming media companies use users' smartphones as edge devices to upload portions of video resources to other users. Is this the client-server model?

It is P2P.

©2026 Raj Jain

Peer-to-Peer

- Does not require always-on servers
- Hosts communicate directly
  ⇒ Peers
- Hosts may come on or may go off at any time
- Examples: File Sharing (Bit Torrent, eMule, LimeWire), Telephony (Skype)
- Highly scalable
- Highly symmetric traffic
  ⇒ ISP unfriendly
- Difficult to authenticate ⇒ Insecure
- Need incentives to share



Student Questions

- Does the IPFS p2p protocol solve the authenticity problem by hashing the files?

Don't know much about authenticity in IPFS. Hashing is not an authentication mechanism. It is an integrity mechanism. To be discussed in Chapter 8.

- When would client-server be more appropriate than P2P?

Web server, DNS, Mail Server, etc.

- Since web 3.0 is decentralized. Is it going to be a P2P network?

Web3=Web using blockchain

- Based on my understanding, is it correct to say that the Bitcoin network is an example of a peer-to-peer network?

Bitcoin does not have a network. It uses the Internet. Mining nodes are a kind of server. So not P2P.

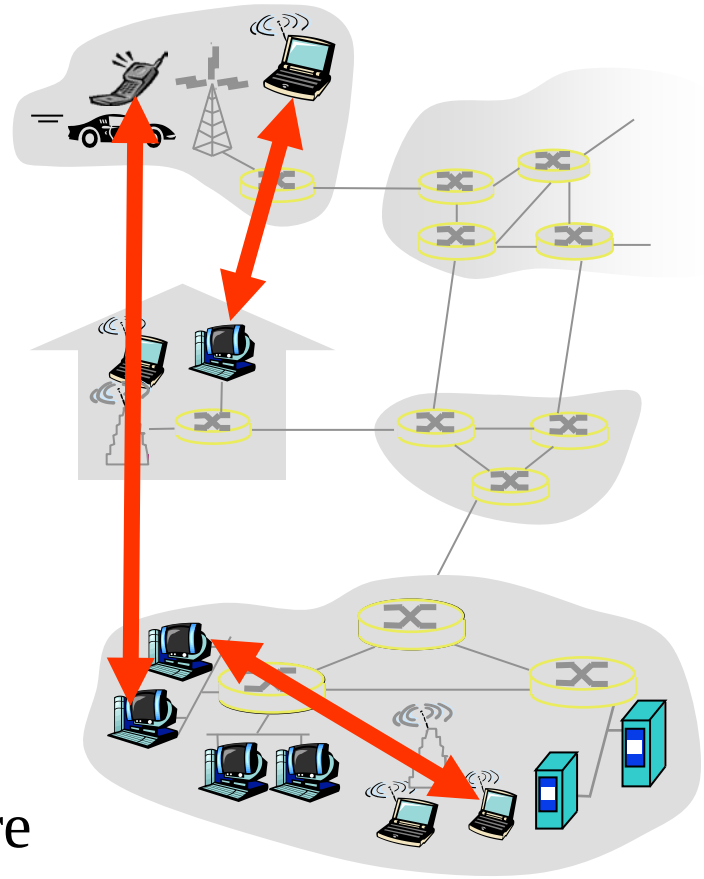
- Did P2P only serve file transfer?

No. You can make almost any other application.

©2026 Raj Jain

Peer-to-Peer

- Does not require always-on servers
- Hosts communicate directly
  ⇒ Peers
- Hosts may come on or may go off at any time.
- Examples: File Sharing (Bit Torrent, eMule, LimeWire), Telephony (Skype)
- Highly scalable
- Highly symmetric traffic
  ⇒ ISP unfriendly
- Difficult to authenticate ⇒ Insecure
- Need incentives to share



Student Questions

- Are there any applications that make use of both peer-to-peer and client-server?

File transfer

- The peer-to-peer seems more efficient than the client-server. Is that true?

Yes. See Slide 2-52

- How could torrenting and peer-to-peer be made safer?

Using authenticated peers

- Does Peer-to-Peer require a server (to see who is seeding or sending which part of the file etc.)?

The servers are used to discover seeders and leechers, but they do not keep track of the part of the file. If there are no servers, one could discover by broadcast, but that would result in too much traffic.

- Is it true that P2P does not involve any servers at all?

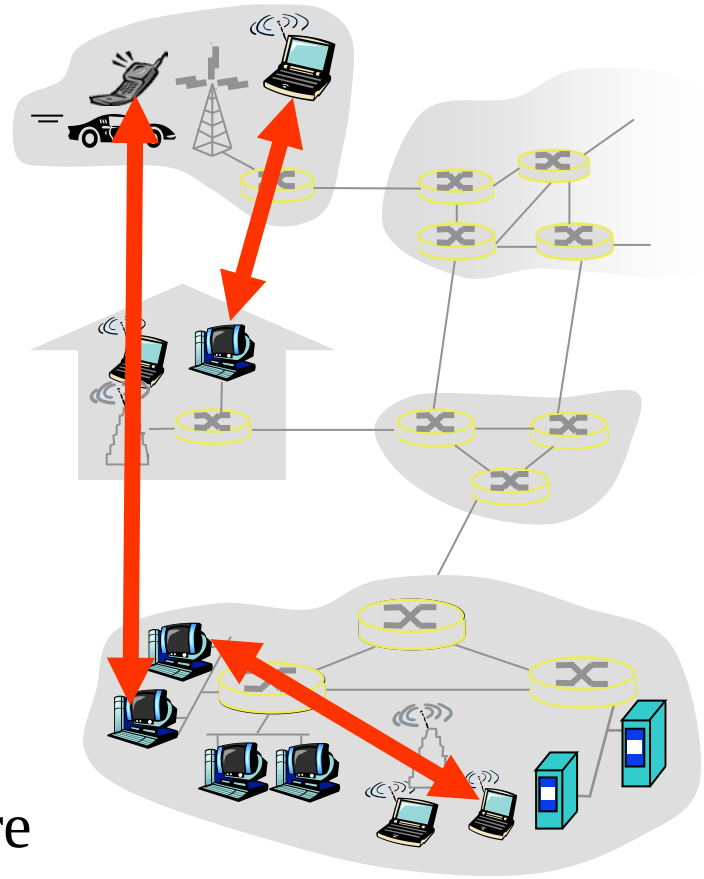
P2P=Anyone can serve anyone.

- Is zoom considered to be P2P or client-to-server?

Client to Server

Peer-to-Peer

- Does not require always-on servers
- Hosts communicate directly ⇒ Peers
- Hosts may come on or may go off at any time.
- Examples: File Sharing (Bit Torrent, eMule, LimeWire), Telephony (Skype)
- Highly scalable
- Highly symmetric traffic ⇒ ISP unfriendly
- Difficult to authenticate ⇒ Insecure
- Need incentives to share



Student Questions

- Why are ISPs throttling P2P traffic, and is there any way to get around it?

VPN

- If the host you want to receive information from is off, what can you do?

Try later. There are ways to turn it on remotely if it is your computer.

- Is P2P more scalable than client-server? Why?

It is proven later.

- In P2P architecture, the links between clients don't go through servers but through routers.

Yes.

- While downloading a file via torrent from peers, do I give my downloaded part to others?

Yes

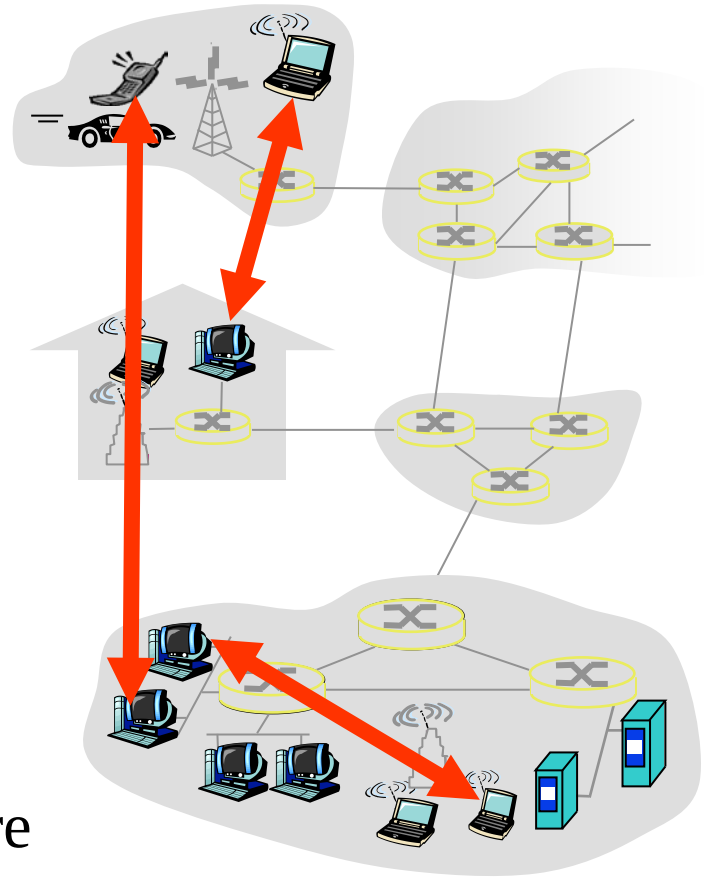
- If some part of the file could not be found in all currently available peers, do I have to wait for others to become online to download the entire file?

Yes.

©2026 Raj Jain

Peer-to-Peer

- Does not require always-on servers
- Hosts communicate directly
  ⇒ Peers
- Hosts may come on or may go off at any time.
- Examples: File Sharing (Bit Torrent, eMule, LimeWire), Telephony (Skype)
- Highly scalable
- Highly symmetric traffic
  ⇒ ISP unfriendly
- Difficult to authenticate ⇒ Insecure
- Need incentives to share



Student Questions

- Why is P2P symmetrical traffic not friendly for ISP?

ISP systems are designed for high download speed and low upload speed.

- Can you elaborate on exactly how the bottleneck appears in peer-to-peer?

Torrent (Peer-to-Peer) traffic requires uploading. Carrier networks are not designed for uploading.

- Would an example of a peer-to-peer server be like using Box or Duopush to be authenticated while not on a WashU Wi-Fi network?

No. Airdrop is an example of peer-to-peer. But does not go through the carrier network. So, OK.

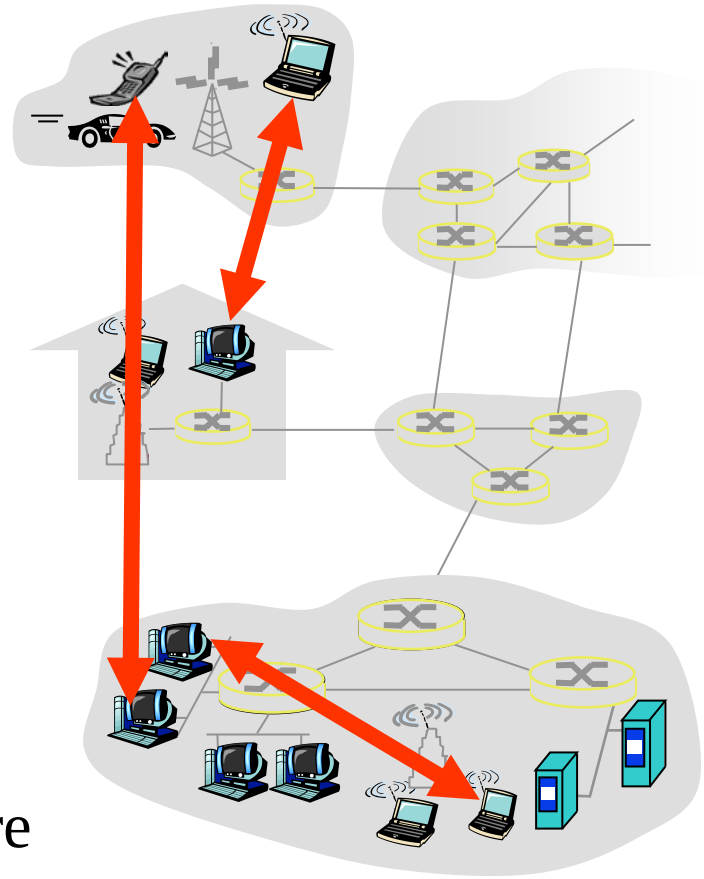
- How is Skype P2P, but Zoom is C-S?

In the old days, carrying voice calls was illegal unless you had a carrier license. So Skype was designed not to require a permanent server.

©2026 Raj Jain

Peer-to-Peer

- Does not require always-on servers
- Hosts communicate directly
  ⇒ Peers
- Hosts may come on or may go off at any time.
- Examples: File Sharing (Bit Torrent, eMule, LimeWire), Telephony (Skype)
- Highly scalable
- Highly symmetric traffic
  ⇒ ISP unfriendly
- Difficult to authenticate ⇒ Insecure
- Need incentives to share



Student Questions

- What do you mean by "Highly symmetric traffic"? And why is ISP unfriendly?

Symmetric: Upload=Download
Carrier networks are designed for low upload speed.

- Peer-to-Peer is illegal. Right?

No. Its use to distribute copyrighted material is illegal. Windows uses P2P for updates.

- Would it be possible to set up your own P2P network for calling/messaging and eliminate the need to rely on ISPs entirely? If so, how difficult would that be?

It's easy. Now, It is done frequently.

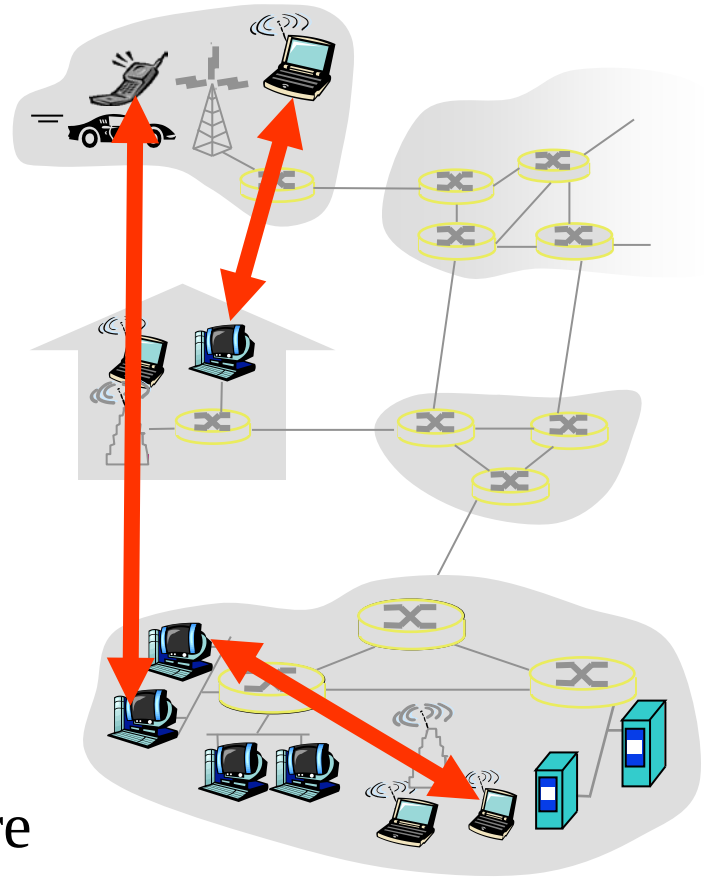
- What makes peer-to-peer networks more scalable?

Distributed serving

©2026 Raj Jain

Peer-to-Peer

- Does not require always-on servers
- Hosts communicate directly
  ⇒ Peers
- Hosts may come on or may go off at any time.
- Examples: File Sharing (Bit Torrent, eMule, LimeWire), Telephony (Skype)
- Highly scalable
- Highly symmetric traffic
  ⇒ ISP unfriendly
- Difficult to authenticate ⇒ Insecure
- Need incentives to share



Student Questions

- Why is peer to peer hard to validate and have there been any advances in making it more secure recently?

You are talking to new random nodes constantly.

- Is end-to-end encryption a type of peer-to-peer application?

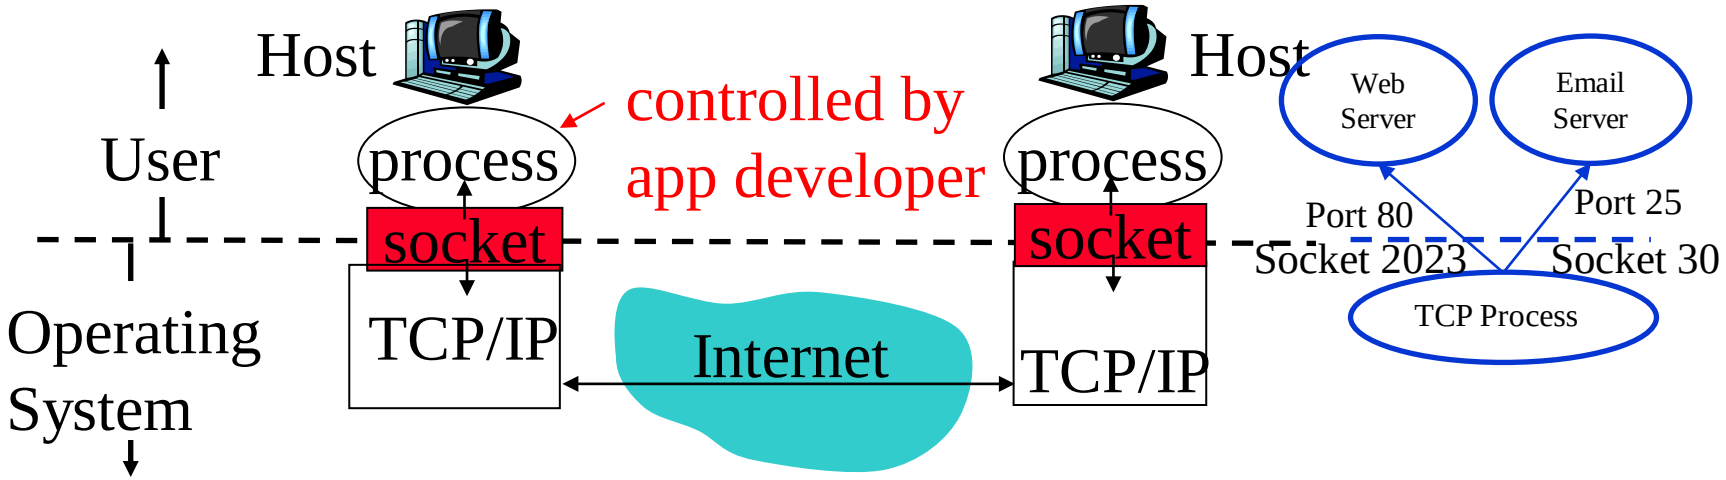
No, P2P means pieces of file are coming from many places. Entire file from one place is not P2P.

©2026 Raj Jain

Process Communications

- Inter-Process Communication on the Same Host
  ⇒ Operating system provides message passing.
- Unix provides an application programming interface called "sockets."
- Inter-Process Communication on Different Hosts
  ⇒ Network provides message passing.





Student Questions

- Can you explain again what ports are and how they are used?

Ports are Networking/TCP concepts. A socket is an OS/Unix concept. Ports are standardized so that applications can run on any computer. Sockets are not standard. Each process opens sockets and numbers them as needed. For example, port 80 is for the Web. Web running on different servers may use different sockets but the same port, 80.
See above.

- On pages 2-8, where is the port on that picture? In the lecture, you pointed to it, but we can't see where you pointed since this is just a screen recording.

A new picture has been added.

- Is the socket an SDU or PDU?

The socket is the Unix name for API

- How can we know which port of a server is open?

There are utilities/tools.

- Which layer is a socket in?

Socket is not a networking term.

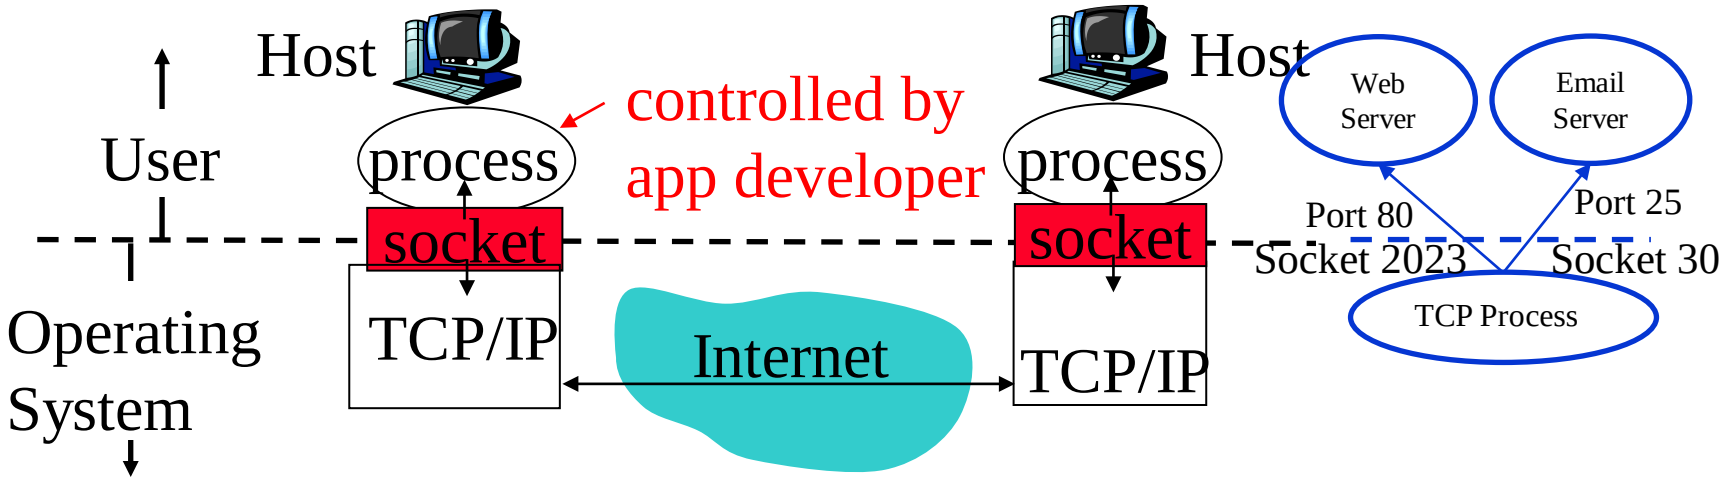
- Since Windows is not in the Unix family, does it have "sockets"? *Yes.*

©2026 Raj Jain

Process Communications

- Inter-Process Communication on the Same Host
  ⇒ Operating system provides message passing.
- Unix provides an application programming interface called "sockets."
- Inter-Process Communication on Different Hosts
  ⇒ Network provides message passing.





Student Questions

- Can we expose multiple applications to the same port? *No.*
- Are programming languages responsible for giving developers the socket interface to make web requests? If so, how do languages support different OSs?

Unix provides a Socket interface.

- Will the port open for TCP/IP protocols or applications or both? *Both. Any OS service.*
- I saw port 22 is always being used. Could you explain why port 22 is usually preferred?

Most apps used port numbers standardized by the Internet Engineering Task Force (IETF).

- Is the number of sockets the same as the computers?

No. The number of sockets is related to the number of interprocess links.

- Two processes communicating in the same host use socket; if they are in different hosts, the communication needs a port?

The socket is an operating system term. Port is a networking term.

- How is the port different from the sockets?

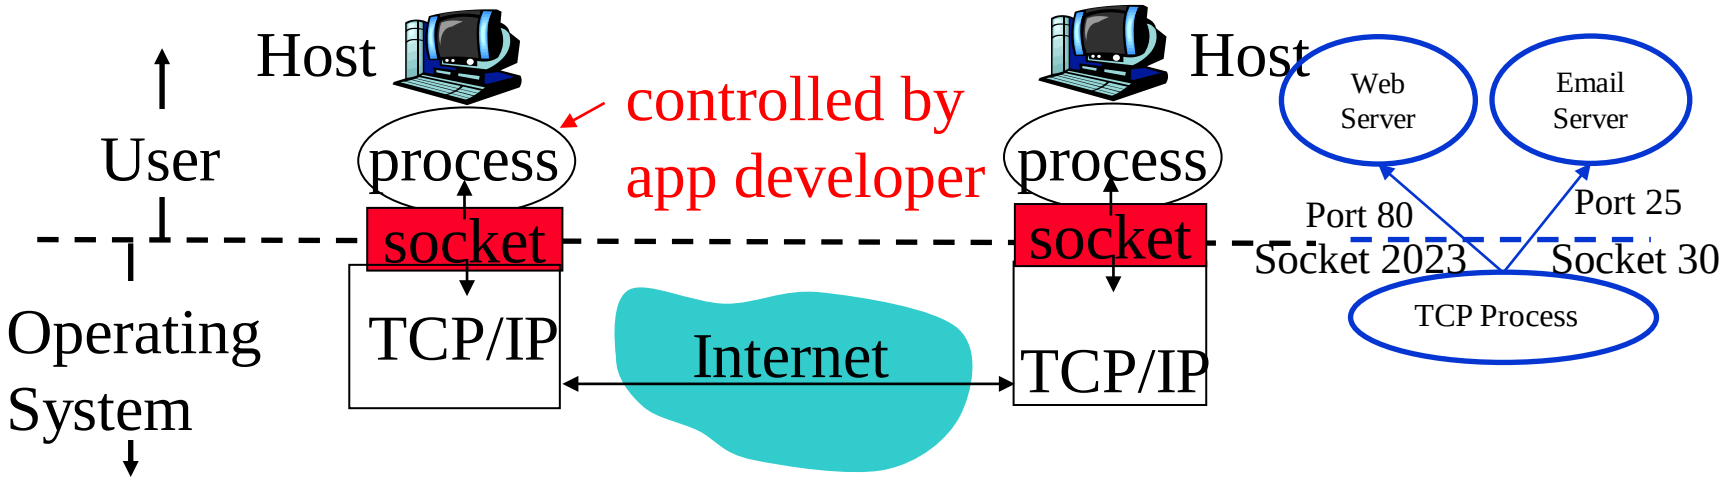
See above.

©2026 Raj Jain

Process Communications

- Inter-Process Communication on the Same Host
  ⇒ Operating system provides message passing.
- Unix provides an application programming interface called "sockets."
- Inter-Process Communication on Different Hosts
  ⇒ Network provides message passing.





Student Questions

- Are sockets used for communications between programs on the same computer?

Yes. But if the processes are on a different computer, they use the TCP process on each computer to communicate.

- Sockets used for communications between processes not in the same host system?

In that case, the sockets connect to a local networking (TCP) process, which provides the connection to the remote system.

- Can multiple websites transmit data to a client simultaneously through the same port?

Multiple hosts can send data to a server on a single port. This is quickly done using UDP. However, TCP requires a separate connection before data transmission.

- Are these Unix sockets similar to/identical to the web sockets?

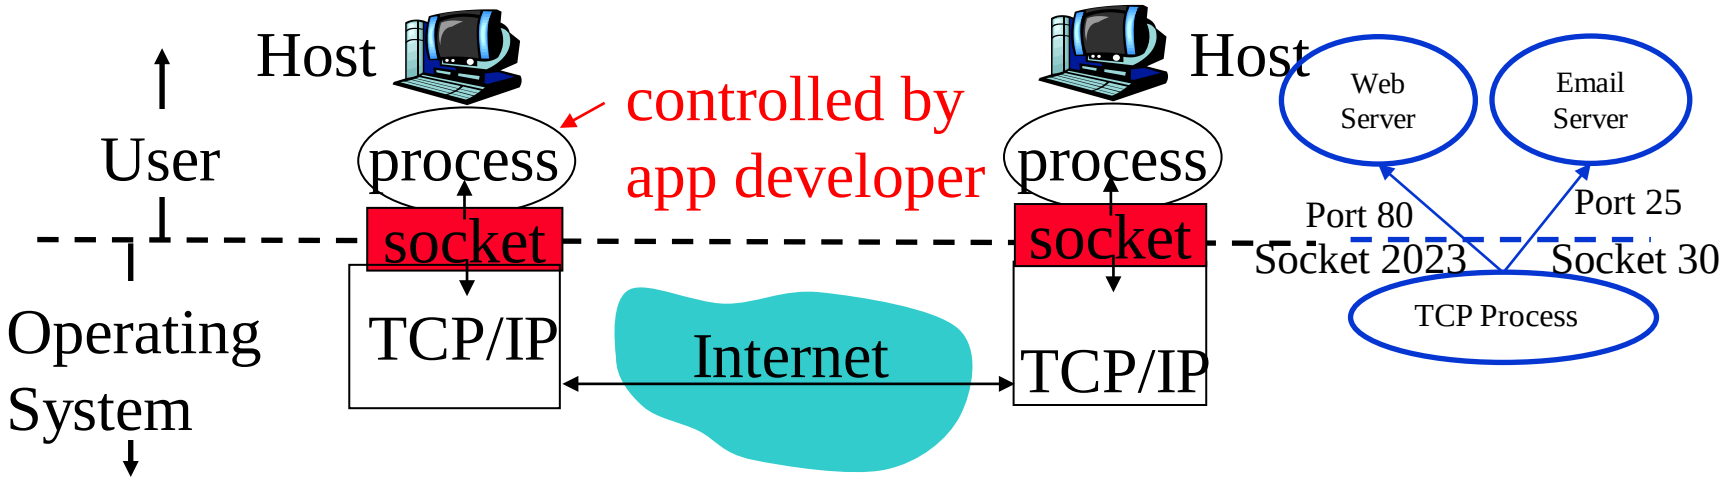
No. WebSocket was an application layer protocol designed by W3C. It is now an IETF standard specified in RFC 6455.

©2026 Raj Jain

Process Communications

- Inter-Process Communication on the Same Host
  ⇒ Operating system provides message passing.
- Unix provides an application programming interface called "sockets."
- Inter-Process Communication on Different Hosts
  ⇒ Network provides message passing.





Student Questions

- How do DNS servers translate domain names into IP addresses? Are there servers that manage some key-value pairing of domain names to IPs? Is that the sole purpose of DNS servers?

Yes. Yes. Yes.

- May I understand sockets and ports like this: Applications exchange messages using operating system sockets, and the Transport layer exchanges messages using ports?

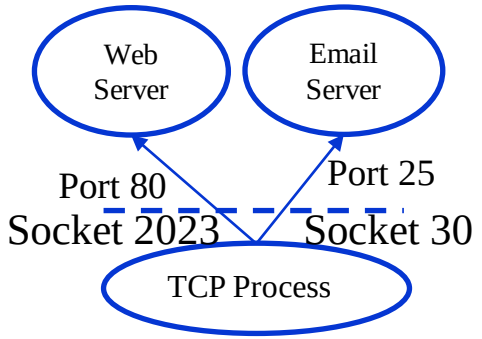
Yes. Instead of applications, I would say "processes," which could be application or system processes.

©2026 Raj Jain

Names, Addresses, Ports

- Domain Name System: www.google.com
- IP Address: 209.85.225.147
- 4 decimal numbers less than 256=8 bits each
  ⇒ 32-bits
- Ports: Entry point (Transport service access points)
- 21=FTP, 80=HTTP





Student Questions

- Are sockets only used for IPC? Does each process have its own fixed set of sockets, or is it a pool shared by processes? Are ports only used for inter-host communication (as opposed to IPC)?

Yes, sockets are used for inter-process communication. Even non-networking processes need and use sockets. Ports are exclusively for networking. The networking module on the computer knows which port is handled by which process on what socket.

- Are all the ports identical, other than the number?

No, each port provides a different service.

- You mentioned that we could change the default port for HTTP from 80 to whatever we want. Is that also true for other applications? For example, can we change the port for FTP from 21 to something else?

Yes. The clients must be told to connect to the new port numbers.

- How does a name turn into an IP address?

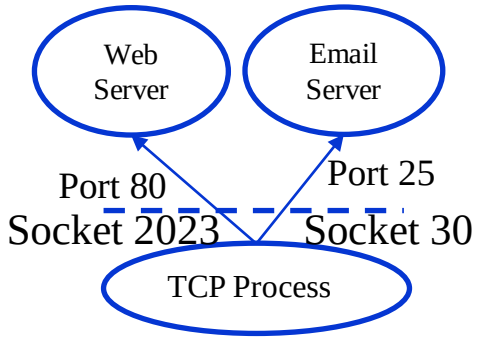
Names are easier to remember and static. Numbers are dynamic.

©2026 Raj Jain

Names, Addresses, Ports

- Domain Name System: www.google.com
- IP Address: 209.85.225.147
- 4 decimal numbers less than 256=8 bits each
  ⇒ 32-bits
- Ports: Entry point (Transport service access points)
- 21=FTP, 80=HTTP





Student Questions

- Who picked these commonly used port numbers? (80 for HTTP, 443 for HTTPS) *IETF*
- What are the differences between a port and a socket?

The socket is the Unix name for a port.

- What is the difference between using the default port vs. any other port numbers? Does the user have to specify the port at any time?
- Is there a predetermined number of ports on a computer? Does that number change between computers?

1 to 2^{16}-1

- Are all the ports exist virtually, or do they physically exist on chips or other hardware?

They are just numbers.

- What is the Physical Address (MAC) in our IP information? Is that identical to any computer in the world?

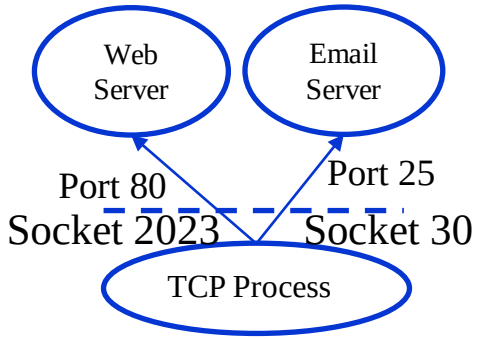
Discussed in Chapter 6.

©2026 Raj Jain

Names, Addresses, Ports

- Domain Name System: www.google.com
- IP Address: 209.85.225.147
- 4 decimal numbers less than 256=8 bits each
  ⇒ 32-bits
- Ports: Entry point (Transport service access points)
- 21=FTP, 80=HTTP





Student Questions

- The data are sent to the computer's port by socket, then sent to an external TCP/IP port?

Each port has a socket assigned to it.

- For HTTPS, isn't the standard port 443, and you can connect to that port directly instead of connecting to port 80 and then upgrading the security?

Sure. Both are OK.

- How can you connect to multiple websites on different tabs all through port 80?

Like party-line group calls.

- Are socket numbers allocated randomly or following a rule?

Socket numbers are generally sequential but can be random.

- Can multiple processes accept connections from the same port?

No. In each computer, each port is opened by only one process.

- Can a single process accept connections from multiple ports?

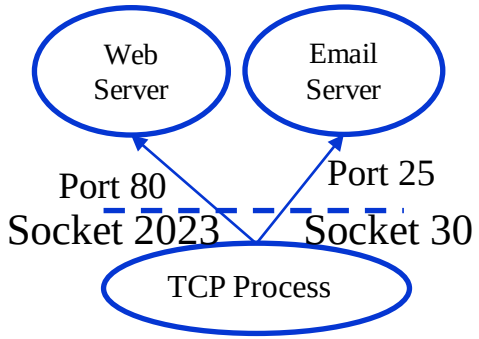
Yes.

©2026 Raj Jain

Names, Addresses, Ports

- Domain Name System: www.google.com
- IP Address: 209.85.225.147
- 4 decimal numbers less than 256=8 bits each
  ⇒ 32-bits
- Ports: Entry point (Transport service access points)
- 21=FTP, 80=HTTP





Student Questions

- Are port uses standardized? How was the standardization established?

Standardized by Internet Engineering Task Force (IETF)

- Is there any other special number for ports except for 21 and 80? What are they used for?

Each application has its own ports.

- In the picture of real ports shown in the slide, they're connected by land. Is each port connected somewhere?

Yes, it is connected to a process.

- Is "opening the port" the equivalent of having a process listening to that port?

Yes.

- Are there any additional security procedures on the operating system's level?

No. But the process can secure its services.

- What is the difference between a physical address and an IP address?

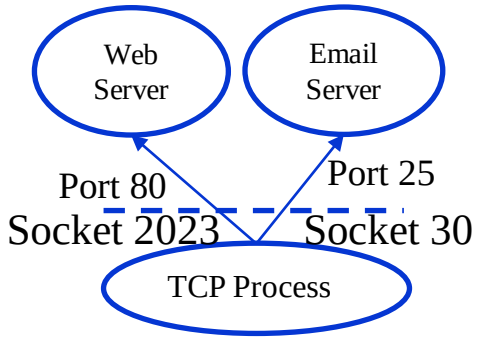
Physical: Room 23, top rack
IP: 209.85.225.147

©2026 Raj Jain

Names, Addresses, Ports

- Domain Name System: www.google.com
- IP Address: 209.85.225.147
- 4 decimal numbers less than 256=8 bits each
  ⇒ 32-bits
- Ports: Entry point (Transport service access points)
- 21=FTP, 80=HTTP





Student Questions

- Since all tabs from HTTP go through port 80, will having multiple tabs slow things down?

Yes. Each tab is a sub-process (thread) and uses resources.

- What to do if all ports have been taken up?

There are 65535 ports.

- Can protocols like HTTP and FTP have to be on ports 80 and 21, respectively, or is this just convention?

These are defaults only. You can use any other port.

- Is there a way to increase the number of ports you have?

Max 2^{16}-1=64000 ports

- How does browser determine what port it should use (i.e. HTTP vs HTTPS)?

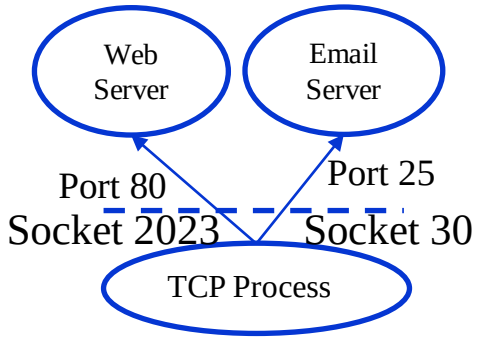
Some servers will accept only HTTPS connections.

©2026 Raj Jain

Names, Addresses, Ports

- Domain Name System: www.google.com
- IP Address: 209.85.225.147
- 4 decimal numbers less than 256=8 bits each
  ⇒ 32-bits
- Ports: Entry point (Transport service access points)
- 21=FTP, 80=HTTP CHECKSUM





Student Questions

- Can your IP Address change, or is it set for all devices and servers?

For each device, network manager can decide whether its IP address changes or remains fixed.

- If IP address changes, does the information in the hosts file also change automatically?

No. IP addresses are found by DNS. Whose entries change automatically. Host file is only a small part of DNS.

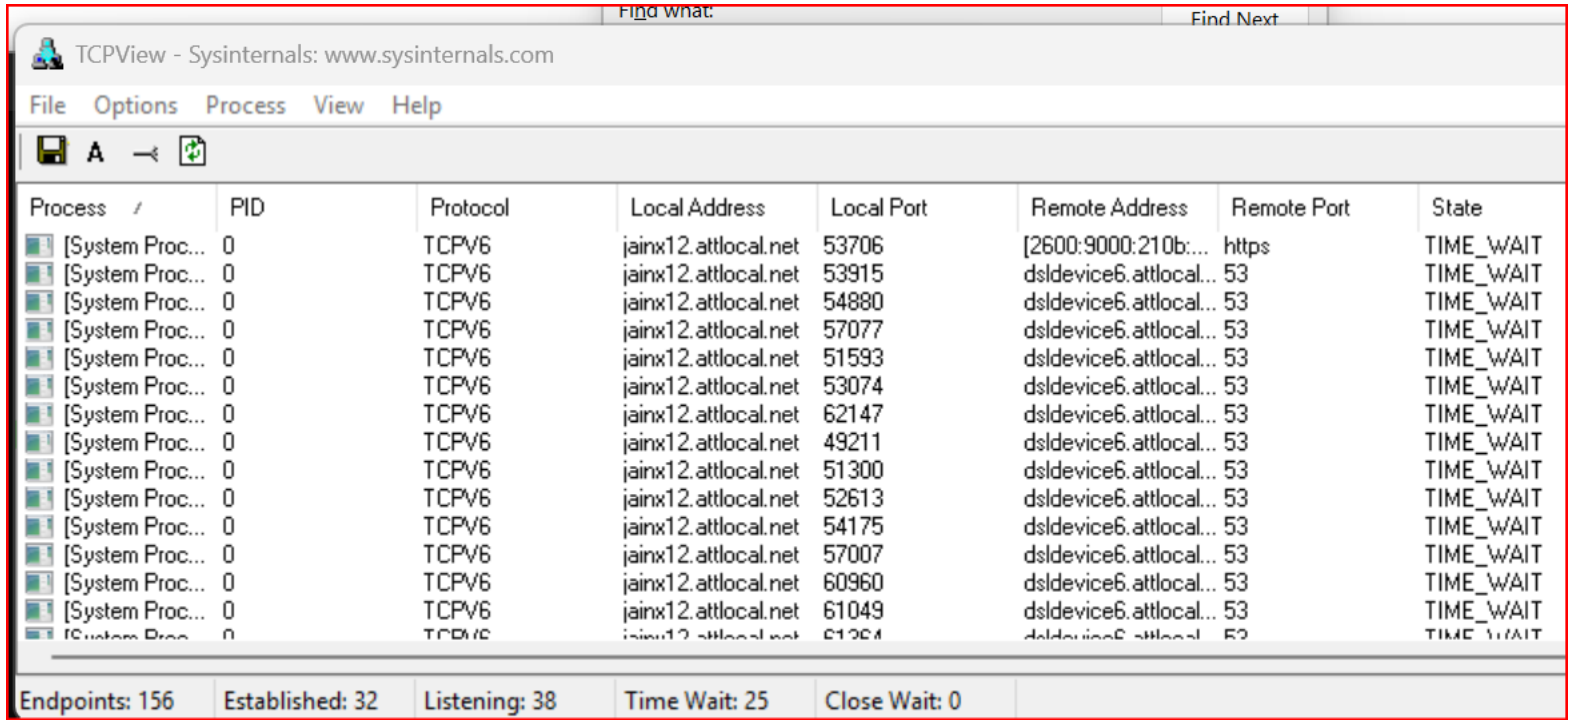
- Why can the same host run multiple services at the same time, but they require different port numbers?

Different services use different port numbers.

©2026 Raj Jain

Network Utilities

- TCPview: Shows active ports on your system
  https://docs.microsoft.com/en-us/sysinternals/downloads/tcpview



Student Questions

©2026 Raj Jain

Transports

TCP	UDP
Reliable data transfer	Unreliable Data Transfer
Packet Sequence # required	No Sequence #s
Every packet is acked	Not Acked
Lost packets are retransmitted	No Retransmission
May cause long delay	Quick and Lossy
Connection-oriented service	Connection-less Service
Good for Reliable and delay-insenstive applications	Good for loss-tolerant and delay sensitive applications
Applications: email, http, ftp, Remote terminal access	Telephony, Streaming Multimedia

Student Questions

- Will TCP wait to transmit packet #6 until receipt of packet #5 is acknowledged?

No. We will discuss the details in Chapter 3.

- What makes TCP reliable for data transfer?

Discussed in Chapter 3

- Are transports a single layer up from the physical media used for communication?

Transports are layer 4
Physical media is layer 1.

- How exactly do transports work?

Discussed in Chapter 3

- Does connection-oriented mean the destination for the packets is the same, but the path for each packet to get to the destination might be different?

Connection ⇒Setup before talking

- Do calling applications that use the internet (like Discord) use TCP or UDP?

UDP

- Is the TCP communicating by sending PDUs?

Yes.

- What exactly is acked?

Receipt of the packet #x

©2026 Raj Jain

Transports

TCP	UDP
Reliable data transfer	Unreliable Data Transfer
Packet Sequence # required	No Sequence #s
Every packet is acked	Not Acked
Lost packets are retransmitted	No Retransmission
May cause long delay	Quick and Lossy
Connection-oriented service	Connection-less Service
Good for Reliable and delay-insenstive applications	Good for loss-tolerant and delay sensitive applications
Applications: email, http, ftp, Remote terminal access	Telephony, Streaming Multimedia

Student Questions

- What happens if TCP packets arrive out of order

They may be cached or dropped.

- Can TCP and UDP be used together in the same application?

Yes.

- Can UDP transmit more data than TCP in same time since it doesn't ack packets?

Yes. However, some data may not make it.

- What is the difference in delay between TCP and UDP?

Retransmissions may delay TCP transfers.

- What do you mean by “Every packet is acked"?

The destination sends an acknowledgment that it has received the n^{th} packet.

- The book says streaming multimedia (YouTube) uses TCP nowadays. Why TCP if UDP is better for streaming?

If you design both ends of a connection, you can design it delay-tolerant by receiving and storing it early.

©2026 Raj Jain

Application Layer Protocols

- HTTP: HyperText Transfer Protocol
- FTP: File Transfer Protocol
- SMTP: Simple Mail Transfer Protocol
- DNS: Domain Name Server (Control Plane Application)
- P2P: Peer-to-Peer Applications (Class of applications)
- Skype
- Each application has its own protocol, message format, the semantics of fields.

Student Questions

- Is the control plane mentioned here the same as an SDN control plane?

Yes. But the Control plane existed much before SDN.

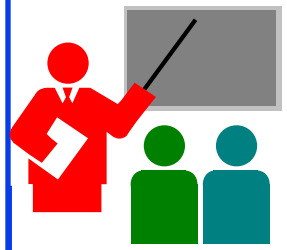
- What do you mean by DNS is a control plane application?

It is done by the network and not directly by the users.

- How Skype is an application layer protocol? Is it like Zoom or Discord?

Yes, Zoom/Discord/WhatsApp offer better user interface.

©2026 Raj Jain



Application Arch: Summary

1. P2P applications are **more scalable** than client-server
2. Applications exchanges messages using operating system **sockets**
3. Applications communicate using host **names**, **addresses**, and **ports**
4. Applications use transports: **TCP**, **UDP**, …
5. TCP is used for **reliable** communication
   UDP for **loss-tolerant delay-sensitive** applications

Student Questions

- UDP is unreliable. Does that mean the APPs using UDP are unreliable?

No. Sometimes it is better to lose some packets than wait for everything, e.g., video.

- What is the difference between a segment and a datagram?

TCP PDUs are called Segments and IP PDUs are called datagrams. Shown in Slide 1.43a.

Ref: Read Section 2.1 full. Try R1-R9

©2026 Raj Jain



HTTP

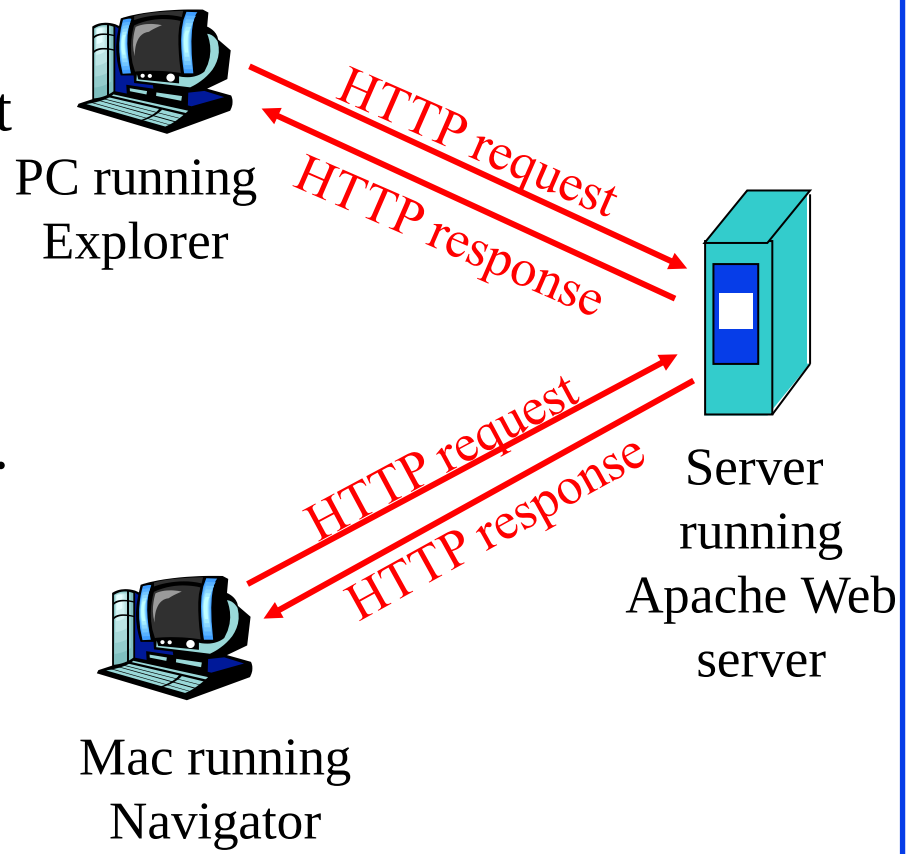
1. Concepts
2. Sample Web Page
3. HTTP Messages
4. Cookies
5. Proxy Servers
6. Conditional GET

Student Questions

©2026 Raj Jain

HTTP Concepts

- **Client**=Browser, e.g., Internet Explorer, Firefox
- **HTTP Server**, e.g., Microsoft Internet Information Service (IIS), Apache
- **Web Page**=Group of objects
- **Object**=Text, Images, files, ...
- **URL**: Uniform Resource Locator http://www.cse.wustl.edu/~jain/cse473-09/sample.htm



Student Questions

- When building a web page, Apache can be downloaded to a device in order to serve the web page. Does this download establish a connection?

Apache is an open-source web server software. Firefox is an example of an open-source web client.

- If the server does not respond, what would I see on the Explorer?

Timeout or unreachable

- What is the difference between URI and URL?

URI=Universal Resource Identifiers
URL=Universal Resource Locators

- Can sockets be apart physically?

Sockets are used for communications inside one system.

- Is the URL beyond cse.wustl.edu.... (e.g.,/~jain/cse473-09/sample.html) also done with a DNS page, or is it within the group of Web Pages such as in Apache Web Server.

DNS deals only with the host name and location. The operating system does rest.

©2026 Raj Jain

HTTP

- Uses TCP
- **Stateless**: The server does not remember the previous history
- **Non-Persistent**: Open a new TCP connection, get one object, and close
- **Persistent**: Open one TCP connection, get all objects, and close The server leaves the connection open after sending an object and closes on the timeout.
- Web pages are written in HyperText Markup Language (**HTML**)

Student Questions

- Can you explain what Persistent and Non-Persistent are?

Non-Persistent = Connection is closed immediately after a query has been answered. The next query will require a new TCP connection.

- For persistent HTTP: even if the client has closed their end of the connection, does the server still see the connection as open prior to timeout? If one end closes the connection, do both ends see it? Is it the client or the server that decides whether the TCP connection is persistent or non-persistent?

Closing a TCP connection will be discussed in Chapter 3. Any side can tell the other side and close a connection.

- Does this mean that HTTP can be switched between Persistent and Non-Persistent?

Yes. Originally it was non-persistent. Now it is mostly persistent.

- Does HTTPS have states since it is secure?

Yes.

©2026 Raj Jain

HTTP

- Uses TCP
- **Stateless**: The server does not remember the previous history
- **Non-Persistent**: Open a new TCP connection, get one object, and close
- **Persistent**: Open one TCP connection, get all objects, and close The server leaves the connection open after sending an object and closes on the timeout.
- Web pages are written in HyperText Markup Language (**HTML**)

Student Questions

- In general, what happens during the process of the server forgetting its previous history (becoming stateless)?

It will not remember your name, address, or credit card numbers.

- Is non-persistent safer than persistent?

Security is independent of persistence.

- Can you give some more examples of non-persistent and persistent HTTP?

Persistent=Get many objects
Non-persistent=Get one object

- What's the difference between HTTP and HTTPS?

HTTPS is secure

- Do all web page transfers use persistent TCP?

Yes. Non-persistent is old.

- What happens on the server side when you access a website and your login is saved? How does it remember the user if the connection is stateless?

Cookies are discussed on Slide 2.23.

- How does HTTP's stateless protocol harmonize with persistent connections for efficient web content delivery?

HTTP 1.1 provides persistent connection.

©2026 Raj Jain

HTTP

- Uses TCP
- **Stateless**: The server does not remember the previous history
- **Non-Persistent**: Open a new TCP connection, get one object, and close
- **Persistent**: Open one TCP connection, get all objects, and close The server leaves the connection open after sending an object and closes on the timeout.
- Web pages are written in HyperText Markup Language (**HTML**)

Student Questions

- If HTTP messages over the internet are packet-switched, what is the benefit of opening a virtual persistent "connection?

So that various packets can be sent to the same process in the computer.

- How do persistent HTTP servers protect against DoS attacks without making timeouts too short?

When a server receives many requests quickly, it denies them.

- If we had computing resources, would making HTTP stateful advantageous?

Most applications use only stateful HTTP connections.

- What occasions are suitable for persistent and non-persistent modes respectively?

Persistent is useful for multiple related transactions.

- How does persistent HTTP improve performance compared to non-persistent HTTP?

Avoids 4-way handshakes for each object on the page.

©2026 Raj Jain

HTTP

- Uses TCP
- **Stateless**: The server does not remember the previous history
- **Non-Persistent**: Open a new TCP connection, get one object, and close
- **Persistent**: Open one TCP connection, get all objects, and close
  The server leaves the connection open after sending an object and closes on the timeout.
- Web pages are written in HyperText Markup Language (**HTML**)

Student Questions

- When would you still use non-persistent connections?

Simple queries expecting just one object.

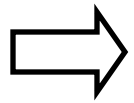
- Are the Non-Persistent and Persistent modes of HTTP switched dynamically or statically?

The application decides.

©2026 Raj Jain

Sample Web Page

```
<HTML>
<HEAD>
</HEAD>
<BODY>
<img src=jain.jpg>
<BR>
Raj Jain
</BODY>
</HTML>
```





Raj Jain

Student Questions

- How much HTML experience do we need for this course?

Not much.

- Can one HTTP session have multiple open TCP connections?

They generally do.

- Why does persistent HTTP need to close on timeout? Is it due to issues with too many people trying to connect to the server?

No, the client may have walked out. We need the resources that were committed for the client.

- What kinds of HTTP connections are non-persistent, and what kinds are persistent?

Very few elements on a page

⇒ *Nonpersistent is OK*

©2026 Raj Jain

Sample HTTP Request Message

GET /~jain/cse473-16/sample.htm HTTP/1.1
Host: www.cse.wustl.edu
Connection: close
User-agent: Mozilla/4.0
Accept-Language: en

- **Method** = Get
- **Object** = /~jain/cse473-16/sample.htm
- **Host** = www.cse.wustl.edu
- **Version** = HTTP/1.1
- **Header Fields** = Host, Connection, User-agent, …

Student Questions

- Does a URL never include www.cse.wustl.edu (using this slide as an example)?

URL consists of two parts: Host and Directory. These are separately indicated in HTTP. The host is case-insensitive. The directory is case-sensitive.
URL=Uniform Resource Locator.
Here directory is the URL/location on the host.

- It seems like there are many versions of HTTP, like 0.9,1.0,1.1, and 2.0. What would happen if I used another version than 1.1? Can my browser still interpret this response? What stops people from using the latest version besides habits?

Servers and Browsers try to be up-to-date and backward compatible.

- Can I confuse the User-agent with the wrong parameters? For example, I use Google Chrome to browse, but in the message, I send Edge.

Yes. You can do that. The response may be customized for Edge.

©2026 Raj Jain

Sample HTTP Request Message

GET /~jain/cse473-16/sample.htm HTTP/1.1
Host: www.cse.wustl.edu
Connection: close
User-agent: Mozilla/4.0
Accept-Language: en

- **Method** = Get
- **Object** = /~jain/cse473-16/sample.htm
- **Host** = www.cse.wustl.edu
- **Version** = HTTP/1.1
- **Header Fields** = Host, Connection, User-agent, …

Student Questions

- In the HTTP request message part, I see the URL does not include host name, so does URL only indicate directory name?

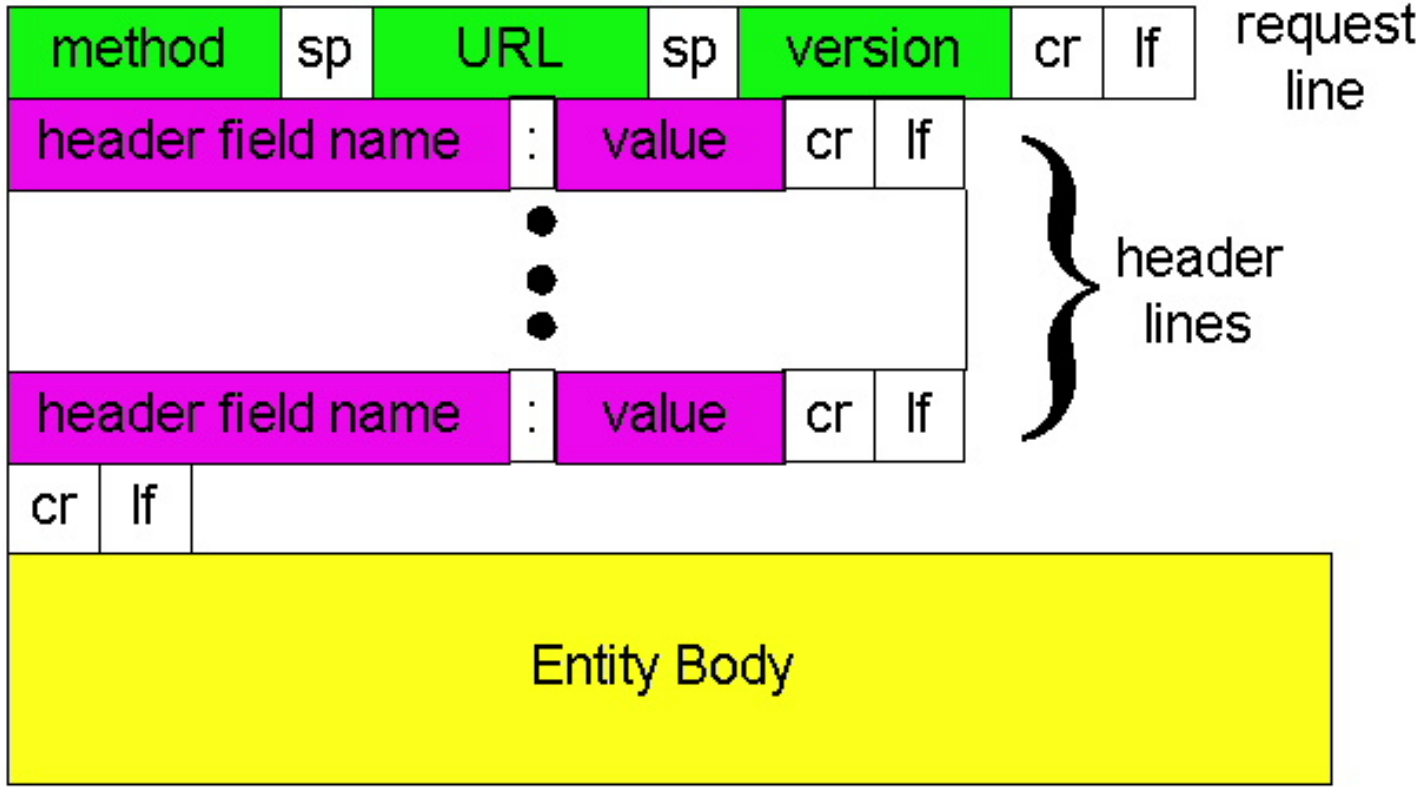
The host is in the next line.

- Where do we type these kind of HTTP request messages? In a browser or in CMD?

It used to work with Telnet. Blocked on most servers now.

©2026 Raj Jain

HTTP Request Message Format



Student Questions

- Why hasn't there been a push to remove legacy characters such as CR and LF from the HTTP Messaging Format?

You can manually type the whole request using a keyboard. This is a feature, not a bug.

- Is there a default if a version is not provided?

No. If a version-specific feature is used, it is not clear what the server/browser will do.

- Why does the HTTP request message format require "cr", "lf", or "sp" to separate the information?

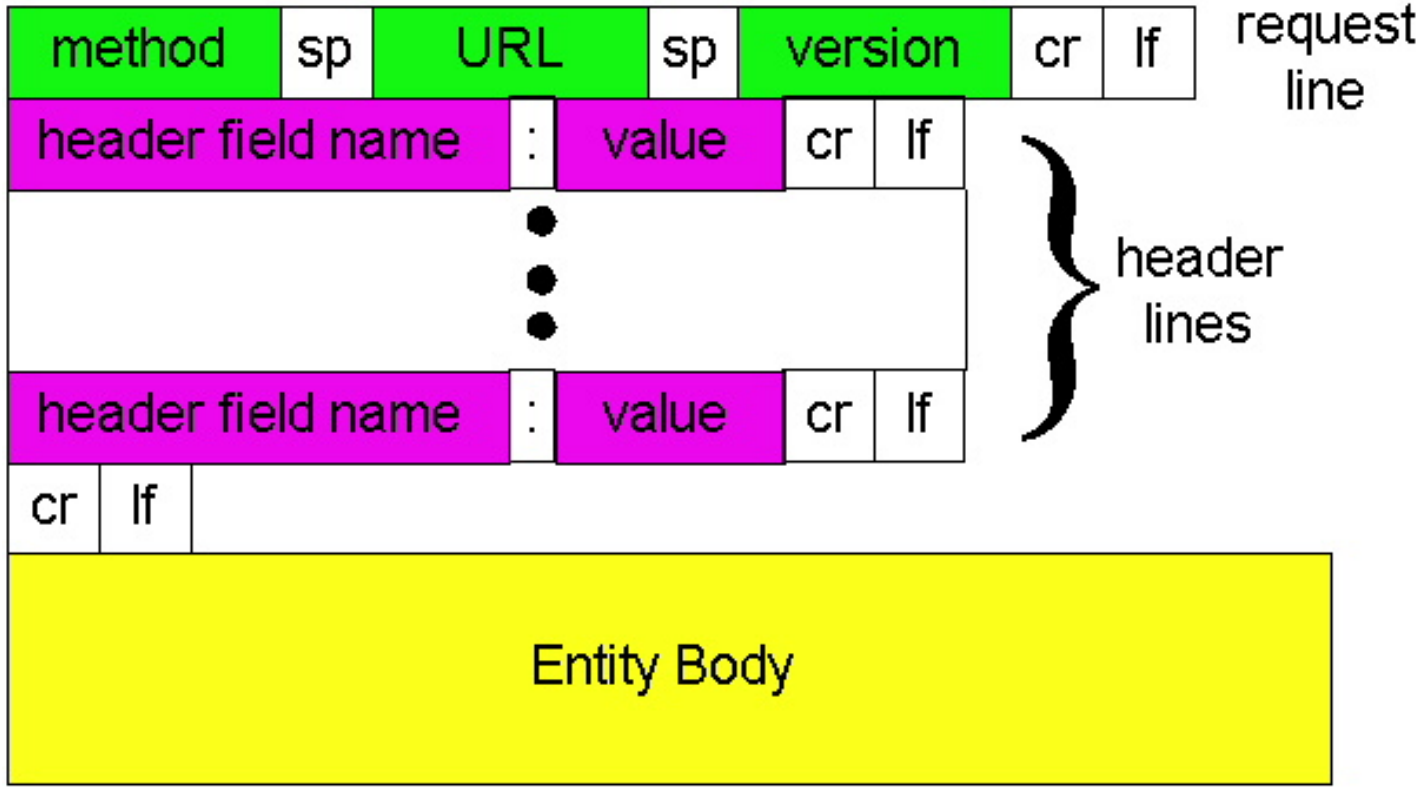
CR=Carriage Return
LF=Line Feed
SP=Space bar
WEB was designed by non-CS (physics) people. They wanted it as close to English as possible. It was one of the first ASCII protocols. Now there are many. They are easier to debug.

- What is the purpose of cr lf? *End of the field.*
- What would be in the entity-body? How is this different than an HTML body element?

The entity body is the server response. HTML describes the entire page, which may require many requests/responses.

©2026 Raj Jain

HTTP Request Message Format



Student Questions

- How do URLs with a non-latin scripts/characters/emojis work? Are they valid names for web servers, or do they get traslated into latin scripts? Does this add another layer to the request and response process?

The Unicode standard unifies character encoding. UTF-8, UTF-16, are examples of Unicodes. HTTP now allows and encourages Unicodes in body and allows it in URLs.
Ref: https://en.wikipedia.org/wiki/Unicode

- Have there been attempts to get away from HTTP since it is a text-based protocol?

No. HTTP resulted in many followers.

©2026 Raj Jain

Sample HTTP Response Message

HTTP/1.1 200 OK
Connection: close
Date: Tue, 09 Sept 2009 13:00:15 GMT
Server: Apache/1.3.0 (Unix)
Last-Modified: Sun, 6 May 2009 09:23:24 GMT
Content-Length: 6500
Content-Type: Text/html

Data…

Status Codes:

- 200 OK
- 301 Moved Permanently
- 400 Bad Request
- 404 Not Found
- 505 HTTP Version Not Supported

Student Questions

- What's the difference between a host and a server? Does the host decide what type of server it will use?

Host = Client or server

- What is 200 OK and 301 Moved Permanently?

Error messages have two parts: The first part can be read by computers, 2nd the part by humans.

- Is the Content-Length the number of bytes or the number of lines?

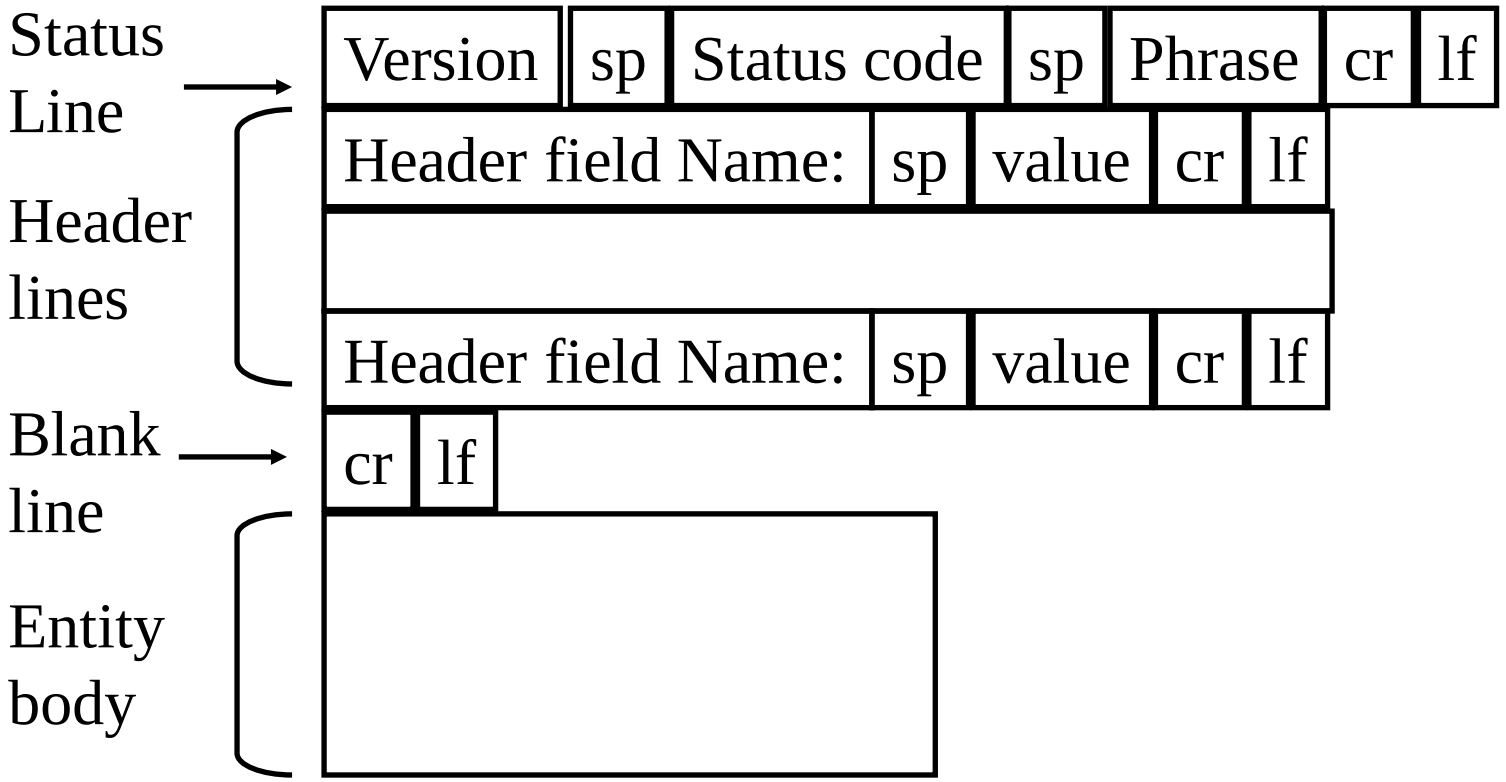
Number of bytes

- Are status codes similar to sockets/ports in that there are well-defined ones, but developers have the ability to customize them to suit their needs?

W3C defines status codes.Not sure if there is small space left for user defined codes.

©2026 Raj Jain

HTTP Response Message Format



Student Questions

©2026 Raj Jain

Hands-on HTTP

telnet www1.cse.wustl.edu 80
GET /~jain/cse473-19/sample.htm HTTP/1.1
Host: www1.cse.wustl.edu

HTTP/1.1 200 OK
Date: Tue, 13 Sep 2019 23:39:53 GMT
Server: Apache/2.2.3 (CentOS)
Accept-Ranges: bytes
Content-Length: 233
Connection: close
Content-Type: text/html; charset=ISO-8859-1

<HTML>
<head>
</head>
<body>
This is a sample text.
</body>
</html>

This is a sample text.

NOTE: Many servers no longer allow telnet access and so this may not work with those servers.

Student Questions

- Where does the content length start and end? Does the content length only refer to the requested data or also include the rest of the response message?

"Content-Length" does not include anything other than the content (data).

- What does "Connection: close" mean

Please close the TCP connection after responding to this request.

- What does "Accept-Ranges: bytes" mean?

Unit of length

- How do images/sound work?

They are binary files or streams.

- In the first line(telnet www1.cse.wustl.edu 80), can we access a different port by physically entering, let's say, 123 instead of 80? If that's not possible, how can we access a different port on someone else's machine, and how do ports ensure security?

Initially, telnet access was allowed on all ports. Now it is blocked for security reasons.

- Why don't some servers allow telnet access?

Telnet is insecure.

©2026 Raj Jain

Hands-on HTTP

telnet www1.cse.wustl.edu 80
GET /~jain/cse473-19/sample.htm HTTP/1.1
Host: www1.cse.wustl.edu

HTTP/1.1 200 OK
Date: Tue, 13 Sep 2019 23:39:53 GMT
Server: Apache/2.2.3 (CentOS)
Accept-Ranges: bytes
Content-Length: 233
Connection: close
Content-Type: text/html; charset=ISO-8859-1

<HTML>
<head>
</head>
<body>
This is a sample text.
</body>
</html>

This is a sample text.

NOTE: Many servers no longer allow telnet access and so this may not work with those servers.

Student Questions

- What is Apache?

Apache is an open-source web server.

- How is Telnet access different from a standard browser?

Telnet is a remote terminal protocol. It gives almost complete access to the computer.

- What's the difference between the Telnet and Host line if it specifies the same link?

Telnet requires only the host, not the rest of the URL.

- For webpages with images, videos, sounds, or other multimedia, will the content type still be text/HTML?

No, e.g., audio/mpeg. [RFC 6938]

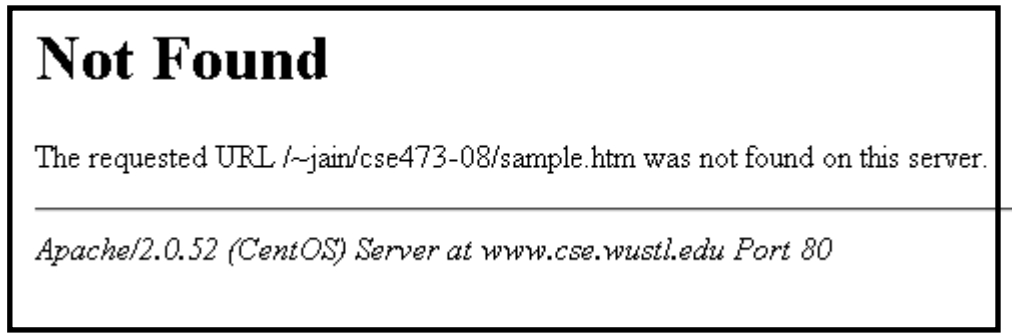
- Is there any way to prevent/catch small typos in the URL from returning 404, such as redirecting or checking capitalization?

Yes. But this is not a part of the protocol. This can be done by the application clients (Firefox, Chrome,…). The directory part of the URL is case-sensitive on Linux and Case insensitive on Windows. People make money by buying misspelled domain names, e.g., citibnk.com.
⇒ Credit Alert Systems

©2026 Raj Jain

Hands-on HTTP (cont)

telnet www1.cse.wustl.edu 80
GET /~jain/cse473-08/sample.htm HTTP/1.1
Host: www1.cse.wustl.edu



HTTP/1.1 404 Not Found
Date: Tue, 13 Sep 2019 23:42:48 GMT
Server: Apache/2.2.3 (CentOS)
Content-Length: 307
Connection: close
Content-Type: text/html; charset=iso-8859-1

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>404 Not Found</title>
</head><body>
<h1>Not Found</h1>
<p>The requested URL /~jain/cse473-08/sample.htm was not found on this server.</p>
<hr>
<address>Apache/2.2.3 (CentOS) Server at www1.cse.wustl.edu Port 80</address>
</body></html>
```

Student Questions

- Can someone modify or add cookies manually to access other users' data?

You can modify cookies. But they are generally encrypted for security. If you modify any encrypted information, the receiver will know that the message has been modified and will reject it.

- How does an ad appear on a server when I searched for the thing on another server previously?

They all share the info like credit reports

- Do different websites use different cookies?

Yes.

- What could we use, If TelNet is not allowed?

You can monitor the traffic or write your own browser.

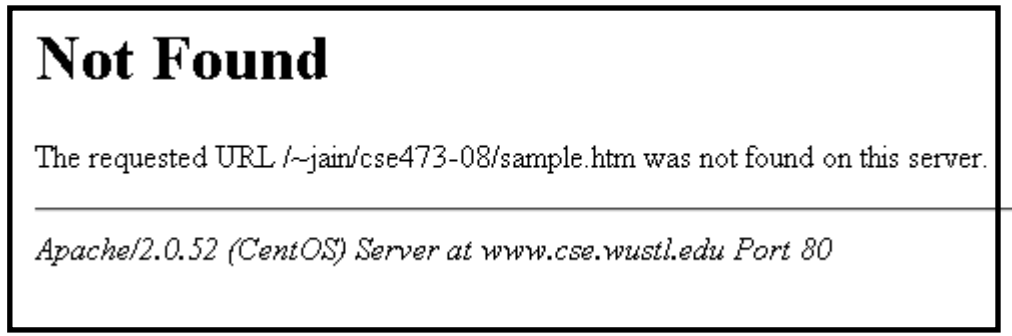
- Does "Content-Length: 233" means 233 Bytes or 233 characters?

Bytes

©2026 Raj Jain

Hands-on HTTP (cont)

telnet www1.cse.wustl.edu 80
GET /~jain/cse473-08/sample.htm HTTP/1.1
Host: www1.cse.wustl.edu



HTTP/1.1 404 Not Found
Date: Tue, 13 Sep 2019 23:42:48 GMT
Server: Apache/2.2.3 (CentOS)
Content-Length: 307
Connection: close
Content-Type: text/html; charset=iso-8859-1

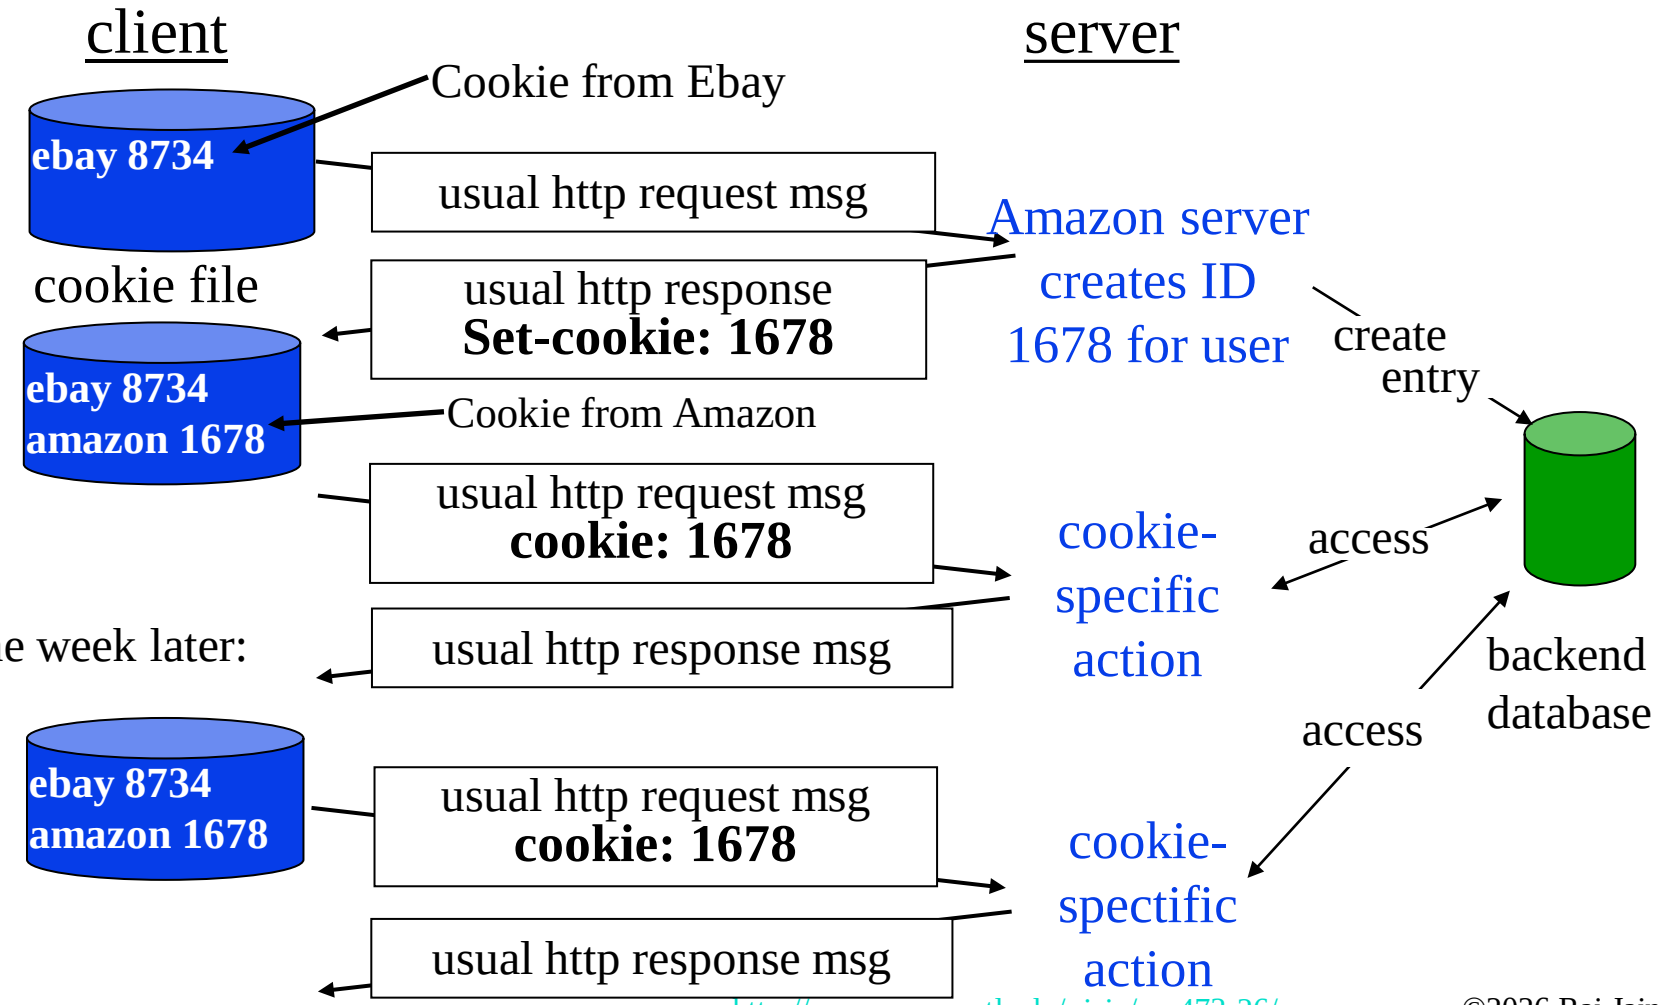
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>404 Not Found</title>
</head><body>
<h1>Not Found</h1>
The requested URL /~jain/cse473-08/sample.htm was not found on this server.
<hr>
<address>Apache/2.2.3 (CentOS) Server at www1.cse.wustl.edu Port 80</address>
</body></html>

Student Questions

©2026 Raj Jain

Cookies

- Allow servers to remember previous information
  C:\Documents and Settings\Raj Jain\Cookies\



Student Questions

- Can only the server determine when it wants to delete a cookie, and are cookies stored on the browser level or website level?

Cookies are stored on your computer. Cookies are used by servers as one factor to identify you. But they are misused by advertisers to track you. You can delete them anytime or set the browser option to not store them. But then you may have to do two-factor authentication on some servers.

- Are there web servers that require a username and password that do not use cookies? If so, how do those servers "remember" user information?

They start without prior information. That is when they will ask you to verify using 2-factor authentication.

- What do you think about Google's decision to move away from cookies? *They use cookies.*
- Where do clients store cookies?

Wherever browsers store bookmarks

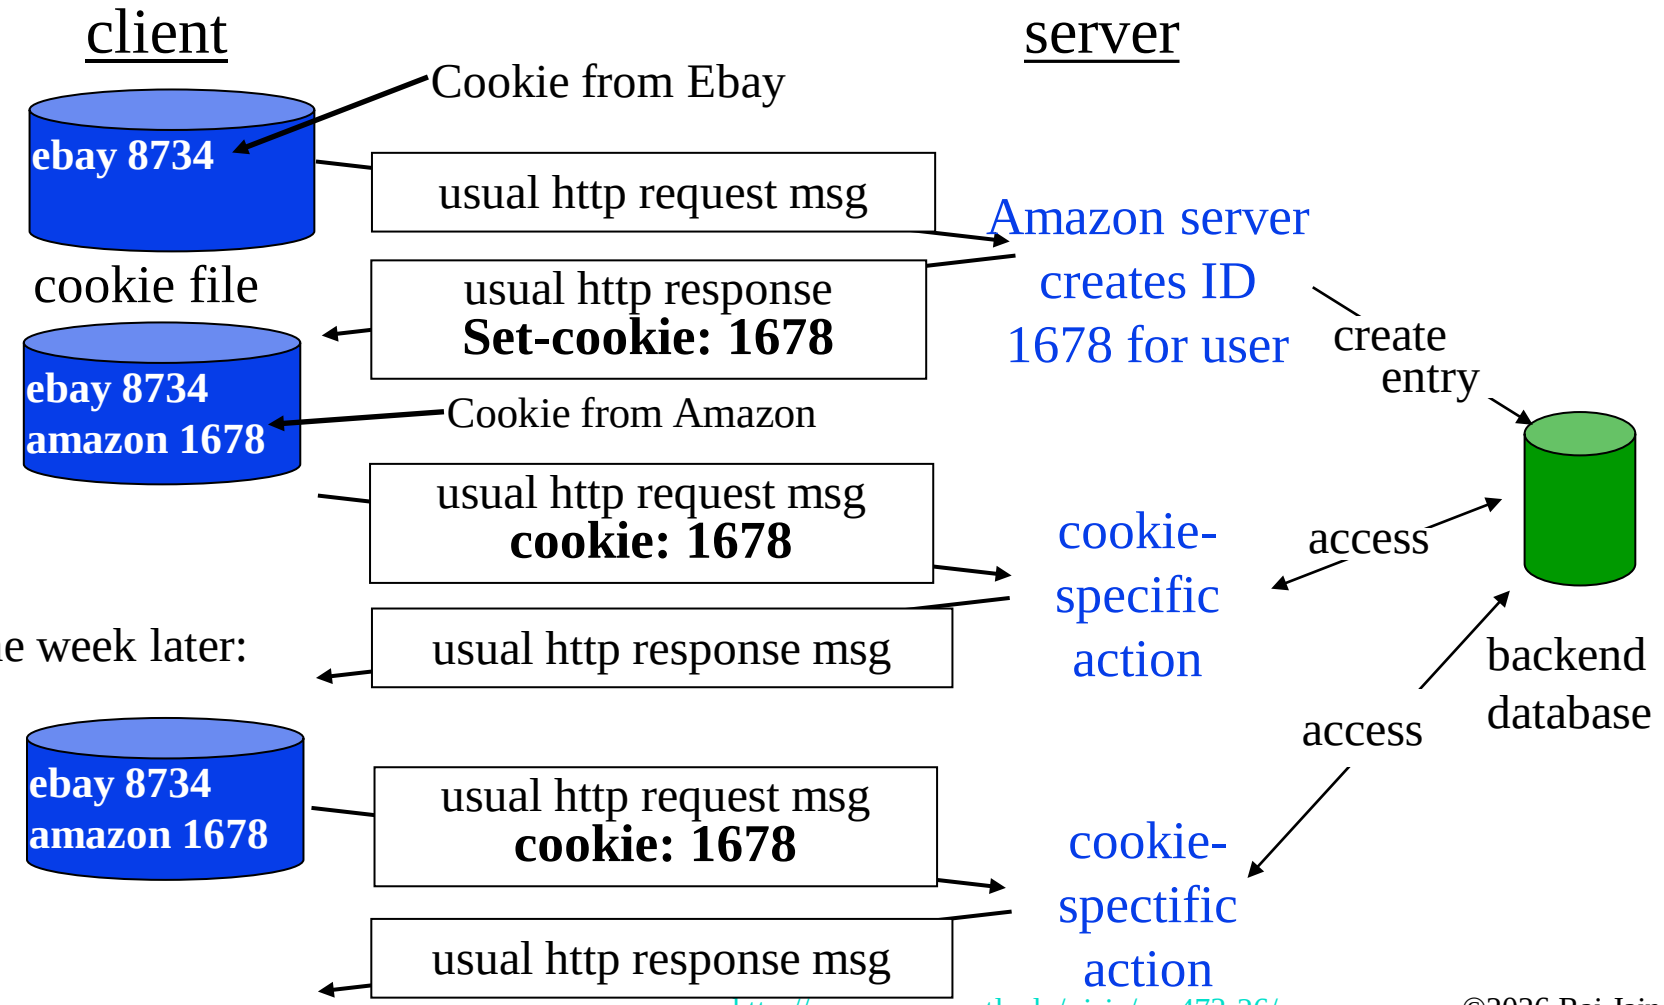
- Does each website own a cookie, or does some specific content on the website own a cookie?

Each server issues a cookie. A page may have info from many servers, e.g., an Ad server.

©2026 Raj Jain

Cookies

- Allow servers to remember previous information
  C:\Documents and Settings\Raj Jain\Cookies\



Student Questions

- Is sensitive information (passwords, etc.) ever stored in a cookie? Is it possible for a malicious user to create a "fake" cookie to impersonate another user?

Passwords are never transmitted on the wire. There are methods that allow password authentication without transmitting the password.

- Are there any mistakes on this slide?

Why do you think so?

- How does the server decide what value to set for the cookie?

Unique value that can be looked up to find the user again.

- What happens when the user has disabled cookies?

The website will not work.

- Should we always decline unnecessary cookies from websites?

Some plugins remove advertising cookies.

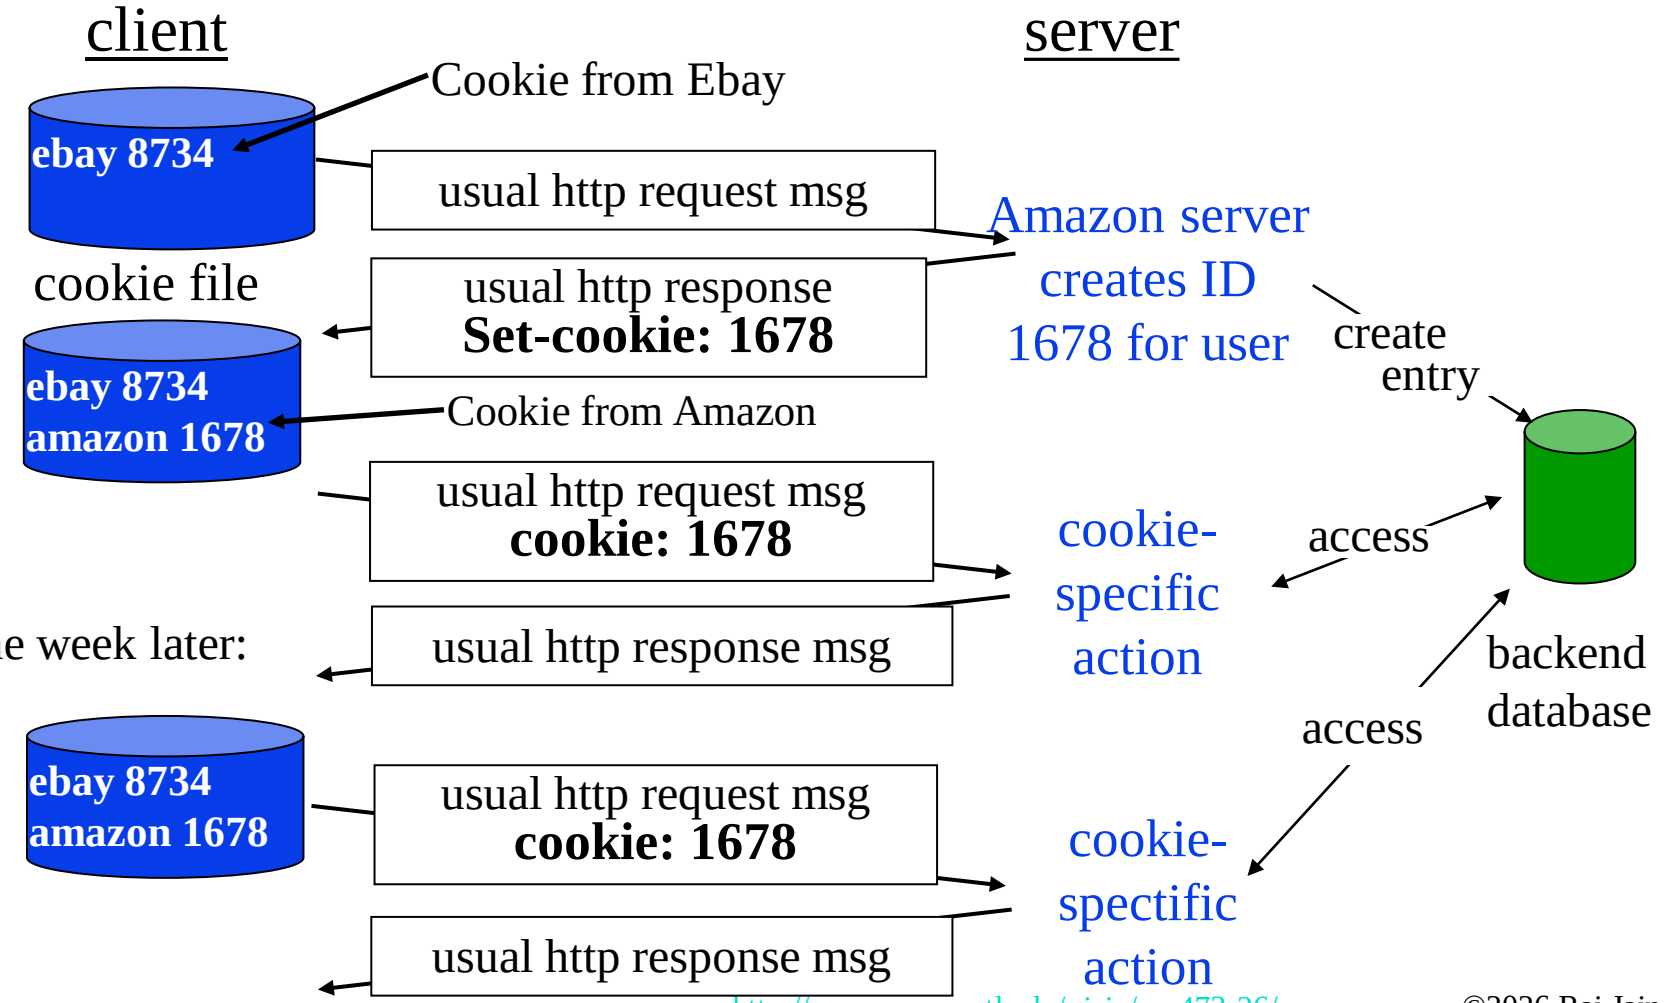
- What additional information do client cookies store except for their IDs?

There is no limit. They can store time, location, authentication, etc.

©2026 Raj Jain

Cookies

- Allow servers to remember previous information
  C:\Documents and Settings\Raj Jain\Cookies\



Student Questions

- Can the server request cookies from other servers? Or can it only access cookies from its own?

Most large websites use a server farm. Servers in a farm act as one server and use the same cookie format, and can interpret them.

- What are the disadvantages of accepting cookies when browsing all websites?

They track you.

- What about Facebook? Do they use cookies to check what posts I viewed and decide what to show me next? *Yes.*
- What if cookies never existed? What would be the drawbacks of not having cookies?

Difficult to do business on the internet.

- Are cookies just like tokens? Is user-related information kept in the server database? *Yes.*

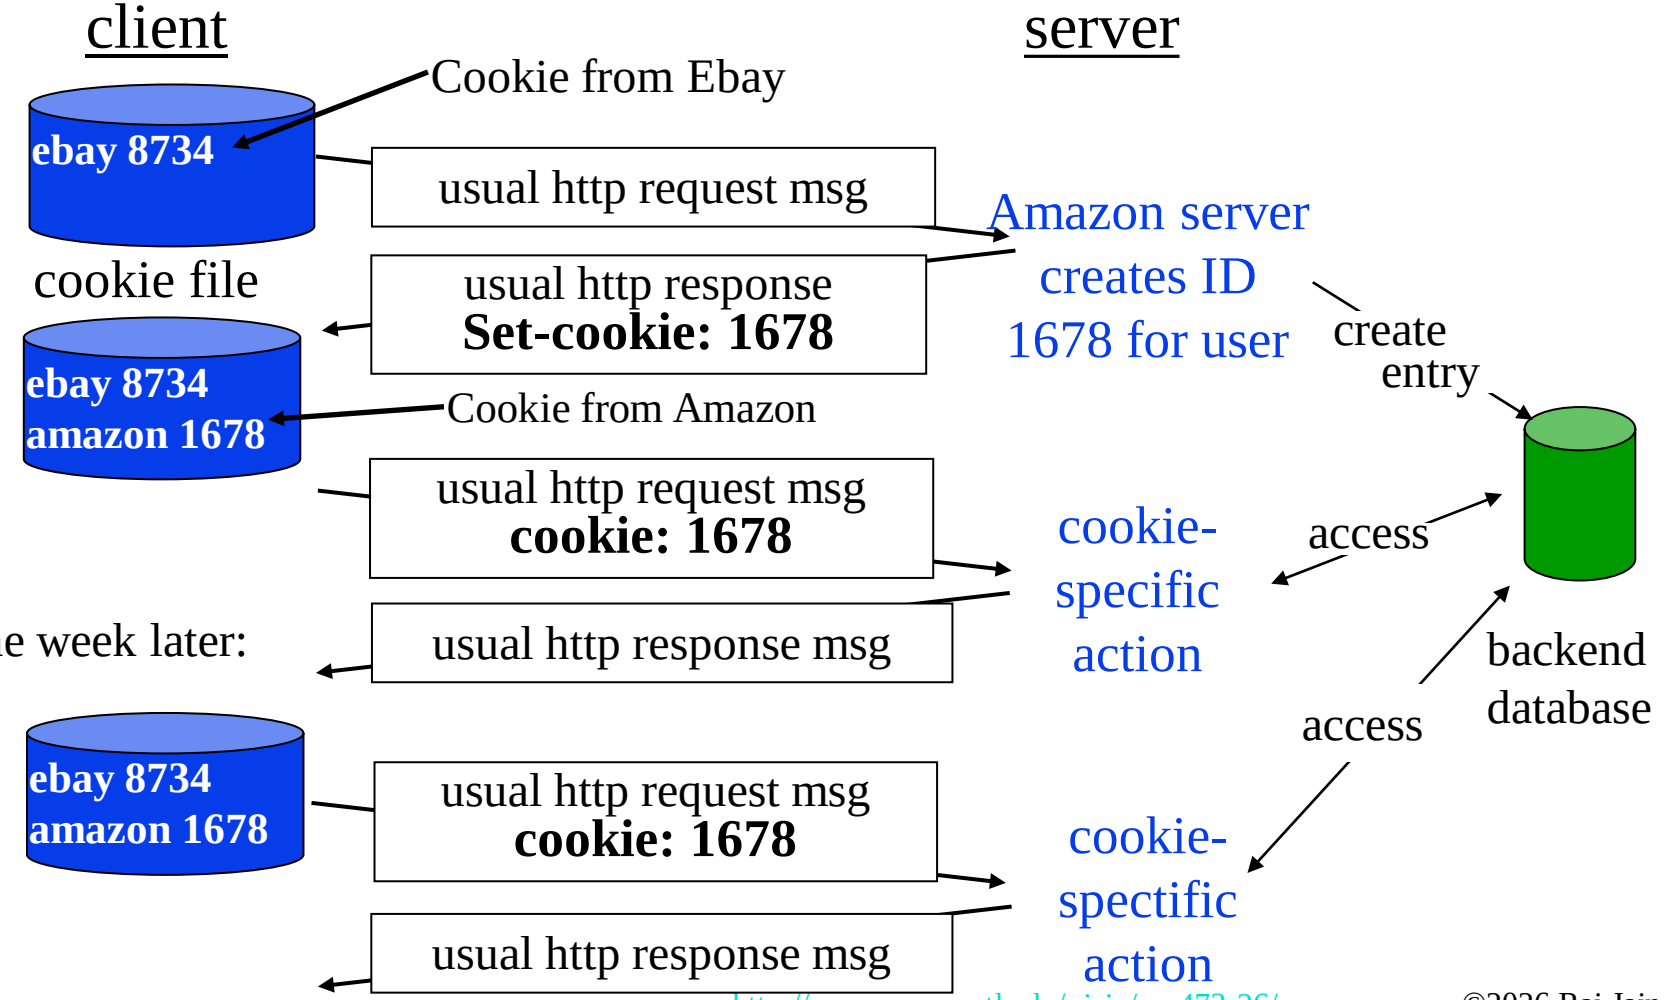
- Are cookies usually sent encrypted?

No. But they can be.

©2026 Raj Jain

Cookies

- Allow servers to remember previous information
  C:\Documents and Settings\Raj Jain\Cookies\



Student Questions

- Can you manually create and set your cookie instead of having the server give it to you?

Yes, but it may reject it if the server does not find it in its list.

- Shouldn't cookies always be key-value pairs?

Yes. But the value can be pretty complex.

- What information do websites require that necessitates cookies?

To relate your current query with your previous queries.

- How do websites identify you again when you block cookies?

They will treat you as if they don't know you. Re-authentication.

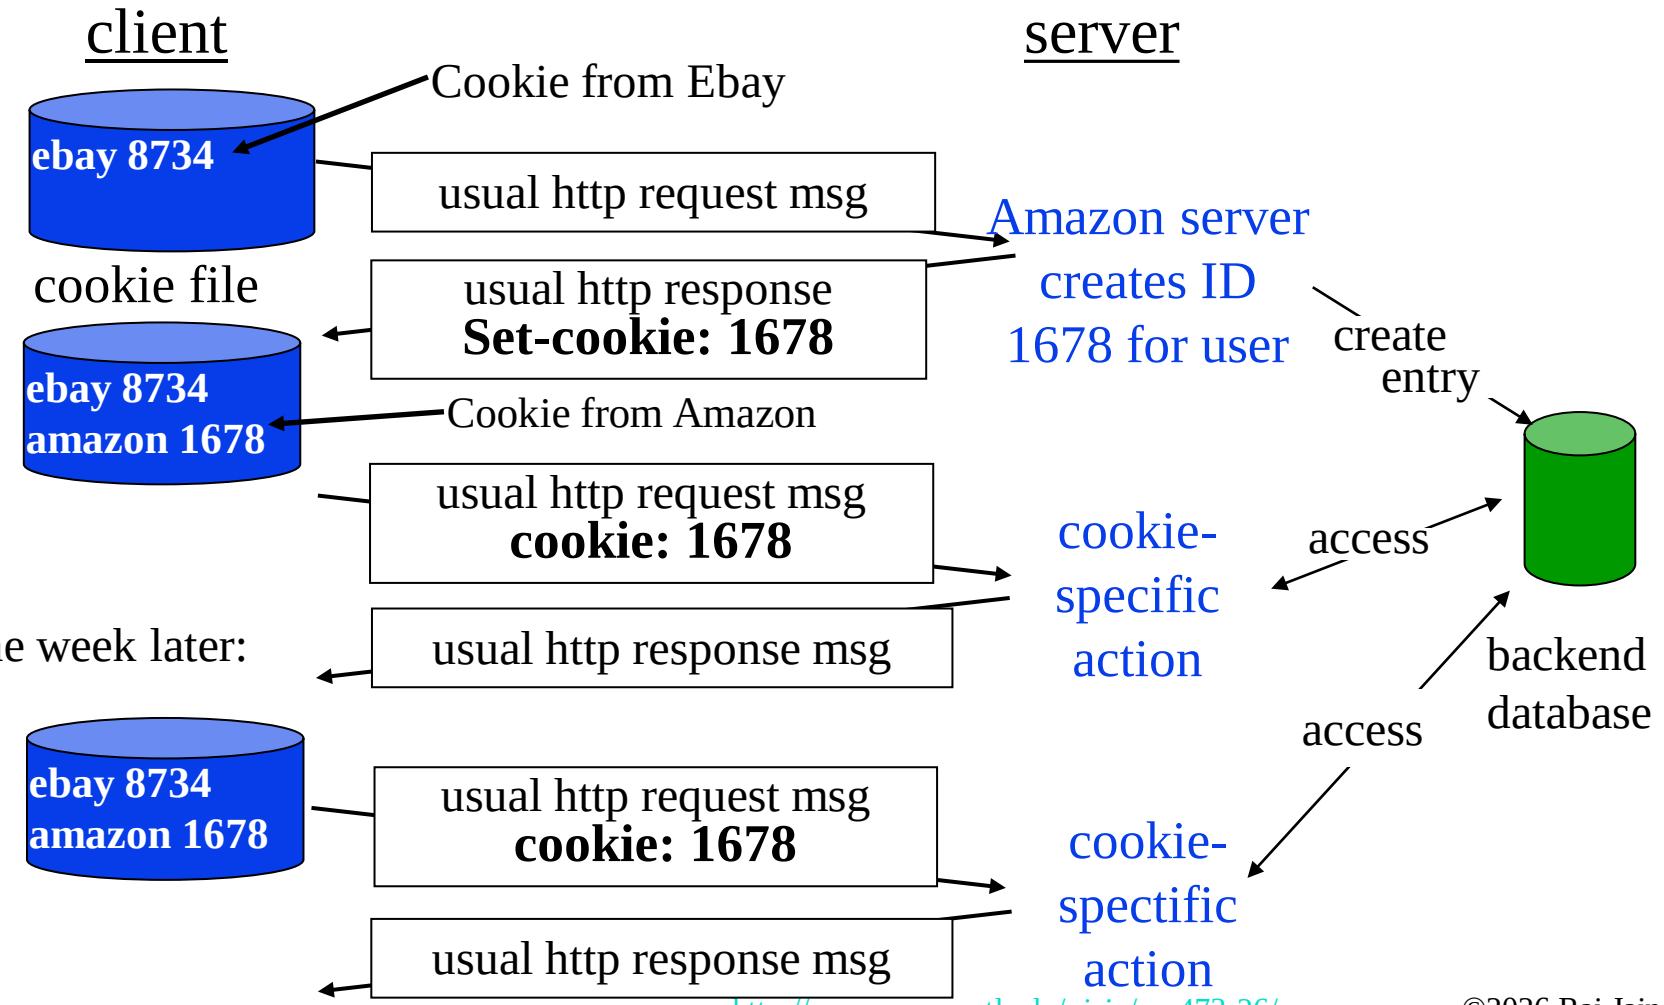
- Are cookies in the client to allow the client to cache or remember previous information?

No. To allow the server to locate the previous information.

©2026 Raj Jain

Cookies

- Allow servers to remember previous information
  C:\Documents and Settings\Raj Jain\Cookies\



Student Questions

- If I switch browsers, does the cookie change as well?

Yes. Browsers currently don't share cookies.

- I've heard that Google wants to remove cookies from Chrome; what could replace them?

Google is against third-party cookies to enhance privacy.

- Where are cookies stored? I understand that the client IDs are stored in the website's backend DB on the server side but how are the client cookies stored? I see a DB on the diagram but what DB is that?

Each browser has a local directory or file to store this.
%APPDATA%\Mozilla\Firefox\Profiles\cookies.sqlite for Firefox.

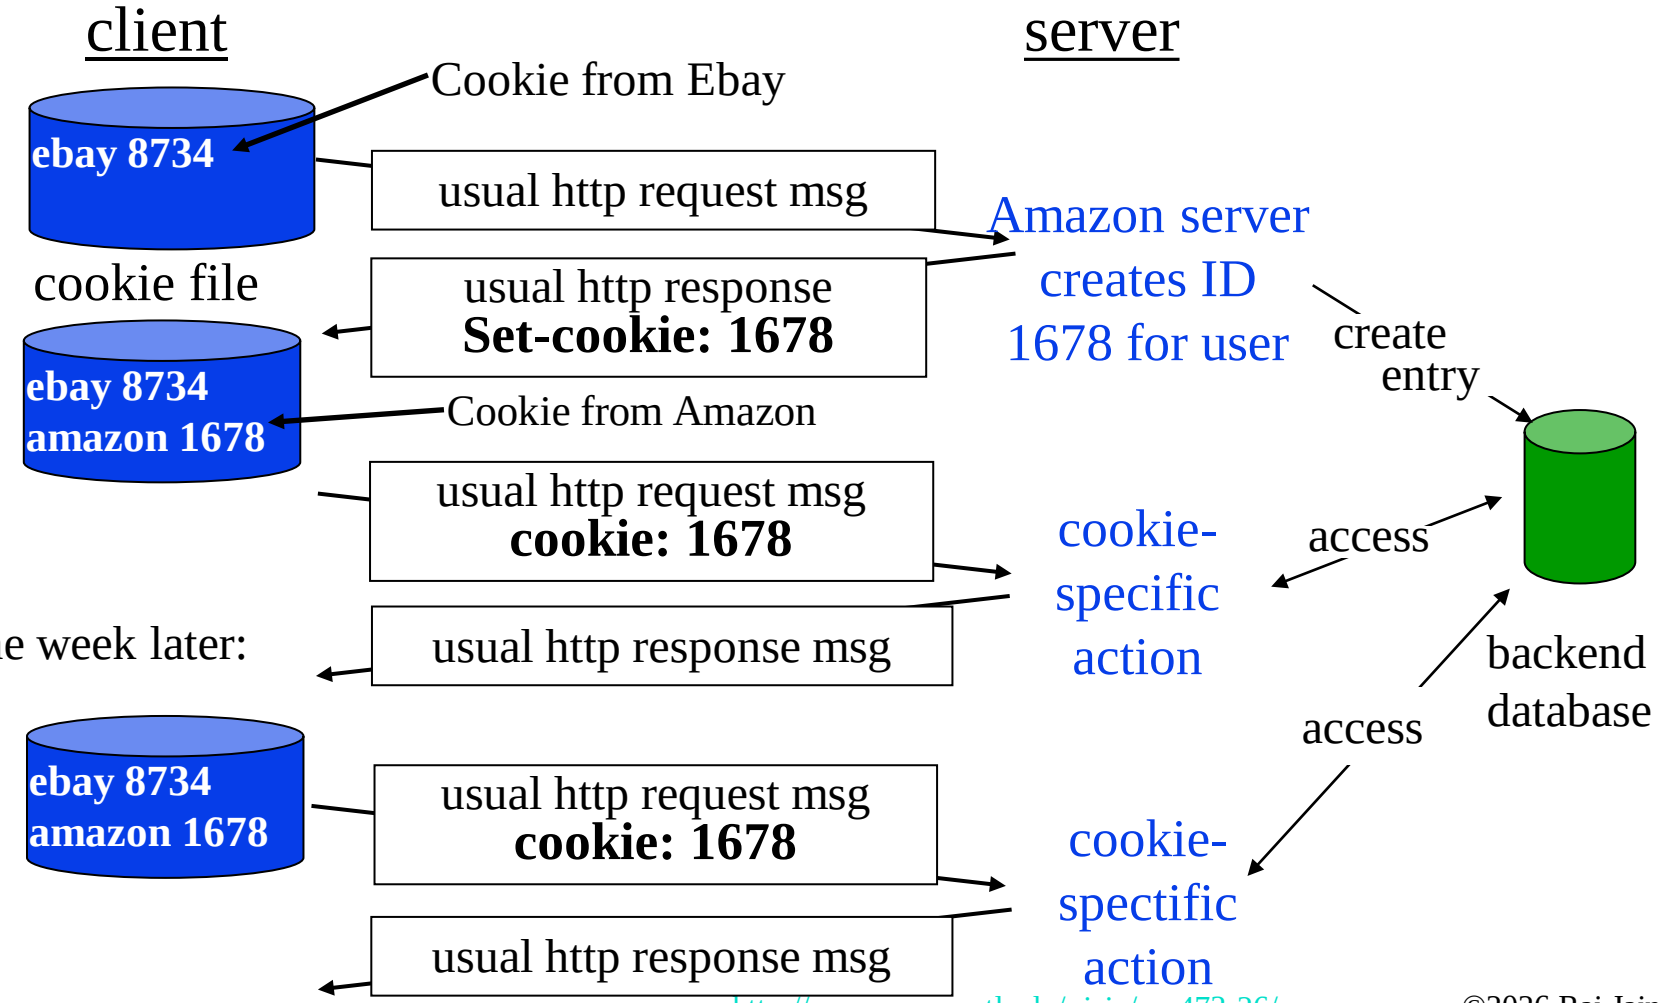
- Do cookies ever present any security concerns?

Privacy concerns since they tell the servers what you did before. Most commercial servers sell this info.

©2026 Raj Jain

Cookies

- Allow servers to remember previous information
  C:\Documents and Settings\Raj Jain\Cookies\



Student Questions

- Why do we see so much controversy on the use of cookies if it isn't storing sensitive information?

See above.

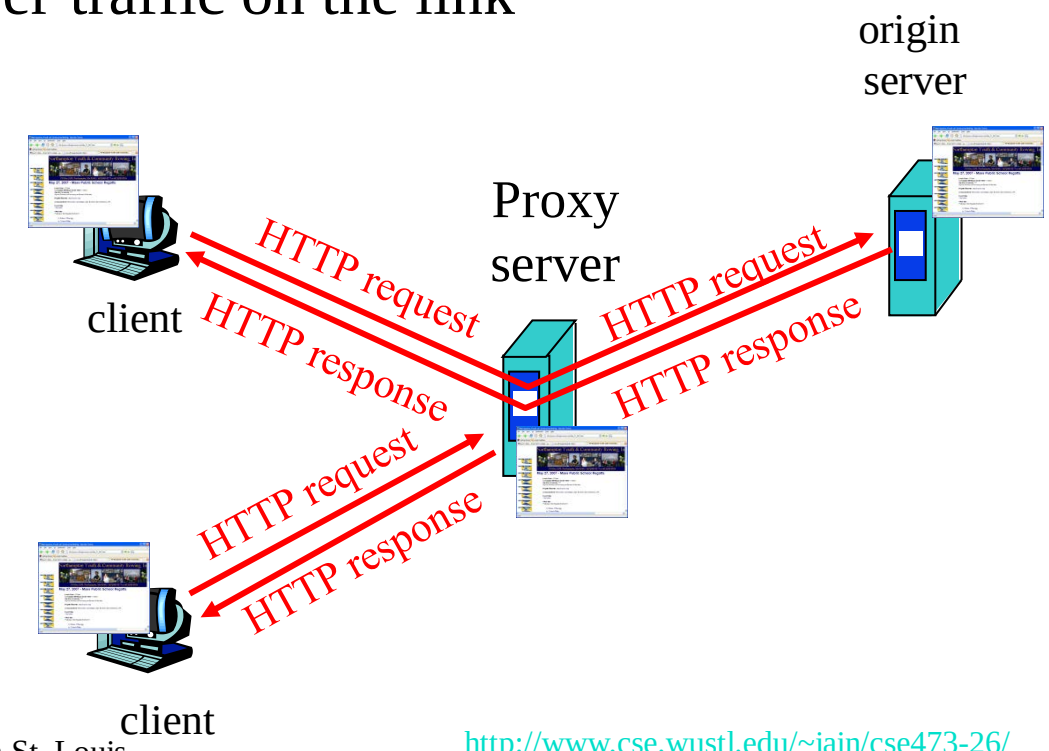
- Can you explain the diagram in more detail?

Sure.

©2026 Raj Jain

Proxy Server: Web Caching

- All requests are sent to proxy server
- Proxy server caches objects
- Only new objects are requested from origin server
- Fast, Lower traffic on the link



Student Questions

- Is there an effect or significance to URLs starting with 'www'?

WWW is for convenience. HTTP does not require it. You can use any name for your web server. http://raj.jain.com is a valid URL for a web server.

- So if we use a proxy server, the data saved by the proxy server has already been got by us before?

It could be for someone else that needed the same info.

- Is the proxy server the local cache of the origin server? *Yes*
- Is being fast a requirement for a proxy server, or is it just desirable? If a proxy server is slow, would it not become a bottleneck?

Anything on the path between the client and server could become a bottleneck if slow.

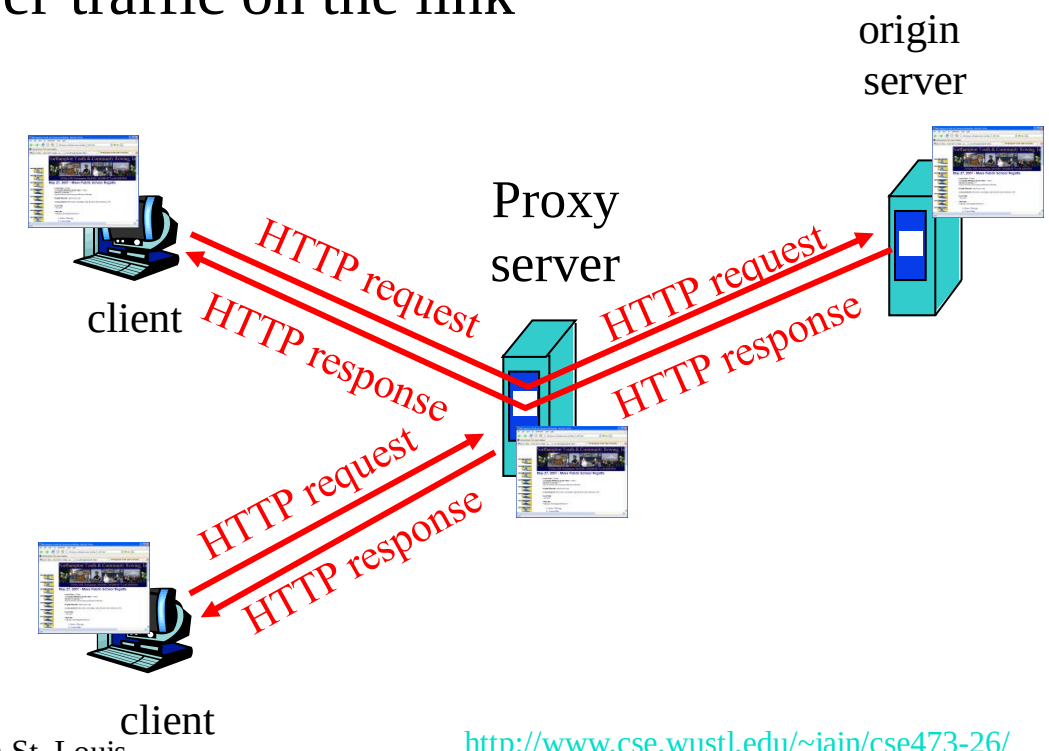
- [Book p. 108] When a proxy server stores objects, does this include cookies too? What security protocols are in place to prevent critical information from being stored and stolen?

Critical information would usually be encrypted and not stored or interpretable on the way.

©2026 Raj Jain

Proxy Server: Web Caching

- All requests are sent to proxy server
- Proxy server caches objects
- Only new objects are requested from origin server
- Fast, Lower traffic on the link



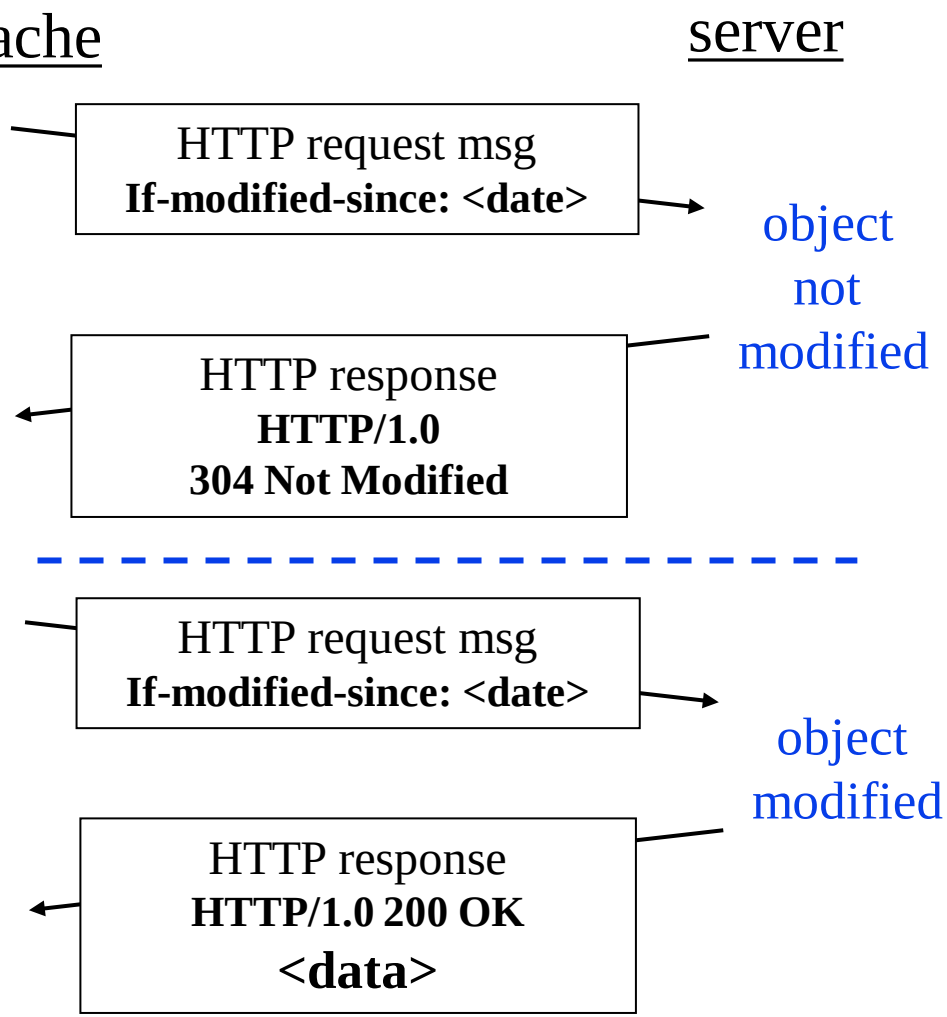
Student Questions

- How does the proxy server know when the page has been updated so it can refresh?

Every object is returned with a lifetime.

©2026 Raj Jain

Conditional GET



Student Questions

- Can a conditional get be used to launch a DoS attack? You mentioned that multiple get requests could be sent, and only a few might return (or return with the intended response).

Clients can retransmit the request.

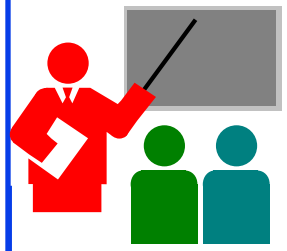
- Are there also conditional POSTs?

Yes. Ref: *HTTP conditional requests - HTTP | MDN (mozilla.org)*

- Chapter 2.2, page 142, conditional GET: does a proxy server send a conditional GET to the underlying server every time it receives a request for a cached object? If not, how can it guarantee that the resource has not been modified?

Each response has a lifetime field.

©2026 Raj Jain



HTTP: Summary

1. HTTP is a **client-server** protocol.
   Uses text-based messages
2. Web pages are generally written in **HTML.**
3. HTTP uses **non-persistent/persistent** TCP connections
4. Cookies allow servers to maintain a **state.**
5. Proxy servers improve performance by **caching** frequently used pages
6. **Conditional gets** allow proxy servers to reduce Internet traffic

Ref: Read Section 2.2 Full. Try R10-R14.

Student Questions

- Are conditional GETs only used with proxy servers?

This is not necessary.

- What does the "state" mean in message 4?

State=Memory

- [Book p. 166, R13] How web caching reduce the delay in receiving a requested object? Will it reduce delay for all objects?

Shorter round trip for cached objects. Lower traffic load for other objects.

- With advances in quantum computing, current asymmetric encryption algorithms may become obsolete. What are some alternatives to protecting data in transit?

Quantum-safe encryption

- How does a conditional get reduce internet traffic?

Information that hasn't changed is not sent.

©2026 Raj Jain

Homework 2A: HTTP

[10 points] The text below shows the reply sent from the server in response to the HTTP GET message. Answer the following questions, indicating where in the message below you find the answer.

HTTP/1.1 200 OK
Date: Tue, 07 Mar 2019 14:39:45GMT
Server: Apache/2.0.52 (Fedor)
Last-Modified: Sat, 5 Jan 2019 19:27:46 GMT
Etag: "526c3-f22-a88a4c80"
Accept-ranges: bytes
Content-Length: 5071
Keep-Alive: timeout=max=100
Connection: Keep-Alive
Content-Type: text/html; charset=ISO-8859-1

<!doctype html publi "-//w3c//dtd html 4.0 transitional//en">
<html>
<head>
<much more document text following here (not shown)>

A. Was the server able to successfully find the document or not?
B. What time was the document reply provided?
C. When was the document last modified?
D. How many bytes are there in the document being returned?
E. What are the first 5 bytes of the document being returned?
F. Did the server agree to a persistent connection?

Student Questions

- Is the end of HTML content indicated by End-of-File or CR/LF

Two CR/LFs.

©2026 Raj Jain

Lab 2A: Domains

[10 points] Submit answers for the following: (See hints in the parenthesis.)

1. Find the IP addresses of www.google.com and www.yahoo.com (ping)
2. Modify the hosts file to map www.google.com to yahoo's IP address and ping to www.google.com. Notice what address it is pinging to. Remove the modification to the host file, open a new command window and repeat. (Windows: c:\windows\system32\drivers\etc\hosts
   Mac: /private/etc/hosts)
3. Find the domain name and country of 128.252.165.7 (https://lookup.icann.org/en/lookup)
4. Find the owner of the wustl.edu domain (http://www.networksolutions.com/whois/index.jsp https://net.educause.edu/whois.htm)
5. Find the name server of the wustl.edu domain (http://www.networksolutions.com/whois/index.jsp)

Submit the screen snapshot showing the command used and the output. (Use Alt-Print-screen to capture the window to the clipboard and then paste it to Word). Submit either the Word file or a PDF of the Word file.

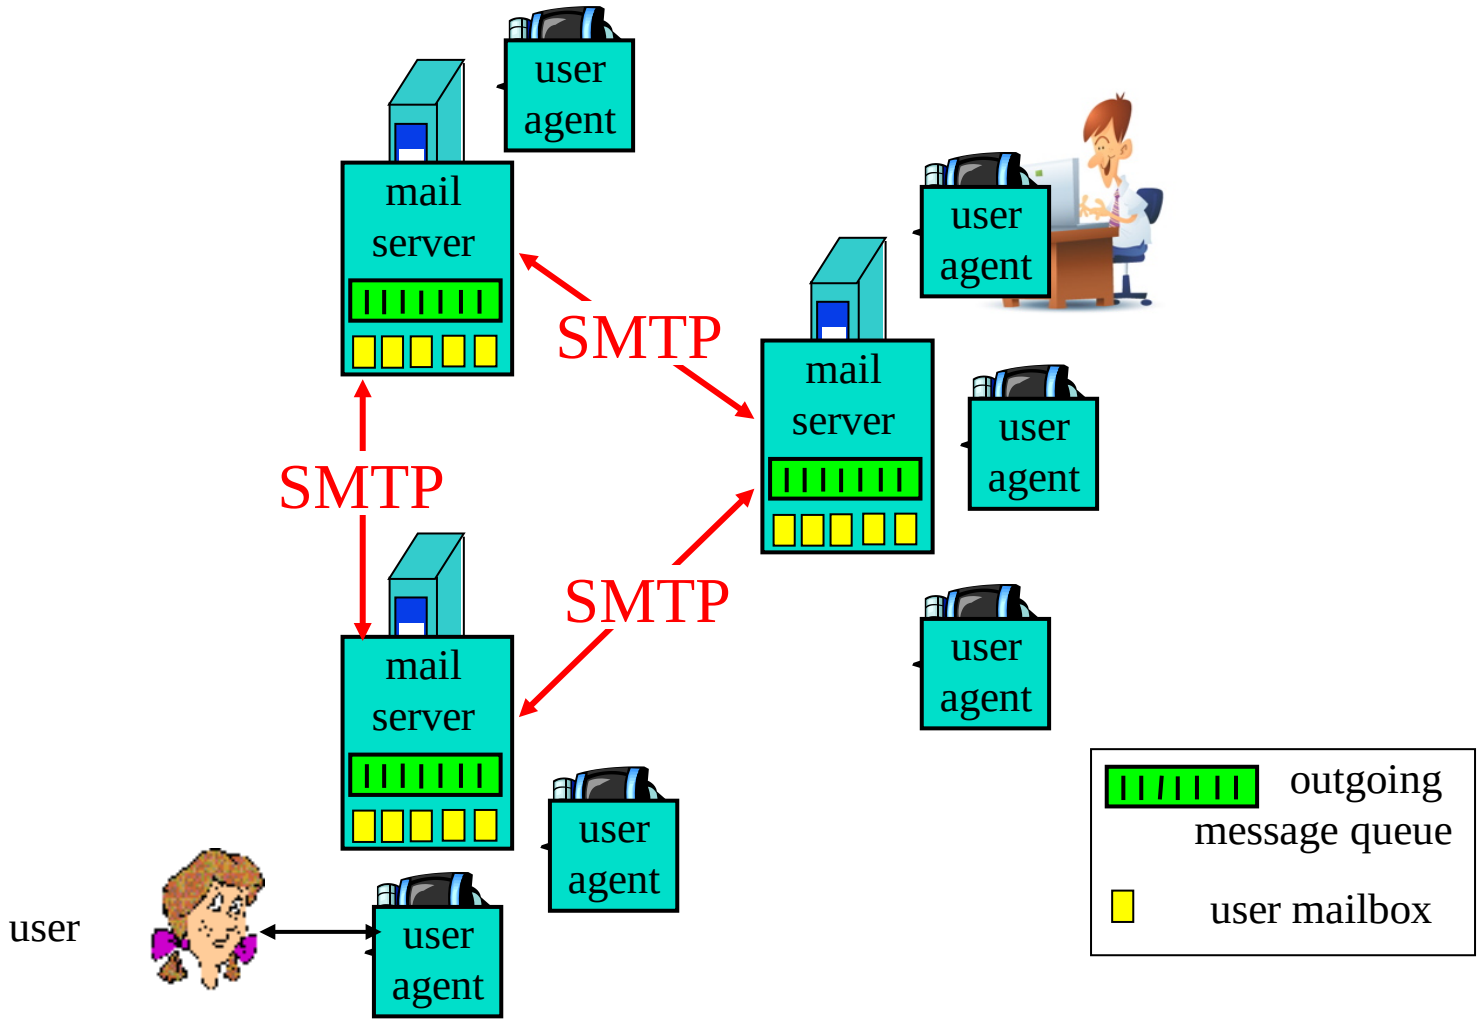
Student Questions

- Do we need to submit screenshots of our work like with lab 1?

No

©2026 Raj Jain

Electronic Mail



Student Questions

- How does SMTP know how to deliver mail to the correct server?

Mail servers use DNS to find the destination mail server. DNS is discussed later in this chapter.

- What protocols are used by social media to send messages, such as Facebook, Whatsapp, etc.?

Each app uses many protocols, including some proprietary protocols.

- Does SMTP also require the destination's IP address to deliver emails?

SMTP requires the destination mail server's IP address. Mail applications find that out from t

- What do mail servers do exactly? Is gmail.com the domain name for a mail server?

Similar to US Post office. Gmail.com is a generic domain. Google uses it to allocate email addresses. Mail servers could have any name.
he destination user's email address and DNS.

- What happens in mail servers when a user sends an email to an invalid email address?

The mail is returned to the sender.

©2026 Raj Jain

SMTP

- Simple Mail Transfer Protocol
- Old Protocol: Allows only 7-bit ASCII messages
- All binary objects have to be converted to ASCII
- Uses port 25 at the server

Student Questions

- Why hasn't SMTP been upgraded to accept more than 7 bit ASCII messages?

There are mappings that allow representing all characters in all languages to ASCII, e.g., UTF-8.

- Has SMTP been innovated over the years? Or has it mostly stayed the same?

All protocols are continuously evolving and updated.

- Do we have to remember the port number for each protocol?

No, but we usually remember 21, 25, and 80.

©2026 Raj Jain

Sample SMTP Exchange

C: telnet mail.seas.wustl.edu 25
S: 220 POSTOFFICE.seas.wustl.edu Microsoft ESMTP MAIL Service, Version: 6.0.3790.46 75 ready at Tue, 13 Sep 2011 18:34:56 -0500
C: HELO acm.org
S: 250 POSTOFFICE.seas.wustl.edu Hello [128.252.19.232]
C: MAIL FROM: jain@acm.org
S: 250 2.1.0 jain@acm.org....Sender OK
C: RCPT TO: jain@wustl.edu
S: 250 2.1.5 jain@wustl.edu
C: DATA
S: 354 Start mail input; end with <CRLF>.<CRLF>
C: This is test email.
This serves as an exmaple for CS4703 class.
.
S: 250 2.6.0 <MAIL2j97vPYGrN7kf0V00000aff@POSTOFFICE.seas.wustl.edu> Queued mail for delivery
C: QUIT
S: 221 2.0.0 POSTOFFICE.seas.wustl.edu Service closing transmission channel

NOTE: Many servers no longer allow telnet access and so this may not work with those servers.

Student Questions

- Do the attached images also go through SMTP protocol or require some other protocol?

Bit strings can also be mapped as ASCII.

- Why is telnet considered insecure? Why has it not been updated and made more secure?

Telnet is no longer used. Similar and better functionality is offered by newer protocols.

- Using telnet, in the MAIL FROM section, can I type in any random email address? How do we verify the sender since we are not logging in?

Email senders are not verified or authenticated unless you use secure email.

- Can we fake our identity?

Yes, you could fake in the past. But now, most email is secure and authenticated.

- What has replaced Telnet?

Remote desktop.

- Why is telnet iffy? Is it deprecated?

Yes. Insecure.

©2026 Raj Jain

Sample SMTP Exchange

C: telnet mail.seas.wustl.edu 25
S: 220 POSTOFFICE.seas.wustl.edu Microsoft ESMTP MAIL Service, Version: 6.0.3790.46
75 ready at Tue, 13 Sep 2011 18:34:56 -0500
C: HELO acm.org
S: 250 POSTOFFICE.seas.wustl.edu Hello [128.252.19.232]
C: MAIL FROM: jain@acm.org
S: 250 2.1.0 jain@acm.org....Sender OK
C: RCPT TO: jain@wustl.edu
S: 250 2.1.5 jain@wustl.edu
C: DATA
S: 354 Start mail input; end with <CRLF>.<CRLF>
C: This is test email.
This serves as an exmaple for CS4703 class.
.
S: 250 2.6.0 <MAIL2j97vPYGrN7kf0V00000aff@POSTOFFICE.seas.wustl.edu> Queued mail for delivery
C: QUIT
S: 221 2.0.0 POSTOFFICE.seas.wustl.edu Service closing transmission channel

NOTE: Many servers no longer allow telnet access and so this may not work with those servers.

Student Questions

©2026 Raj Jain

HTTP vs. SMTP

HTTP	SMTP
Persistent/Non-Persistent TCP	Persistent TCP
Mostly Pull	Mostly Push
Accepts binary objects	Accepts only 7-bit ASCII
One Object/response	Multiple objects/message

Student Questions

- The table indicates that HTTP only allows one object per response. It was also suggested previously that one website is composed of several objects: sound, image, video, text, etc. Then does it mean that one website might require several HTTP responses to completely transfer the data objects to the client?

Persistent connections return multiple responses. One object per response.

- Can we use HTTP to replace SMTP?

Yes. Now a days, HTTP is used to transport any protocol. Discussed in advanced networking courses.

- Is SMTP considered persistent because it sends multiple back-and-forth communications that remember each other, while HTTP must use cookies to do this?

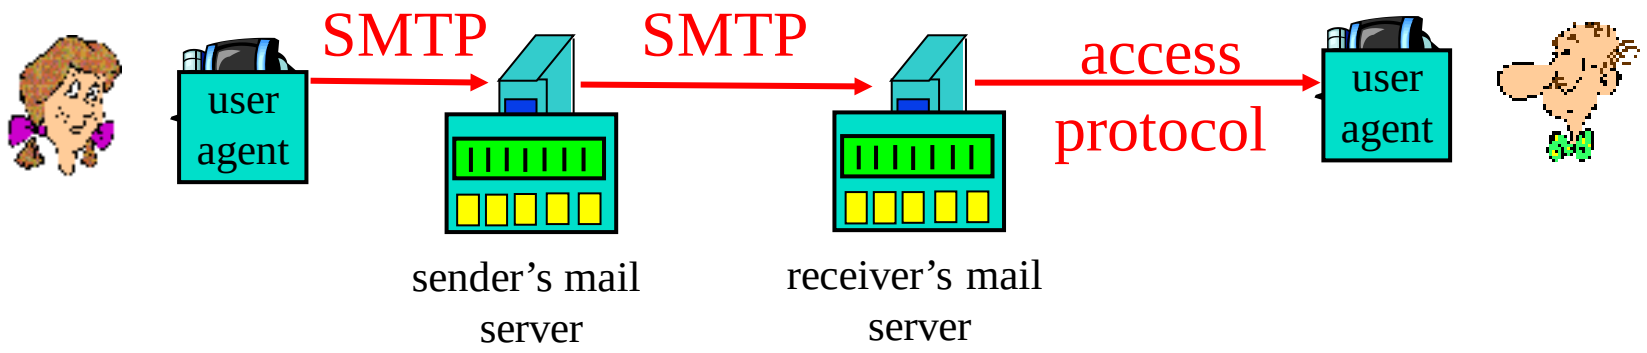
Yes. Different groups designed them.

©2026 Raj Jain

End of Part 1

Mail Access Protocols

- SMTP can be used to send messages to destination user agent
  ⇒ Requires destination to be always accessible
- Post Office Protocol - Version 3 (POP3)
- Internet Mail Access Protocol (IMAP)
- HTTP



Student Questions

- Is SMTP now obsolete?

No. I use it every day.

- Is HTTP also a mailing protocol, or does it work alongside/on top of a mailing protocol to deliver emails?

HTTP is used underneath the email protocol.

- Do we still use SMTP or POP3 these days or all switch to HTTP?

Most email programs, such as Outlook and smartphone emails, use POP3/IMAP. HTTP is just an alternative for these programs.

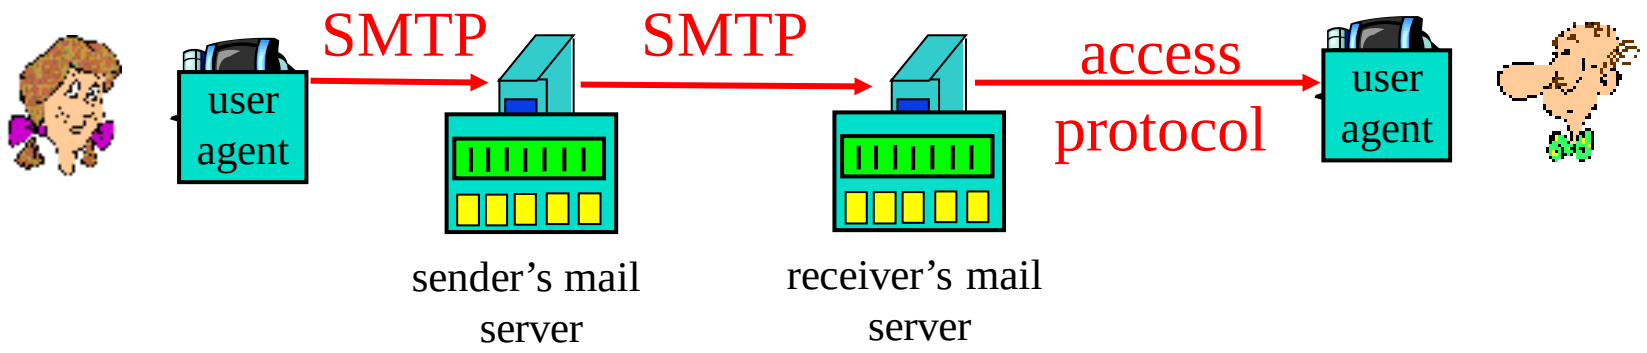
- What are the differences between the three protocols for receiving mail, and when would you choose one over the others?

POP3 allows you to delete emails from the server. IMAP allows you to read emails on multiple devices, and a copy is always kept on the server. HTTP is just an interface. Browsers may use IMAP.

©2026 Raj Jain

Mail Access Protocols

- SMTP can be used to send messages to destination user agent
  ⇒ Requires destination to be always accessible
- Post Office Protocol - Version 3 (POP3)
- Internet Mail Access Protocol (IMAP)
- HTTP



Student Questions

©2026 Raj Jain

POP3 protocol

Authorization phase

Transaction phase

S: +OK POP3 server ready
C: user bob
S: +OK
C: pass hungry
S: +OK user successfully logged on

C: list
S: 1 498
S: 2 912
S: .
C: retr 1
S: <message 1 contents>
S: .
C: dele 1
C: retr 2
S: <message 2 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off

Student Questions

- Are there any common ways of cross-protocol communication?

Not common. But yes. See https://en.wikipedia.org/wiki/Cross-layer_optimization

- Is there any reason that there is only one protocol for sending emails but multiple for receiving?

SMTP has been updated many times as the needs change. Receiving required new protocols to satisfy different people's needs.

- So between mail servers, they use SMTP, and between a user and mail server, POP is used?

POP is for receiving by the user.

- In the textbook, it says a user can send an email by using HTTP and also retrieve an email using HTTP, but that only happens between a user agent and the mail server. Why can it not happen between two mail servers?

SMTP is much faster and more efficient. HTTP is easier for the user, but it requires a lot more bits.

©2026 Raj Jain

POP3 protocol

Authorization phase

S: +OK POP3 server ready
C: user bob
S: +OK
C: pass hungry
S: +OK user successfully logged on

Transaction phase

C: list
S: 1 498
S: 2 912
S: .
C: retr 1
S: <message 1 contents>
S: .
C: dele 1
C: retr 2
S: <message 2 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off

Student Questions

- For POP3 protocol, does the server always need to send back "OK" to clients' requests, or it's optional?

For servers, it is required. If the client doesn't wait for it, some messages may be lost.

- Since I had already deleted the first message, the second message became the current first message. Why would the client still ask for "retr 2"?

The serial numbers are not changed in the middle of a connection.

- Does POP have any previous versions? If it has any predecessor, what is the main drawback?

POP3 is 3rd version of POP.

- For a server, how to know the address of the server serves the user that the outgoing email is sent to? Is this using the feature of aliasing of DNS protocol? If so, it seems the DNS server needs to maintain a highly verbose list of aliases because there exist huge numbers of email accounts.

Yes. DNS can resolve all addresses in the world. See Slide 2-42

©2026 Raj Jain

POP3 protocol

Authorization phase

S: +OK POP3 server ready
C: user bob
S: +OK
C: pass hungry
S: +OK user successfully logged on

Transaction phase

C: list
S: 1 498
S: 2 912
S: .
C: retr 1
S: <message 1 contents>
S: .
C: dele 1
C: retr 2
S: <message 2 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off

Student Questions

- Is the user urgent for the client in this condition?

In what condition?

- Are POP3 requests in ASCII like HTML or a different format?

Yes, ASCII

- Can we do cross-protocol communication?

No. Both servers and clients have to speak the same protocol.

- What will happen if a client tries to retrieve a message that doesn't exist? (e.g., retr 3 for this case)

Get an error code.

- In POP3, does the download and delete mode automatically delete the message after downloading, or does the user have to do it manually? Also, does the download and keep mode allow users to delete messages manually, or is that forbidden?

Download alone will not delete. Download and delete, or delete alone will delete.

©2026 Raj Jain

IMAP

- Internet Mail Access Protocol
- More sophisticated than POP3
- Allows users to maintain folders on the server
- Messages can be moved from one folder to another
- Users can get only headers or other components of the message
- Official IMAP site: www.imap.org

Student Questions

- When a message is deleted on the phone, can we still access the message on the laptop?

You can program each receiver to either delete the email or keep it even after reading it. So if you program your phone to not delete it from the server, it will be there.

- Does more advanced protocol mean lower efficiency? *No.*
- Can you go over IMAP again? *Sure.*
- Who decides which protocol (POP3/IMAP) to use when receiving e-mails?

The server offers choices for users to select one.

- Does an email server on the client end keep sending emails?

No. Only as requested by the user.

- Does the messaging protocol for IMAP look similar to POP3's messaging protocol?

No. It is very different. Both PoP and IMAP use ASCII.

©2026 Raj Jain

IMAP

- Internet Mail Access Protocol
- More sophisticated than POP3
- Allows users to maintain folders on the server
- Messages can be moved from one folder to another
- Users can get only headers or other components of the message
- Official IMAP site: www.imap.org

Student Questions

- Is IMAP managing on the server and POP3 managing on the device?

Both have management on client and server. In POP the functions on the server side are limited.

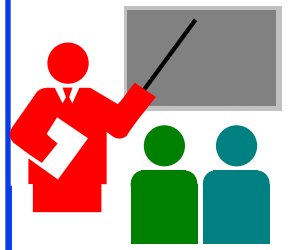
- What's the advantage of using IMAP or POP3?

IMAP is multi-device. POP is single device. IMAP can store messages for ever and can do many functions such as search, classify on the server.

- Why would someone want to use POP3 over IMAP?

They may not want to keep the email on the server forever.

©2026 Raj Jain



Mail: Summary

1. SMTP is the protocol for **sending** email
2. SMTP uses only **7-bit ASCII** messages
3. POP3, IMAP, or HTTP is used to **receive** email

Student Questions

- Why are SMTP TCP connections persistent? Shouldn't email messages only just send the one item?

SMTP sets up a TCP connection, sends all pending emails, and then disconnects. Each mail is often several megabytes, so it is not a good idea to disconnect after each segment. Therefore, the connections are persistent, but they are not on when there is no email to send.

- Can you explain R20 from this chapter.

To see email headers in outlook:
Open message. File->Properties

- [Book p. 121] where does the email header go in the protocol? Under DATA?

No header is handled separately by the protocol.

- Does any mail application use HTTP to receive email? Is it secure to use HTTP?

The user decides to use a browser. The mail servers then decide to use HTTPS.

Ref: Read Section 2.3 Full. Try R15-R20

©2026 Raj Jain

Homework 2B: Mail

[12 points] Consider accessing your e-mail with POP3.

a) Suppose you have configured your POP mail client to operate in the download and keep mode. Complete the following transaction to retrieve both messages, and sign off. Show the complete sequence of messages. (Fill in ? and successive messages)

C: list

S: 1 500

S: 2 901

S: .

C: retr 1

S: blah blah …

S: … Blah

S: .

?

?

b) Repeat part a if you have programmed your POP client in download and delete mode.

c) Suppose five minutes later, you again access POP to retrieve a new e-mail. Suppose that in the five-minute interval, no new messages have been sent to you. Provide a transcript of this second POP session for both options a and b above.

Student Questions

- For problem c, do we need to include the authorization phase? Or just the transaction phase?

Just the transaction phase.

- Does POP involve other more "complex" commands/programming, such as conditionals or loops?

No. But the client software can do all that.

- Do we have to remember the wire protocol of SMTP for exam?

Homework is the primary source of exam questions.

- To complete this homework, will we want to run our own POP server so that we can see the exact print out?

No. This is not the entire protocol. Just a few commands.

- How can we finish Homework 2B? Should we search the Internet for POP3 commands, or should we build a server to test it ourselves?

This is not a lab. You need to answer after reading the book and the slides.

©2026 Raj Jain

Homework 2B: Mail

[12 points] Consider accessing your e-mail with POP3.

a) Suppose you have configured your POP mail client to operate in the download and keep mode. Complete the following transaction to retrieve both messages, and sign off. Show the complete sequence of messages. (Fill in ? and successive messages)

C: list

S: 1 500

S: 2 901

S: .

C: retr 1

S: blah blah …

S: … Blah

S: .

?

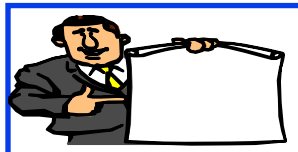
?

b) Repeat part a if you have programmed your POP client in download and delete mode.

c) Suppose five minutes later, you again access POP to retrieve a new e-mail. Suppose that in the five-minute interval, no new messages have been sent to you. Provide a transcript of this second POP session for both options a and b above.

Student Questions

©2026 Raj Jain



Domain Name Service

1. DNS Hierarchy
2. How DNS Works?
3. DNS Records
4. DNS Message Format
5. DNS Registration
6. DNS Vulnerability

Student Questions

- How does DNS work when the service is distributed across multiple servers? How does it know to resolve to the closest?

See Slide 2-42

- How does DNS translate a hostname to IP?

See Slide 2-42

- To clarify, can you send an email from a non-existing mail server (i.e., president@whitehouse.com)?

Yes. But nowadays, the servers authenticate the user and verify the domain name.

- Some mail servers still don't care what from address is used. If I claim to have sent mail from an address, would the actual owner of that address be able to find out (without the recipient telling them)?

Yes, only the name is faked. The IP address is still your IP address.

- Is it risky to change DNS? *Yes.*
- Does DNS perform a handshake?

No. But Secure DNS uses TSL that may do a handshake.

©2026 Raj Jain

DNS

- Domain Name Service
- DNS servers translate a hostname to IP address
  E.g., www.wustl.edu ⇒ 128.252.87.149
- A distributed database of all hosts in the universe
- Other Services:
  - **Host Aliasing**: www.rajjain.com or www.cse.wustl.edu/~jain/
  - **Mail Server Aliasing**: MX record (e.g., jain@wustl.edu)
  - **Load Distribution**: Multiple addresses, rotated

Student Questions

- Does host Aliasing save different URLs that direct to the same website?

Host aliasing simply provides the same IP address for different names.

- Do DNS servers violate your privacy by knowing every site you wish to visit?

Yes.

- Could you please explain more about the idea of a distributed database?

See Slide 2-41

- We can get an IP address by DNS. But how can we get the port number? Should I use the default port number if I want to talk with a server?

Yes. Unless the serving company told you to use another port #.

- Are the host aliasing and load distribution mean that the hostname and IP address are not one-to-one mappings?

Yes. Google.com is thousands of servers.

- Are DNS numbers fixed?

What is a DNS number?

- Why do websites have multiple IP addresses?

For load distribution

©2026 Raj Jain

DNS

- Domain Name Service
- DNS servers translate a hostname to IP address
  E.g., www.wustl.edu ⇒ 128.252.87.149
- A distributed database of all hosts in the universe
- Other Services:
  - **Host Aliasing**: www.rajjain.com or www.cse.wustl.edu/~jain/
  - **Mail Server Aliasing**: MX record (e.g., jain@wustl.edu)
  - **Load Distribution**: Multiple addresses, rotated

Student Questions

- Does host Aliasing save different URLs that direct to the same website?

Host aliasing simply provides the same IP address for different names.

- Do DNS servers violate your privacy by knowing every site you wish to visit? *Yes.*
- Could you please explain more about the idea of a distributed database?

See Slide 2-41

- We can get an IP address by DNS. But how can we get the port number? Should I use the default port number if I want to talk with a server? *Yes. Unless the serving company told you to use another port #.*
- Can you elaborate on how DNS helps Mail Server Aliasing?

DNS also provides the name of the mail server

- How does DNS resolve one hostname to different IP addresses in the context of load distribution?

DNS rotates among multiple addresses

- Is DNS an injective function for hostname and IP address? Can two different hostnames be translated into the same IP address? *Yes*

©2026 Raj Jain

DNS

- Domain Name Service
- DNS servers translate a hostname to IP address
  E.g., www.wustl.edu ⇒ 128.252.87.149
- A distributed database of all hosts in the universe
- Other Services:
  - **Host Aliasing**: www.rajjain.com or www.cse.wustl.edu/~jain/
  - **Mail Server Aliasing**: MX record (e.g., jain@wustl.edu)
  - **Load Distribution**: Multiple addresses, rotated

Student Questions

- Can one use a random email address and self-created for aliasing

It has to be registered on DNS

- Why are there different versions of google.com within St. Louis? What are the differences?

For load distribution. There are no differences.

- Who maintains/owns the DNS? How was it created?

Every organization owns a DNS for nodes on their network.

- Will the IP address be depleted or run out of? How can we make sure there are enough IP addresses we can use?

IP addresses were depleted in 1991. New methods have been devised, and IPv6 has been designed to provide more addresses.

- Can you override a DNS result by modifying the /etc/hosts file?

Yes. The host file is the first step in DNS.

- Are there other systems besides DNS for translating host names to IP addresses?

I am not aware of any.

©2026 Raj Jain

DNS

- Domain Name Service
- DNS servers translate a hostname to IP address
  E.g., www.wustl.edu ⇒ 128.252.87.149
- A distributed database of all hosts in the universe
- Other Services:
  - **Host Aliasing**: www.rajjain.com or www.cse.wustl.edu/~jain/
  - **Mail Server Aliasing**: MX record (e.g., jain@wustl.edu)
  - **Load Distribution**: Multiple addresses, rotated

Student Questions

- How does DNS update its database?

All changes travel from leaves towards the root.

- Sometimes, a hostname cannot be retrieved from the correct IP address in certain countries; we call this DNS poisoning.

DNS poisoning is the act of putting misinformation in DNS. Attackers (and, in some cases, governments) use this.

©2026 Raj Jain

DNS Example

F:\>**nslookup www.wustl.edu**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Name: www.wustl.edu
Address: 128.252.87.149

F:\>**nslookup www.google.com**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Non-authoritative answer:
Name: www.l.google.com
Addresses: 74.125.225.48, 74.125.225.52, 74.125.225.50, 74.125.225.49
74.125.225.51
Aliases: www.google.com

Student Questions

- Is 'www' the hostname of the DNS server?

No. www is usually the hostname of a Web server. It could be anything else too.

- Can you explain the difference between sockets and ports again? Are they on different layers?

Sockets are in the operating system. Ports are for the transport layer.

- What part is the www in www.facebook or wustl in wustl.instructure.com? Is it another child authoritative server below facebook or an instructure authoritative server?

These are names in those domains. Wustl.instructure.com is a name in instructure.com. Several names may map to a single server at instructure.com

- A data center has a lot of hosts. Do these hosts share the same IP?

Not necessarily. Address allocation is discussed in Chapter 4.

- Where can I find my public IP?

Whatismyip.com

©2026 Raj Jain

DNS Example

F:\>**nslookup www.wustl.edu**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Name: www.wustl.edu
Address: 128.252.87.149

F:\>**nslookup www.google.com**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Non-authoritative answer:
Name: www.l.google.com
Addresses: 74.125.225.48, 74.125.225.52, 74.125.225.50, 74.125.225.49
74.125.225.51
Aliases: www.google.com

Student Questions

- Do all hosts under a Wi-Fi network share the same public IP?

Not necessarily. Address allocation is discussed in Chapter 4.

- If I do nslookup when I connect to my Mobile Hotspot, the server address does not look like an IP. It is not 32 bits. Why is that?

IPv6 addresses are 128-bit long and written in hex.

- Is there anything useful we can do with the server address?

You need the address to get there and do useful things.

- What is "nslookup" used for?

Name Server Lookup

- Is "www" same as "ns00.ip‘’? If so, why don’t we use "ns00.ip" instead of “www?”

A computer can have many names. WWW is another name for the same computer. You can use any.

- Why are there so many IP addresses for www.google.com?

It is a group of servers distributed all over the world.

©2026 Raj Jain

DNS Example

F:\>**nslookup www.wustl.edu**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Name: www.wustl.edu
Address: 128.252.87.149

F:\>**nslookup www.google.com**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Non-authoritative answer:
Name: www.l.google.com
Addresses: 74.125.225.48, 74.125.225.52, 74.125.225.50, 74.125.225.49
74.125.225.51
Aliases: www.google.com

Student Questions

- What happens if an authoritative DNS record changes, rendering a non-authoritative cache entry obsolete?

Your connection fails, and you query DNS again.

- Is there a system for recursive query or timed expiration of cached records, or is it recalculated when the connection to the obsolete address fails?

See above

- What is the main purpose for using www1, www2, www3, etc. instead of just www in a domain?

If you have multiple servers, you can name them www1 or www2…. You can also use one name and distribute queries through some method.

©2026 Raj Jain

DNS Example

F:\>**nslookup www.wustl.edu**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Name: www.wustl.edu
Address: 128.252.87.149

F:\>**nslookup www.google.com**
Server: ns00.ip.wustl.edu
Address: 128.252.0.1

Non-authoritative answer:
Name: www.l.google.com
Addresses: 74.125.225.48, 74.125.225.52, 74.125.225.50, 74.125.225.49
74.125.225.51
Aliases: www.google.com

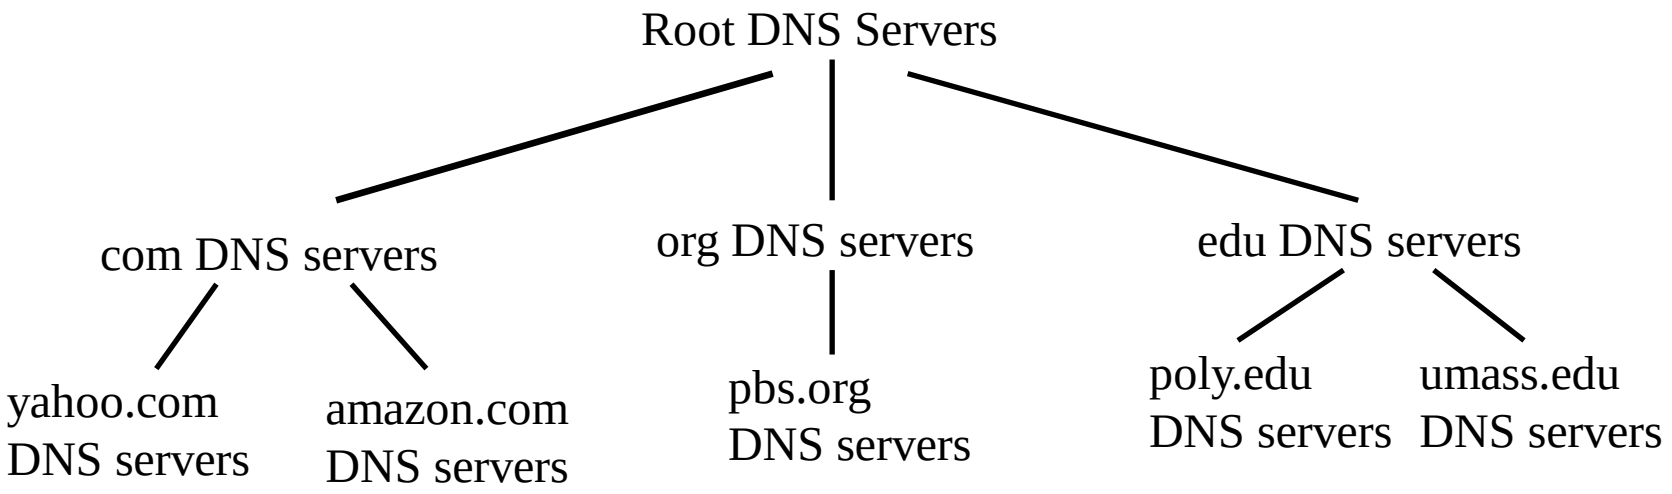
Student Questions

- Are the 256^4 addresses enough for the servers' public addresses across the world? If not, how do we solve this problem?

Discussed in Chapter 4-5 on IP, when we talk about addressing.

©2026 Raj Jain

DNS Hierarchy



- Root DNS Servers
- Top-level Domain (TLD) servers
- Authoritative DNS Servers

Student Questions

- What is the "seas." prefix that's discussed?

SEAS=School of Engineering and Applied Science at WUSTL. Now renamed to McKelvey School of Engineering.

- Why is the authoritative DNS server umass? How does that relate to WUSTL?

Umass will be the authoritative server for names ending in Umass.edu and not for wustl.edu.

- Do TLDs actually change anything about how the protocol or data is managed? Like, are .gov and .edu servers managed differently?

No. They manage more names, and so need more computer capacity.

- What organization manages the root DNS server?

Internet Assigned Numbers Authority (IANA)

- So if we want to give a name like 'seas.', should we have a DNS server?

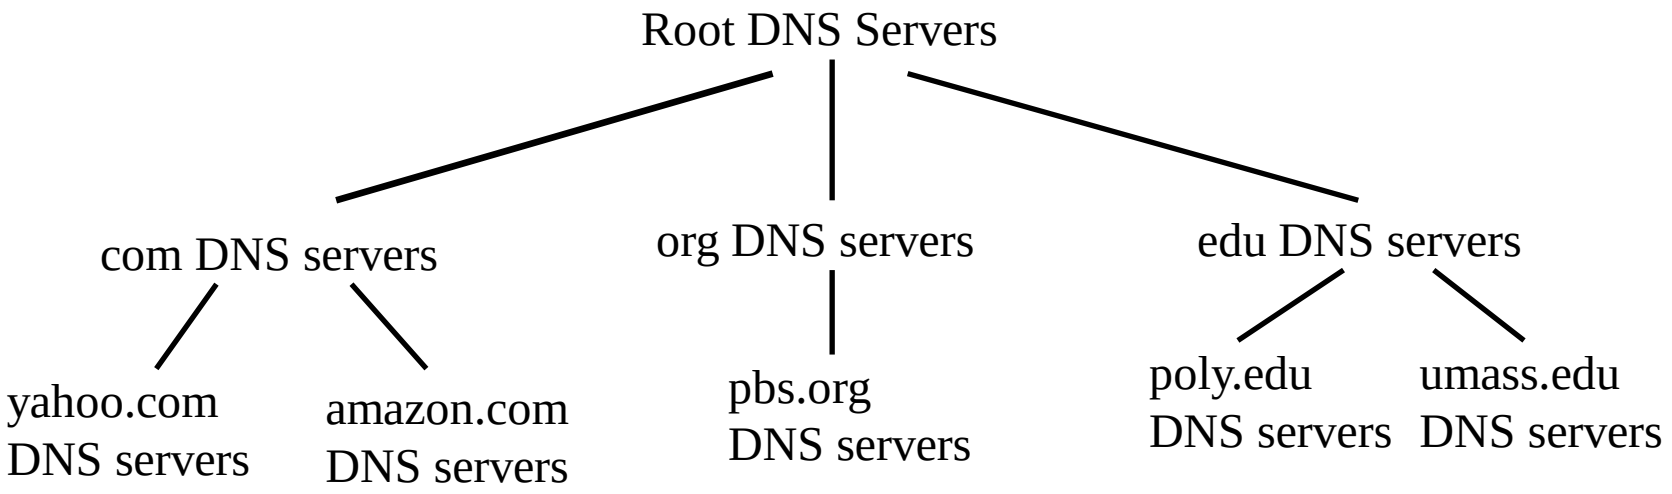
Seas.wustl.edu would resolve all names ending in seas.wustl.edu

- Where are the TLD servers located in DNS hierarchy

Thirteen places which volunteered to provide the service.

©2026 Raj Jain

DNS Hierarchy



- Root DNS Servers
- Top-level Domain (TLD) servers
- Authoritative DNS Servers

Student Questions

- Will the 'poly.edu' DNS server still can have its children's DNS server?

Yes. There are no restrictions on the levels of hierarchy.

- Do DNS servers in the same hierarchy interact with each other? Is it only at the root level?

Anyone can talk to anyone without going through the root.

- The textbook shows that the countries with the most root servers are the US, Canada, Brazil, India, and Germany as of 2020. What factors decide where the root servers go?

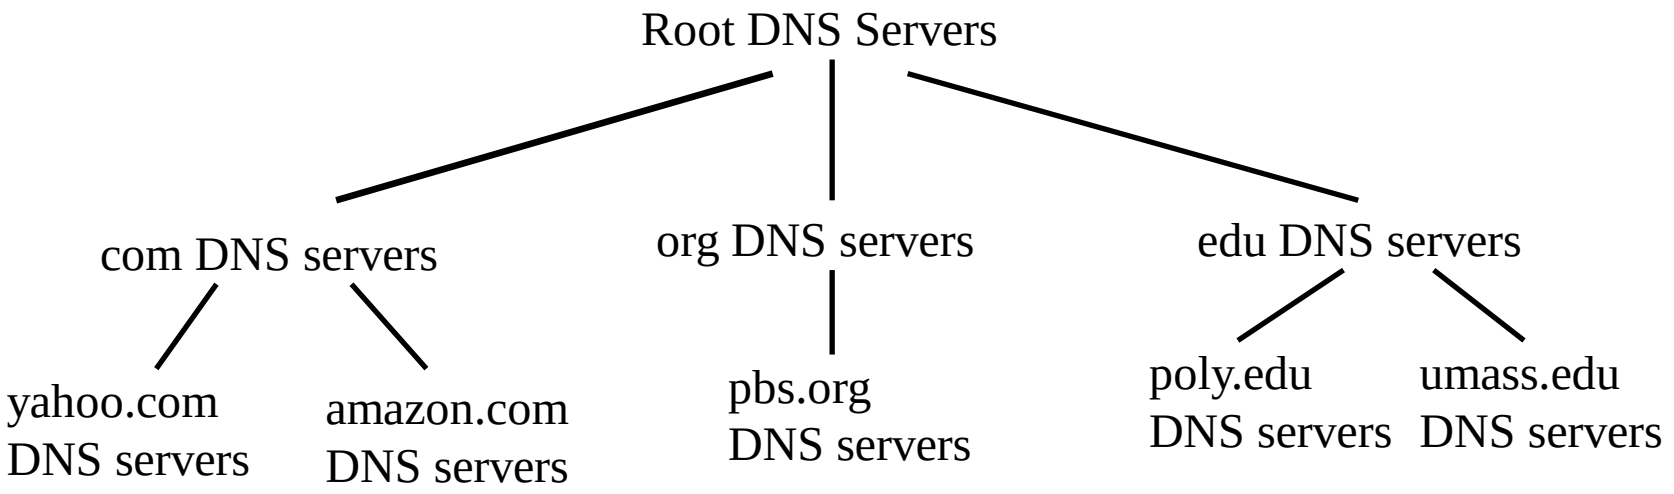
Internet Assigned Number Authority (IANA) decides.

- What is the difference between com DNS servers and other DNS servers, like edu and org? Why the com DNS servers are the most common ones?

No difference. Classes were based on the type of organization: Commercial, educational, etc.

©2026 Raj Jain

DNS Hierarchy



- Root DNS Servers
- Top-level Domain (TLD) servers
- Authoritative DNS Servers

Student Questions

- Are com DNS servers, ... edu DNS servers, and yahoo.com DNS servers, ... umass.edu DNS servers listed in the graph all TLD servers? What are considered higher and lower-level DNS servers?

.com and .edu are TLD. Umass.edu is not. The root is the highest level.

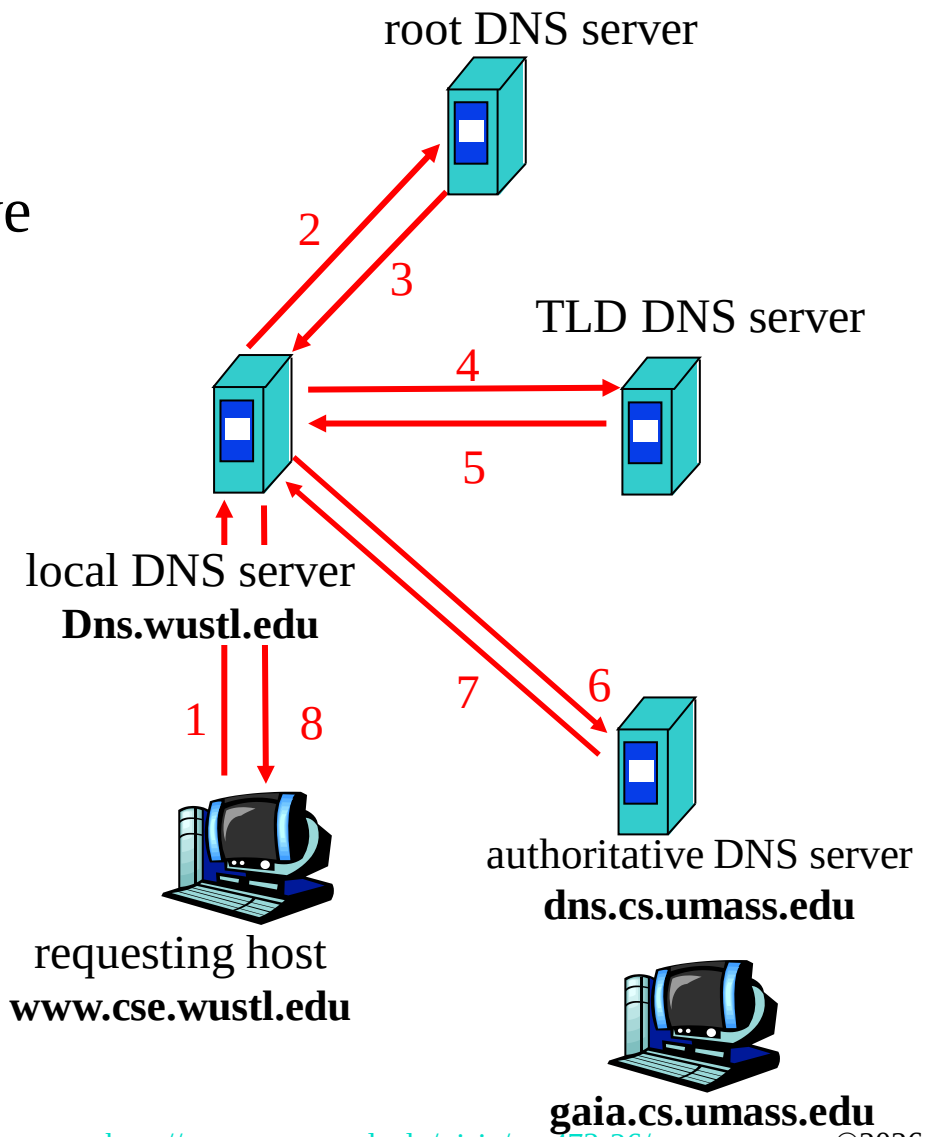
- Do people need approval for the use of specific TLDs? How does reputability work to ensure .gov or .edu is accurate?

You need to register the name with an authorized registrar. They will not let you register in unauthorized domains.

©2026 Raj Jain

How DNS Works?

- Redirects
- **Recursive queries**: Give me an answer
- **Iterative queries**: Give me an answer or a hint
- DNS responses are cached



Student Questions

- What's the typical cache policy of DNS servers? *Each entry has a lifetime.*
- Can you explain the example of the picture again? What exactly is being requested, and why does UMass provide an answer to WUSTL?

A server in WUSTL needs the address of gaia.cs.umass.edu.

- How does the computer choose between sending recursive queries or iterative queries?

Lower-level users just want the answer. Higher-level DNS servers may use the hint.

- Why would you want a hint and not an answer? Can you explain this part again?

You may want to resolve many addresses in that domain.

- What may cause iterative queries?

You may want to resolve many addresses in that domain.

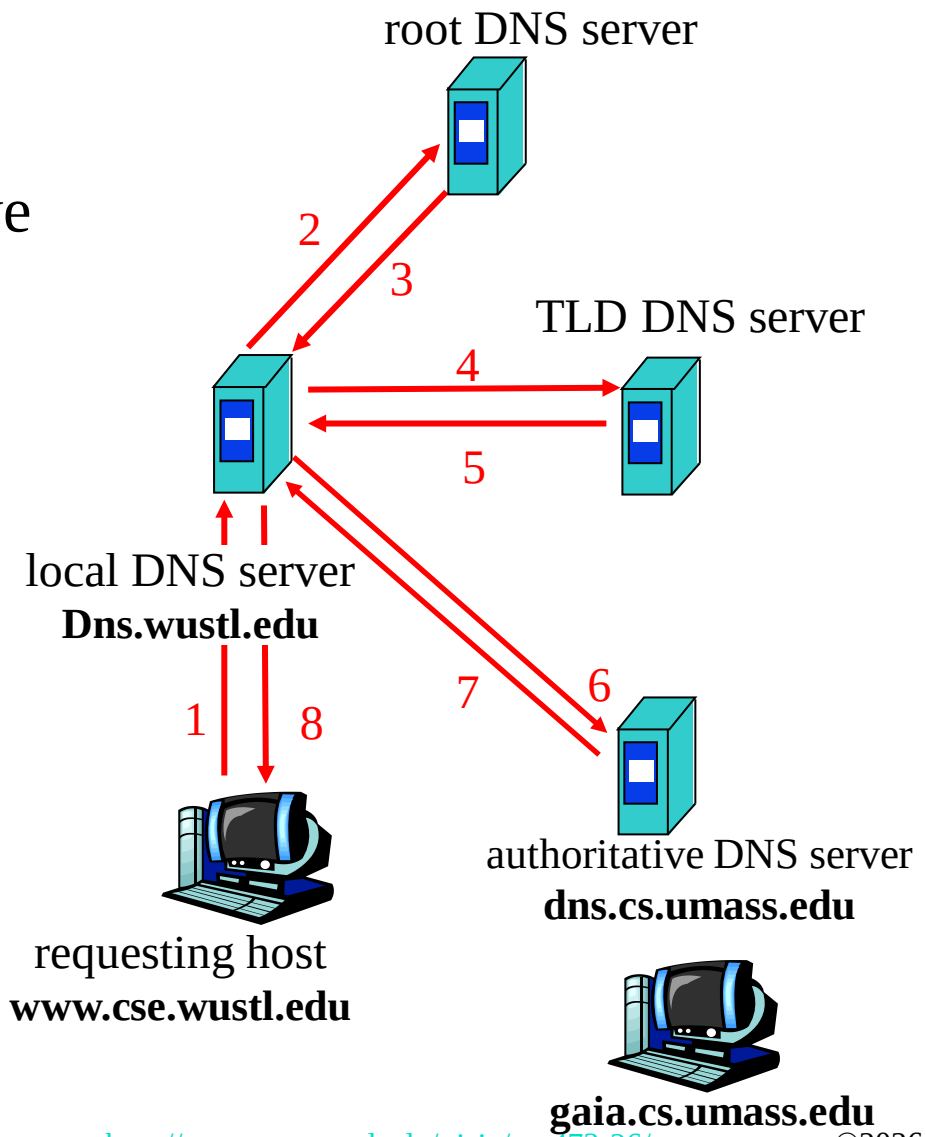
- What's the cache policy of the DNS server?

Controlled by the network manager. Refresh after a while or refresh only if unreachable.

©2026 Raj Jain

How DNS Works?

- Redirects
- **Recursive queries**: Give me an answer
- **Iterative queries**: Give me an answer or a hint
- DNS responses are cached



Student Questions

- What's the difference between a local DNS server and an authoritative DNS server?

Local is near you. Authoritative is near the name you are trying to resolve.

- Which of the requests on the picture are recursive, and which are iterative?

Anyone can be iterative or recursive. Most clients will make recursive requests. Most servers will make an iterative request.

- Can edu hosts query edu DNS servers for xx.com hostnames? Do they only query xx.edu?

Anyone can query any server

- By recursive, does that mean (let's say) Server A querying Server B will cause Server B to make its query to Server C and so on?

Yes.

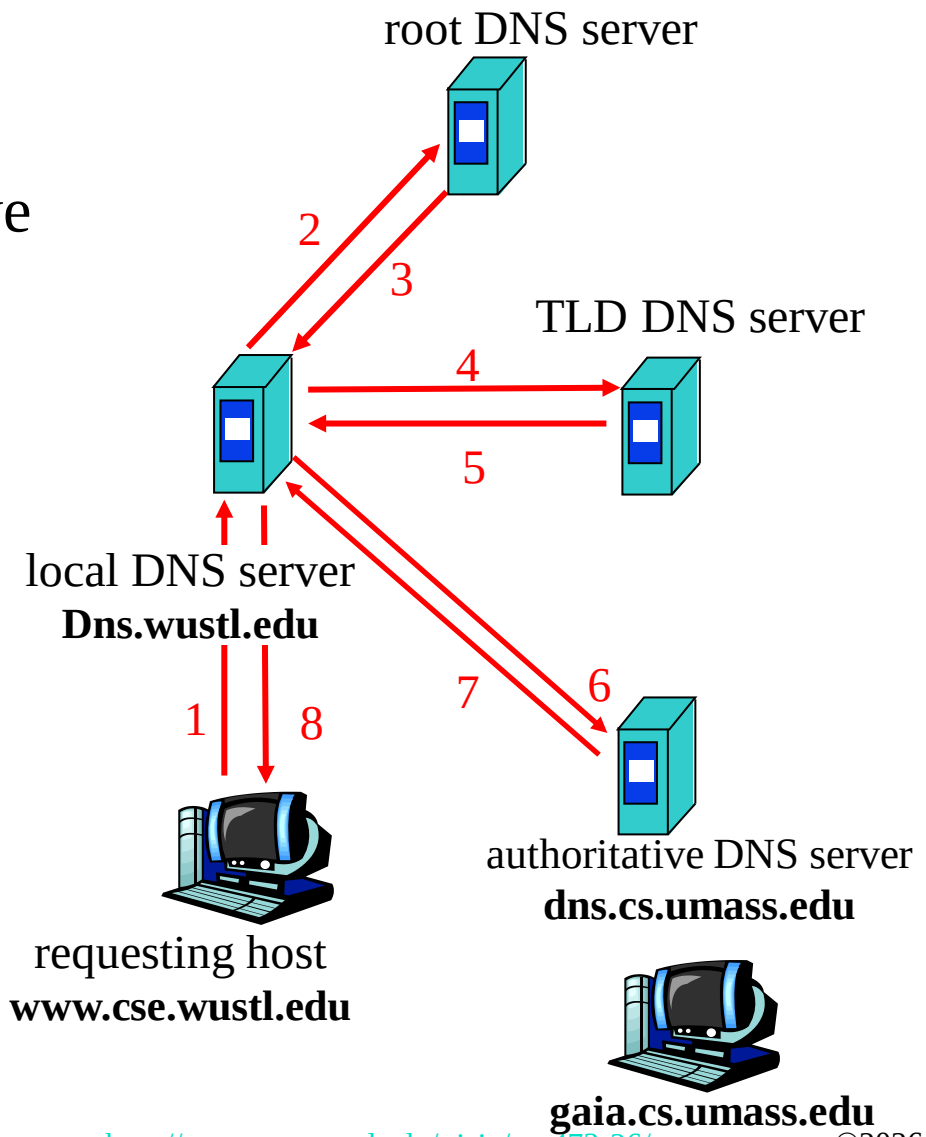
- [Book p.130] What determines whether DNS queries are recursive or iterative?

Low-power devices prefer recursive. DNS servers may use iterative.

©2026 Raj Jain

How DNS Works?

- Redirects
- **Recursive queries**: Give me an answer
- **Iterative queries**: Give me an answer or a hint
- DNS responses are cached



Student Questions

- In what situations are iterative queries used over recursive queries? Why is it better at all?

Servers generally use iterative queries. End clients generally use recursive queries. A server may reach there faster than waiting for another server.

- Does DNS mix recursive queries and iterative queries?

Yes. See above.

- Can there be multiple instances of a TLD server or root server, say one for each region?

Yes, always.

- What are the differences and interactions among Redirects, Recursive queries, and Iterative queries in the DNS resolution process?

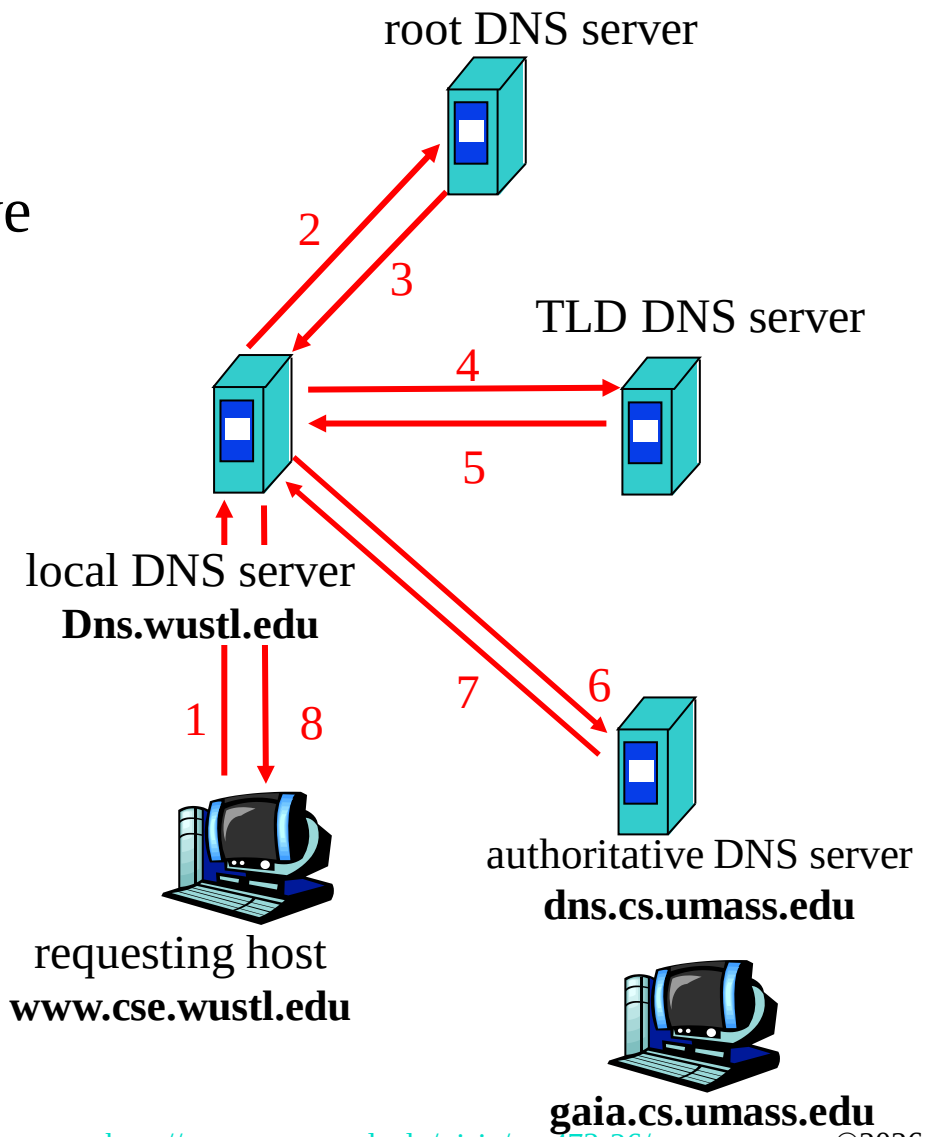
See the first question above for Interactive vs recursive. Redirect is a response, not a query.

- Is the local DNS server hosted by an ISP?

Not necessarily. Anyone can host it. Users decide which DNS to use.

How DNS Works?

- Redirects
- **Recursive queries**: Give me an answer
- **Iterative queries**: Give me an answer or a hint
- DNS responses are cached



Student Questions

- In the pictures, which queries are recursive and which are iterative? Can these two types of queries be converted into one another?

1 is recursive. 2, 4, 6 are iterative.

- I try to use nslookup on wustl.edu and washu.edu, but I find that I can't go to website through the IPv4 address 23.185.0.3 and also the IPv6 address, why?

Multiple domains are often hosted on a single IP address. The server uses the domain name to determine which website to serve. Accessing the IP address directly won't provide the necessary domain information.

©2026 Raj Jain

DNS Records

- Resource Records=(Name, Value, Type, TTL)
- Type=A: IP Address for the host name
- Type=NS: Name server for the domain name
- Type=CNAME: Canonical name for a host name
- Type=MX: Canonical name of mail server

Student Questions

- Say if a mail system and web page both use DNS. Then how does DNS differentiate between these two requests?

The answer is returned to the sending port.

- Is there a specific reason mail servers use type MX? If it is a canonical name, why not just use type CNAME?

The mail server for cse.wustl.edu may be different from the Web server. Each computer has a canonical name.

- How can DNS records keep updated?

Routers that assign addresses automatically keep them updated.

- What is the difference between a domain name and a hostname?

A domain may have many subdomains and hosts.
DELLXPS.Home.rajjain.com

- Can A record provide the hostname of the mail server? [Book p. 132]

No MX records.

©2026 Raj Jain

DNS Records

- Resource Records=(Name, Value, Type, TTL)
- Type=A: IP Address for the host name
- Type=NS: Name server for the domain name
- Type=CNAME: Canonical name for a host name
- Type=MX: Canonical name of mail server

Student Questions

- I heard I cannot set a CNAME record on the zone's apex (or root). Why?

A domain can have a subdomain—for example, cse.wustl.edu. Wustl.edu is the root or apex. You can redirect cse.wustl.edu to another name, e.g., compsci.wustl.edu, and then move compsci servers around as needed. DNS does not allow such redirections for root names. But you can move them by changing their IP addresses using type A records.

- What do you mean by name server?

DNS is the protocol. A name server is a node.
wustl.edu=23.185.0.3
ns.wuslt.edu=128.252.120.1

©2026 Raj Jain

DNS Message Format

- **Questions**: Name, type
- **Answers**: Name, type, value, TTL
- **Authority**: Other authoritative servers
- **Additional**: Other information, e.g., IP address of canonical name in MX response

Identification	Flags	12 Bytes	12:34	01:00
# of Questions	# of Answer RRs		00:01	00:00
# of Authority RRs	# of Additional RRs		00:00	00:00
Questions			07:65:78:61:6D:70:6C:65:03:63:6F:6D:00 :00 01:00 01 (Name: example.com, Type A, Class Internet)	
Answers				
Authority				
Additional Information			Additional Information	

Student Questions

- What is the authority section used for in most cases? What is an example?

Seas.wustl.edu the authority for all seas.wustl.edu nisames. But WUSTL.edu may also serve as an authority.

- What is the relationship between DNS Records and the DNS Message Format?

Records are kept on the disk. Messages are sent over the network. Messages need additional headers.

- Is this the format of the query or the response

Both

- Is it possible for multiple DNS servers to return different answers, and if so, how is that resolved?

Authoritative server's answer is final.

- The first 12 bytes are fixed. So the rest can be any amount?

16 bits ⇒ Max 2^{16}-1 RRs

- What is meant by "resource record"?

Items

©2026 Raj Jain

DNS Message Format

- **Questions**: Name, type
- **Answers**: Name, type, value, TTL
- **Authority**: Other authoritative servers
- **Additional**: Other information, e.g., IP address of canonical name in MX response

Identification	Flags	12 Bytes	12:34	01:00
# of Questions	# of Answer RRs		00:01	00:00
# of Authority RRs	# of Additional RRs		00:00	00:00
Questions			07:65:78:61:6D:70:6C:65:03:63:6F:6D:00 :00 01:00 01 (Name: example.com, Type: A, Class: Internet)	
Answers				
Authority				
Additional Information			Additional Information	

Student Questions

- Do we have to have all the abbreviations memorized?

Yes, common ones.

- Authority is another authoritative server, so when do we need to put other authoritative servers here except our main authoritative server?

When there are many servers, or when you get the answer from another server.

- How many bytes is each field?

You will need to refer to DNS RFC.

- Is there an example of what this DNS messages might look like in real life, or could you provide an example of what could be filled in the boxes?

See the example on the left.

- Can the fields of a DNS message be empty?

Yes. But often, it will have an error code rather than empty.

©2026 Raj Jain

DNS Registration

- Many Registrars
- Internet Corporation for Assigned Names and Numbers (ICANN) accredits registrars
- www.internic.net
- Registrars provide authoritative name servers, A and MX records for the domain.

Student Questions

- Can you re-explain DNS cache poisoning?

Some server gives wrong answers to a DNS request. The client sends its packets to the wrong address until it decides to check the address again.

- Does it cost money for ICANN to accredit registrars or for the Registrars to provide you with authoritative name servers?

There are always management and administration costs for all services.

©2026 Raj Jain

DNS Vulnerability

- Distributed Denial of service attack on Name server
- DNS cache poisoning – A server gives the wrong answer

Student Questions

- Could you explain how DNS cache poisoning works again?

You get a wrong answer from someone. Then you keep sending your packets to the wrong place.

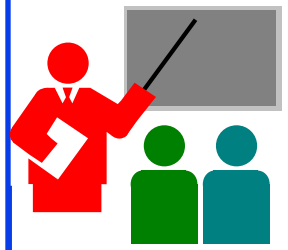
- How could we avoid it

Secure DNS

- Are there any prerequisites when carrying out a DNS cache poisoning attack?

Now, it isn't easy since everyone is authenticated.

©2026 Raj Jain



DNS: Summary

1. DNS is used to **resolve names** to IP address
2. Also provides Name aliasing (CNAME), Mail Server (**MX**) records
3. DNS is a distributed database
   ⇒ Servers ask other servers for answers when needed
4. **Recursive** (answer only) or **iterative** (answer or hint) queries
5. **Root** Servers, **Top level domain** servers, **Authoritative** servers

Ref: Read Section 2.4 Full.

Student Questions

- Does DNS consist of a single database, or could there be multiple databases?

Multiple databases. However, there is only one authoritative database.

- What would happen if I changed the DNS of my PC?

You will be sending your queries to this different server.

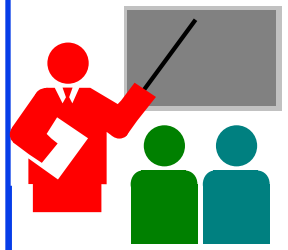
- As far as I understand, every OS has a DNS cache. Would disabling it also disable access to the world wide web (meaning the browsers or something else strongly depends on the cache), or would it just make accessing it a lot harder?

You can resolve names without caching them just by asking someone else.

- Is there any enforcement mechanism ensuring DNS servers translate names correctly? *Secure DNS*
- Why we could change our DNS

Because some DNSs are faster.

©2026 Raj Jain



DNS: Summary

1. DNS is used to **resolve names** to IP address
2. Also provides Name aliasing (CNAME), Mail Server (**MX**) records
3. DNS is a distributed database
   ⇒ Servers ask other servers for answers when needed
4. **Recursive** (answer only) or **iterative** (answer or hint) queries
5. **Root** Servers, **Top level domain** servers, **Authoritative** servers

Ref: Read Section 2.4 Full.

Student Questions

- When the requesting host sends a request, does it send it to only one local DNS server, or does it send multiple queries and wait for the first response to come back? If not multiple, how does requesting host know which local DNS server to ask?

Every computer has a default DNS server and a backup server.

- Is the Authoritative Server like the original author of a masterpiece instead of a holder of a copy of the work?

Yes.

- How does the distributed structure of DNS enhance domain resolution efficiency?

Impossible to centralize it since every company maintains their networks, computers, and name.

©2026 Raj Jain

Homework 2C: DNS

- [4 points]
- Is it possible for an organization's web server and mail server to have exactly the same hostname?
- What would be the type of RR that contains the host name of the mail server?

Student Questions

- What is RR?

Resource record

©2026 Raj Jain



Peer-to-Peer Applications

1. Client Server vs. P2P Scalability
2. P2P File Distribution (BitTorrent)

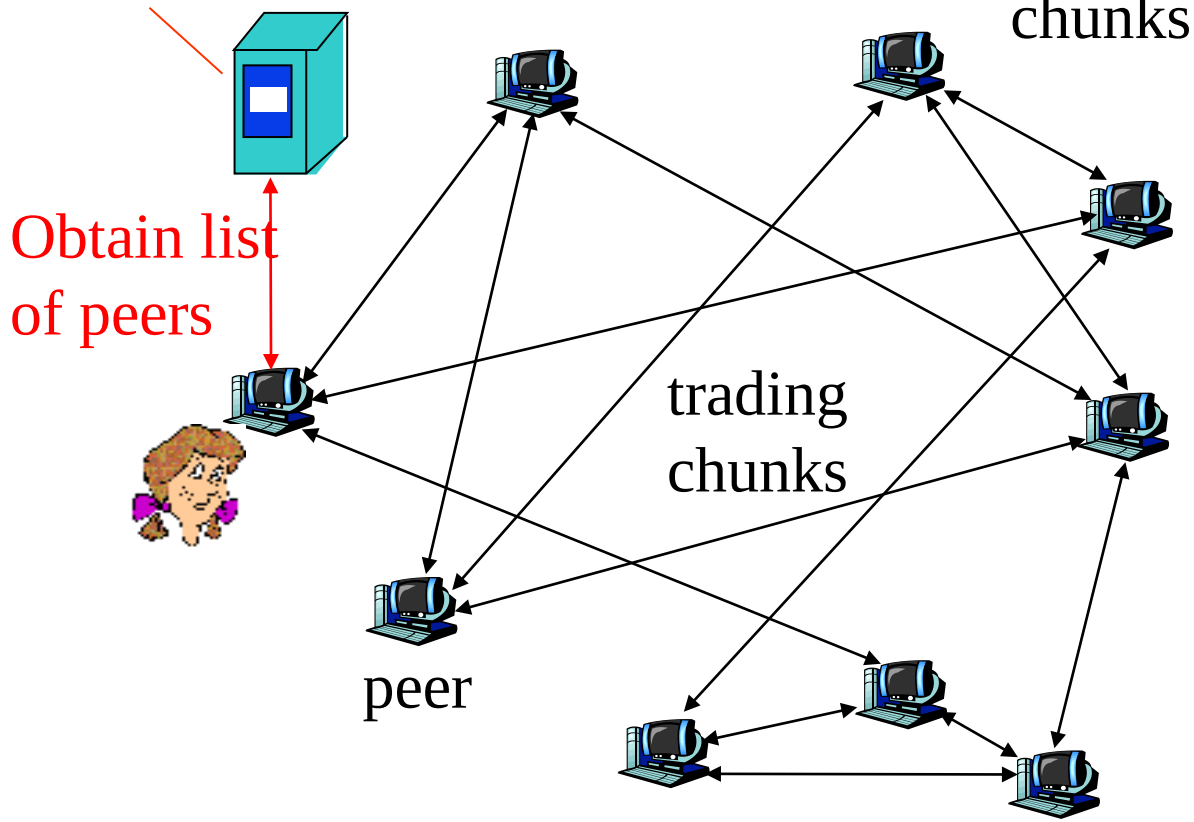
Student Questions

©2026 Raj Jain

P2P File Distribution (BitTorrent)

Tracker: tracks peers participating in torrent

Torrent: group of peers exchanging chunks of a file



Student Questions

- Can the communication in p2p be blocked or monitored by other clients?

Yes. They can become members of the torrent but throw away your requests.

- Is P2P vulnerable to attack because of its decentralized nature? *Security is independent of the P2P vs. Client-Server.*
- For peer-to-peer, if you have the file, can anyone get the file from you automatically, or do you have to take some action to make the file available to others? *You have to put it in the directory where your torrent client can access it. All files in the directory where downloaded files are kept are automatically served.*

- Can you further explain what BitTorrent is? I don't understand what that is.

It is a protocol to get files from the internet, generally movies, songs, or applications.

- Wouldn't the files from different nodes be slightly different (particular version of a book, movie in a different language, etc.)?

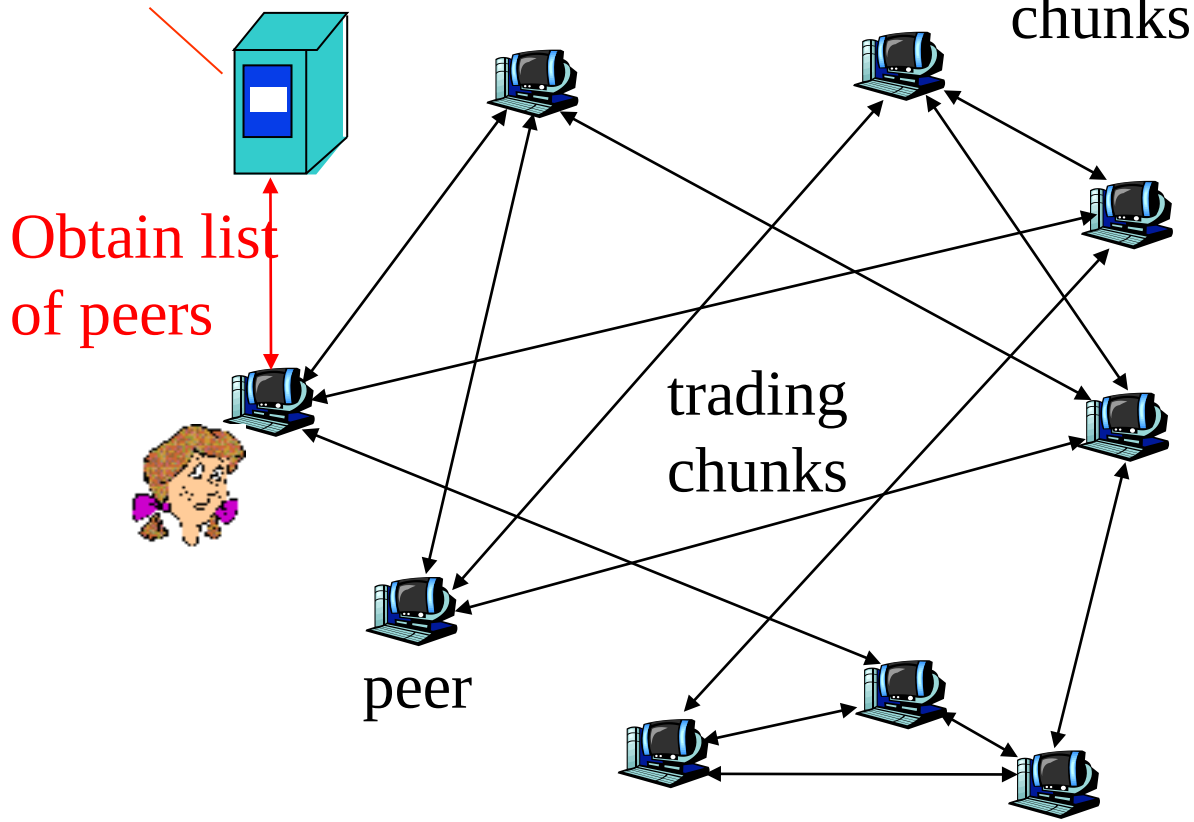
No. All files should have precisely the same hash. One byte change would make it a different torrent.

©2026 Raj Jain

P2P File Distribution (BitTorrent)

Tracker: tracks peers participating in torrent

Torrent: group of peers exchanging chunks of a file



Student Questions

- How do torrents differentiate between different hosts within the same private subnet? They share the same public IP address and may share other attributes like User-Agent as well.

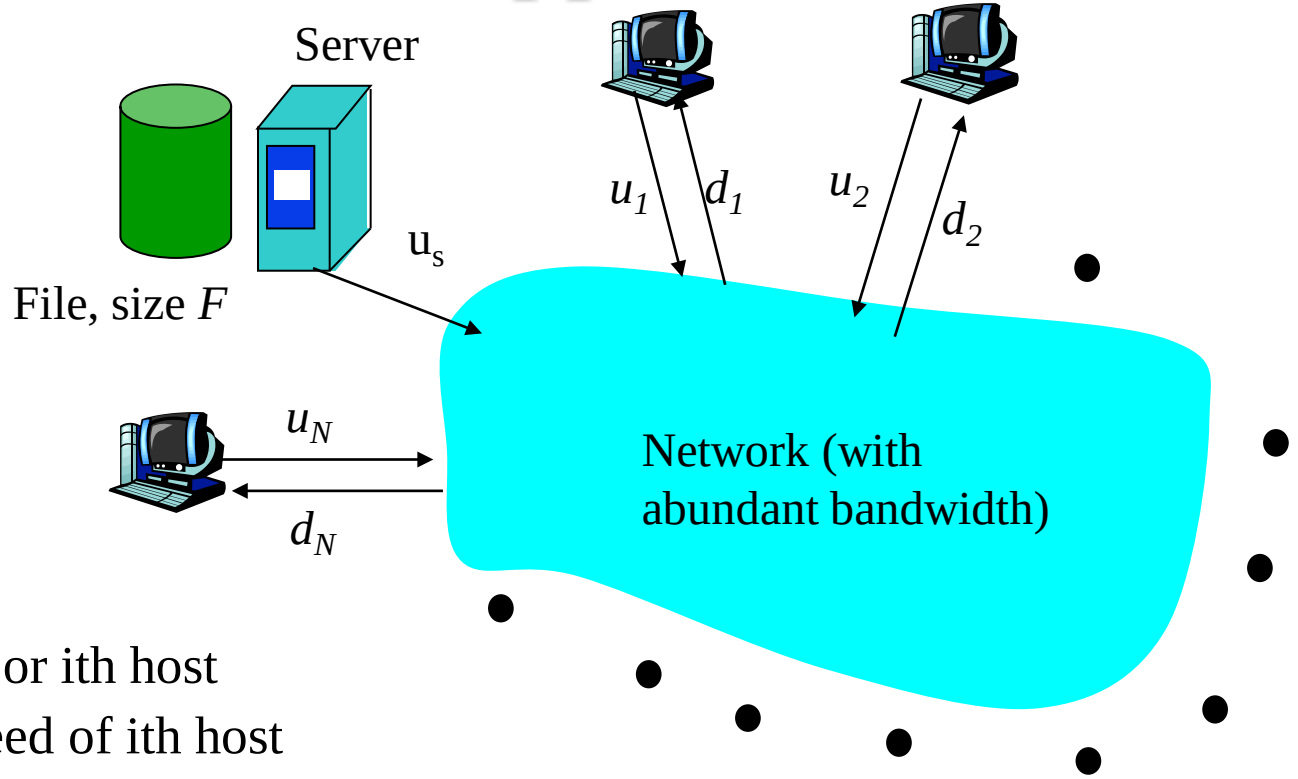
Port reservation guarantees that only one host gets the messages.

- How does the tracker in the BitTorrent process optimize file distribution efficiency among peers and manage potential challenges in peer-to-peer file sharing networks?

Load balancing via random distribution. Via geographical locations if known.

©2026 Raj Jain

Peer-to-Peer Applications



- N= # of peers
- F= File Size
- u_i= uplink speed or ith host
- d_i= downlink speed of ith host
- $d_{min}=\min\{d_1, d_2, ..., d_N\}$
- File distribution time using central server $D_{cs} \geq \max\{ NF/u_s, F/d_{min}\}$
  Server, Client
- File distribution time using P2P $D_{P2P} \geq \max\{F/u_s, F/d_{min}, NF/(u_s+\Sigma u_i)\}$
  Server, Client, Network

Student Questions

- Are we required to know the Asymptotic Complexities of Server Distribution Algorithms vs. P2P Algorithms?

Yes. These formulas should be on your cheat sheet.

- Why is it not NF/u_s + F/d_min ?

Which one?

- The book said D_cs is the distribution time; is this the same meaning as a total delay?

Yes. Total delay = time for giving the file to everyone = Distribution time

---

- If the file is uploaded once and downloaded by N users, why is DCS based on N * F/uplink speed? Shouldn't it be N * F / downlink speed since N users want to download the same file?

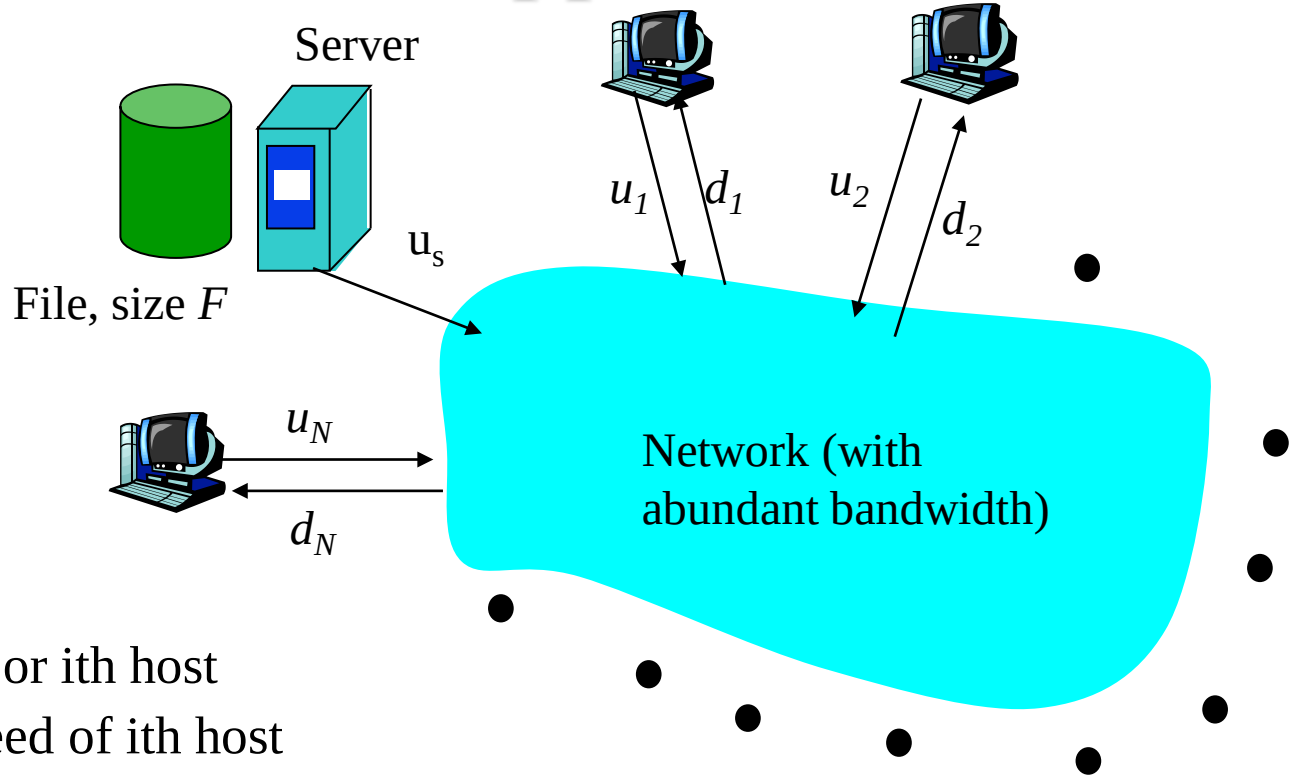
N users are downloading simultaneously. Only one user is uploading.

- Are we required to know the Asymptotic Complexities of Server Distribution Algorithms vs. P2P Algorithms?

Yes. These formulas should be on your cheat sheet.

©2026 Raj Jain

Peer-to-Peer Applications



- N= # of peers
- F= File Size
- u_i= uplink speed or ith host
- d_i= downlink speed of ith host
- $d_{min}=\min\{d_1, d_2, ..., d_N\}$
- File distribution time using central server $D_{cs} \geq \max\{ NF/u_s, F/d_{min}\}$
  Server, Client
- File distribution time using P2P $D_{P2P} \geq \max\{F/u_s, F/d_{min}, NF/(u_s+\Sigma u_i)\}$
  Server, Client, Network

Student Questions

- Why is it not NF/u_s + F/d_min ?

Which one?

- The book said D_cs is the distribution time; is this the same meaning as a total delay?

Yes. Total delay = time for giving the file to everyone = Distribution time

- For D_{P2P}, should it be NF/u_s instead of F/u_s?

No. See the last term. Upload is in parallel.

- Could you explain more about D_{cs} and D_{P2P}?

See updated text.

- Can you go over the math again for us?

Sure.

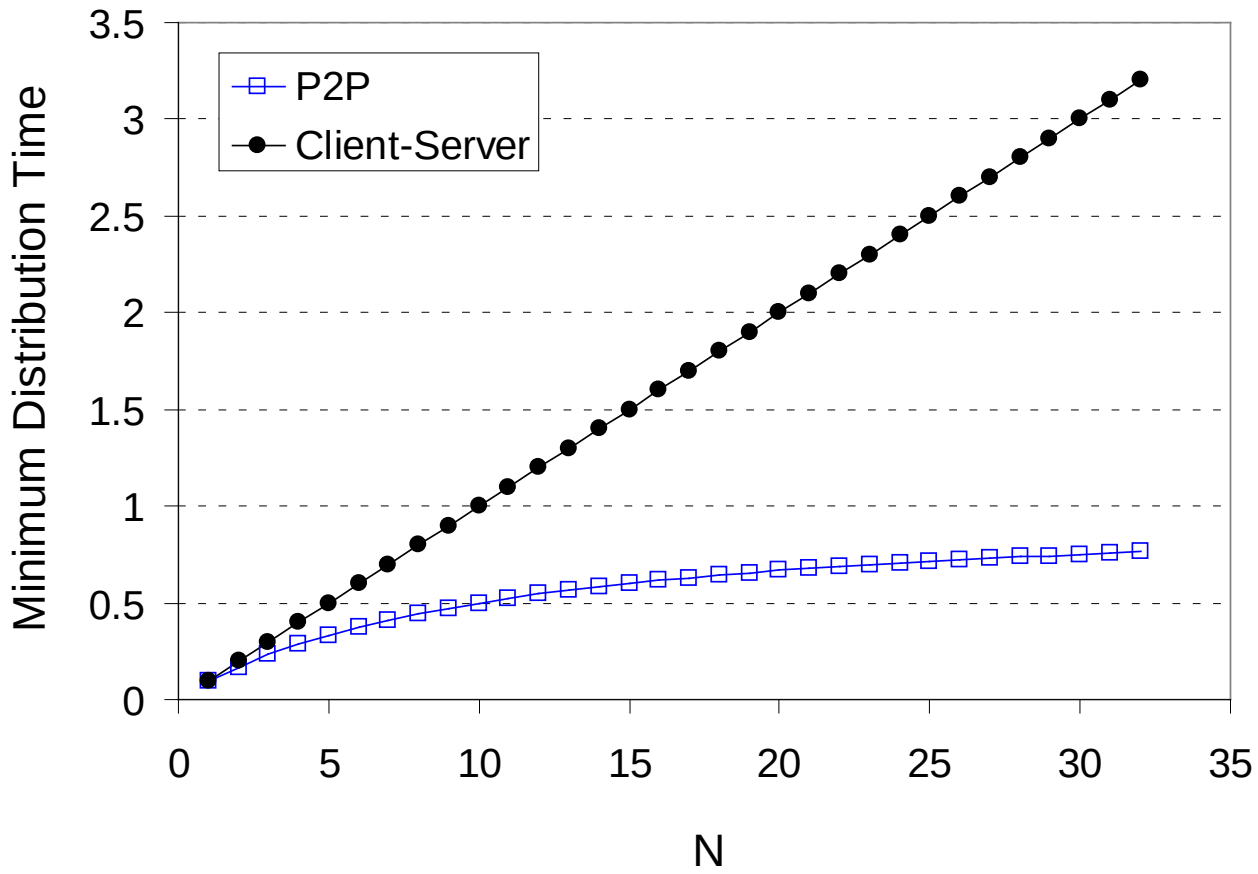
- Why does the distribution time take the maximum value of the server's upload time and peer's download time but not the minimum?

Distribution = last file completion

©2026 Raj Jain

Client Server vs. P2P Scalability

Client upload rate = u, F/u = 1 hour, $u_s = 10u$, $d_{min} \geq u_s$



Student Questions

- Can you go over scenarios in which it would be advantageous to use a client/server design rather than P2P?
- If Peer to peer networks are able to scale much better than the client-server model, Why do we still use the client-server model?

For companies, it is easier to manage central servers. Charging and billing are also easier with C/S than with P2P.

- P2P has a high scalability benefit over client-server. Is P2P only less popular because it has been used for nefarious purposes? *See above*
- What exactly is meant by "Distribution time" on the y-axis of the graph?

Time to finish the distribution of files to all clients.

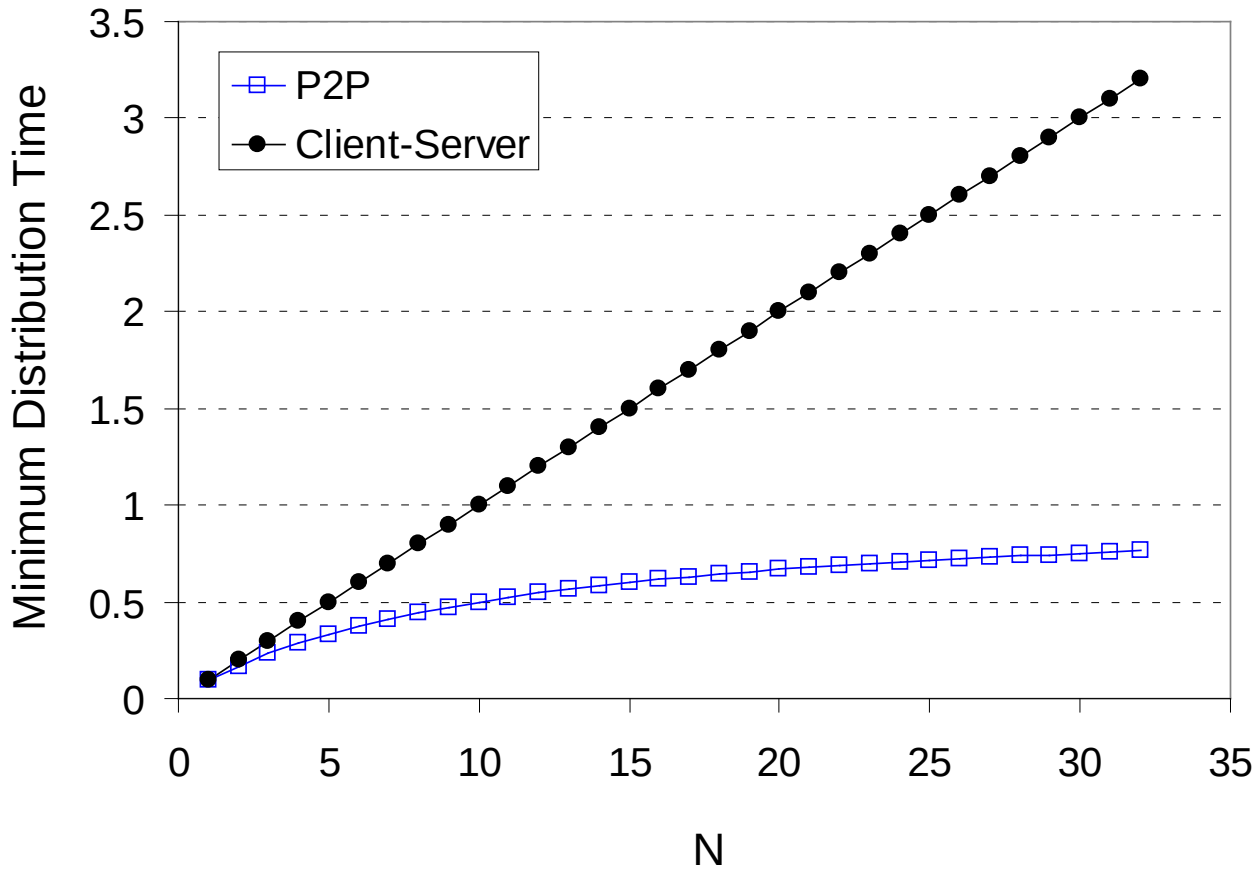
- Does P2P have a similar idea to blockchain?

Both are distributed but different.

©2026 Raj Jain

Client Server vs. P2P Scalability

Client upload rate = u, F/u = 1 hour, $u_s = 10u$, $d_{min} \geq u_s$

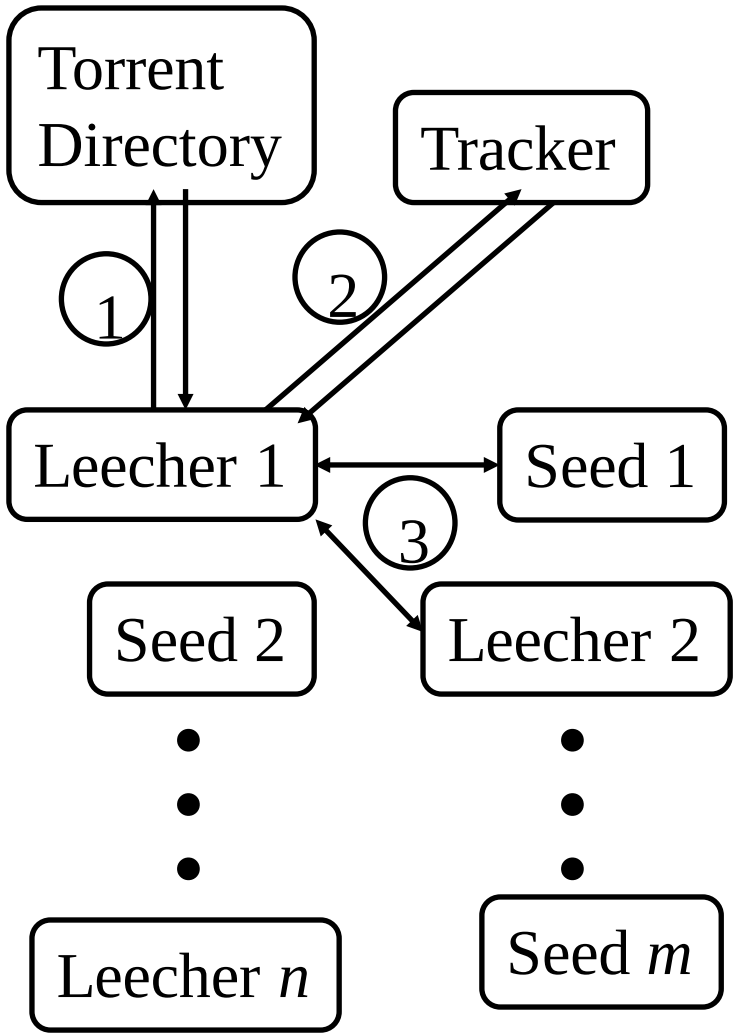


Student Questions

©2026 Raj Jain

BitTorrent P2P File Distribution

- **Peers**=nodes participating in a file distribution
- **Torrent**=Set of all peers
- **Torrent File** =a file containing information about the tracker, object ID, and file
- Files are segmented into equal size **chunks** (256KB)
- **Seeds**=Peers that have the complete file
- **Leechers**=Peers that have incomplete file
- **Tracker**=Has list of all peers



Student Questions

- Like this figure on the left, do Leechers have different roles in different positions?

Leechers are also servers.

- For a particular movie file, is the list of all peers changing once there are new seeds or leechers? If so, will the tracker change? And if the tracker changes, will we have different torrent files for this particular movie?

The torrent file is simply the "ID" of the file and a list of seeds + leechers. ID is simply a hash of the content.

- I am still a little bit confused about the function of the torrent directory and tracker.

Torrent directory list many files. The info stays the same as the list of peers for one file.

- How do P2P tracker records stay efficiently synced in real-time for large networks?

All clients inform tracker.

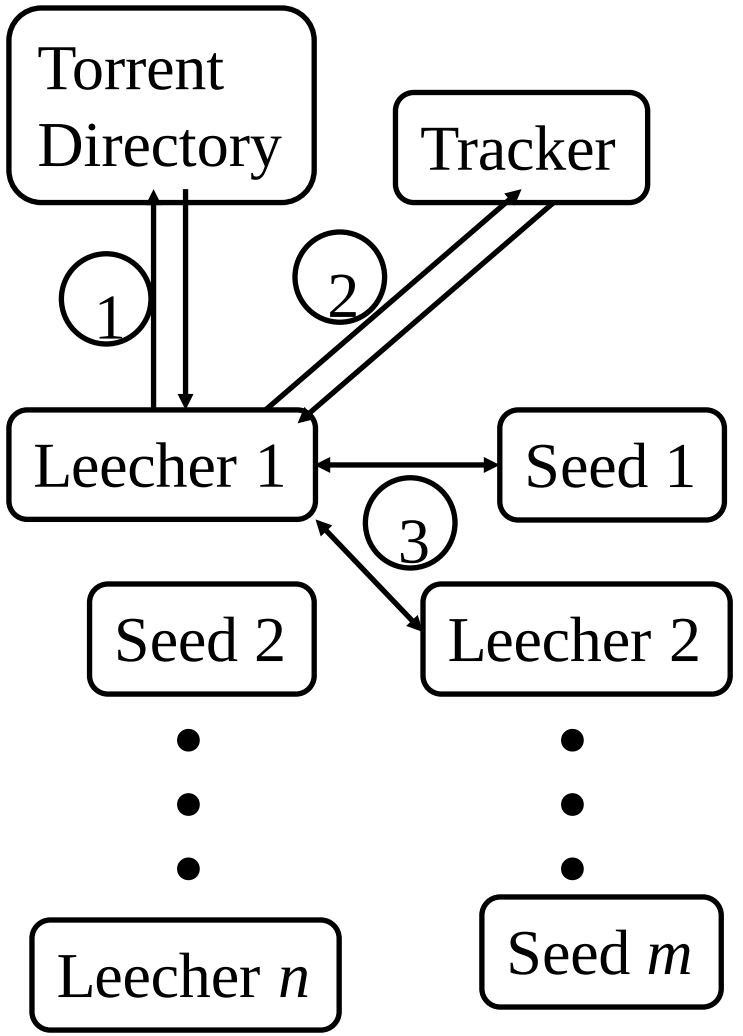
- Are seeds, leechers, and trackers all considered peers?

Seeds and leechers are peers.

©2026 Raj Jain

BitTorrent P2P File Distribution

- **Peers**=nodes participating in a file distribution
- **Torrent**=Set of all peers
- **Torrent File** =a file containing information about the tracker, object ID, and file
- Files are segmented into equal size **chunks** (256KB)
- **Seeds**=Peers that have the complete file
- **Leechers**=Peers that have incomplete file
- **Tracker**=Has list of all peers



Student Questions

- Why would someone be a leecher as opposed to a seed? *Leechers need complete files.*
- What if the file size is not the whole number of 256 when they are segmented into chunks?

Chunk size will indicate what the length is.

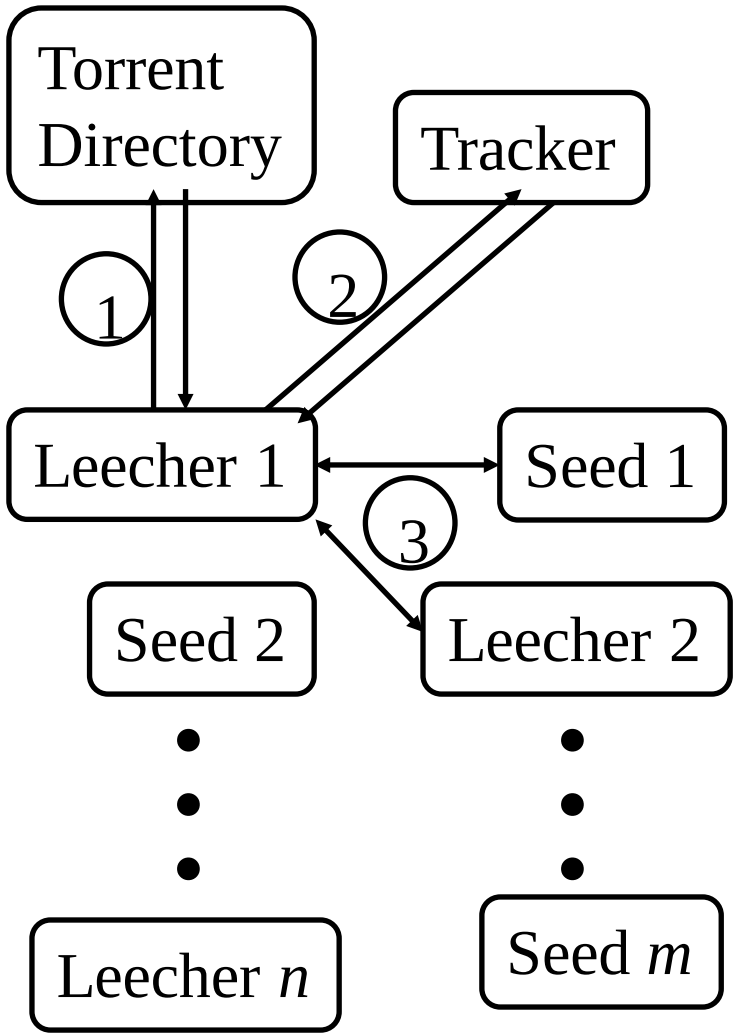
- Do seeds just let the application stay in memory all day so it can send files to others? *Yes, they want to be nice to others, or they may be downloading other files.*
- Is k 1000, and K is 1024 *Yes.*
- Will leechers become seeds after exchanging data with other leechers? *Yes.*
- Is the typical size 256kB or 256KB?

The chunk size is variable and is specified in the torrent file. A torrent creator could specify any size, a multiple of 256b since SHA-256 hash (which requires eight 32-bit words) is used. So both 256,000 (256kB) and 262,144 (256KB) are valid. However, since it is storage, 256KB would be more appropriate. Note that b is for bits. k and K, although both stand for kilo, are NOT the same. k=1000, K=1024.

©2026 Raj Jain

BitTorrent P2P File Distribution

- **Peers**=nodes participating in a file distribution
- **Torrent**=Set of all peers
- **Torrent File** =a file containing information about the tracker, object ID, and file
- Files are segmented into equal size **chunks** (256KB)
- **Seeds**=Peers that have the complete file
- **Leechers**=Peers that have incomplete file
- **Tracker**=Has list of all peers



Student Questions

- Please explain the graph again. Where is the tracker stored? Can it be used to trace back who is sharing illegal copyright?

Yes. Tracker is not sharing, but it is helping. Carriers can track who is downloading by simply looking at the protocol messages. However, everything is now encrypted.

- Can leechers also serve files, and if so how do other leechers know which part of the file they are being served?

You ask a leecher for its list of chunks.

- How Microsoft can get information from other computers near you when you download something?

All computers are using Windows that has P2P embedded.

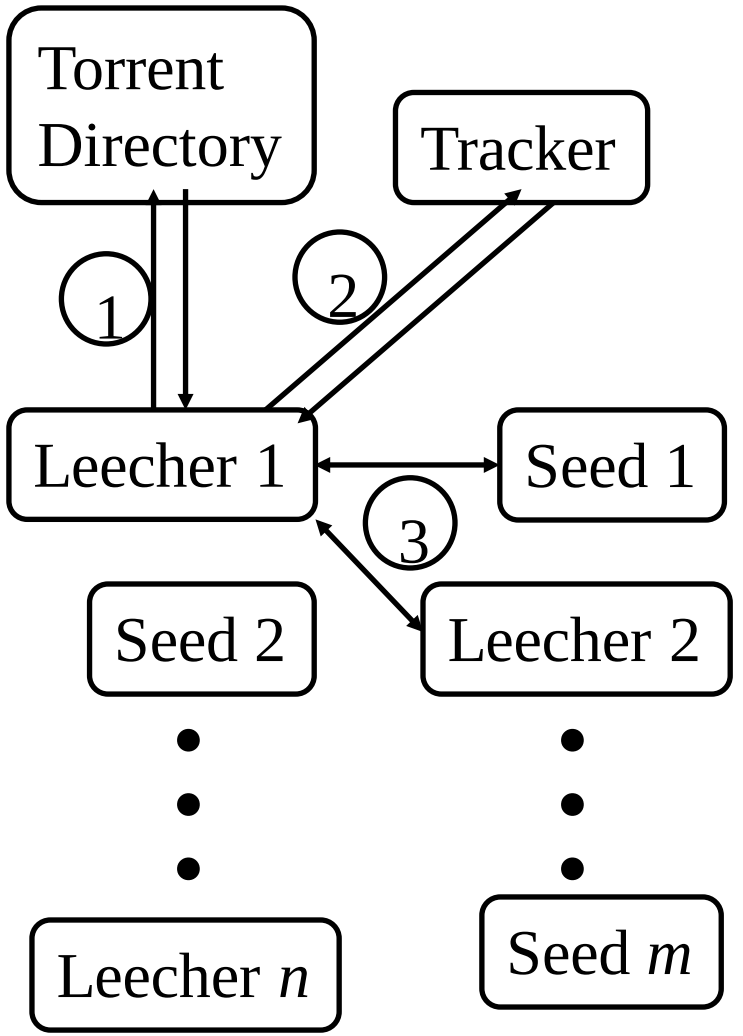
- Do leechers usually stay leechers or is it just a state until they have all the data needed to become a seed?

All leechers should be serving as well even if partial.

©2026 Raj Jain

BitTorrent P2P File Distribution

- **Peers**=nodes participating in a file distribution
- **Torrent**=Set of all peers
- **Torrent File** =a file containing information about the tracker, object ID, and file
- Files are segmented into equal size **chunks** (256KB)
- **Seeds**=Peers that have the complete file
- **Leechers**=Peers that have incomplete file
- **Tracker**=Has list of all peers



Student Questions

- Do most leechers typically complete their downloads and become seeds, or do some never fully obtain the complete file?

Some become seeds. Some exit w or w/o complete file. It is up to the user.

- How do you define complete and incomplete files?

Obvious.

- How do trackers keep the seeds and leechers updated in real time?

Leechers go to trackers first and report periodically.

©2026 Raj Jain

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers. The tracker may provide a random subset (say, 50) of peers.
3. Alice sets up TCP connections with these peers in parallel and gets a map of available chunks

- Requests least available chunks first (**rarest first**)
- Every 10 seconds, Alice calculates the receiving rates
- Sends to (**Unchokes**) the top 4 senders
- Every 30 seconds, Alice sends to one randomly selected peer (**optimistically unchokes**)
  $\Rightarrow$ Helps find high-rate neighbors.
- Ref: www.bittorrent.org http://en.wikipedia.org/wiki/BitTorrent_(protocol)

Student Questions

- Can you explain unchoking/choking?

Chocking = Not sending the packets

- Is Alice sending the same file as what she is receiving?

It may have other files that other users want.

- Can you elaborate on what "Unchokes"

Unchoke=Starts sending.

- Does "top 4 senders" refer to the top 4 senders to that particular peer or the top 4 senders in the whole torrent?

To that peer.

- Is Alice finding high-speed neighbors for herself or the tracker?

For herself.

- Can we go over Slide 2-54 again, specifically regarding choking/unchoking and how this relates to file distribution?

Sure.

- How does the tracker keep track of all the available peers?

Whoever asks for information gets into the database.

©2026 Raj Jain

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers. The tracker may provide a random subset (say, 50) of peers.
3. Alice sets up TCP connections with these peers in parallel and gets a map of available chunks

- Requests least available chunks first (**rarest first**)
- Every 10 seconds, Alice calculates the receiving rates
- Sends to (**Unchokes**) the top 4 senders
- Every 30 seconds, Alice sends to one randomly selected peer (**optimistically unchokes**)
  $\Rightarrow$ Helps find high-rate neighbors.
- Ref: www.bittorrent.org http://en.wikipedia.org/wiki/BitTorrent_(protocol)

Student Questions

- According to the slides, Alice continuously sends data to the top 4 senders. However, in a scenario where we only have 4 peers, and one peer is not sending any data to the other peers. Would it still continuously receive data from Alice? Previously in the hw, we saw that if a peer (Bob) does not upload any data to any other peers, he was still able to get the file due to optimistic unchoking. However, it could take him a long time to receive the file because the process of receiving it was random. However, since now we only have 4 peers in total, would Bob receive data at the same rate as the other peers who are uploading data?

In this case, random selection will result in the 4th user being selected every time, and yes, it will get the file.

- Is this describing a continuous process? Is Alice requesting the file?

Yes

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers. The tracker may provide a random subset (say, 50) of peers.
3. Alice sets up TCP connections with these peers in parallel and gets a map of available chunks

- Requests least available chunks first (**rarest first**)
- Every 10 seconds, Alice calculates the receiving rates
- Sends to (**Unchokes**) the top 4 senders
- Every 30 seconds, Alice sends to one randomly selected peer (**optimistically unchokes**)
  $\Rightarrow$ Helps find high-rate neighbors.
- Ref: www.bittorrent.org http://en.wikipedia.org/wiki/BitTorrent_(protocol)

Student Questions

- I don't understand what Alice is "sending" or to whom. It wouldn't make sense for her to be sending chunks of the same file back to senders.

She may have other files that those senders may want.

- What happens if everyone is missing one of the chunks? Would the chunk be sent from the tracker?

Trackers do not keep any files.

- How does the machine know which package is the rarest?

You ask 50 people about 20 different chunks. The rarest is the smallest number of people having it.

- Why does BitTorrent request the least available chunk first?

Optimal strategy

- Why would the seeders not have the whole file? Why is there an availability problem only for specific chunks?

There may be few seeders, or the seeder may not want to serve you.

©2026 Raj Jain

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers. The tracker may provide a random subset (say, 50) of peers.
3. Alice sets up TCP connections with these peers in parallel and gets a map of available chunks

- Requests least available chunks first (**rarest first**)
- Every 10 seconds, Alice calculates the receiving rates
- Sends to (**Unchokes**) the top 4 senders
- Every 30 seconds, Alice sends to one randomly selected peer (**optimistically unchokes**)
  $\Rightarrow$ Helps find high-rate neighbors.
- Ref: www.bittorrent.org http://en.wikipedia.org/wiki/BitTorrent_(protocol)

Student Questions

- *Is it always the top 4 senders, or is it just 4 in this example?*

Four is in the protocol.

- What if a malicious user chokes everyone to inflate the availability falsely?

If you check everyone, they will all choke you.

- What does it mean by parallel, and parallel in what way?

Send and receive in parallel.

- If optimistically unchokes, why doesn't it send them in increasing order of receiving rate?

It unchokes a node that is not sending currently.

- Every 10s, the user will try a member in the choke group,p but every 30,s will directly promote a member in the choke group to the unchoke group.?

Every 10s, the user selects the top 4 friends. Every 10s, the user selects one unknown as a new friend.

- What is optimistic unchokes?

Hoping it will supply some chunks.

©2026 Raj Jain

BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers. The tracker may provide a random subset (say, 50) of peers.
3. Alice sets up TCP connections with these peers in parallel and gets a map of available chunks

- Requests least available chunks first (**rarest first**)
- Every 10 seconds, Alice calculates the receiving rates
- Sends to (**Unchokes**) the top 4 senders
- Every 30 seconds, Alice sends to one randomly selected peer (**optimistically unchokes**)
  $\Rightarrow$ Helps find high-rate neighbors.
- Ref: www.bittorrent.org http://en.wikipedia.org/wiki/BitTorrent_(protocol)

Student Questions

- Why does Alice need to send anything if she's trying to receive the 'Raj Jain lecture'?

Everyone has to upload.

- Are the top 4 senders the people upload to most files?

Yes. These have sent you the most chunks.

- Does this mean that if we want to receive a file we should upload many to increase discoverability?

Yes, to receive fast.

- Is the process automatic, or could Alice control where to get the files first?

Apps and their settings control them.

- What is meant by "optimistically unchokes?"

Randomly selects, hoping that the host will serve it too.

- Could you please explain unchoking part of the file distribution?

Unchoke = Serve Choke=Do not serve

- Are the top 4 senders defined by those sending Alice the most chunks or the whose with highest sending rates?

Highest sending rates.

©2026 Raj Jain

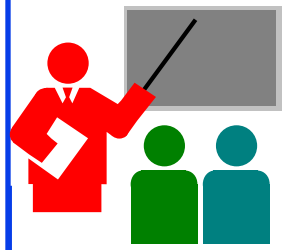
BitTorrent File Distribution (Cont)

1. Alice uses torrent directories (search engines) to find a torrent for "Raj Jain's Lecture"
2. Alice contacts the tracker to get the current list of peers. The tracker may provide a random subset (say, 50) of peers.
3. Alice sets up TCP connections with these peers in parallel and gets a map of available chunks

- Requests least available chunks first (**rarest first**)
- Every 10 seconds, Alice calculates the receiving rates
- Sends to (**Unchokes**) the top 4 senders
- Every 30 seconds, Alice sends to one randomly selected peer (**optimistically unchokes**)
  $\Rightarrow$ Helps find high-rate neighbors.
- Ref: www.bittorrent.org http://en.wikipedia.org/wiki/BitTorrent_(protocol)

Student Questions

©2026 Raj Jain



P2P Applications: Summary

1. P2P applications are more scalable
   ⇒ **More efficient** when the number of peers is large
2. BitTorrent has **peers, trackers, seeds**, and **leechers**
3. BitTorrent unchokes 4 top uploaders and one random node for **load balancing**

Ref: Read Section 2.5 full. Try R21-R23.

Student Questions

- Are there ways to create a P2P network that does not allow communications to be blocked or monitored? Would a form of E2E encryption and a mechanism to reject "bad" peers accomplish this?

Yes, encryption helps. But determining a bad peer is difficult.

©2026 Raj Jain

Homework 2D: P2P

[4 points] P26. Suppose Bob joins a BitTorrent torrent, but he does not want to upload any data to any other peers (so-called free-riding).

A. Bob claims that he can receive a complete copy of the file that is shared by the swarm. Is Bob's claim possible? Why or Why not?

B. Bob further claims that he can further make his "free-riding" more efficient by using a collection of multiple computers (with distinct IP addresses) in the computer lab in his department. How can he do that?

Student Questions

©2026 Raj Jain

Streaming Video

- Video traffic is 80% of consumer traffic
- Video: 25-30 Frames/sec
- Video can be compressed:
  - Spatial: next pixel is similar to this in this frame
  - Temporal: Pixels in the next frame is similar to this
- Variable bit rate (VBR)/Constant bit rate (CBR)
  - Motion Picture Expert Group (MPEG) 1: 1.5 Mbps
  - MPEG2: 3-6 Mbps
  - MPEG4 (.mp4): Less than 1 Mbps



Ref: Cisco Visual Networking Index: Forecast and Methodology, 2014-2019 White Paper, http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html

Student Questions

- What is the difference between VBR and CBR?

The bit rate depends on the changes in the frame and, therefore, on the motion. So the choice is either to keep decreasing the quality (CBR) or to increase the bit rate (VBR) during motion.

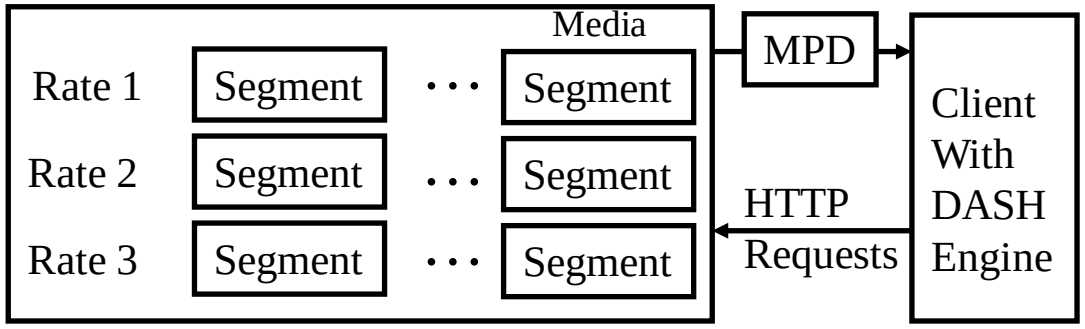
- What's the difference between spatial and temporal compression?

Space = pixels in one frame
Time = successive frames

©2026 Raj Jain

Dynamic Adaptive Streaming over HTTP (DASH)

- DASH provides an efficient method for video streaming
- Standard Web Servers: No changes are required to servers, Content Distribution Networks (CDN), or HTTP protocol.
- Mobile client controls what is downloaded using a "**media presentation description (MPD)**" file defined by DASH.
- MPD contains URLs for segments
- The client measures throughput and requests segments as needed. Allows fast forward, rewind, etc.



Student Questions

- For the DASH protocol, I'm confused about what benefits using multiple URLs brings. Homework 2E implies that every segment of every video quality has a distinct URL, but wouldn't a single URL per video quality be enough?
  Page 149 of the textbook states:
  "With DASH, each video version is stored in the HTTP server, each with a different URL...
  The client then selects one chunk at a time by specifying a URL and a byte range in an HTTP GET

Each video is stored in multiple rates, and each copy is partitioned into multiple segments. Each segment needs a URL.

# of URLs = # of rates × # of segments/rate

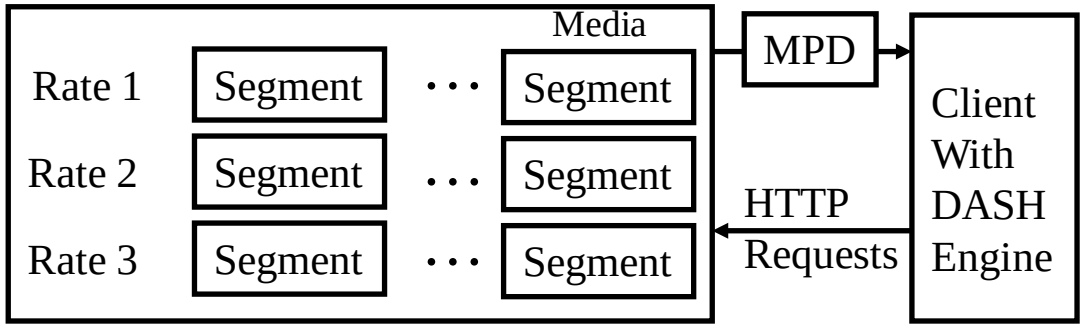
- What is the relationship between MPD and the manifest file?

Same.

©2026 Raj Jain

Dynamic Adaptive Streaming over HTTP (DASH)

- DASH provides an efficient method for video streaming
- Standard Web Servers: No changes are required to servers, Content Distribution Networks (CDN), or HTTP protocol.
- Mobile client controls what is downloaded using a "**media presentation description (MPD)**" file defined by DASH.
- MPD contains URLs for segments
- The client measures throughput and requests segments as needed. Allows fast forward, rewind, etc.



Student Questions

- If it happens that a high-quality video will drop to a low-quality video. Would video drop from 720p to 360p. Are we intentionally dropping some packets of the video at each frame or simply loading frames with lower quality? *Both.*
- Can you go over this slide again, please?

Sure.

- Is this why video streaming rewind functions are usually in preset chunks of time? Ex. hitting rewind goes back 5, 10, 15, etc., seconds at a time? *Yes.*

- If we refresh when streaming a video, does it reload to the beginning of the chunk, or how does the browser know where to continue from?

Location of playing is stored in the app/browser separately than the media.

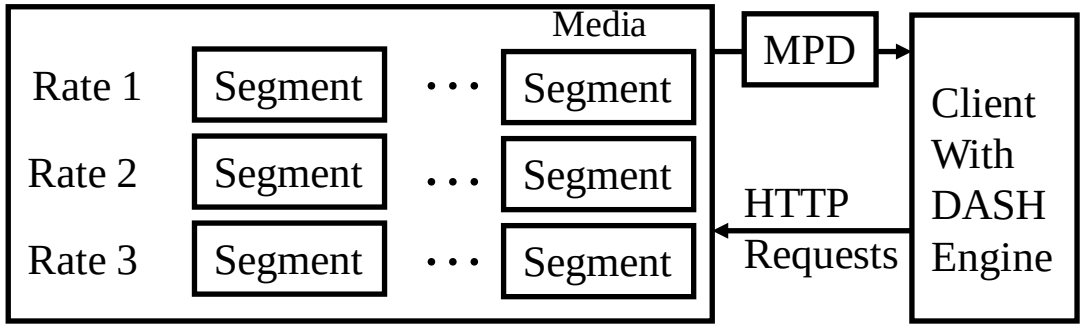
- Why is it necessary for each segment of a video to have a different URL? Can one URL not handle variable rates of multiple video segments?

So you can change the rate in the middle of a movie w/o reloading the entire movie.

©2026 Raj Jain

Dynamic Adaptive Streaming over HTTP (DASH)

- DASH provides an efficient method for video streaming
- Standard Web Servers: No changes are required to servers, Content Distribution Networks (CDN), or HTTP protocol.
- Mobile client controls what is downloaded using a "**media presentation description (MPD)**" file defined by DASH.
- MPD contains URLs for segments
- The client measures throughput and requests segments as needed. Allows fast forward, rewind, etc.

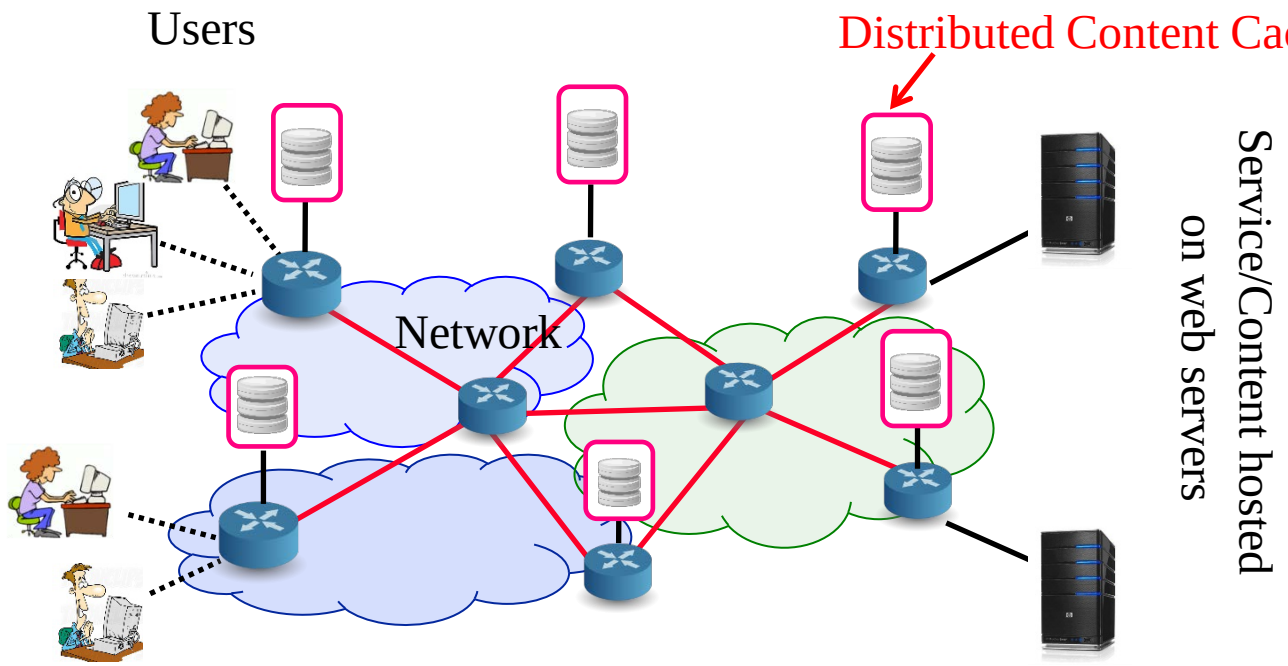


Student Questions

©2026 Raj Jain

Content Distribution Networks (CDN)

- To reduce latency to worldwide users, the data is replicated at many sites
- Users are directed to the nearby site by DNS
- netflix.com -> cdn_stl.com or cdn_sfo.com, …



Student Questions

- What is the difference between CDNs and Load balancers?

CDNs are distributed copies.

- Do CDNs produce similar sites with different names on different servers? *Yes*
- Are CDNs synonymous with network proxy/caches, and if not, what specifically makes them different?

Client organizations set proxies. CDNs are done by serving organizations.

- Can you elaborate on how DNS redirects users to nearby regions? What will happen if I use public DNS such as 8.8.8.8?

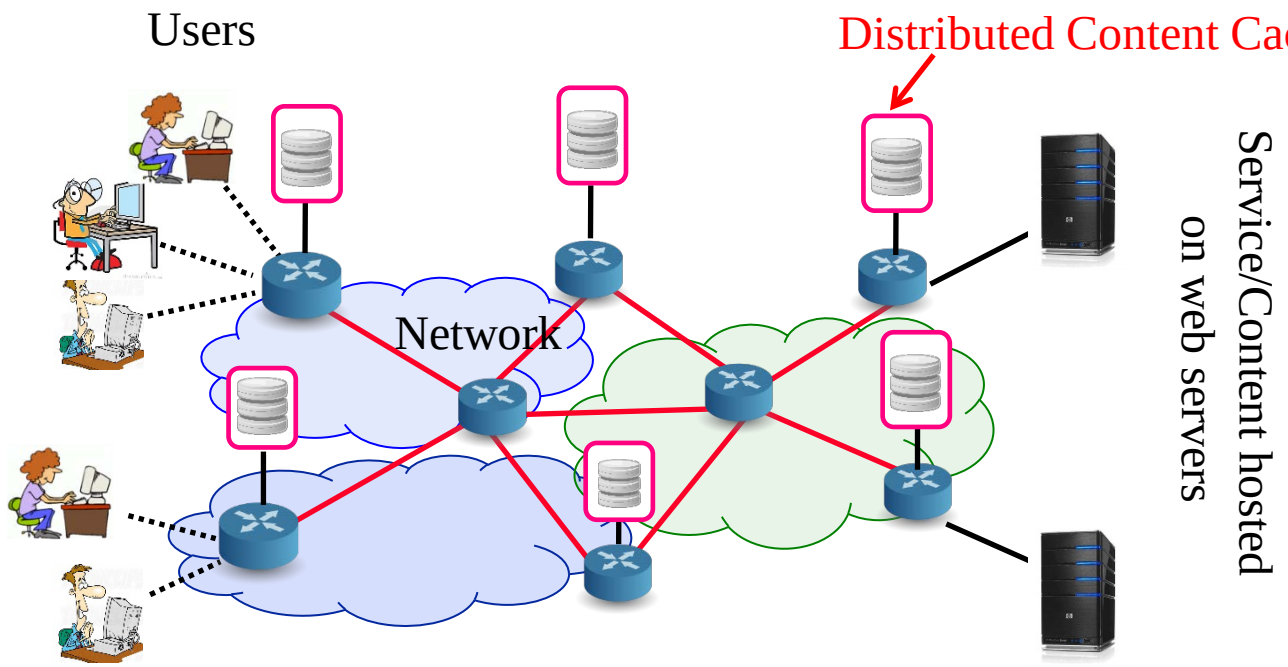
The local DNS server has local name-to-IP translations. Public DNS is also distributed like Google.

- Are there any limitations on the content a CDN can deliver? *No.*
- I've noticed that the WUSTL sites (WebStac, course listings, etc.) are extremely slow to load from East Asia. Could this be because the data is only at the WashU server and not replicated through CDNs? *Yes.*

©2026 Raj Jain

Content Distribution Networks (CDN)

- To reduce latency to worldwide users, the data is replicated at many sites
- Users are directed to the nearby site by DNS
- netflix.com -> cdn_stl.com or cdn_sfo.com, …



Student Questions

- Can you go through more examples of CDN networks besides video streaming?

Most websites use CDNs now.

- How do CDNs decide what data to host?

Companies pay them to host.

©2026 Raj Jain

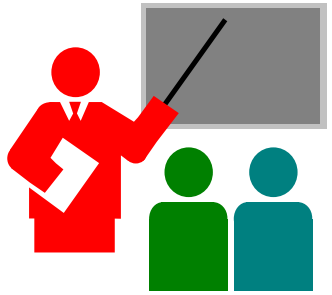
Homework 2E: DASH

- [2 points] A DASH system stores video at 5 different qualities (rates) and 10-minute segments. How many URLs will be required for a 3-hour movie?

Student Questions

©2026 Raj Jain

Application Layer: Summary



1. Applications use TCP/UDP **ports** for communication.
2. HTTP/FTP/SMTP are **client-server** protocols and use TCP connections
3. HTTP is **stateless,** but cookies allow servers to maintain states
4. **Proxy** servers improve performance by caching
5. BitTorrent is a **P2P** file distribution protocol and uses trackers to keep a list of peers.
6. **DASH** allows clients to request different video segments as needed
7. **CDN** directs users to nearby copies via DNS

Ref: In addition to previous readings, read Sections 2.6.1-2.6.3. Try R24-R25.

Student Questions

- Why email and file transfers are using different protocols? Why not treat email also as a file?

Files are usually "pulled," while emails are usually "pushed."

- Since proxy servers cache from the original server, does this mean it's a secure way to hide actions on the Internet?

Nothing is hidden. The fact that the proxy server took it from the original server is visible to the world.

- Can you go over what cache is?

Computer Science 101.

- Are all five homework questions and two labs due on Monday?

Yes.

©2026 Raj Jain

Lab 2B: UDP Pinger

- [50 points] In this lab, you will learn the basics of socket programming for UDP in Python. You will learn how to send and receive datagram packets using UDP sockets and also how to set a proper socket timeout. Throughout the lab, you will gain familiarity with a Ping application and its usefulness in computing statistics such as packet loss rate.
- You will first study a simple Internet ping server written in Python and implement a corresponding client. The functionality provided by these programs is similar to the functionality provided by standard ping programs available in modern operating systems. However, these programs use a simpler protocol, UDP, rather than the standard Internet Control Message Protocol (ICMP) to communicate with each other. The ping protocol allows a client machine to send a packet of data to a remote machine and have the remote machine return the data back to the client unchanged (an action referred to as echoing). Among other uses, the ping protocol allows hosts to determine round-trip times to other machines.
- You are given the complete code for the Ping server below. Your task is to write the Ping client.

Student Questions

- Will the exam involve the use of the Socket APIs?

The exam does not require any coding. So no APIs. But it may ask about standard port numbers.

- *Is the ping client also required to support command line arguments to specify the IP address and port number?*

No, sending packets to 127.0.0.1:12000 is sufficient.

- How precise should the time be?

The most accurate clock time available on the computer. For example, in Windows, Time command gives me 9:04:49.91, which is accurate to 1/100th of a second.

- When are Lab 1 and Lab 2 due? Where do we submit them?

Monday, February 7th. Submit on Canvas.

©2026 Raj Jain

Lab 2B (Cont)

Server Code

The following code fully implements a ping server. You need to compile and run this code before running your client program. *You do not need to modify this code.*

In this server code, 36% of the client's packets are simulated to be lost. You should study this code carefully, as it will help you write your ping client.

```
# UDPPingerServer.py
# We will need the following module to generate randomized lost packets
import random
from socket import *
# Create a UDP socket
# Notice the use of SOCK_DGRAM for UDP packets
serverSocket = socket(AF_INET, SOCK_DGRAM)
# Assign IP address and port number to socket
serverSocket.bind(('', 12000))
```

Student Questions

- Is there a place to download the server code as a file or is it just copy paste from the slides?

There is link to a file on the website. Please right-click and "save as" to save the file.

©2026 Raj Jain

Lab 2B (Cont)

```
while True:
# Generate random number in the range of 0 to 10
rand = random.randint(0, 10)
# Receive the client packet along with the address it is coming from
message, address = serverSocket.recvfrom(1024)
# Capitalize the message from the client
message = message.upper()
# If rand is less is than 4, we consider the packet lost and do not respond
if rand < 4:
continue
# Otherwise, the server responds
serverSocket.sendto(message, address)
```

The server sits in an infinite loop listening for incoming UDP packets. When a packet comes in and if a randomized integer is greater than or equal to 4, the server simply capitalizes the encapsulated data and sends it back to the client.

Student Questions

©2026 Raj Jain

Lab 2B (Cont)

Packet Loss

UDP provides applications with an unreliable transport service. Messages may get lost in the network due to router queue overflows, faulty hardware or some other reasons. Because packet loss is rare or even non-existent in typical campus networks, the server in this lab injects artificial loss to simulate the effects of network packet loss. The server creates a variable randomized integer that determines whether a particular incoming packet is lost or not.

Client Code

You need to implement the following client program.

The client should send 10 pings to the server. Because UDP is an unreliable protocol, a packet sent from the client to the server may be lost in the network or vice versa. For this reason, the client cannot wait indefinitely for a reply to a ping message. You should get the client to wait up to one second for a reply; if no reply is received within one second, your client program should assume that the packet was lost during transmission across the network. You will need to look up the Python documentation to find out how to set the timeout value on a datagram socket.

Student Questions

©2026 Raj Jain

Lab 2B (Cont)

Specifically, your client program should

(1) send the ping message using UDP (Note: Unlike TCP, you do not need to establish a connection first since UDP is a connectionless protocol.)

(2) print the response message from the server, if any

(3) calculate and print the round trip time (RTT), in seconds, of each packet, if server responses

(4) otherwise, print "Request timed out"

During development, you should run the UDPPingerServer.py on your machine and test your client by sending packets to *localhost* (or, **127.0.0.1**). After you have fully debugged your code, you should see how your application communicates across the network with the ping server and ping client running on different machines.

Message Format

The ping messages in this lab are formatted in a simple way. The client message is one line consisting of ASCII characters in the following format:

Ping *sequence_number time*

where *sequence_number* starts at 1 and progresses to 10 for each successive ping message sent by the client, and *time* is the time when the client sends the message.

Student Questions

- Should we do time since the program started or the standard time (e.g., 11:38 PM)? Also, how precise should it be?

The clock time up to the precision given by your operating system. For example, Time command on Windows gives up to 100th of a second, e.g., 9:04:49.91

©2026 Raj Jain

Lab 2B (Cont)

What to Hand in

You will hand in the complete client code and screenshots at the client, verifying that your ping program works as required.

Student Questions

- When is the lab due?

Next Monday.

MIME Encoding

```
From: "Saved by Internet Explorer 11"
Subject:
Date: Wed, 3 Feb 2021 12:08:58 -0600
MIME-Version: 1.0
Content-Type: multipart/related;
                type="text/html";
                boundary="----=_NextPart_000_0000_01D6FA25.59EF7E90"
X-MimeOLE: Produced By Microsoft MimeOLE

This is a multi-part message in MIME format.

------=_NextPart_000_0000_01D6FA25.59EF7E90
Content-Type: text/html;
                charset="Windows-1252"
Content-Transfer-Encoding: quoted-printable
Content-Location: file://D:\u\jain.htm

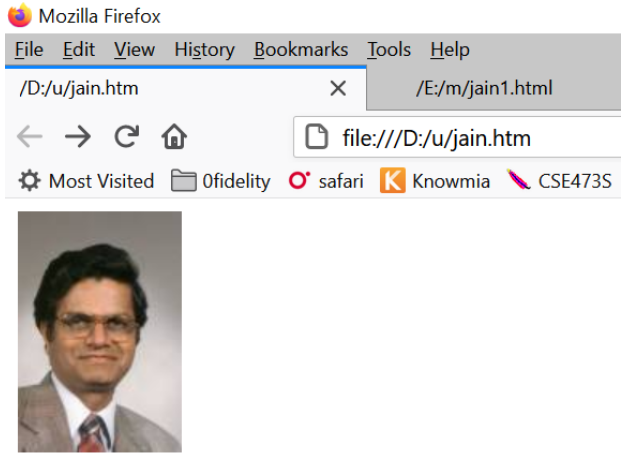
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.0 Transitional//EN">
<HTML><HEAD><META content=3D"IE=3D5.0000" =
http-equiv=3D"X-UA-Compatible">

<META http-equiv=3D"Content-Type" content=3D"text/html; =
charset=3Dwindows-1252">
<META name=3D"GENERATOR" content=3D"MSHTML 11.00.10570.1001"></HEAD>=20
<BODY><IMG src=3D"file:///D:/u/jain.jpg"> <BR>Raj Jain </BODY></HTML>
```

Student Questions

MIME = Multipurpose Internet Mail Extensions

MHTML = MIME Encoded HTML



Raj Jain

©2026 Raj Jain

MIME Encoding (Cont)

------=_NextPart_000_0000_01D6FA25.59EF7E90
Content-Type: image/jpeg
Content-Transfer-Encoding: **base64**
Content-Location: file:///D:/u/**jain.jpg**

/9j/4AAQSkZJRgABAQEAYABgAAD/2wBDAAoHBwgHBgoICAgLCgoLDhgQDg0NDh0VFhEYIx8lJCIf
IiEmKzcvJik0KSEiMEExNDk7Pj4+JS5ESUM8SDc9Pjv/2wBDAQoLCw4NDhwQEBw7KCIoOzs7Ozs7
Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozs7Ozv/wAARCACWAGYDASIA
AhEBAxEB/8QAGwAAAQUBAQAAAAAAAAAAAAAAAABgACAwQFAQf/xAA3EAABAwIFAQYCCQQDAAAAAAAB
AAIDBBEFBhIhMUETIlFhcYGRsRQjMjNCUqHB0RVi4fAHNFP/xAAZAQADAQEBAAAAAAAAAAAAAAAA
AQMEAgX/xAAiEQACAgICAgIDAAAAAAAAAAAAAQIRAyEEMRJBEyIUMmH/2gAMA5QAAhEDEQA/AK1l
…….
0zhMzuPdcSDp6rQgZdhdYEHghOMYvpIuPMJUkbaeU6S7S7hl9h7KRQY8Fsl9PeHCuUcvZVUM4P3b
D3Wk/uLg+6Kc5SXqaMX5Dv2QkKB+IYfiFG5319O/6TE//wBBax+Q91r1lccToMHncbyPhJd6iwJ+
IKb/AF2cV9qRXgLo3gg7cLQhDnEaXb+XRUnAM72nU09FagmDC022U0M26bFZ6ezJ7ys8fxD+VrMn
jnj7SJ4cCh0EEBzRcLjJ5KWXXESL8joV0AQE7pKrT1bapmoCzhy3wSTEBjagmudDb7LNTT5qYu2A
aLW3v6pJKZQjcdZueCo2m5HTdJJIaJK6ok/psVGx7mtlmOq3UBpNvksmW0H1cY07bkJLjJ2jdxk
vBv+mtgwL5mNebmQGMn+1wt/HwUD6V+GYjFhxk1tpqcAEeJJcfmkkqQf0ISWTRba8OdpI2sugN
hfYgljr2APCSSEQZfw6QvYn4/Qvk0AtC0MvISWTRba8OdpI2sugN

------=_NextPart_000_0000_01D6FA25.59EF7E90--

Student Questions

- What does the first line of the MIME message signify? (containing NextPart + some numbers)

Not really sure. The numerical part could be a unique serial number in hex. Such numbers are generally obtained by hashing.

Acronyms

- ASCII — American Standard Code for Information Interchange
- CBR — Constant bit rate
- CDN — Content Distribution Network
- DASH — Dynamic Adaptive Streaming
- DNS — Domain Name System
- FTP — File Transfer Protocol
- GMT — Greenwich Mean Time
- HTML — Hyper-Text Markup Language
- HTTP — Hyper-Text Transfer Protocol
- ICANN — International Corporation for Assigned Names and Numbers
- ID — Identifier
- IMAP — Internet Message Access Protocol
- IP — Internet Protocol
- ISO — International Standards Organization
- ISP — Internet Service Provider
- kB — Kilo Byte

Student Questions

©2026 Raj Jain

Acronyms (Cont)

- MPD Media Presentation Description
- MPEG Moving Picture Expert Group
- NAT Network Address Translator
- NS Name Service
- PC Personal Computer
- POP Point of Presence
- RR Resource Record
- SMTP Simple Mail Transfer Protocol
- TCP Transmission Control Protocol
- TLD Top Level Domain
- TTL Time to Live
- UDP Universal Data Protocol
- URL Uniform Resource Locator
- VBR Variable bit rate

Student Questions

©2026 Raj Jain

Scan This to Download These Slides





Raj Jain
http://rajjain.com

http://www.cse.wustl.edu/~jain/cse4703-26/i_2app.htm

Student Questions

- Can you post the slides with the fixes from last year? You say in the video that slide 51 has incorrect point arrows.

All slides are the latest, with all the known errors correct. The slide had an error before the video was recorded. All errors are fixed on the same day.

- I'm still confused about one-to-one and many-to-many questions. For many-to-many, can one question have multiple answers and vice versa?

There are no "many-to-many" questions. Only single-choice or multiple-choice questions. We do not use multiple-choice questions in the exams.

- When are Lab 1 and Lab 2 due? Where do we submit them?

Monday, February 7th. Submit on Canvas.

- Can you post the slides with the fixes from last year? You say in the video that slide 51 has incorrect point arrows.

All slides are the latest, with all the known errors correct. The slide had an error before the video was recorded. All errors are fixed on the same day.

©2026 Raj Jain

Scan This to Download These Slides





Raj Jain

http://rajjain.com

http://www.cse.wustl.edu/~jain/cse4703-26/i_2app.htm

Student Questions

- How does the TOR Network protect/hide users IP addresses and deal with cookies?

TOR is related to routing. Too advanced to be covered in this course.

- Is there a spot to check our answers to the review questions that were called out in the lecture?

Video questions are related to statements in the video and just binary questions. Feedback is immediate.

- What are the limitations of DNSSEC (DNS Security Extensions) in preventing DNS-based attacks?

DNSSEC is not a part of this course.

- How to prevent DNS hijacking attacks?

Not discussed in this course.

©2026 Raj Jain

Related Modules



CSE 567: The Art of Computer Systems Performance Analysis
https://www.youtube.com/playlist?list=PLjGG94etKypJEKjNAa1n_1X0bWWNyZcof

CSE473S: Introduction to Computer Networks (Fall 2011),
https://www.youtube.com/playlist?list=PLjGG94etKypJWOSPMh8Azcgy5e_10TiDw





CSE 570: Recent Advances in Networking (Spring 2013)
https://www.youtube.com/playlist?list=PLjGG94etKypLHyBN8mOgwJLHD2FFIMGq5

CSE571S: Network Security (Spring 2011),
https://www.youtube.com/playlist?list=PLjGG94etKypKvzfVtutHcPFJXumyyg93u





Video Podcasts of Prof. Raj Jain's Lectures,
https://www.youtube.com/channel/UCN4-5wzNP9-ruOzQMs-8NUw

Student Questions

- Should I be reading the chapters before or after the videos and class or does the order not matter as long as I follow with the book?

You must read the book before doing the homework.

©2026 Raj Jain