

Course on Computer Communication and Networks

Lecture 3 Chapter 2: Application-layer

EDA344/DIT 420, CTH/GU

Based on the book Computer Networking: A Top Down Approach, Jim Kurose, Keith Ross, Addison-Wesley.

Chapter 2: Application Layer

Chapter goals:

- conceptual + implementation aspects of network application protocols
 - client server paradigm
- specific protocols:
 - http, smtp, pop, dns,
 - *p2p & streaming, CDN: later in the course)*

Applications and application-layer protocols

Application: communicating, distributed processes

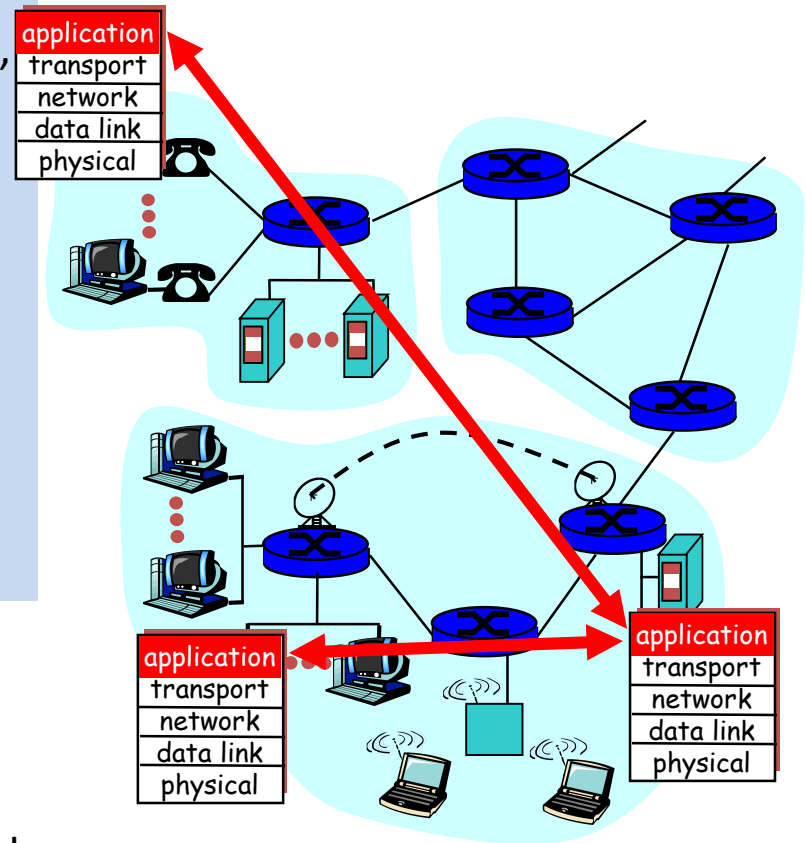
- running in network hosts in “user space”
- e.g., email, file transfer, the Web

Application-layer protocols


- Are only one “piece” of an application - others are e.g. **user agents**.
 - Web: browser
 - E-mail: mail reader
 - streaming audio/video: media player

App-layer protocols:

- **define** messages exchanged and actions taken
- **use services** provided by lower layer protocols



Client-server architecture



server:


- always-on
- permanent host address
- clusters of servers for scaling

clients:


- communicate with server
- may be intermittently connected
- may have dynamic host addresses
- do not communicate directly with each other

Peer2Peer architecture

- *no* always-on server
- peers request service from other peers, provide service in return
- peers are intermittently connected and may change addresses
 - complex management



Roadmap



- Addressing and Applications needs from transport layer
- Http
 - General description and functionality
 - Authentication, cookies and related aspects
 - Caching and proxies
- SMTP (POP, IMAP)
- DNS

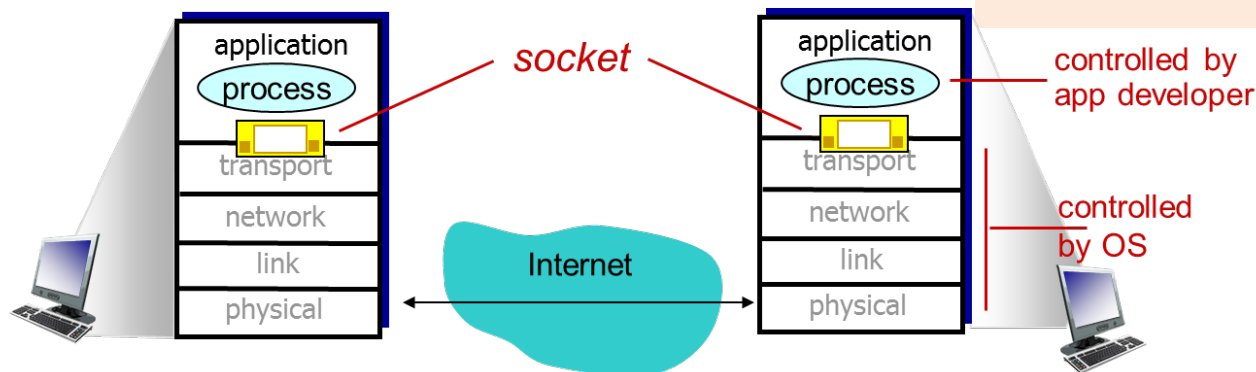
Addressing, sockets

socket: Internet application programming interface


- 2 processes communicate by sending data into socket, reading data out of socket (like sending out, receiving in via doors)

Q: how does a process “identify” the other process with which it wants to communicate?

- **IP address** (unique) of host running other process
- “**port number**” - allows receiving host to determine to which local process the message should be delivered



Roadmap



- Addressing and Applications needs from transport layer
- Http
 - General description and functionality
 - Authentication, cookies and related aspects
 - Caching and proxies
- SMTP (POP, IMAP)
- DNS

Transport service requirements: common apps

| application | data loss | throughput | time sensitive |
|-----------------------|---------------|---|----------------|
| file transfer | no loss | elastic | no |
| e-mail | no loss | elastic | no |
| Web documents | no loss | elastic | no |
| real-time audio/video | loss-tolerant | audio: 5kbps-1Mbps video: 10kbps-5Mbps | y, 100' s msec |
| interactive games | loss-tolerant | few kbps up | y, 100' s msec |
| text messaging | no loss | elastic | yes and no |

Services to upper layer by Internet transport protocols

TCP service:



- *connection-oriented reliable transport* between sending and receiving process
 - correct, in-order delivery of data
 - setup required between client, server
- *does not provide*: timing, bandwidth guarantees

UDP service:




- *connectionless*
- *Unreliable, “best-effort” transport* between sending and receiving process
- *does not provide*: timing, or bandwidth guarantee

Internet apps: application, transport protocols

| application | application layer protocol | underlying transport protocol |
|------------------------|---|--------------------------------------|
| e-mail | SMTP [RFC 2821] | TCP |
| remote terminal access | Telnet [RFC 854] | TCP |
| Web | HTTP [RFC 2616] | TCP |
| file transfer | FTP [RFC 959] | TCP |
| streaming multimedia | HTTP (e.g., YouTube), RTP [RFC 1889] | TCP or UDP |
| Internet telephony | SIP, RTP, proprietary (e.g., Skype) | TCP or UDP |

Roadmap



- Addressing and Applications needs from transport layer
- Http
 - General description and functionality
 - Authentication, cookies and related aspects
 - Caching and proxies
- SMTP (POP, IMAP)
- DNS

Web and HTTP

First, some jargon...

- *web page* consists of *objects*
- object can be HTML file, JPEG image, Java applet, audio file,...
- web page consists of *base HTML-file* which includes *several referenced objects*
- each object is addressable by a *URL*, e.g.,

`www.someschool.edu/someDept/pic.gif`

host name

path name

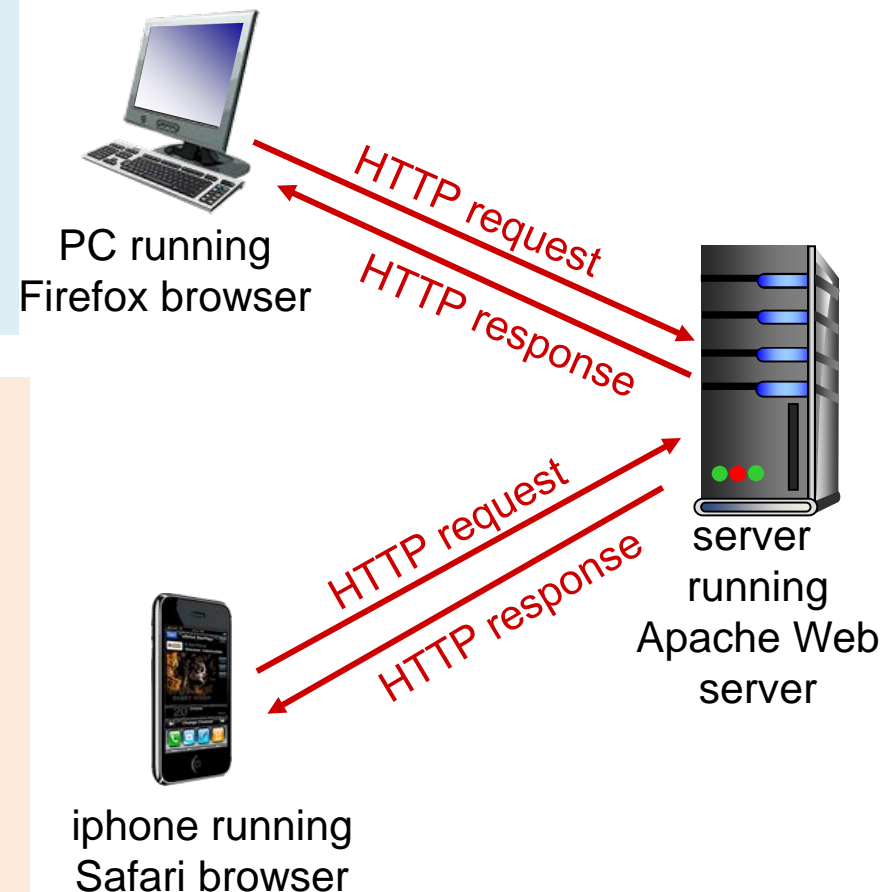
HTTP: **hypertext transfer protocol** overview

Web' s application layer protocol

- *http client*: web browser; requests, receives, displays Web objects
- *http server*: Web server sends objects

uses TCP:

- client initiates TCP connection to server, port 80
- server accepts TCP connection
- HTTP messages (application-layer protocol messages) exchanged
- TCP connection closed



http example

user enters URL

eg `www.someSchool.edu/someDepartment/home.index`

(contains text,
references to 10
jpeg images)

1a. http client initiates TCP connection to http server (process) at `www.someSchool.edu`. Port 80 is default for http server.

1b. http server at host `www.someSchool.edu` waiting for TCP connection at port 80. “accepts” connection, notifying client

2. client sends http *request message* (containing URL) into TCP connection socket

3. server receives request, forms *response message* with requested object (`someDepartment/home.index`), sends message into socket

4. client receives response msg with file, displays html. Parsing html file, finds 10 referenced jpeg objects

4a. server closes TCP connection.

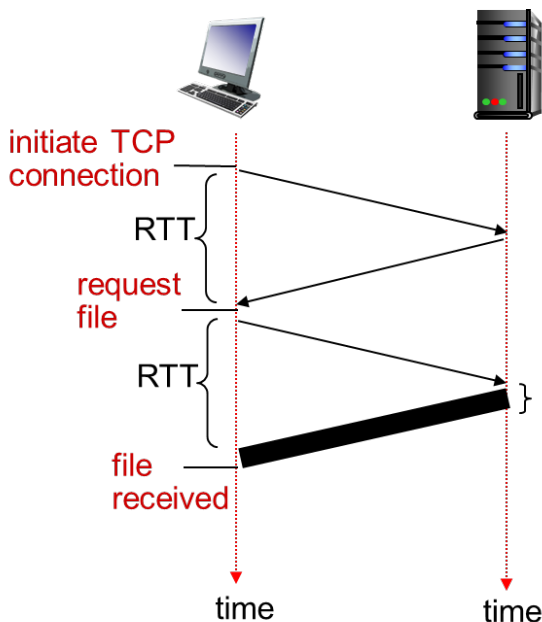
6. Steps 1-5 repeated for each of 10 jpeg objects

time
↓

Non-persistent and persistent http

Non-persistent (http1.0)

- server parses request, responds, closes TCP connection
- *non-persistent HTTP response time = $2RTT + \text{file transmission time}$*
- **new TCP connection for each object**
=> extra overhead per object



Persistent

- **on same TCP connection:** server parses request, responds, parses new request, ..
- Client sends requests for all referenced objects as soon as it receives base HTML;
- Less overhead per object
- **Objects are fetched sequentially (http 1.1)**
 - update http/2: fetches in priority ordering

With both, browsers can open parallel sessions

HTTP request message

- two types of HTTP messages: *request, response*
- **HTTP request message:**
 - ASCII (human-readable format)

The diagram illustrates the structure of an HTTP request message. It shows a sequence of lines: a request line, followed by header lines, and a final line indicating the end of the header lines. Annotations with arrows point to specific parts of the message:

- request line (GET, POST, HEAD commands)**: Points to the first line of the message: `GET /index.html HTTP/1.1\r\n`.
- header lines**: A bracket on the left side groups the lines from `Host:` to `Connection:`.
- carriage return, line feed at start of line indicates end of header lines**: Points to the `\r\n` sequence at the end of the `Connection:` line.
- carriage return character**: Points to the `\r` character in the `\r\n` sequence at the end of the request line.
- line-feed character**: Points to the `\n` character in the `\r\n` sequence at the end of the request line.

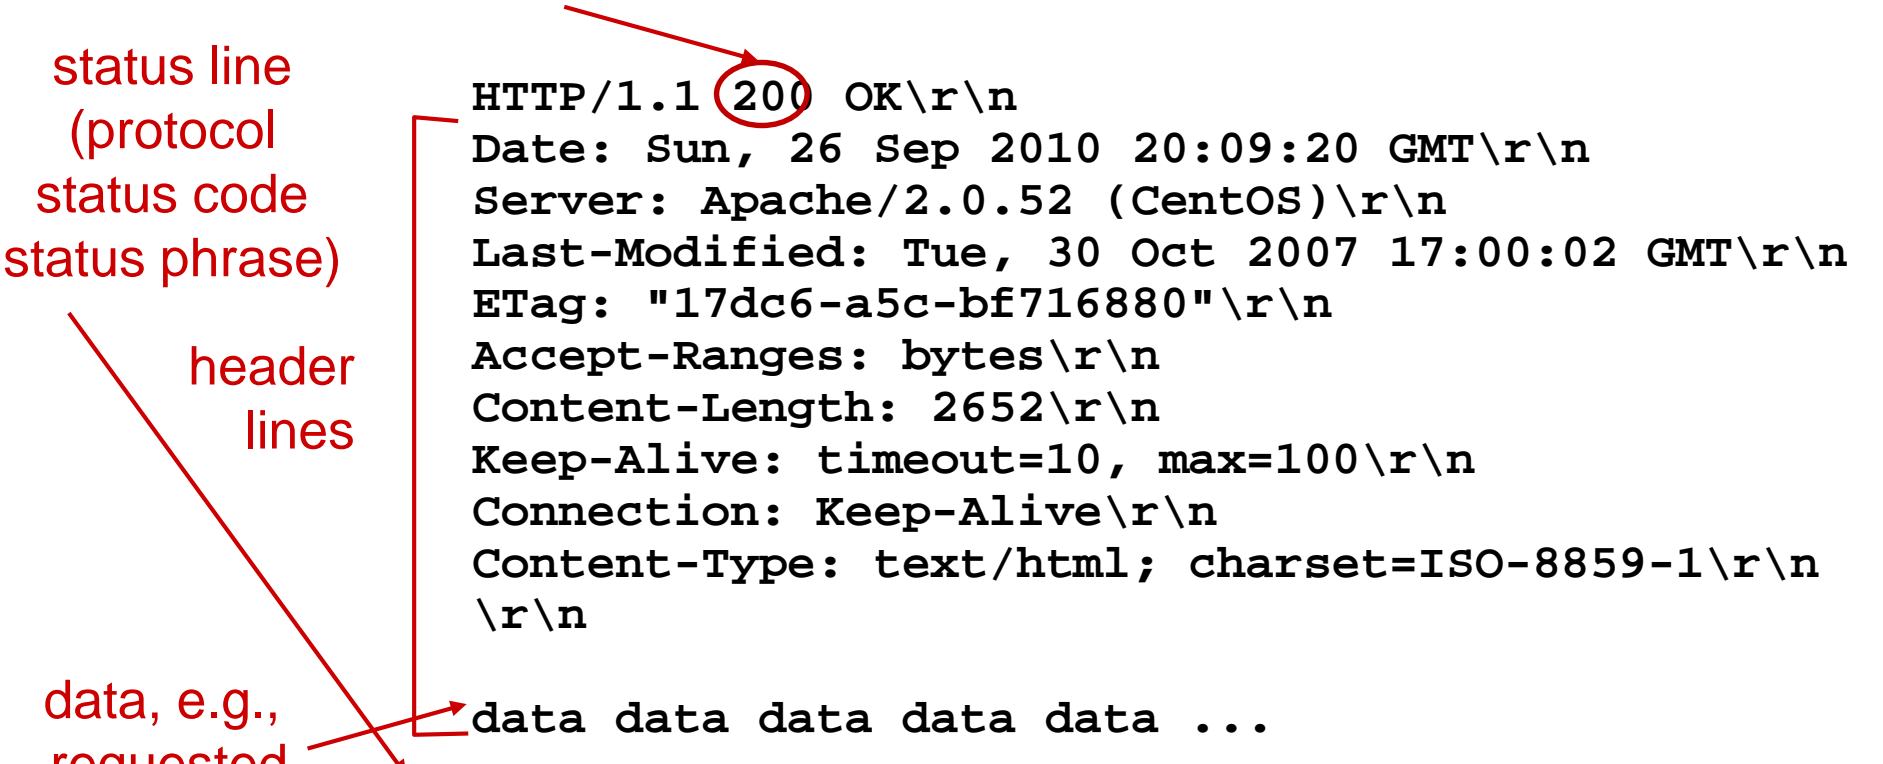
```
GET /index.html HTTP/1.1\r\nHost: www-net.cs.umass.edu\r\nUser-Agent: Firefox/3.6.10\r\nAccept: text/html,application/xhtml+xml\r\nAccept-Language: en-us,en;q=0.5\r\nAccept-Encoding: gzip,deflate\r\nAccept-Charset: ISO-8859-1,utf-8;q=0.7\r\nKeep-Alive: 115\r\nConnection: keep-alive\r\n\r\n
```

HTTP response message

status line
(protocol
status code
status phrase)

header
lines

data, e.g.,
requested
HTML file



```
HTTP/1.1 200 OK\r\n
Date: Sun, 26 Sep 2010 20:09:20 GMT\r\n
Server: Apache/2.0.52 (CentOS)\r\n
Last-Modified: Tue, 30 Oct 2007 17:00:02 GMT\r\n
ETag: "17dc6-a5c-bf716880"\r\n
Accept-Ranges: bytes\r\n
Content-Length: 2652\r\n
Keep-Alive: timeout=10, max=100\r\n
Connection: Keep-Alive\r\n
Content-Type: text/html; charset=ISO-8859-1\r\n
\r\n
```

data data data data data ...

200 OK: request succeeded, requested object in this msg

301 Moved Permanently: requested object moved, new location specified later in this message (Location:)

400 Bad Request: request message not understood

404 Not Found: requested document not found on this server

505 HTTP Version Not Supported

Trying out HTTP (client side) for yourself

1. Telnet to a Web server:

```
telnet cis.poly.edu 80
```

opens TCP connection to port 80
(default HTTP server port) at cis.poly.edu.
anything typed in sent
to port 80 at cis.poly.edu

2. type in a GET HTTP request:

```
GET /~ross/ HTTP/1.1  
Host: cis.poly.edu
```


by typing this in (hit carriage
return twice), you send
this minimal (but complete)
GET request to HTTP server

3. look at response message sent by HTTP server!

Topic of the programming assignment

tation
nended! You

Roadmap



- Addressing and Applications needs from transport layer
- Http
 - General description and functionality
 - Authentication, cookies and related aspects
 - Caching and proxies
- SMTP (POP, IMAP)
- DNS

HTTP is “stateless”

HTTP server maintains no information about past client requests

aside
protocols that maintain
“state” are complex!

- ❖ past history (state) must be maintained
- ❖ if server/client crashes, their views of “state” may be inconsistent, must be reconciled

Q: how do web applications keep state though?

Cookies: keeping “state”

client



server



cookie file



ebay 8734
amazon 1678


usual http request msg

Amazon server
creates ID
1678 for user

usual http response
set-cookie: 1678

create
entry

backend
database



usual http request msg
cookie: 1678

cookie-
specific
action

access

usual http response msg

one week later:



ebay 8734
amazon 1678

usual http request msg
cookie: 1678

cookie-
specific
action

access

usual http response msg


Cookies (continued)

cookies can bring:

- authorization
- shopping carts
- recommendations
- user session state

- aside —
- ## Cookies and privacy:
- ☐ cookies permit sites to learn a lot about you
 - ☐ you may supply name and e-mail to sites
 - ☐ search engines use cookies to learn yet more
 - ☐ advertising companies obtain info across sites

Roadmap




- Addressing and Applications needs from transport layer
- Http
 - General description and functionality
 - Authentication, cookies and related aspects
 - Caching and proxies
- SMTP (POP, IMAP)
- DNS

Web Caches (proxy server)

Goal: satisfy client request without involving origin server

- user configures browser: Web accesses via web cache
- client sends all http requests to web cache; the cache(proxy) server acts as a usually caches do
- Hierarchical, cooperative caching, ICP: Internet Caching Protocol (RFC2187)




Why Web Caching?

Assume: cache is close to client
(e.g., in same network)


- smaller response time
- decrease traffic to distant servers
 - link out of institutional/local ISP network can be bottleneck
- Important for large data applications (e.g. video,...)

Performance effect:



$$E(\text{delay}) = \text{hitRatio} * \text{LocalAccDelay} + (1 - \text{hitRatio}) * \text{RemoteAccDelay}$$

Roadmap



- Addressing and Applications needs from transport layer
- Http
 - General description and functionality
 - Authentication, cookies and related aspects
 - Caching and proxies
- SMTP (POP, IMAP)
- DNS

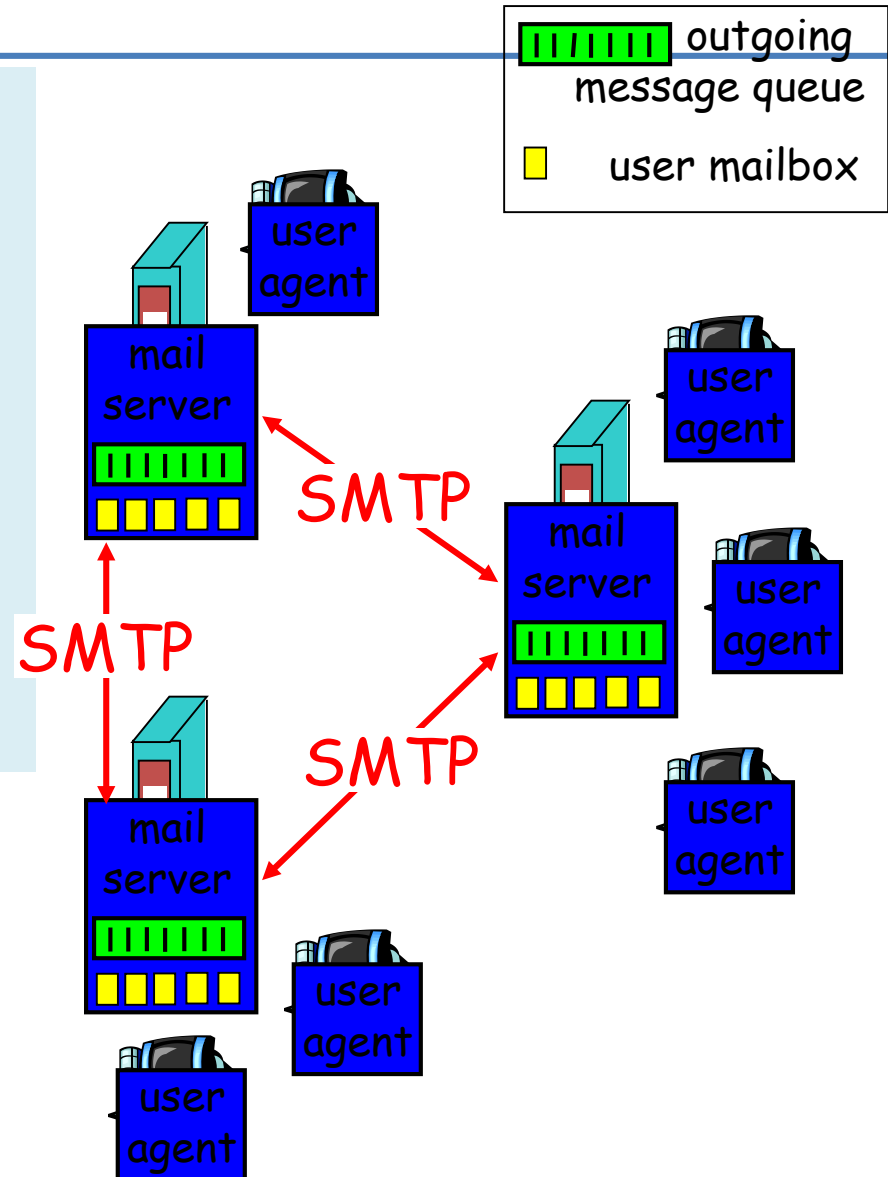
Electronic Mail

User Agent

- a.k.a. “mail reader: composing, editing, reading mail messages -e.g., Outlook,

Mail Servers

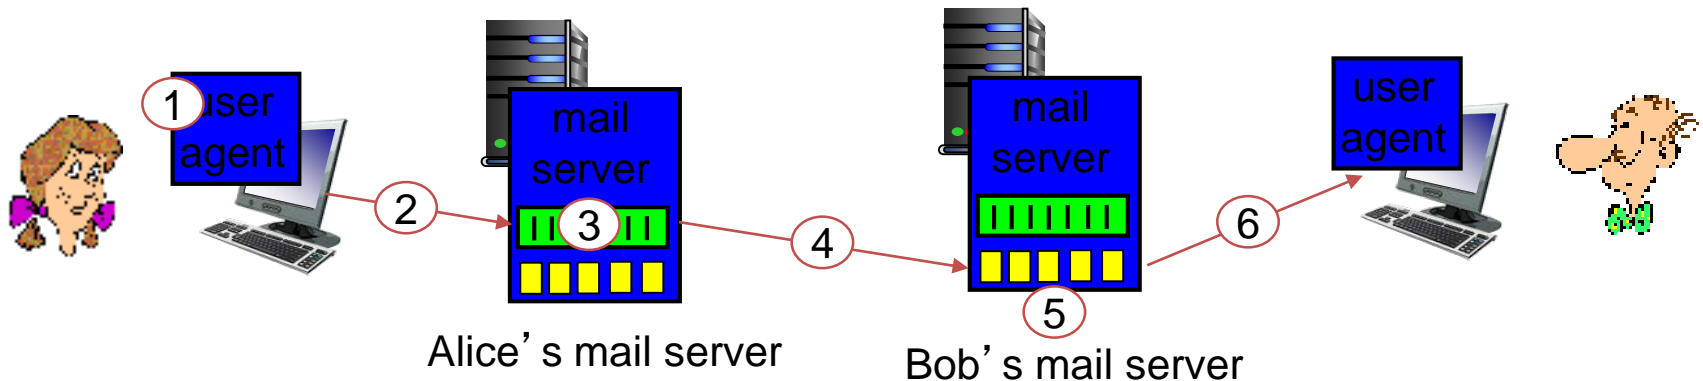
- **SMTP protocol** between mail servers to send email messages
 - client: sending mail server
 - “server”: receiving mail server



Scenario: Alice sends message to Bob

- 1) Alice, UA: message “to” bob@someschool.edu
- 2) Alice, UA: sends message to her mail server’s queue
- 3) Alice, mail server: TCP connection with Bob’s mail server (acting as a client of SMTP)

- 4) Alice’s mail server sends Alice’s message over the TCP connection
- 5) Bob’s mail server places the message in Bob’s mailbox
- 6) Bob invokes his UA to read message



Sample SMTP interaction

```
S: 220 hamburger.edu
C: HELO crepes.fr
S: 250 Hello crepes.fr, pleased to meet you
C: MAIL FROM: <alice@crepes.fr>
S: 250 alice@crepes.fr... Sender ok
C: RCPT TO: <bob@hamburger.edu>
S: 250 bob@hamburger.edu ... Recipient ok
C: DATA
S: 354 Enter mail, end with "." on a line by itself
C: Do you like ketchup?
C: How about pickles?
C: .
S: 250 Message accepted for delivery
C: QUIT
S: 221 hamburger.edu closing connection
```

- SMTP (RFC 2821) uses TCP, port 25
- three phases
 - handshaking (greeting)
 - transfer of messages
 - closure


You can try it out through
telnet servername 25

- see 220 reply from server
- enter HELO, MAIL FROM, RCPT TO, DATA, QUIT commands

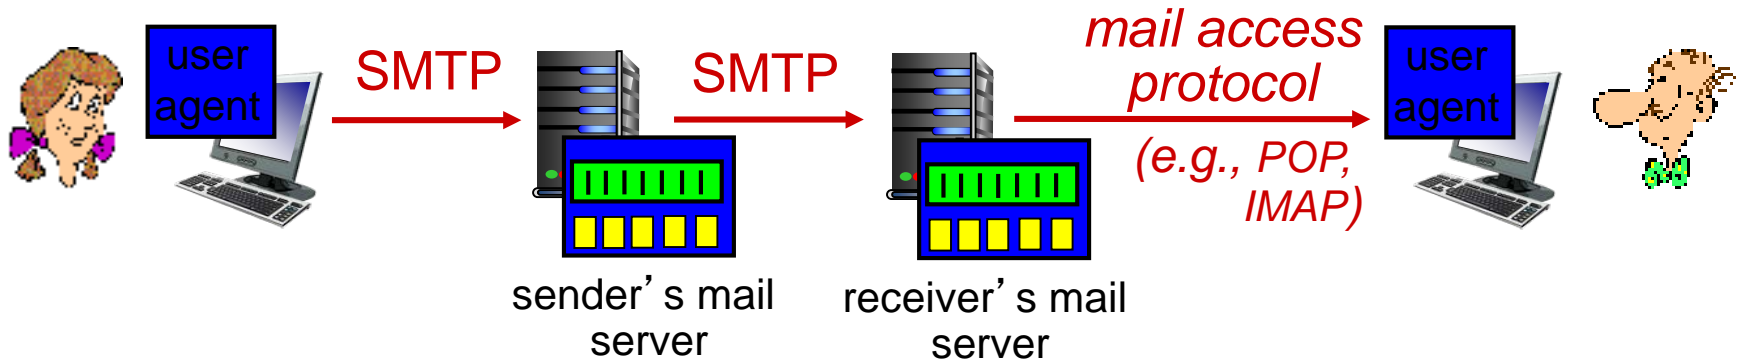
SMTP & Mail message format

RFC 822: standard for text message format:

- header lines, e.g.,
 - To:, From:, Subject:*different from* SMTP MAIL
FROM, RCPT TO:
commands
- Body: the “message”
 - ASCII 7-bit characters only




Mail access protocols



- **SMTP:** delivery/storage to receiver's server
- mail access protocol: retrieval from server
 - **POP:** Post Office Protocol [RFC 1939]: authorization, download
 - **IMAP:** Internet Mail Access Protocol [RFC 1730]: more features, including manipulation of stored msgs on server

Roadmap



- Addressing and Applications needs from transport layer
- Http
 - General description and functionality
 - Authentication, cookies and related aspects
 - Caching and proxies
- SMTP (POP, IMAP)
- DNS

DNS: Domain Name System

People: many identifiers:

- SSN, name, Passport #

Internet hosts, routers: IP address (32 bit) - used for addressing datagrams (129.16.237.85)

- name (alphanumeric addresses) hard to process @ router
- “name”, e.g., (www.cs.chalmers.se)- used by humans

Q: map between IP addresses and name ?

Hostname to IP address translation

- **Example:** `www.chalmers.se` `129.16.71.10`
- File with mapping may be edited on the system
 - Unix: `/etc/hosts`
 - Windows: `c:\windows\system32\drivers\etc\hosts`
 - Example of an entry manually entered in the file:
 - `"129.16.20.245 fibula.ce.chalmers.se fibula"`

Does not scale for all possible hosts,
hard to change

- All hosts need one copy of the file
Impossible on the Internet

Alternative: DNS, a large
distributed database
DNS – **D**omain **N**ame **S**ystem

DNS: services, structure

DNS services


- hostname to IP address translation
- host aliasing
 - canonical, alias names
- mail server aliasing
- load distribution
 - replicated Web servers: many IP addresses correspond to one name

why not centralize DNS?

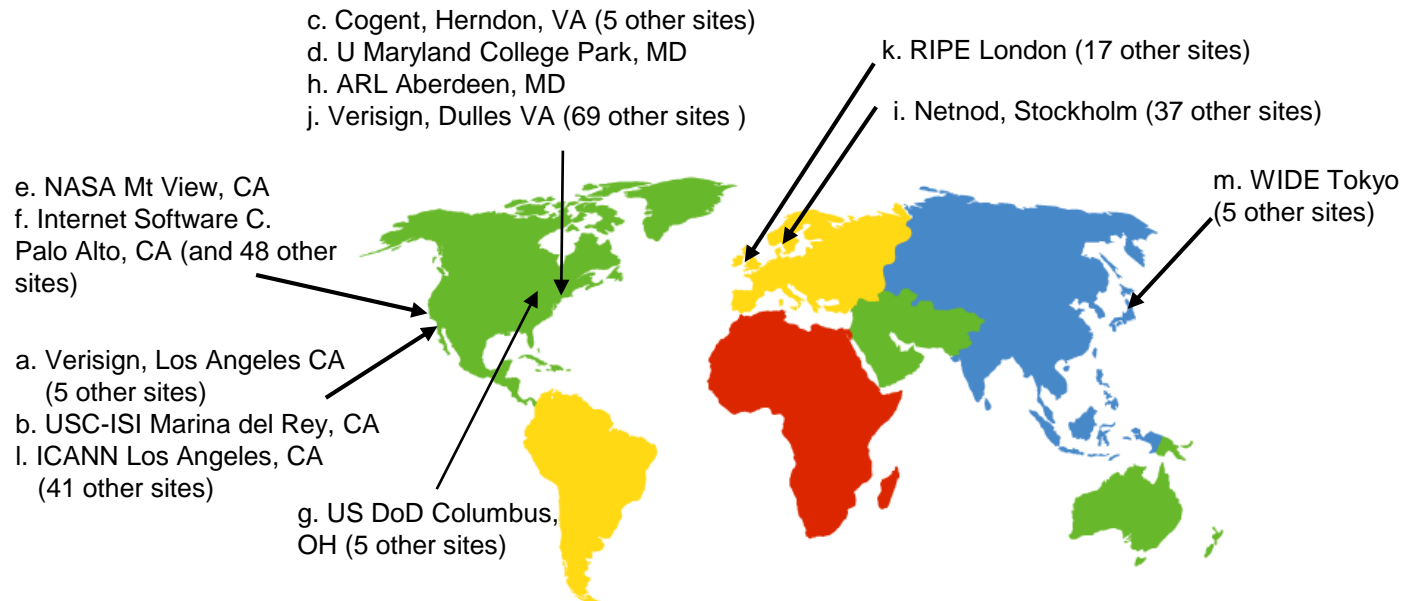
- single point of failure
- traffic volume
- maintenance

A: would not scale!

DNS: a distributed, hierarchical database



DNS: root name servers



*13 logical root name “servers”
worldwide - each “server”
replicated many times
<http://www.root-servers.org/>*

TLD, authoritative, local servers

top-level domain (TLD) servers:

- responsible for com, org, net, edu, aero, jobs, all top-level country domains, e.g.: uk, fr, ca, jp
- Network Solutions maintains servers for .com TLD
- Educause for .edu TLD

authoritative DNS servers:

- organization's own DNS server(s), providing authoritative hostname to IP mappings for organization's named hosts
- can be maintained by organization or service provider

Local name server


- acts as proxy for clients, caches entries for TTL
- Sends queries to DNS hierarchy
- each ISP has one

DNS name resolution example

- host at cis.poly.edu wants IP address for gaia.cs.umass.edu

iterated query:


- contacted server replies with name of server to contact
- “I don’t know this name, but ask this server”



DNS name resolution example

recursive query:

- puts burden of name resolution on contacted name server
- heavy load at upper levels of hierarchy?



DNS: caching, updating records

- once (any) name server learns mapping, it *caches* mapping
 - cache entries timeout after some time (TTL)
 - TLD servers typically cached in local name servers
 - thus root name servers not often visited
- cached entries may be *out-of-date* (best effort name-to-address translation!)
 - if name host changes IP address, may not be known Internet-wide until all TTLs expire
- update/notify mechanisms proposed IETF standard
 - RFC 2136

DNS records

DNS: distributed db storing resource records (**RR**)

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

type=A

- **name** is hostname
- **value** is IP address

type=NS

- **name** is domain (e.g., foo.com)
- **value** is hostname of authoritative name server for this domain

type=CNAME

- **name** is alias name for some “canonical” (the real) name
- Eg **www.ibm.com** is really **servereast.backup2.ibm.com**
- **value** is canonical name

type=MX

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

DNS protocol, messages

- *query* and *reply* messages (use UDP), both with same *message format*

identification: 16 bit #

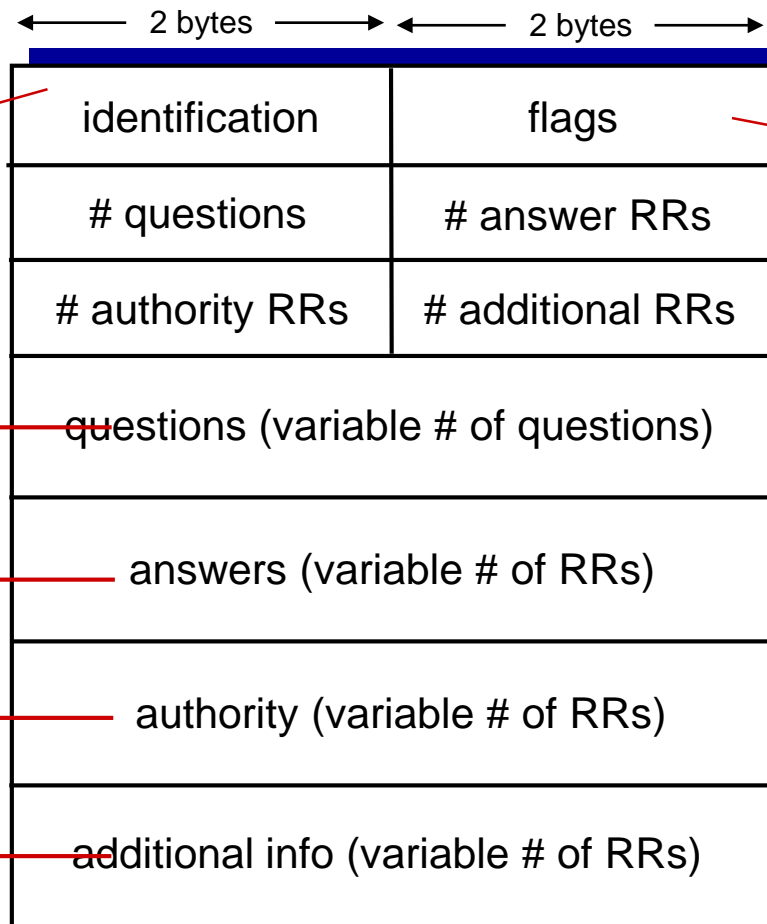
for query, reply to
query uses same #

name, type fields
for a query

RRs in response
to query

records for
authoritative servers

additional “helpful”
info that may be used



flags:

- query or reply
- recursion desired
- recursion available
- reply is authoritative

Inserting records into DNS

- example: new startup “Network Utopia”
- register name networkutopia.com at *DNS registrar* (e.g., Network Solutions)
 - provide names, IP addresses of *authoritative name* server (primary and secondary)
 - registrar inserts two RRs into .com TLD server:
(networkutopia.com, dns1.networkutopia.com, NS)
(dns1.networkutopia.com, 212.212.212.1, A)
- Adding a *new host/service to domain*:
 - *Add to authoritative name server*
 - type A record for www.networkutopia.com
 - type MX record for networkutopia.com (mail)

DNS and security risks

DDoS attacks

- Bombard root servers
 - Mitigation (it actually works 😊): local DNS servers cache IPs of TLD servers, allowing root server bypass
- Bombard TLD servers
 - Potentially more dangerous

Redirect attacks

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

Exploit DNS for DDoS


- Send queries with spoofed source address: target IP

Summary

- Addressing and Applications needs from transport layer
- application architectures
 - client-server
 - *(p2p: will study later in the course, after the layers-centered study)*
- specific protocols:
 - Http (connection to programming assignment)
 - Caching etc
 - SMTP (POP, IMAP)
 - DNS

Coming soon, after a first pass of the 4 top layers

- P2P applications
- video streaming and content distribution networks



Resources

Reading list main textbook:

- Study: 5/e,6/e: 2.2, 2.4-2.5, 5/e: 2.7-2.9, 6/e: 2.7-2.8, 7/e: 2.2-2.4
- Quick reading: 5/e,6/e: 2.1, 2.3, 2.6, 7/e: 2.1, 2.5

Review questions from the book, useful for summary study

- Chapter 2: 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 16, 19, 20

Example review question

Properties of transport service of interest to the app.

Data loss

- some apps (e.g., audio) can tolerate some loss
- other apps (e.g., file transfer, telnet) require 100% reliable data transfer
- In-order vs arbitrary-order delivery

Bandwidth, Timing, Security

- ❑ some apps (e.g., multimedia, interactive games) require minimum amount of **bandwidth**, and/or low **delay** and/or low **jitter**
- ❑ other apps (elastic apps, e.g. file transfer) are ok with any bandwidth, timing they get

Some apps also require **confidentiality**, **integrity** (more in network security)

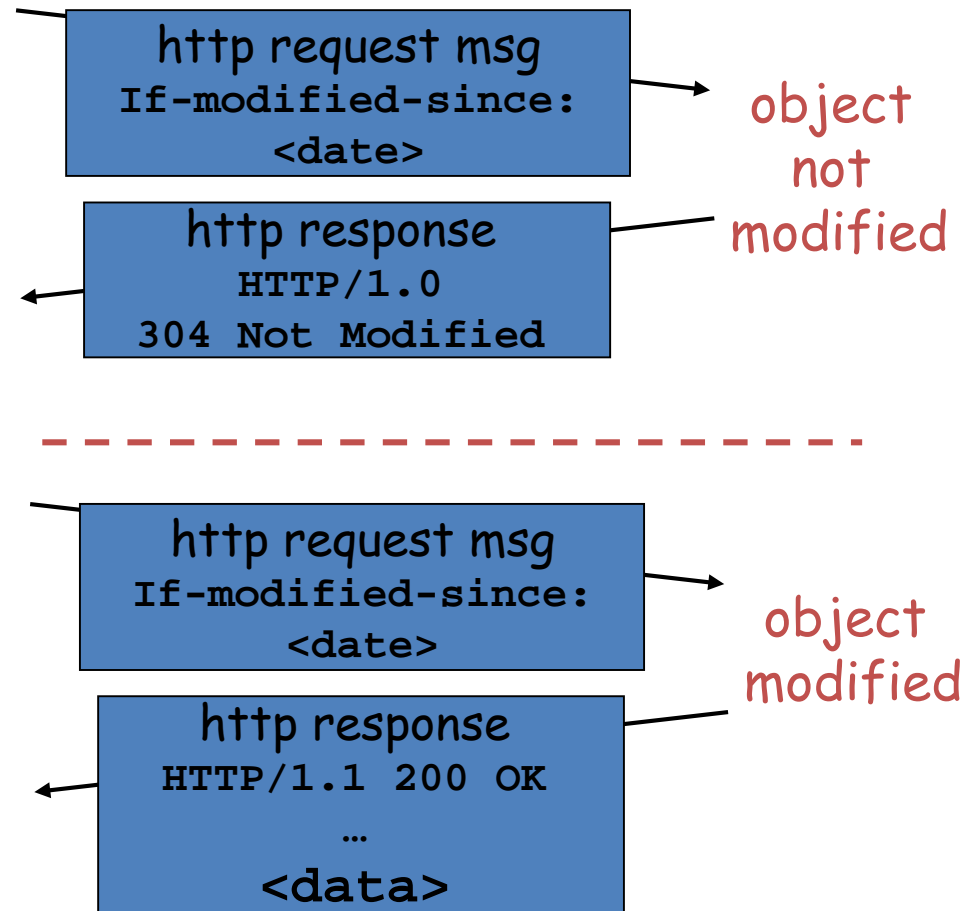
Extra slides/notes

HTTP: Conditional GET: client-side caching

- **Goal:** don't send object if client has up-to-date stored (cached) version
- client: specify date of cached copy in http request
`If-modified-since: <date>`
- server: response contains no object if cached copy up-to-date:
`HTTP/1.0 304 Not Modified`

client

server




Caching example:

assumptions:

- ❖ avg object size: 100K bits
- ❖ avg request rate from browsers to origin servers: 15/sec
 - ❖ i.e. avg data rate to browsers: 1.50 Mbps
- ❖ RTT from institutional router to any origin server: 2 sec
- ❖ access link rate: 1.54 Mbps

consequences:

- ❖ LAN utilization: 1.5% *problem!*
- ❖ access link utilization = **99%**
- ❖ total delay = Internet delay + access delay + LAN delay
= 2 sec + minutes + quite_small




Caching example: faster access link

assumptions:

- ❖ avg object size: 100K bits
- ❖ avg request rate from browsers to origin servers: 15/sec
 - ❖ i.e. avg data rate to browsers: 1.50 Mbps
- ❖ RTT from institutional router to any origin server: 2 sec
- ❖ access link rate: ~~1.54 Mbps~~

consequences:

- ❖ LAN utilization: 1.5%
- ❖ access link utilization = ~~99%~~ → 9.9%
- ❖ total delay = Internet delay + access delay + LAN delay
= 2 sec → msec + minutes + usecs



Cost: increased access link speed (not cheap!)

Caching example: install local cache

assumptions:


- ❖ avg object size: 100K bits
- ❖ avg request rate from browsers to origin servers: 15/sec
 - ❖ i.e. avg data rate to browsers: 1.50 Mbps
- ❖ RTT from institutional router to any origin server: 2 sec
- ❖ access link rate: 1.54 Mbps

consequences:

- ❖ LAN utilization: 1.5%
- ❖ access link utilization ?
- ❖ total delay ?

How to compute link utilization, delay?


Cost: web cache (cheap!)



Caching example: install local cache


Calculating access link utilization, delay with cache:

- suppose cache hit rate is 0.4
 - 40% requests satisfied at cache, 60% requests satisfied at origin
- ❖ access link utilization:
 - 60% of requests use access link
- ❖ data rate to browsers over access link = $0.6 * 1.50 \text{ Mbps} = .9 \text{ Mbps}$
 - utilization = $0.9 / 1.54 = .58$
- ❖ total delay
 - = $0.6 * (\text{delay from origin servers}) + 0.4 * (\text{delay when satisfied at cache})$
 - = $0.6 (2.01) + 0.4 (\sim \text{msecs})$
 - = $\sim 1.2 \text{ secs}$
 - less than with 154 Mbps link (and cheaper too!)



FTP: separate control, data connections

- FTP client contacts FTP server at port 21, using TCP
- client authorized over control connection
- client browses remote directory, sends commands over control connection
- when server receives file transfer command, *server* opens 2nd TCP data connection (for file) to client
- after transferring one file, server closes data connection



- ❖ server opens another TCP data connection to transfer another file
- ❖ control connection: “*out of band*”
- ❖ FTP server maintains “state”: current directory, earlier authentication

FTP commands, responses

sample commands:

- sent as ASCII text over control channel
- **USER *username***
- **PASS *password***
- **LIST** return list of file in current directory
- **RETR *filename*** retrieves (gets) file
- **STOR *filename*** stores (puts) file onto remote host

sample return codes

- status code and phrase (as in HTTP)
- **331 Username OK, password required**
- **125 data connection already open; transfer starting**
- **425 Can't open data connection**
- **452 Error writing file**

POP3 protocol

authorization phase

- client commands:
 - **user**: declare username
 - **pass**: password
- server responses
 - +OK
 - -ERR

transaction phase, client:

- **list**: list message numbers
- **retr**: retrieve message by number
- **dele**: delete
- **quit**

```
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 1 contents>
S: .
C: dele 2
C: quit
S: +OK POP3 server signing off
```

POP3 (more) and IMAP

more about POP3

- previous example uses POP3 “download and delete” mode
 - Bob cannot re-read e-mail if he changes client
- POP3 “download-and-keep”: copies of messages on different clients
- POP3 is stateless across sessions

IMAP

- keeps all messages in one place: at server
- allows user to organize messages in folders
- keeps user state across sessions:
 - names of folders and mappings between message IDs and folder name