

Digital Communication in the Modern World

Application Layer cont. DNS, SMTP

<http://www.cs.huji.ac.il/~com1>
com1@cs.huji.ac.il

Some of the slides have been borrowed from:
Computer Networking: A Top Down Approach Featuring the Internet,
2nd edition.
Jim Kurose, Keith Ross
Addison-Wesley, July 2002.

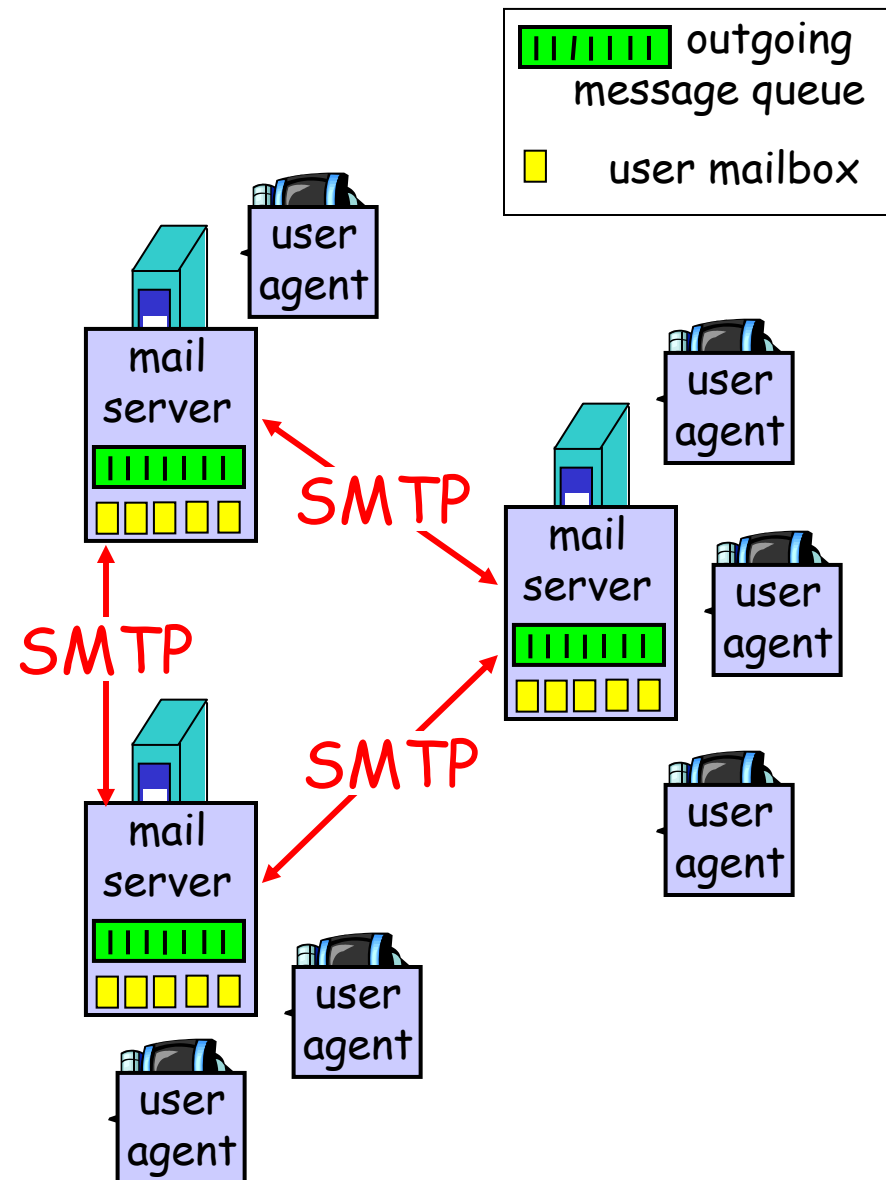
Electronic Mail

Three major components:

- ❑ user agents (clients)
- ❑ mail servers
- ❑ simple mail transfer protocol: SMTP

User Agent


- ❑ a.k.a. "mail reader"
- ❑ composing, editing, reading mail messages
- ❑ e.g., Eudora, Outlook, elm, Netscape Messenger, PINE
- ❑ outgoing, incoming messages stored on server



Electronic Mail: mail servers

Mail Servers

- ❑ **mailbox** contains incoming messages for user
- ❑ **message queue** of outgoing (to be sent) mail messages
- ❑ **SMTP protocol** between mail servers to send email messages
 - client: sending mail server
 - "server": receiving mail server

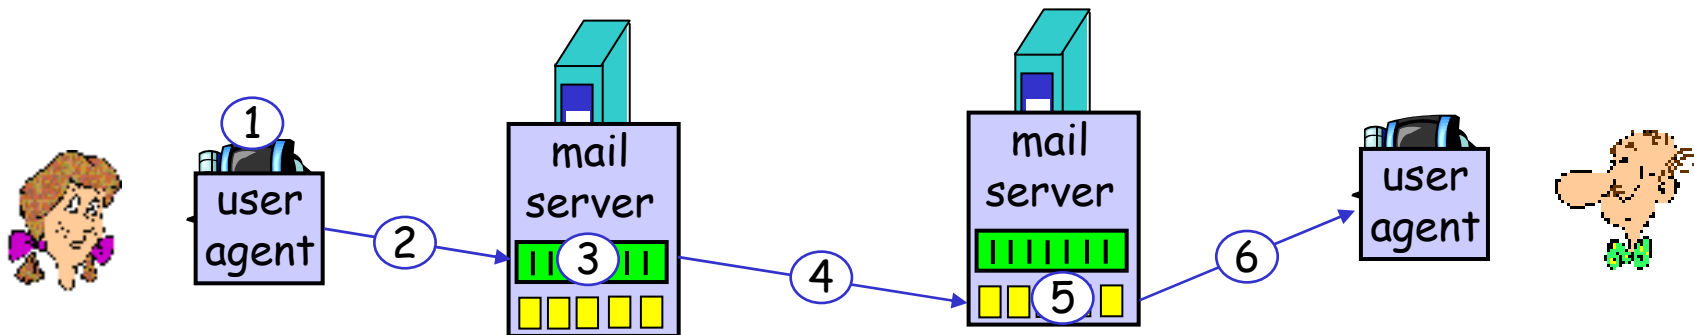


Electronic Mail: SMTP [RFC 2821]

- ❑ uses TCP to reliably transfer email message from client to server, port 25
- ❑ direct transfer: sending server to receiving server
- ❑ three phases of transfer
 - handshaking (greeting)
 - transfer of messages
 - closure
- ❑ command/response interaction
 - **commands**: ASCII text
 - **response**: status code and phrase
- ❑ messages must be in 7-bit ASCII

Scenario: Alice sends message to Bob

- 1) Alice uses UA to compose message and "to" bob@someschool.edu
- 2) Alice's UA sends message to her mail server; message placed in message queue
- 3) Client side of SMTP opens TCP connection with Bob's mail server
- 4) SMTP client sends Alice's message over the TCP connection
- 5) Bob's mail server places the message in Bob's mailbox
- 6) Bob invokes his user agent to read message



Sample SMTP interaction

```
S: 220 mail.cs.huji.ac.il
C: HELO mail.cs.huji.ac.il
S: 250 Hello mail.cs.ac.il, pleased to meet you
C: MAIL FROM: <falafel@cs.ac.il>
S: 250 falafel@cs.ac.il... Sender ok
C: RCPT TO: <sabih@pita.com>
S: 250 sabih@pita.co ... Recipient ok
C: DATA
S: 354 Enter mail, end with "." on a line by itself
C: Do you want with hilbe?
C:   How about amba?
C: .
S: 250 Message accepted for delivery
C: QUIT
S: 221 mail.cs.huji.ac.il closing connection
```

Try SMTP interaction for yourself:

- ❑ `telnet servername 25`
- ❑ see 220 reply from server
- ❑ enter HELO, MAIL FROM, RCPT TO, DATA, QUIT commands

above lets you send email without using email client (reader)

SMTP: final words

- ❑ SMTP uses persistent connections
- ❑ SMTP requires message (header & body) to be in 7-bit ASCII
- ❑ SMTP server uses CRLF.CRLF to determine end of message

Comparison with HTTP:

- ❑ HTTP: pull
- ❑ SMTP: push
- ❑ both have ASCII command/response interaction, status codes
- ❑ HTTP: each object encapsulated in its own response msg
- ❑ SMTP: multiple objects sent in multipart msg

Mail message format

SMTP: protocol for exchanging email msgs

RFC 822: standard for text message format:


□ header lines, e.g.,

- To:
- From:
- Subject:

different from SMTP commands!

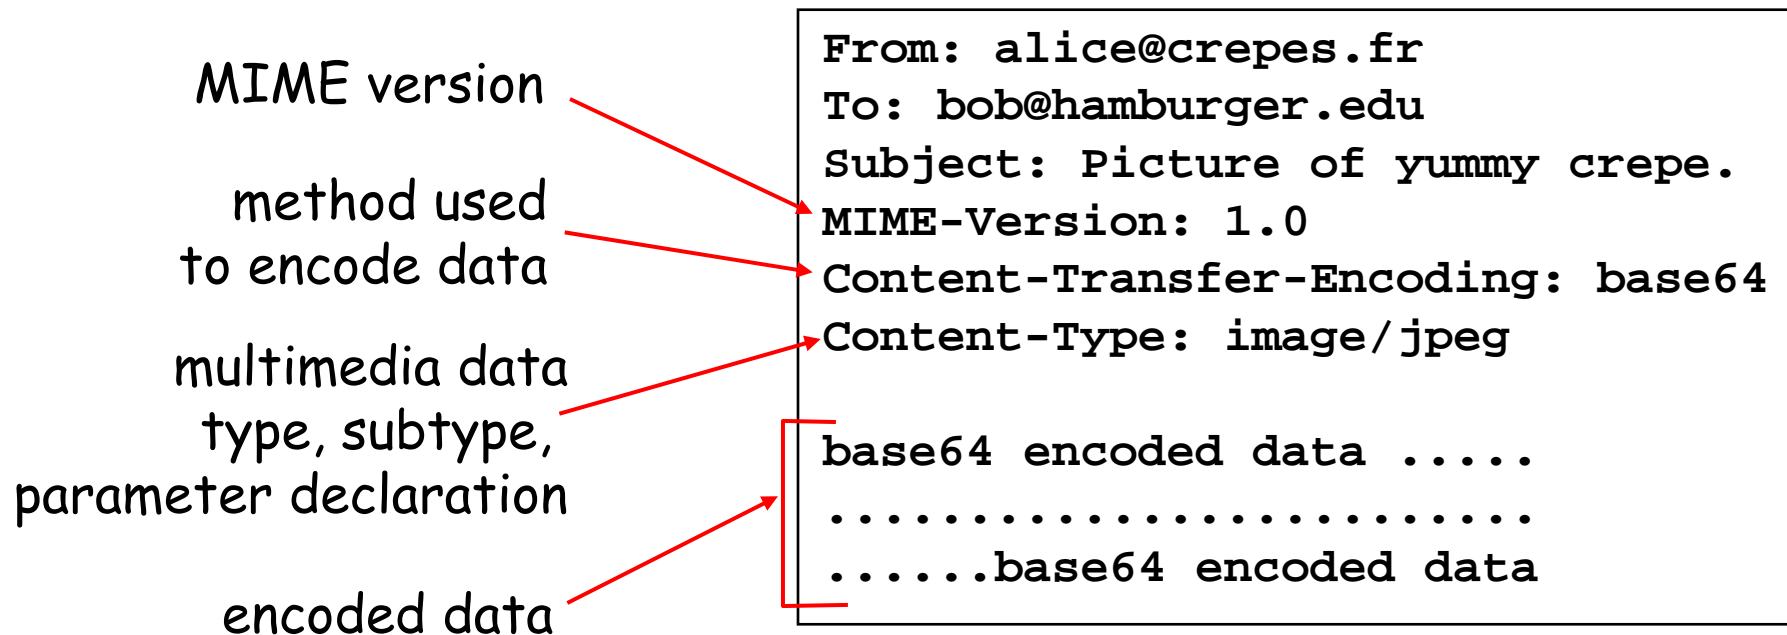
□ body

- the "message", ASCII characters only



Message format: multimedia extensions

- ❑ MIME: multimedia mail extension, RFC 2045, 2056
- ❑ additional lines in msg header declare MIME content type



MIME types

Content-Type: type/subtype; parameters

Text

- ❑ example subtypes: plain, html

Video

- ❑ example subtypes: mpeg, quicktime

Image

- ❑ example subtypes: jpeg, gif

Application

- ❑ other data that must be processed by reader before "viewable"
- ❑ example subtypes: msword, octet-stream

Audio

- ❑ example subtypes: basic (8-bit mu-law encoded), 32kadtcm (32 kbps coding)

Multipart Type

From: alice@crepes.fr
To: bob@hamburger.edu
Subject: Picture of yummy crepe.
MIME-Version: 1.0
Content-Type: multipart/mixed; boundary=StartOfNextPart

--StartOfNextPart

Dear Bob, Please find a picture of a crepe.

--StartOfNextPart

Content-Transfer-Encoding: base64

Content-Type: image/jpeg

base64 encoded data


.....

.....base64 encoded data

--StartOfNextPart

Do you want the recipe?

Mail access protocols



- ❑ SMTP: delivery/storage to receiver's server
- ❑ Mail access protocol: retrieval from server
 - POP: Post Office Protocol [RFC 1939]
 - authorization (agent <-->server) and download
 - IMAP: Internet Mail Access Protocol [RFC 1730]
 - more features (more complex)
 - manipulation of stored msgs on server
 - HTTP: Hotmail , Yahoo! Mail, Gmail, etc.

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

IMAP

- ❑ Keep all messages in one place: the server
- ❑ Allows user to organize messages in folders
- ❑ IMAP keeps user state across sessions:
 - names of folders and mappings between message IDs and folder name

DNS: Domain Name System

People: many identifiers:

- SSN, name, passport #

Internet hosts, routers:

- IP address (32 bit) - used for addressing datagrams
- "name", e.g., gaia.cs.umass.edu - used by humans

Q: map between IP addresses and name ?

Domain Name System:

- *distributed database*
implemented in hierarchy of many *name servers*
- *application-layer protocol*
host, routers, name servers to communicate to *resolve* names (address/name translation)
 - note: core Internet function, implemented as application-layer protocol
 - complexity at network's "edge"

DNS name servers

Why not centralize DNS?

- ❑ single point of failure
- ❑ traffic volume
- ❑ distant centralized database
- ❑ maintenance

doesn't scale!

- ❑ no server has all name-to-IP address mappings

local name servers:

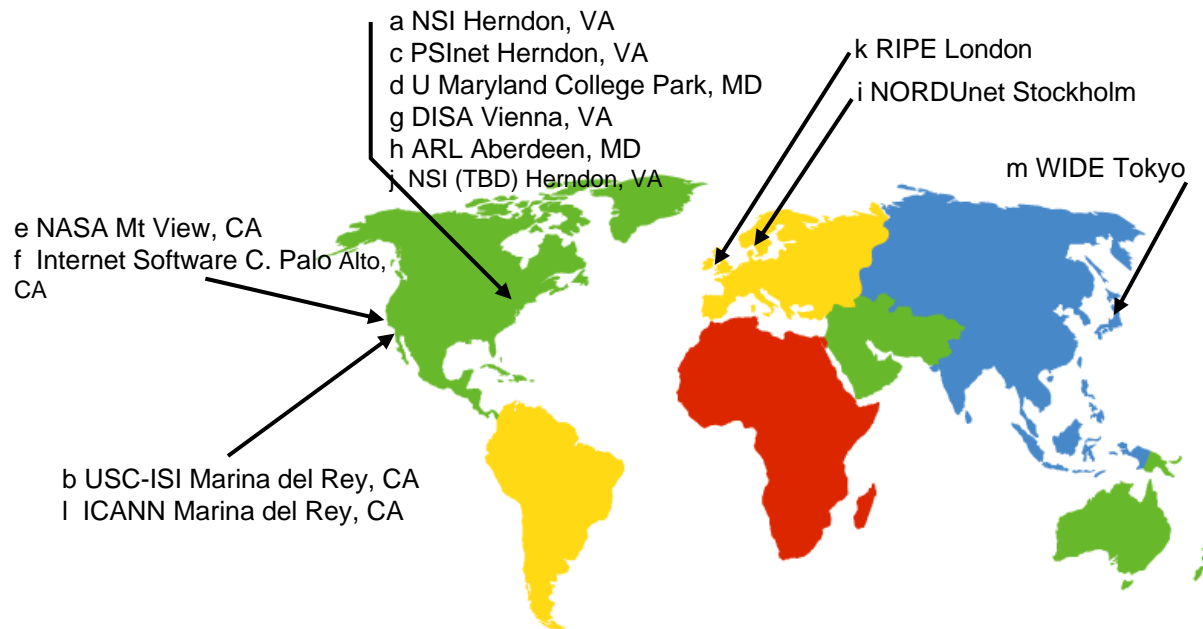
- each ISP, company has *local (default) name server*
- host DNS query first goes to local name server

authoritative name server:

- for a host: stores that host's IP address, name
- can perform name/address translation for that host's name

DNS: Root name servers

- ❑ contacted by local name server that can not resolve name
- ❑ root name server:
 - contacts authoritative name server if name mapping not known
 - gets mapping
 - returns mapping to local name server




13 root name
servers worldwide

Simple DNS example

host `surf.eurecom.fr`
wants IP address of
`gaia.cs.umass.edu`


1. contacts its local DNS server, `dns.eurecom.fr`
2. `dns.eurecom.fr` contacts root name server, if necessary
3. root name server contacts authoritative name server, `dns.umass.edu`, if necessary



DNS example

Root name server:

- ❑ may not know authoritative name server
- ❑ may know *intermediate name server*: who to contact to find authoritative name server




DNS: iterated queries

recursive query:

- ❑ puts burden of name resolution on contacted name server
- ❑ heavy load?

iterated query:

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



DNS: caching and updating records

- ❑ once (any) name server learns mapping, it *caches* mapping
 - cache entries timeout (disappear) after some time
- ❑ update/notify mechanisms under design by IETF
 - RFC 2136
 - <http://www.ietf.org/html.charters/dnsind-charter.html>

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 IP address of authoritative name server for this domain

□ Type=CNAME

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

□ Type=MX

- value is name of mailserver associated with name

DNS protocol, messages

DNS protocol : *query* and *reply* messages, both with same *message format*

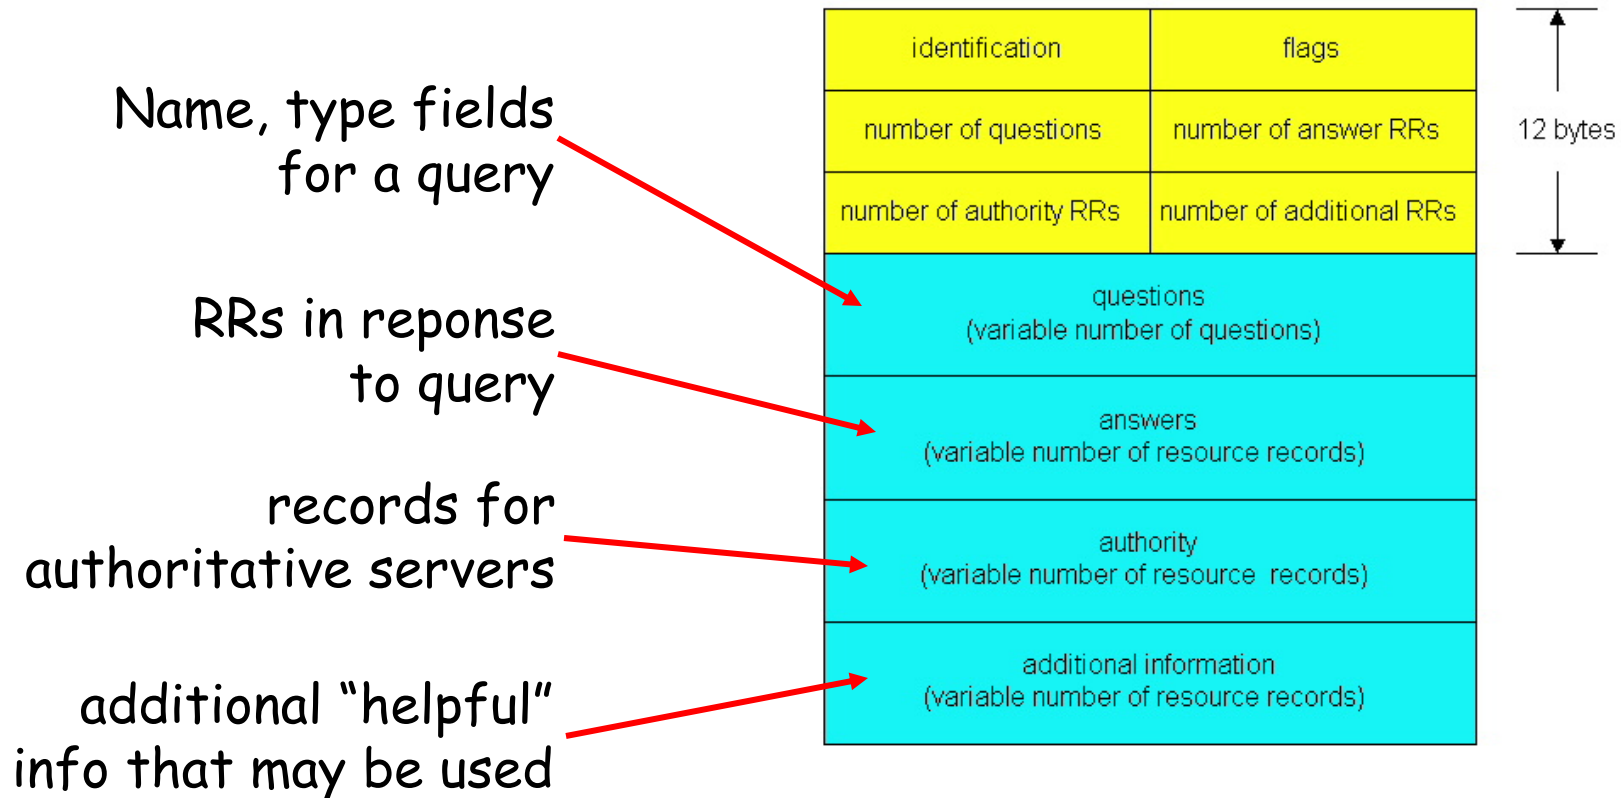
msg header

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

identification	flags
number of questions	number of answer RRs
number of authority RRs	number of additional RRs
questions (variable number of questions)	
answers (variable number of resource records)	
authority (variable number of resource records)	
additional information (variable number of resource records)	

↑
12 bytes
↓


DNS protocol, messages



Web caches (proxy server)

Goal: satisfy client request without involving origin server

- user sets browser: Web accesses via cache
- browser sends all HTTP requests to cache
 - object in cache: cache returns object
 - else cache requests object from origin server, then returns object to client



More about Web caching

- ❑ Cache acts as both client and server
- ❑ Cache can do up-to-date check using If-modified-since HTTP header
 - Issue: should cache take risk and deliver cached object without checking?
 - Heuristics are used.
- ❑ Typically cache is installed by ISP (university, company, residential ISP)

Why Web caching?

- ❑ Reduce response time for client request.
- ❑ Reduce traffic on an institution's access link.
- ❑ Internet dense with caches enables "poor" content providers to effectively deliver content


Caching example (1)

Assumptions

- ❑ average object size = 100,000 bits
- ❑ avg. request rate from institution's browser to origin servers = 15/sec
- ❑ delay from institutional router to any origin server and back to router = 2 sec

Consequences

- ❑ utilization on LAN = 15%
- ❑ utilization on access link = 100%
- ❑ total delay = Internet delay + access delay + LAN delay
= 2 sec + minutes + milliseconds




Caching example (2)

Possible solution

- increase bandwidth of access link to, say, 10 Mbps

Consequences

- utilization on LAN = 15%
- utilization on access link = 15%
- Total delay = Internet delay + access delay + LAN delay
= 2 sec + msecs + msecs
- often a costly upgrade




Caching example (3)

Install cache

- suppose hit rate is .4

Consequence

- 40% requests will be satisfied almost immediately
- 60% requests satisfied by origin server
- utilization of access link reduced to 60%, resulting in negligible delays (say 10 msec)
- total delay = Internet delay + access delay + LAN delay
$$= .6 * 2 \text{ sec} + .6 * .01 \text{ secs} + \text{milliseconds} < 1.3 \text{ secs}$$




Content distribution networks (CDNs)

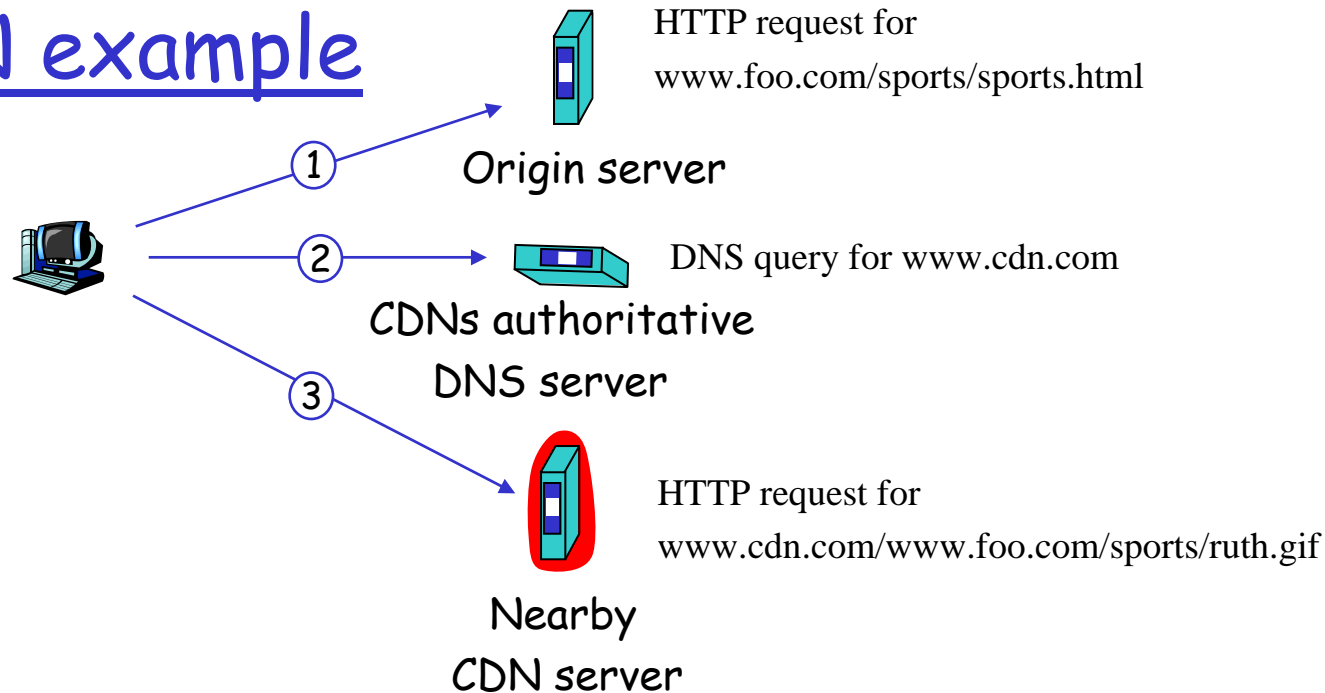
- ❑ The content providers are the CDN customers.

Content replication

- ❑ CDN company installs hundreds of CDN servers throughout Internet
 - in lower-tier ISPs, close to users
- ❑ CDN replicates its customers' content in CDN servers. When provider updates content, CDN updates servers



CDN example



origin server

- ❑ www.foo.com
- ❑ distributes HTML
- ❑ Replaces:
http://www.foo.com/sports.ruth.gif
with

http://www.cdn.com/www.foo.com/sports/ruth.gif

CDN company

- ❑ cdn.com
- ❑ distributes gif files
- ❑ uses its authoritative
DNS server to route
redirect requests

More about CDNs

routing requests

- ❑ CDN creates a “map”, indicating distances from leaf ISPs and CDN nodes
- ❑ when query arrives at authoritative DNS server:
 - server determines ISP from which query originates
 - uses “map” to determine best CDN server

not just Web pages

- ❑ streaming stored audio/video
- ❑ streaming real-time audio/video
 - CDN nodes create application-layer overlay network