

Week 2 Application Layer



These slides are modified from the slides made available by Kurose and Ross.

Computer Networking: A Top Down Approach Featuring the Internet, 2nd edition.
Jim Kurose, Keith Ross
Addison-Wesley, July 2002.

2: Application Layer 1

Week 2: Application Layer

Our goals:

- conceptual, implementation aspects of network application protocols
 - transport-layer service models
 - client-server paradigm
 - peer-to-peer paradigm
- learn about protocols by examining popular application-level protocols
 - HTTP
 - FTP
 - SMTP / POP3 / IMAP
 - DNS
- programming network applications
 - socket API

2: Application Layer 2

Network applications: some jargon

Process: program running within a host.

- within same host, two processes communicate using **interprocess communication** (defined by OS).
- processes running in different hosts communicate with an **application-layer protocol**

user agent: interfaces with user "above" and network "below".

- implements user interface & application-level protocol
 - Web: browser
 - E-mail: mail reader
 - streaming audio/video: media player

2: Application Layer 3

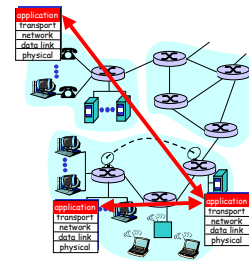
Applications and application-layer protocols

Application: communicating, distributed processes

- e.g., e-mail, Web, P2P file sharing, instant messaging
- running in end systems (hosts)
- exchange messages to implement application

Application-layer protocols

- one "piece" of an app
- define messages exchanged by apps and actions taken
- use communication services provided by lower layer protocols (TCP, UDP)



2: Application Layer 4

App-layer protocol defines

- Types of messages exchanged, eg, request & response messages
- Syntax of message types: what fields in messages & how fields are delineated
- Semantics of the fields, ie, meaning of information in fields
- Rules for when and how processes send & respond to messages

Public-domain protocols:

- defined in RFCs
- allows for interoperability
- eg, HTTP, SMTP

2: Application Layer 5

Client-server paradigm

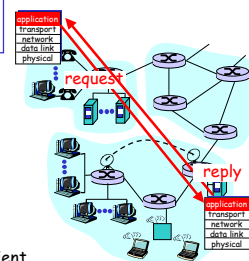
Typical network app has two pieces: **client** and **server**

Client:

- initiates contact with server ("speaks first")
- typically requests service from server,
- Web: client implemented in browser; e-mail: in mail reader

Server:

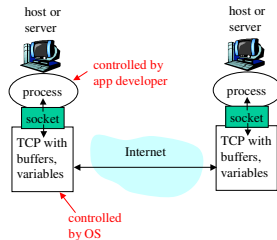
- provides requested service to client
- e.g., Web server sends requested Web page, mail server delivers e-mail



2: Application Layer 6

Processes communicating across network

- process sends/receives messages to/from its socket
- socket analogous to door
 - sending process shoves message out door
 - sending process assumes transport infrastructure on other side of door which brings message to socket at receiving process
- API: (1) choice of transport protocol; (2) ability to fix a few parameters (lots more on this later)



2: Application Layer 7

Addressing processes:

- For a process to receive messages, it must have an identifier
- Every host has a unique 32-bit IP address
- Q: does the IP address of the host on which the process runs suffice for identifying the process?
- Answer: No, many processes can be running on same host
- Identifier includes both the IP address and **port numbers** associated with the process on the host.
- Example port numbers:
 - HTTP server: 80
 - Mail server: 25
- More on this later

2: Application Layer 8

What transport service does an app need?

Data loss

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

Timing

- some apps (e.g., Internet telephony, interactive games) require low delay to be "effective"

Bandwidth

- some apps (e.g., multimedia) require minimum amount of bandwidth to be "effective"
- other apps ("elastic apps") make use of whatever bandwidth they get

2: Application Layer 9

Transport service requirements of common apps

Application	Data loss	Bandwidth	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	yes, 100's msec
stored audio/video	loss-tolerant	same as above	yes, few secs
interactive games	loss-tolerant	few kbps up	yes, 100's msec
instant messaging	no loss	elastic	yes and no

2: Application Layer 10

Internet transport protocols services

TCP service:

- connection-oriented:** setup required between client and server processes
- reliable transport** between sending and receiving process
- flow control:** sender won't overwhelm receiver
- congestion control:** throttle sender when network overloaded
- does not providing:** timing, minimum bandwidth guarantees

UDP service:

- unreliable data transfer between sending and receiving process
- does not provide: connection setup, reliability, flow control, congestion control, timing, or bandwidth guarantee

Q: why bother? Why is there a UDP?

2: Application Layer 11

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	proprietary (e.g. RealNetworks)	TCP or UDP
Internet telephony	proprietary (e.g., Dialpad)	typically UDP

2: Application Layer 12

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**
- Example URL:

`www.cs.rit.edu/somcourse/pic.gif`

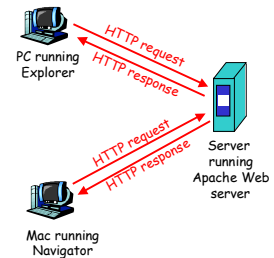
host name path name

2: Application Layer 13

HTTP overview

HTTP: hypertext transfer protocol

- Web's application layer protocol
- client/server model
 - **client**: browser that requests, receives, "displays" Web objects
 - **server**: Web server sends objects in response to requests
- HTTP 1.0: RFC 1945
- HTTP 1.1: RFC 2068



2: Application Layer 14

HTTP overview (continued)

Uses TCP:

- client initiates TCP connection (creates socket) to server, port 80
- server accepts TCP connection from client
- HTTP messages (application-layer protocol messages) exchanged between browser (HTTP client) and Web server (HTTP server)
- TCP connection closed

HTTP is "stateless"

- 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

2: Application Layer 15

HTTP connections

Nonpersistent HTTP

- At most one object is sent over a TCP connection.
- HTTP/1.0 uses nonpersistent HTTP

Persistent HTTP

- Multiple objects can be sent over single TCP connection between client and server.
- HTTP/1.1 uses persistent connections in default mode

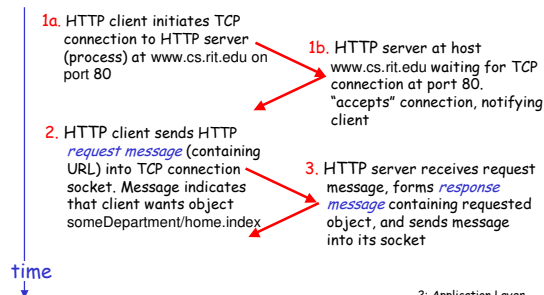
2: Application Layer 16

Nonpersistent HTTP

Suppose user enters URL

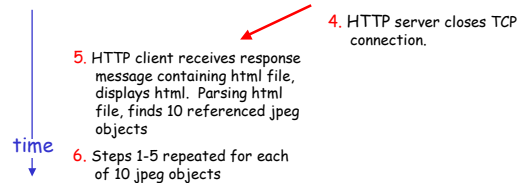
`www.cs.rit.edu/someDepartment/home.index`

(contains text, references to 10 jpeg images)



2: Application Layer 17

Nonpersistent HTTP (cont.)



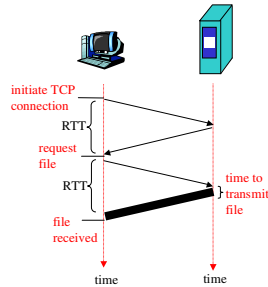
2: Application Layer 18

Response time modeling

Definition of RTT: time to send a small packet to travel from client to server and back.

Response time:

- one RTT to initiate TCP connection
 - one RTT for HTTP request and first few bytes of HTTP response to return
 - file transmission time
- total = 2RTT + transmit time**



2: Application Layer 19

Persistent HTTP

Nonpersistent HTTP issues:

- requires 2 RTTs per object
- OS must work and allocate host resources for each TCP connection
- but browsers often open parallel TCP connections to fetch referenced objects

Persistent HTTP

- server leaves connection open after sending response
- subsequent HTTP messages between same client/server are sent over connection

Persistent without pipelining:

- client issues new request only when previous response has been received
- one RTT for each referenced object

Persistent with pipelining:

- default in HTTP/1.1
- client sends requests as soon as it encounters a referenced object
- as little as one RTT for all the referenced objects

2: Application Layer 20

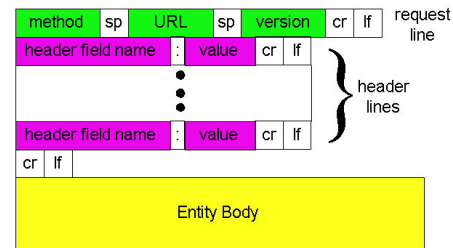
HTTP request message

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

request line (GET, POST, HEAD commands) → GET /somedir/page.html HTTP/1.1
 Host: www.someschool.edu
 User-agent: Mozilla/4.0
 Connection: close
 Accept-language: fr
 header lines
 Carriage return, line feed (extra carriage return, line feed) indicates end of message

2: Application Layer 21

HTTP request message: general format



2: Application Layer 22

Uploading form input

Post method:

- Web page often includes form input
- Input is uploaded to server in entity body

URL method:

- Uses GET method
- Input is uploaded in URL field of request line:

www.somesite.com/animalsearch?monkeys&banana

2: Application Layer 23

Method types

HTTP/1.0

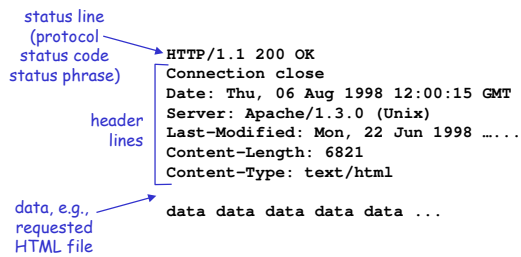
- GET
- POST
- HEAD
 - asks server to leave requested object out of response

HTTP/1.1

- GET, POST, HEAD
- PUT
 - uploads file in entity body to path specified in URL field
- DELETE
 - deletes file specified in the URL field

2: Application Layer 24

HTTP response message



2: Application Layer 25

HTTP response status codes

In first line in server->client response message.
A few sample codes:

- 200 OK**
 - request succeeded, requested object later in this message
- 301 Moved Permanently**
 - requested object moved, new location specified later in this message (Location:)
- 400 Bad Request**
 - request message not understood by server
- 404 Not Found**
 - requested document not found on this server
- 505 HTTP Version Not Supported**

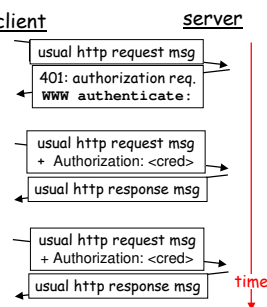
2: Application Layer 26

User-server interaction: authorization

Authorization: control access to server content

- authorization credentials: typically name, password
- stateless:** client must present authorization in *each* request
 - authorization: header line in each request
 - if no **authorization:** header, server refuses access, sends

WWW authenticate: header line in response



2: Application Layer 27

Cookies: keeping "state"

Many major Web sites use cookies

Four components:

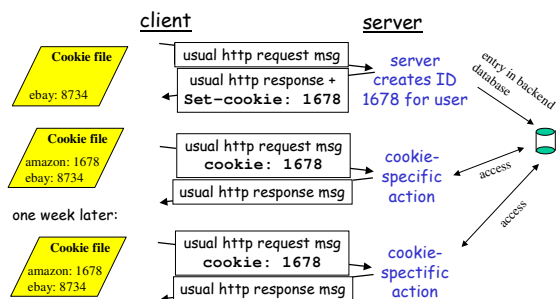
- cookie header line in the HTTP response message
- cookie header line in HTTP request message
- cookie file kept on user's host and managed by user's browser
- back-end database at Web site

Example:

- Susan access Internet always from same PC
- She visits a specific e-commerce site for first time
- When initial HTTP requests arrives at site, site creates a unique ID and creates an entry in backend database for ID

2: Application Layer 28

Cookies: keeping "state" (cont.)



2: Application Layer 29

Cookies (continued)

What cookies can bring:

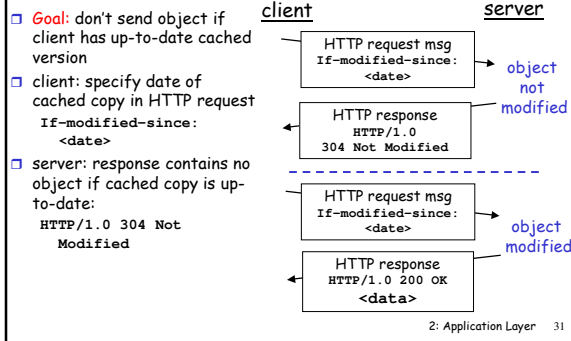
- authorization
- shopping carts
- recommendations
- user session state (Web e-mail)

Cookies and privacy:

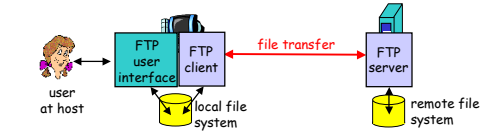
- cookies permit sites to learn a lot about you
- you may supply name and e-mail to sites
- search engines use redirection & cookies to learn yet more
- advertising companies obtain info across sites

2: Application Layer 30

Conditional GET: client-side caching



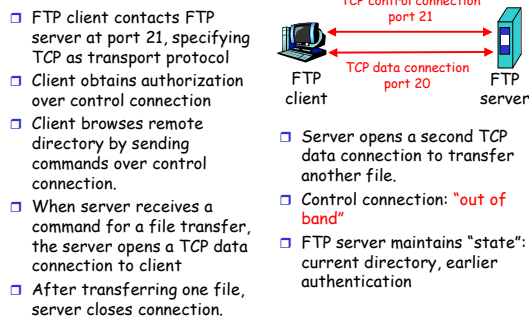
FTP: the file transfer protocol



- transfer file to/from remote host
- client/server model
 - client: side that initiates transfer (either to/from remote)
 - server: remote host
- ftp: RFC 959
- ftp server: port 21

2: Application Layer 32

FTP: separate control, data connections



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

2: Application Layer 34

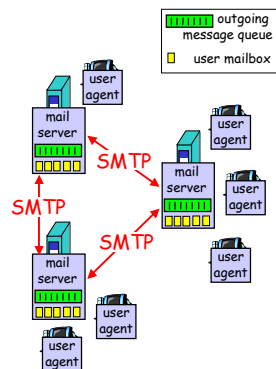
Electronic Mail

Three major components:

- user agents
- 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
- outgoing, incoming messages stored on server

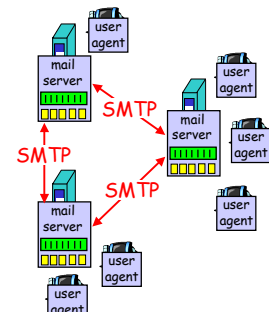


2: Application Layer 35

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



2: Application Layer 36

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

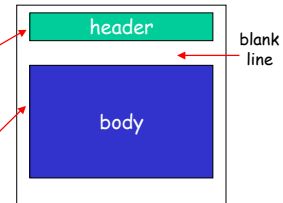
2: Application Layer 37

Mail message format

SMTP: protocol for exchanging email msgs

RFC 822: standard for text message format:

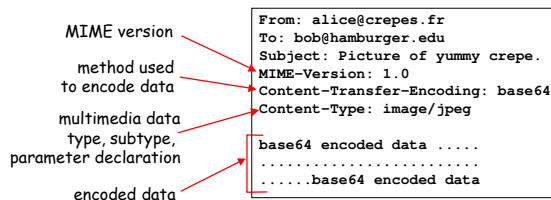
- header lines, e.g.,
 - To:
 - From:
 - Subject:*different from SMTP commands!*
- body
 - the "message", ASCII characters only



2: Application Layer 38

Message format: multimedia extensions

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



2: Application Layer 39

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)

2: Application Layer 40

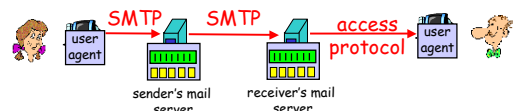
Multipart Type

```
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?
```

2: Application Layer 41

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, etc.

2: Application Layer 42

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"

2: Application Layer 43

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

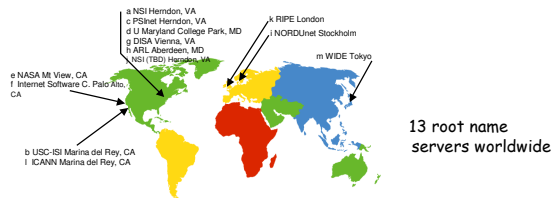
authoritative name server:

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

2: Application Layer 44

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

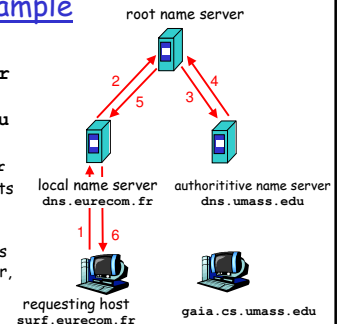


2: Application Layer 45

Simple DNS example

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

- contacts its local DNS server, dns.eurecom.fr
- dns.eurecom.fr contacts root name server, if necessary
- root name server contacts authoritative name server, dns.umass.edu, if necessary

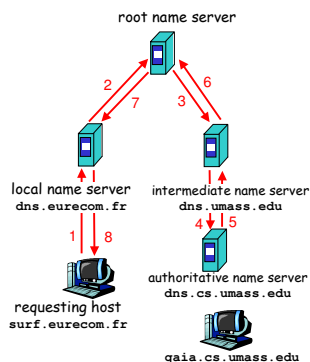


2: Application Layer 46

DNS example

Root name server:

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



2: Application Layer 47

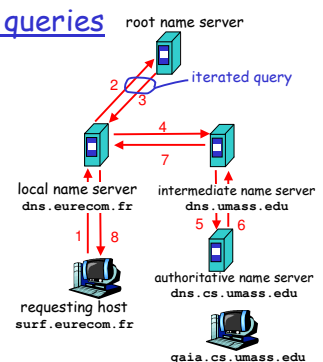
DNS: iterated queries

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"



2: Application Layer 48

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>

2: Application Layer 49

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=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=NS
 - name is domain (e.g. foo.com)
 - value is IP address of authoritative name server for this domain
- Type=MX
 - value is name of mailserver associated with name

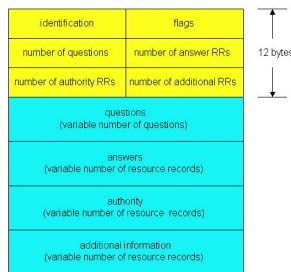
2: Application Layer 50

DNS protocol, messages

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

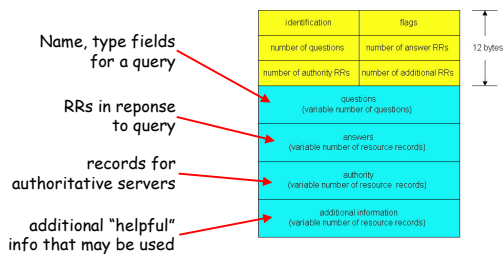
msg header

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



2: Application Layer 51

DNS protocol, messages



2: Application Layer 52

Socket programming

Goal: learn how to build client/server application that communicate using sockets

Socket API

- introduced in BSD4.1 UNIX, 1981
- explicitly created, used, released by apps
- client/server paradigm
- two types of transport service via socket API:
 - unreliable datagram
 - reliable, byte stream-oriented

socket

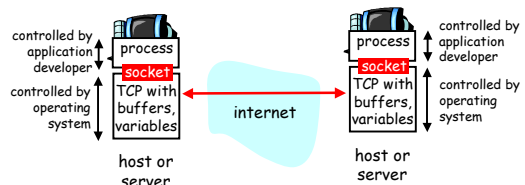
a *host-local, application-created, OS-controlled* interface (a "door") into which application process can both send and receive messages to/from another application process

2: Application Layer 53

Socket-programming using TCP

Socket: a door between application process and end-end-transport protocol (UDP or TCP)

TCP service: reliable transfer of *bytes* from one process to another



2: Application Layer 54

Socket programming *with TCP*

Client must contact server

- server process must first be running
- server must have created socket (door) that welcomes client's contact

Client contacts server by:

- creating client-local TCP socket
- specifying IP address, port number of server process
- When **client creates socket**: client TCP establishes connection to server TCP

- When contacted by client, **server TCP creates new socket** for server process to communicate with client
 - allows server to talk with multiple clients
 - source port numbers used to distinguish clients (*more in Chap 3*)

application viewpoint

TCP provides reliable, in-order transfer of bytes ("pipe") between client and server

2: Application Layer 55

Stream jargon

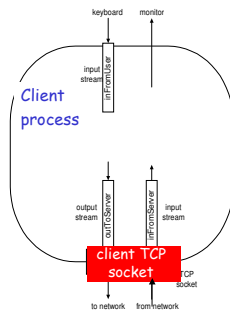
- A **stream** is a sequence of characters that flow into or out of a process.
- An **input stream** is attached to some input source for the process, eg, keyboard or socket.
- An **output stream** is attached to an output source, eg, monitor or socket.

2: Application Layer 56

Socket programming with TCP

Example client-server app:

- client reads line from standard input (`inFromUser` stream), sends to server via socket (`outToServer` stream)
- server reads line from socket
- server converts line to uppercase, sends back to client
- client reads, prints modified line from socket (`inFromServer` stream)

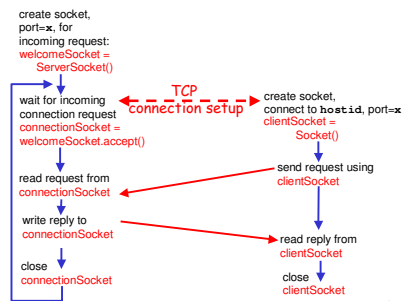


2: Application Layer 57

Client/server socket interaction: TCP

Server (running on `hostid`)

Client



2: Application Layer 58

Example: Java client (TCP)

```
import java.io.*;
import java.net.*;
class TCPClient {

    public static void main(String argv[]) throws Exception
    {
        String sentence;
        String modifiedSentence;

        Create input stream -> BufferedReader inFromUser =
            new BufferedReader(new InputStreamReader(System.in));

        Create client socket, connect to server -> Socket clientSocket = new Socket("hostname", 6789);

        Create output stream attached to socket -> DataOutputStream outToServer =
            new DataOutputStream(clientSocket.getOutputStream());
```

2: Application Layer 59

Example: Java client (TCP), cont.

```
        Create input stream attached to socket -> BufferedReader inFromServer =
            new BufferedReader(new InputStreamReader(clientSocket.getInputStream()));

        sentence = inFromUser.readLine();

        Send line to server -> outToServer.writeBytes(sentence + '\n');

        Read line from server -> modifiedSentence = inFromServer.readLine();
        System.out.println("FROM SERVER: " + modifiedSentence);

        clientSocket.close();
    }
}
```

2: Application Layer 60

Example: Java server (TCP)

```
import java.io.*;
import java.net.*;

class TCPServer {

    public static void main(String argv[]) throws Exception
    {
        String clientSentence;
        String capitalizedSentence;

        ServerSocket welcomeSocket = new ServerSocket(6789);

        while(true) {
            Socket connectionSocket = welcomeSocket.accept();

            BufferedReader inFromClient =
                new BufferedReader(new
                    InputStreamReader(connectionSocket.getInputStream()));

```

Annotations for the TCP server code:

- Create welcoming socket at port 6789**: points to `ServerSocket welcomeSocket = new ServerSocket(6789);`
- Wait, on welcoming socket for contact by client**: points to `Socket connectionSocket = welcomeSocket.accept();`
- Create input stream, attached to socket**: points to `BufferedReader inFromClient = new BufferedReader(new InputStreamReader(connectionSocket.getInputStream()));`

2: Application Layer 61

Example: Java server (TCP), cont

```
        DataOutputStream outToClient =
            new DataOutputStream(connectionSocket.getOutputStream());

        clientSentence = inFromClient.readLine();

        capitalizedSentence = clientSentence.toUpperCase() + '\n';

        outToClient.writeBytes(capitalizedSentence);
    }
}
```

Annotations for the TCP server continuation:

- Create output stream, attached to socket**: points to `DataOutputStream outToClient = new DataOutputStream(connectionSocket.getOutputStream());`
- Read in line from socket**: points to `clientSentence = inFromClient.readLine();`
- Write out line to socket**: points to `outToClient.writeBytes(capitalizedSentence);`
- End of while loop, loop back and wait for another client connection**: points to the closing brace of the `while` loop.

2: Application Layer 62

Socket programming *with UDP*

UDP: no "connection" between client and server

- no handshaking
- sender explicitly attaches IP address and port of destination to each packet
- server must extract IP address, port of sender from received packet

UDP: transmitted data may be received out of order, or lost

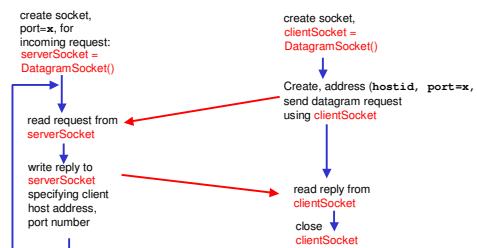
application viewpoint
UDP provides *unreliable* transfer of groups of bytes ("datagrams") between client and server

2: Application Layer 63

Client/server socket interaction: UDP

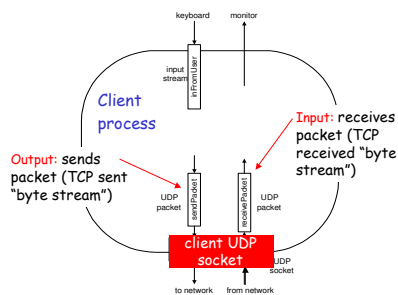
Server (running on hostid)

Client



2: Application Layer 64

Example: Java client (UDP)



2: Application Layer 65

Example: Java client (UDP)

```
import java.io.*;
import java.net.*;

class UDPClient {
    public static void main(String args[]) throws Exception
    {
        BufferedReader inFromUser =
            new BufferedReader(new InputStreamReader(System.in));

        DatagramSocket clientSocket = new DatagramSocket();

        InetAddress IPAddress = InetAddress.getByName("hostname");

        byte[] sendData = new byte[1024];
        byte[] receiveData = new byte[1024];

        String sentence = inFromUser.readLine();
        sendData = sentence.getBytes();
    }
}
```

Annotations for the UDP client code:

- Create input stream**: points to `BufferedReader inFromUser = new BufferedReader(new InputStreamReader(System.in));`
- Create client socket**: points to `DatagramSocket clientSocket = new DatagramSocket();`
- Translate hostname to IP address using DNS**: points to `InetAddress IPAddress = InetAddress.getByName("hostname");`

2: Application Layer 66

Example: Java client (UDP), cont.

```

Create datagram with data-to-send, length, IP addr, port → DatagramPacket sendPacket = new DatagramPacket(sendData, sendData.length, IPAddress, 9876);

Send datagram to server → clientSocket.send(sendPacket);

Read datagram from server → DatagramPacket receivePacket = new DatagramPacket(receiveData, receiveData.length);
                           clientSocket.receive(receivePacket);

String modifiedSentence = new String(receivePacket.getData());
System.out.println("FROM SERVER:" + modifiedSentence);
clientSocket.close();
}

```

2: Application Layer 67

Example: Java server (UDP)

```

import java.io.*;
import java.net.*;

class UDPServer {
    public static void main(String args[]) throws Exception {
        Create datagram socket at port 9876 → DatagramSocket serverSocket = new DatagramSocket(9876);

        byte[] receiveData = new byte[1024];
        byte[] sendData = new byte[1024];

        while(true) {
            Create space for received datagram → DatagramPacket receivePacket = new DatagramPacket(receiveData, receiveData.length);
            Receive datagram → serverSocket.receive(receivePacket);
        }
    }
}

```

2: Application Layer 68

Example: Java server (UDP), cont

```

String sentence = new String(receivePacket.getData());
Get IP addr, port #, of sender → InetAddress IPAddress = receivePacket.getAddress();
                                int port = receivePacket.getPort();

String capitalizedSentence = sentence.toUpperCase();

sendData = capitalizedSentence.getBytes();

Create datagram to send to client → DatagramPacket sendPacket = new DatagramPacket(sendData, sendData.length, IPAddress, port);

Write out datagram to socket → serverSocket.send(sendPacket);

End of while loop, loop back and wait for another datagram

```

2: Application Layer 69

Building a simple Web server

- handles one HTTP request
 - accepts the request
 - parses header
 - obtains requested file from server's file system
 - creates HTTP response message:
 - header lines + file
 - sends response to client
- after creating server, you can request file using a browser (eg IE explorer)
 - see text for details

2: Application Layer 70

Socket programming: references

C-language tutorial (audio/slides):

- "Unix Network Programming" (J. Kurose), <http://manic.cs.umass.edu/~amldemo/courseware/intro>.

Java-tutorials:

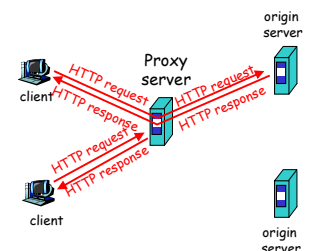
- "All About Sockets" (Sun tutorial), <http://www.javaworld.com/javaworld/jw-12-1996/jw-12-sockets.html>
- "Socket Programming in Java: a tutorial," <http://www.javaworld.com/javaworld/jw-12-1996/jw-12-sockets.html>

2: Application Layer 71

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



2: Application Layer 72

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

2: Application Layer 73

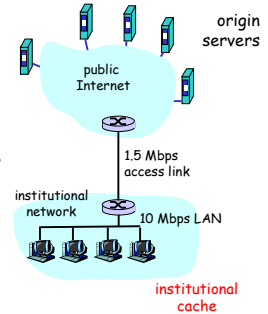
Caching example (1)

Assumptions

- average object size = 100,000 bits
- avg. request rate from institution's browser to origin server = 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



2: Application Layer 74

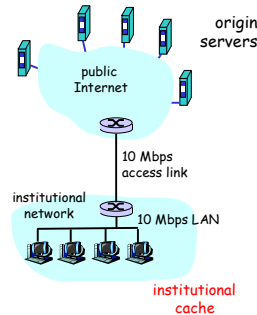
Caching example (2)

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 + msec + msec
- often a costly upgrade



2: Application Layer 75

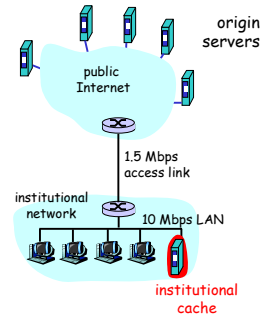
Caching example (3)

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 sec + .6 * .01 sec + milliseconds < 1.3 secs



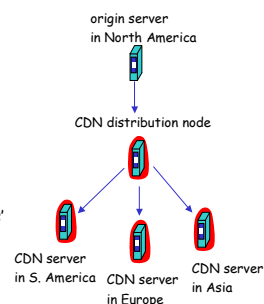
2: Application Layer 76

Content distribution networks (CDNs)

- The content providers are the CDN customers.

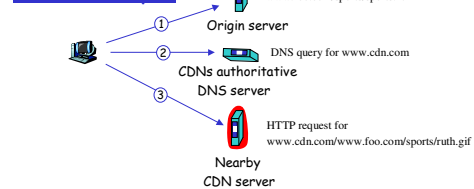
Content replication

- 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



2: Application Layer 77

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

2: Application Layer 78

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

2: Application Layer 79

P2P file sharing

Example

- Alice runs P2P client application on her notebook computer
- Intermittently connects to Internet; gets new IP address for each connection
- Asks for "Hey Jude"
- Application displays other peers that have copy of Hey Jude.

- Alice chooses one of the peers, Bob.
- File is copied from Bob's PC to Alice's notebook: HTTP
- While Alice downloads, other users uploading from Alice.
- Alice's peer is both a Web client and a transient Web server.

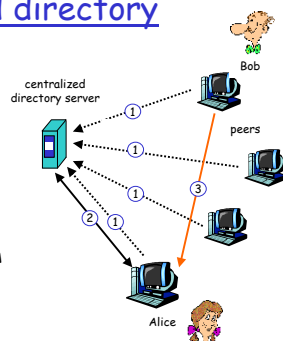
All peers are servers = highly scalable!

2: Application Layer 80

P2P: centralized directory

original "Napster" design

- 1) when peer connects, it informs central server:
 - IP address
 - content
- 2) Alice queries for "Hey Jude"
- 3) Alice requests file from Bob



2: Application Layer 81

P2P: problems with centralized directory

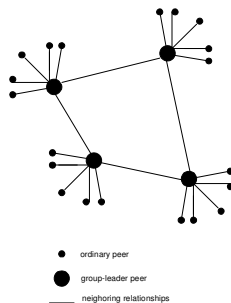
- Single point of failure
- Performance bottleneck
- Copyright infringement

file transfer is decentralized, but locating content is highly decentralized

2: Application Layer 82

P2P: decentralized directory

- Each peer is either a group leader or assigned to a group leader.
- Group leader tracks the content in all its children.
- Peer queries group leader; group leader may query other group leaders.



2: Application Layer 83

More about decentralized directory

overlay network

- peers are nodes
- edges between peers and their group leaders
- edges between some pairs of group leaders
- virtual neighbors
- bootstrap node
 - connecting peer is either assigned to a group leader or designated as leader

advantages of approach

- no centralized directory server
 - location service distributed over peers
 - more difficult to shut down

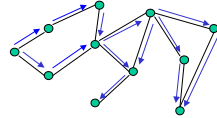
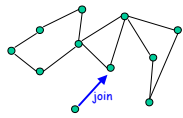
disadvantages of approach

- bootstrap node needed
- group leaders can get overloaded

2: Application Layer 84

P2P: Query flooding

- Gnutella
- no hierarchy
- use bootstrap node to learn about others
- join message
- Send query to neighbors
- Neighbors forward query
- If queried peer has object, it sends message back to querying peer



2: Application Layer 85

P2P: more on query flooding

Pros

- peers have similar responsibilities: no group leaders
- highly decentralized
- no peer maintains directory info

Cons

- excessive query traffic
- query radius: may not have content when present
- bootstrap node
- maintenance of overlay network

2: Application Layer 86