

College of Information Studies

University of Maryland Hornbake Library Building College Park, MD 20742-4345

Email

Session 5 INST 346 Technologies, Infrastructure and Architecture

Muddiest Points

Format of the HTTP messages
 What GET, HEAD, POST actually do

• Who creates proxy servers?

• How to create a Web server

Goals for Today

- Finish Email
 - Review SMTP
 - POP3 and IMAP

• Learn socket programming

• Getahead: DNS (maybe!)

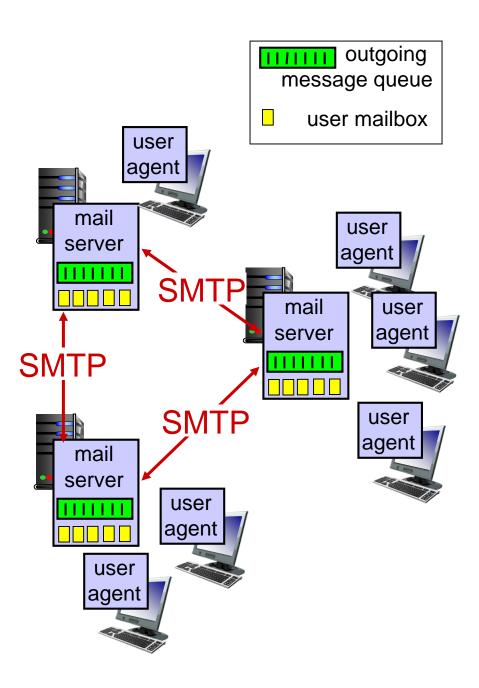
Email

Three major components:

- user agents ("mail reader")
- mail servers
- simple mail transfer protocol: SMTP

User Agent

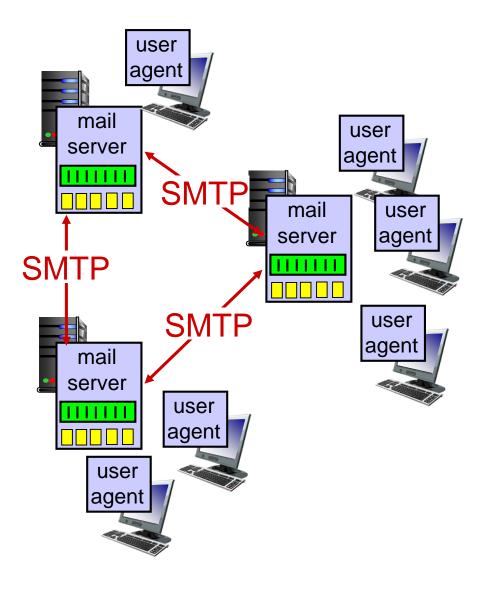
- composing, editing, reading email messages
- e.g., Outlook, Thunderbird, iPhone mail client
- outgoing, incoming messages stored on server



Email: 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

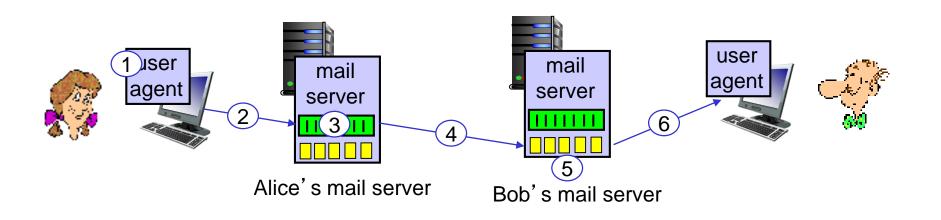


- uses TCP to reliably transfer email message from client to server, port 25
- direct transfer: sending server to receiving server
- three phases of transfer
 - handshaking (greeting)
 - transfer messages
 - close
- command/response interaction (like HTTP)
 - commands: ASCII text
 - response: status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

Scenario: Alice sends message to Bob

- I) Alice uses UA to compose message "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

Mail server (client) at crepes.fr has mail to send Client initiates connection to hamburger.edu port 25

- 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: 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 message
- SMTP: multiple objects sent in multipart message

Mail message format



- Body: the "message"
 - ASCII characters only

Mail access protocols



- SMTP: delivery/storage to receiver's mail server
- mail access protocol: upload to and download from a mail server
 - POP: Post Office Protocol [RFC 1939]: authorization, download
 - IMAP: Internet Mail Access Protocol [RFC 1730]: more features, including manipulation of stored messages in folders on the mail server

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 ussrid S: +OK C: pass hungry password 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

Comparing POP3 and IMAP

more about POP3

- previous example uses POP3 "download and delete" mode
 - Bob cannot re-read email if he changes client
- POP3 "download-andkeep": 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

Socket programming

goal: learn how to build client/server applications that communicate using sockets

socket: outbox/inbox between application process and end-end-transport protocol



Socket programming


Two socket types for two transport services:

- UDP: unreliable datagram
- TCP: reliable, byte stream-oriented

Application Example:

- client reads a line of characters (data) from its keyboard and sends data to server
- 2. server receives the data and converts characters to uppercase
- 3. server sends modified data to client
- 4. client receives modified data and displays line on its screen

Client/server socket interaction: UDP



Example app: UDP client

Python UDPClient

include Python's socket from socket import * library serverName = 'localhost' serverPort = 12000create UDP socket for _____ clientSocket = socket(AF_INET, server SOCK_DGRAM) get user keyboard input _____ message = input('Input lowercase sentence:') Attach server name, port to clientSocket.sendto(message.encode(), message; send into socket (serverName, serverPort)) read reply characters from — modifiedMessage, serverAddress = socket into string clientSocket.recvfrom(2048) print out received string ----> print(modifiedMessage.decode()) and close socket clientSocket.close()

Example app: UDP server

Python UDPServer

from socket import *

serverPort = 12000

serverSocket = socket(AF_INET, SOCK_DGRAM) create UDP socket

bind socket to local port serverSocket.bind((", serverPort))

print ("The server is ready to receive")

loop forever

number 12000

while True:

Read from UDP socket into message, getting client's address (client IP and port)

send upper case string back to this client

message, clientAddress = serverSocket.recvfrom(2048) modifiedMessage = message.decode().upper() serverSocket.sendto(modifiedMessage.encode(),

clientAddress)

Running Python

• Install the latest Python 3 from:

– https://www.python.org/downloads/

• Download the programs

- Materials used in class link from schedule

- Open two shell windows
 - On a PC, type "cmd" in the search box
 - On a Mac, open a terminal
- In one shell, type:
 - python udpserver.py
- In the other, type:
 - python udpclient.py

Socket programming with TCP

client must contact server

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

client contacts server by:

- Creating 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 that particular 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 byte-stream transfer ("pipe") between client and server

Client/server socket interaction: TCP



client



Example app:TCP client

Python TCPClient

from socket import * serverName = 'localhost' serverPort = 12000create TCP socket for server, remote port 12000 clientSocket = socket(AF_INETCSOCK_STREAM) clientSocket.connect((serverName,serverPort)) sentence = input('Input lowercase sentence:') No need to attach server clientSocket.send(sentence.encode()) name, port modifiedSentence = clientSocket.recv(1024)print ('From Server:', modifiedSentence.decode()) clientSocket.close()

Example app:TCP server

Python TCPServer

from socket import *

create TCP welcoming socket

server begins listening for incoming TCP requests

loop forever server waits on accept() for incoming requests, new socket created on return

read bytes from socket (but not address as in UDP)

close connection to this client (but not welcoming socket)

serverPort = 12000serverSocket = socket(AF_INET,SOCK_STREAM) serverSocket.bind(('',serverPort)) serverSocket.listen(1) print('The server is ready to receive') while True: connectionSocket, addr = serverSocket.accept() sentence = connectionSocket.recv(1024).decode()

capitalizedSentence = sentence.upper() connectionSocket.send(capitalizedSentence.

encode())

connectionSocket.close()

Getahead: DNS

DNS: domain name system

people: many identifiers:

• SSN, name, passport #

Internet hosts, routers:

- IP address (32 bit) used for addressing datagrams
- "name", e.g.,
 www.yahoo.com used by humans
- Q: how to map between IP address and name, and vice versa ?

Domain Name System:

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

DNS: services, structure

DNS services

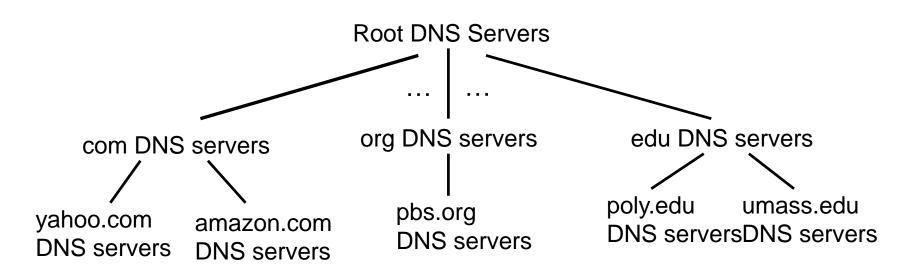
- hostname to IP address translation
- host aliasing
 - canonical, alias names
- mail server aliasing
- Ioad distribution
 - replicated Web servers: many IP addresses correspond to one name

why not centralize DNS?

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

A: doesn't scale!

DNS: a distributed, hierarchical database

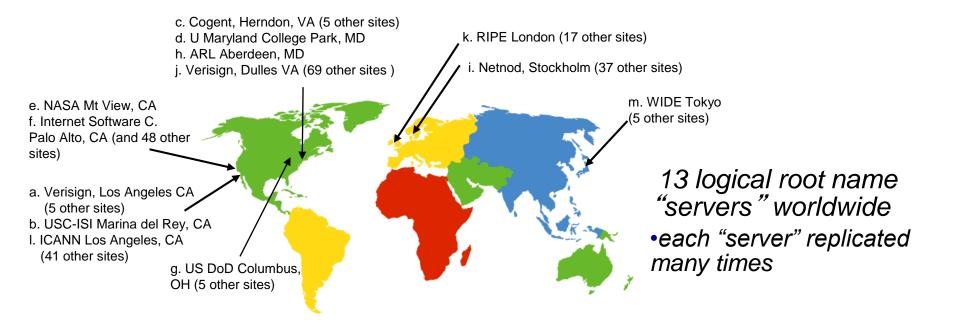


client wants IP for www.amazon.com; 1st approximation:

- client queries root server to find com DNS server
- client queries .com DNS server to get amazon.com DNS server
- client queries amazon.com DNS server to get IP address for www.amazon.com

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



TLD, authoritative servers

top-level domain (TLD) servers:

- responsible for com, org, net, edu, aero, jobs, museums, and 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 DNS name server

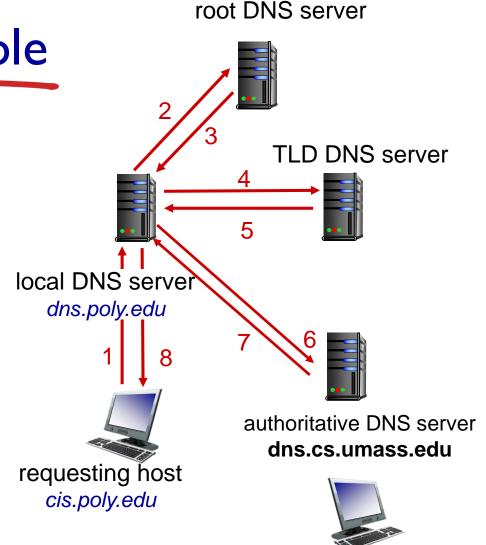
- does not strictly belong to hierarchy
- each ISP (residential ISP, company, university) has one
 - also called "default name server"
- when host makes DNS query, query is sent to its local DNS server
 - has local cache of recent name-to-address translation pairs (but may be out of date!)
 - acts as proxy, forwards query into hierarchy

DNS name resolution example

 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"

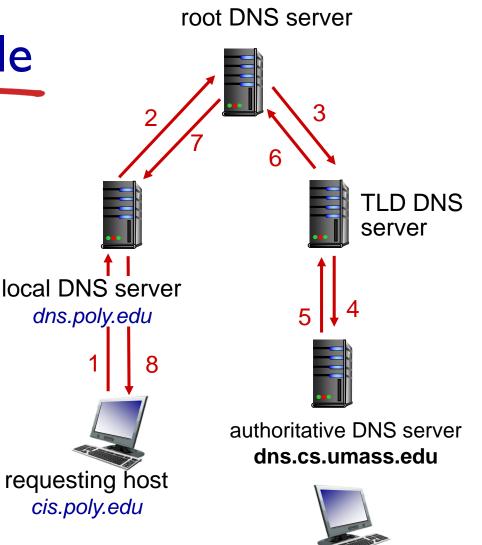


gaia.cs.umass.edu

DNS name resolution example

recursive query:

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



gaia.cs.umass.edu

DNS: caching, updating records

- once (any) name server learns mapping, it caches mapping
 - cache entries timeout (disappear) 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: distributed database storing resource records (RR)

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



- 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
- www.ibm.com is really servereast.backup2.ibm.com
- value is canonical name

<u>type=MX</u>

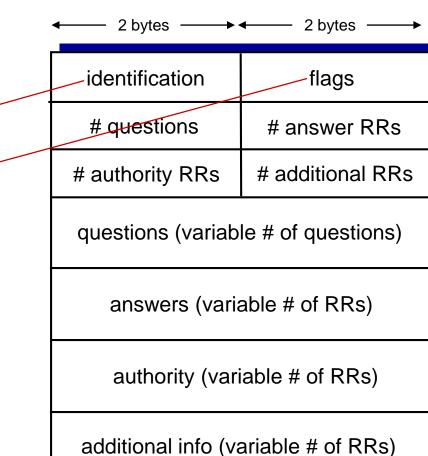
 value is name of mailserver associated with name

DNS protocol, messages


query and reply messages, both with same message format

message header

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



DNS protocol, messages



Inserting records into DNS

- example: new startup "Network Utopia"
- register name networkuptopia.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)
- create authoritative server type A record for www.networkuptopia.com; type MX record for networkutopia.com

Before You Go

On a sheet of paper, answer the following (ungraded) question (no names, please):

What was the muddlest point in today's class?