TCP / IP

What is a Network?

- Computer network
  - a set of computers using common protocols to communicate over connecting transmission media.
- Protocol
  - a formal description of message formats and the rules two or more machines follow to exchange messages.

Protocols

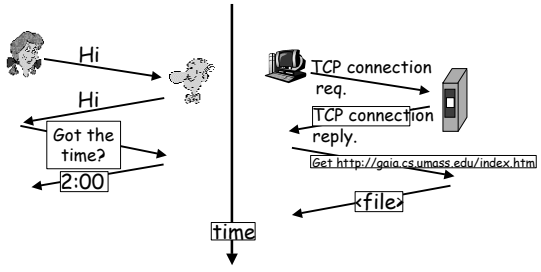

The ISO OSI Model

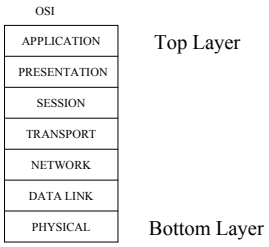

The TCP/IP Protocol Suite

- TCP/IP is a set of protocols that were created specifically to allow development of network and internetwork communications on a global scale
- TCP/IP is the most commonly used protocols within the internet.
- TCP/IP is normally considered to be a four-layer system.

The TCP/IP Protocol Suite

Application	Telnet, FTP, e-mail, etc.
Transport	TCP, UDP
Network	IP
Link	device driver and interface card

The TCP/IP Reference Model

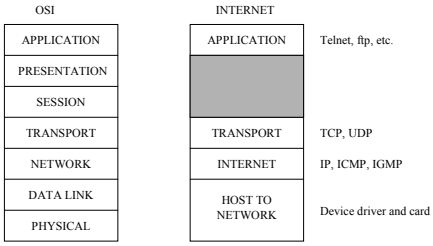

Standardization

- Who controls the TCP/IP protoocol suite?
  - The *Internet Society* (ISOC) a professional society to facilitate, support, and promote the evolution and growth of the Internet
  - The *Internet Architecture Board* (IAB) the technical oversight and coordination body
  - The *Internet Engineering Task Force* (IETF) is the near-term, standards-oriented group.
  - The *Internet Research Task Force* (IRTF) pursues long-term research projects.

RFCs

- All official standards in the internet community are published as a *Request for Comments*, or *RFC*.
- All RFCs are available at no charge through electronic mail, FTP, or the Web.
- A nice place to get RFCs is at
  - http://www.rfc-editor.org/

IP: Internet Protocol

- IP is the workhorse protocol of the TCP/IP protocol suite
- IP provides an unreliable, connectionless, datagram delivery service
- RFC791 is the official specification of IP
- Sits in the Internet Layer of the TCP/IP Model

Addressing

- A distinction is made between names, addresses, and routes
  - A name indicates what we seek
  - An address indicates where it is
  - A route indicates how to get there
- IP deals primarily with addresses. It is the task of higher level protocols to make the mapping from names to addresses.

IP Addresses

- Every host on the internet must have a unique *Internet Address* (an IP address)
- IP addresses are 32-bit numbers and are divided into two components: the host address and the network address
  - The number of bits assigned to the host and network varies depending on the class of the address

Dotted Decimal Notation

- IP addresses are normally written as four numbers (octets), one for each byte of the address.
  - 129.21.38.169

Dotted Decimal Notation

- Two sections
  - Net
    - identifies the network to which a computer belongs
    - Will always contain the first octet
  - Host
    - Identifies an individual machine
    - Will always contain the last octet

IP address classes

- Class A
  - For very large networks
  - 1st octet identifies net (1-126) / Octet 2-4 identifies host
  - Can support 16million+ (2^{24}-2) hosts
- Class B
  - For medium sized networks (like college campuses)
  - Octet 1-2 identifies net / Octet 3-4 identifies host
    - Octet 1 = 128 – 191
  - Can support 65K hosts (2^{16}-2)

IP address classes

- Class C
  - For small businesses
  - Octet 1-3 identifies net / Octet 4 identifies host
    - 1st octet = 192 – 223
  - Can support 254 hosts
- Class D / Class E
  - Multicast

IP Address Classes

- The easiest way to differentiate between the classes is to look at the first number

Class	*Range*
A	0.0.0.0 to 127.255.255.255
B	128.0.0.0 to 191.255.255.255
C	192.0.0.0 to 223.255.255.255
D	224.0.0.0 to 239.255.255.255
E	240.0.0.0 to 247.255.255.255

Special IP addresses

- 0.0.0.0
  - The default network for any machine
- 127.0.0.1
  - Loopback address – to send messages to yourself
- 255.255.255.255
  - Broadcast – send messages to all on a network.

Assigning IP Addresses

- Since every interface must have a unique IP address, there must be a central authority for assigning numbers
- That authority is the *Internet Network Information Center*, called the InterNIC.
- The InterNIC assigns only network ids, the assignment of host ids is up to the system administrator

IP Addresses

- Note that each piece of network hardware also has a unique "Ethernet (MAC) Address"
  - For IP over ethernet, conversion from IP to MAC address must be made.
  - Conversion table usually kept in cache
  - Address Resolution Protocol (ARP) used to query hardware for addresses
    - NOT a TCP/IP Protocol!

IP

- Routing and delivery of "packets"
  - Only responsible for sending "packets" from one point to another
  - Doesn't care what's in the packet or if packets are part of a larger message.

Other Internet Protocols

- ICMP
  - Internet Control Message Protocol
  - Supports packets containing error, control, and informational methods.
  - Defined on top of IP
  - E.g. used by ping.
- IGMP
  - Internet Group Management Protocol
  - Standard for IP multicasting over the Internet
  - Defined on top of IP

Transport Layer

- The transport layer is responsible for
  - Disassembling and assembling streams of data
  - Addressing packets (and send to IP)
  - Error checking

Types of Transfer

- Networks typically provide two types of transfer
  - Connection-oriented
    - often reliable
    - stream based
    - Point to point – like phone call
  - Connectionless
    - often unreliable
    - datagram based
    - Sends independent packets of data – like Postal Mail
      - Order of delivery is not important
      - Delivery not guaranteed,

Connection-oriented Transfer

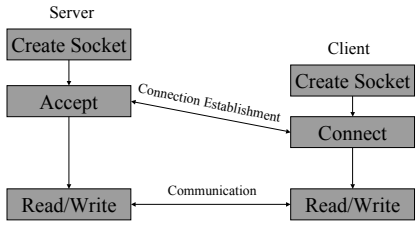

Connectionless Transfer

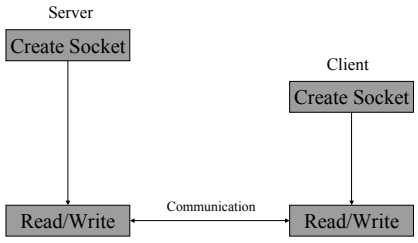

Transport Layer

- TCP
  - Transmission Control Protocol
  - Connection based protocol that provides a reliable flow of data
- UDP
  - User Datagram Protocol
  - Sends independent packets of data with no guarantee of arrival

The TCP/IP Protocol Suite

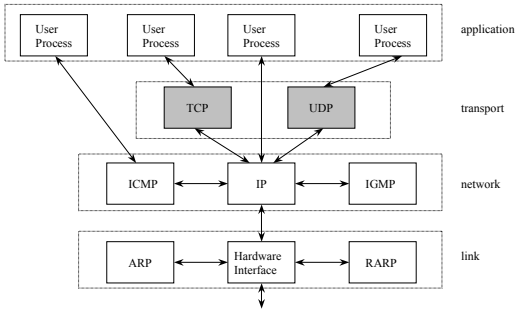

Transmission Control Protocol

- TCP provides a connection-oriented, reliable, byte stream service (RFC793)
- TCP is an independent, general purpose protocol that can be adapted for use with delivery systems other than IP.

TCP Streams

- A stream of 8-bit bytes is exchanged across a TCP connection.
- The treatment of the byte stream by TCP is similar to the treatment of a file by the UNIX operating system.
- Connections provided by TCP allow concurrent transfer in both directions. Such connections are called *full duplex*.

TCP Ports

- TCP uses protocol port numbers to identify the ultimate destination within a machine.
- How does one determine the port to communicate with?
  - Well-known Ports
  - Randomly Assigned Ports
- Questions

User Datagram Protocol

- UDP is a simple, unreliable, datagram-oriented, transport layer protocol (RFC768).

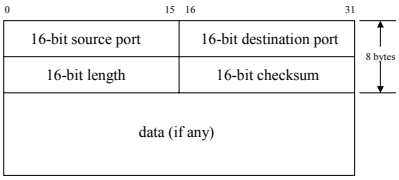

User Datagram Protocol

- Designed for applications where streams of data need not be disassembled or assembled
  - Messages that fit in a single packet
- Will not keep track of what is send nor resend if not received

UNIX Networking

- In the early 80s UNIX was being developed by several organizations
- One of the most influential development groups was UCB
  - 4BSD provided support for the DARPA Internet networking protocols, TCP/IP
- Some consider 4BSD responsible for the popularity of the TCP/IP protocols

Application Layer: Sockets

- Berkeley sockets are one of the most widely used communication APIs
- A socket is an object from which messages are sent and received

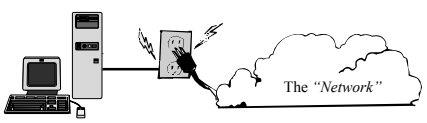

Socket

- A socket represents a connection (line of communication) between 2 processes:
  - Usually on different machines, but can be on the same machine
  - Communication is usually bi-directional
    - Based on some protocol
    - Different Sockets will implement different protocols
  - Analogous to `ChatSession` object in project.

Sockets

- But more on this next week when we talk about networking in Java.
- Questions?

Summary

- TCP/IP
  - IP – Internet Layer
  - TCP / UDP – Transport Layer
  - Sockets – Application Layer
- Thursday: Exam 2 Review