AP/ITEC 2210 3.0 A: System Administration Fall 2023

Instructor: Jamon Camisso
ITEC 2210 Chat: mattermost.itec2210.ca
Email: jamon@yorku.ca
Website: https://eclass.yorku.ca/

Date/Time: Wednesday, 19:00-22:00
Location: Zoom / ACE 003
Office hours: Via Mattermost any time

Class 4 - TCP/IP - Internet

- Midterm:
  - October 18, online during our 7-10pm time slot
  - We'll be using Zoom to collaborate on answers
  - Hopefully a change from the usual midterm tedium

Class 4 - TCP/IP - Internet

- Lab 1/Assignment 1:
  - Clock is ticking. October 11 @ 23:59 due date.

Class 4 - TCP/IP - Internet

- Readings PSNA CH23 Network Architecture
  - p399-401 OSI model
  - p404 VLAN Myths
  - p408-422 Sections 23.5-23.8 inclusive
  - P425-430 Sections 23.10-23.12 inclusive
  - This Class will span into next week, a lot to go over

Class 4 - TCP/IP - Internet

- So what are we talking about?
- OSI Model
- Open Systems Interconnection
- Prescriptive framework

Table 23.1: The OSI Network Model

Layer	Name	Description
1	Physical	The physical connection between devices: copper, fiber, radio, laser
2	Data link	Interface (or MAC) addressing, flow control, low-level error notification
3	Network	Logical addressing (e.g., IP addresses) and routing (e.g., RIP, OSPF, IGRP)
4	Transport	Data transport, error checking and recovery, virtual circuits (e.g., TCP sessions)
5	Session	Communication-session management (e.g., AppleTalk name binding, or PPTP)
6	Presentation	Data formats, character encoding, compression, encryption (e.g., ASCII, Unicode, HTML, MP3, MPEG)
7	Application	Application protocols (e.g., SMTP for email, HTTP for web, and FTP for file transfer)

Class 4 - TCP/IP - Internet

- TCP/IP model
- Compare to OSI
- TCP/IP is a bit more descriptive of how things work, more fuzzy
- Both models are widely used

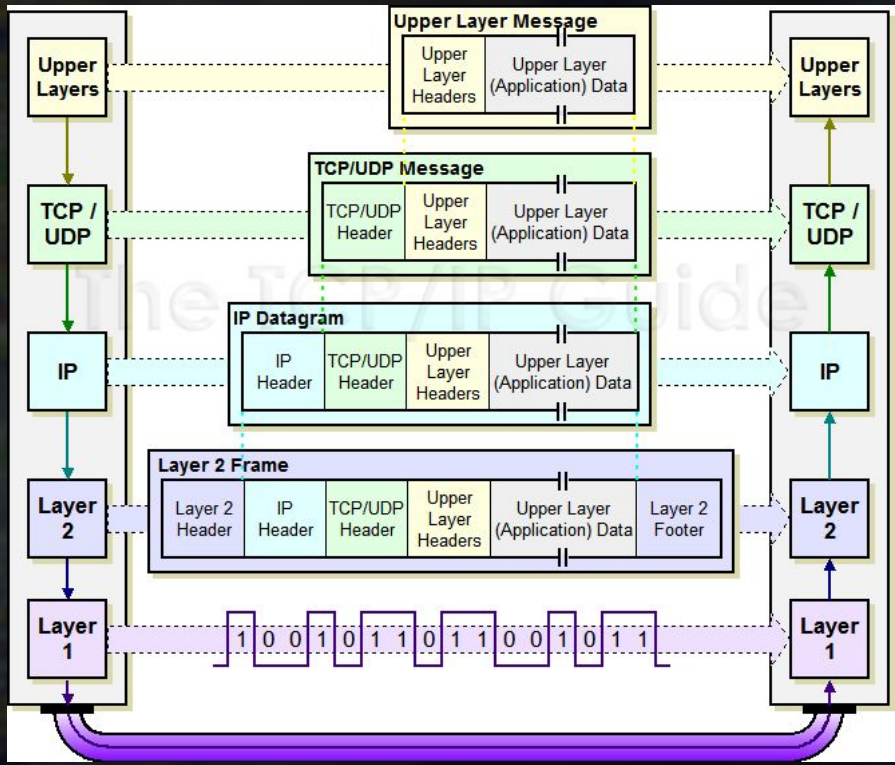

tcpipguide.com/free/t_IPDatagramEncapsulation.htm

Class 4 - TCP/IP - Internet

- As a sysadmin you'll work in depth with both models
  - Network troubleshooting, TCP/IP
  - Web, mail, print etc. troubleshooting, OSI
  - Tools to troubleshoot lower layers (usually) let you examine upper layer data as well
  - Best tools: Wireshark (GUI) tcpdump (CLI)

Class 4 - TCP/IP - Internet

- Acronyms
  - RFC
  - TCP
  - IP
  - MTU
  - UDP
  - CIDR
  - BGP
  - AS/ASN
  - OSI

Class 4 - TCP/IP - Internet

- **RFC - Request for Comments**
  - “Memos in the **Requests for Comments (RFC)** document series contain technical and organizational notes about the Internet. They cover many aspects of computer networking, including protocols, procedures, programs, and concepts, as well as meeting notes, opinions, and sometimes humor”
  - IPoAC - https://tools.ietf.org/html/rfc1149
  - HTCPCP - https://tools.ietf.org/html/rfc2324

Class 4 - TCP/IP - Internet

https://gizmodo.com/this-is-the-room-where-the-internet-was-born-1527205592

Class 4 - TCP/IP - Internet

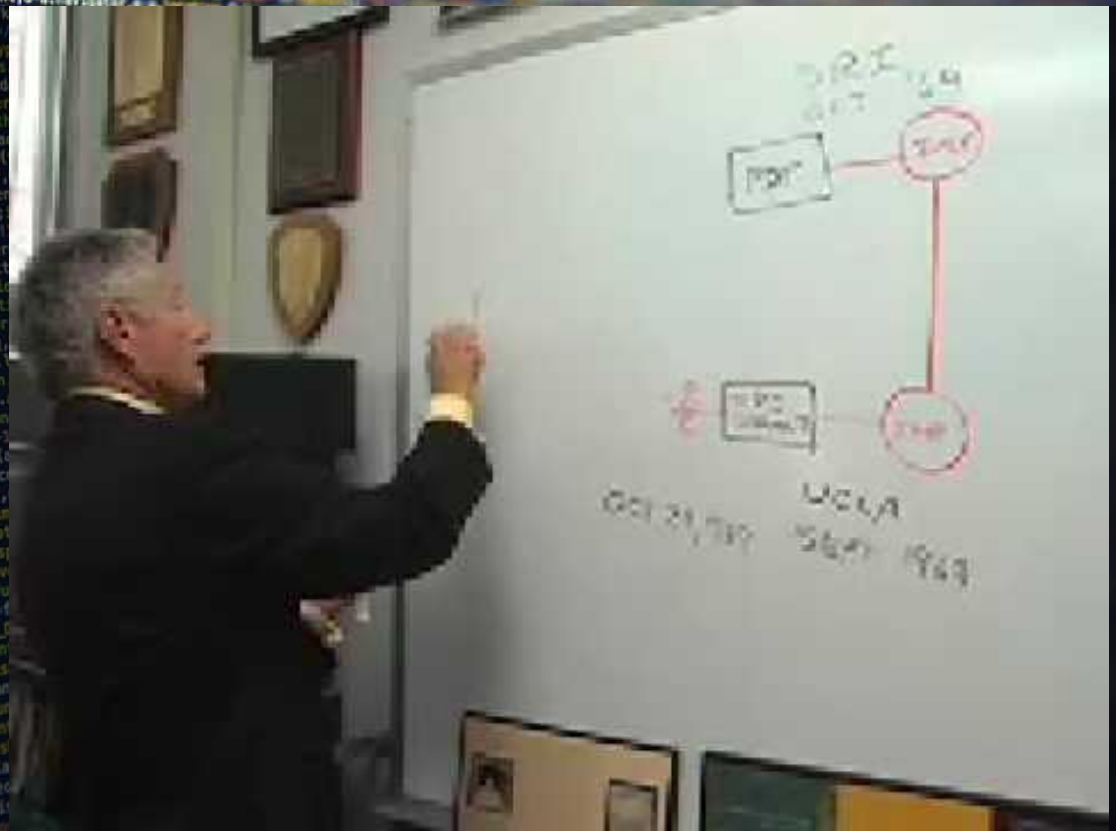

Class 4 - TCP/IP - Internet

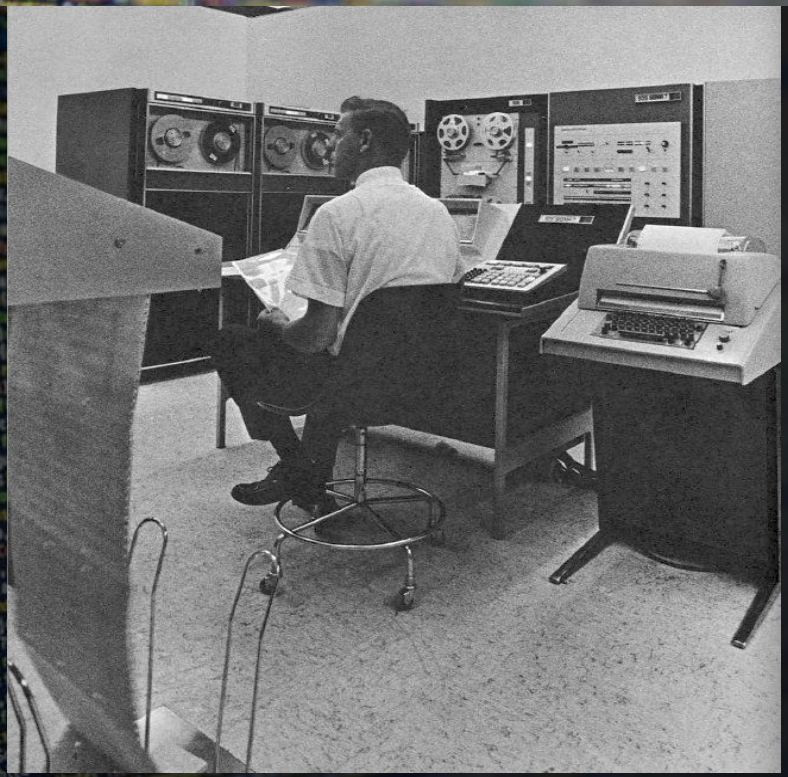

http://www.feb-patrimoine.com/projet/10070/cae_10070.htm

Class 4 - TCP/IP - Internet

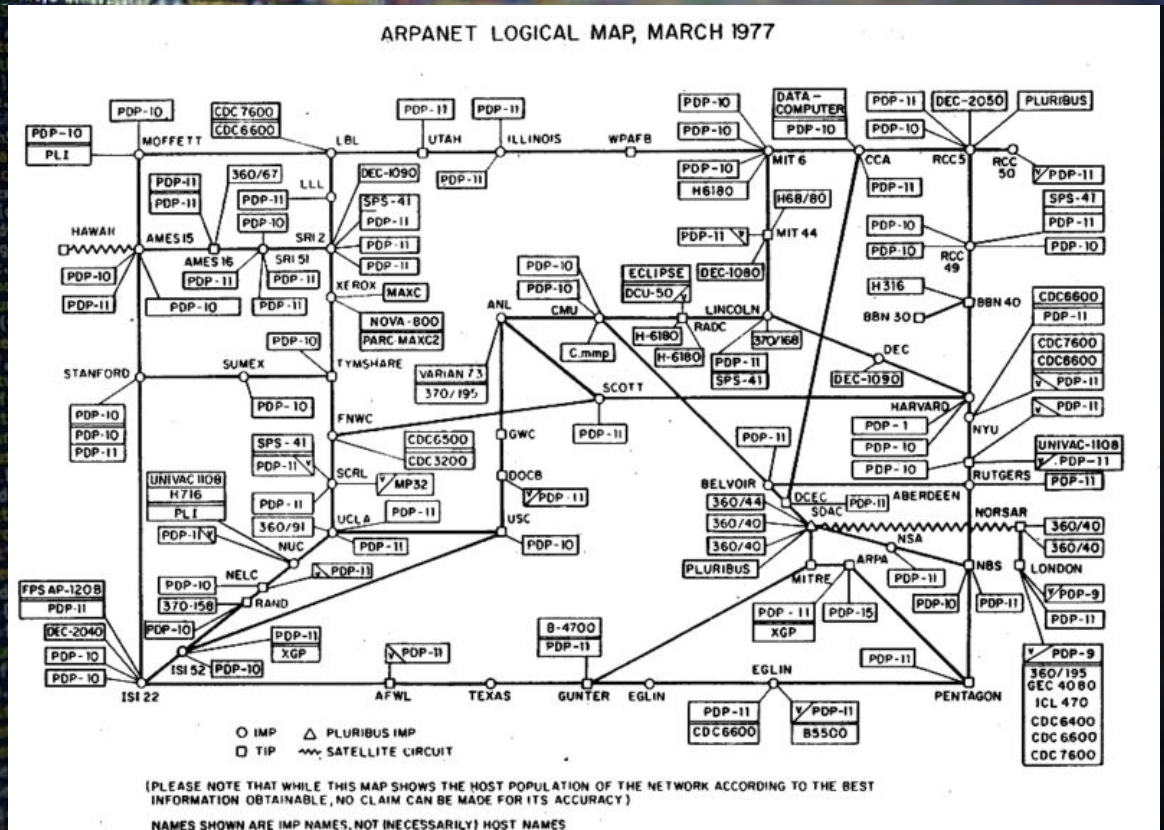

Class 4 - TCP/IP - Internet

- **ARPAnet** continued growing and growing with more IMPs and hosts, in public and private institutions
- Network participants gradually built on previous agreed upon and best practices with new RFCs until..
- Flag day: Jan 1, 1983
- ARPAnet changed over every system from NCP to TCP/IP
- Internet grew up

Class 4 - TCP/IP - Internet

- **IP - RFC 791 - Internet Protocol - Sept 1981**
- The internet protocol implements two basic functions: **addressing** and **fragmentation**
- **Addressing**:
  - “The internet modules use the addresses carried in the internet **header** to transmit internet datagrams toward their destinations. The selection of a path for transmission is called **routing**”

Class 4 - TCP/IP - Internet

- **IP - RFC 791 - Internet Protocol - Sept 1981**
- **Addressing (IPv4)**:
  - Addresses are fixed length of four octets (32 bits)
  - IPv4 range is from 0.0.0.0 - 255.255.255.255
  - Maximum number of IPv4 IP addresses is:
  - 2^32 = 4,294,967,296 (unsigned long, 4 bytes)

Class 4 - TCP/IP - Internet

- **IP - RFC 791 - Internet Protocol - Sept 1981**
- **Addressing (IPv4)**:
  - So why octets (or bytes) of 256 values you ask?
  - 4 bytes, each byte represents a 0-255 value
  - 2^8 = 11111111 in binary bits = 256. 8 bits = 1 byte
  - 0 is a value, so 0-255 range = 256 possible values

Class 4 - TCP/IP - Internet

- **IP - RFC 791 - Internet Protocol - Sept 1981**
- **Addressing (IPv4)**:
  - Byte 1: 0, Byte 2: 0.0, Byte 3: 0.0.0, Byte 4: 0.0.0.0
  - 0.0.0.0 - 255.255.255.255
  - . demarcates boundary between bytes for humans

Class 4 - TCP/IP - Internet

- **IP - RFC 791 - Internet Protocol - Sept 1981**
- **Routing:**
  - IP gateway hosts (IMPs in ARPAnet) are special in that they maintain lists of known routes to other hosts.
  - Keep in mind IP routing occurs at Layer 3, so packets are not changed (except to increment counters, TTL)
  - IP routing is now done with **CIDR (RFC 1519)**

Class 4 - TCP/IP - Internet

- **IP - RFC 791 - Internet Protocol - Sept 1981**
- **Routing:**
  - Main idea with routing and the Internet in general, is datagrams are passed from host to host, the goal being to get a packet closer to the destination
  - Optimal route selection is a separate issue
  - Protocols for sharing/discovery: **EIGRP**, **OSPF**, **BGP**

Class 4 - TCP/IP - Internet

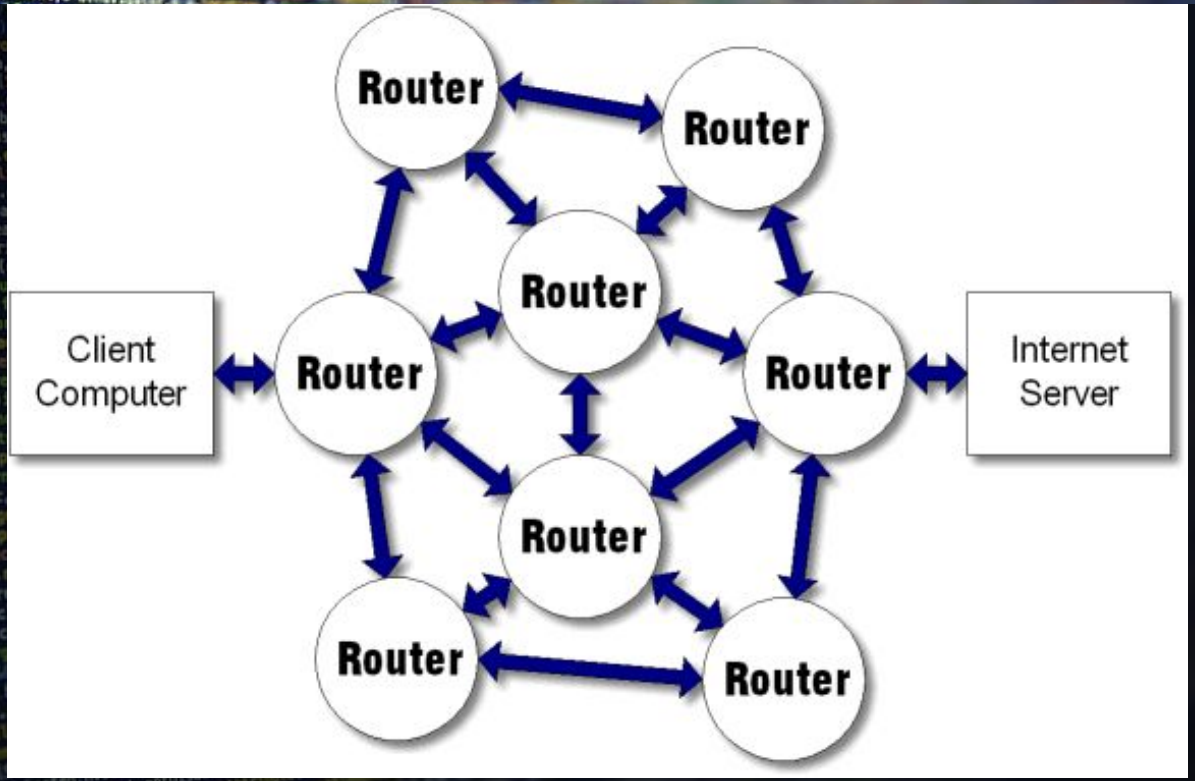

http://dtucker.cs.edinboro.edu/IPRouting.htm

Class 3 - TCP/IP - Internet

- **IP - RFC 791 - Internet Protocol - Sept 1981**
- **Routing**:
  - Book is right to gloss over routing - there are full courses and certifications that cover it in depth
  - Takeaway for us: **Autonomous Systems (ASes)** keep a globally up to date table of routes - currently 943749 routes, shared by 74267 ASes. Kept up to date in real time via BGP (external interfaces), OSPF & EIGRP (internal network interfaces). Other protocols too.

Class 4 - TCP/IP - Internet

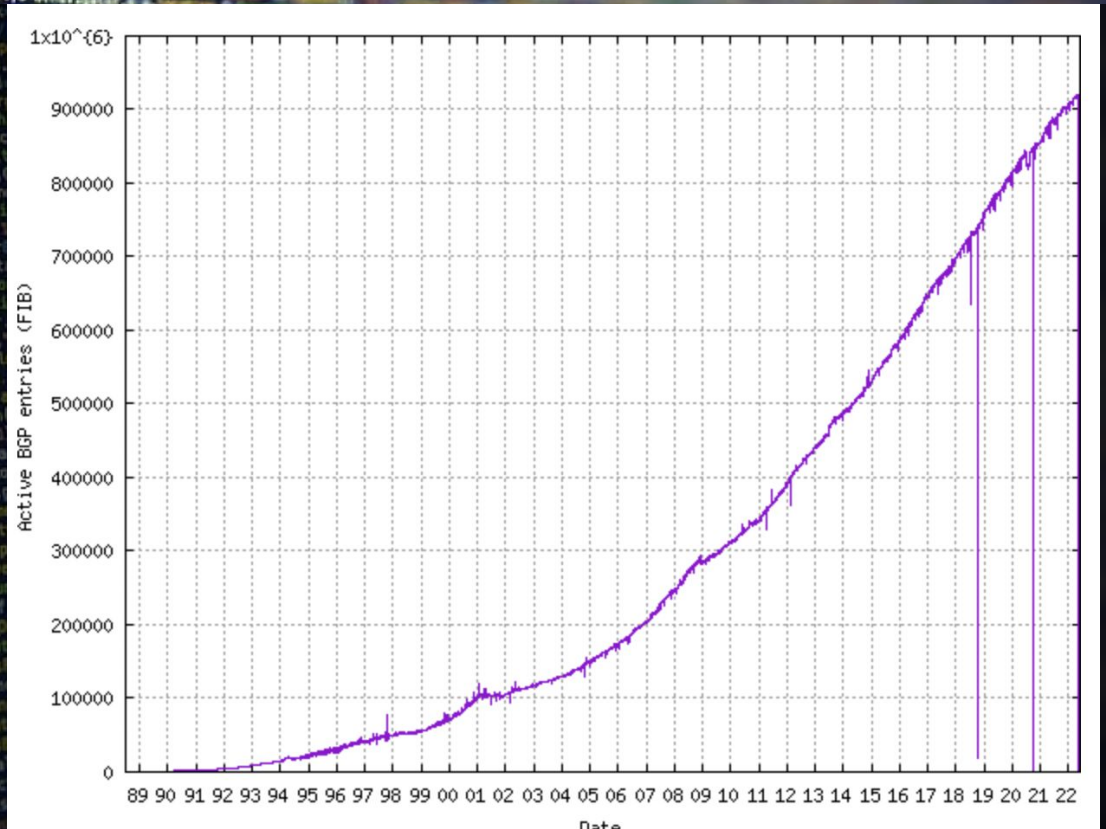

Class 4 - TCP/IP - Internet

- IP - **RFC 791 - Internet Protocol - Sept 1981**
- **Fragmentation**:
  - “The internet modules use fields in the internet header to **fragment** and **reassemble** internet datagrams when necessary for transmission through "small packet" networks”
  - How does a computer know when to fragment?

Class 4 - TCP/IP - Internet

- IP - **RFC 791 - Internet Protocol**
- **MTU - Maximum Transmission Unit**:
  - “If the total length is less than or equal the maximum transmission unit then submit this datagram to the next step in datagram processing; otherwise cut the datagram into two fragments, the first fragment being the maximum size, and the second fragment being the rest of the datagram”

Class 4 - TCP/IP - Internet

- IP - **RFC 791 - Internet Protocol**
- **MTU - Maximum Transmission Unit**:
  - MTU is related to, but not the same as maximum size of an ethernet frame
  - RFC 894 explains how to encapsulate IP datagrams in ethernet frames
  - Largest normal ethernet MTU is 1500

Class 4 - TCP/IP - Internet

- IP - **RFC 791 - Internet Protocol**
- **MTU - Maximum Transmission Unit**:
  - Tradeoff with MTU between speed and efficiency
  - Smaller sized packets are faster to process, but less efficient because it takes more packets to send the same amount of data as larger packets. However:
  - Larger packets tie up network resources to process

Class 4 - TCP/IP - Internet

- IP - **RFC 791 - Internet Protocol**
- Maximum datagram Size is 65,535 octets
- That's 2^16 octets, or bytes, e.g. 64kb
- Also an unsigned short int. **2 bytes** represent entire datagram length

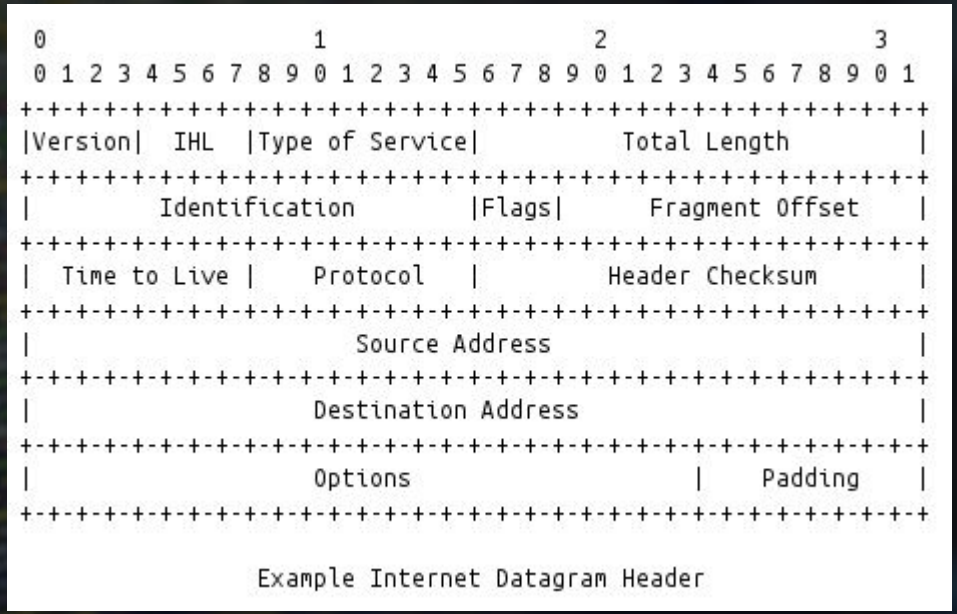

Example Internet Datagram Header

Class 4 - TCP/IP - Internet

- IP - **RFC 791 - Internet Protocol**
- Lots more to it
  - ICMP - Internet Control Message Protocol (RFC 792)
    - Part of IP protocol, but uses IP for datagrams
    - Ping uses ICMP
  - Path MTU (RFC 1191) - find the smallest MTU in a route/path
  - IGMP, anycast, multicast routing
  - Everything above, do it again for IPv6! Bigger this time!

Class 4/5 - TCP/IP - Internet

- Tools to put it all together:
  - traceroute -A google.com
    - Note different tries result in different routes
  - mtr -z google.com
    - Better traceroute (mostly)
  - whois
    - Combine with traceroute/mtr output: whois AS3

Class 4 - TCP/IP - Internet

- TCP - **RFC 793 - Transmission Control Protocol**
- Quoting: “A connection progresses through a series of states during its lifetime. The states are:
- LISTEN, SYN-SENT, SYN-RECEIVED, ESTABLISHED, FIN-WAIT-1, FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK, TIME-WAIT, and the fictional state CLOSED.
- CLOSED is fictional because it represents the state when there is no TCB, and therefore, no connection.”

Class 4 - TCP/IP - Internet

- TCP - **RFC 793 - Transmission Control Protocol**

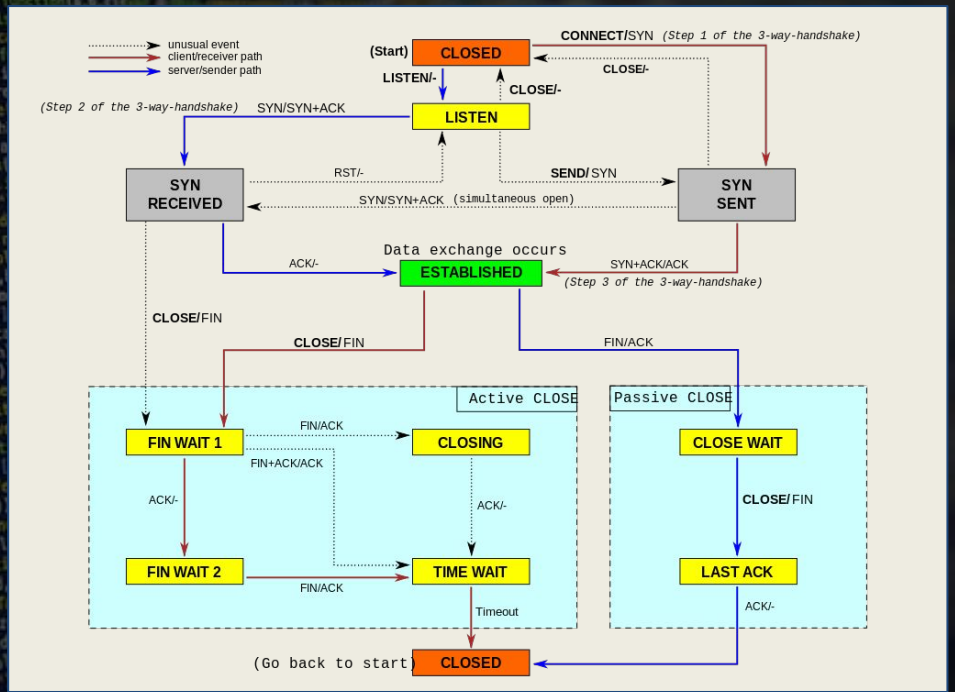

Class 4 - TCP/IP - Internet

- TCP - **RFC 793 - Transmission Control Protocol**
- Connection setup: The Three-way Handshake
- Establishes host sequence numbers
  - 1) A --> B SYN my sequence number is X
  - 2) A <-- B ACK your sequence number is X
  - 3) A <-- B SYN my sequence number is Y
  - 4) A --> B ACK your sequence number is Y
- 2& 3 are combined into SYN-ACK, hence 3 way handshake

Class 4 - TCP/IP - Internet

- TCP - **RFC 793 - Transmission Control Protocol**
- The Three-way Handshake

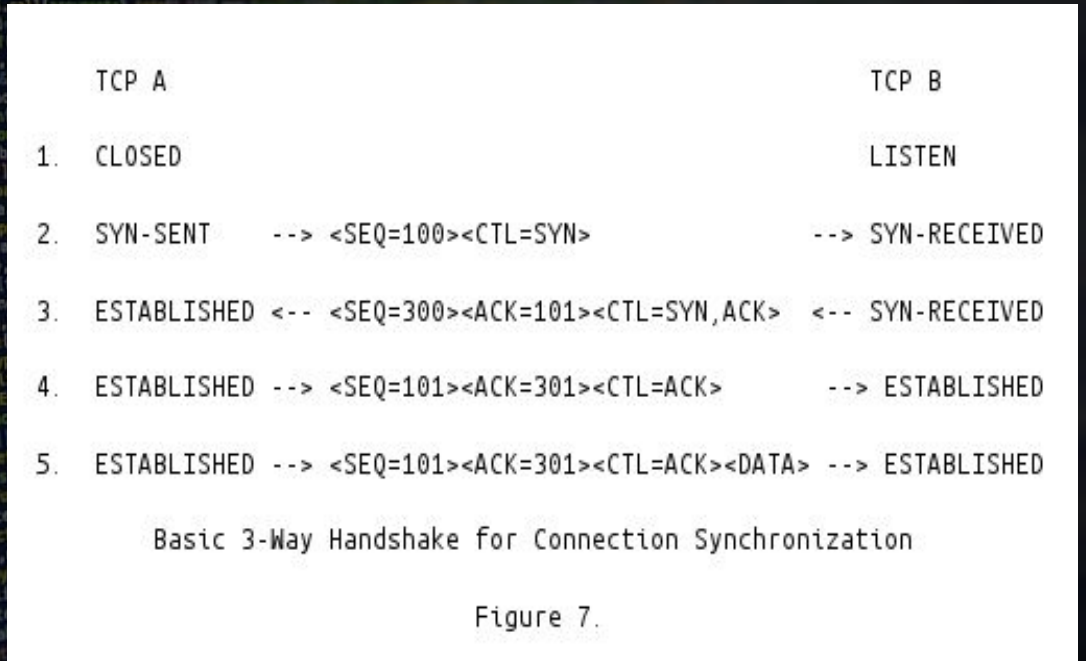

```
    TCP A                                                TCP B

1.  CLOSED                                               LISTEN

2.  SYN-SENT    --> <SEQ=100><CTL=SYN>               --> SYN-RECEIVED

3.  ESTABLISHED <-- <SEQ=300><ACK=101><CTL=SYN,ACK>  <-- SYN-RECEIVED

4.  ESTABLISHED --> <SEQ=101><ACK=301><CTL=ACK>       --> ESTABLISHED

5.  ESTABLISHED --> <SEQ=101><ACK=301><CTL=ACK><DATA> --> ESTABLISHED

        Basic 3-Way Handshake for Connection Synchronization

                                Figure 7.
```

Class 4 - TCP/IP - Internet

- TCP - **RFC 793 - Transmission Control Protocol**
- Connection close: Another 3-way handshake

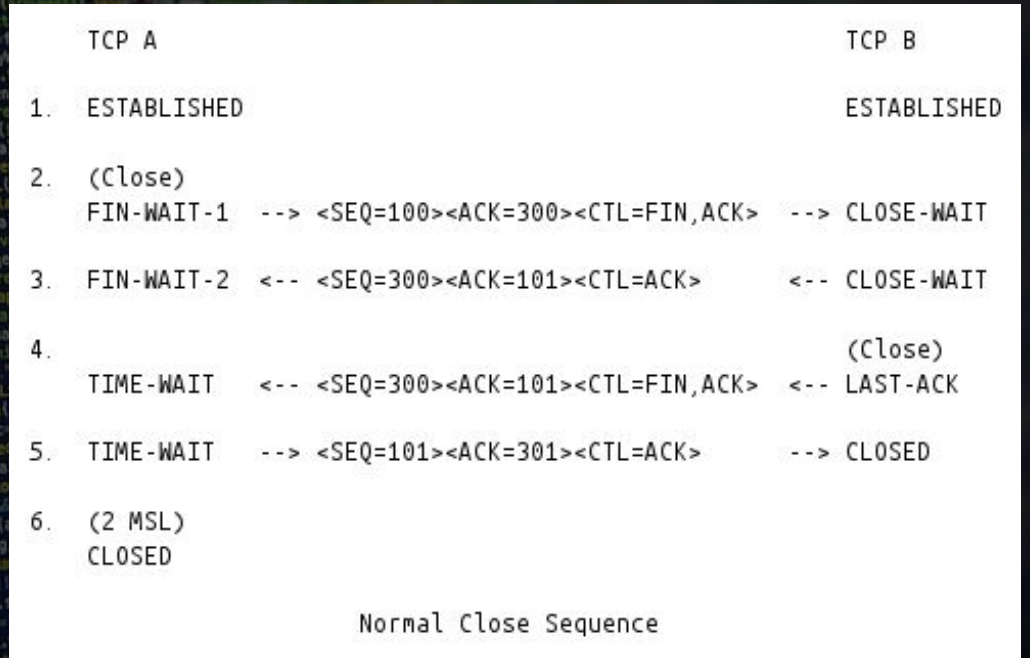

Class 4 - TCP/IP - Internet

- TCP - **RFC 793 - Transmission Control Protocol**
- Lots more to it:
  - Keepalive tuning
  - Buffer allocations
  - Congestion control
  - Window sizing
  - Syn cookies

Class 4 - TCP/IP - Internet

Hello, would you like to hear a TCP joke?

Yes, I'd like to hear a TCP joke.

OK, I'll tell you a TCP joke

OK, I'll hear a TCP joke.

Are you ready to hear a TCP joke?

Yes, I am ready to hear a TCP joke.

Class 4 - TCP/IP - Internet

OK, I'm about to send the TCP joke. It will last 10 seconds, it has two characters, it does not have a setting, it ends with a punchline.

OK, I'm ready to hear the TCP joke that will last 10 seconds, has two characters, does not have a setting and will end with a punchline.

I'm sorry, your connection has timed out...

...

Hello, would you like to hear a TCP joke?

Source: https://twitter.com/casheeew/status/481438450848395264

Class 4 - TCP/IP - Internet

- UDP - **RFC 768 - User Datagram Protocol**
  - “This protocol provides a procedure for application programs to send messages to other programs with a minimum of protocol mechanism. The protocol is transaction oriented, and delivery and duplicate protection are not guaranteed. Applications requiring ordered reliable delivery of streams of data should use the Transmission Control Protocol (TCP)“

Class 4 - TCP/IP - Internet

- UDP - **RFC 768 - User Datagram Protocol**
  - Key points: **delivery**, and **duplication** protection are **not** guaranteed. Transactional means connectionless
    - Delivery corollary: order is not guaranteed either
  - Thus, UDP applications must implement logic to handle ordering, data reassembly, as well as reliability
  - Essentially, fire and forget

Class 4 - TCP/IP - Internet

- UDP - **RFC 768 - User Datagram Protocol**
- Small size, 16 bits for source, 16 for destination, length, checksum
- Total of 8 bytes, the rest is data
- Checksum is optional
- Can be fragmented! Careful

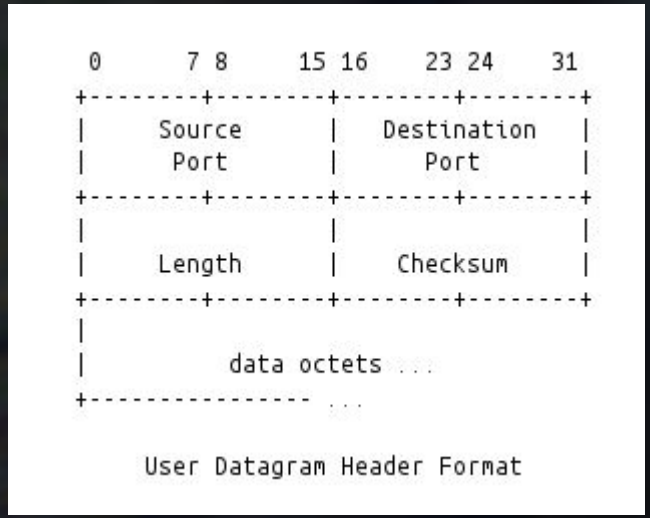

User Datagram Header Format

Class 4 - TCP/IP - Internet

- UDP - **RFC 768 - User Datagram Protocol**
  - More widely used than you might think
  - Many VoIP protocols use it, since small losses of data are perceptually hard to detect
  - Online gaming, video streaming for the same perceptual reasons, and because of less overhead

Class 4 - TCP/IP - Internet

- UDP - **RFC 768 - User Datagram Protocol**
  - I'd tell you a UDP joke, but you might not get it.

Class 4 - TCP/IP - Internet

- TCP & UDP ports: **Wikipedia**
- A host has 65,535 ports where a program can listen for connections
- On Linux, /etc/services contains a list of programs and ports. E.g. HTTP on port 80, SMTP on port 25
- Ports handle multiple sessions (1000's+). Useful metaphor: IP is a building address, Port is a mailbox in the building

Class 4 - TCP/IP - Internet

- Lots more to it!
  - Layer 1, yes, you can spend a lot of time (and money) here, things like 100Gbit fibre optic transceivers
  - Layer 2, where to begin
    - VLANs
    - Jumbo frames
    - Switch topology
    - Spanning tree
    - MAC tables, and ACLs
  - Layer 2/3 - ARP

Class 4 - TCP/IP - Internet

- Demo time!
  - UDP: nc -u 127.0.0.1 53 & tcpdump -nei any port 53
  - TCP: nc -l 1337 & nc 127.0.0.1 1337 & tcpdump -nei any port 1337
  - Note the difference in number of packets sent
  - Note hosts also have a local address: 127.0.0.1 for IPv4, and ::1 for IPv6. Used for communicating locally, e.g. OpenProject & PostgreSQL
  - Run `ip address show lo` to see the interface

Class 4 - TCP/IP - Internet

- Demo time!
  - HTTP with curl: curl -si 127.0.0.1 & tcpdump -nei any port 80
  - DNS with dig: dig A yorku.ca & tcpdump -nei any port 53
  - DNS with dig & TCP: dig +tcp A yorku.ca & tcpdump -nei any port 53
  - Run `ip address show lo` to see the interface