

# ECE 435 – Network Engineering

## Lecture 13

Vince Weaver

<http://web.eece.maine.edu/~vweaver>

[vincent.weaver@maine.edu](mailto:vincent.weaver@maine.edu)

17 October 2017

# Announcements

- HW#6 was posted  
No second part?



# HW#5 Review

- Header length was the most trouble, top 4 bits of nibble (0x8)
- It's a web request
- Size:  $0x46 = 70$  bytes,  $4/70 = 5.7\%$
- 3-way handshake SYN/SYN+ACK/ACK
- Sends hi / ack / sends back HI / ack. Note PSH sent so that it doesn't wait and piggyback
- Closing connection. FIN/ACK+FIN/ACK



Many people missed this question, poor placement of pagebreak?



# IPv4 Addresses

- Each IP address is 32-bits and has network address and host ID
- Unique to \*interface\* not necessarily to \*host\*.
- Who hands these out? ICANN and various regional authorities Internet Corporation for Assigned Names and Numbers  
Internet Assigned Numbers Authority (IANA)
- Regional Internet Registrars
  - AfriNIC (Africa)



- ARIN (N America),
- APNIC (Asia-pacific)
- LACNIC (latin america),
- RIPE NCC (Europe and rest)
- Can write many ways (all equivalent) but most common is dotted decimal



# Subnets

- Number of hosts available can be larger than possible
- Divide network into subnets
- All hosts on subnet have the same prefix (left bits)
- Use subnet mask indicating the leftmost bits to use as subnet
- Can look like 255.255.255.0 meaning only bottom 8 bits are for host

Alternately can write this as 192.168.8.0/24 (24 is number of leading binary 1s in mask)



# Classful IP Routing (No Longer Used)

- Routers just shifted right for A, B, and C class. Looked up A and B in table, C in hash table to find where to send
- Has a routing entry for each Class A (256), an entry for each class B (16k). Class C (2 million) a bit much, so hash table.
- Why so simple? In 80s memory and processors were expensive!
- Original classful addressing scheme (not necessarily used





anymore)

- Class A: 8 bit network (high bit 0) (24 bits of hosts) 0.0.0.0 to 127.255.255.255
- Class B: 16 bit network, (high bits 10) 128.0.0.0 to 191.255.255.255
- Class C: 24 bit network (high bits 110) 192.0.0.0 to 223.255.255.255
- Class D: multicast (high bits 1110) 224.0.0.0 to 239.255.255.255
- Class E: reserved (high bits 1111) 240.0.0.0 to 255.255.255.255



- Special cases
  - 0.0.0.0/8 reserved for current network (RFC 6890)
  - 10.0.0.0/8 private network (RFC 1918)
  - 100.64.0.0/10 shared address space (RFC 6598)
  - 127.0.0.0/8 loopback (RFC 6890)
  - 169.254.0.0/16 link-local (RFC 3927)
  - 172.16.0.0/12 private network (RFC 1918)
  - 192.0.0.0/24 IETF (RFC 6890)
  - 192.0.2.0/24 test (RFC 5737)
  - 192.88.99.0/24 IPv6 to IPv4 relay (RFC 3068)
  - 192.168.0.0/16 Private Network (RFC 1918)



- 224.0.0.0/4 IP Multicast (class D) (RFC 5771)
- 240.0.0.0/4 Reserved (class E) (RFC 1700)
- 255.255.255.255 Broadcast (RFC 919)
- .0 represents a subnet
- .1 is often (but not always) a router
- it all host bits 1, broadcast for that subnet
- 255.255.255.255 is broadcast for device that doesn't know own IP yet (DHCP)



# Classless Inter-Domain Routing (CIDR)

- Running out (have run out) of network addresses
- For many groups, Class-A too big, Class-C too small (three bears problem?)
- Merge neighboring class-C together
- RFC 1519
- Scalability problem: each network takes up space in routing table
- Solution, group neighboring class Cs together
- With CIDR bit more complex.



- Triplet with IP address, subnet mask, outgoing line.
- In theory has to scan all. If multiple matches, one with longest mask is used.
- There are algorithms to make this go faster.
- Example – from 444 in Tannenbaum



# Local IP Routing

- If on same subnet, send packet directly to destination
- Otherwise, send on outgoing. See Linux route command. Often a "default router" 0.0.0.0/0. If doesn't match any other, sent out over default route
- Algorithm: if to same host, skip network. If to same subnet, deliver directly (Ethernet, etc) otherwise, send to default router
- If multiple network interfaces: If to this machine, deliver it, If to directly connected subnet, directly deliver, else



deliver to next hop router

- How do we know if on network? If  $((\text{hostIP XOR destip}) \& \text{subnetmask}) == 0$
- If local, how do we map IP to MAC? We'll see that in a minute.
- Due to CIDR, longest prefix matching. If match both a /21 and /24 then 24 is the one to send to as it's the longest.
- Data structures. Hashes? Trie?



# Linux/UNIX routing setup

- “route” command
- `route add default gateway sets default gateway (router) for packets leaving the local network`
- also set up local subnets you are on, those packets don't need a router
- more complicated if you are configuring your Linux box to *\*be\** a router





# ICMP

- Internet Control Message Protocol
- Carried as a payload in an IP packet
- Type set to 1
- Codes
  - DESTINATION UNREACHABLE, Also if MTU is too small but do-not-fragment set
  - SOURCE QUENCH – should slow transmission rate (congestion), This is now usually done in transport layer



- REDIRECT – try the other router path
- TIME EXCEEDED – exceeded TTL, traceroute uses this
- PARAMETER PROBLEM – illegal value in header
- ECHO, ECHO\_REPLY – see if machine is up
- TIMESTAMP, TIMESTAMP\_REPLY – performance debug
- Some sysadmins block ICMP. Why?



# ping

- Mike Muuss in 1983  
`http://ftp.arl.army.mil/~mike/ping.html`
- Like sonar ping (Hunt for Red October), not any of the backronyms you might find.
- Ping the duck
- ICMP ECHO packet, waits for ECHO reply. Prints timing info, etc.
- Ping of death
- Ping flood



- Broadcast ping to x.x.x.255 (no longer works)
- Used to just say “host is alive”. People would make machines called elvis.



# traceroute

- Van Jacobson in 1987 (also wrote tcpdump)
- Uses ICMP
- \*not\* tracer-t
- Send packet with  $TTL=1$ , when sends ICMP error message know where first hop is
- Send packet with  $TTL=2$ , find next
- Linux traceroute sends UDP packets as originally ICMP requests weren't supposed to generate ICMP errors

