ECE 435 – Network Engineering
Lecture 22

Vince Weaver

https://web.eece.maine.edu/~vweaver

vincent.weaver@maine.edu

24 March 2025

Announcements

- HW#7 was graded
- HW#8 will be posted (short with close deadline)
- Remember project topics are due 28th (Friday)
- Hand back and go over midterms (Average 85%)

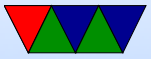

HW#7 Review – IPv4 Header

```
0x000e:  4500 = version(4), header length(5)=20 bytes
                ToS=0
0x0010:  0038 = packet length (56 bytes)
0x0012:  572a = identifier
0x0014:  4000 = fragment 0100 0000 0000 0000 =
                do not fragment, offset 0
0x0016:  40 =    TTL = 64 hops
0x0017:  06 =    Upper layer protocol (6=TCP)
0x0018:  69cc = checksum
0x001a:  c0a80833 = source IP 192.168.8.51
0x001e:  826f2e7f = dest IP 130.111.46.127
```

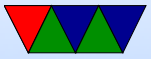

HW#7 Review – IPv4 Addresses

- Valid IPs
  - 1.1.1.1 = Y
  - 123.267.67.44 = N
  - 192.168.8.1 = Y
  - 3232237569 = 192.168.8.1 = Y
  - 0xc0a80801 = 192.168.8.1 = Y
- A class-A allocation is roughly $2^{24}/2^{32}$ which is 0.39%

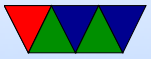

HW#7 Review – Subnetting

- 192.168.13.0/24. subnet 255.255.255.0, lowest ip 192.168.13.1, highest 192.168.13.254 (traditionally can't use the host values with all 0 or 1s (so .0 and .255 on a /24). There's a push to allow .0 on Linux
- First hop not local (how to tell?) goes to router Otherwise go direct (can you go direct? how).

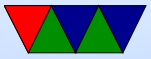

HW#7 Review – IPv4 ping/traceroute

- Ping google. 1e100.net?
- Traceroute. Some routers block?
  Used to pass through Neville hall
- Interesting, people tracerouting umaine from spectrum have packets going via chicago and boston bngrme/sebgme/rochny/chgil

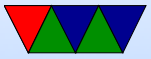

HW#7 Review – NAT

- No 192.168.8.x should not be able to connect to outside directly.
- NAT is happening.
- Why is nat showing UNREPLIED? TCP vs UDP difference. Can you detect when TCP connection is closed? Yes. Can you detect when UDP connection is done? No. Must keep port open a bit in case reply. How long. Forever? What goes wrong with that?

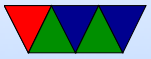

Implementations

- Actual Router
- Can install on your Linux machine
- Zebra was traditional, discontinued
- Quagga
- BIRD
- OpenBGPD and OpenOSPFD
- Potentially dangerous to mess around with unless you isolate your network well

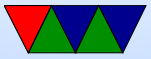

Other types of Routing

- Mobile – what if machines can come and go?
  have a "home" location. Packets go there. When you get on network, update with actual location. Network gets packets at home location, encapsulates and sends to actual location
- Ad Hoc / Mesh Routing
  Bunch of machines in an area, routers and devices can come or go more or less randomly.
  route discovery, pass packets to neighbors in hope it gets

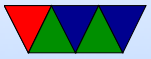

it closer to router

- Peer to Peer File Sharing
  - Centralized server? Napster? Easy to take down.
  - Want Distributed, no central control.
  - Flooding: connect to one other connected node. Floods requests (sort of like broadcast) until it finds who has file, then direct connect to transfer.

  - distributed hash table
- Secret routing
  TOR / The Onion Router

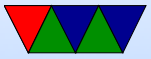

Packet encrypted multiple times, in layers. Randomly sent to next machine which decrypts that layer, passed on

At end comes out random "exit node" and drops onto regular internet

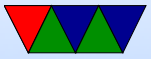

Broadcast Routing

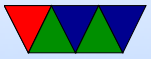

Casting

- Unicast – 1:1 – one sender, one destination
- Broadcast – 1:all
- Multicast – 1:many – specify a subset of all
- Anycast – a set of equivalent hosts, which one gets the packet depends on something like closeness / latency
- Geocast – broadcast to limited geographic area

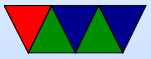

Anycast

- Can have multiple servers with same IP address
- Routers (via BGP) configured to send you to the closest one
- Companies can have datacenters on east cost / west coast/ europe, etc, and using same IP get to closest

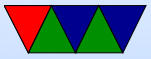

Examples

- DNS root servers are load-balanced via anycast these days
- IPv6 6to4 tunneling uses anycast with 192.88.99.1?
- Content Delivery Networks (CDN) can load balance heavily used sites this way

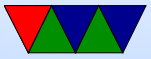

Anycast Downsides

- If you have a TCP connection and somehow your routing changes to a different server location your connection will break
- Can hijack connection if you can get your fake routing info into a server.

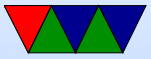

IPv6 Anycast

- IPv6 you can do things at BGP level like IPv4
- IPv6 also has anycast built in for on same subnet
- Can assign special anycast address and the routers responsible for handling it
- Lowest address ::0/124 and also top 128 ffff:ffff:ffff:fff80/121

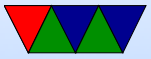

Unicast/Multicast/Broadcast

- Unicast – send from one machine to another
- What if want to send to multiple?
  - Multi-unicast – open direct connection to each destination. Inefficient
  - Broadcast – send to *every* destination? Waste bandwidth, but also need to know all possible destinations
  - Flooding? Also too much bandwidth
  - Multi-destination routing

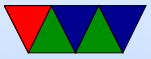

Multicast Goals

- Only send to users who want it
- Each member only receives one copy
- No loops
- Path traveled should be optimal

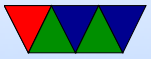

Multicast Structure

- Spanning tree – tree with source as root and members as leaves
- Reverse-path forwarding

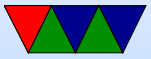

Why would you multicast?

- Live streams? Backups?
- Why not just multi-unicast?
  - More work on sender, many more packets sent
  - Latency between first and last packet sent

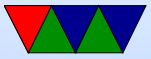

Multicast IP

- For IP, just join a class D network
- To both sender and receiver it's like sending/receiving a unicast packet
- all the hard work done by routers
- How do you join a multicast group?
- Router two tasks: group membership management, packet delivery.

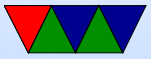

Group Management

- IGMP (Internet Group Management Protocol)
  - IGMPv3 RFC 3376
  - query, report, leave
  - querier and noquerier
  - router with lowest IP is querier
  - no real controls on who can join or send

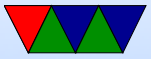

Multicast Trees

- Steiner tree – NP complete, no one uses
- Heuristics, but none generate entire tree as need centralized and global knowledge
- DVMRP (Distance-Vector Routing Protocol)
  original protocol, MBONE
- Reverse path Forwarding – flood packet out all interfaces except one it came in on. Can have loops; drop dupes. Then forward on the one that has traveled the shortest path.

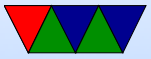

Is running the routing table backwards

- Reverse path Broadcast – avoid getting multiple packets
- Protocol Independent Multicast (PIM)
  DVRMP not scalable for multicast groups with sparse members
- MOSPF
- CBT

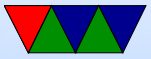

MBone

- Attempt to make multicast useful in mid 1990s
- Concerts over the Mbone

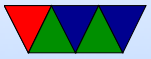

Local Network Broadcasts

- 224.0.0.0/4 was reserved from Class D for multicast
- 224.0.0.0 to 224.0.0.255 for local network broadcasts
- Things like cluster stats (ganglia, can never get to work?)
- Routing info protocol (RPIv2) OSPF, mDNS, etc.

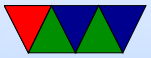

mDNS

- Multicast local network hostname resolution
- Bonjour (mac), Avahi (Linux)
- Multicast to 224.0.0.251 (ipv4) or ff02::fb (ipv6)
- Issue if two machines have same name
- Broadcast name as connect to network, all devices on local net subscribe to broadcast at that address