

ECE 435 – Network Engineering

Lecture 16

Vince Weaver

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

`vincent.weaver@maine.edu`

12:30pm, Barrows 125

19 March 2024

Announcements

- Reminder HW#7 will be graded soon
- Remember project topics are due 26th (next Tues)
- Hand back and go over midterms (Average 83%)
- Due to faculty interviews, Monday office hours cancelled this week and next
- Joseph Olivas from Intel talk, Friday 1pm Hill Auditorium



Aside on New Pi4 Cluster

- “What I did over spring break”
- Five pi4 nodes (on 8GB, rest 4GB)
- Gigabit ethernet, pi4 fast enough to keep up
- Power over Ethernet
- USB ethernet for external connection on head node, NAT
- Did not use DHCP or network authentication due to low number of nodes
- Did set up NFS file sharing



- ssh-agent for control of nodes
- ganglia multicast



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

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
 - 123.267.67.44 = N
 - 8.8.8.8 = Y
 - 3232237569 = 192.168.8.1
 - 0xc0a80801 = 192.168.8.1
- 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?



HW#7 Review – IPv6 Addresses

- 2607:f8b0:4009:0801:0000:0000:0000:200e – OK
- 2607:f8b0:4009:801::200e – OK
- 2607:f8b0::4009:801::200e – can you have two colons?
- 123.45.67.18 – ipv4



HW#7 Review – IPv6 Packet

0x000e: 6002 2618 :

6 = IPv6

00 = traffic class

2618 = flow label

0x0012: 0031 = payload length, 49

0x0014: 11 = next header = 0x11, UDP

0x0015: 40 = hop limit 0x40, 64

0x0016: 2610 0048 0100 08da 0230 18ff feab 1c39
source address

0x0026: 2001 4860 4860 0000 0000 0000 0000 8844
destination address



HW#7 Review – Traceroute

- internet2
- bost/hart/newy probably boston, hartford, new york
- lon2.uk London
- janet is british academic network
- 6→7 across ocean
- 80ms = ?? speed of light
- $80 \times 10^{-3} \text{s} \times 3 \times 10^8 \text{m/s} = 24000 \text{km?}$ $5500 \text{km} = 1/4$ speed of light?




HW#7 Review – Traceroute6

- different hops? IP6 different? random chance
- hop 5→6
- Washington? internet 2?
- Abilene was the predecessor to internet2
- fra.de Frankfurt Germany probably not France
- latency 133ms rather than 106ms



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 problems/benefits

- Can spread server load around (DNS servers, web servers, netflix servers)
- Can hijack connection if you can get your fake routing info into a server.



IPv6 Anycast

TODO: move that here



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 (RIPv2) 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



The Physical Layer

See next lecture

