

ECE 435 – Network Engineering

Lecture 23

Vince Weaver


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

vincent.weaver@maine.edu

27 March 2026

Announcements

- Don't forget HW#7
- Project ideas due
- HW#8 will be posted



Peering

- How companies agree to connect their networks together.
- There's not really a master connection, but instead companies agree to have routers talk to each other via BGP.
- Types
 - Transit – pay money to pass through network.
 - Peering – In many cases no money changes hands. Why? Well if you have a lot of users, but no content, people won't stay with you. Same if you have content



but no access to users. Averages out and is mutually beneficial.




Peering Analogy

- You and your neighbor both do online games
- You both pay the ISP (maybe per byte) for this, but latency as it goes out to the ISP network and back
- Would it make sense to just run a wire between your houses and have a direct connection for game data?
This is peering and probably mutually beneficial
- Even if it did, would you still maybe not want your neighbor be able to get all their internet access through your network?



Reasons to Peer

- Increased redundancy
- Increased capacity/performance
- Increased routing control
- Fame (high-tier network)
- Ease of requesting aid (?)
- Avoid tromboning (without peering, your connection might go from UMaine to New York, then back to Orono to your apartment if UMaine and your local provider don't peer)



Peering Locations

- Peering locations, often in large data centers.
- Internet Exchange Points (IXP)
- At one point there were 4 major ones (Metropolitan Area Exchange) MAE-East (Virginia) [in basement of parking garage, at one point half of internet went through here], Chicago, NY, SF. All defunct now
- Exchange map: <https://www.internetexchangemap.com/>
- PNI – instead of IXP can just have a direct connection between two networks



Peering Tiers

- Tier 1 network is one that can reach rest of internet without paying for transit;
- Tier 2 peers with some but purchases for other;
- Tier 3 only purchases



Depeering


- If you think you aren't getting a good deal, break up
- Some situations there is a fight, a hope that the customers lose enough performance will have to repeer.
- Can be a lot of drama
- Large companies depeering and instead charging for traffic. Germany, Vodafone?

<https://news.ycombinator.com/item?id=45848484>



Net Neutrality

- This is a related issue
- Should content providers have to pay ISPs for carrying their packets
- Can ISPs prioritize packets from content providers willing to pay more



IPv6 Peering Issues


- IPv6 Peering issues – see

https://www.theregister.co.uk/2018/08/28/ipv6_peering_squabbles/



Routing Security Issues

- Problems – routing black hole, use BGP to send addresses intentionally to 0.0.0.0 and get dropped. BGP will propagate
- router update mistakes can accidentally blackhole parts of the internet
- BGP Hijacking – taking over another group's network addresses via BGP
 - Normally if you have AS you announce prefixes that you originate, for example 130.111.218.23/23



- Can hijack by claiming you originate a network you don't
- Can claim you have a shorter route to the network
- Can announce a more specific CIDR prefix than the real one



Routing Security – Events

- In 2008 Pakistan was trying to blackhole Youtube and accidentally announced to world via BGP and took it down world wide
- <https://arstechnica.com/information-technology/2019/06/bgp-mishap-sends-european-mobile-traffic-t>
- https://www.theregister.co.uk/2019/06/24/verizon_bgp_misconfiguration_cloudflare/
Verizon accidentally routed a lot of internet through Pittsburgh Steel Mill
- October 4 2021 – Facebook dropped off internet for 6 hours, DNS took down the BGP links. Had trouble



getting back up, including story that they couldn't get card access to datacenter due to internet being down

- March 2022, part of twitter routed through Russia



Routing Security – Countermeasures

- Filtering – filter BGP requests to only let valid requests escape
- Resource Public Key Infrastructure (RPKI) – digitally sign address ranges and only allow authorized origins (Route Origin Authorizations)
- Work is ongoing on RPKI but not everyone has it implemented yet



Routing in the News

- <https://blog.cloudflare.com/route-leak-incident-january-22-2026/>



Implementations

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



Other types of Routing – Mobile

- 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



Other types of Routing – Ad Hoc / Mesh

- 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




Other types of Routing – 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



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




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 coast / 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
- Story about mbone at work in the 90s



Why Did Broadcast Never Really Take Off?



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

