# ECE 435 – Network Engineering
# Lecture 21

Vince Weaver

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

vincent.weaver@maine.edu

14 March 2025

# Announcements

- Don't forget HW#7
- Midterms still being graded

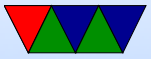

# Hierarchical Routing

- Would you want to have all routers in network on flat network?
  Routing table would be a bit complex
- Split into a hierarchy
- Network made up of Autonomous Systems (AS)

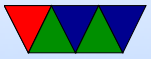

# Autonomous System (AS)

- A network under control of one group, with one routing policy
- Inside an AS, interior routing, between is exterior routing.
- Usually you need to be a large enough group and be able to get some network connectivity
- Then you need to convince your ISP to add a route to your AS

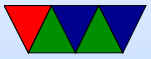

# Autonomous System Numbers

- How to get ASN (number for AS)? Similar to getting IP addresses.
- Traditionally were 16-bit numbers, but ran out. In 2007 expanded to 32-bit. X.Y (dotted decimal). Old 16-bit are 0.X
- Can look up, UMaine is AS557
  `https://bgpview.io/asn/557`

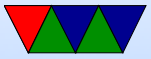

# Routing

- Systems under same command (same ISP) use intra-domain routing protocol, or interior gateway protocol (IGP)
- Border routers connect to border routers of others
- Inter-domain routing, EGP (exterior gateway protocol)

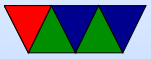

# AS Types

- Stub AS – like ISP with customers, one gateway to internet
- Multihomed AS – multiple gateways (why?) redundancy. traffic generally doesn't flow through
- Transit AS – traffic can flow through network
- Internet Exchange Point (IX/IXP) – where networks can meet up

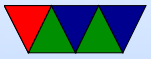

# Diagram

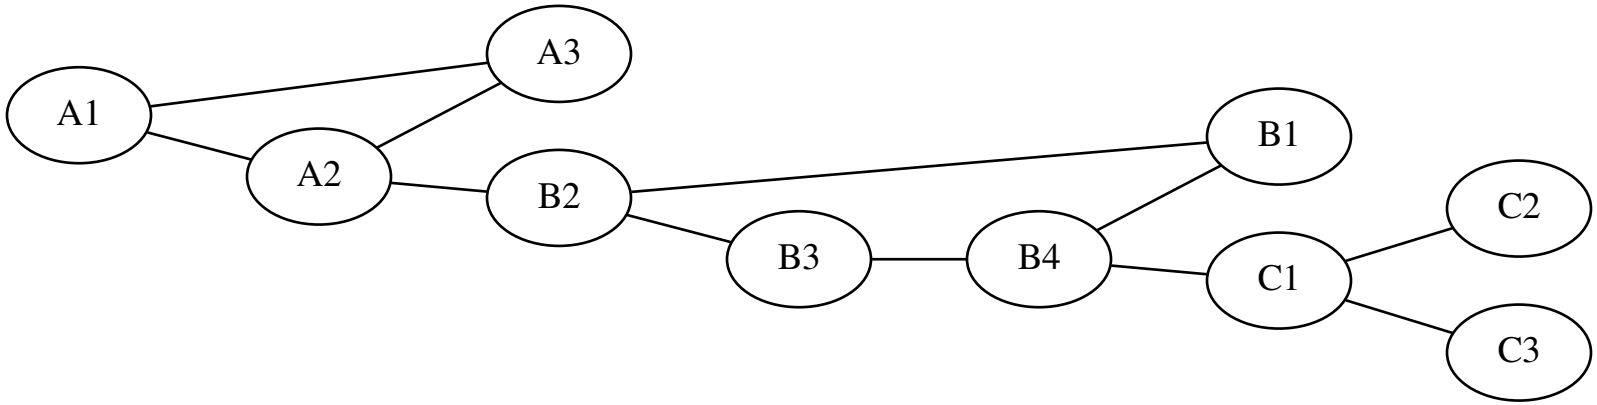


- Packet A1 - A3 internal A1 - B2 goes to border router and across, then local A1 - C2 goes to border router to B network, across local to B/C border, then finally to C
- If flat network, need to know 10 machines in routing table

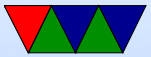

- In hierarchical only need to communicate to 2-3 other routers, find way to border router

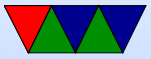

# Intra-Domain Routing / Interior Gateway Protocols

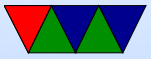

# Historical – RIP (Routing Information Protocol)

- Used by ARPAnet until 1979(?)
- by Xerox, included in BSD, routed RFC 2543
- distance vector routing, with hop count, max 15 hops
  - RIP advertisements over UDP port 52
  - Send advertisement every 30s, or when changes
  - Only sends to neighbors
- Routing table: dest, next hop, distance
- Algorithm

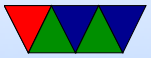

- Get table update
- Increment all hops by 1 (you're one hop away)
- Go down list.

  If route not in table, add it

  If route there, and next hop same (but cost diff), replace it as this is new info

  If route there but cost less, replace it

- On power up, comes up with hard-coded routes and values of 1 and no next-hop. Can send packet to request immediate update from neighbors.
- Packet description

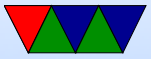

- Timers
  - Periodic timer, technically 30s, reality randomized between 25 and 35 (why?)
  - Expiration timer – 180s. If no update in this time, problem, hop count set to 16 (unreachable)
  - Garbage collection – 120s – once unreachable, advertise it as such for a while before removing so others notice
- Issues
  - Slow Convergence – a change in routing tables takes 30s per hop to propagate through

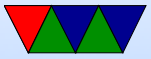

Part of why limited to 15 hops

- Instability – packets can be caught in loops. Ways to fix:

  Triggered update – send update info immediately, not wait 30s

  Split Horizon – if a router sends you update info, don't send this back to it Poison Reverse – like split horizon, but when send back, mark as 16 the routes received from that interface.

- There was a RIP2

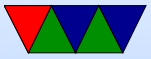

# OSPF (Open Shortest Path First)

- successor to RIP. RFC2328 (5340 for IPv6)
- Idea of Areas inside of an AS. Split up into areas.
- Each area connected by backbone router
- Router on two areas is area border router
- Link-state Routing
  - State is flooded: when a change happens (and only then) it sends this state to all neighbors, which send to all neighbors, until the whole network receives it
  - each router uses Dijkstra to find least cost for self,

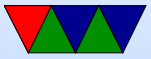

builds table

- Types of link
  - Point-to-point – routers directly connected
  - Transient Link – network with several routers can be simplified?
  - Stub Link – a network connected to only one router
  - Virtual Link – a path between two routers that traverses other routers
- load balancing – supports equal-cost multipath routing (can equally use equal cost routes)
- supports CIDR routing

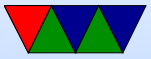

- support available for multicast
- 8-byte password for authentication
- supports hierarchical
- example? complex!

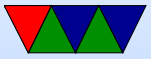

# Inter-Domain Routing

- Can be complicated.
- Say company with network, and two connections to outside X,Y. Don't want to send packets out and back even if it looks like lower cost.
- Also don't want to transit packets between X and Y for outsiders. Policy.

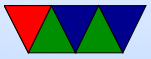

# BGP (border gateway protocol)

- Intro in 1989 (sketched on two napkins?), started being used 1994
- On fourth version – BGP4 RFC 4271
  support for CIDR and route aggregation
- Uses TCP (reliable) port 179.
- Works for both IPv4 and IPv6 (the latter as an extension)
- Uses path vector rather than distance vector
  - full path, not just next-hop
  - exchanges info with neighbor, but includes complete

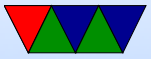

path info to avoid looping.
  - Each AS has unique number, so if it sees itself in the path knows there is a loop.
  - Policy routing – can also reject new route based on policy
- Four types of messages – open, update, keepalive, notification
- Whole table not passed around (Due to size), only updates
- Due to size of internet, uses distance vector over link state.

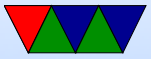

- Keeps track of all feasible paths, but only advertises the "best" one

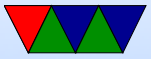

# Interior / Exterior BGP

- Interior
  - Interior is a full mesh
  - iBGP makes sure that the setups for multiple gateway routers are kept synchronized
  - Can have a Route Reflector (RR) to avoid the overhead of full mesh
- Exterior
  - eBGP used to talk between other exterior routers at peers.

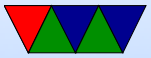

# Routing table Size

- Example. Full BGP of internet backbone router might have more than 300,000 entries (2010) now over 965,000 (2024)
- `http://bgp.potaroo.net/`
- Some routers had limit of 512k so on August 12 2014 part of internet went down when crossed the border.
- Ipv6 currently around 205k (march 2024)

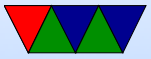

# Peering

- 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

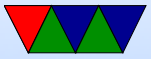

but no access to users. Averages out and is mutually beneficial.

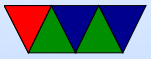

# Peering Analogy

- 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?

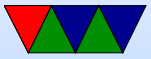

# Reasons to Peer

- 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)

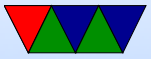

# Peering Locations

- 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

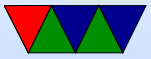

between two networks

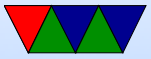

# Peering Tiers

- 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

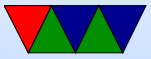

# Depeering

- 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

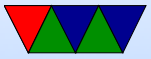

# Net Neutrality

- 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

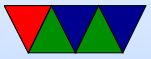

# IPv6 Peering Issues

- IPv6 Peering issues – see

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

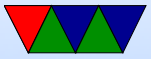

# Routing Security Issues

- 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 groups network addresses via BGP
  - Normally if you have AS you announce prefixes that you originate, for example 130.111.218.23/23

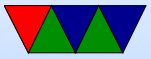

- Can hijack by claiming you originate a network you don't
- Can claim you have a shorter route to the network
- Can announce a more specific CIDR prefix than the real one

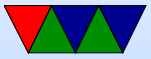

# Routing Security – Events

- 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

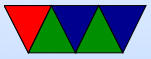

getting back up, including story that they couldn't get card access to datacenter due to internet being down

- March 2022, part of twitter routed through Russia

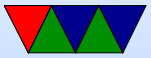

# Routing Security – Countermeasures

- 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

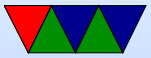

# Implementations

- 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

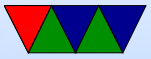

# Other types of Routing

- Mobile – what do you do when 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 Routing
  Bunch of machines in an area, routers and devices can come or go more or less randomly.

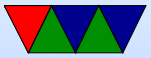

route discovery

- 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

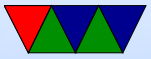

TOR / The onion Router

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

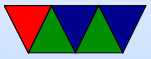

# Broadcast Routing

See next lecture