ECE 435 – Network Engineering
Lecture 15

Vince Weaver

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

vincent.weaver@maine.edu

24 October 2017

Announcements

- Project
- RFC791 post-it note:

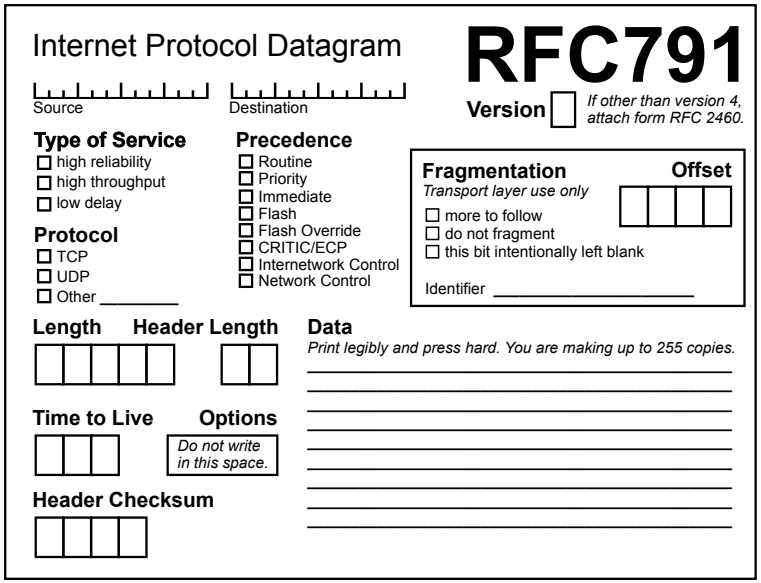

for more info, check IPv4 specifications at http://www.ietf.org/rfc/rfc0791.txt

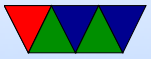

HW#6 Review

- Decoding header, much better than last homework
- Network connections
  - CLOSE-WAIT: received a FIN and ACKed it, waiting to close

    Only a few, https and imap
  - ESTAB: established, a few ssh, https, imap connections
  - SYN-RECV: way too many, SYN flood
  - TIME-WAIT: connection closed, waiting a bit before

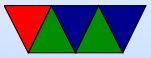

re-using port
  - UNCONN – UDP listening. 789? ipp, mdns (multicast DNS, bonjour, can find names on network w/o running DNS), lsof -i udp:789, rpcbind
  - LISTEN – listening. Can see ipp (CUPS printing), netbios/microsoft, apparently have SAMBA running,
- Synflood, by default Linux uses SYN cookies to defend against this

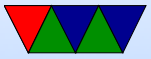

Project

- Can work in groups
- Do something interesting network related.
- Can use any operating system and written in any language (asm, C, python, C++, Java, etc.)
- Coding, benchmarking
- Past projects: network games, firewall config, network attached storage, mesh networks
- Will be a final writeup, and then a 10 minute presentation and demo in front of the class during last week of classes.

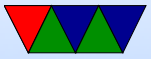

Hierarchical Routing

- Would you want to have all routers in network on flat network? Crazy routing table
- Split into a hierarchy
- Autonomous System (AS) – a network under control of one group.
  Inside an AS, interior routing, between is exterior routing.
- Systems under same command (same ISP) use intra-domain routing protocol, or interior gateway protocol (IGP)

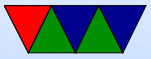

- Border routers connect to border routers of others
- Inter-domain routing, EGP (exterior gateway protocol)
- types
  - Stub AS – like ISP with customers, one gateway to internet
  - Multihomed AS – multiple gateways (why?) redundancy. traffic generally doesn't flow through
  - Transit AS – traffic can flow through network

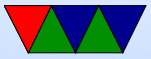

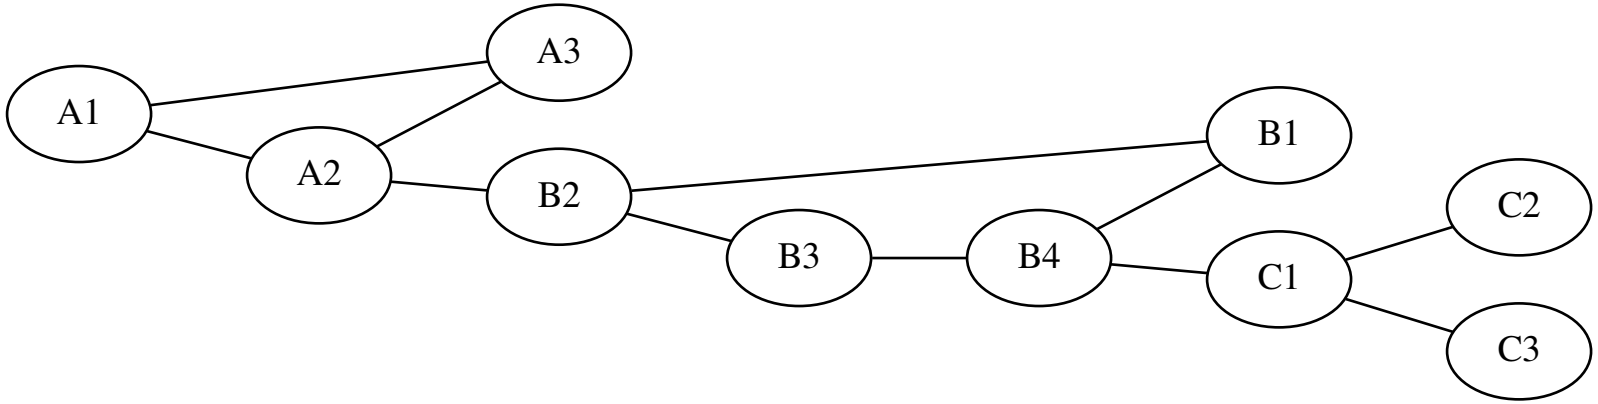

- Packet A1 - A3 internal A1 - B2 goes to border router and across, then local A1 - C2 goes to border router to B network, across local to B/C border, then finally to C
- If flat network, need to know 10 machines in routing table
- In hierarchical only need to communicate to 2-3 other routers, find way to border router

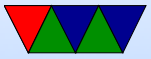

Intra-Domain Routing / Interior Gateway Protocols

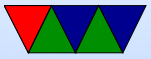

RIP (Routing Information Protocol)

- 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 ever 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)

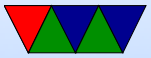

- Go down list.
  If route not in table, add it
  If route there, and next hop same (but cost diff), replace it as this is new info
  If route there but cost less, replace it
- On power up, comes up with hard-coded routes and values of 1 and no next-hop. Can send packet to request immediate update from neighbors.
- Packet description
- Timers
  - Periodic timer, technically 30s, reality randomized

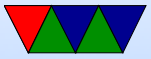

between 25 and 35 (why?)
  - Expiration timer – 180s. If no update in this time, problem, hop count set to 16 (unreachable)
  - Garbage collection – 120s – once unreachable, advertise it as such for a while before removing so others notice
- Issues
  - Slow Convergence – a change in routing tables takes 30s per hop to propagate through
    Part of why limited to 15 hops
  - Instability – packets can be caught in loops. Ways to

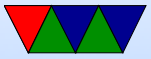

fix:

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

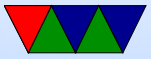

OSPF (Open Shortest Path First)

- successor to RIP. RFC2328 (5340 for IPv6)
- Idea of Areas inside of an AS. Split up into areas. Each area connected by backbone router
- 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

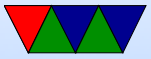

- Types of link
  - Point-to-point – routers directly connected
  - Transient Link – network with several routers can be simplified?
  - Stub Link – a network connected to only one router
  - Virtual Link – a path between two routers that traverses other routers
- load balancing – supports equal-cost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

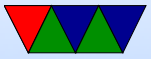

- 8-byte password for authentication
- supports hierarchical
- example? complex!

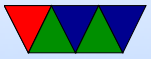

Inter-Domain Routing

- 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.

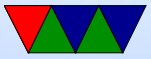

BGP (border gateway protocol)

- Intro in 1989, four versions – BGP4 RFC 4271
- 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.

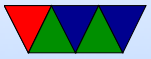

  - Policy routing – can also reject new route based on policy
- Four types of messages – open, update, keepalive, notification
- Whole table not passed around (Due to size), only updates
- Due to size of internet, uses distance vector over link state.
- interior and exterior BGP
- iBGP makes sure that the setups for multiple gateway routers are kept synchronized

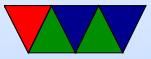

- eBGP used to talk between other exterior routers at peers.
- Keeps track of all feasible paths, but only advertises the "best" one

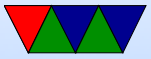

Routing table Size

- Example. Full BGP of internet backbone router might have more than 300,000 entries (2010) now over 700,000.
- `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 only around 20k

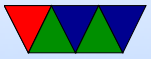

Peering

- 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.

- 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

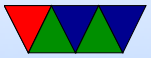

access to users. Averages out and is mutually beneficial.

  - Increased redundancy
  - Increased capacity
  - Increased routing control
  - Improved performance
  - Fame (high-tier network)
  - Ease of requesting aid (?)

- Customer – you buy an internet connection

- Peering locations, often in large data centers.
  At one point there were 4 major ones (Metropolitan Area

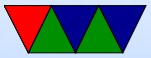

Exchange) MAE-East (Virginia) [in basement of parking garage, at one point half of internet went through here], Chicago, NY, SF. All defunct now

- 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.
- Related – net neutrality
- Tiers – Tier 1 network is one that can reach rest of internet without paying for transit; Tier 2 peers with

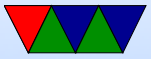

some but purchases for other; Tier 3 only purchases

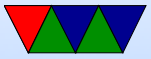

Routing Security Issues

- 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
- In 2008 Pakistan was trying to blackhole Youtube and accidentally announced to world via BGP and took it down world wide

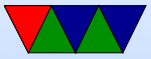

Implementations

- Actual Router
- Can install on your Linux machine
- Zebra was traditionally, discontinued
- Quagga
- BIRD
- OpenBGPD and OpenOSPFD

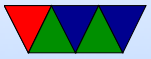

- Potentially dangerous to mess around with unless you isolate your network well