Networking for people who don’t run Networks

A.k.a How to talk to Network Engineers without sounding clueless

Mike Lovell – OpenWest 2017

Purpose

The IP Network Model

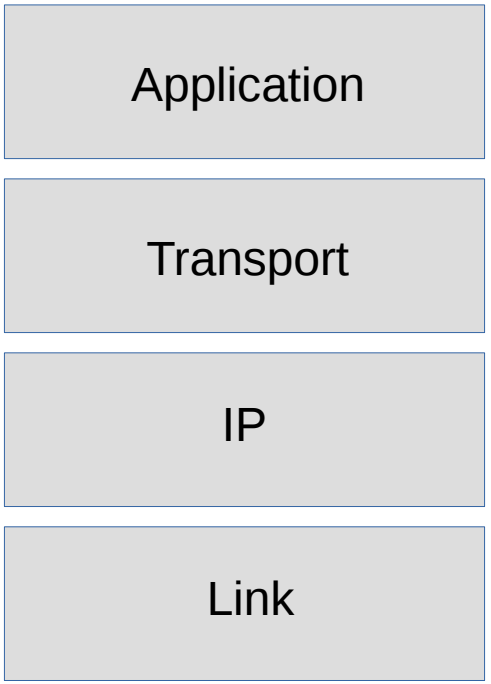

HTTP, SSH, Telnet, Email, Memes

TCP, UDP, ICMP, SCTP, GRE, etc

IPv4, IPv6

Ethernet, WiFi, VPN, Serial, Infiniband, etc.

IP Addresses

- Unique identifier of a node in an address scope
- 32 bits in IPv4
  - 192.168.1.1
- 128 bits in IPv6
  - fd00::1
  - 2607:f8b0:4005:806::200e
- Nodes can have multiple addresses
  - Different interfaces and/or Different Scopes
- Address has 2 parts. Network Prefix and Host Address

Subnets, Prefixes, and CIDR

- Addresses are grouped together on base 2 boundaries
- Groups are called a Subnet or a Prefix
- All Addresses within a Subnet are assumed to be on the same Link
- Subnets are signified by the number of bits in a network bitmask
  - Mask of the highest bits in the address
  - Old Method: 255.255.255.0
  - New Method: /24

Subnet Calculations

- Number of total addresses

  2^(Address Length – Subnet Length)

  IPv4 /24: 2^(32-24) = 2^8 = 256

  IPv6 /64: 2^(128-64) = 2^64 = A really big number
- ipcalc is a useful tool
- A few addresses are unusable for nodes
  - Network Address is first address of a subnet. i.e. 192.168.1.0/24
  - Broadcast address is the last address. i.e. 192.168.1.255

IPv6 Addresses

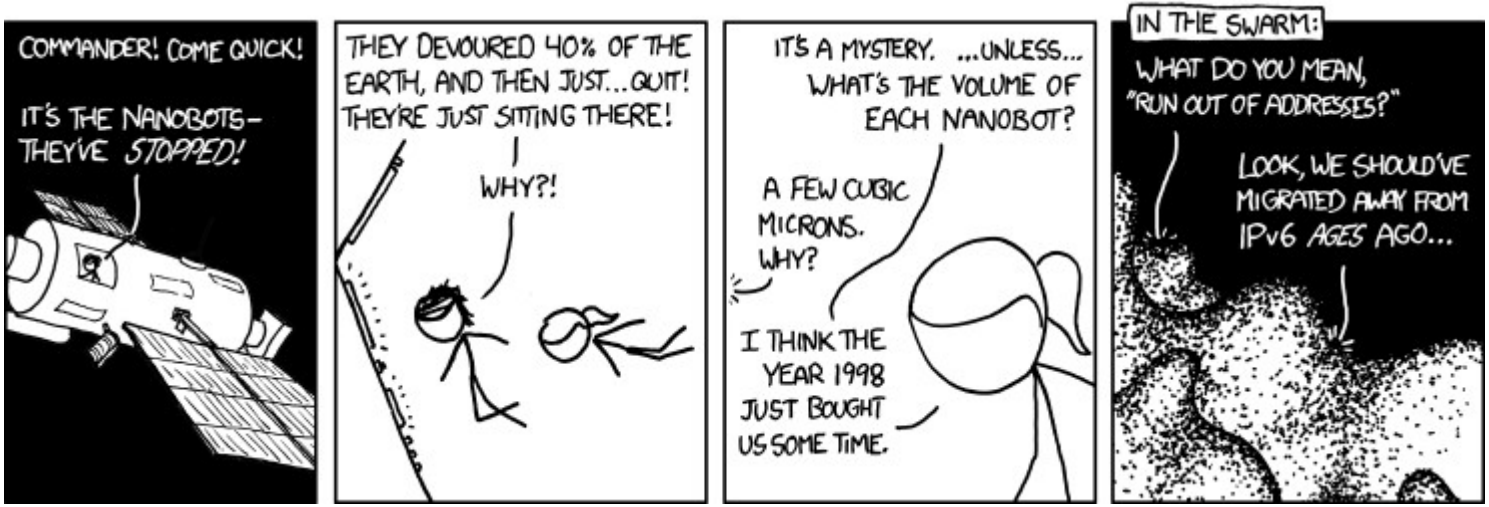

Credit: Randall Munroe https://xkcd.com/865/

See https://www.youtube.com/watch?v=4I_9iVjWi8c for a previous talk about IPv6

Link Layer - Ethernet

- Originally built as a shared, broadcast medium or bus
  - All nodes shared one wire
  - Packets got delivered to all hosts
- Destination, Source, Ethertype, Payload, and Checksum are the only parts of the original structure
  - No TTL
- Later extended to Repeated and then Switched networks, without changing the frame format
  - Oh Spanning Tree, how I love and despise your simplicity

Link Layer - Ethernet

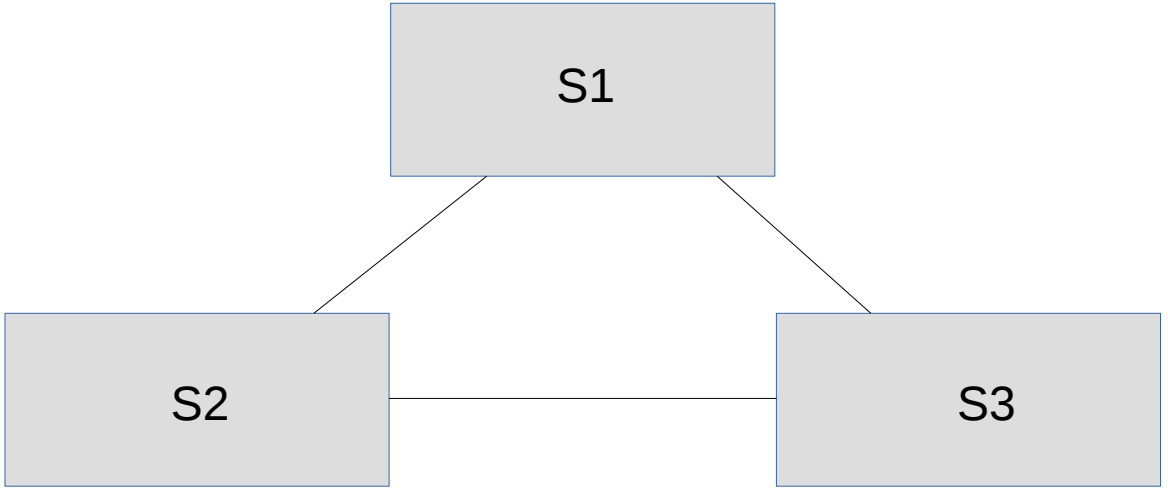

Link Layer - Ethernet

I think that I shall never see
A graph more lovely than a tree.
A tree whose crucial property
Is loop-free connectivity.
A tree which must be sure to span
So packets can reach every LAN.
First the Root must be selected
By ID it is elected.
Least cost paths from Root are traced
In the tree these paths are placed.
A mesh is made by folks like me
Then bridges find a spanning tree.

Algorhyme – Radia Perlman

Link Layer - Ethernet

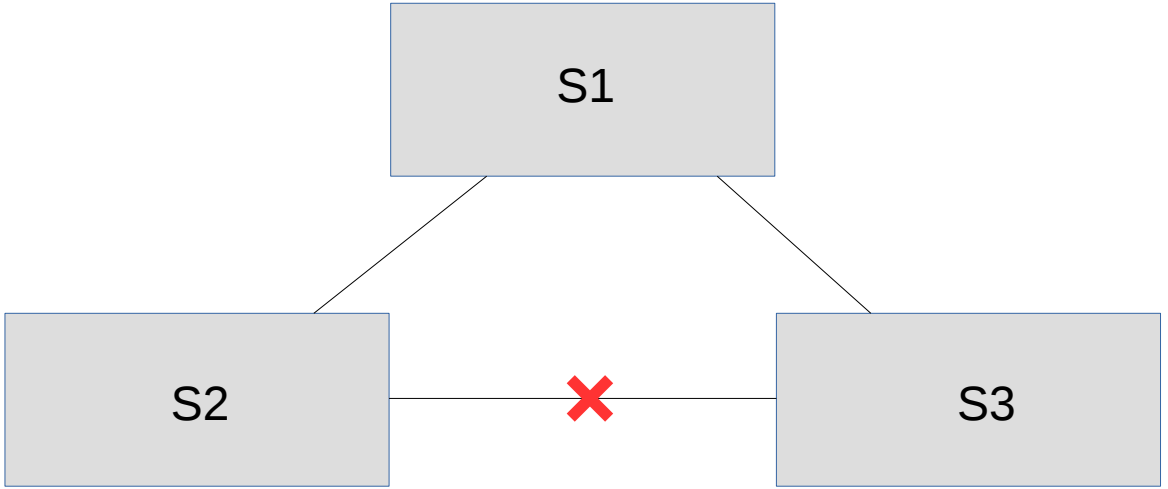

Link Layer - Ethernet

- Having separate physical networks for each logical separation is painful
- VLANs added to allow for multiple logical networks to share a common physical network
  - Mainly used on links between switches
  - Virtualization has brought it to the server level
- Uses special Ethertype, 0x8100, followed by a 2 byte VLAN header
  - 12 bits for VLAN Identifier, 3 bit QOS value, and a flag

Link Layer - Ethernet

- Sometimes more bandwidth is needed than a single link has
- Link Redundancy is also desired
- Both solved with 802.3ad and Link Aggregation Control Protocol
  - Multiple links are combined into one logical link
  - Packets of a similar nature are sent down the same link to prevent packets being delivered out of order
  - Single Session isn’t going to go faster than the speed of 1 link

Link Layer - Ethernet

- Need to determine Ethernet address for a given IP address
- Address Resolution Protocol (ARP) - IPv4
  - Send an Ethernet and IP Broadcast Packet
    - Specify what IP address you need to lookup
    - Specify who to respond to
  - All other nodes receive and the host with the address replies
- Neighbor Discovery Protocol (NDP) – IPv6
  - Similar but uses Multicast
  - Combined with Router Advertisements and other features

TCP and UDP

- IP Addresses are supposed to identify hosts on a network
  - What exactly constitutes a host? Hrm.
- TCP and UDP have separate identifiers, called port numbers, that allow for communicating with an individual service on that host
  - Frequently identified in the form of tcp/80 or udp/123
  - Assignments done by the Internet Assigned Numbers Authority but are treated more like guidelines

UDP

- Header includes Source Port, Destination Port, Length, and maybe a Checksum
  - Checksum is optional in IPv4 but required in IPv6
- The rest is up to you

I’d tell you a UDP joke but you might not get it

TCP

- Header has a lot of fields including Sequence Number and Acknowledgement Number
  - Allows for each side to indicate how much data has been sent and how much has been received
- 3-way Handshake allows for initial connection set up
  - Client sends packet with SYN flag set
  - Server responds with SYN and ACK flags set
  - Client responds with ACK flag set
- FIN and ACK used to gracefully close a connection
- RST used to half close a connection or be a jerk

ICMP

- Internet Control Message Protocol
- Allows for exchanging messages to indicate various aspects of the network
- Includes a Type and Code field
  - Type 0: Echo Reply
  - Type 8: Echo Request
  - Type 11: TTL Exceeded
  - Type 3: Unreachable, see code for more details
    - Code 4: Fragmentation Required
    - Code 10: Host Administratively Prohibited

Routing

- Only devices on a given Link and Subnet are assumed to be able to talk directly with each other
- Talking between Subnets requires one or more intermediaries, called Routers
- Have at least one Address in each Subnet they’re connected to
- May or may not modify packet
  - Network Address Translation

Routing

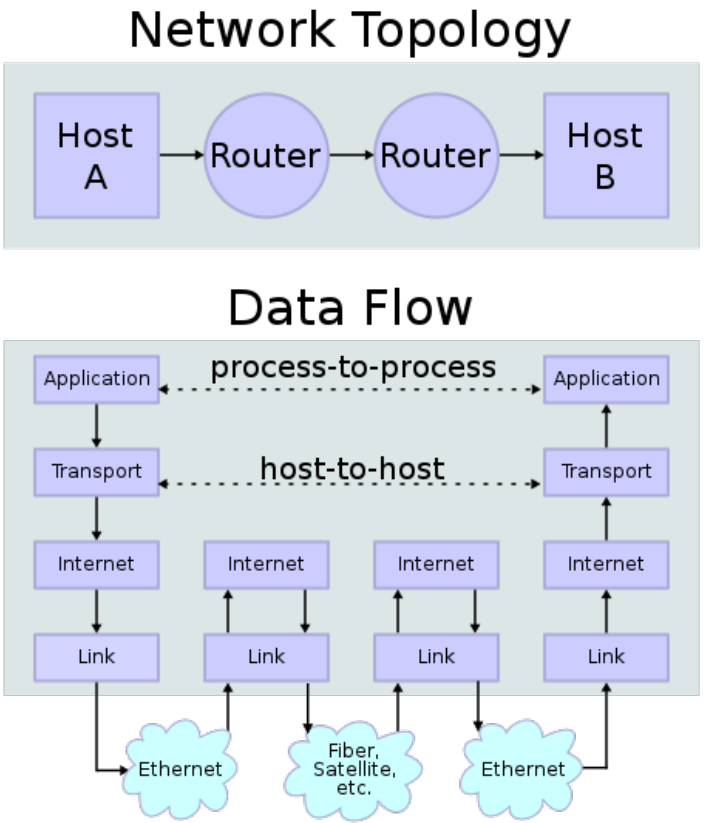

Credit: Wikipedia User Kbrose
License: CC BY-SA
https://en.wikipedia.org/wiki/Internet_protocol_suite
https://commons.wikimedia.org/wiki/File:IP_stack_connections.svg

Firewalls

- Used to apply policy on traffic
- Zone-based and/or Interface-based
- Stateless
  - Based on static rules about particular fields in a packet
  - i.e. Packets from 192.168.1.0/24 to 0.0.0.0/0 are allowed
- Stateful
  - Examines packets in relation to each other and the context in which they were received
  - i.e. After a SYN from 192.168.1.5 to 10.1.2.10, allow packets from 10.1.2.10 to 192.168.1.5

Firewalls

NAT itself is not a security measure

Troubleshooting Tools

- ping and ping6
  - Checks if an ICMP packet make to back and forth between hosts
- traceroute and traceroute6
  - Looks for the route between hosts
- tcpdump
  - Displays traffic crossing a network interface
- iproute2 Tools
  - Tool set for showing and setting network configuration on recent Linux distributions

Troubleshooting Tools

- Basic Questions
  - Is the network interface up?
  - Are the IP settings correct?
  - Can the gateway be reached?
- Commands to use
  - ‘ip address’
  - ‘ip route get #.#.#.#’
  - ‘ip neighbor’
  - ping

Troubleshooting Tools

- What is the Source and Destination Address of the packet?
- What is the Source and Destination Ports of the packet?
  - Network Engineers are gonna want at least both sets of these
- Can hosts talk to a third party?
- Can the source talk to other hosts on the same Subnet as the destination?
- Try `traceroute` and `traceroute -I`
  - Using -I requires escalated permissions but uses ICMP packets vs UDP packets with TTLs