Introduction to Networking

Chendi Luo

Housekeeping

- Friday: Halloween special topics: Haunted Smart Homes
- Sunday: Connections Museum Trip
- Next Wednesday: XSS Exploits

whoami

- Chendi Luo
- Sophomore in CS/Art
- Resident NixOS evangelist

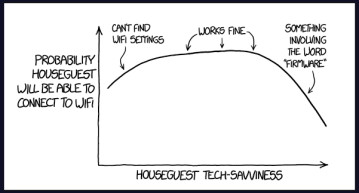

https://xkcd.com/1785/

What is Networking?

- I have multiple devices
- I want them to share data
- They communicate, forming a network
- `ping google.com`

A Practical Example

- I have a computer
- I enter `google.com` in the browser of my computer
- The browser makes an HTTP GET request to `google.com`
- `google.com` sends back a webpage

How did this happen?

- How did we find `google.com`?
- How did we reach the internet?
- How does the connection work?

How did we find google.com?

URLs (Universal Resource Locator)

`https://www.google.com/`

- Protocol: `https`
- Host: `www.google.com`
  - Top Level Domain: `com`
  - Domain: `google`
  - Subdomain: `www`
- Resource Path: `/`
- Browsers fill in the protocol and path for you
- Networks don't understand URLs

IP (Internet Protocol) Addresses

- Unique address for a device
- Network address followed by the host address
- IPv4: 4 byte address, formatted as decimal numbers
  - e.g. `192.168.0.0`
- IPv6: 16 byte address, formatted as hex
  - e.g `2001:0db8:0000:0000:0000:8a2e:0370:7334`
- Can find an IP using `nslookup`

Hosts File

- Mapping of names to IP addresses
- `/etc/hosts` on Linux
- Can be manually configured by the user
- Typically defaults to just localhost

Domain Name System (DNS)

- How domain names are organized globally
- DNS server stores mappings of domain names to IP addresses
- Devices usually preconfigured with a DNS server's IP e.g. `1.1.1.1`
- Also get a DNS when you get IP from the router

DNS Resolution

- Goal is to find the Authoritative DNS server for `google.com`
- Two kinds of requests
  - Recursive: Send requests on my behalf and just give me the final answer
  - Iterative: Just give me the best you have and I'll figure it out
- Device sends a recursive DNS query to a DNS server for `google.com`
- DNS server sends iterative DNS queries to increasingly lower level servers, then returns the final IP address
  - Query root to get the top level domain server, query TLD server to get the domain server, etc

How did we reach the internet?

Local/Wide Area Networks (LAN/WAN)

- Device doesn't connect to "the internet"
- Device connects to a LAN with a router
- Router connects your device to other networks
- WAN used for bigger networks like cellular data, works similarly

WiFi and Ethernet

- How the device connects to the router
- WiFi uses radio waves to communicate
  - Device spams out probe requests to find a network, then waits for replies
- Ethernet uses physical cables to communicate
  - Servers typically have an intermediary layer of a network switch to manage all the ethernet cables and do additional routing

MAC Addresses

- Unique address for a network card assigned by the manufacturer
- Used for local identification within a network since IP addresses change

(you should go enable MAC randomization)

What's my IP?

- Your device has some IP assigned to it by the router using DHCP
  - ...which got a range of IPs from the internet provider
  - ...which got a range of IPs from the regional internet registry
  - ...which got a range of IPs from the Internet Assigned Numbers Authority
- `ifconfig`
- https://www.showmyip.com/

Why are these different?

Subnetting

- Bitmask to show which bits are the network and which are the host
  - Can also be used to describe a network or ip range, rather than specific address
- Historical artifact of classful networking which had a broader way of breaking down the address
- Abbreviated as /x, where x is the number of bits
- Does not allow for more/repeated addresses, just splitting up into more networks

Private IPs

- Special set of IP addresses reserved for private use (i.e. not on the internet)
  - `10.0.0.0/8`
  - `172.16.0.0/12`
  - `192.168.0.0/16`
- Does not include localhost

Dynamic Host Configuration Protocol (DHCP)

- Protocol for dynamically allocating IPs
- Device spams out a discover message to find server
- DHCP server allocates an IP address for a certain amount of time
- On a typical network, router assigns each device a private IP
- Router sends requests using its own public IP, then forwards the response back to the private IP

Virtual Private Networks (VPNs)

- Private network (a local network using private IPs) which is extended virtually (through the internet)
- Often used for large organizations to provide offsite access to internal resources (e.g. Husky OnNet)
- VPN companies are primarily proxies which may happen to work using VPNs
  - "Change your IP" by sending your traffic through another server as sender, then forwarding it back to you

How does the connection work?

Routing

- You have a destination address, but not the route to get there
- Each router knows about its neighbors, but not the full network
- Routing protocols are used to build a routing table and determine where to send traffic
  - Interior gateway protocols (IGP) communicate within a single autonomous system
  - Exterior gateway protocols (EGP) communicate between autonomous systems
- `traceroute`

Autonomous System (AS)

- Collection of IP ranges managed by a single organization
- Uniquely numerically assigned similar to IP addresses
- Your router is assigned as part of an AS run by your ISP along with its IP address

Border Gateway Protocol (BGP)

- Current EGP used to communicate between ASs
- Also used for interior routing, but as a way to communicate external routes to others in the same network

Ports

- Additional specifier for network traffic to distinguish between processes at the same address/device
- Standard ports for different services
  - HTTP 80, HTTPS 443, SSH 22, etc
- Sender will use a random port to receive and send a request to a specific port
- Port forwarding: telling the router to send specific traffic to a local port
- Can scan open ports using `nmap`

Endpoints/Sockets

- Combination of a specific address and port
- Each endpoint has only one listener, but can have multiple connections
- `netstat -anutlp` (all, numeric, udp, tcp, listening, process)
- A connection is identified by two endpoints and the protocol

UDP and TCP

- Both protocols for sending data packets
- UDP
  - Lightweight, minimal error handling, very small headers
  - Ignores data loss
  - Used for low latency applications like voice calls where dropped data is better than lag or only simple updates are needed
- TCP
  - Slower, error handling and retransmission
  - Guarantees sequential requests, requires handshake requests before sending data
  - Used for higher reliability applications where data loss is not acceptable

Transport Layer Security (TLS)

- Layer of encryption used over a connection
- HTTPS is just HTTP data encrypted using TLS

Firewalls

- Set of rules for describing what kinds of traffic to allow
- May be based on IP addresses, ports, protocols
- Commonly used for basic security across a network

Networking, Formally

Networking Models

Arpanet Reference Model (RFC 871)	Internet Standard (RFC 1122)	Internet model (Cisco Academy[58])	TCP/IP 5-layer reference model (Kozierok,[59] Comer[60])	TCP/IP 5-layer reference model (Tanenbaum[61])	TCP/IP protocol suite or Five-layer Internet model (Forouzan,[62] Kurose[63])	TCP/IP model (Stallings[64])	OSI model (ISO/IEC 7498-1:1994[65])
Three layers	*Four layers*	*Four layers*	*Four+one layers*	*Five layers*	*Five layers*	*Five layers*	*Seven layers*
Application/ Process	Application	Application	Application	Application	Application	Application	Application
							Presentation
							Session
Host-to-host	Transport	Transport	Transport	Transport	Transport	Host-to-host or transport	Transport
	Internet	Internetwork	Internet	Internet	Network	Internet	Network
Network interface	Link	Network interface	Data link (Network interface)	Data link	Data link	Network access	Data link
—	—	—	(Hardware)	Physical	Physical	Physical	Physical

https://en.wikipedia.org/wiki/Internet_protocol_suite#Layering_evolution_and_representations_in_the_literature

Networking Models

- Two main models, OSI and TCP/IP
- Both are descriptive, not prescriptive
  - TCP/IP is used to organize the Internet protocol suite standards, but is not a prescriptive standard
- Both use the concept of data passing through layers which do different operations
  - Data goes down as it is sent, then up as it is received
- Layers are not strictly separated, but are helpful as a reference point

Open Systems Interconnection (OSI) Model

- Application: Interface with the user
- Presentation: Encoding data
- Session: Managing the connection between devices
- Transport: Sending data between devices
- Network: Routing data between different networks
- Data Link: Moving data locally
- Physical: Physical transmission of bits

TCP/IP Model

- Application: End user services
- Transport: Sends data between devices
- Internet: Routes packets
- Network Access: Sends packets physically

Questions?