

ECE 435 – Network Engineering

Lecture 14

Vince Weaver

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

`vincent.weaver@maine.edu`

2 March 2023

Announcements

- HW#6 due Friday
- Will try to get HW#5 and HW#6 graded before the midterm
- HW#7 will be posted
- In 2021 Pentagon activated some of its vast IPv4 collection turns out had been unused people using them as unroutable numbers, including China military.

https://www.theregister.com/2021/04/26/defense_department_ipv6/



Midterm Preview

- Can have one page (8.5" x 11") of notes if you want, otherwise closed everything. I do not think you should need a calculator.
- Mostly short answer questions. No long coding exercises or protocol memorization.
- There might be some sockets code, but analyzing it not writing it.
- Know the OSI layers and what each one is for.
- Know at a high level the following protocols:



- WWW/http
- e-mail
- DNS
- Encryption (at a high level)
- UDP + TCP
 - Know the 3-way handshake
 - Know the tradeoffs between UDP and TCP
 - Why does DNS use UDP
 - Why does HTTP1.1 use TCP



Brief HW#5 Review

- source/destination/size/checksum
 - src: a9a0 = 43424
 - dest: 35 = 53 (DNS)
 - size: 2a = 42 bytes
 - yes checksum
 - protocol is DNS (how can you tell?)
- Why use UDP vs TCP
 - lower latency, lower overhead (no need to handshake), simpler



Be careful just saying “faster”, need to explain more what you mean by that.




HW#5 Coding Notes

- Remember to comment your code!
- Getting source port from incoming connection
- Note this is not the IP address
- Getting it from the struct is sort of hard
- Also remember it's in network endian, need to convert with `ntohs()` In general would be an ephemeral port above 40000



The IPv4 Catastrophe



Out of IPv4 Addresses Problem

- IPv4 address exhaustion
- CIDR not enough
- Addresses managed by IANA globally and five regional registrars (RIR)
- Top level ran out in 2011
- All 5 RIRs finally ran out on Nov 25th, 2019
- Why are we out?
 - Always active connections – unlike dialup, many machine are on all the time



- So many devices – 4G mobile devices all have one
- Inefficiencies originally handing out. Companies like Apple, MIT, DEC, all got 16 million address Class A addresses even if didn't need them
- Despite being out, in 2011 reportedly only 14% of addresses being used
- Why not reclaim unused, such as Class E? The bane of network programmers, the out-of-date router that makes assumptions
- Stanford gave back a class A in 2000



Network Address Translation (NAT)


- Private IP ranges, defined in RFC 1918
 - 1 Class A: 10.0.0.0 - 10.255.255.255 (10.0.0.0/8)
 - 16 Class B: 172.16.0.0 - 172.31.255.255 (172.16.0.0/12)
 - 256 Class C: 192.168.0.0 - 192.168.255.255 (192.168.0.0/16)
- Can use for various reasons, most recently due to network depletion
- NAT: map IP addresses from one group to another. often public to private.
- NAT and NAPT (port translation) RFC 3022



- Basic NAT has one to one mapping of external to internal IPs. Each internal host maps to unique external IP



NAT Example



Network Address Port Translation (NAPT)

- NAPT: based on port, only one external IP
 - Full cone – most common
 - once an internal address (iaddr/port) has been mapped to an external (eaddr/port) all packets from iaddr/port are sent out and any incoming are passed through with no additional checks
 - Restricted cone – same as above, but only external that have received packets from internal can send through
 - Port restricted cone – same as above, but also checks



port numbers

- Symmetric – best security – outgoing packets mapped to different eaddr/port if the destination or port differs
- When passing through, NAT needs to re-write dest/source/port and recompute header checksum
- Linux: IP-masquerade/iptables



Many IP people hate NAT

- Violates the IP identifies one machine rule
- Hard to connect two machines if both behind different NATs (NAT transversal)
- Changes IP to be connection oriented, router has to remember connections
- Layering violation, looks at TCP/UDP port numbers
- Only works for TCP/UDP
- Some protocols (like FTP) are even more annoying, send address in plain text in data and that has to be adjusted



too

- Can only NAT up to 64k machines (why? how many ports are there?)



Carrier Grade NAT (CGN, CGNAT, LSN)

- Internal network uses private IP range
- Public facing server channels these through a set of external IP addresses
- NAT444 – potentially traverse 4 different IP (private in home, private in ISP, external IP)
- RFC6598 – allocate 100.64.0.0/10 for this, to avoid complications where internal/external collisions of the RFC1918 ranges
- Downsides



-
- Breaks end-to-end connections
- Stateful
- Doesn't fully solve IPv4 exhaustion problem cases you need a visible IP address (SSL web server?)
- Lots of devices behind a few IPs, what if get banned for spamming/security?
- Breaks port-forwarding for users, as you're in a NAT inside a NAT (port control protocol (PCP) RFC 6887 tries to work around this)



The Internet Protocol v6

- RFC2460 - RFC466
- Started work in 1991
- Many problems with IPv4. Most notable shortage of addresses.
- IPng. (IPv5 was an experimental stream protocol)
- Migration happening, a large amount of web traffic, especially that from phones, is already switched.
- **not** backwards compatible



IPv6 uptake

- As of July 2016 12.5% of traffic is IPv6
- According to Google, March 2022 around 34%
- March 2023, 38%
- <https://www.google.com/intl/en/ipv6/statistics.html>



The Internet Protocol v6 Goals

- Support billions of hosts
- Reduce size of routing tables
- Simplify the protocol (so routers can be faster)
- Better security
- Pay more attention to type of service
- Aid multicasting
- Allow roaming w/o changing address
- Co-exist with existing protocols



The Internet Protocol v6 features

- Address size 128 bits
 - a lot of addresses. 7×10^{23} for ever square meter
- Simpler fixed length header (speeds up processing)
 - Many fields not really used in IPv4 dropped (or made optional)
- Better support for options
- Better security support
 - IPSEC. Originally mandatory, made optional
 - Can encrypt packets at the network layer



- Quality of service (???)
- Anycast (see end of slides)
- Autoconfiguration (like DHCP)
- Minimum fragment size 1280 (up from 576)
- No checksum – was slow and recalculated often




IPv6 header

- Header fixed length of 40 bytes, with Extension headers
- ASCII art from RFC 2460

```
+++++  
|Version| Traffic Class |           Flow Label           |  
+++++  
|           Payload Length           | Next Header | Hop Limit |  
+++++  
|  
+  
|  
+           Source Address           +  
|  
+  
|  
+++++  
|  
+  
|
```





- **Version Number** (1 nibble [4-bits]) = 6
- **Differential Services/Traffic Class** (8-bits) (QoS/congestion control) (6-bit differential services, 2 bits ECN (sort of like recent IPv4?))
- **Flow label** (20-bits) (for streaming?)
- **Payload Length** (16-bits) 64k (header bytes not counted anymore) (if you want longer, extension for Jumbograms (up to 4GB))



- **Next header** (8-bits) If nothing special identifies TCP or UDP If special options (fragmentation, security) indicated
(TCP=0x6, UDP=0x11)
- **Hop Limit** (8-bits) TTL, big debate about whether 8-bits was enough
- **Source Address** (128-bits)
- **Destination Address** (128-bits)
- Why not 64-bit addresses?
- No checksum, link or transport catches issues
What does this mean for UDP?



IPv6 addresses

- 2^{128} is a lot. 7×10^{23} per m^2 of Earth surface
 - Too long for dotted decimal, use colon hexadecimal
 - Why colons? .BE is a valid domain ending for one...
 - X:X:X:X:X:X:X:X where X is 16 bit chunk
 - F000:0123:5678:0000:0000:ABCD:0001:CAFE
 - Can drop leading zeros, as well as groups of zeros
F000:123:5678::ABCD:1:CAFE
- Note, cannot drop two sets of groups of zeros. Why?
Ambiguous.



- Do not have classes RFC 4291
- encoding ipv4 addresses in ipv6?
- autoconfig – generate a unique ID, often based on MAC address and send it to router (Router solicitation). Router sends back router advertisement which has subnet prefix and can be used to generate global address Stateless address auto-config (SLAAC). Can also use DHCPv6
- ::1 ip6-localhost, fe00::0 ip6-localnet



IPv6 Options

- Happen immediately after the header.
- Should occur in numerical order (though routers might be able to handle if they don't)
- Routers should inspect in order as some later might depend on earlier.
- Plain: IPV6:Next=TCP, TCP, Data
- Example: IPV6:Next=Routing, Routing:Next=TCP, TCP, Data
- Various types



- Hop-by-hop (so far only used for jumbo frames, if that set header length is set to 0)
- Source Routing
- Fragmentation
- Authentication
- Encryption




IPv6 fragmentation

- Info is in an extension header
- Routers cannot fragment, only at source
- How can this work when not know MTU?
- MTU is always greater than 1280
- Path MTU discovery protocol to discover MTU along the way (RFC 1981). (IPv4 too, set DNF and get error via ICMP) If too big, sends an error back and source needs to fragment it smaller
- Easier to have source fix things then every router along



the way being able to



IPv6/IPv4 compat

- Dual stack – host runs both IPv4 and IPv6
or internal is IPv6 but router converts to IPv4 before passing on
- Tunneling – encapsulate IPv6 inside of IPv4, tunnels across IPv4, split back out to IPv6 on other side of tunnel
- IPv4 mapped to IPv6 with special 96-bit prefix
there are specs for this, not sure if really implemented




IPv6 Routing

- Much like IPv4
- IPv6 network address, “prefix-length” instead of netmask
- Routing table as before



ICMP6



- Checksum: similar to TCP, also includes pseudo-header
- Type 0, Destination Uncertain
 - Various
- Type 1, Packet too big
- Time Exceeded
- Bad Parameters



- ECHO / ECHO Reply
- Neighbor Discovery Protocol (NDP)
- SEND – Secure Neighbor Discovery Protocol
- Multicast Listener Discovery (MLD)



IPv6 Anycast



IPv6 Manycast (Multicast?)

- Not like IPv4 where sending to special address automatically goes to all
- First send to Link Local “all node” address, ff02::1
- Difficult to set up?



IPv6 Stateless Address Auto-Config (SLAAC)

- IPv6 Stateless Address AutoConfiguration (SLAAC) assumes on /64 subnet (so every subnet contains orders of magnitude more than the total IPv4 space for their own local network)
- Essentially large enough a system could just pick a random address and it would work
- Three ways:
 - EUI-64 (RFC 4291) – based on MAC address



- Stable Private (RFC 7217) – hash based, don't give away MAC
- Privacy Extension Addresses (RFC 4941) – like above but change over time to preserve anonymity
- This is link local address on fe80::/10 with the bottom bits as above
- For security refresh daily, this does happen on OSX/Windows, but not necessarily on Linux
- No broadcast, so nodes coming up listen to multicast address



IPv6 Neighbor Discovery

- Neighbor Solicitation (NS) (RFC 4861) use with SLAAC described in RFC 4862 to get address
- Once has address, sends out a packet to see if collision, if not, it's configured
- Once has link local address, sends out Router solicitation (RS) to multicast address ff02::2
- Router replies with (RA) router advertisement packet with info on router, maybe DNS, etc
- Now needs to get global routable address prefix to use,



either directly or has bits set to indicate it should use
DHCPv6



IPv6 DHCPv6

- Can provide info just like IPv4
- Not just router info, but also things like DNS servers, etc



IPv6 Setup

- I've set up many many IPv4 networks, not any IPv6
- <https://lwn.net/Articles/831854/>
Article by James Bottomley
- With IPv4, DHCP can take care of everything



IPv6 setup issues

- It can be hard to subnet.
- It's recommended an ISP gives you a /56 but often they will just give you a /64
- That's a lot of addresses, but due to SLAAC it's assumed a network has a minimum of 2^{64} addresses so you can't split it up easily
- Annoying if you want multiple subnets at home (for wireless, DMZ, etc)
- Setting up Firewall. Linux has separate ipv4 and ipv6



firewalls

- Having a NAT set up sort of gives you a firewall for free, you don't necessarily get that with IPv6



IPv6 Security Issues

- Shadow Networks – if you have a primarily ipv4 setup but various devices start up IPv6 connections without you realizing it
- Fragmentation – even though only on ends, can still have issues like IPv4 where it's hard to handle fragments as TCP port info and such only in first fragment

