Defending Computer Networks
Lecture 6: TCP and Scanning

Stuart Staniford

Adjunct Professor of Computer Science

Logistics

- Supplementary lecture tonight 6pm

Additional Reading

- Fyodor. *The Art of Port Scanning*. http://nmap.org/p51-11.html.
  - You can skim the code section if time pressed.
- Note again you are being pointed at intro papers that are dated.
  - Have to start somewhere.
  - Practical network attack/defense is not a timeless body of knowledge.
  - Constantly evolving arms race between attackers and defenders coming up with new techniques.

Excellus records hacked; 10.5 million records affected

Steve Orr and Patti Singer, @PattiSingerRoc *7:23 a.m. EDT September 10, 2015*

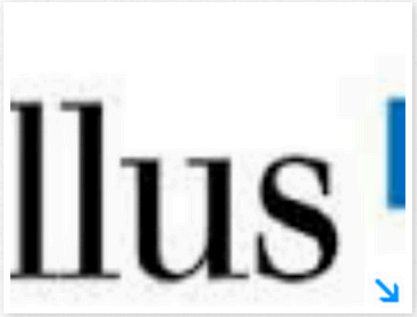

(Photo: Fille image)

Personal information on more than 10 million people, many of them upstate New Yorkers, has been exposed to computer hackers who mounted a sophisticated cyberattack on Rochester-based Excellus BlueCross BlueShield and related companies.

The potential loss of data — which includes names, addresses, telephone numbers, Social Security numbers, financial account information and in some cases sensitive medical information — appears to constitute the biggest known computer hack in local history. It is one of a series of major digital intrusions into Blue Cross affiliates and other health insurers nationwide over the last two years, at least some of which have been tentatively linked to shadowy groups in China.

Main Goals for Today

- Basics of TCP Protocol
- IP Address Space
- Leading to port scanning
  - HW2 will be building a simple TCP parsing application
  - HW3 will build on this to do portscan detection

Refresh on Ether/IP

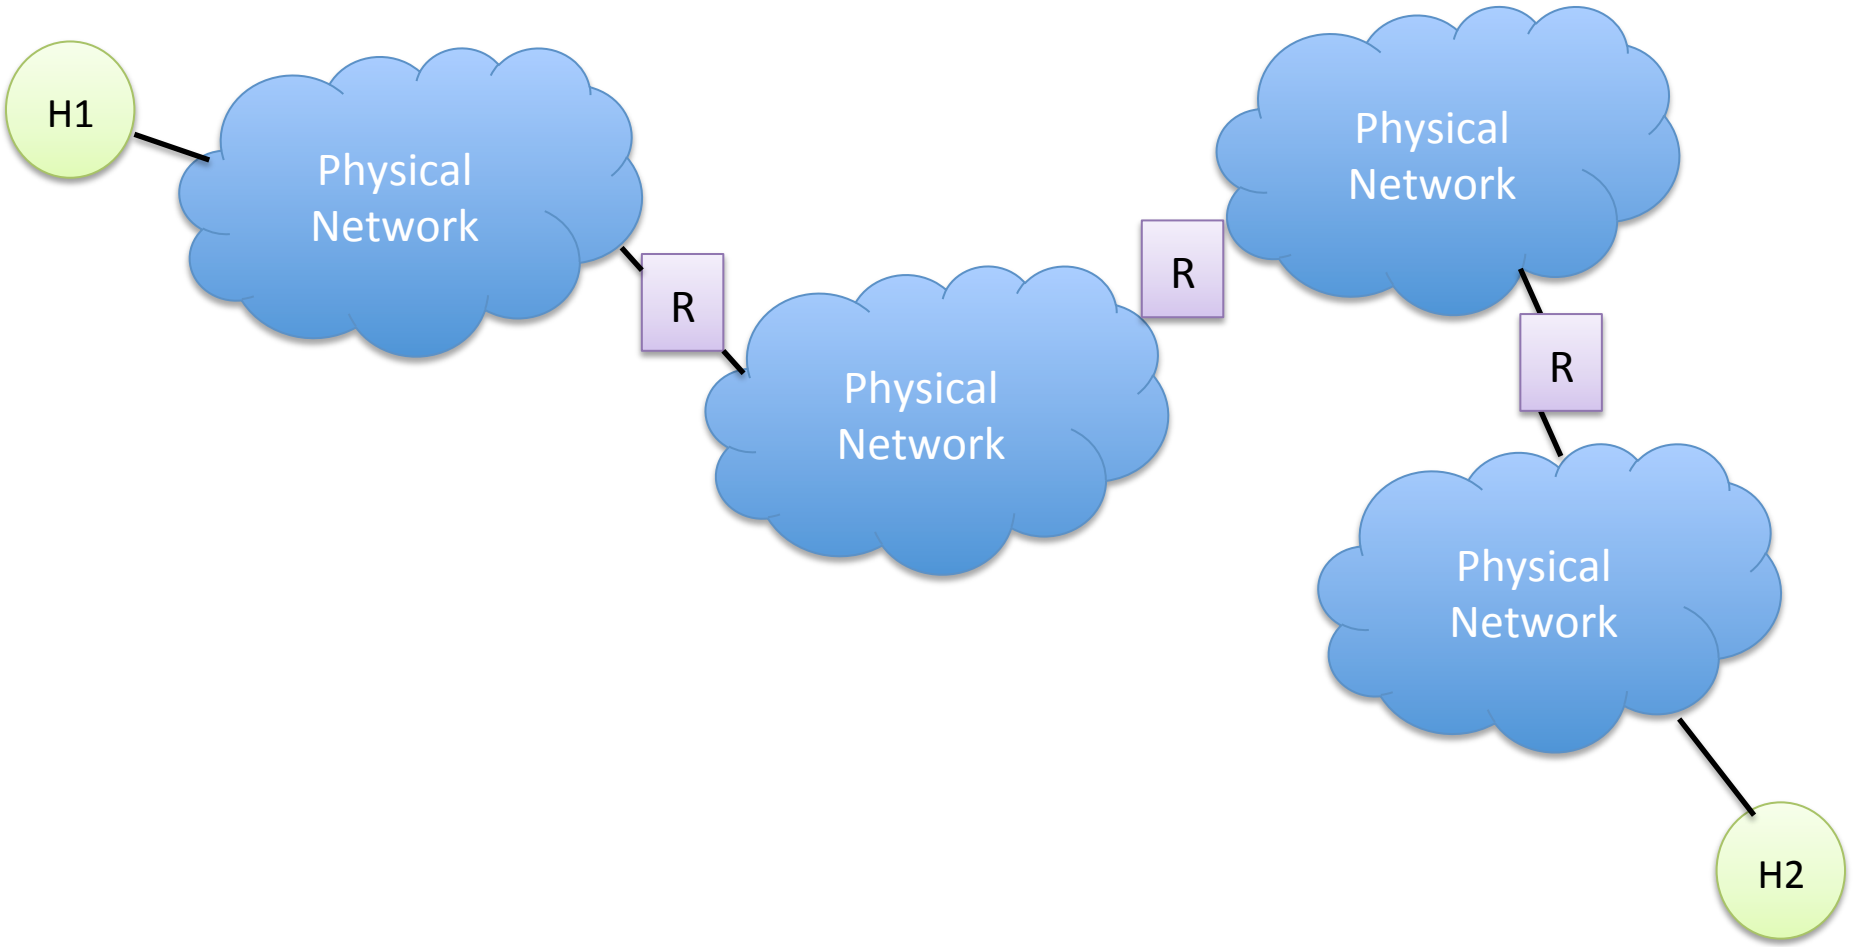

Intro to TCP

- Transmission Control Protocol
- RFC 793 (1981)
- Provides for delivery of stream of data
  - Reliably
  - In order
  - Bi-directionally
  - Between client and server *applications*
    - Not just hosts like IP

Protocol Relationship

- TCP is known as a transport layer protocol
- Goes over the network layer protocol (IP)
  - To provide additional services (reliability, etc)
- Which goes over physical layer (ethernet)
- TCP *segments* are nested inside ip packets
- Nested inside ethernet frames

Ethernet/IP/TCP Nesting

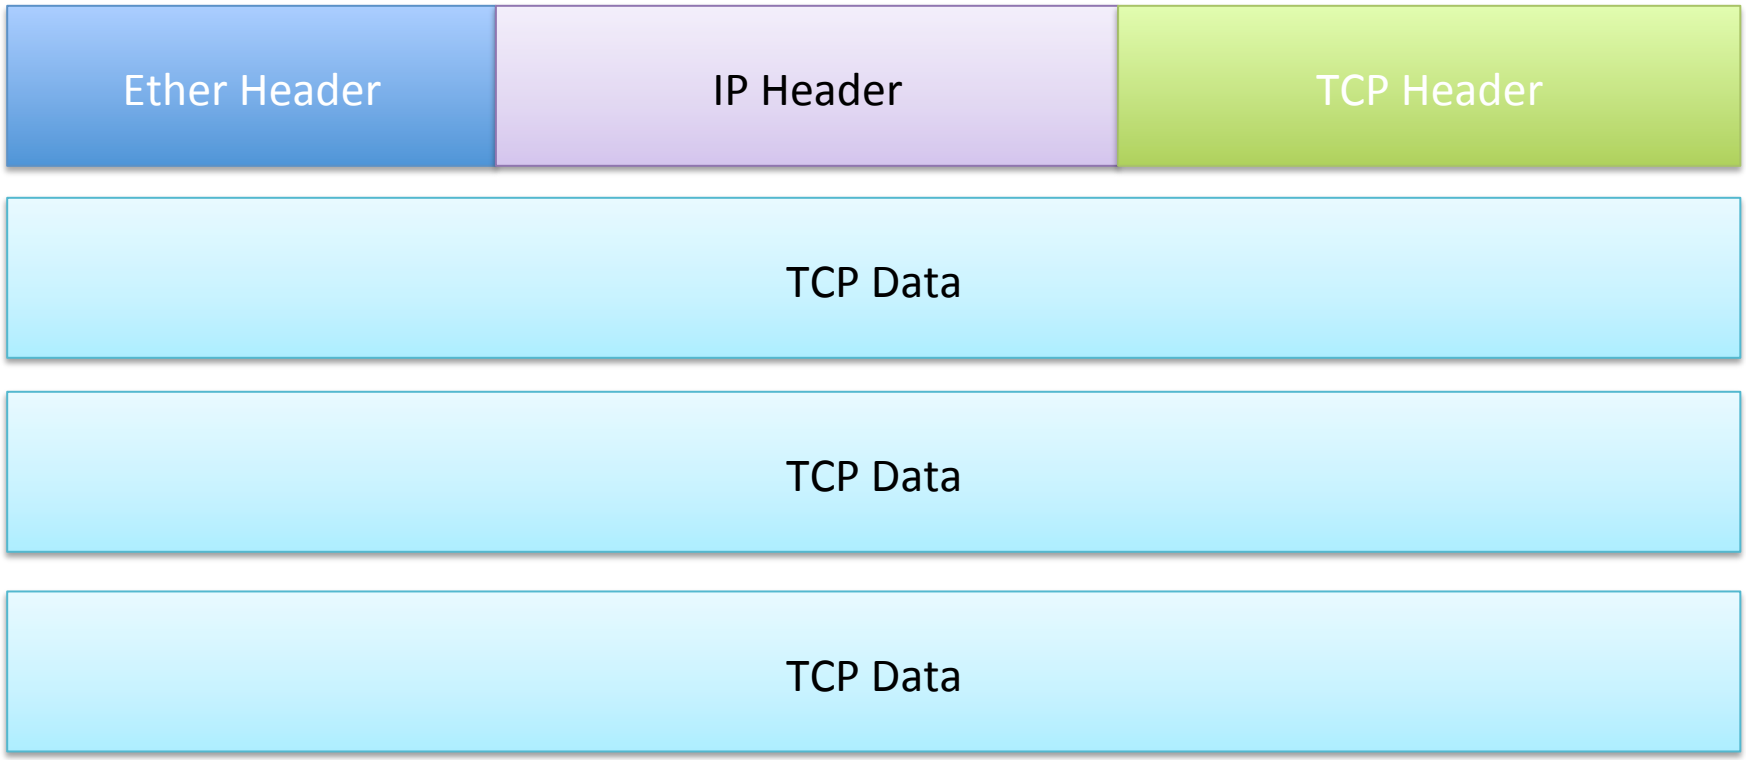

TCP Header Format

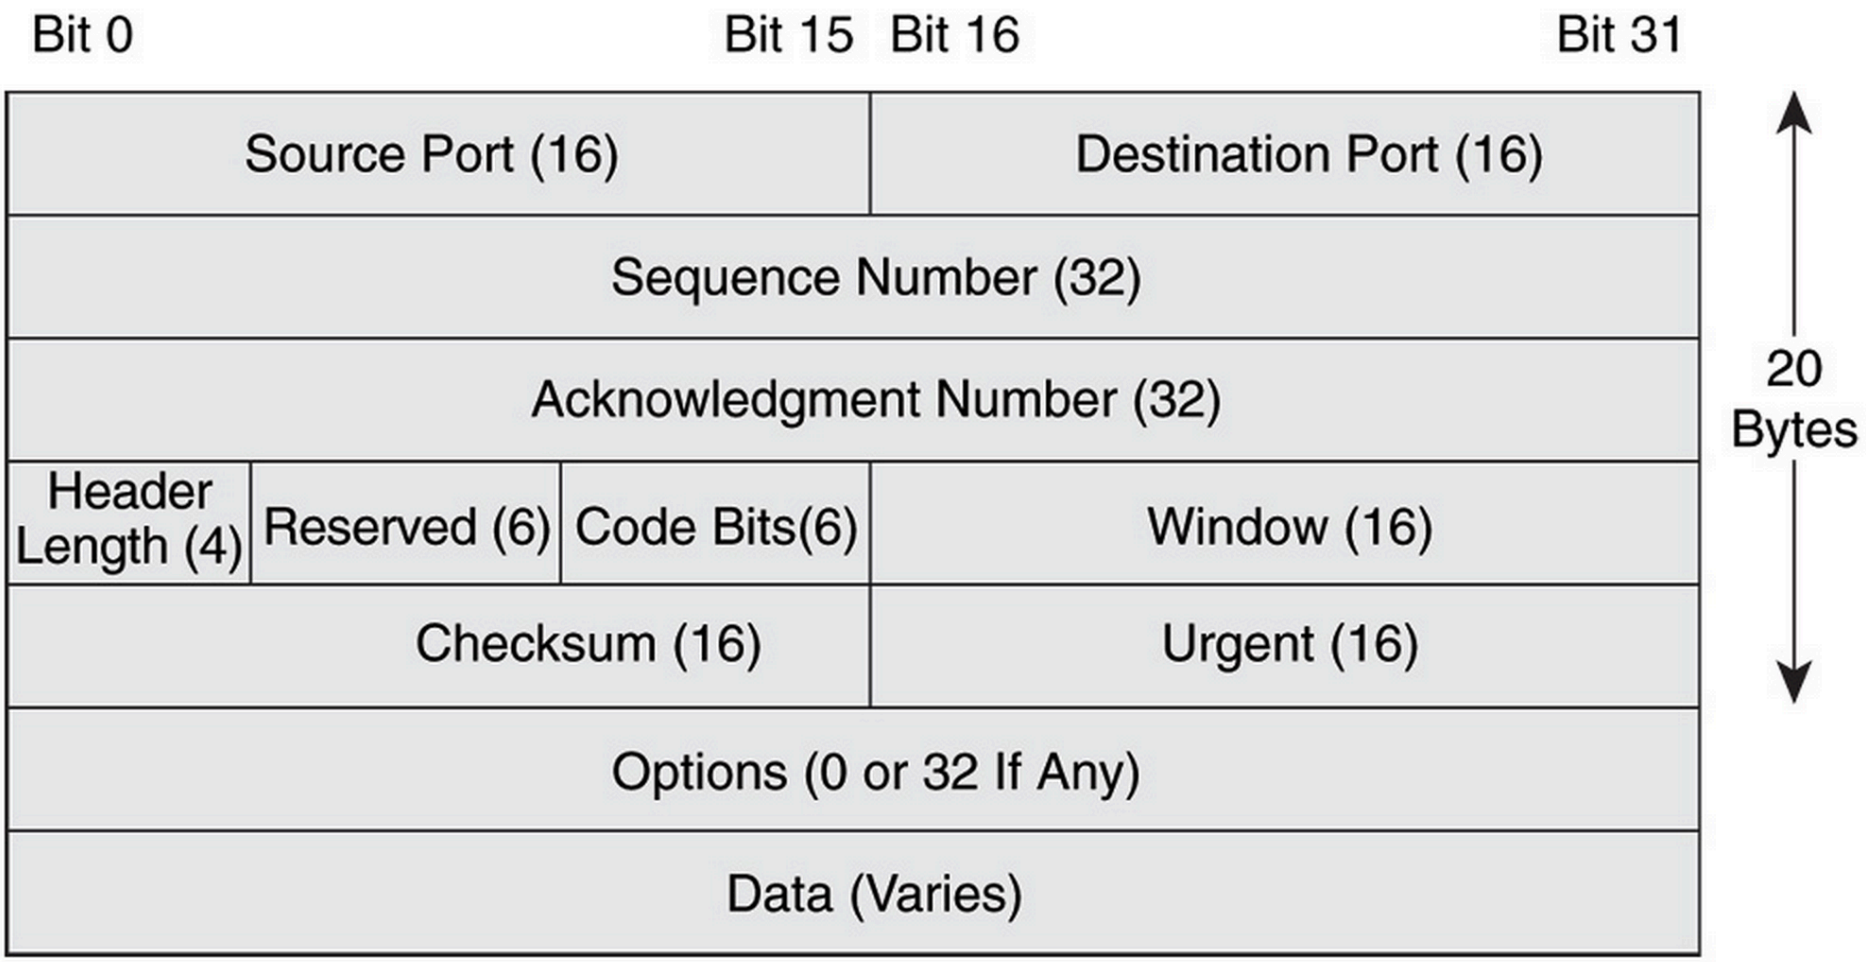

TCP Port Number

- 2 byte quantity (so 65536 possible port #s)
- Server binds to a fixed port
  - Typically "well known" (below 1024):
    - HTTP: 80
    - HTTPS: 443
    - SMTP: 25
    - SSH: 22
- Client typically is assigned a port by OS
  - High numbered
- In some high volume situations port numbers wrap after a while

TCP Connections

- Name for the bidirectional stream
  - Between client application and server application
- Defined by the fivetuple
  - Source IP
  - Source port
  - Dest IP
  - Dest port
  - Protocol

How Reliability/In-Order is done

- Checksums to detect outright transmission error
  - Retransmit if bad
- 32 bit sequence numbers for each byte
  - To detect missing data
  - Retransmit if doesn't show up after a while
- Each segment can
  - Carry some data (indicated by seq number)
  - Acknowledge some data in other direction
    - Ack sequence number

Let's work through an example

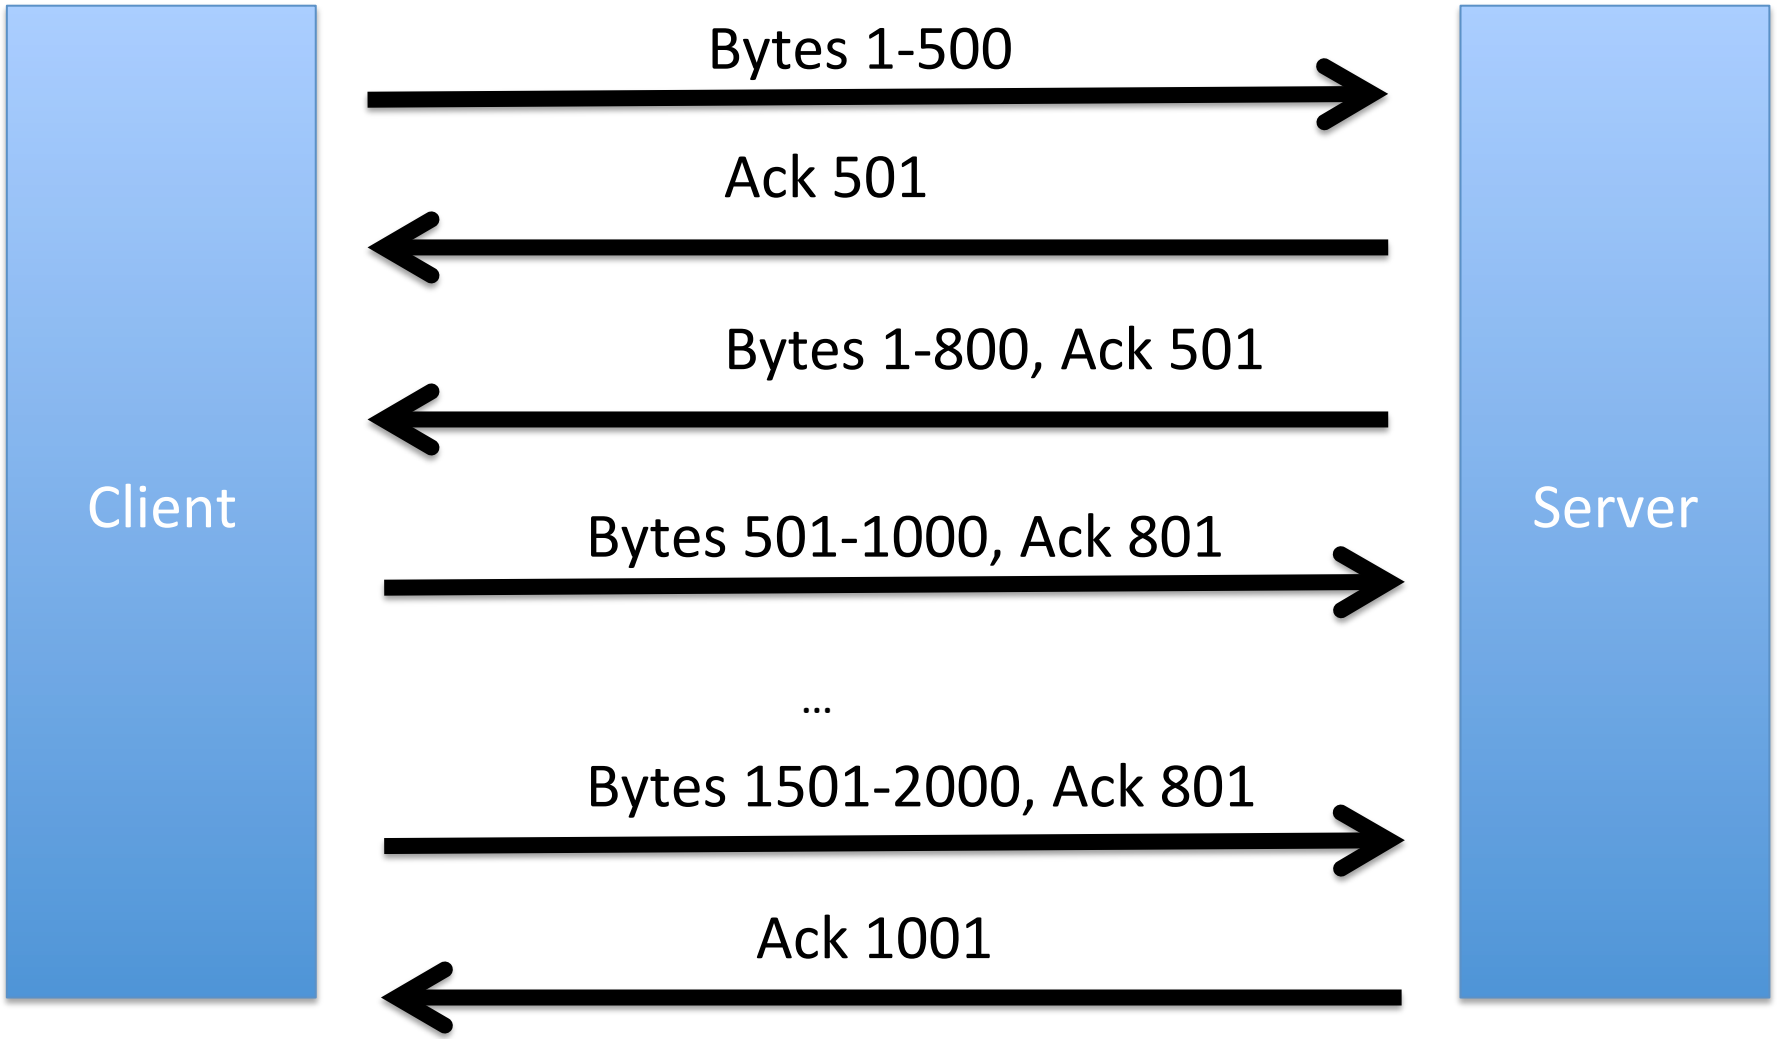

TCP 3-way handshake

- Serves to have client and server agree they are talking
- Also establishes initial sequence numbers
  - In both directions
  - Can't have fixed start (eg zero)
  - Too easy for bad guys if predictable
    - A man in the middle can interfere

Handshake packets

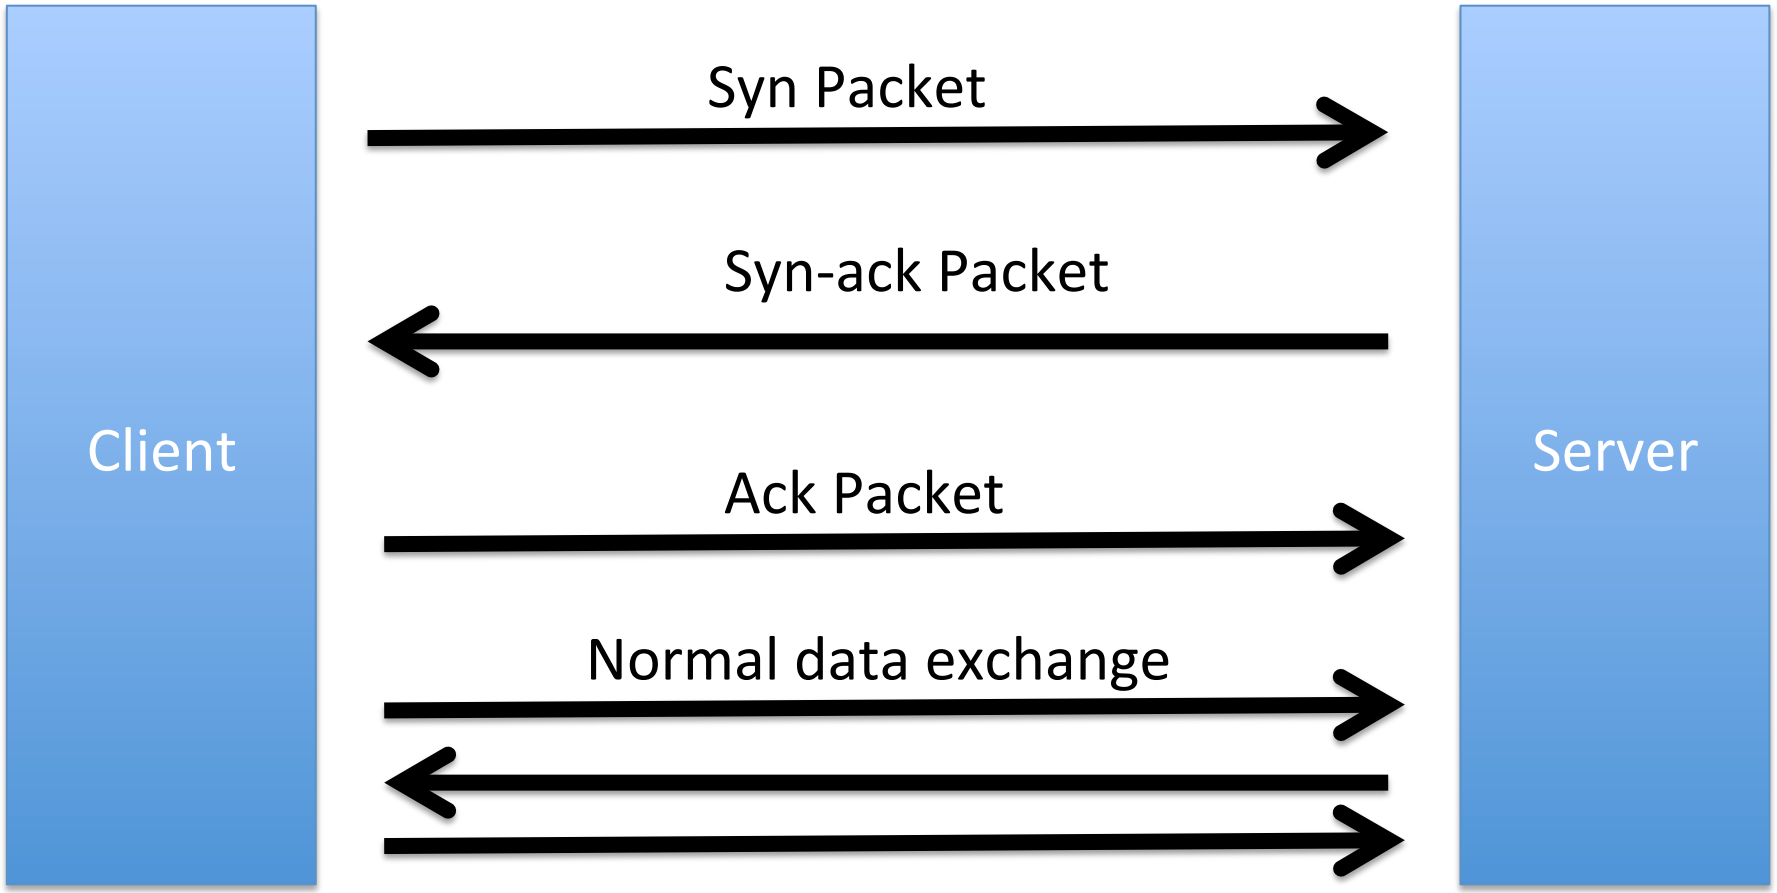

“Syn Packet”

- First packet in handshake
- Makes use of TCP flags byte field in header
  - 0x01 FIN (F)
  - 0x02 SYN (S)
  - 0x04 RST (R)
  - 0x08 PSH (P)
  - 0x10 ACK (A)
  - 0x20 URG (U)
  - 0x40 ECE
  - 0x80 CWR

Syn Packet Layout

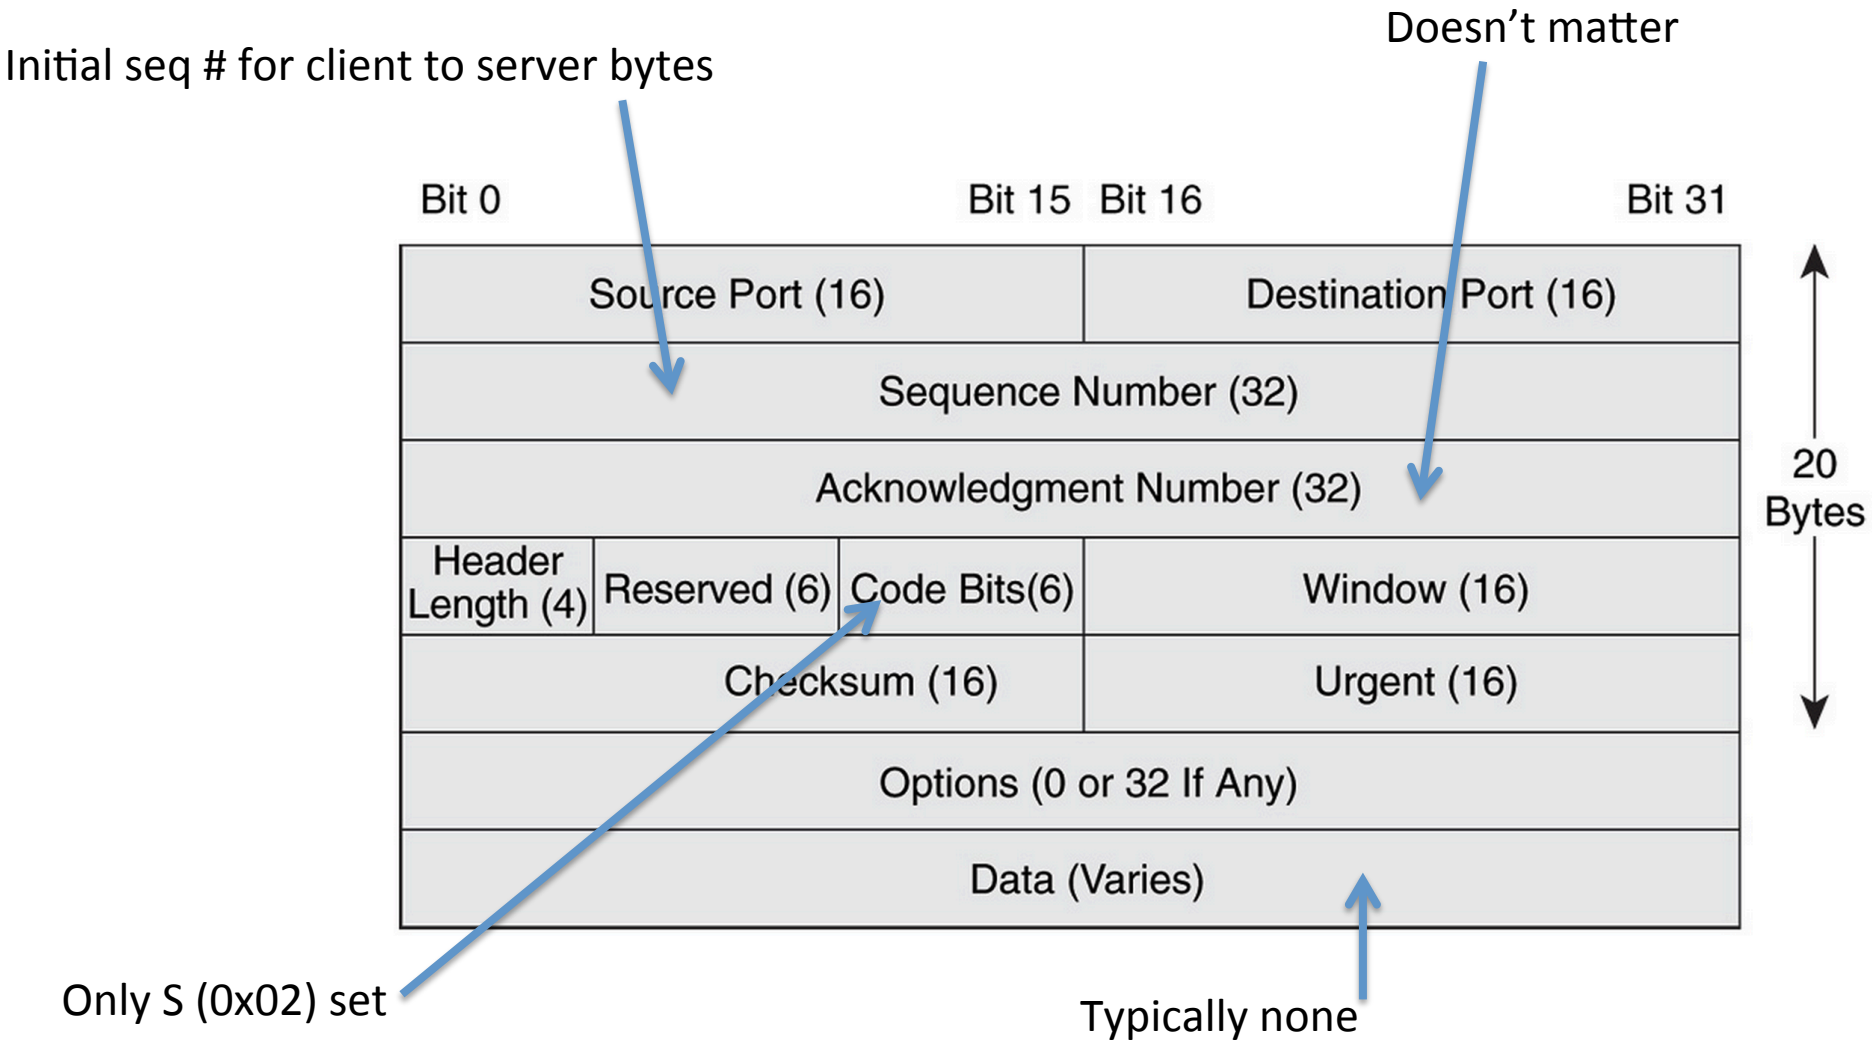

Syn-ack packet

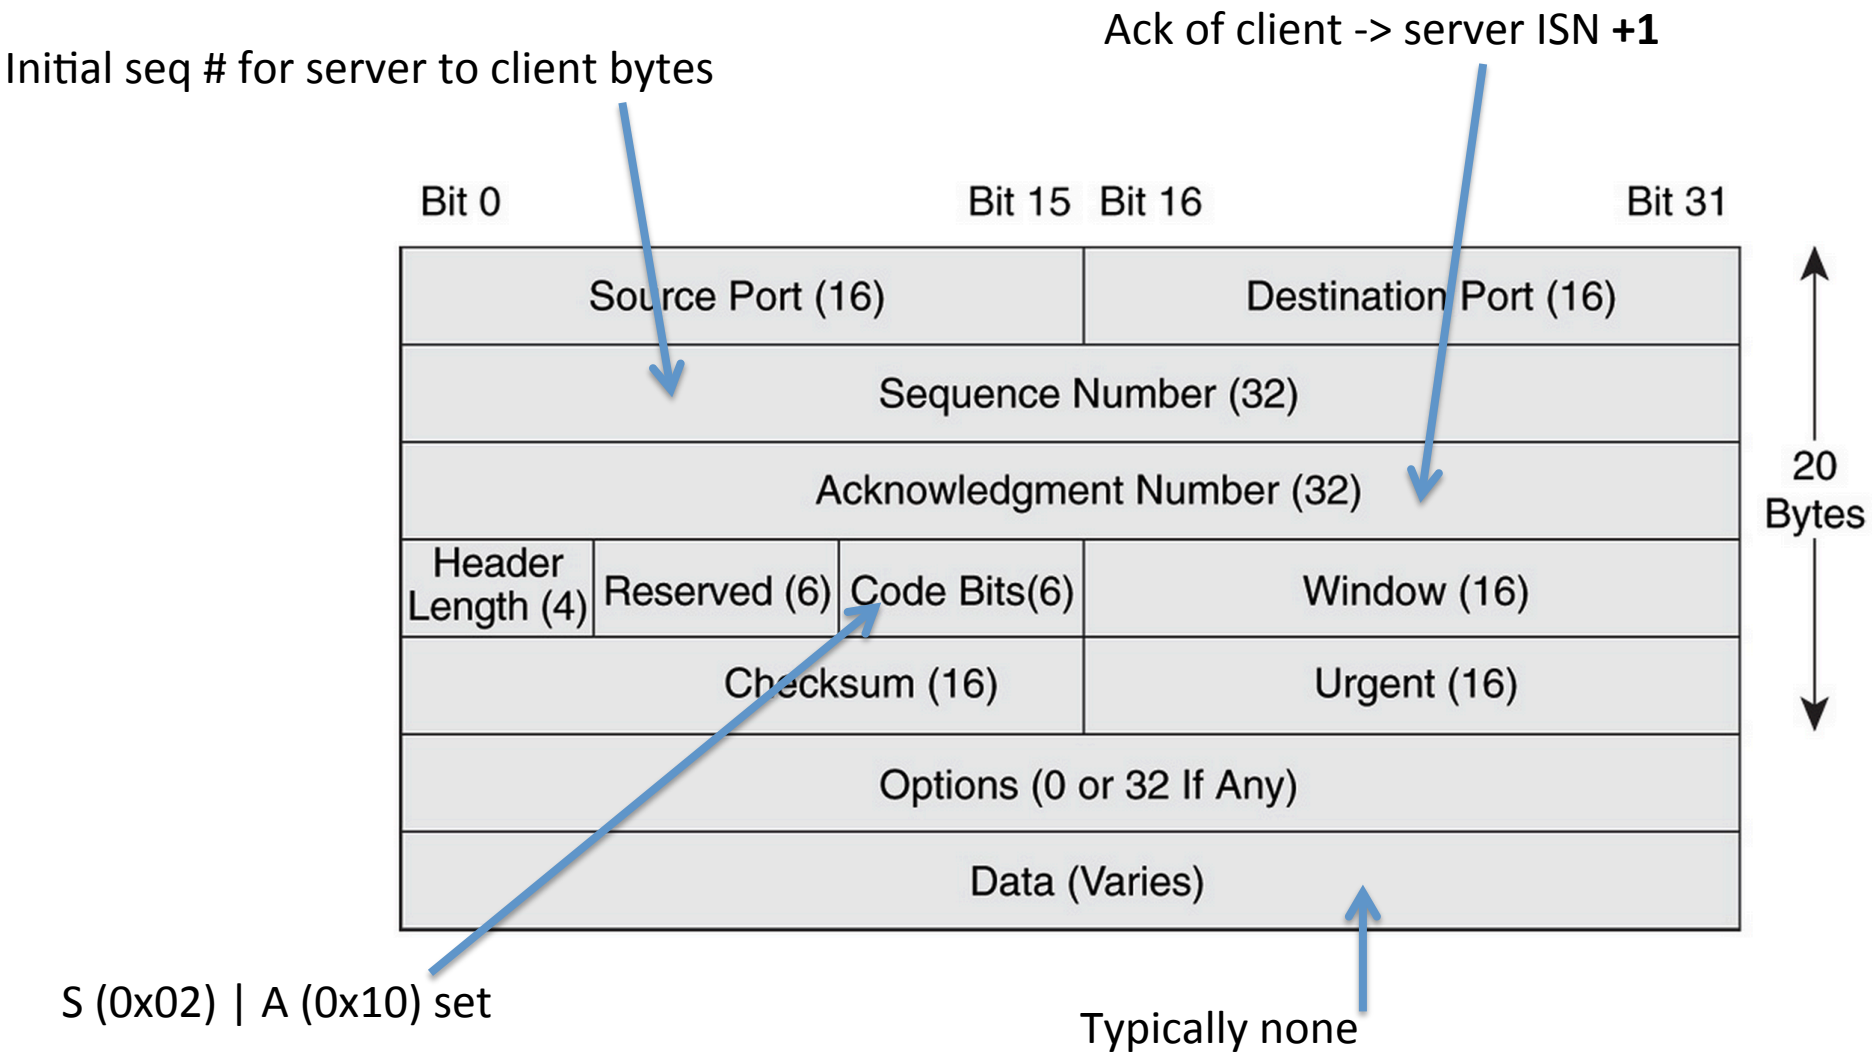

Final Handshake Ack

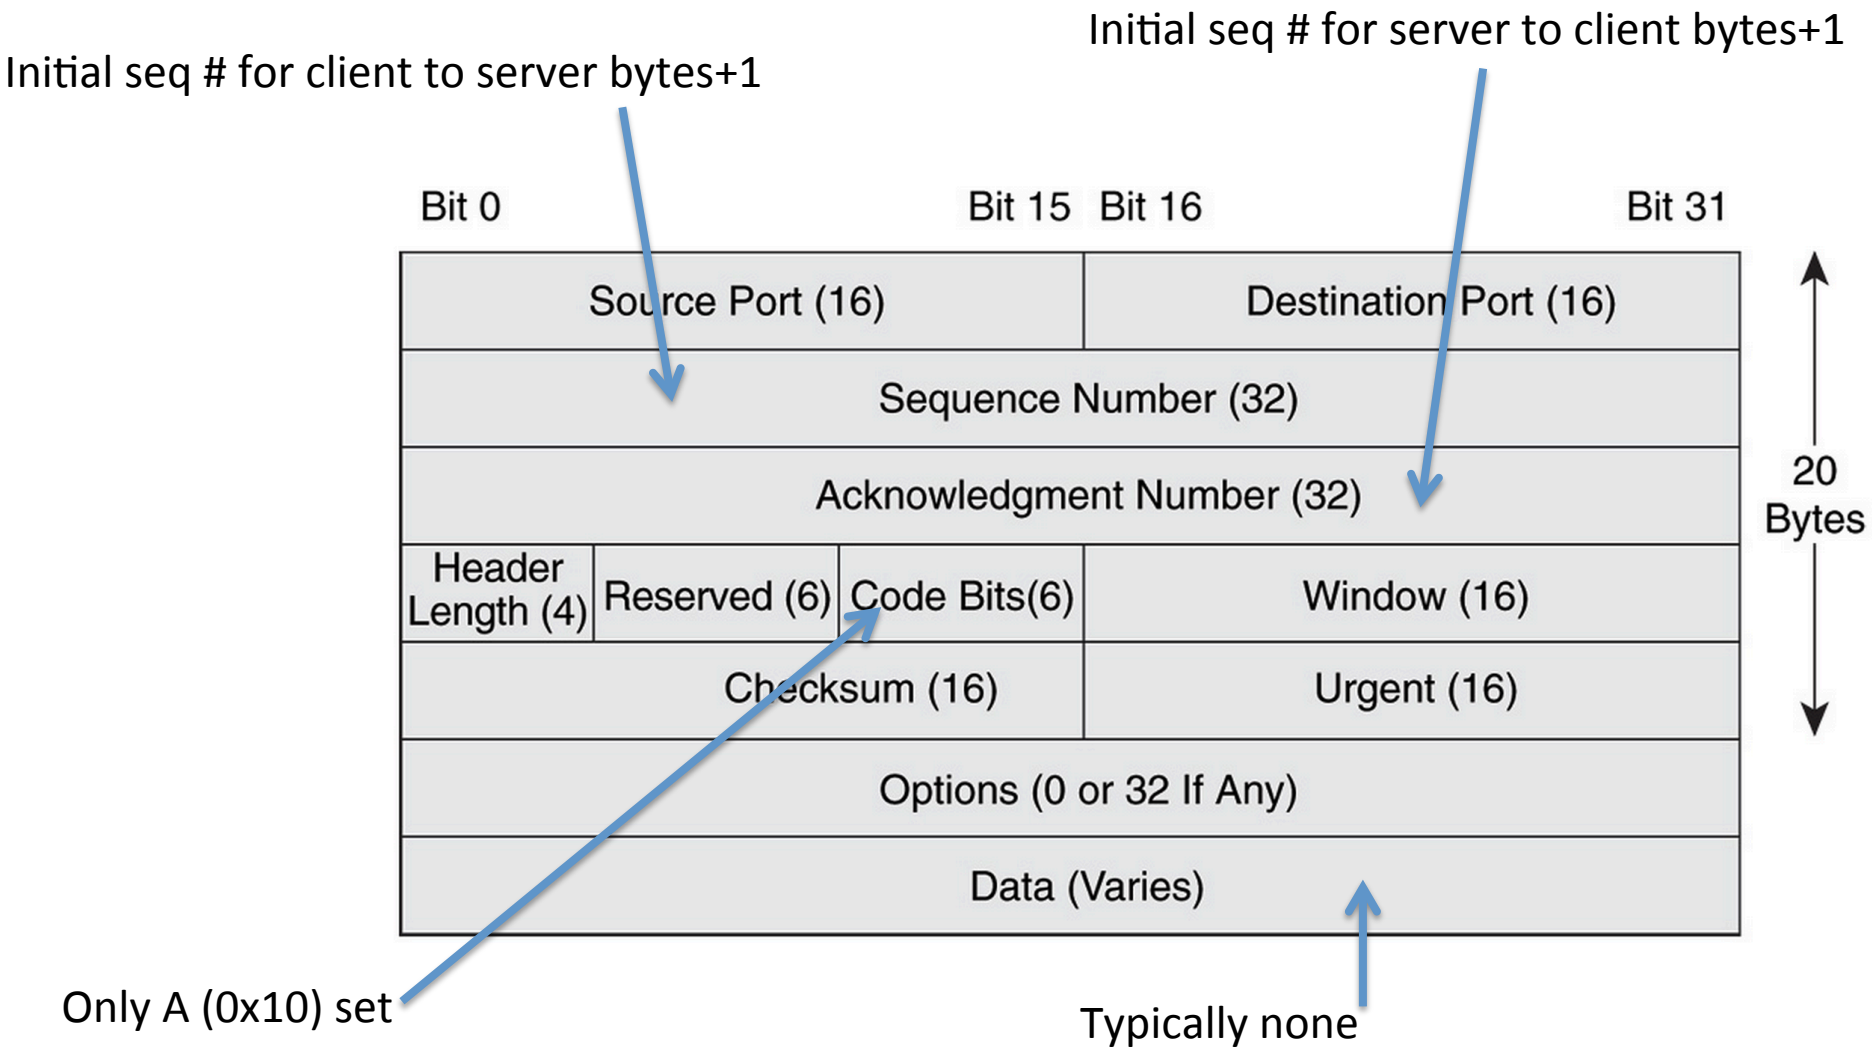

Tcpdump demo

- tcpdump -n -r nytimes.pcap port 58297 |more

IP Address Space

- Different organizations get different amounts
  - Class A: x.0.0.0/8 (2^{24} = 16,777,216)
    - x.1.1.1 is in, as is x.254.254.254)
    - Huge org eg (DOD is 11.0.0.0/8 IBM is 9.0.0.0/8)
  - Class B: x.y.0.0/16 (2^{16} = 65536)
    - Mid-sized organization
      - eg Cornell has 128.253.0.0/16, 128.84.0.0/16, 132.236.0.0/16 and 140.251.0.0/16
  - Class C: x.y.z.0/24 (2^{8} = 256)
    - Small organizations.
  - Can also have intermediate bitmasks.
    - eg /22

Internal Address Spaces

- RFC 1918
  - 10.0.0.0/8
  - 172.16.0.0/12
  - 192.168.0.0/16
- These addresses are not “routable”
- They will not be delivered across the Internet
  - Not allowed on there, technically.
- Need a special translator device at boundary
  - “NAT box” = Network Address Translation
  - Converts them to internet routable addresses

Port Scan Scenarios

- Bad guy wants to map an address space
  - Old style: across the internet
    - Still happens for internet facing servers
    - But rarely can map entire networks any more
  - Newer style: has a compromised machine on an internal network
  - Wants to know “what servers are here?”
  - Specifically, which machines have open ports?

Class B Portscan Example

- 2^16 addresses
- Say bad guy just scans on port 80
  - Eg say he knows an IIS or Apache exploit.
  - Send out 2^{16} syn packets to port 80
    - x.y.0.0, x.y.0.1, x.y.0.2,… x.y.255.254
  - "Horizontal scan on port 80"
  - See who sends back a syn-ack.
    - Means they have a process answering on port 80.
  - Find all the web servers this way.
  - Attack em!
    - Start with sending an ack pkt to establish conn.
    - Or not – if we don't send the 3rd handshake, system typically won't log.
      - Half-open connection

Vertical Port Scan of 1 IP

- Targetting a single IP address.
- Scan all 2^{16} ports.
- Find all ports answering

What Happens if Port Not Open

- No machine at all.
  - Typically get an ICMP response from a router
    - Special protocol for Internet error message packets
    - Saying no host at this address
- Machine but with closed port
  - Typically get a reset packet
  - Like a syn-ack, but with R set instead of S and A
  - Semantics – "stop this immediately"
- Security system (firewall)
  - Silence (depending on configuration)

Visualizing Scans

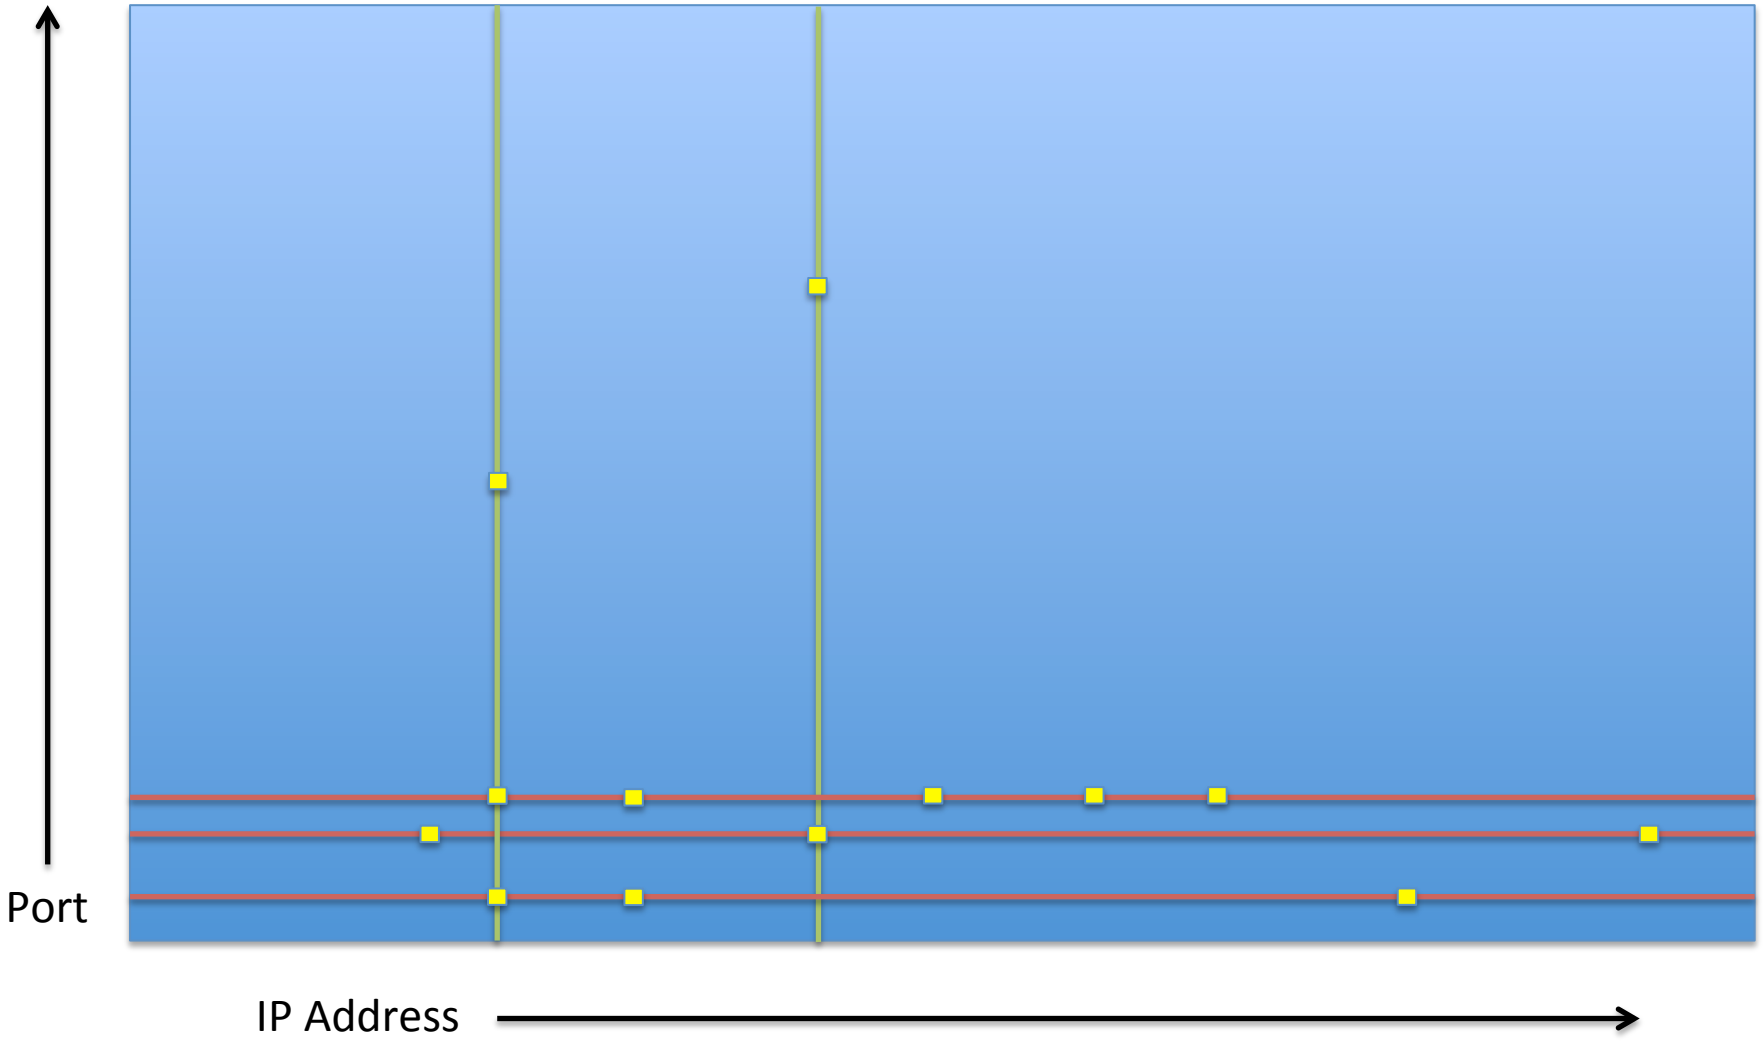

Small Piece of a Large Random Scan