Defending Computer Networks
Lecture 7: Port Scanning

Stuart Staniford

Adjunct Professor of Computer Science

VANCORT CONSULTING
speak up. stand out.

Bio

Eliza VanCort

I grew up in the theater and received a liberal arts education, graduating Phi Beta Kappa from CU Boulder's honors program with a degree in Political Science. After college I worked with at-risk youth for a year, and then attended NYU Law School. Despite a successful first year, I found that I missed the theater, and I realized that law school wasn't the right path for me.

Logistics

- Aim to give out HW2 on Thursday

NEWS ANALYSIS

Tech startups need to get serious about security

The head of the nation's primary consumer protection agency on Wednesday paid a visit to San Francisco, where she called on technology startups to do a better job of incorporating security protections as they race to bring new applications into the market.

Federal Trade Commission Chairwoman Edith Ramirez's comments amplified the agency's "Start With Security" initiative, a program that aims to encourage businesses to prioritize cybersecurity as an integral part of their product development.

Main Goals for Today

- TCP Portscanning
- Detection of Portscanning

Refresh: 3-way handshake

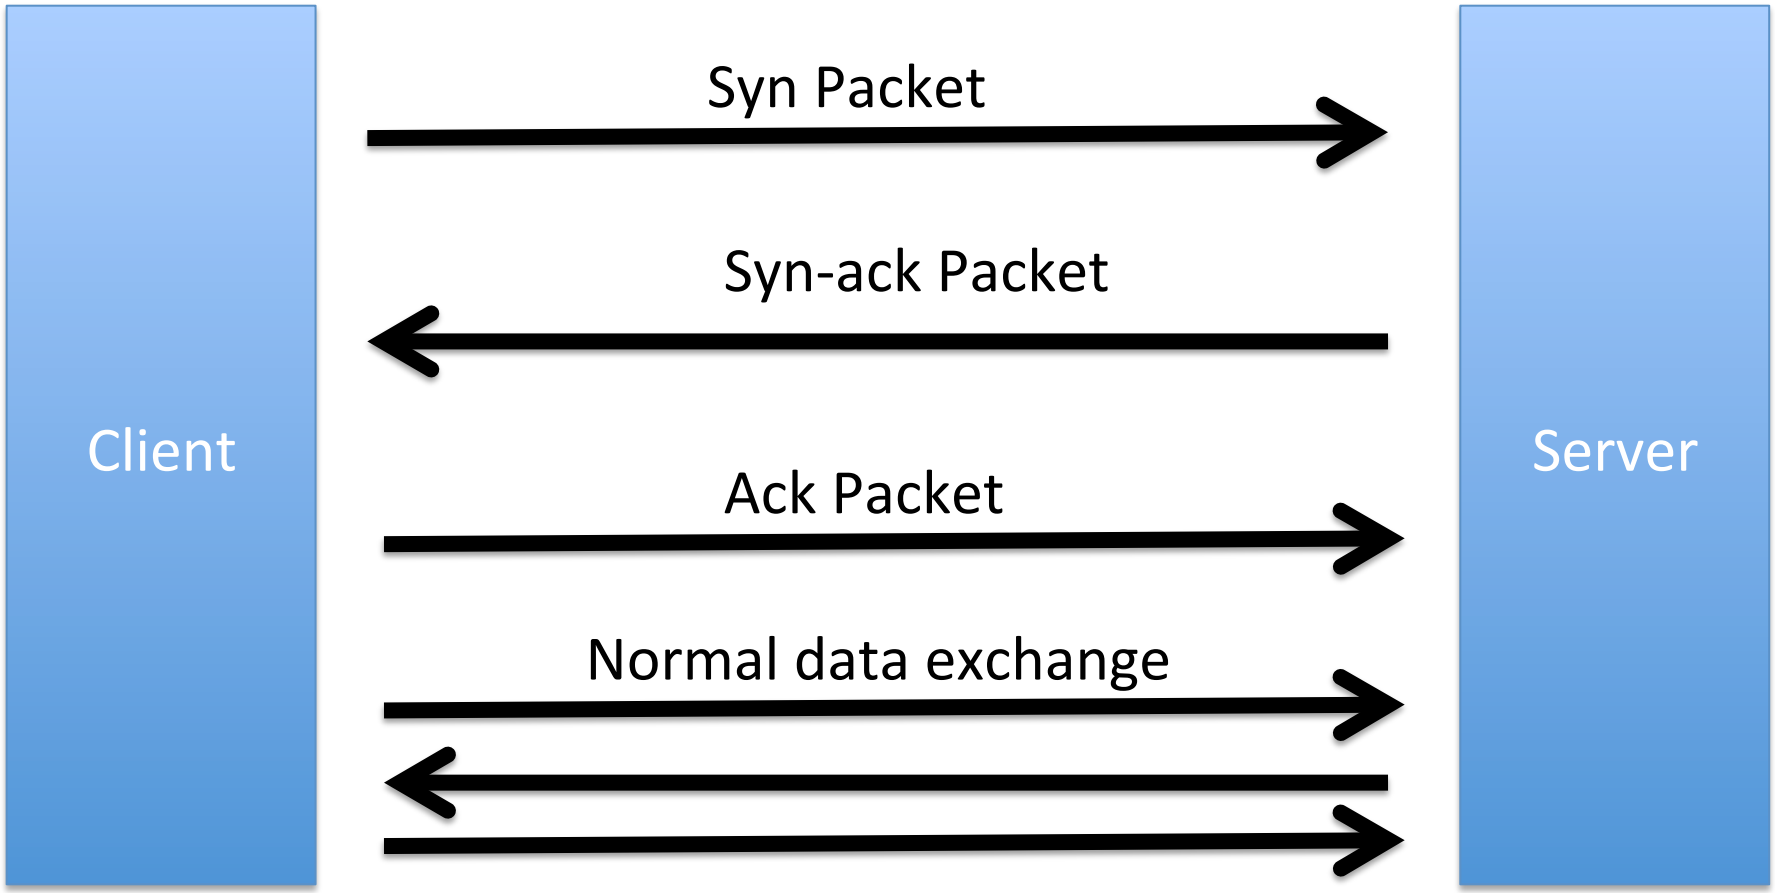

Refresh: IP Address Space

- 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

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

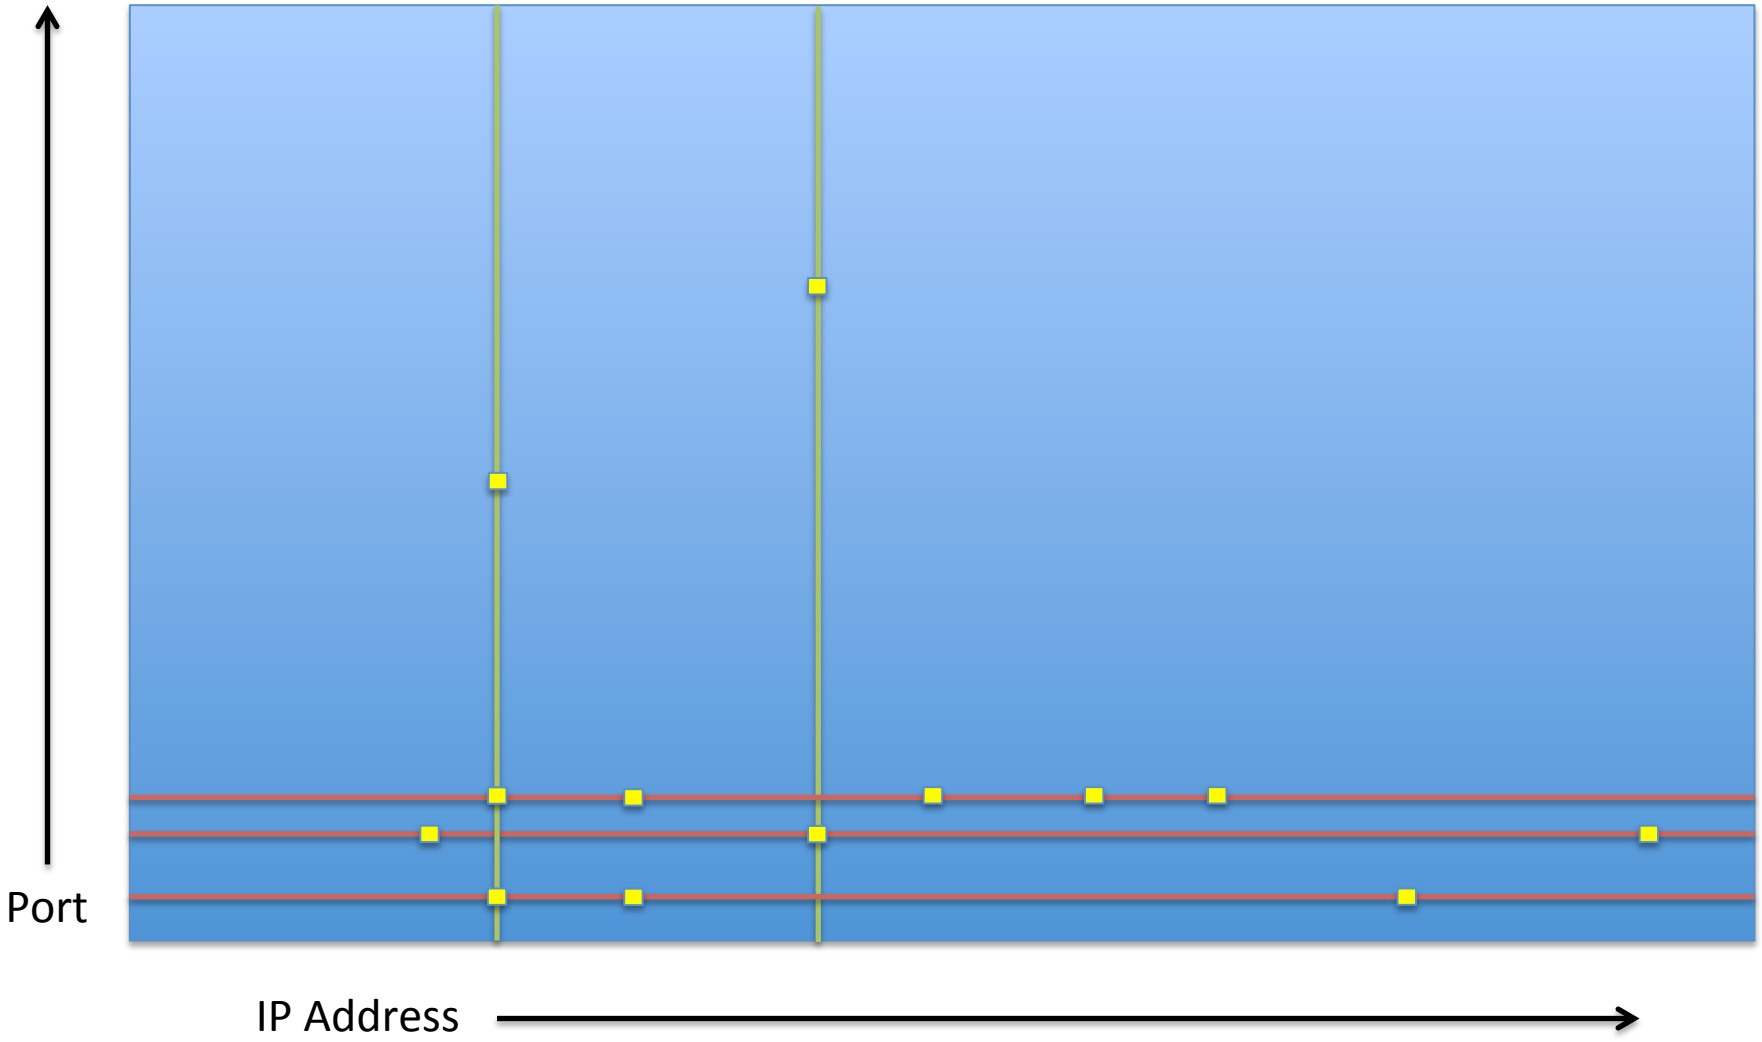

Small Piece of a Large Random Scan

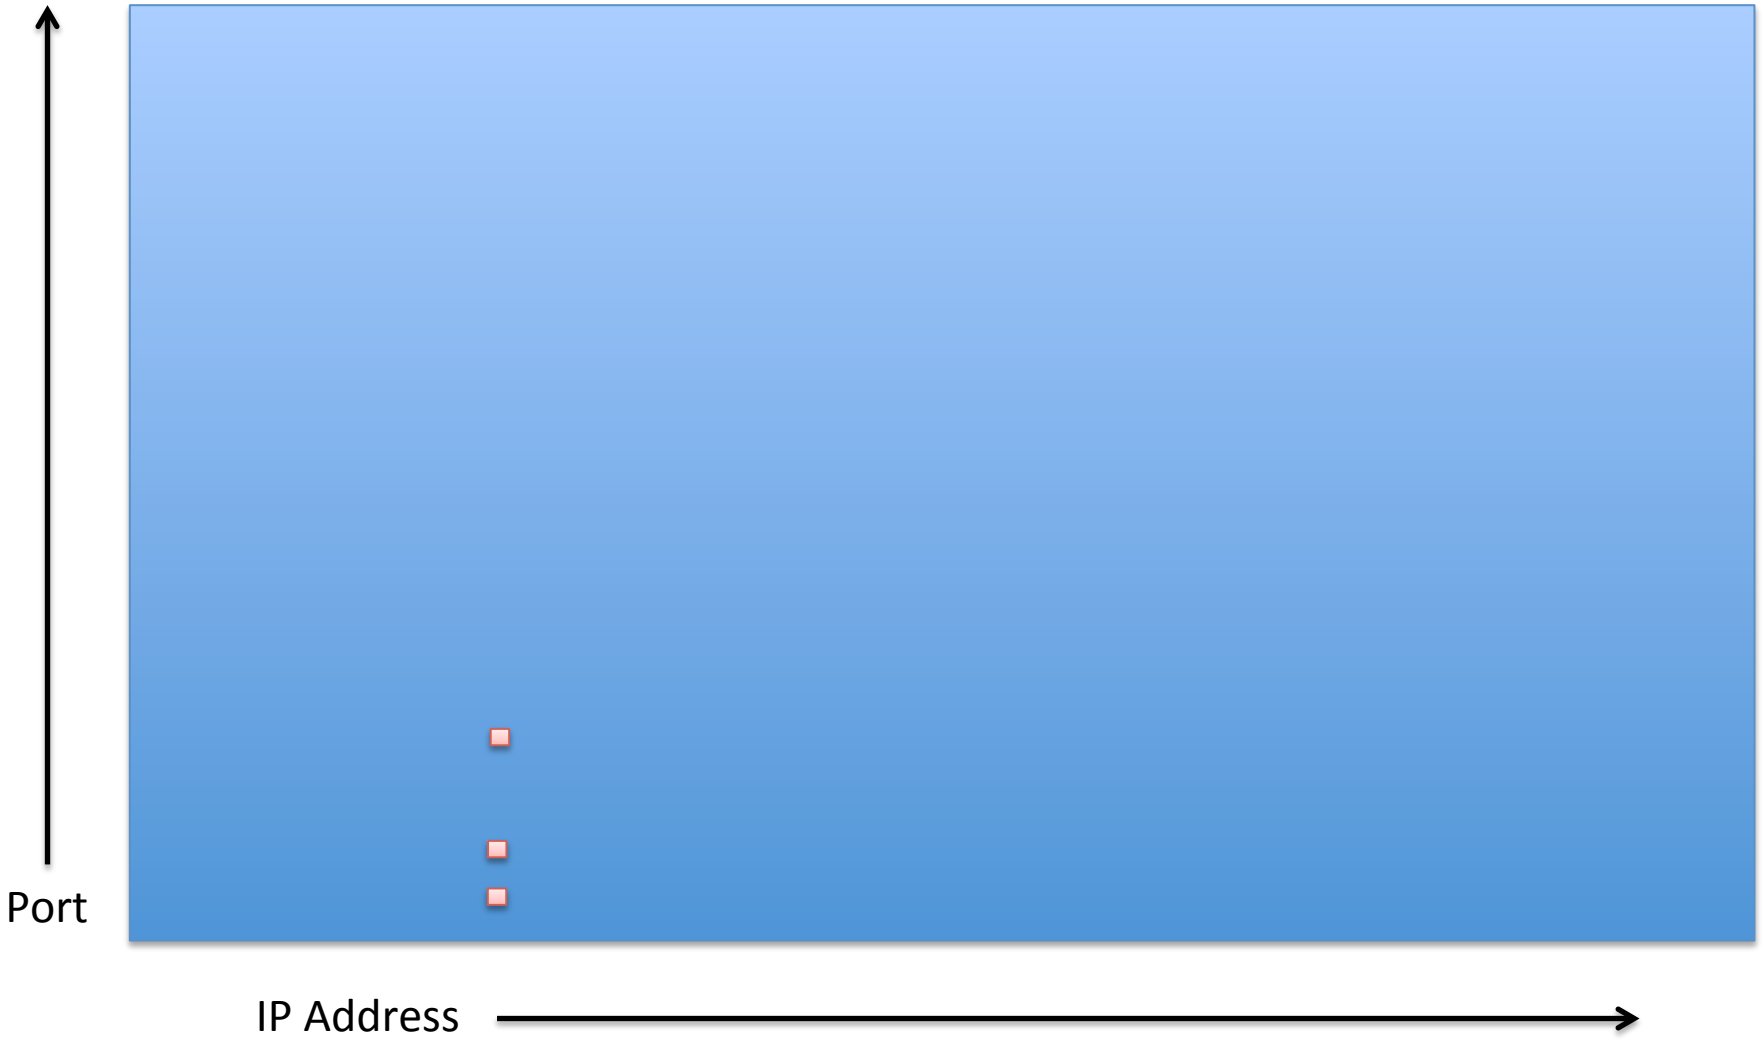

Let’s try it

- sudo nmap -n –sS 10.0.0.2

What’s Happening on The Wire

- sudo tcpdump -n -i en3
- sudo nmap -n –sS 10.0.0.2

Cyberthreat Posed by China and Iran Confounds White House

By **DAVID E. SANGER** SEPT. 15, 2015

President Obama speaking to troops at Fort Meade, Md., home of the National Security Agency and the United States Cyber Command. Stephen Crowley/The New York Times

“Offense is moving a lot faster than defense,” Mr. Obama told troops on Friday at Fort Meade, Md., home of the National Security Agency and the United States Cyber Command. “The Russians are good. The Chinese are good. The Iranians are good.” The problem, he said, was that despite improvements in tracking down the sources of attacks, “we can’t necessarily trace it directly to that state,” making it hard to strike back.

Then he issued a warning: “There comes a point at which we consider this a core national security threat.” If China and other nations cannot figure out the boundaries of what is acceptable, “we can choose to make this an area of competition, which I guarantee you we’ll win if we have to.”

TCP Fin Flag

- Used to indicate orderly close of a connection.
  - Fin (F) 0x0x in TCP header flags field
- Either side may issue a packet with FIN in.
  - Can be a data packet.
- Other side should respond with a FIN pkt.
- Connection is then over and no more pkts should be sent.

FIN Scanning

`-sN`; `-sF`; `-sX` (TCP NULL, FIN, and Xmas scans)

These three scan types (even more are possible with the `--scanflags` option described in the next section) exploit a subtle loophole in the TCP RFC to differentiate between `open` and `closed` ports. Page 65 of RFC 793 says that "if the [destination] port state is CLOSED .... an incoming segment not containing a RST causes a RST to be sent in response." Then the next page discusses packets sent to open ports without the SYN, RST, or ACK bits set, stating that: "you are unlikely to get here, but if you do, drop the segment, and return."

When scanning systems compliant with this RFC text, any packet not containing SYN, RST, or ACK bits will result in a returned RST if the port is closed and no response at all if the port is open. As long as none of those three bits are included, any combination of the other three (FIN, PSH, and URG) are OK. Nmap exploits this with three scan types:

Null scan (`-sN`)

Does not set any bits (TCP flag header is 0)

FIN scan (`-sF`)

Sets just the TCP FIN bit.

Xmas scan (`-sX`)

Sets the FIN, PSH, and URG flags, lighting the packet up like a Christmas tree.

Let's try these and compare

- tcpdump -n -i en0
- nmap -n –sS 10.0.0.2
- nmap -n –sF 10.0.0.2
- If time
  - nmap -n –sX 10.0.0.2
  - nmap -n –sN 10.0.0.2

What is Advantage

- Some early packet filters
  - Network access control devices
  - Would just examine syns to enforce policy
  - Eg if we want to block inbound email,
    - No syns to port 25.
  - Allow all non-syn pkts through
    - on the theory that end-host will not actually allow a connection with no syn.
  - But, end-host might respond to FIN scan, allowing attacker to portscan it through filter.

Let's look at everything nmap can do

- Just for kicks
  - May not work, is slow/flaky at times
- sudo nmap -n –A –T4 10.0.0.2

US agency in charge of power grid and nukes keeps getting breached

The government records show cyberattackers successfully compromised the DOE 159 times between October 2010 and October 2014, and was attacked a total of 1131 times during that period.

USA Today reports that 53 of the successful attacks were root compromises, meaning the attackers had administrator privileges on compromised DOE computer systems.

Of the 159 successful intrusions, 90 compromised the DOE Office of Science, which conducts energy research, and another 19 attacks compromised the National Nuclear Security Administration - the agency in charge of securing the nation's stockpile of nuclear weapons.

https://nakedsecurity.sophos.com/2015/09/11/
us-agency-in-charge-of-power-grid-and-nukes-keeps-getting-breached/

Algorithms to Detect Portscans

- First brush with Network Intrusion Detection
  - General art/science of detecting badness by watching packets fly by.
  - Invented at UC Davis
    - Todd Heberlein et al circa 1989
    - "Network Security Monitor"
  - Portscan detection is a nice sample problem.
  - Illustrates many of the issues in an easy-to-follow context.

Firstly We Need to Get Packets

- Old
  - Promiscuously monitor a hub/wire
- Modern
  - Span port on switch
  - Network tap device
  - Detection device itself inline
    - IPS – Intrusion Prevention System
- For CS 5434 purposes, libpcap
  - 'man pcap' will get you started.

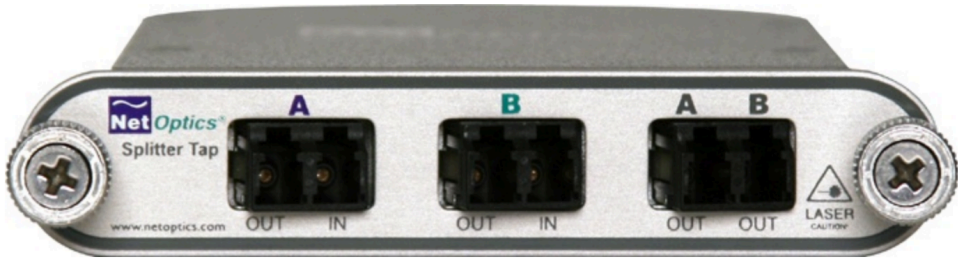

Then we need a data structure

- Simplest possible thing is a hash table
  - keyed on client IP
  - With per-connection counts of relevant stuff
  - Eg just count syns
  - Portscanners will issue more syns than average.
    - Alert when count goes over threshold
    - But what's likely to go wrong?

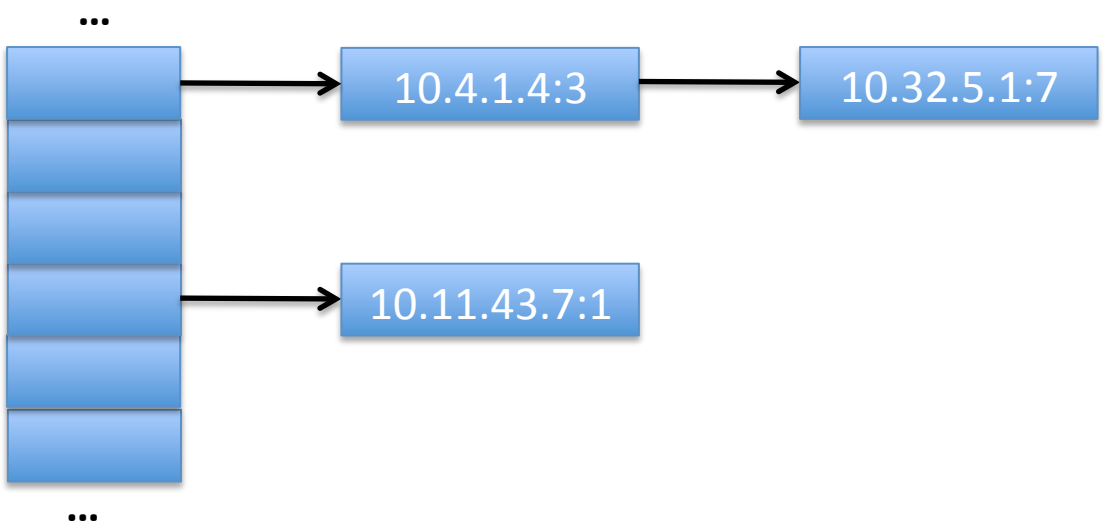

Another possibility

- Look for the actual sequential behavior
  - Syn->10.4.35.1
  - Syn->10.4.35.2
  - Syn->10.4.25.3
  - …
- Implement by having a “last dest” field in table entry
  - Keep counts of “number of increment-by-ones”
- Fragile
  - What could go wrong?

Keep track of unique dests/src?

- Now have to have a way to know
  - what is a unique dst for that src?

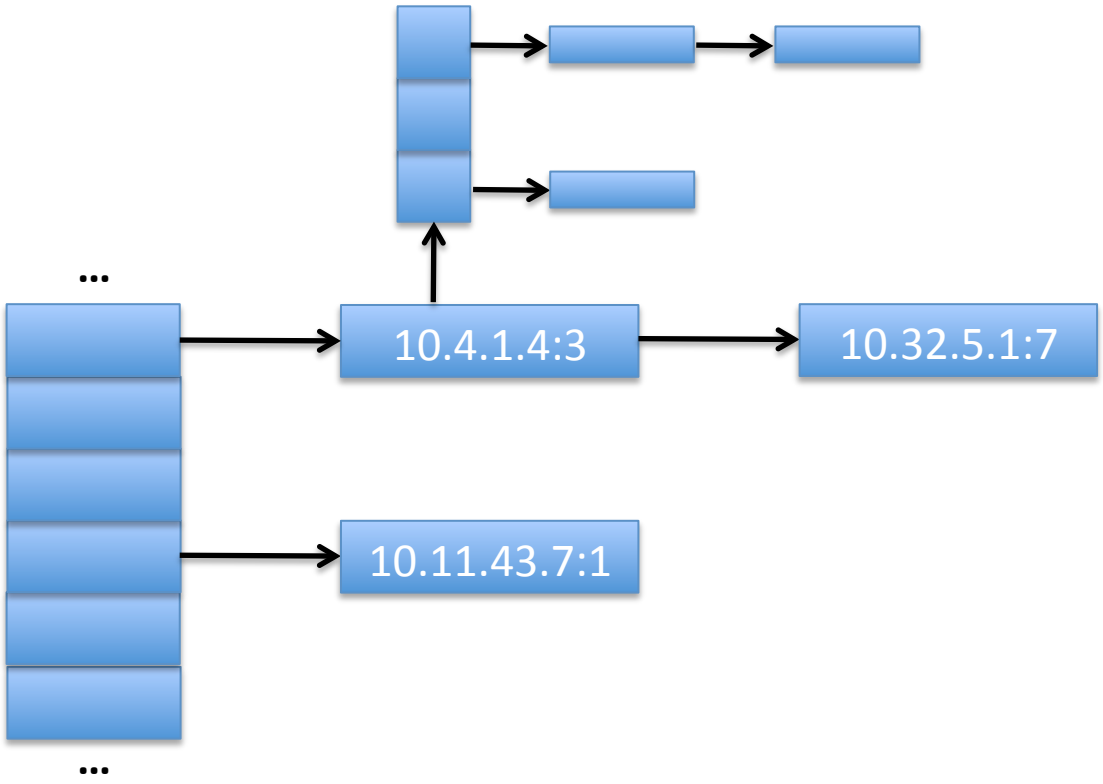

Better Idea

- Key off the idea that port-scanners make a lot of failed connections.
- Legit users make only a few
  - So keep track of "failed-succeed" count
  - Alert when goes over threshold.
- How can the attacker game this?
- Doesn't work in the presence of packet-filter/ firewalls.

Another Idea

- Learn the probability of a syn (say) being to a destination:
  - P(D)
  - Popular servers will have high P(D) (say 5% or 1%)
  - Non-servers will have very low P(D) (1 in 10^6 or 10^9)
  - Take –log(P(D)) and accumulate *that* in hash table
    - Anomaly score
  - Portscanners will accumulate a lot of anomaly score
    - Alert if over a threshold
  - Harder for attackers to game – don't know P(D)
    - Otherwise wouldn't need to portscan

Extending the basic idea

- Keep flow table state
- Know when we see things like unexpected F
- Give that a high anomaly score