Insertion, evasion, denial-of-service, and other network tomfoolery

UNIX process hierarchy

pstree

pstree -u crandall

cs /tmp

wget phrack.org

less index.html

strace -f -o bla.txt wget phrack.org

less bla.txt

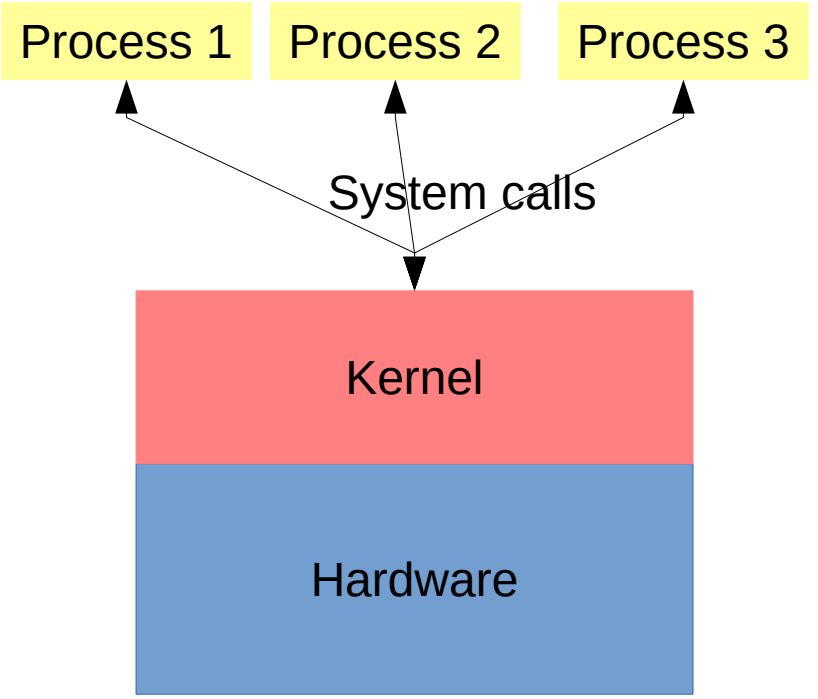

OSI model

- Layer 1: Physical (think Ethernet, 802.11)
- Layer 2: Data Link (think ARP)
- Layer 3: Network (think IP)
- Layer 4: Transport (think TCP)
- Layer 5: Session (think NetBIOS, SOCKS)
- Layer 6: Presentation (think SSL/TLS)
- Layer 7: Application (think HTTP)

TCP Socket 基本流程圖
TCP Socket flow diagram

By OnionBulb - This PNG image was made by OnionBulb. PNG filename originally is "InternetSocketBasicDiagram_zhtw.png"., Public Domain, https://commons.wikimedia.org/w/index.php?curid=11766896

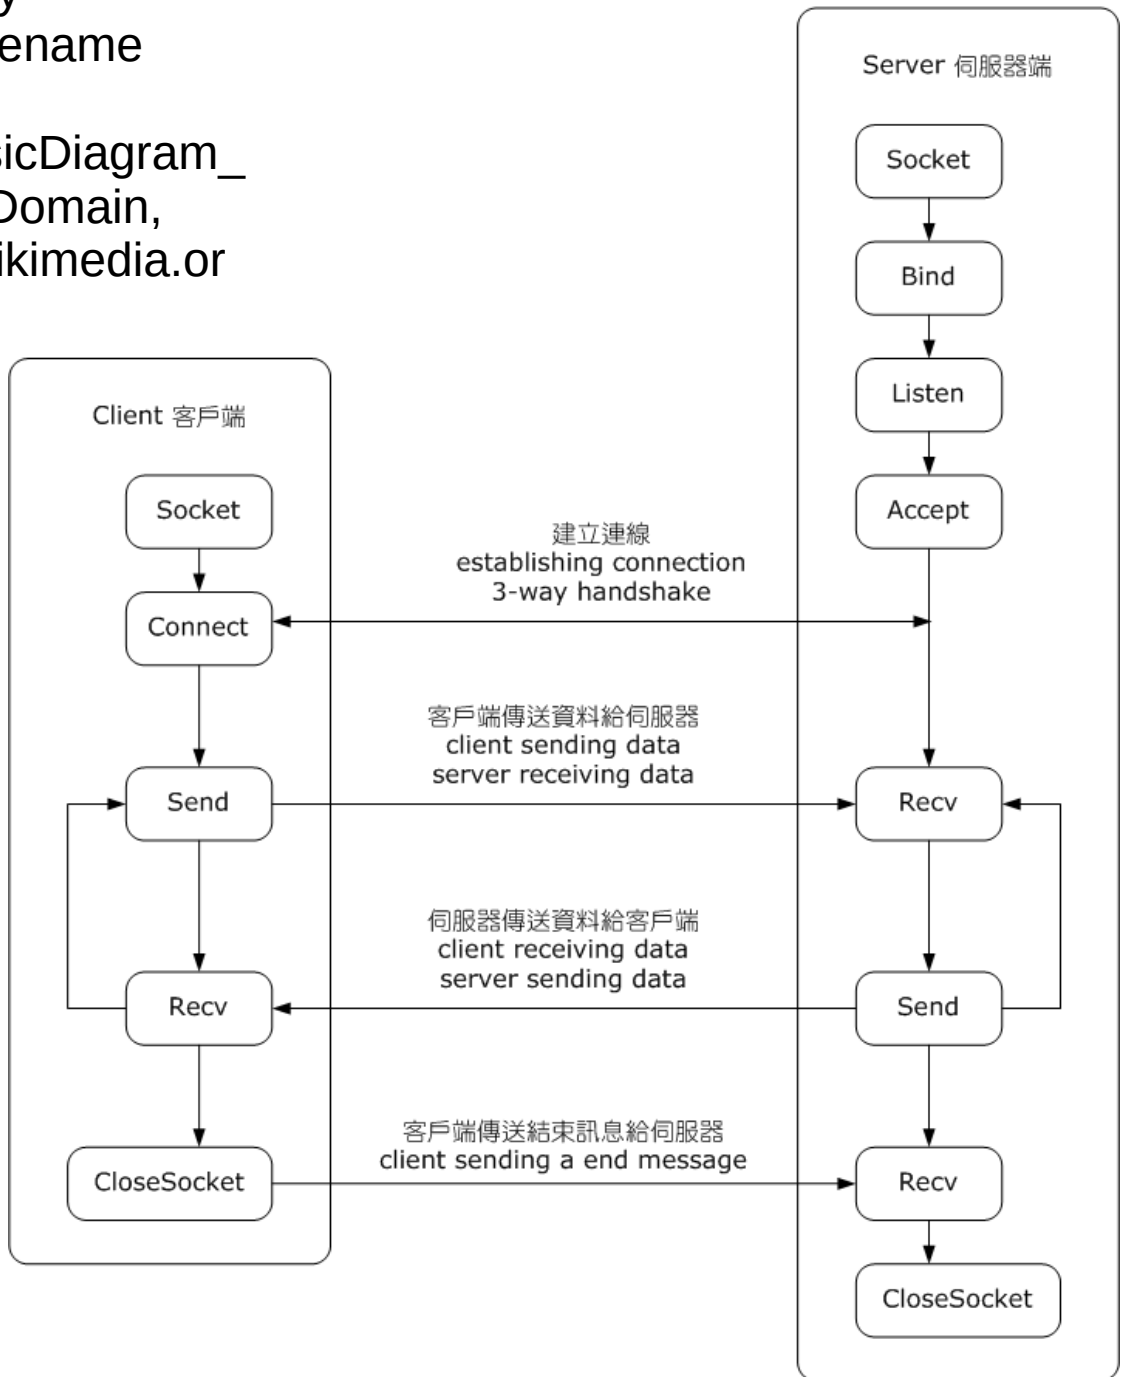

TCP 3-way handshake (review)

- TCP header has flags
  - SYN is “Synchronize”, it means the sequence number has a special meaning
  - ACK is “Acknowledge”, it means the acknowledgment number has meaning
  - RST: “I have no record of such a connection”
  - Also, FIN, CWR, ECN, URG, PUSH

TCP 3-way handshake (review)

- SYN: I'd like to open a connection with you, here's my initial sequence number (ISN)
- SYN/ACK: Okay, I acknowledge your ISN and here's mine
- I ACK your ISN

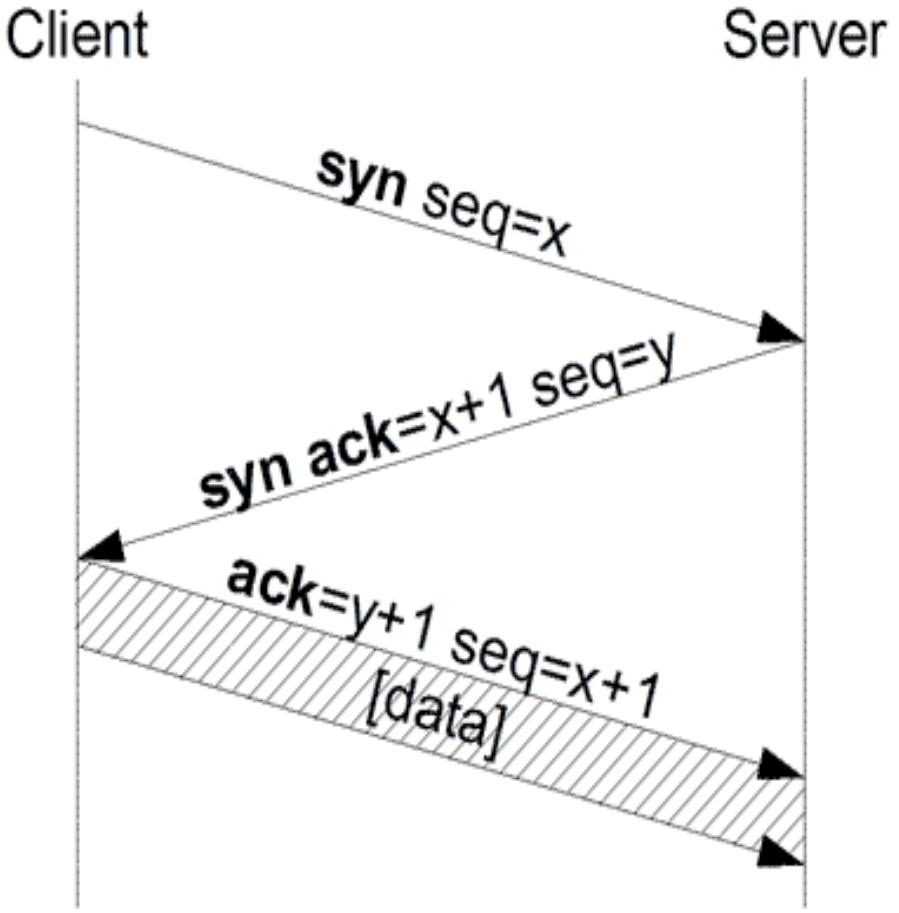

Image from Wikipedia

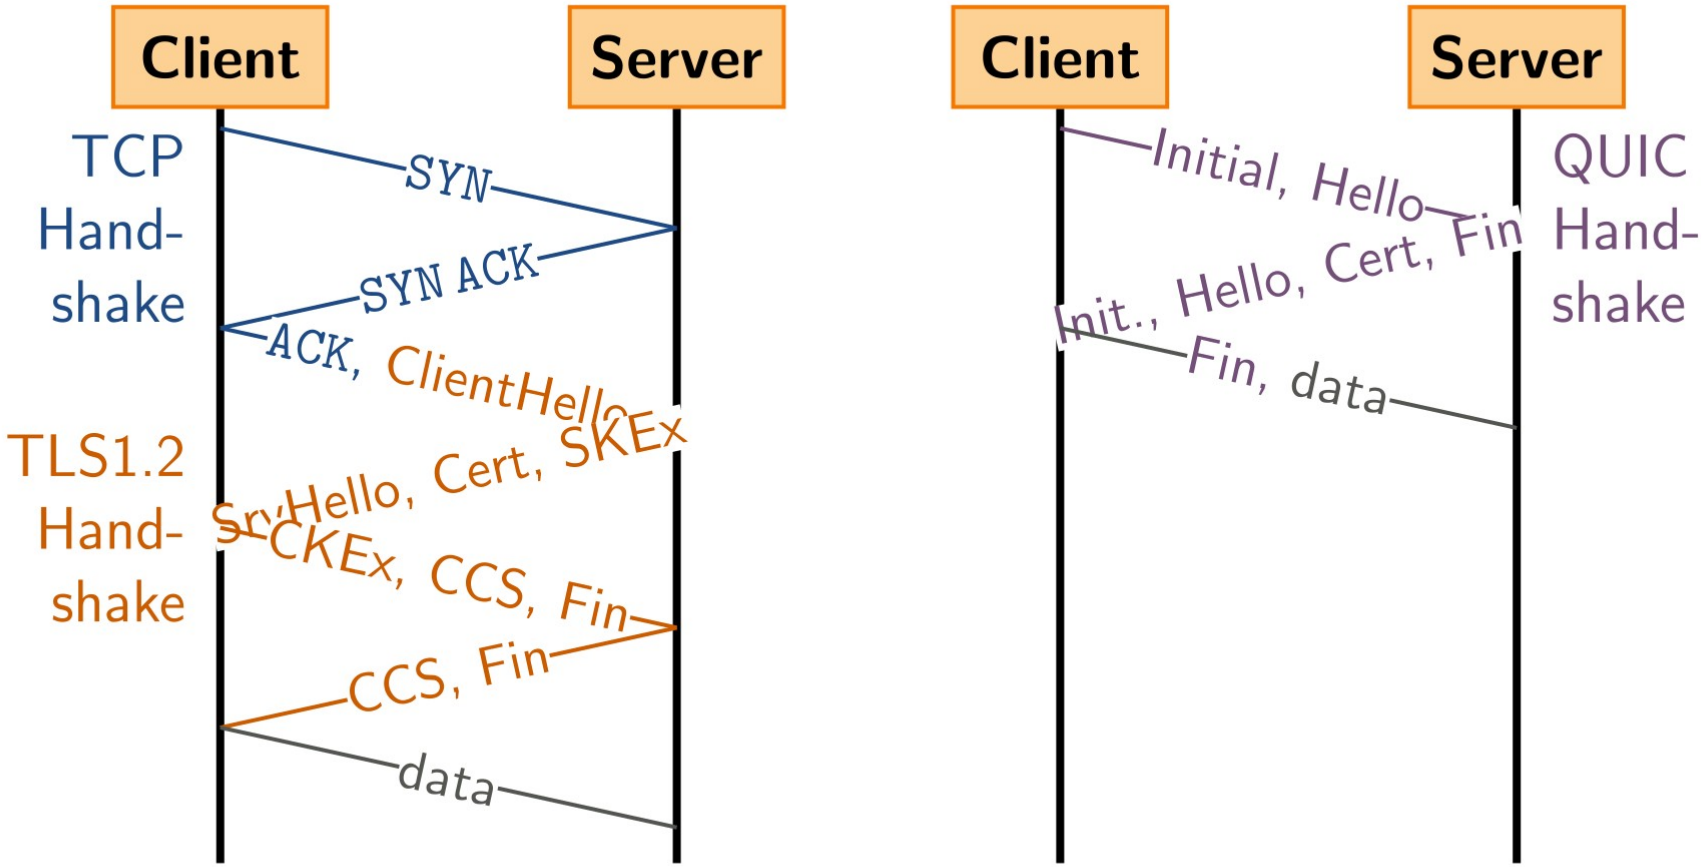

Plagiarized from:
https://en.wikipedia.org/wiki/QUIC

Where do these standards come from?

- IETF = Internet Engineering Task Force
- RFC = Request for Comments
  - MUST, MUST NOT, SHOULD, SHOULD NOT, MAY (RFC 2119)
- “The only laws on the Internet are assembly and RFCs” --Phrack 65
  - Assembly is an abstraction
  - RFCs are not always followed
    - Often ambiguous

TCP 3-way handshake

- TCP header has flags
  - SYN is “Synchronize”, it means the sequence number has a special meaning
  - ACK is “Acknowledge”, it means the acknowledgment number has meaning
  - RST: “I have no record of such a connection”
  - Also, FIN, CWR, ECN, URG, PUSH

Attacks in Layer 1

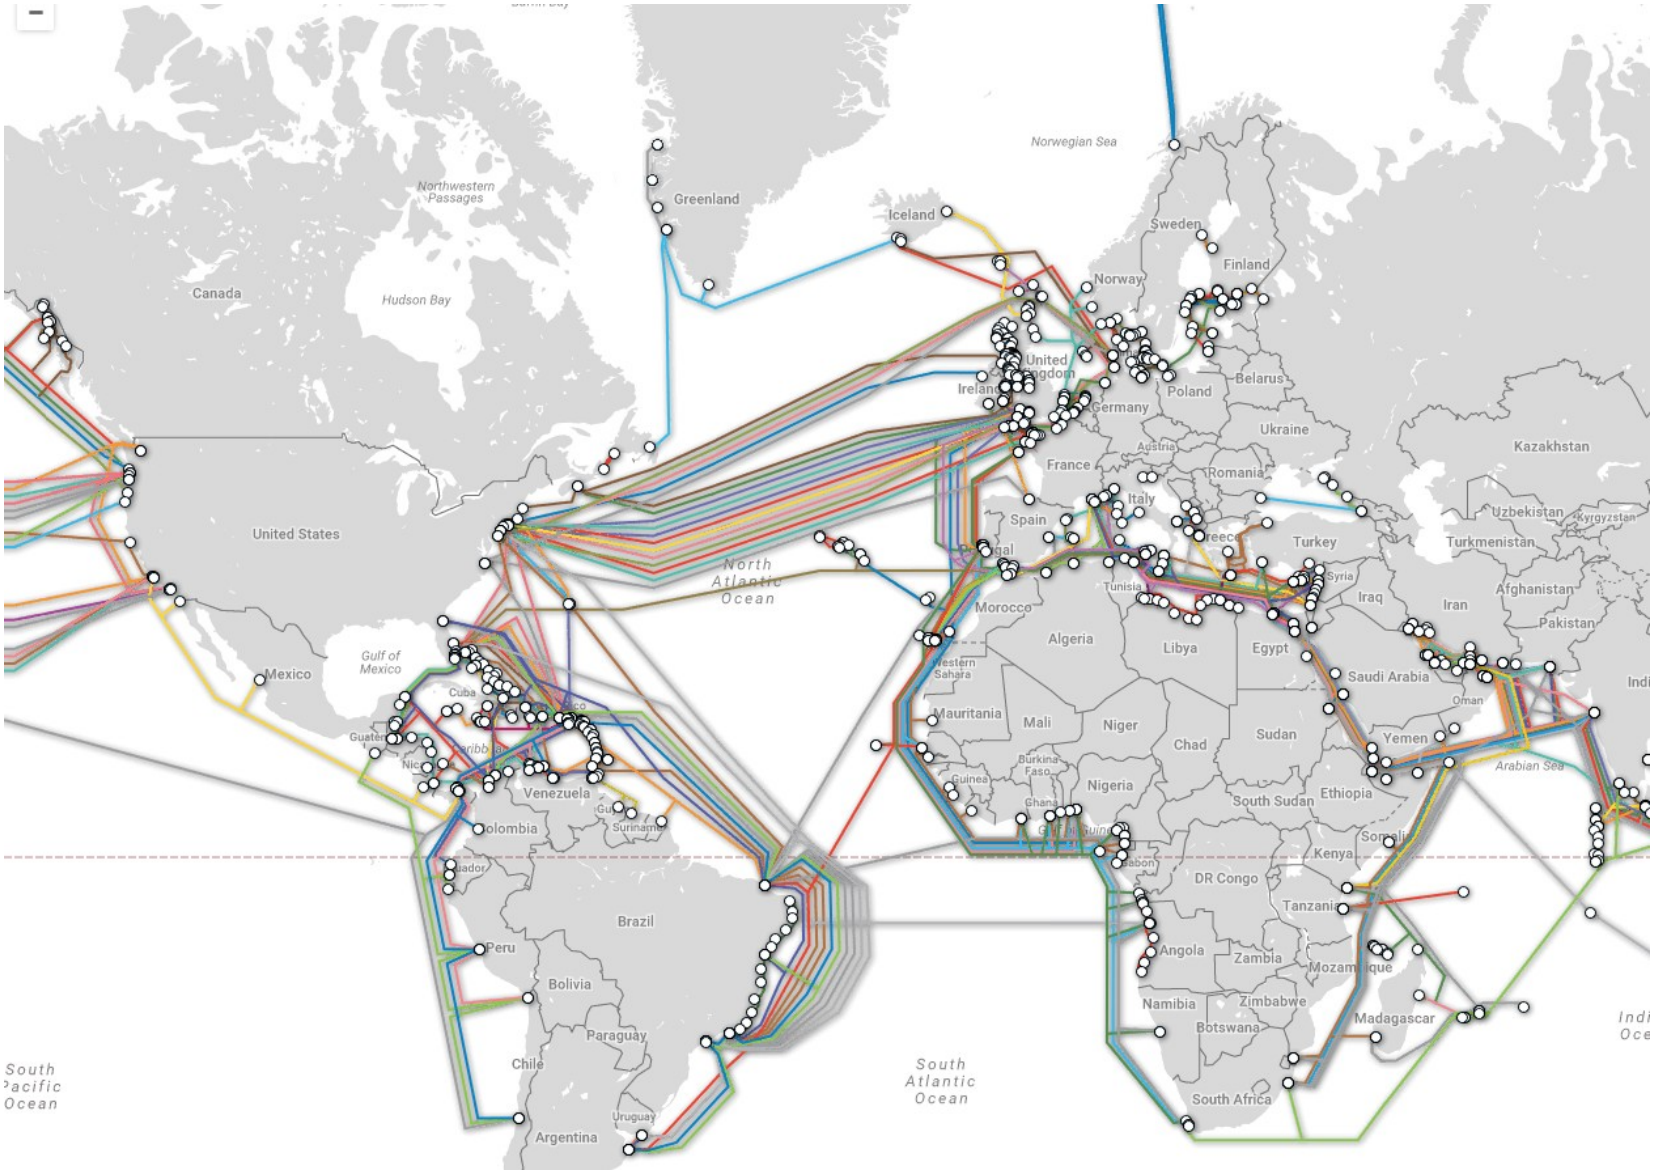

From submarinecablemap.com

Attacks in Layer 1

- Taps are easy
  - Port mirrors on backbone routers literally split light
  - Port is the physical hole in a router, can mirror any of them to get a copy of the traffic
- 802.11 suite of wifi protocols has various issues

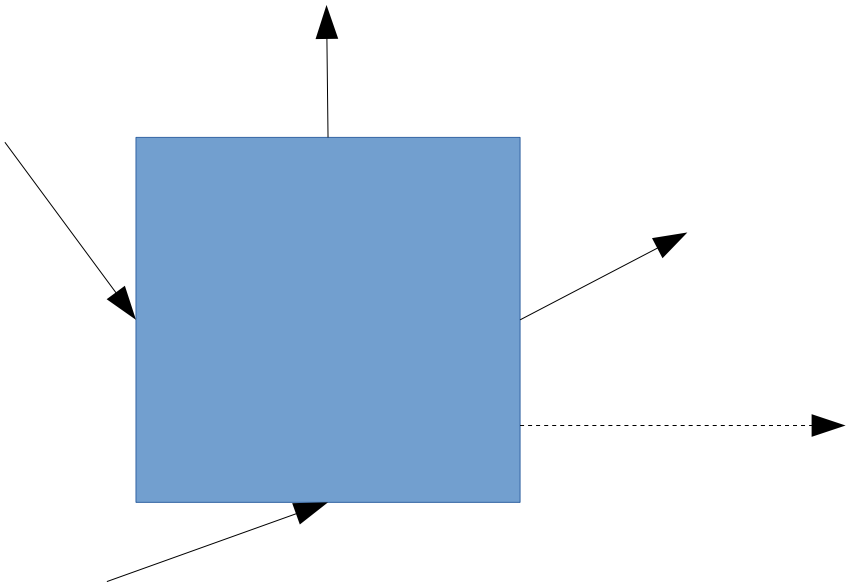

Attacks in Layer 2

- ARP spoofing
- ARP cache poisoning

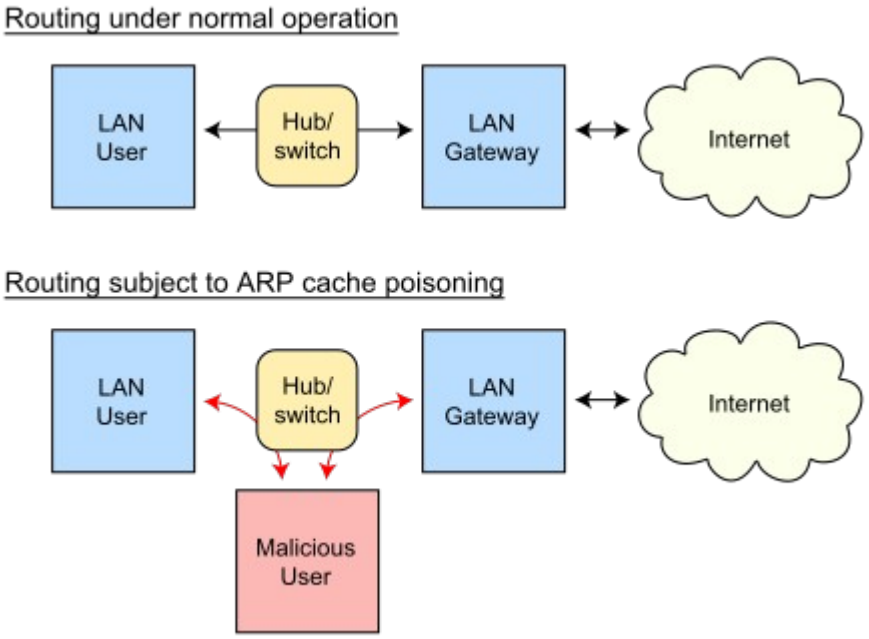

From Wikipedia

Attacks in Layer 3

- Man-in-the-middle
  - Great Cannon is an example (in-path)
- Man-on-the-side
  - Great Firewall of China (GFW, on-path) and NSA QUANTUM are examples
- TTL is a clue, but is easy to hide

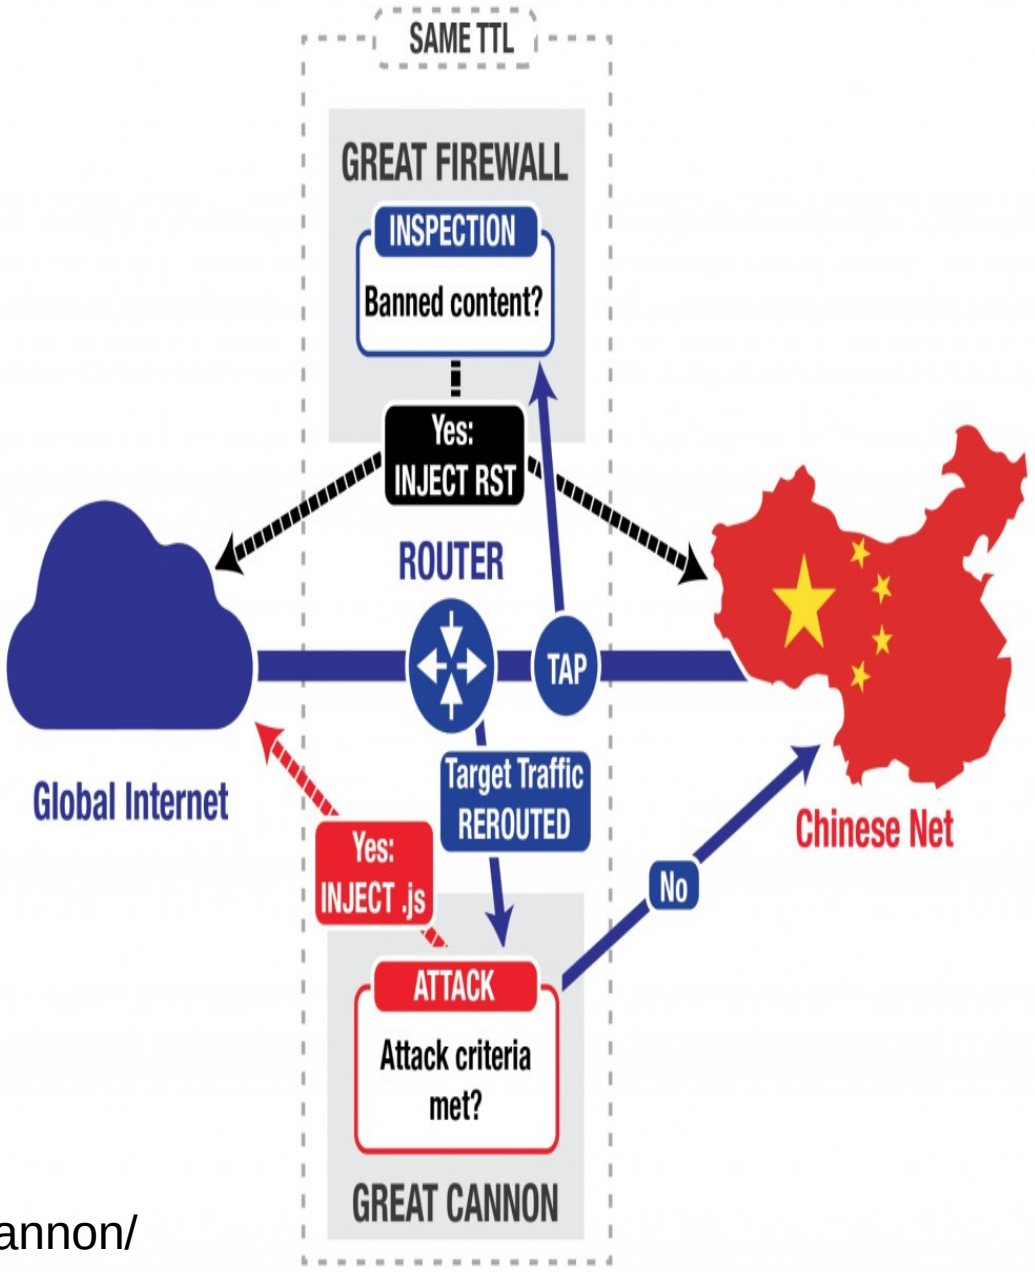

Image reproduced from
https://citizenlab.org/2015/04/chinas-great-cannon/

Intrusion Detection System (examples are Bro or Snort)

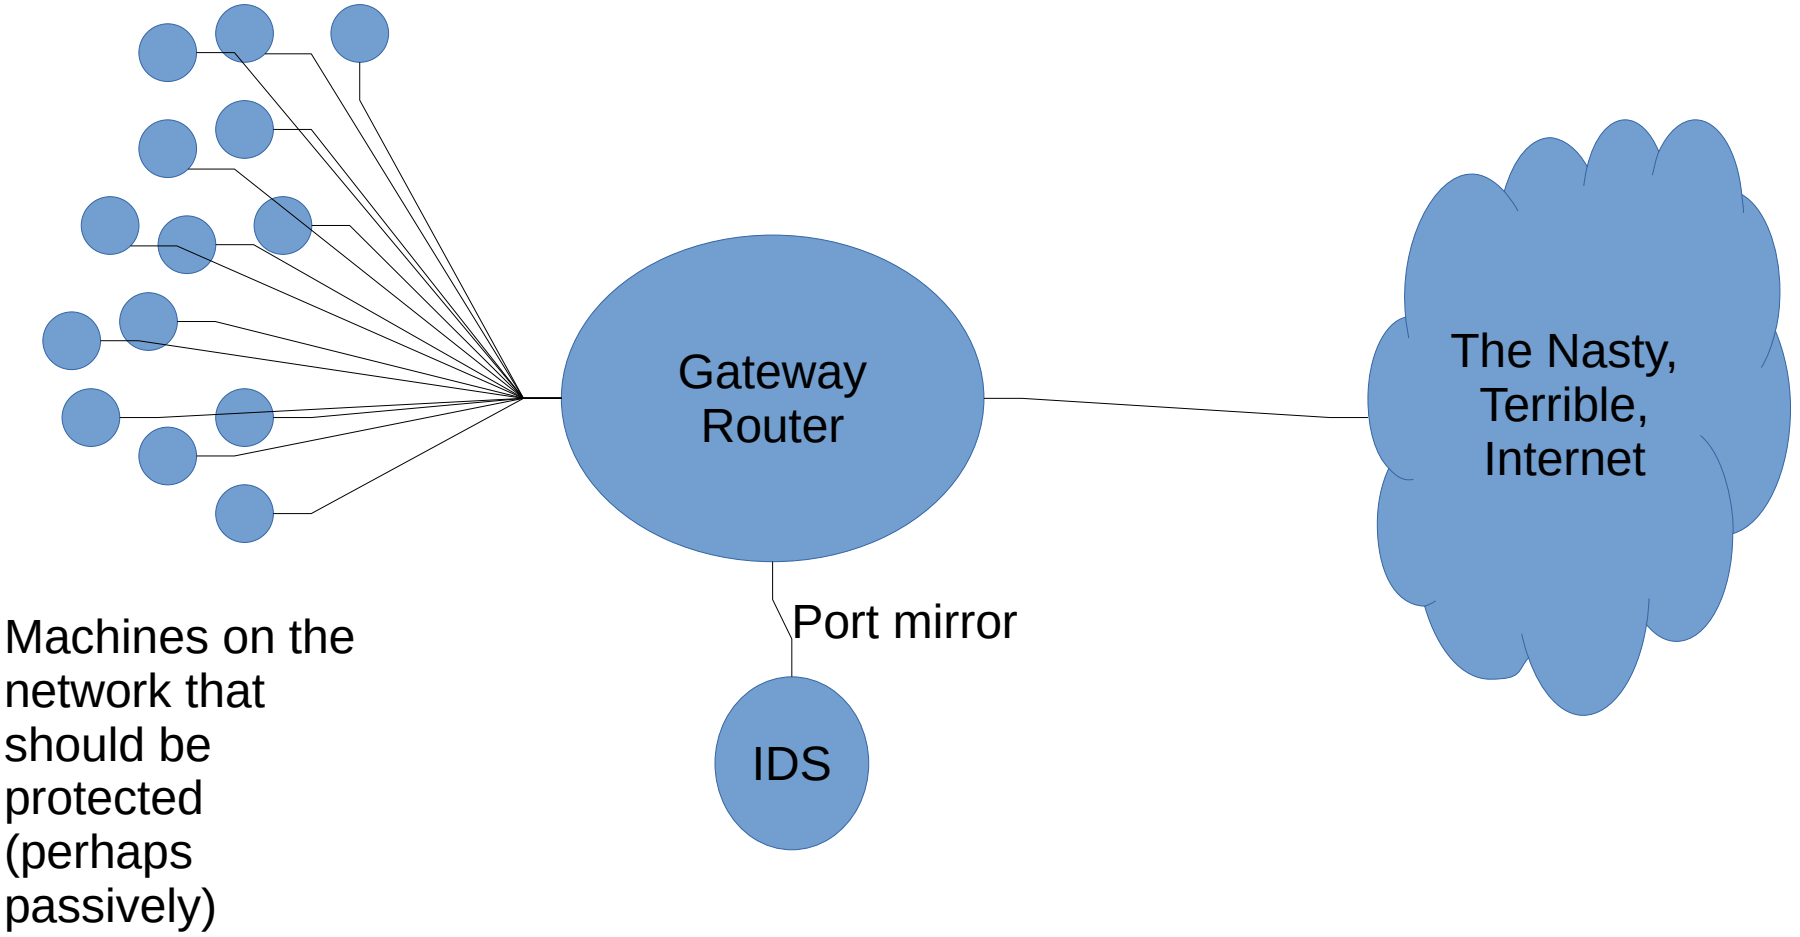

IDS is looking for signatures

- Typically regular expressions, like “`.*<script>.*</script>.*`” appearing in an input to a web form, indicating a Javascript XSS attack.
- How can we (the attacker) get the IDS to see one thing and the victim to see another?
- A stupid example: Great Firewall of China censors “`GET falungong.html`”, but if you send two packets: “`GET fa`” and “`lungong.html`” the endhost reassembles them fine but the GFW is fooled.
- Or, “GET fa%61lungong.html”

A not so useful distinction

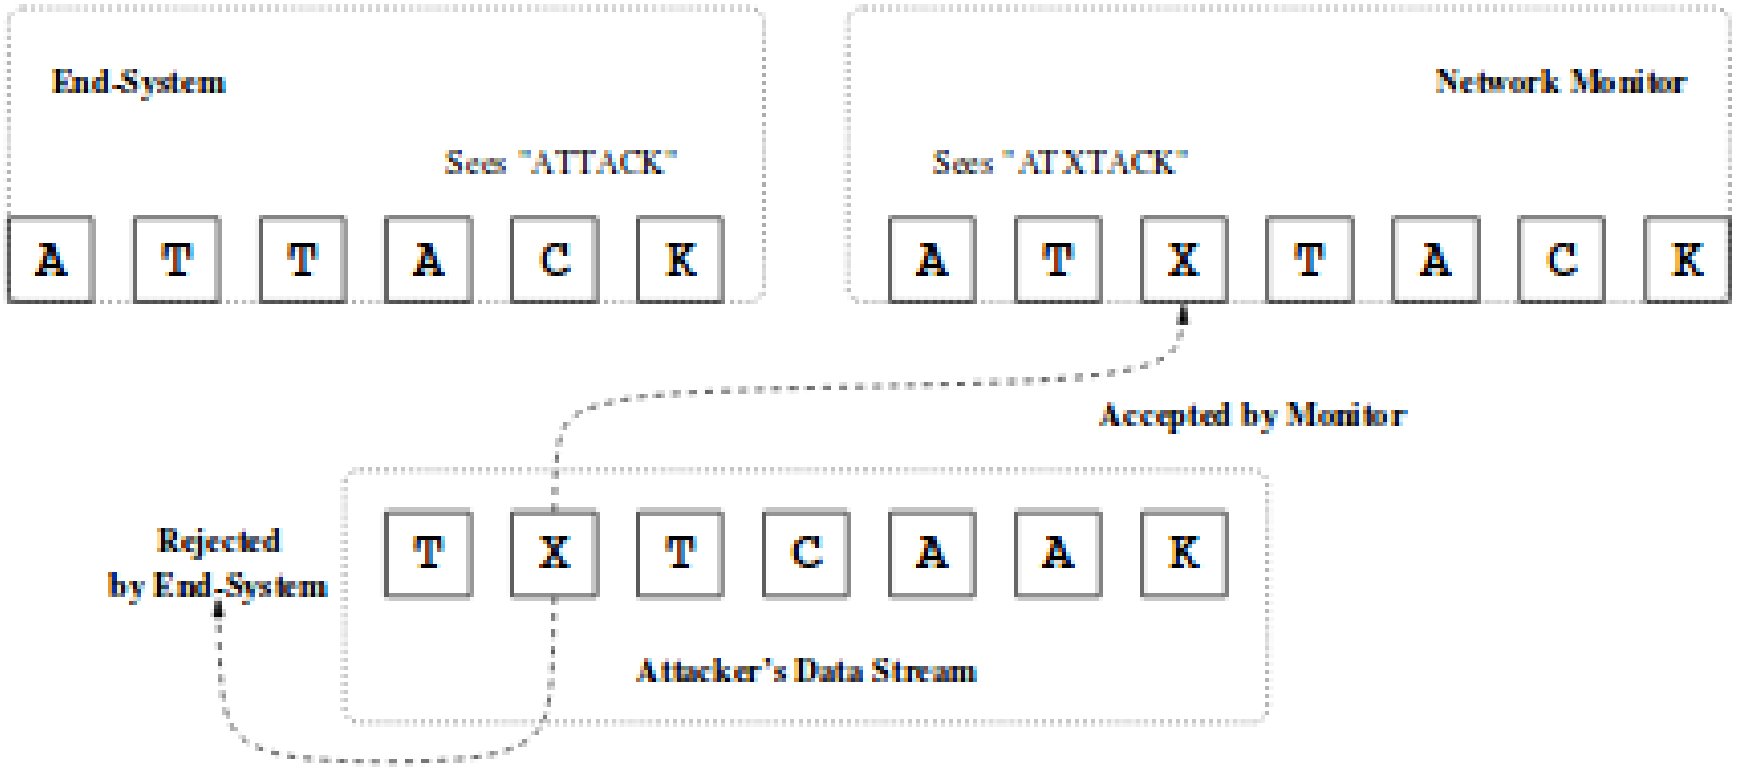

Figure 4: Insertion of the letter 'X'

Figure from Ptacek and Newsham

A not so useful distinction

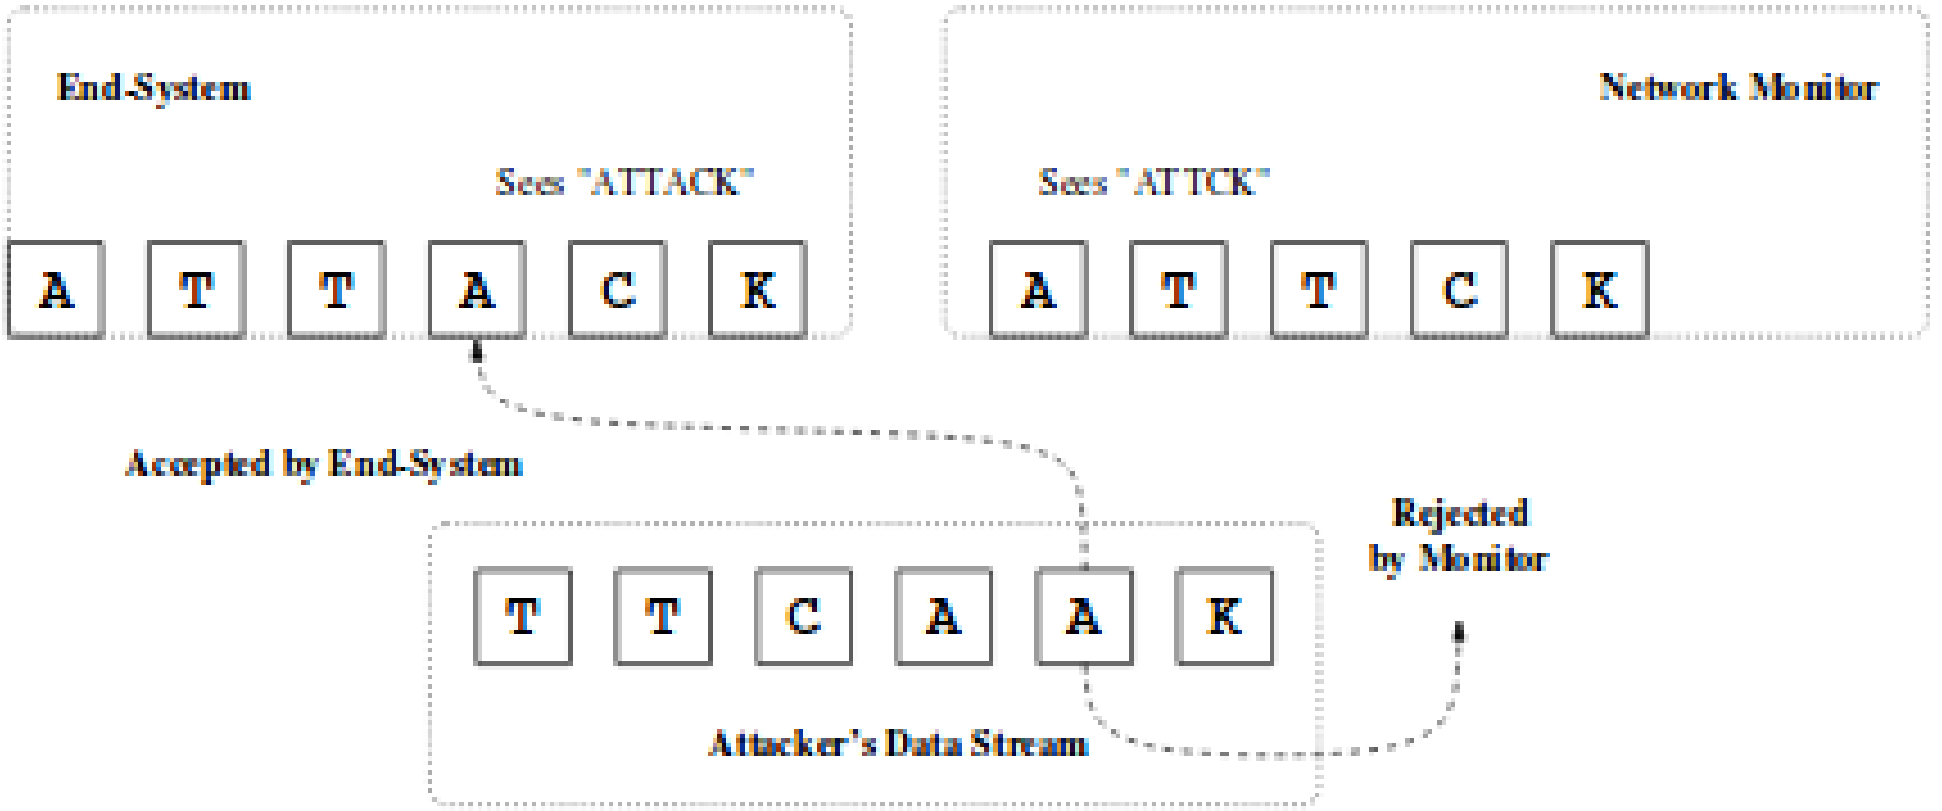

Figure 5: Evasion of the letter 'A'

Figure from Ptacek and Newsham

“Information only has meaning in that it is subject to interpretation”

–Computer Viruses, Theory and Experiments by Fred Cohen, 1984

“The only laws on the Internet are assembly and RFCs”

–Phrack 65 article by julia@winstonsmith.info

“Information is inherently physical”

--(*Lots of people said this, but see Richard Feynman's Lectures on Computation*)

IP reassembly

- Routers (or endhosts, if they want) can break IP packets up into fragments that the receiver has to reassemble
- Ambiguity in the way overlapping IP fragments are put back together into an IP packet
- All of the following images were plagiarized from:

  https://www.sans.org/reading-room/whitepapers/detection/ip-fragment-reassembly-scapy-33969

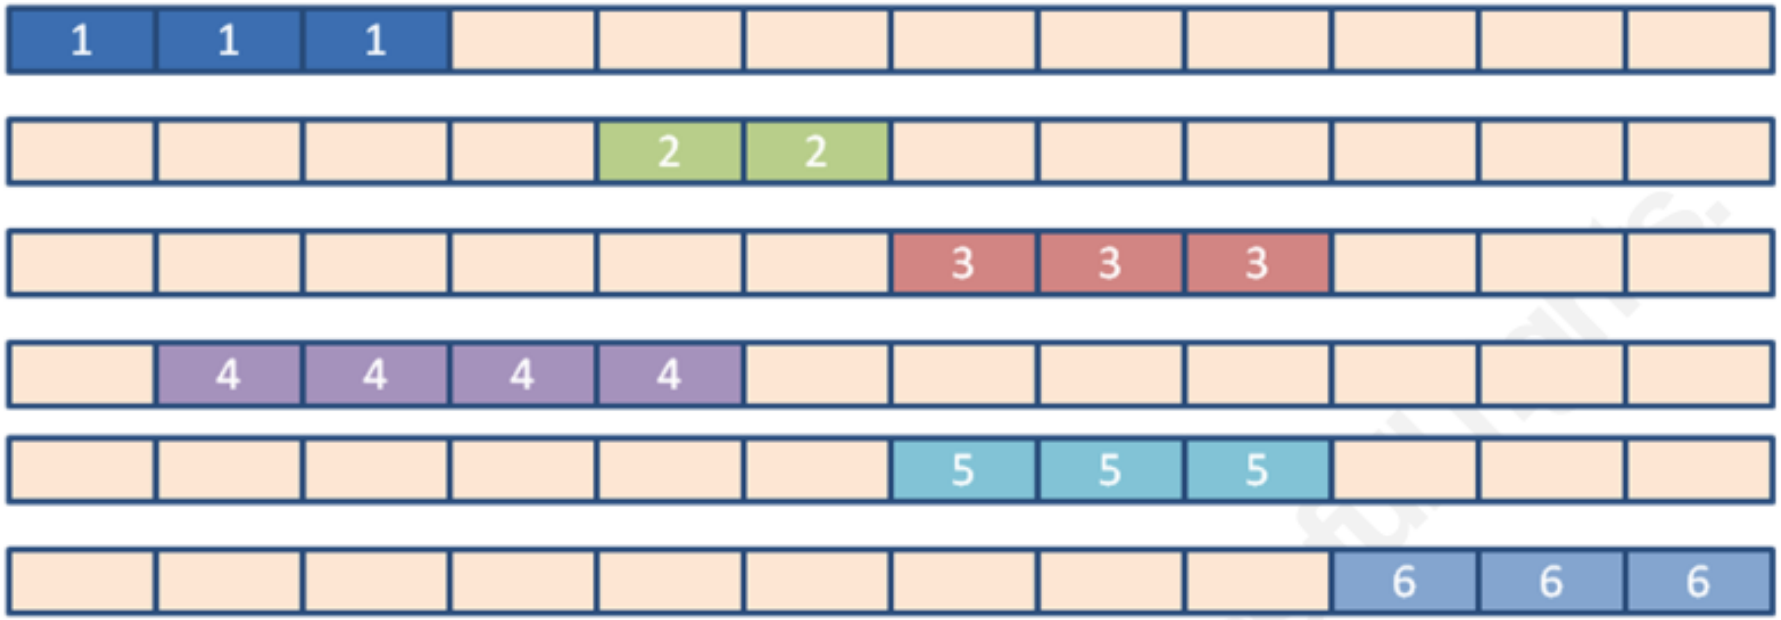

Figure 1: 6 Fragmented Packets (Shankar & Paxson, 2003)(Novak, 2005)

Reassembled using policy: First (Windows, SUN, MacOS, HPUX)

Reassembled using policy: Last/RFC791 (Cisco)

Reassembled using policy: Linux (Linux)

Reassembled using policy: BSD (AIX, FreeBSD, HPUX, VMS)

Reassembled using policy: BSD-Right (HP Jet Direct)

Figure 2: 5 Reassembly Methods (Shankar & Paxson, 2003)(Novak, 2005)

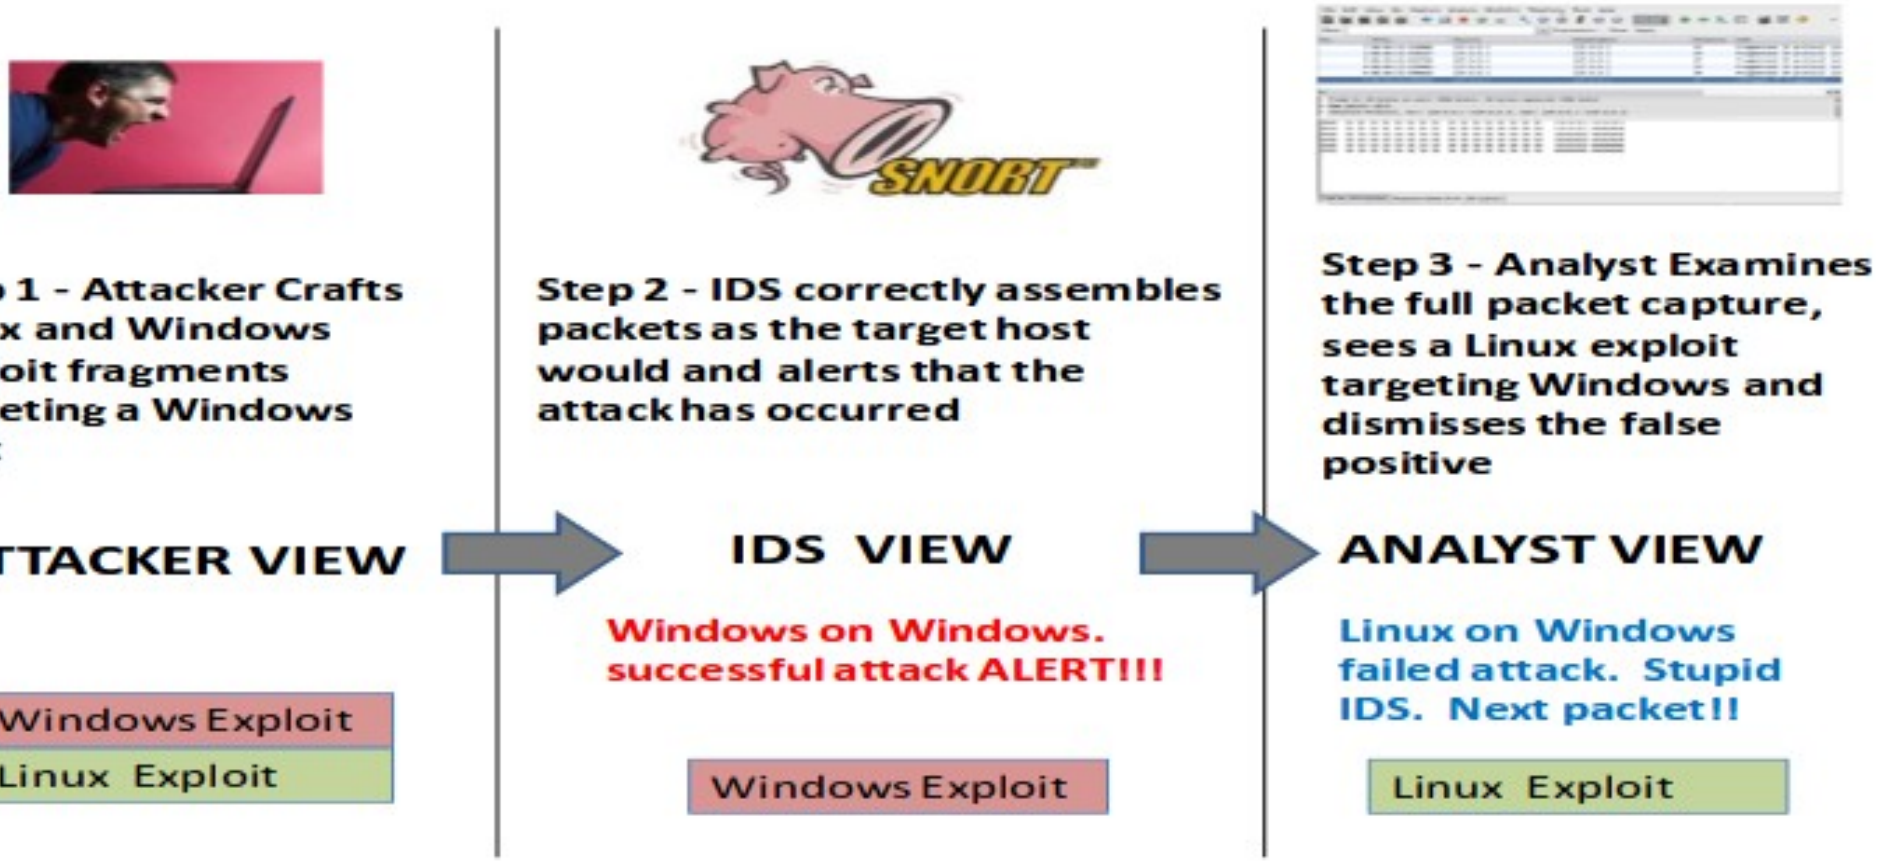

Figure 3: Views of the attacker, IDS and analyst

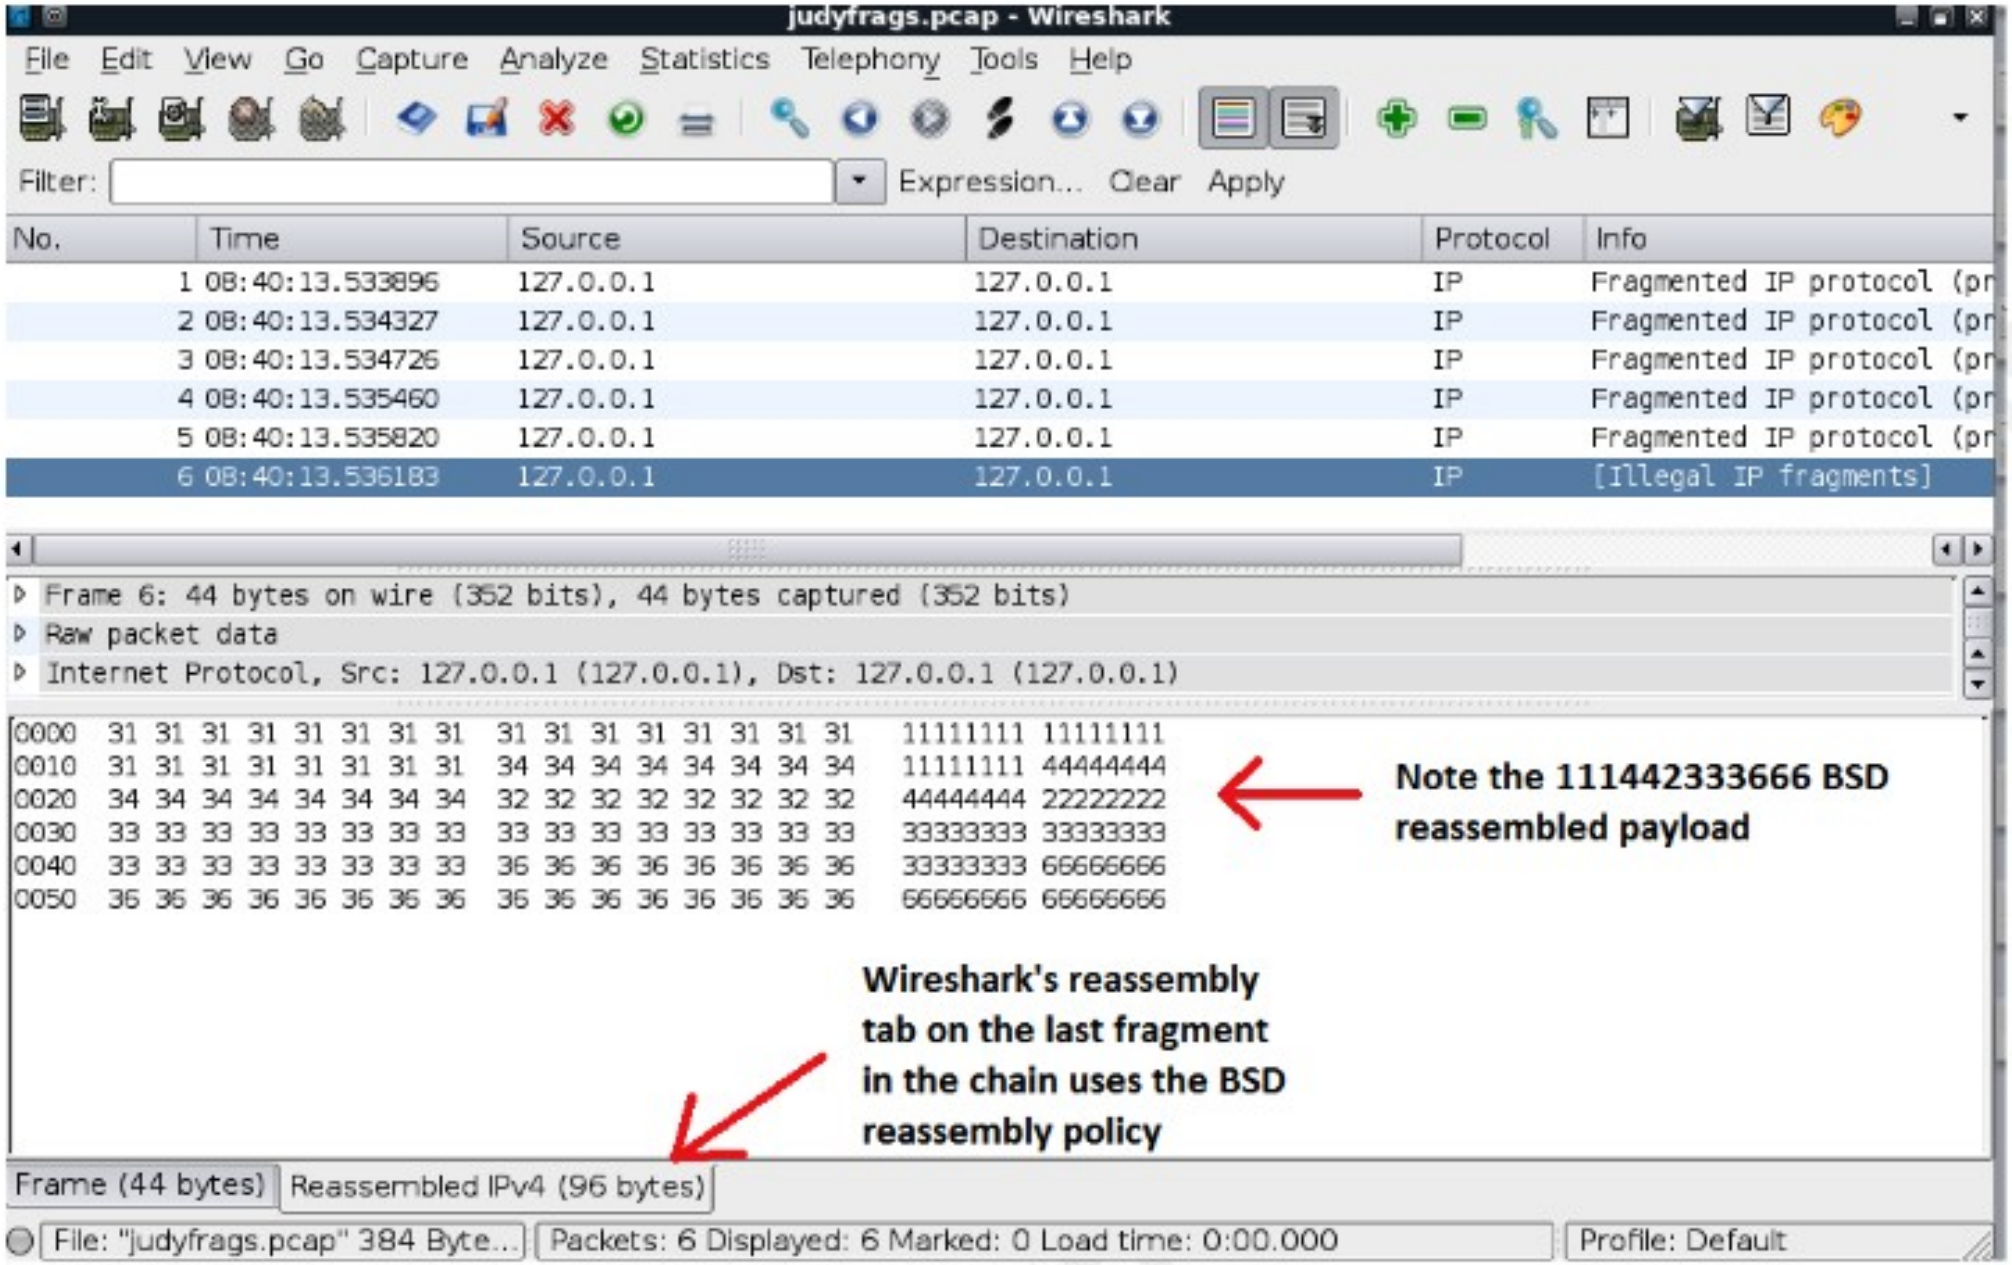

Figure 4: Wireshark uses BSD reassembly technique

TCP is even worse...

- From http://www.icir.org/vern/papers/TcpReassembly/

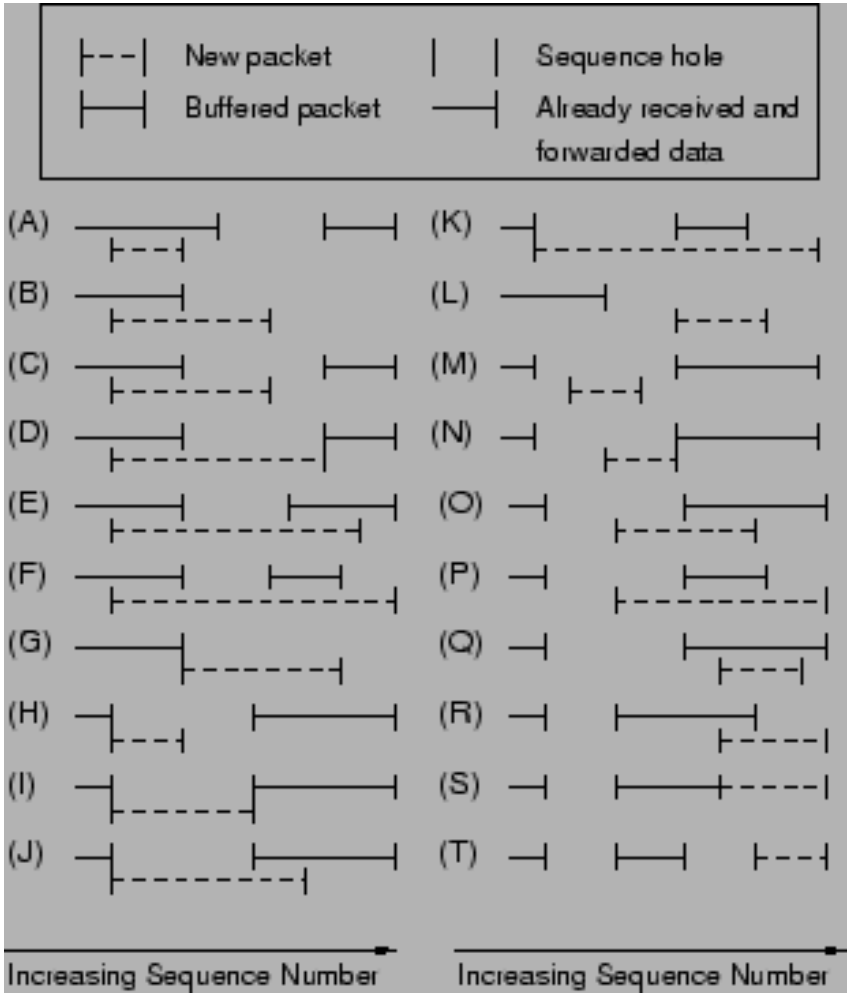

Another example: TTL limiting

- Victim is 10 hops away from you (the attacker)
- IDS is 7 hops away from you, 3 from the victim
- Send a SYN with TTL 64
- Get a SYN/ACK from the victim
- Send a RST with TTL 9
- Send an ACK with TTL 64
- Victim sees SYN, sends SYN/ACK, gets ACK, you have an open connection and can send them data
- IDS sees SYN in one direction, SYN/ACK in the other, then RST and assumes the connection was reset, ACK and all packets that follow (with data) are ignored by the IDS

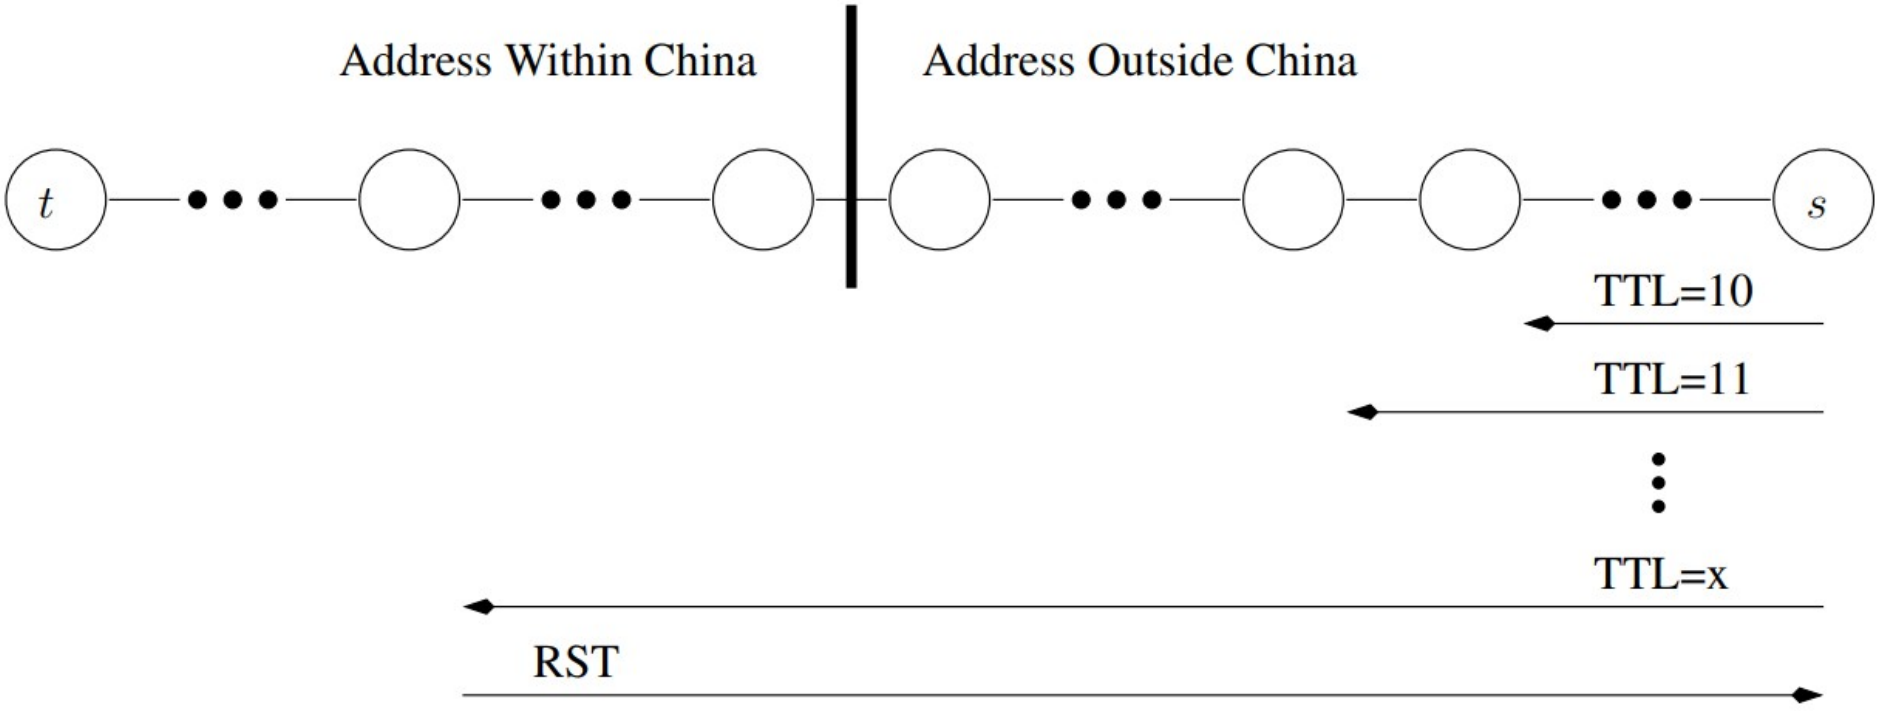

Figure 4: GFC router discovery using TTLs.

Reproduced from:
https://jedcrandall.github.io/concept_doppler_ccs07.pdf

A layer 7 example (XSS) due to Jeff Knockel

- Suppose “`<script>...</script>`” is blacklisted
- Use “`<script>...`” instead, many browsers will happily run the script anyway despite the missing closing tag
- Information only has meaning in that it is subject to interpretation
  - IDS interprets things one way, web browser another

Physical layer injection

- From https://www.usenix.org/legacy/events/woot11/tech/final_files/Goodspeed.pdf

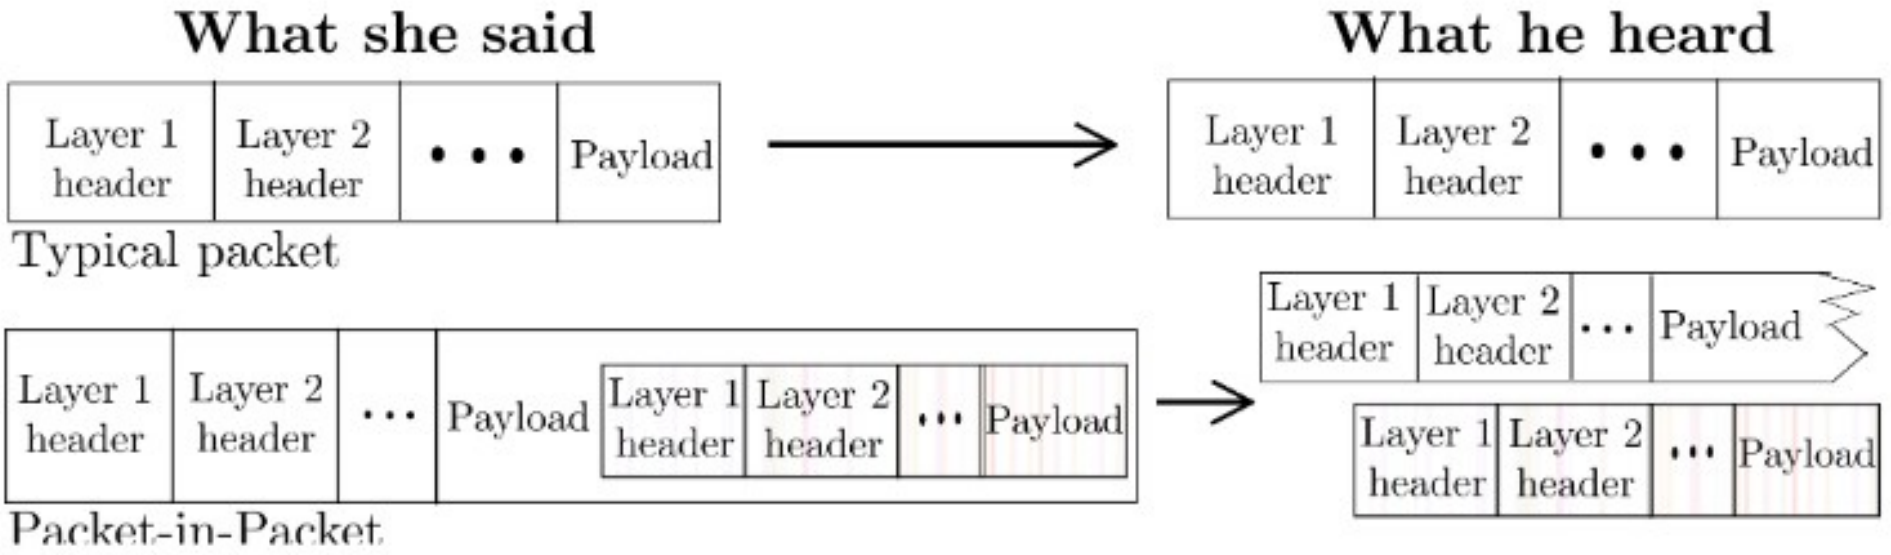

Figure 2: A typical packet's interpretation contrasted with that of a PIP.

Denial-of-Service (DoS) for IDS

- Exhaust the IDS's resources in some way
  - CPU
  - Memory
  - Bandwidth
- Fail-open (just let stuff through) *vs.* fail-closed (slow down the network)
- Example: On accident, “Tony” brought down the UNM Computer Science Dept. network
- Other examples

DoS in general

- Exhaust some kind of resource, *e.g.*:
  - Optimistic ACK to exhaust bandwidth
    - See https://homes.cs.washington.edu/~tom/pubs/CCR99.pdf
  - PING of death (large PING) causes crash
  - Exhaust CPU in layer 7
  - More examples: http://www.isi.edu/~mirkovic/bench/attacks.html
  - SYN flood: Older hosts had either a fixed amount of half-open connections they could keep track of or no limitations at all, attack is to send lots of SYNs and never ACK or RST
    - Defenses: SYN backlog policies and SYN cookies

SYN cookies and SYN backlogs

- SYN cookies
  - Special kind of SYN/ACK
  - See https://cr.yp.to/syncookies.html
  - Can confirm ACK number and reconstruct the necessary state for a connection without having kept any state after sending the SYN cookie
- SYN backlog examples
  - Linux reserves ½, ¼, 1/8th, and so on for successively older SYNs, prunes 5 times a second
  - FreeBSD has 512 buckets of 30, you can't predict what bucket you fall into (in theory)

Coming up...

- Port scans, off-path attacks, and DNS
- BGP and BGP attacks
- Examples of nation-scale NIDS systems (GFW, TSPU, *etc.*)

Resources

- Ptacek and Newsham, Insertion Evasion and Denial of Service: Eluding Network Intrusion Detection