Gotta catch 'em all!

... a field test of portable gigabit taps

SYN-bit
deep traffic analysis

Sake Blok
Relational therapist for computer systems
sake.blok@SYN-bit.nl

How it started…

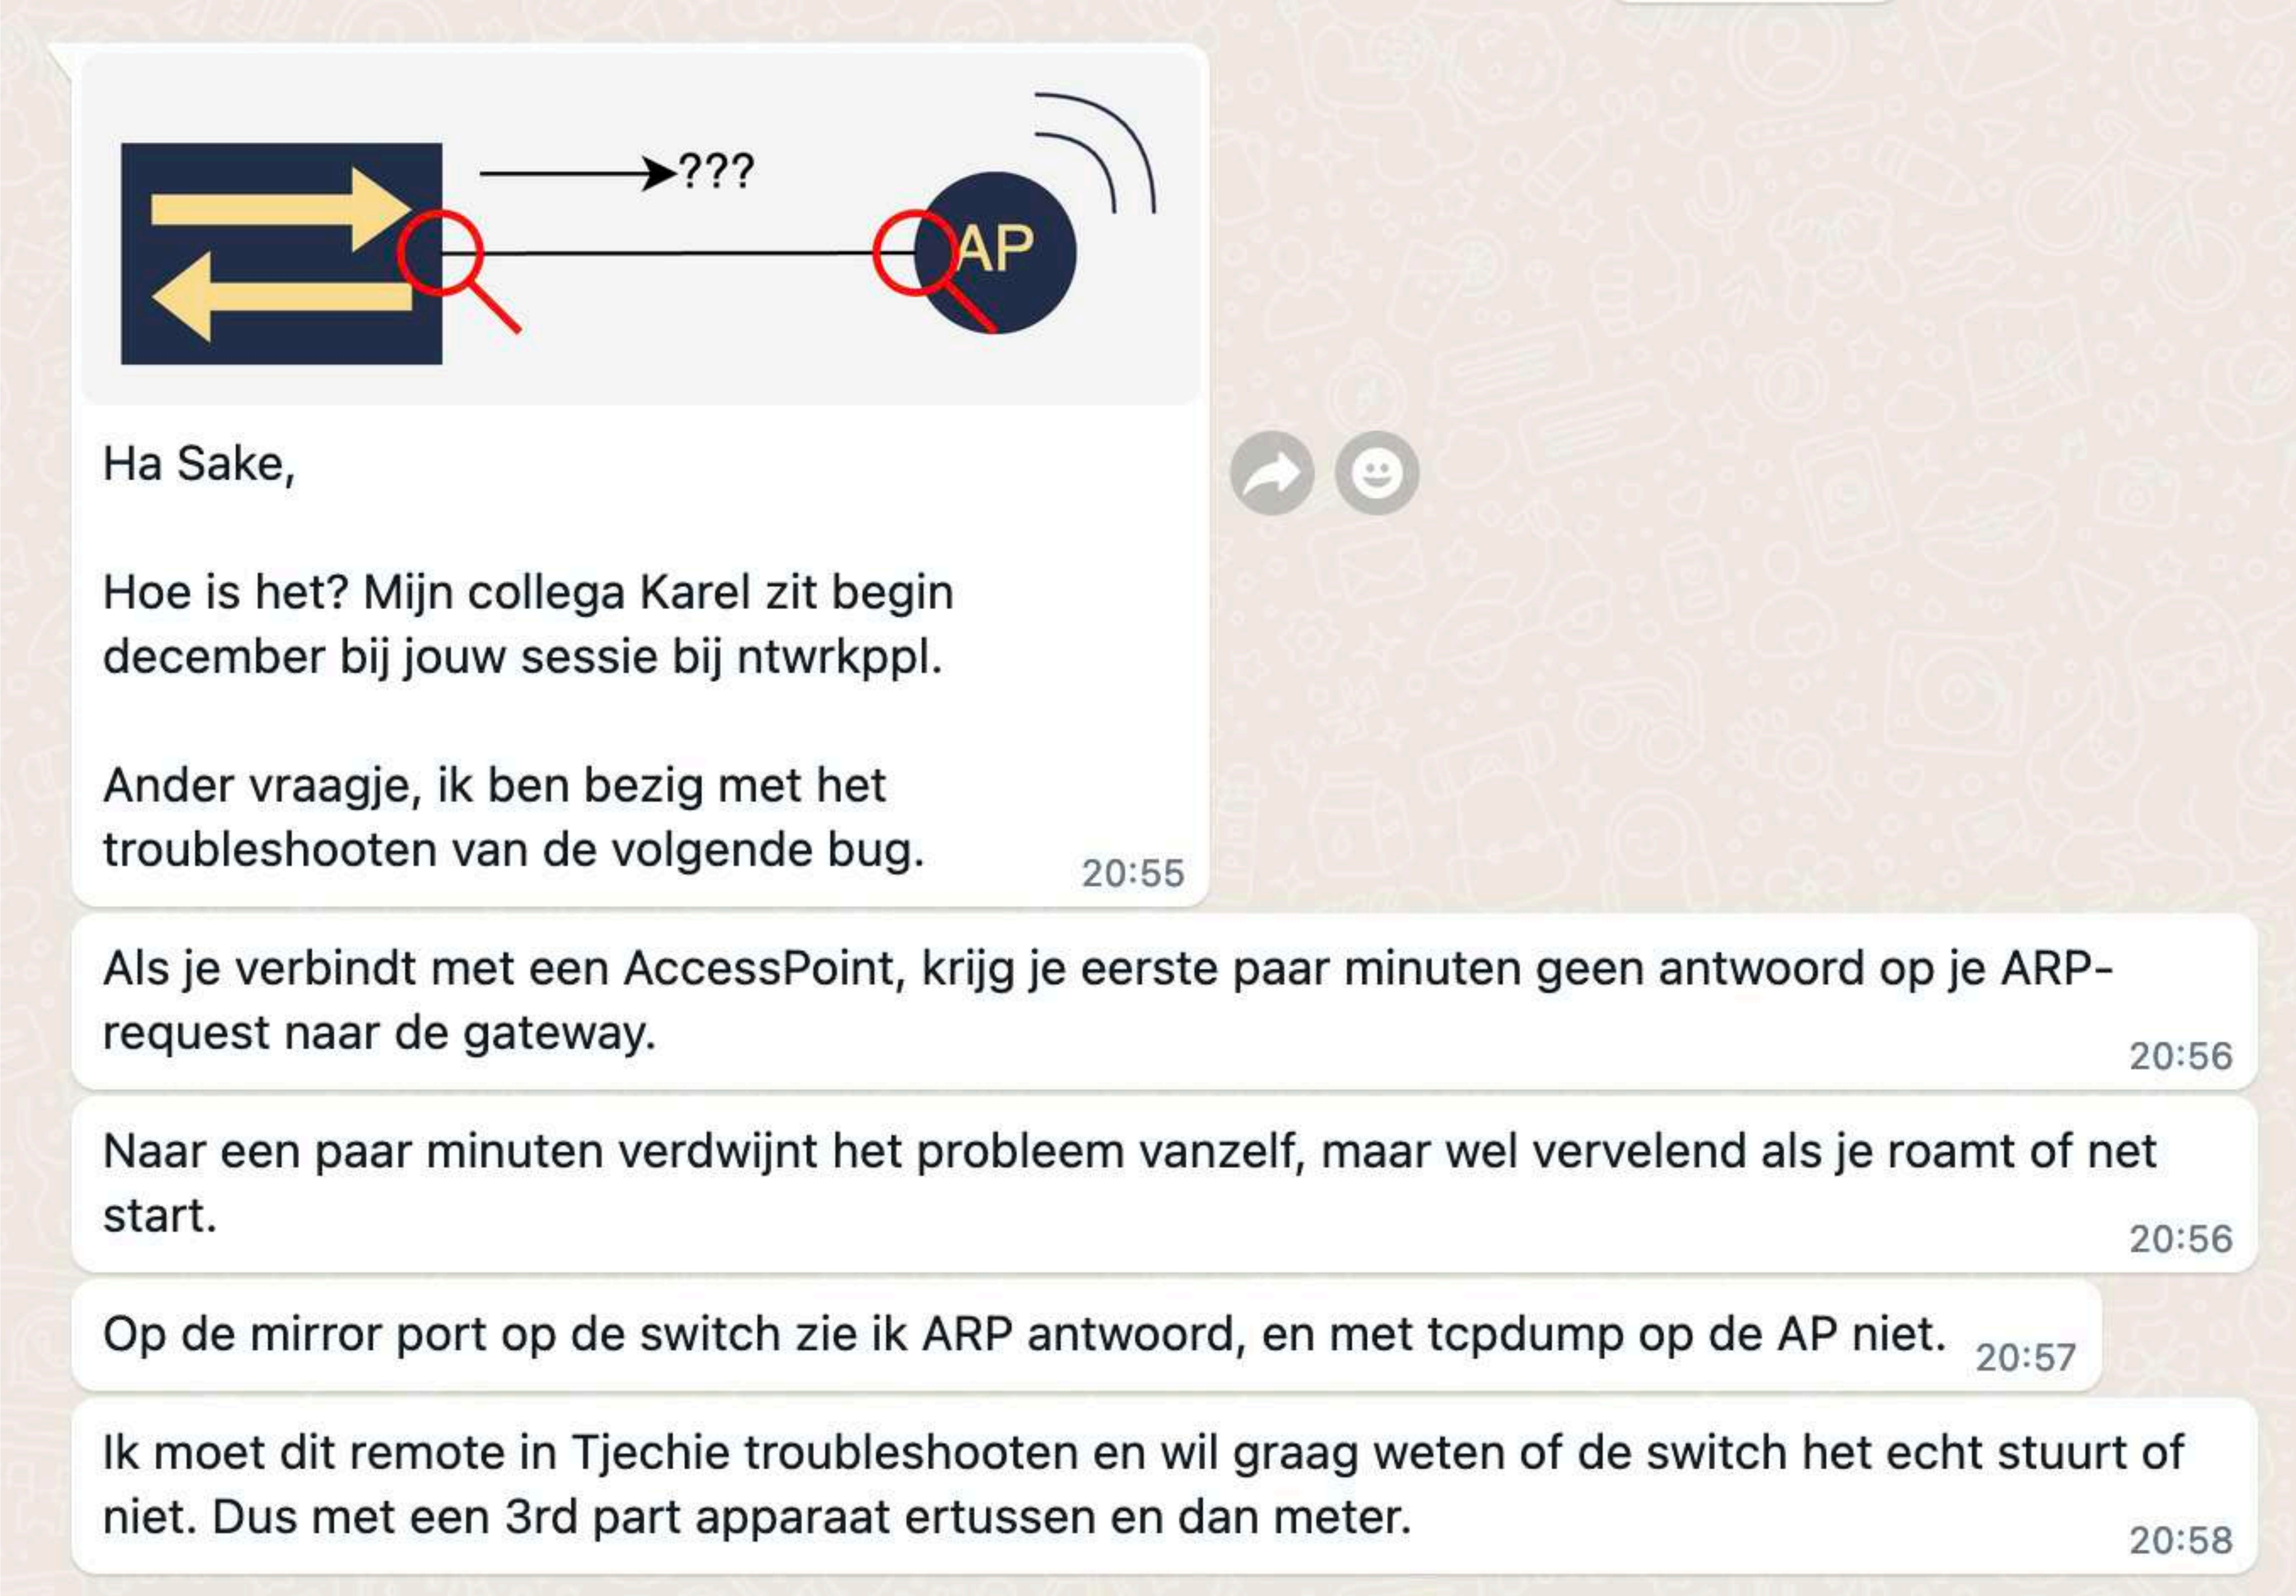

Ha Sake,

Hoe is het? Mijn collega Karel zit begin december bij jouw sessie bij ntwrkppl.

Ander vraagje, ik ben bezig met het troubleshooten van de volgende bug. 20:55

Als je verbindt met een AccessPoint, krijg je eerste paar minuten geen antwoord op je ARP-request naar de gateway. 20:56

Naar een paar minuten verdwijnt het probleem vanzelf, maar wel vervelend als je roamt of net start. 20:56

Op de mirror port op de switch zie ik ARP antwoord, en met tcpdump op de AP niet. 20:57

Ik moet dit remote in Tjechie troubleshooten en wil graag weten of de switch het echt stuurt of niet. Dus met een 3rd part apparaat ertussen en dan meter. 20:58

$ whoami

EURONET INTERNET

ABN·AMRO

ion ip

SYN-bit
deep traffic analysis

SYN-bit

deep traffic analysis

Application and network troubleshooting

Protocol and packet analysis

Training (Wireshark, TCP, SSL)

www.SYN-bit.nl

Agenda

- Many ways to capture packets
- We want a portable troubleshooting TAP!
- All TAPs are created equal, but...
- fast, faster, fastest
- Review & Summary
- Q&A

Many ways to capture packets

- Using the existing infrastructure
  - On one of the endpoints
  - On one of the intermediate devices
  - (R)SPAN / Mirror ports
  - ERSPAN
- Adding capture infrastructure
  - Insert a bridging (linux) system
  - Insert a (real) HUB
  - Insert a switch with SPAN
  - Insert a homemade passive TAP
  - Insert a proper network TAP

https://www.flickr.com/photos/51428653@N06/4742250939/

Using the existing infrastructure

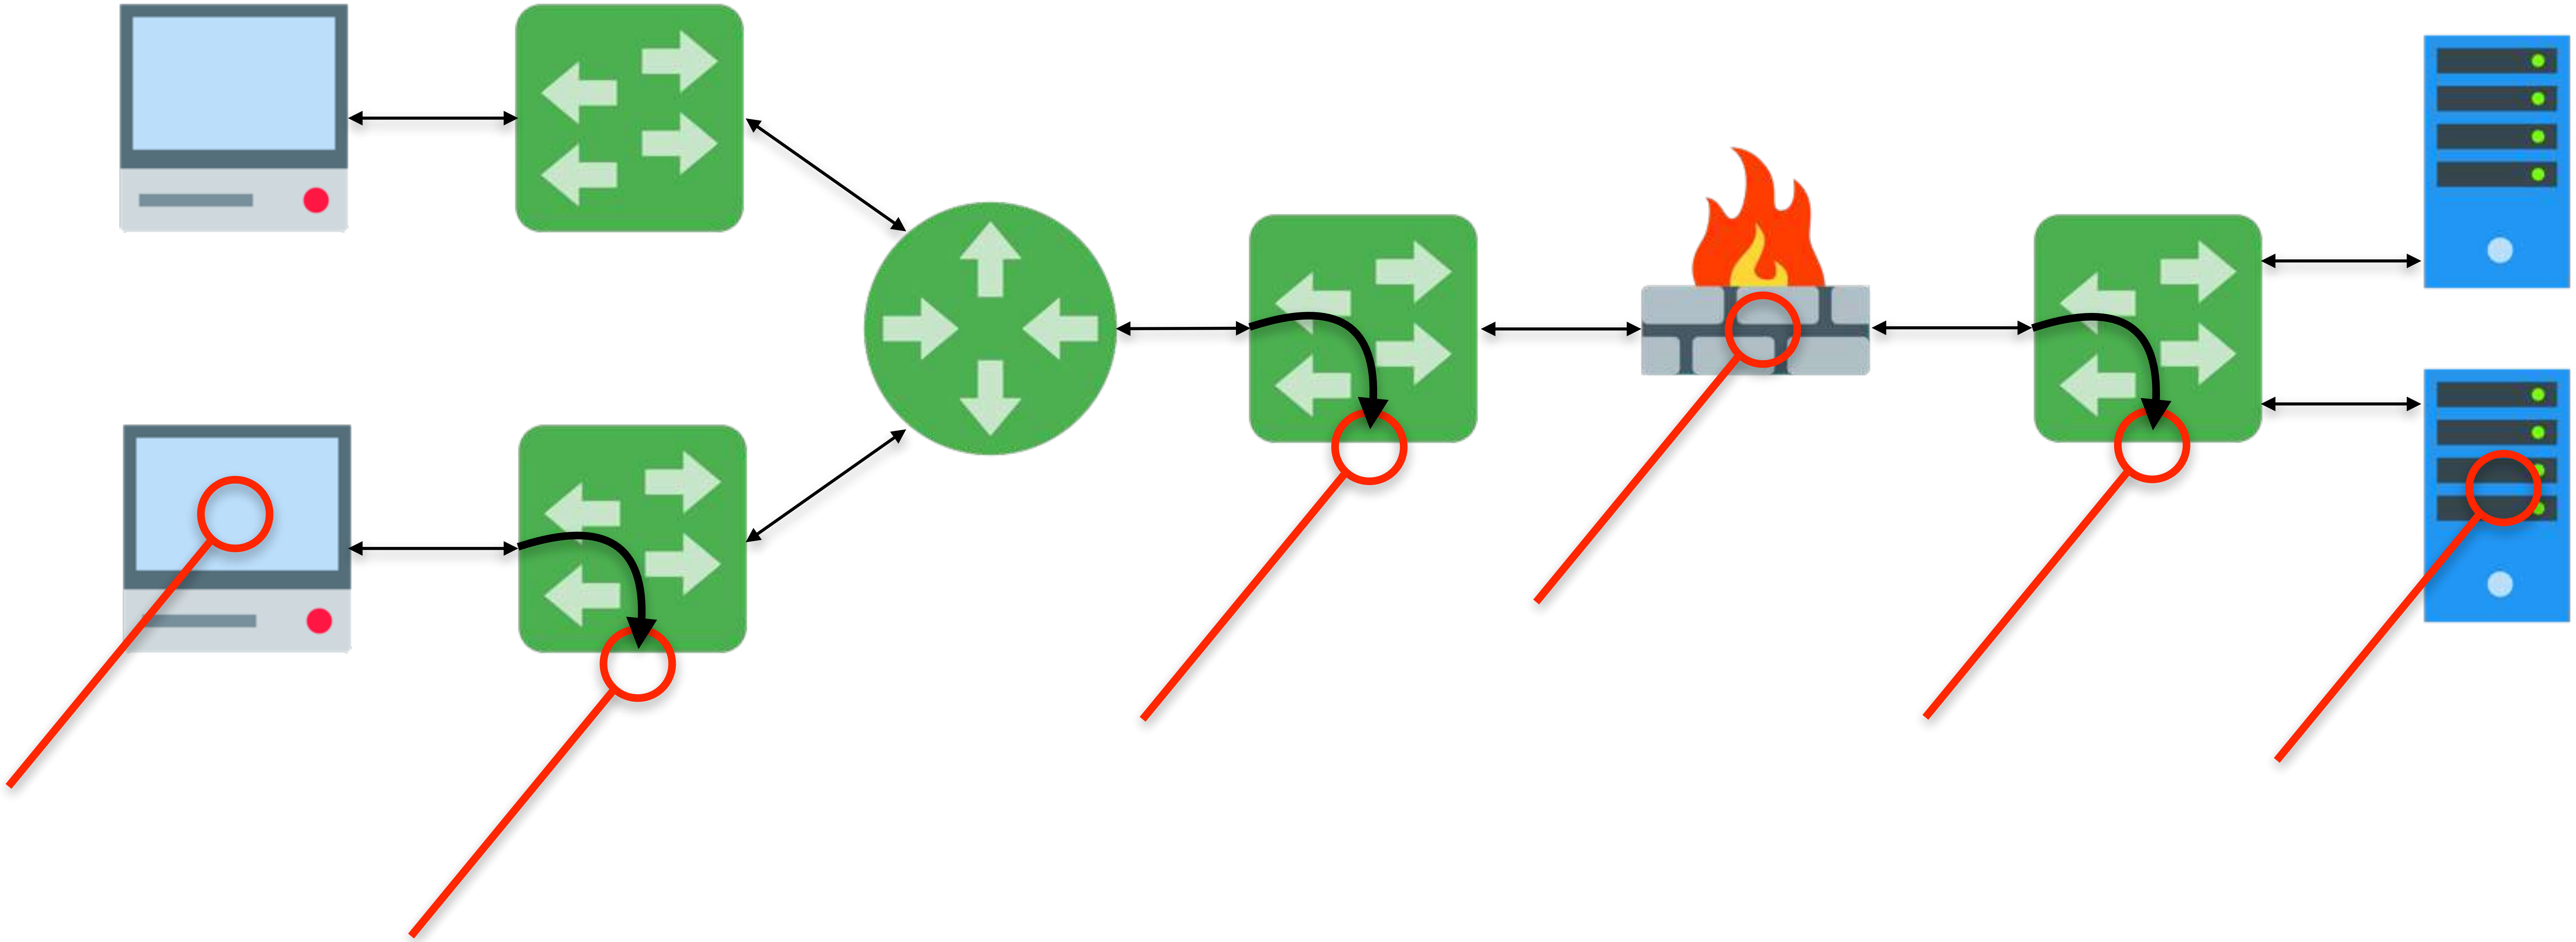

Capturing on an endpoint

- Pros
  - Quick and Easy
    - just (install and) use Wireshark or tcpdump
- Cons
  - Influences the endpoint
    - CPU cycles and disk IO
  - ***Capture is done in the kernel in the middle of the stack***
    - So some processing done for ingress and to do for egress traffic
    - Padding, offloading features (checksum, segmenting, etc)
  - ***3rd party software might be in the way***
    - Host firewall, VPN etc
  - Not possible on all types of endpoints

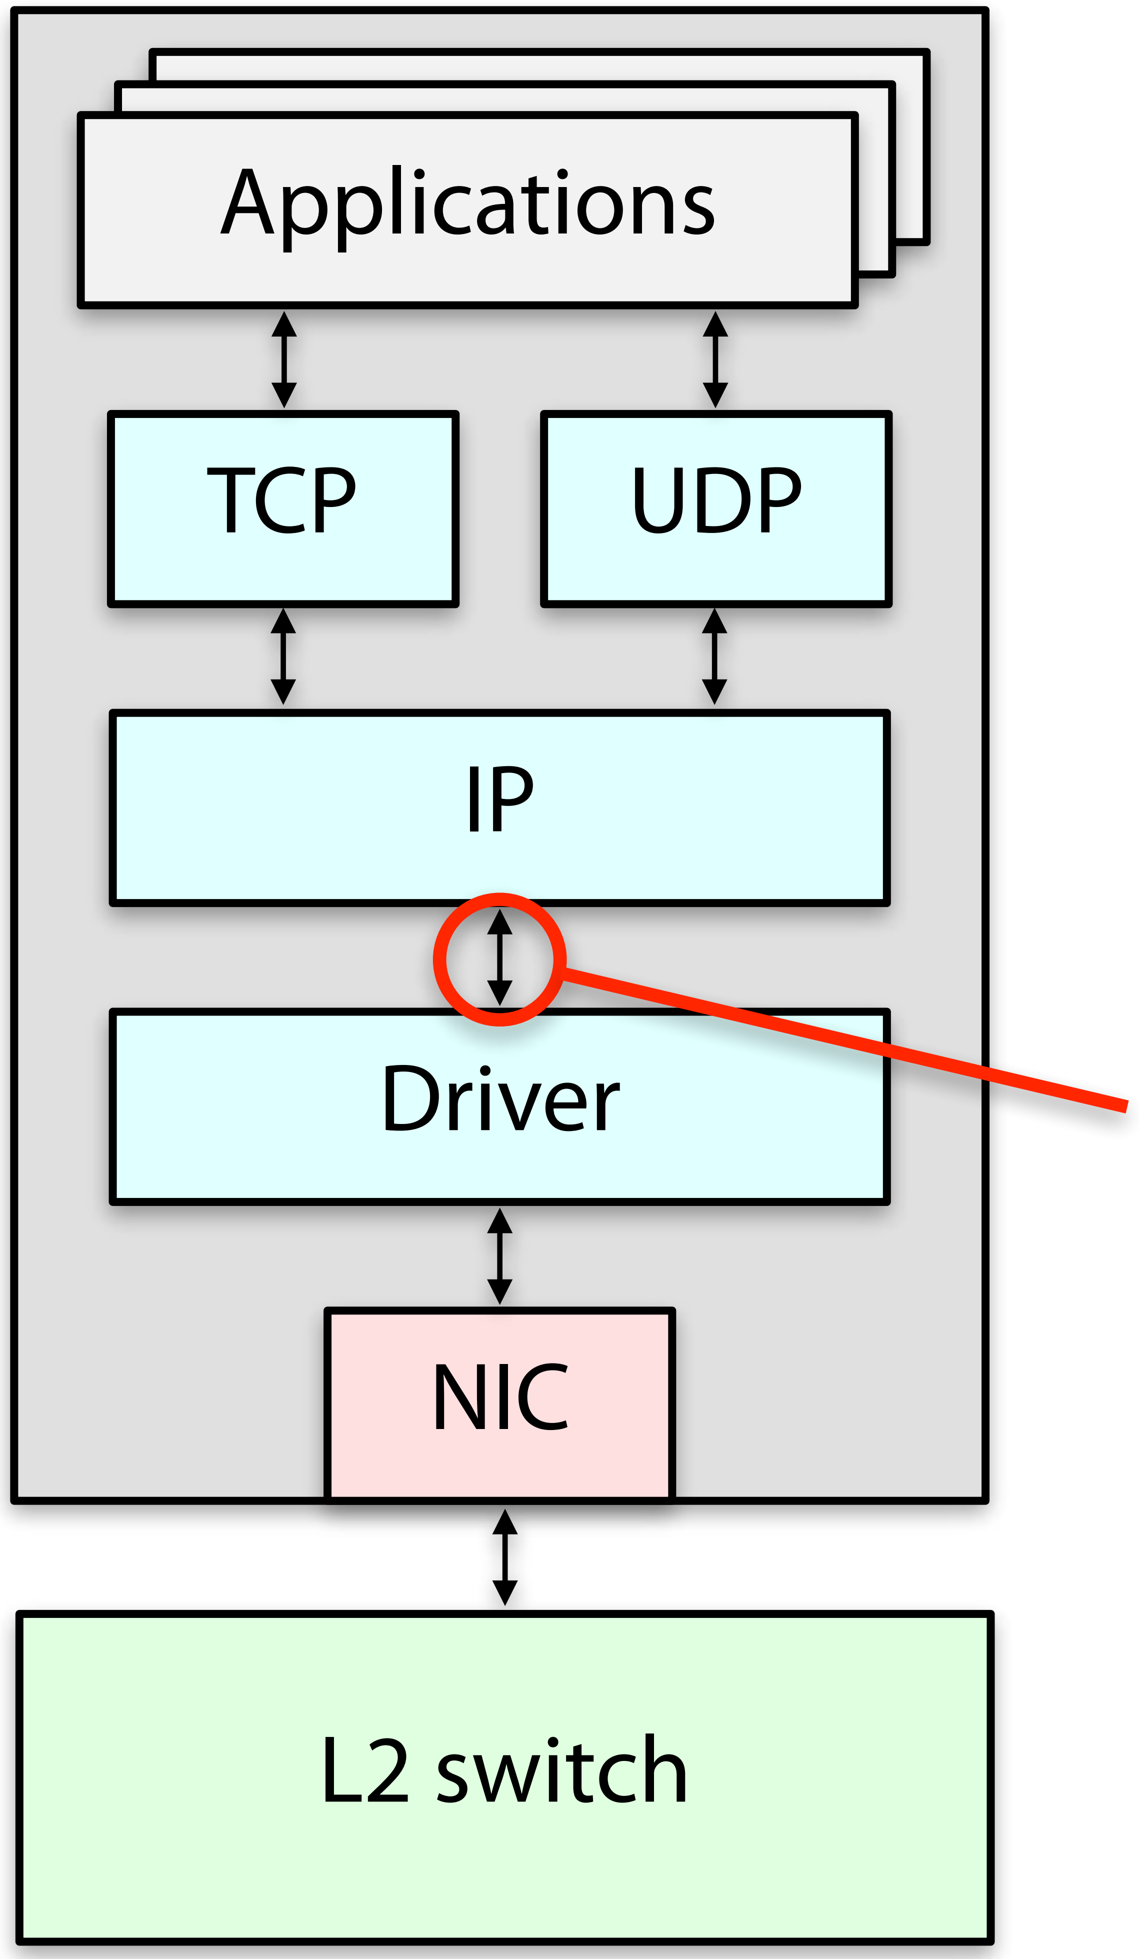

Capturing on an intermediate device

- Pros
  - Quick and Easy
    - Use the on-board capture capabilities (usually tcpdump)
- Cons
  - Influences the intermediate device
    - CPU cycles and disk IO
  - Limited disk-space for capturing
  - ***Capture is done at some point in the device, not on the NIC***
    - So some processing done for ingress and still to do for egress traffic
    - Padding, offloading features (checksum, segmenting, etc)
  - ***Pre/Post NAT differences***
  - ***Hardware offloading (on ASICs) results in missing packets***

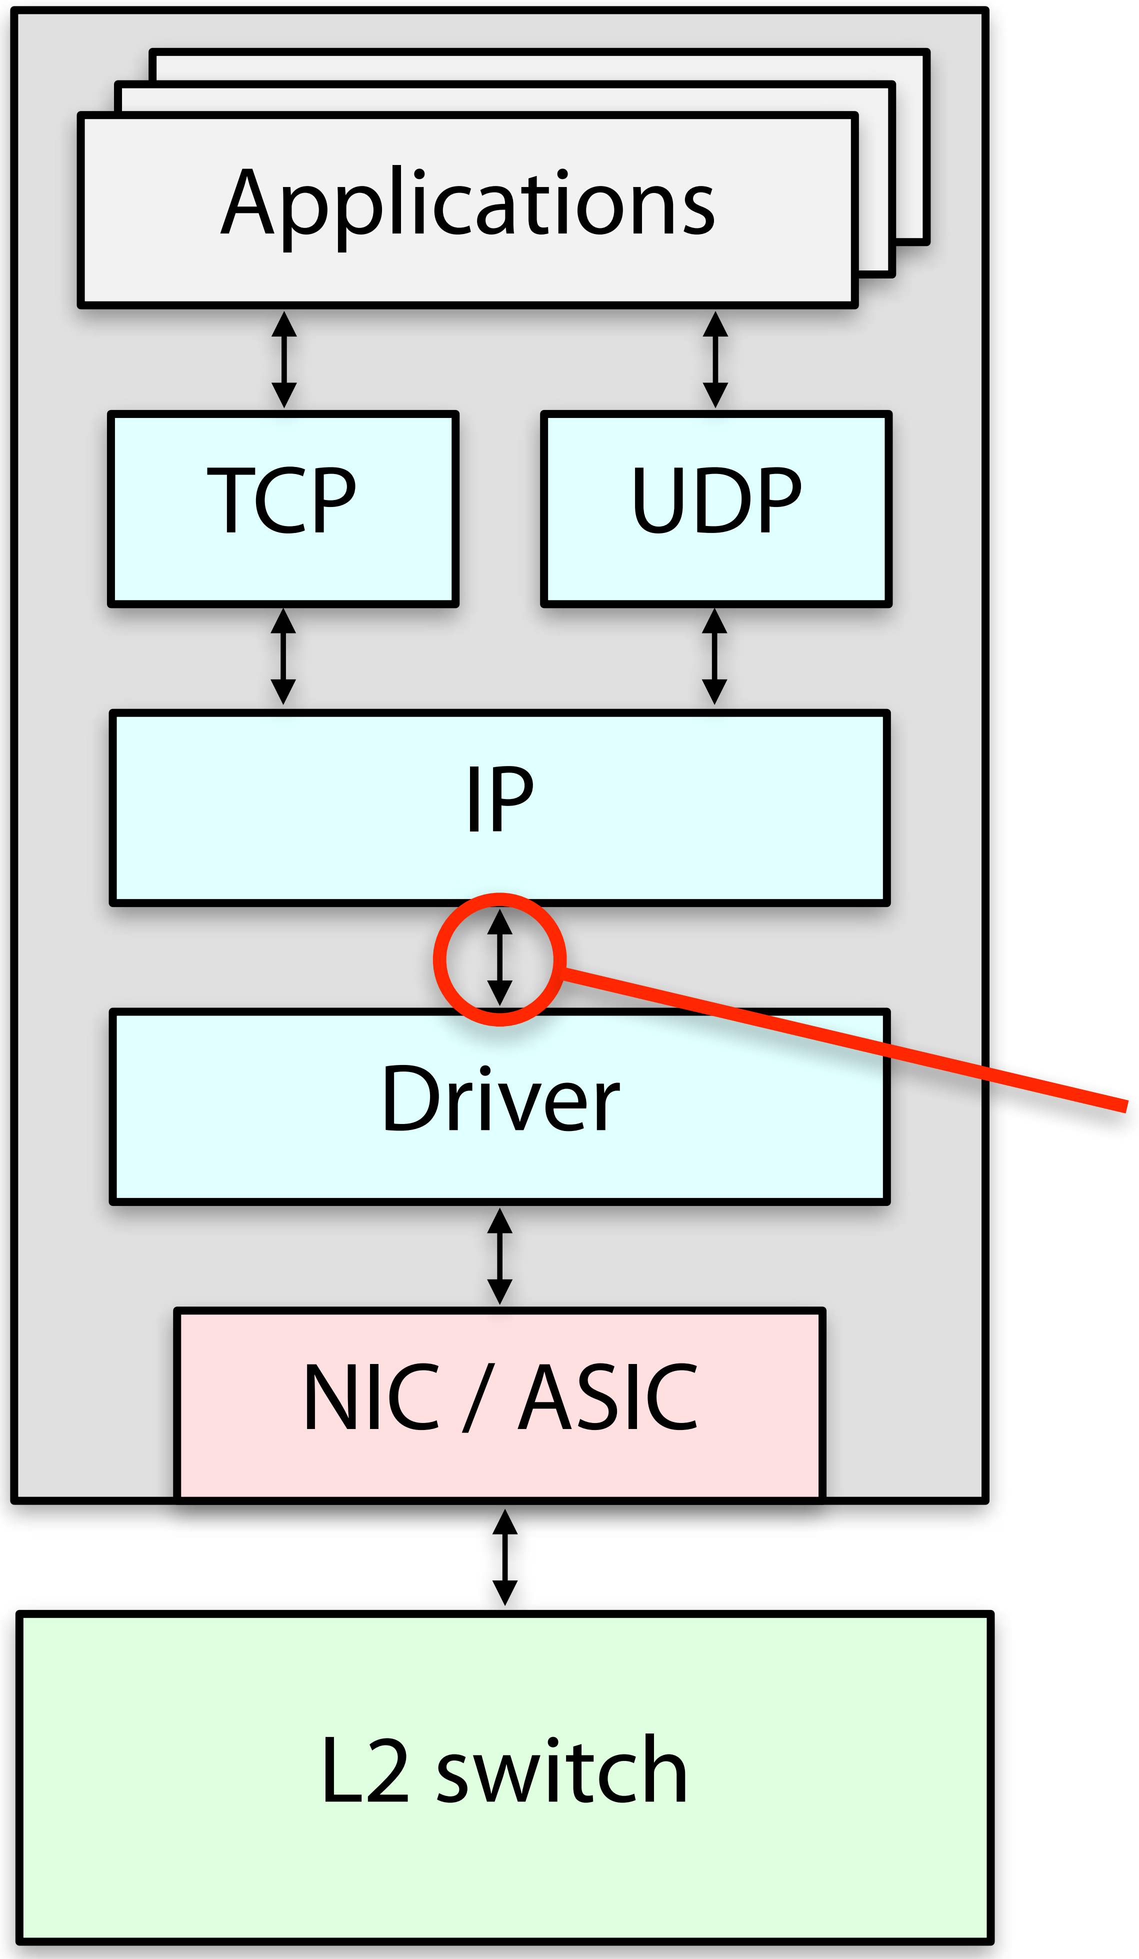

Capturing on a SPAN port

- Pros
  - Not influencing devices under suspicion
  - Usually easy to set up
- Cons
  - Switch configuration access not always easy to get
  - **Switch-generated packets not always mirrored**
  - **Inconsistencies in mirroring in- and egress traffic**
    - Case ubiquity switch: tagged one way, native for the other
  - Oversubscription of capture interface
  - **Not capturing at switchport but somewhere in the datapath**

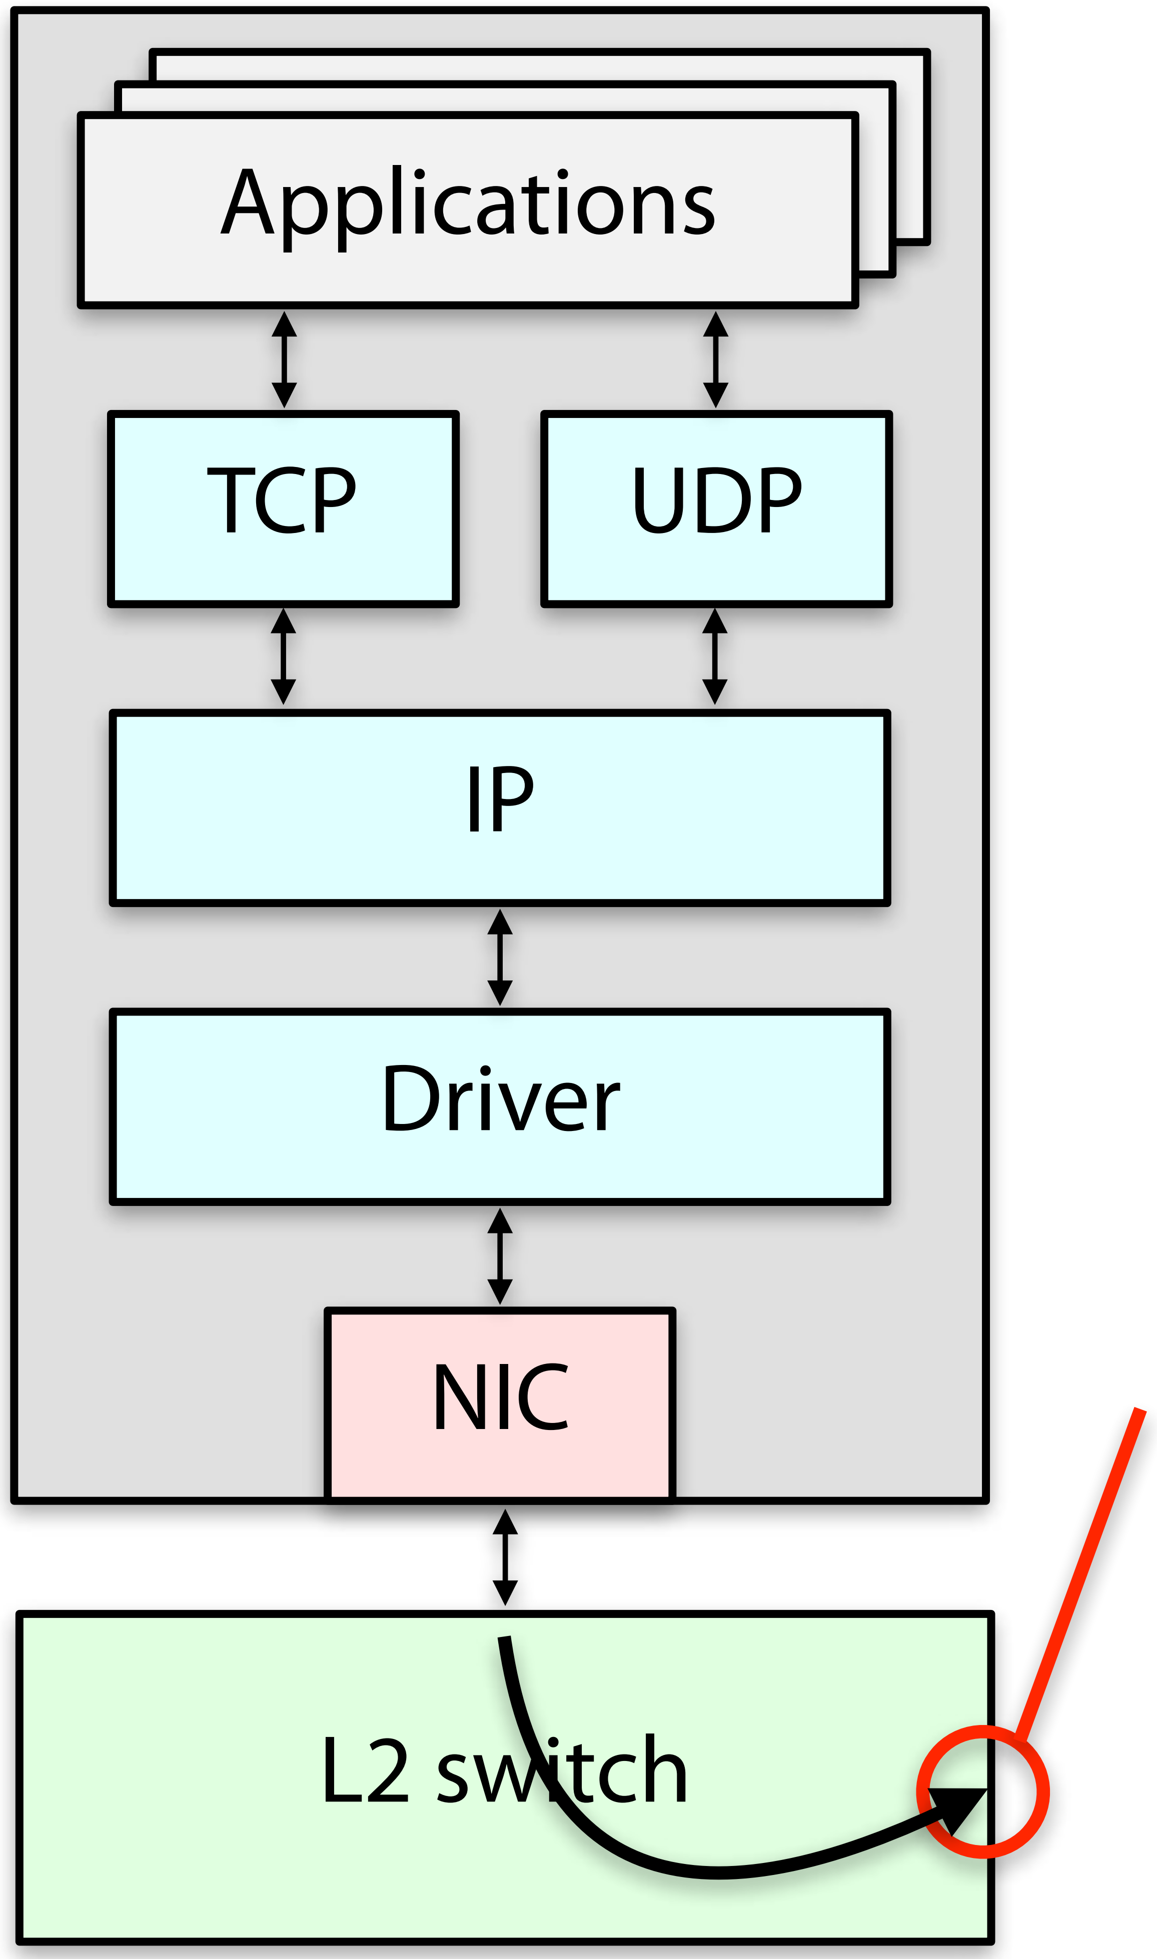

Capturing with ERSPAN

- Pros
  - Not influencing devices under suspicion
  - Usually easy to set up
  - Information on which devices are passed in a fabric
- Cons
  - Switch configuration access not always easy to get
  - ***Switch-generated packets not always mirrored***
  - Adding another data stream to the network
  - Oversubscription of capture interface or network
  - ***Not capturing at switchport but somewhere in the datapath***
  - Capture filters are much harder without decapsulation!

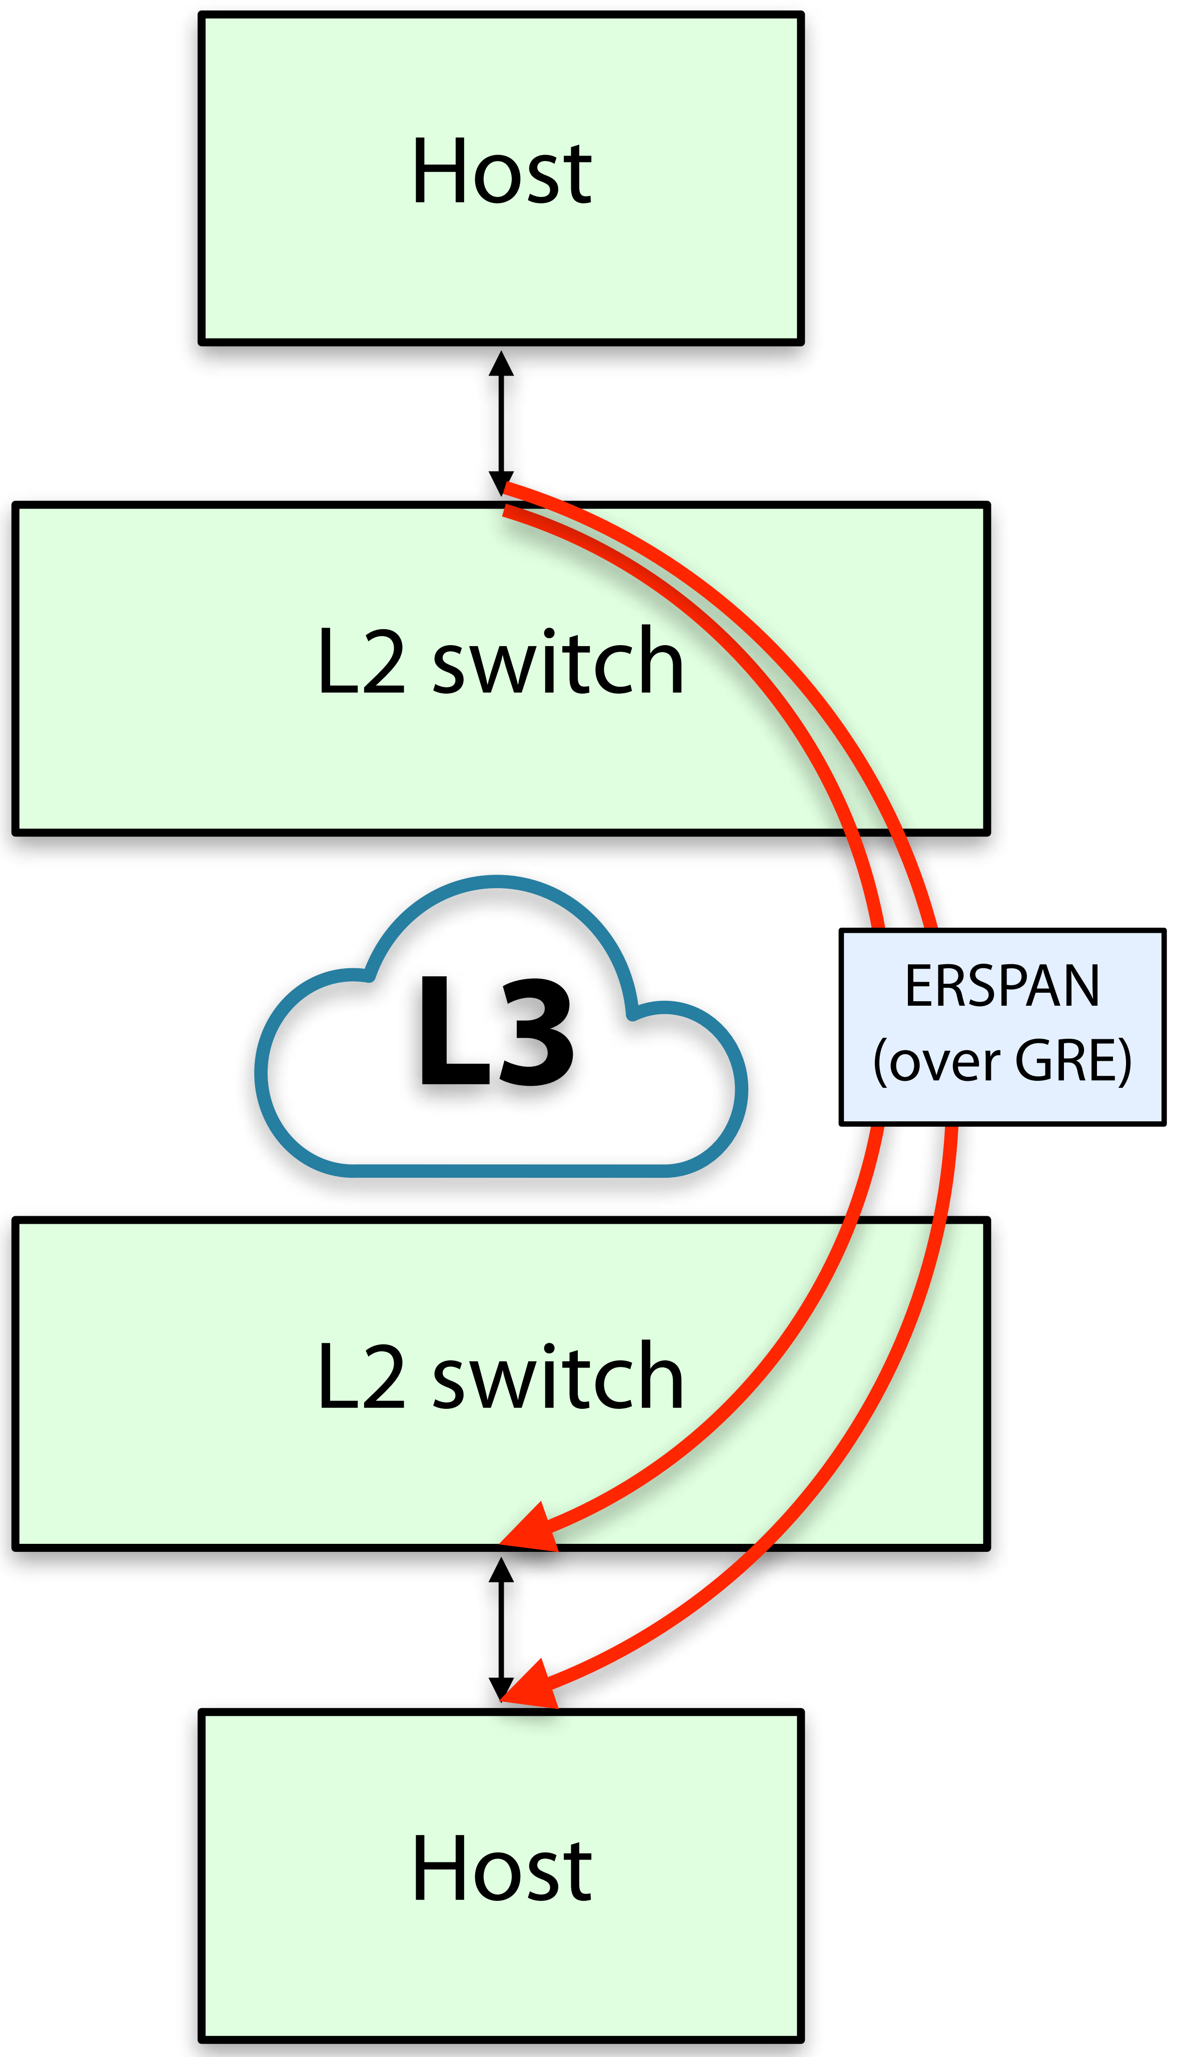

Real Cases

SMB session setup fails over new WAN link

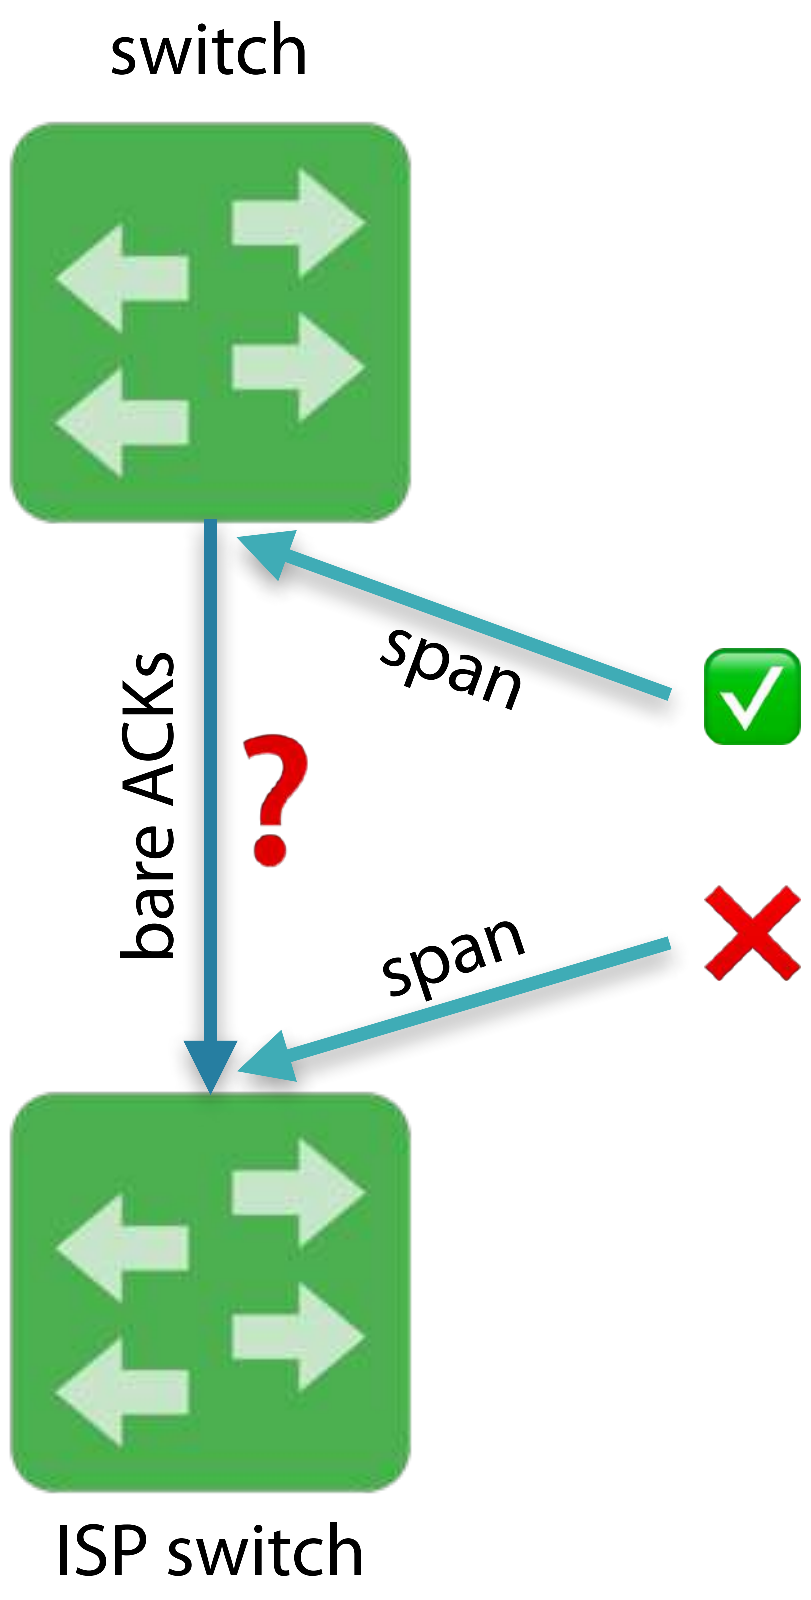

Phone sometimes ends up in guest vlan

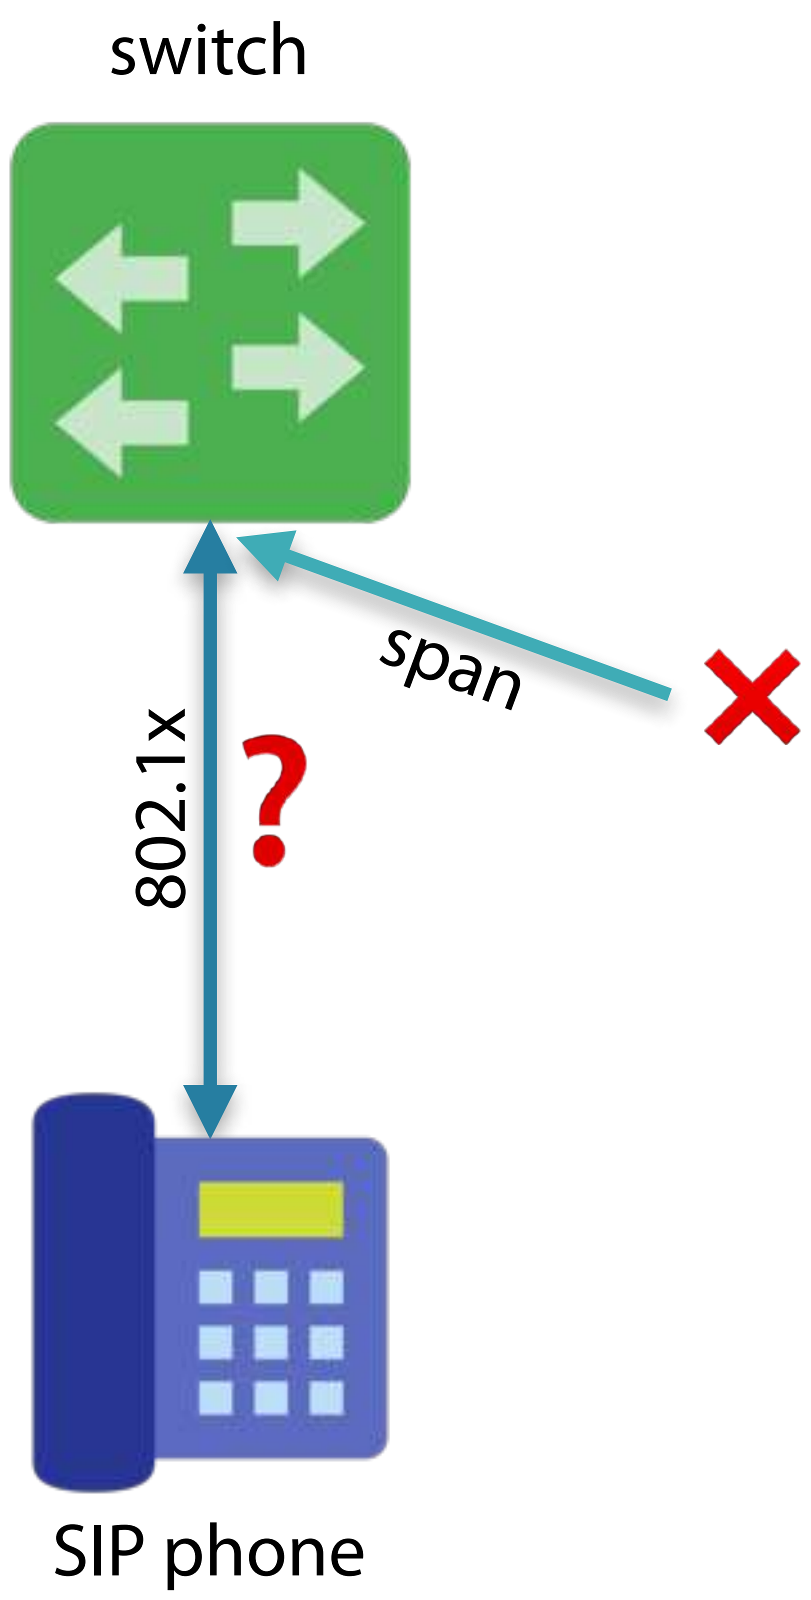

No connectivity after joining SSID

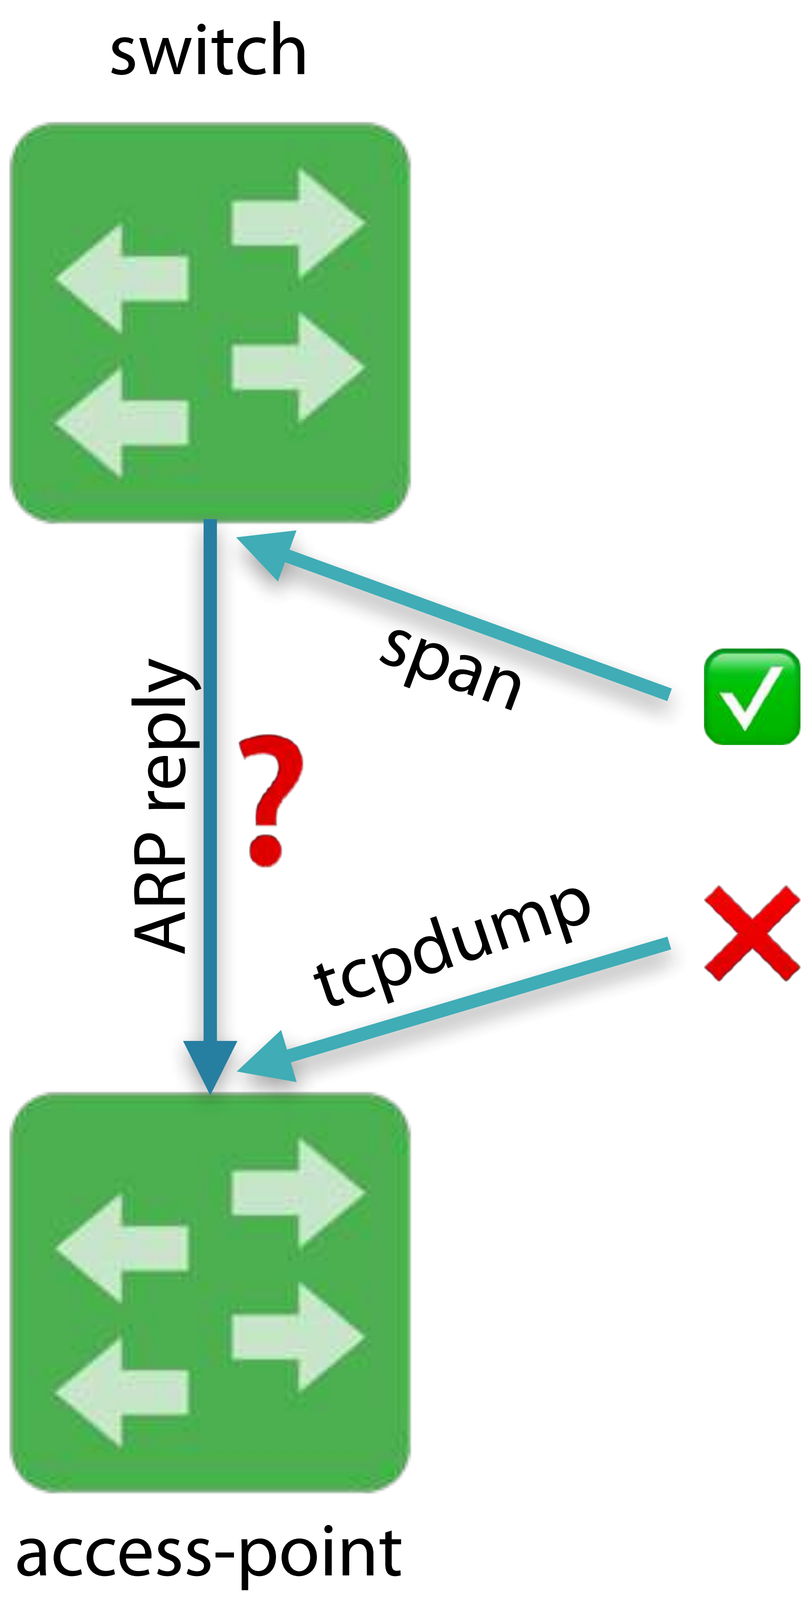

We need proper demarcation!

Baarle-Nassau / Baarle-Hartog

Ulicoten
HEERLE
Baarle Outdoor & Leisure
BOSCHHOVEN
N260
KLEIN-BEDAF
Oordeelsestraat
Niemandsland (enclave)
Baarleseweg
De Bremer
Bremer
VOSKE
GROOT-BEDAF
N135
N12
POPPEL
N639
Baarle-Nassau
Reutsche Loop
REUTH
LOVEREN
Baarle-Hertog
OUDE STRUMPT
Hoogstratensebaan
NEDERLAND
BELGIË
Poolsheining
HOOGEIND
EIKELENBOSCH
TOMMEL
RETH
Kievit
HAL
GORPEIND
SCHALUINEN
parkeerplaats Staatsbosbeheer
Hoogstratensebaan
Het Marksken
Marksken
HEGGE
HEIBRAEK
Straatloop
Ravels Gemeentehuis
ZUIDHEIKANT
GINHOVEN
Schouwloop
ZONDEREIGEN
Kempens Landschap - Wortel-kolonie
WEELDE
WEELDESESTRAAT
N118
De Aa
HEIKANT
PAMPA
BAARLE-NASSAU-GRENS
Weelde Depot
Vliegveld Weelde
Snuffelland
Pups en Zo
Google
STATION
EEL
TULDE
DE G
Lagen

Adding a packet capture infrastructure

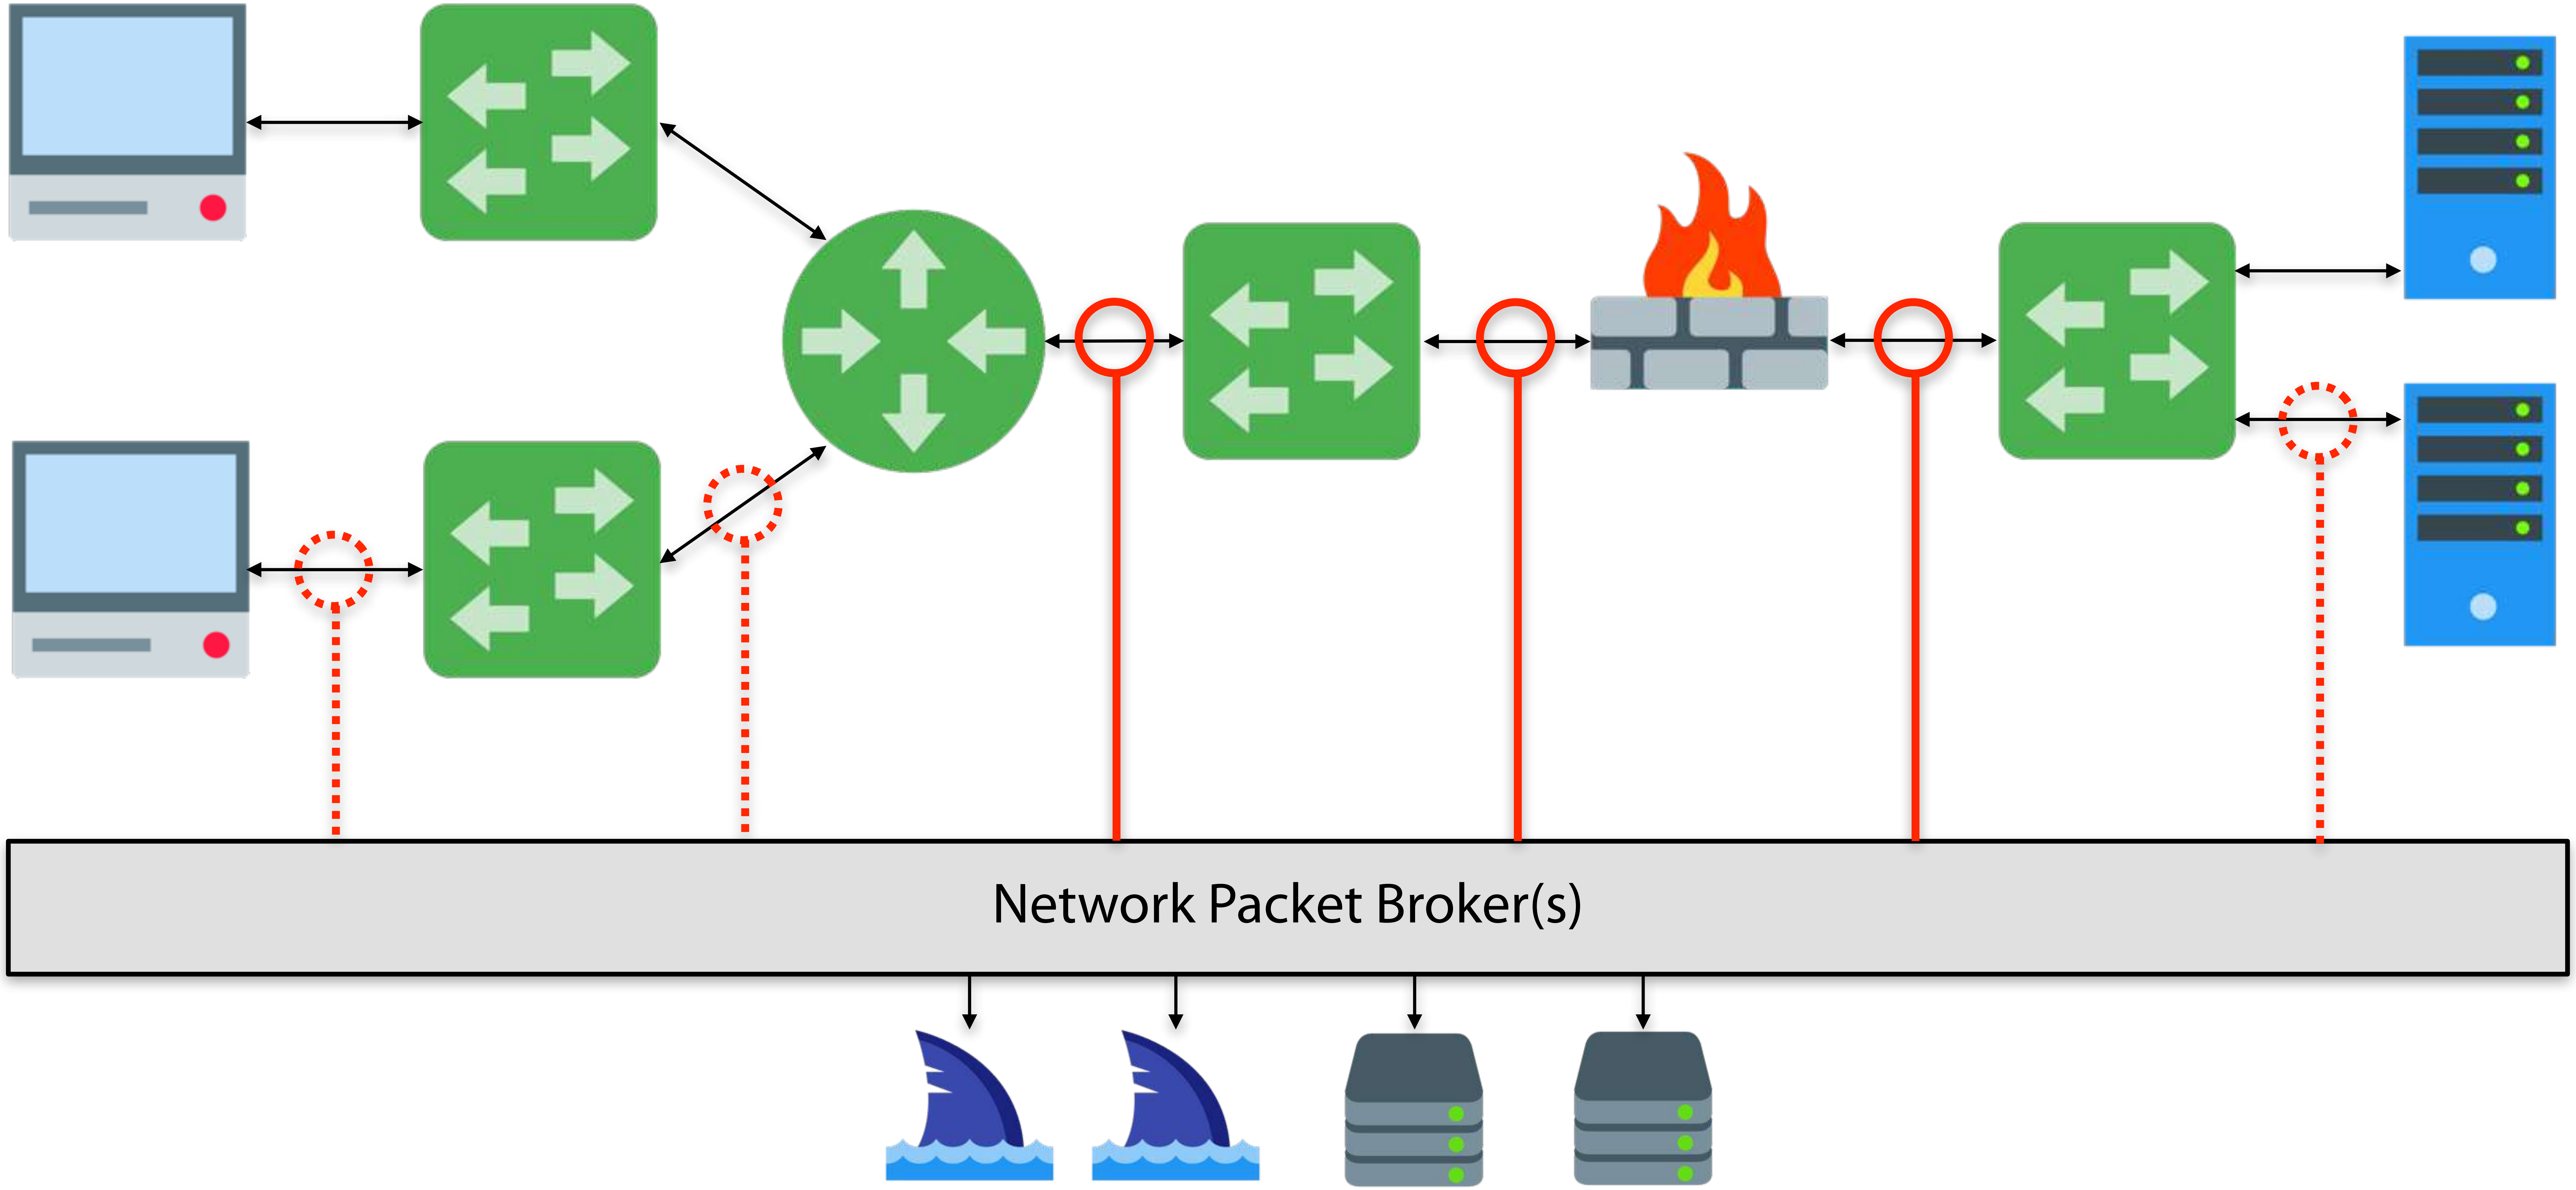

Adding a packet capture infrastructure

- Pros
  - Clear demarcation points between devices
    - so no doubt about effects of internal processing
  - No influence on existing devices
  - Every tool can get to every packet
    - With filtering, deduplication, (optional) hardware timestamping

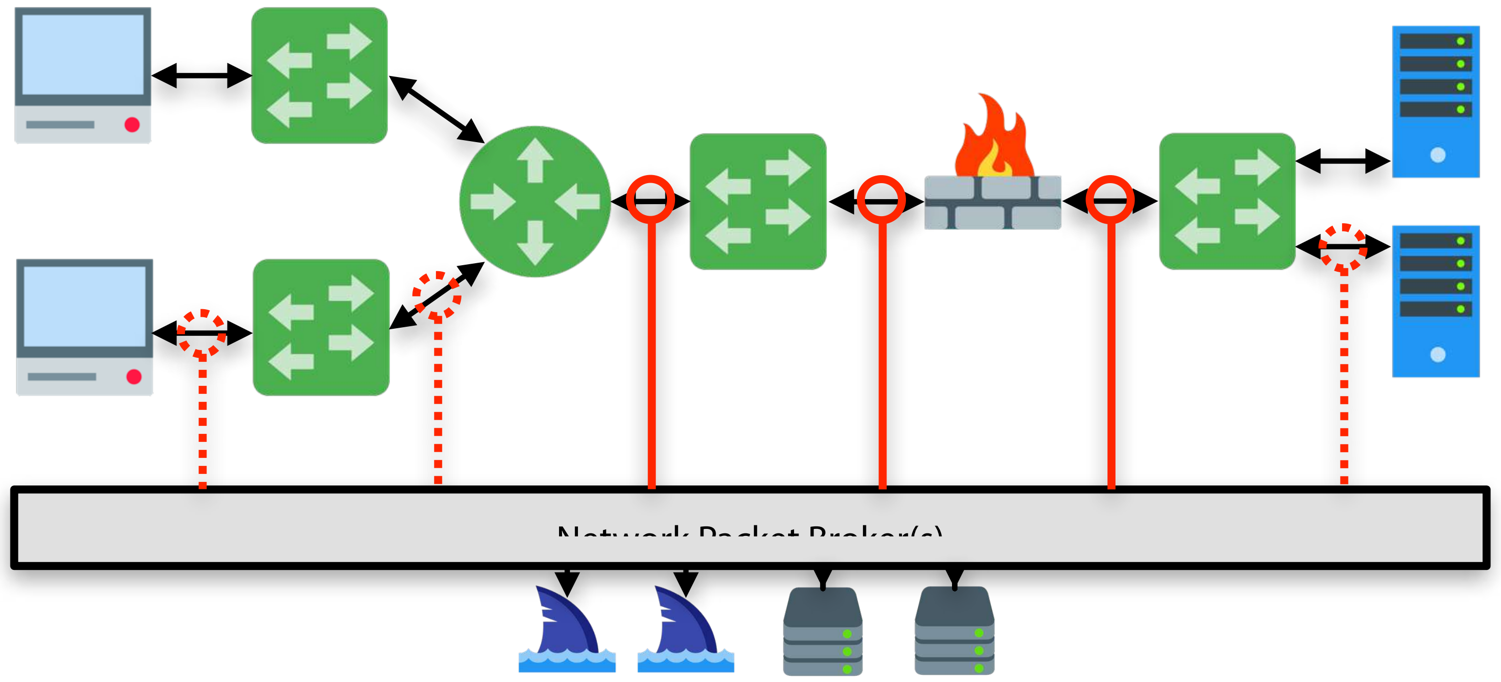

- Cons
  - High cost
    - Need Test Access Points (TAPs) at multiple locations
    - Needs Network Packet Brokers (NPBs)

Or just use one (or a few) points

- Pros
  - Clear demarcation points between devices
    - so no doubt about effects of internal processing
  - No influence on existing devices

- Cons
  - Need to interrupt the connection(s)
  - Cost

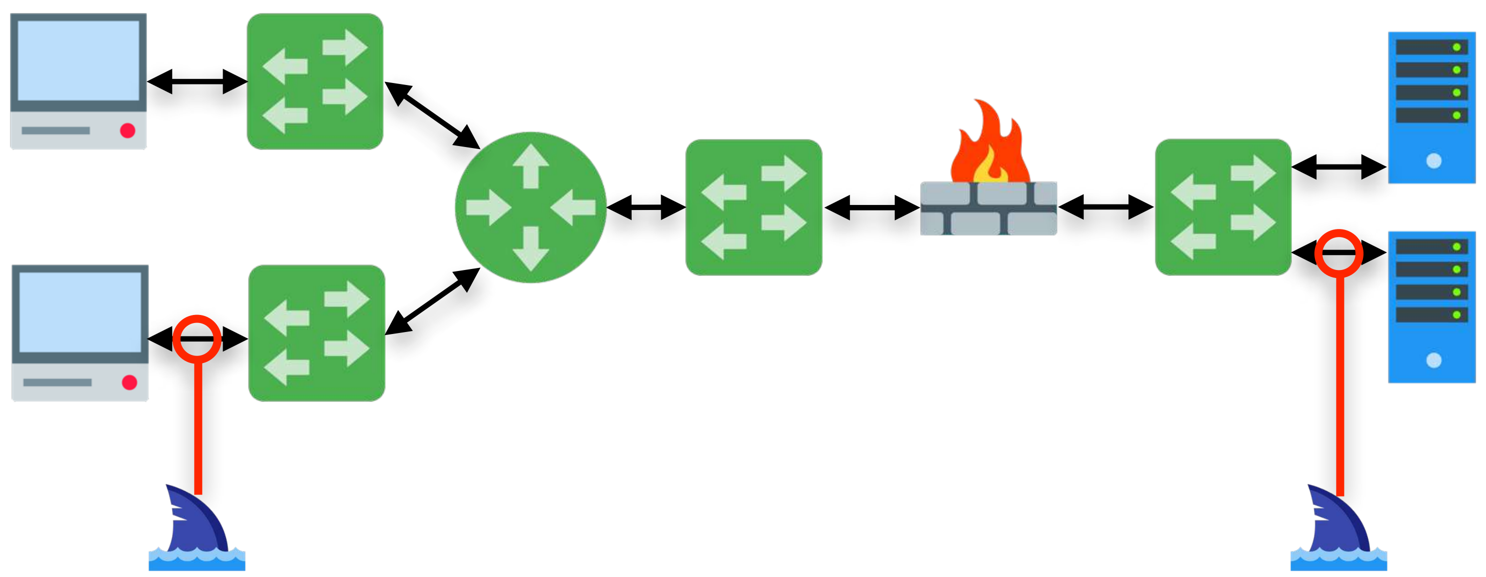

Insert a bridging host

- Pros
  - Quick, Easy and cheap

- Cons
  - All cons of capturing on a host
  - Bridging is not fully transparant
    - Some packets are not bridged by default
  - Inserting/Removing causes interruptions

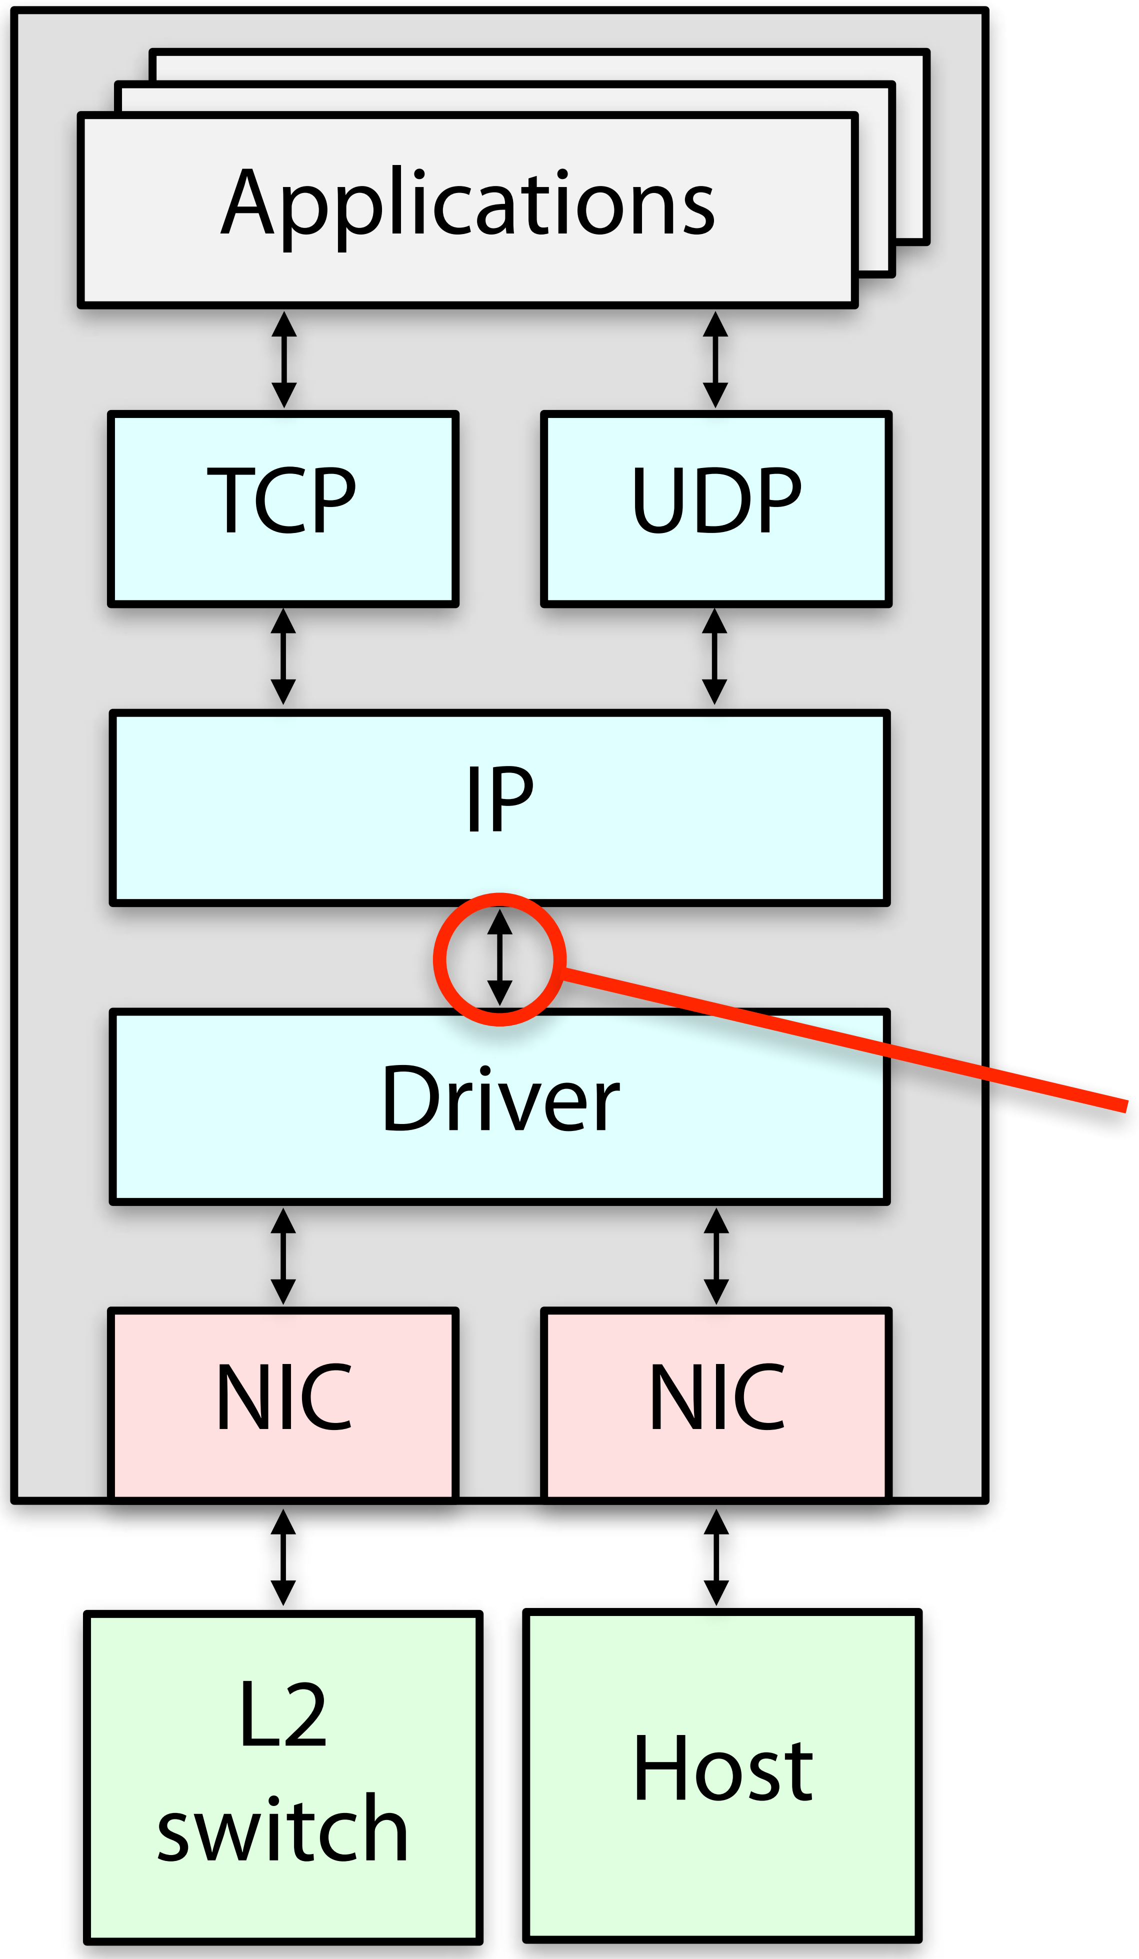

Insert a (real) HUB

- Pros
  - Quick, Easy and cheap
- Cons
  - Changes port speeds to 100 (or 10) Mbps
  - Connection becomes half-duplex
    - so risk of collissions
  - Inserting/Removing causes interruptions

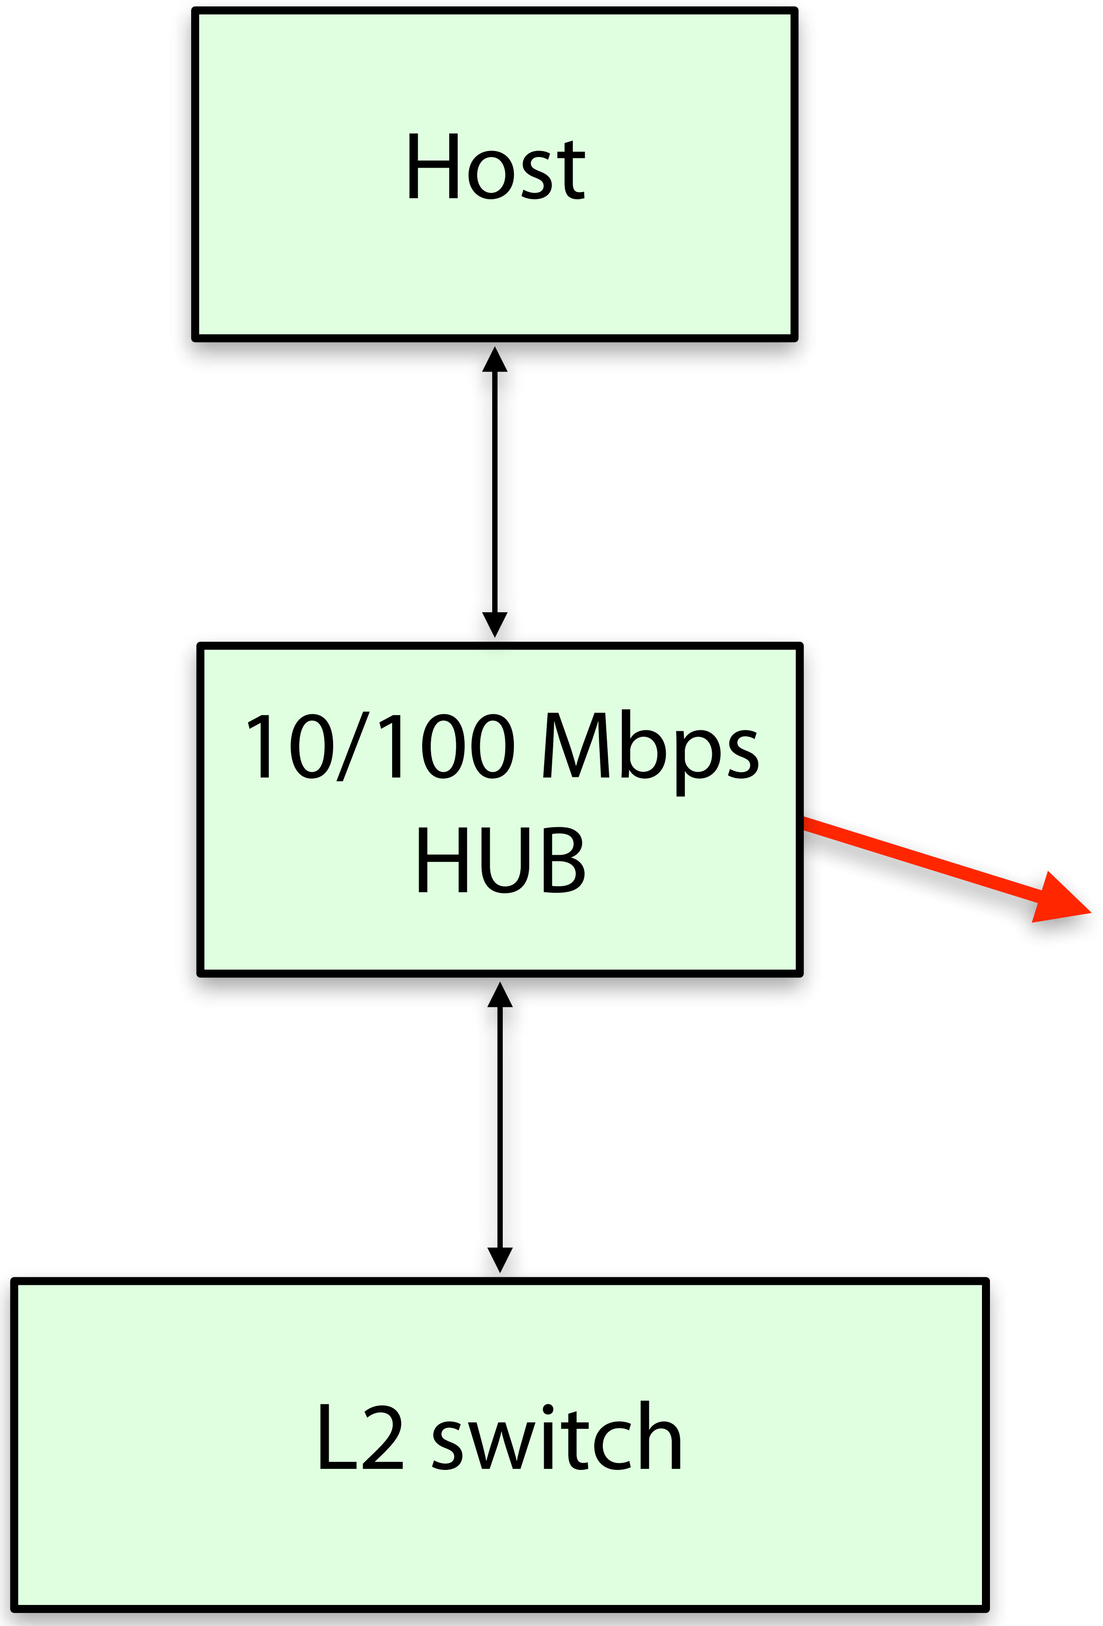

Insert a switch with SPAN

- Pros
  - Quick, Easy

- Cons
  - VLAN configuration must match that of link
  - Switch becomes part of the infrastructure
  - Switches are not fully transparant
    - some (bridge) protocols are not forwarded
  - Inserting/Removing causes interruptions

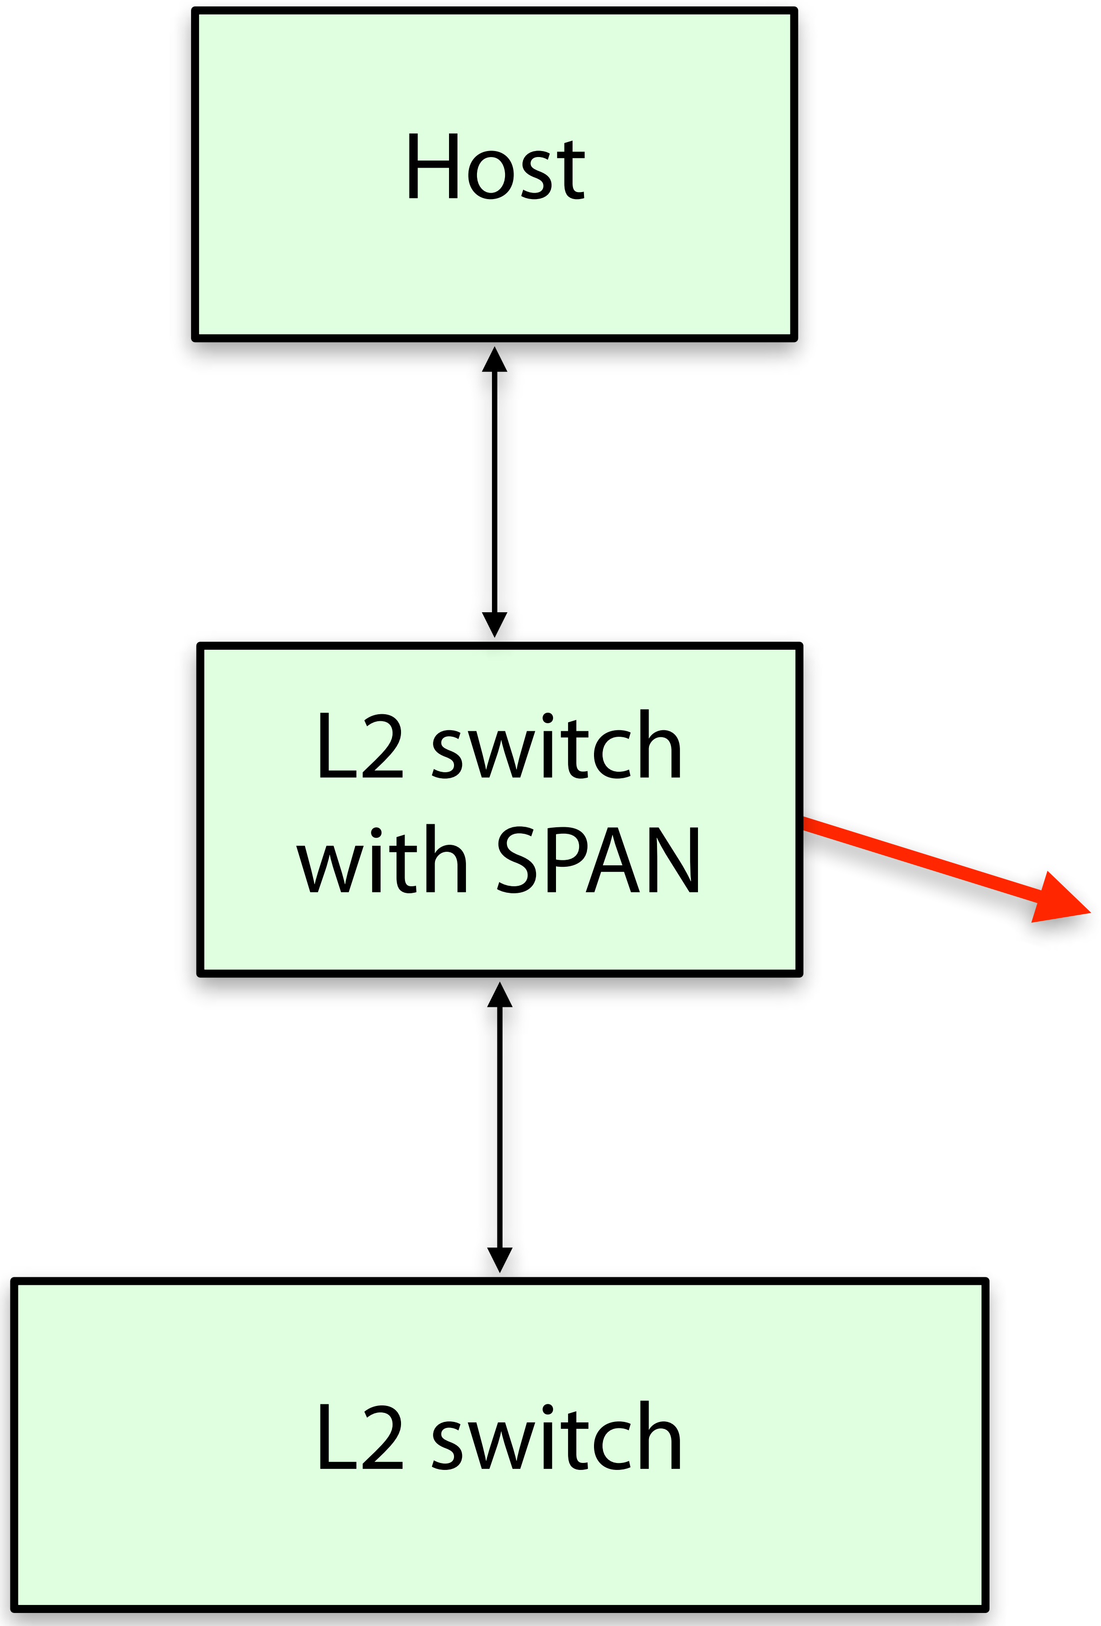

Insert passive TAP

- Pros
  - Quick, Easy and cheap

- Cons
  - Changes port speeds to 100 (or 10) Mbps
  - Need two NICs to capture both directions
  - Inserting/Removing causes interruptions

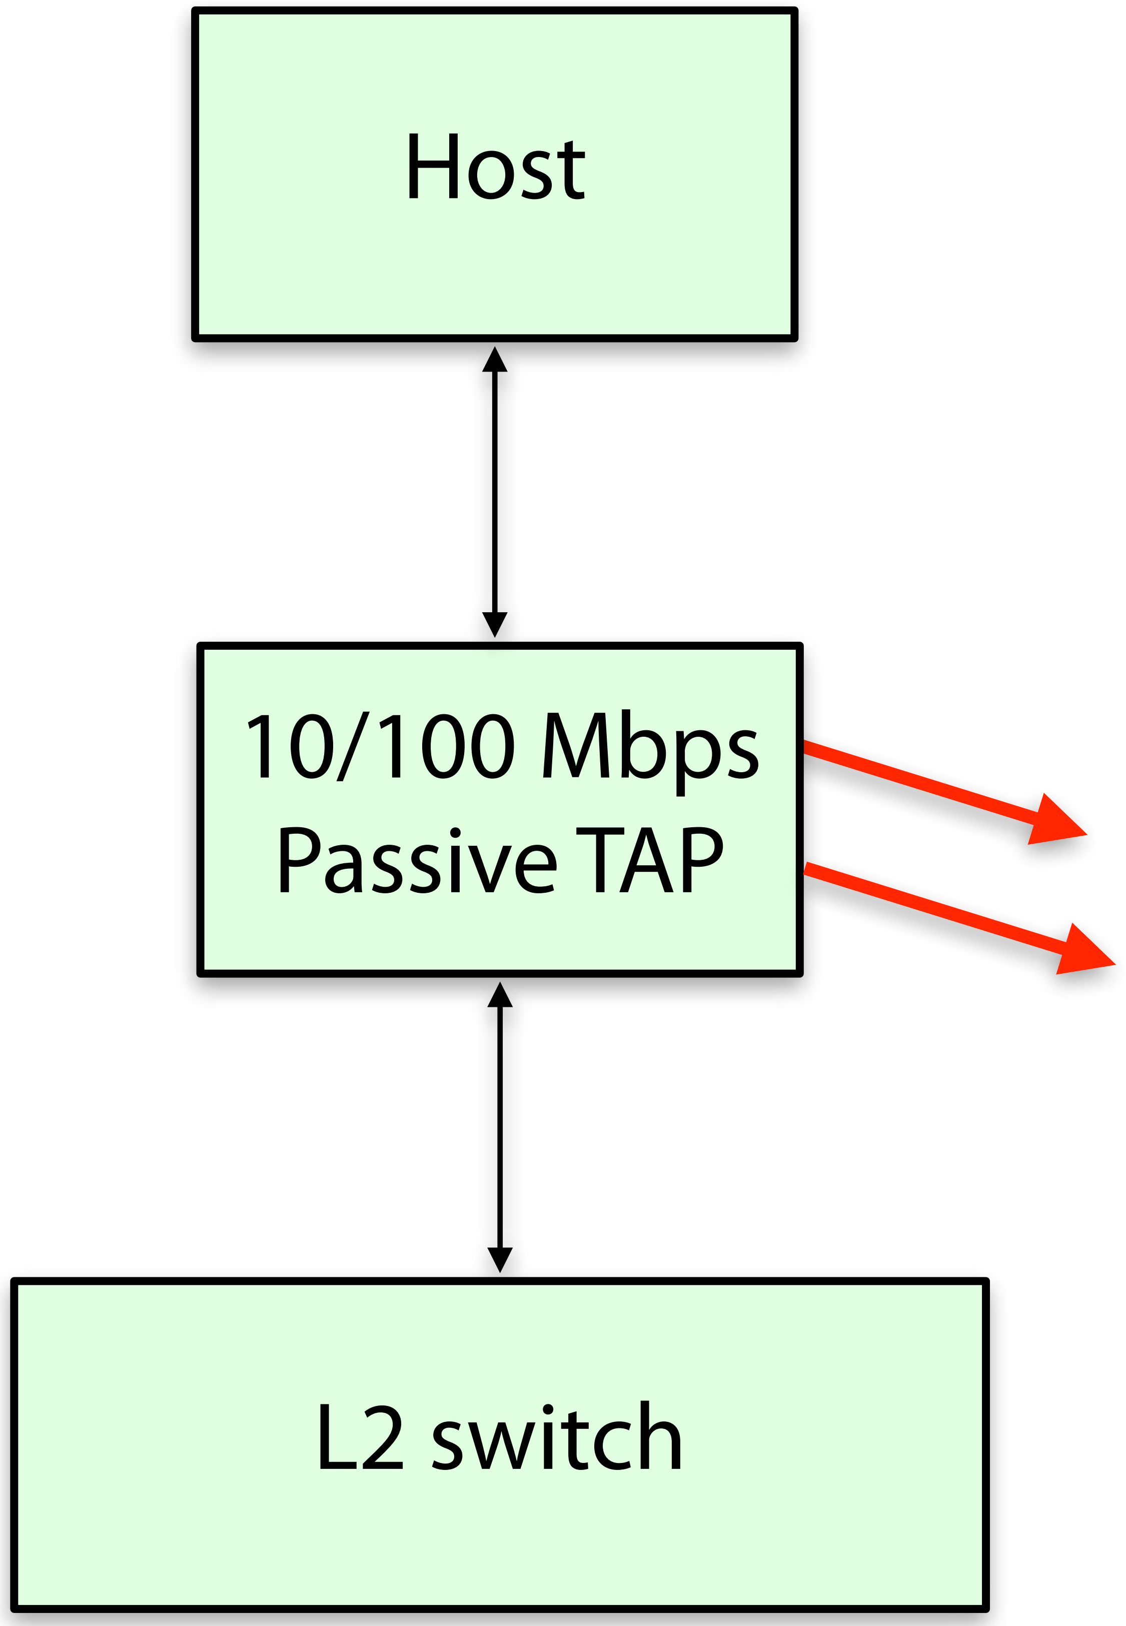

(homemade) passive TAPs

Startech usb32000spt

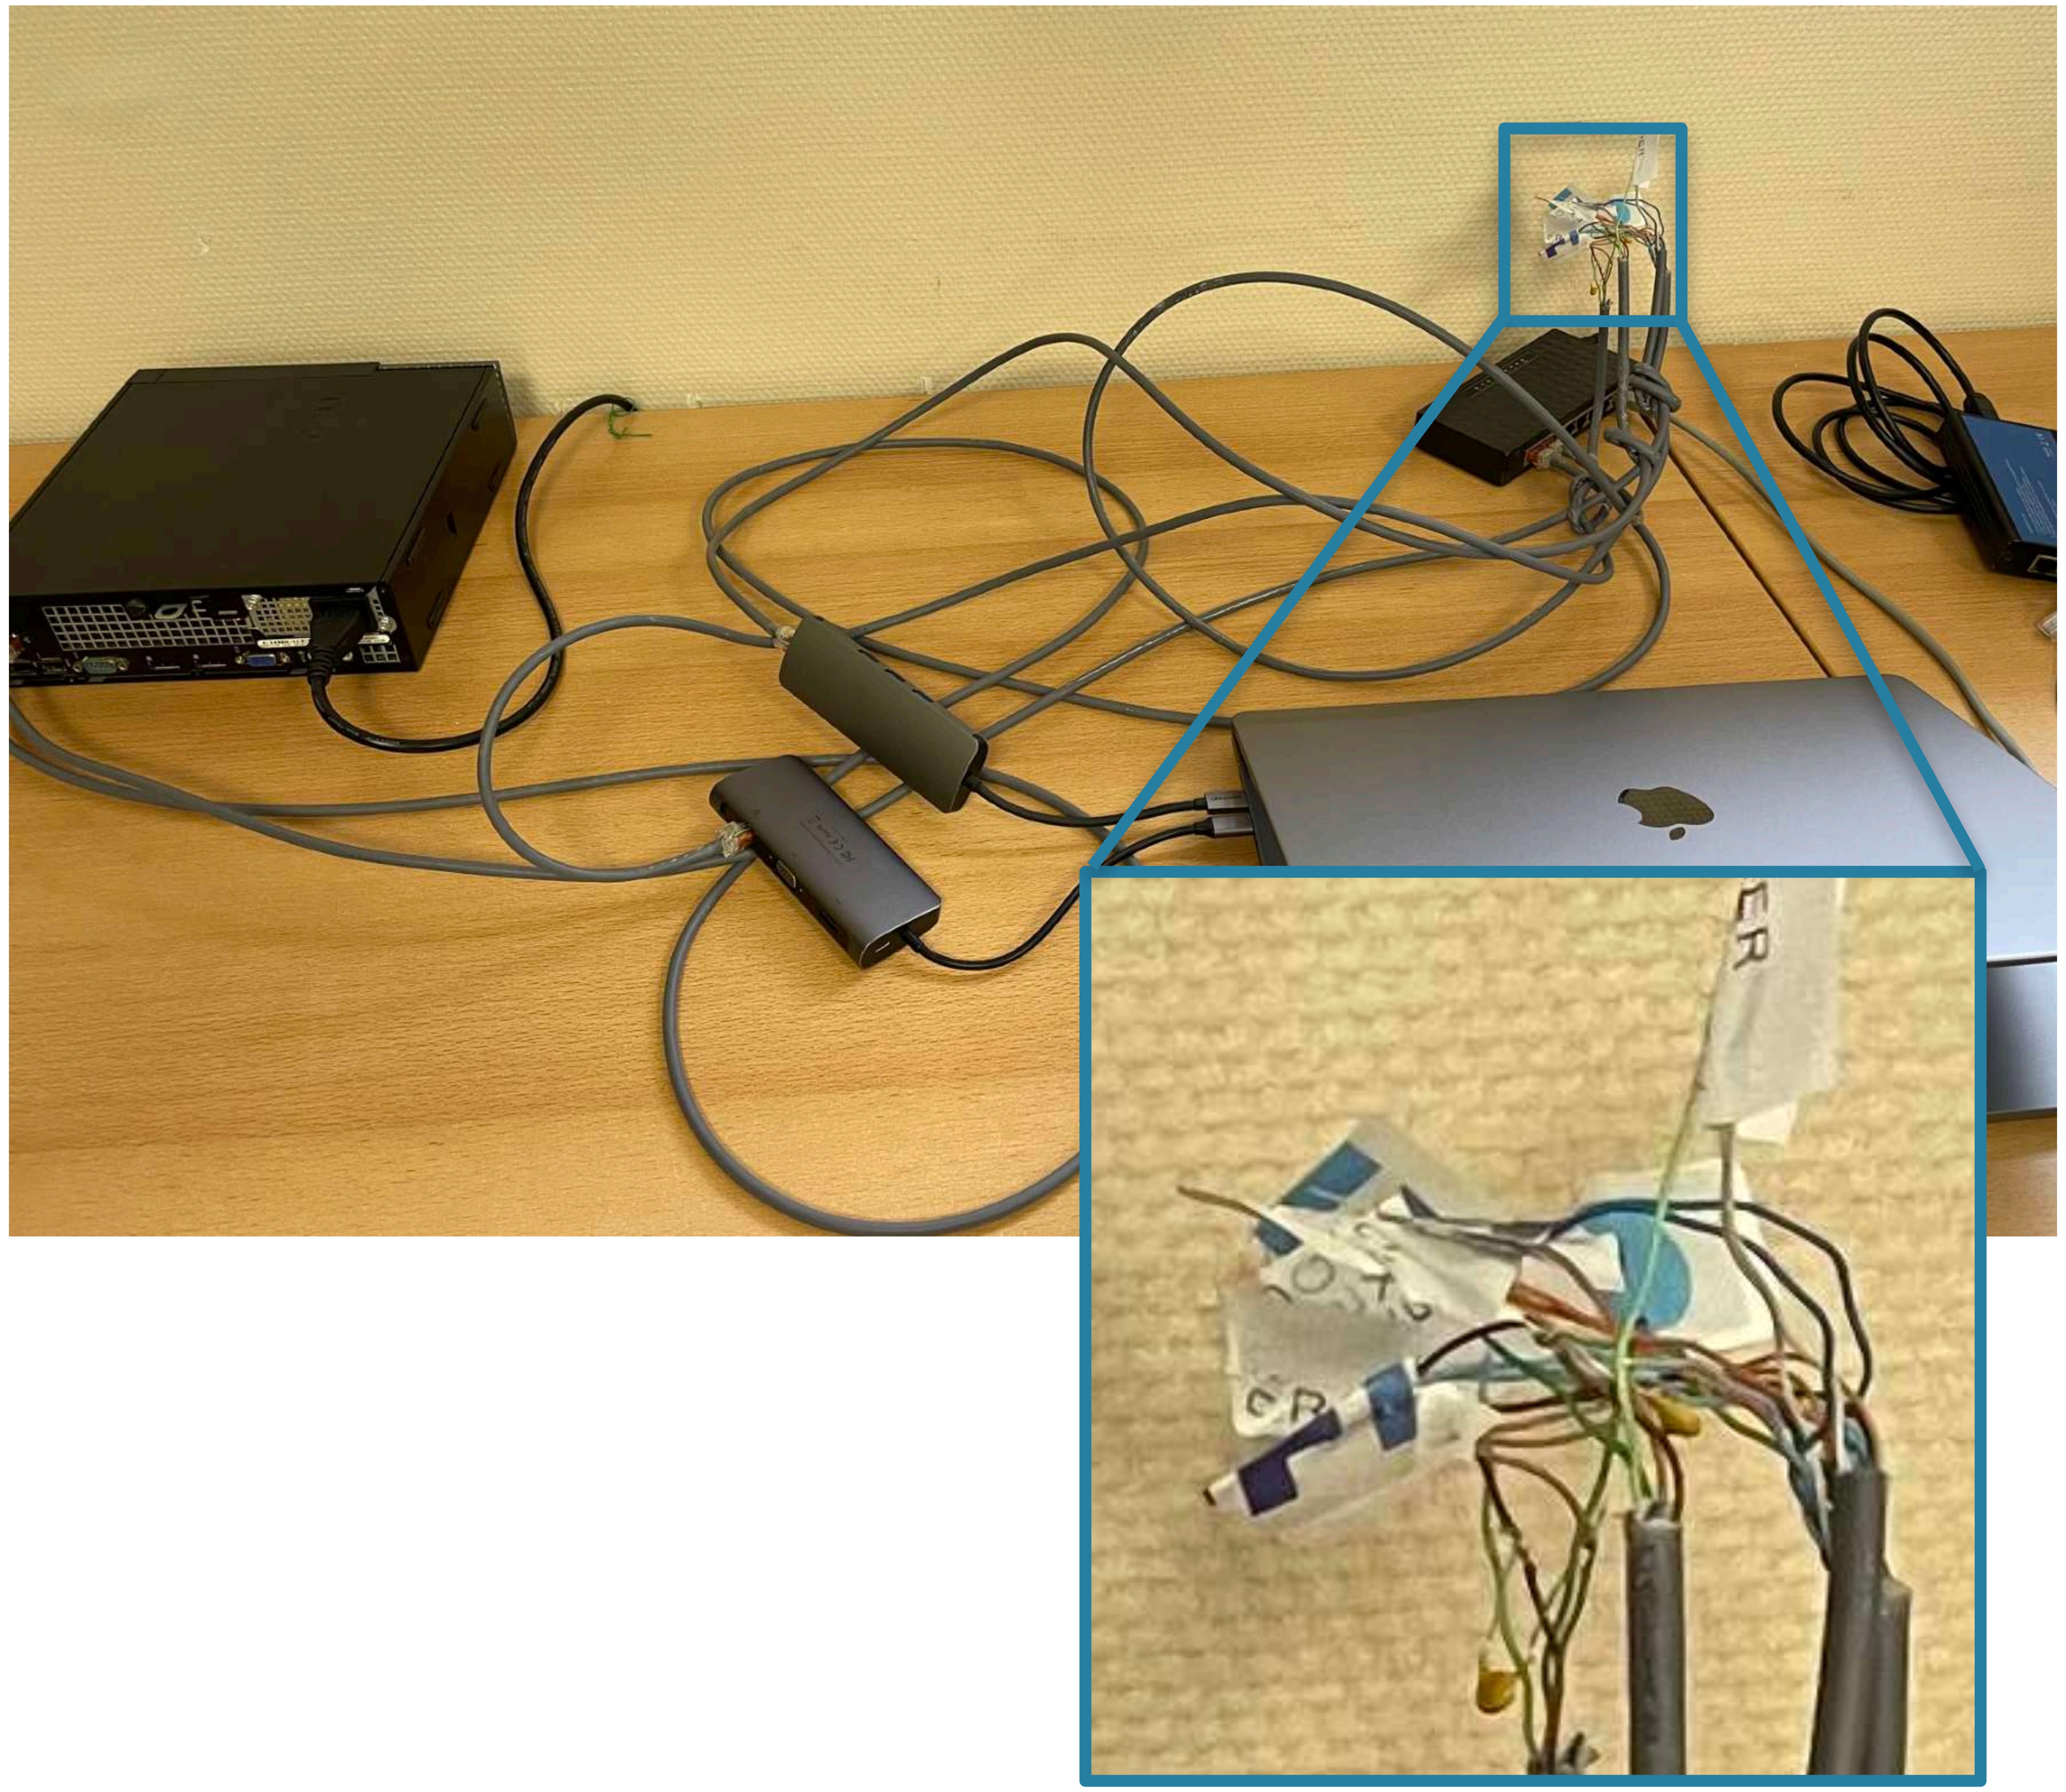

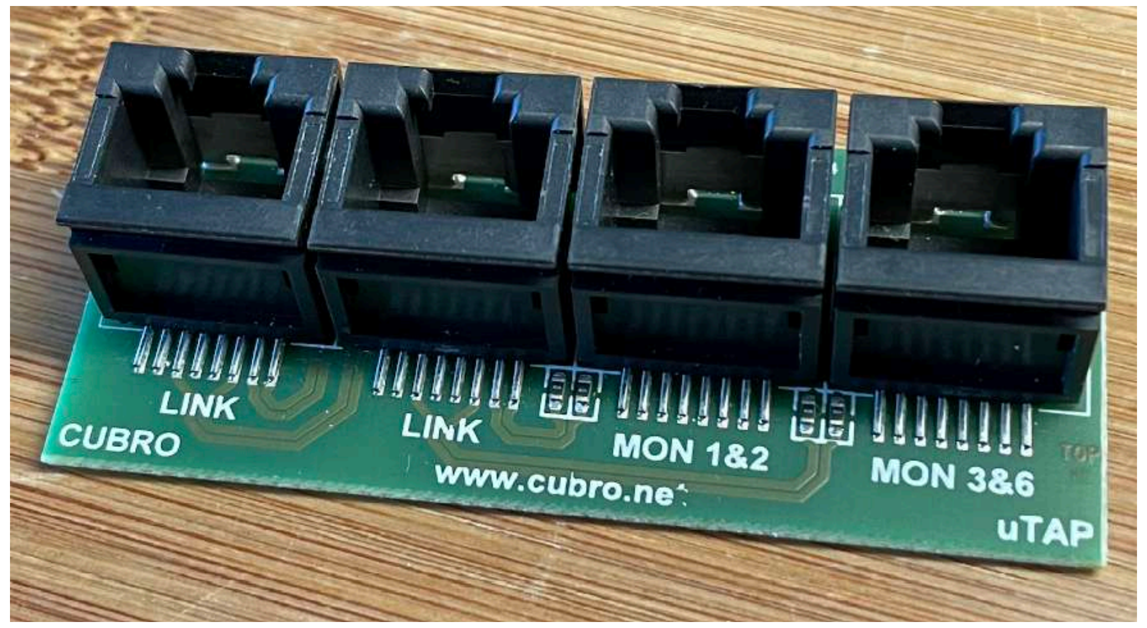

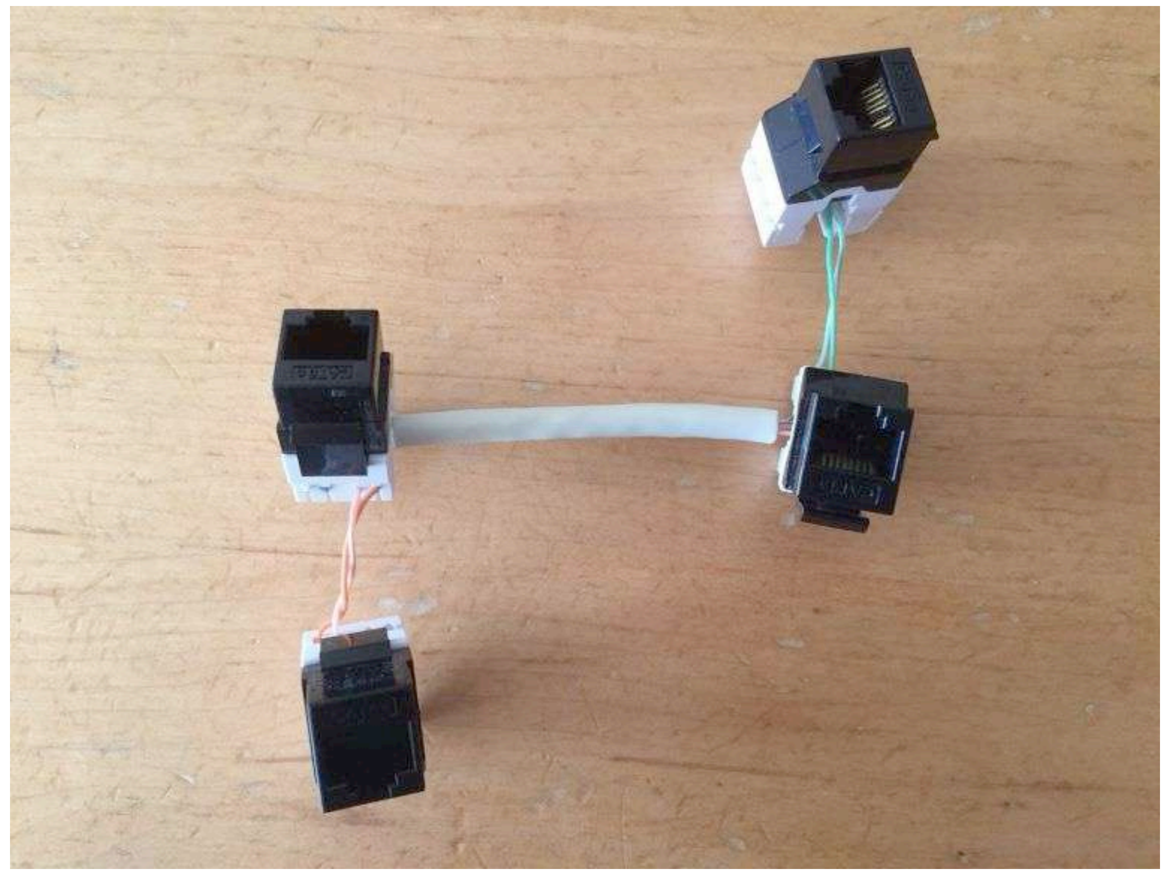

Timestamp frenzy!

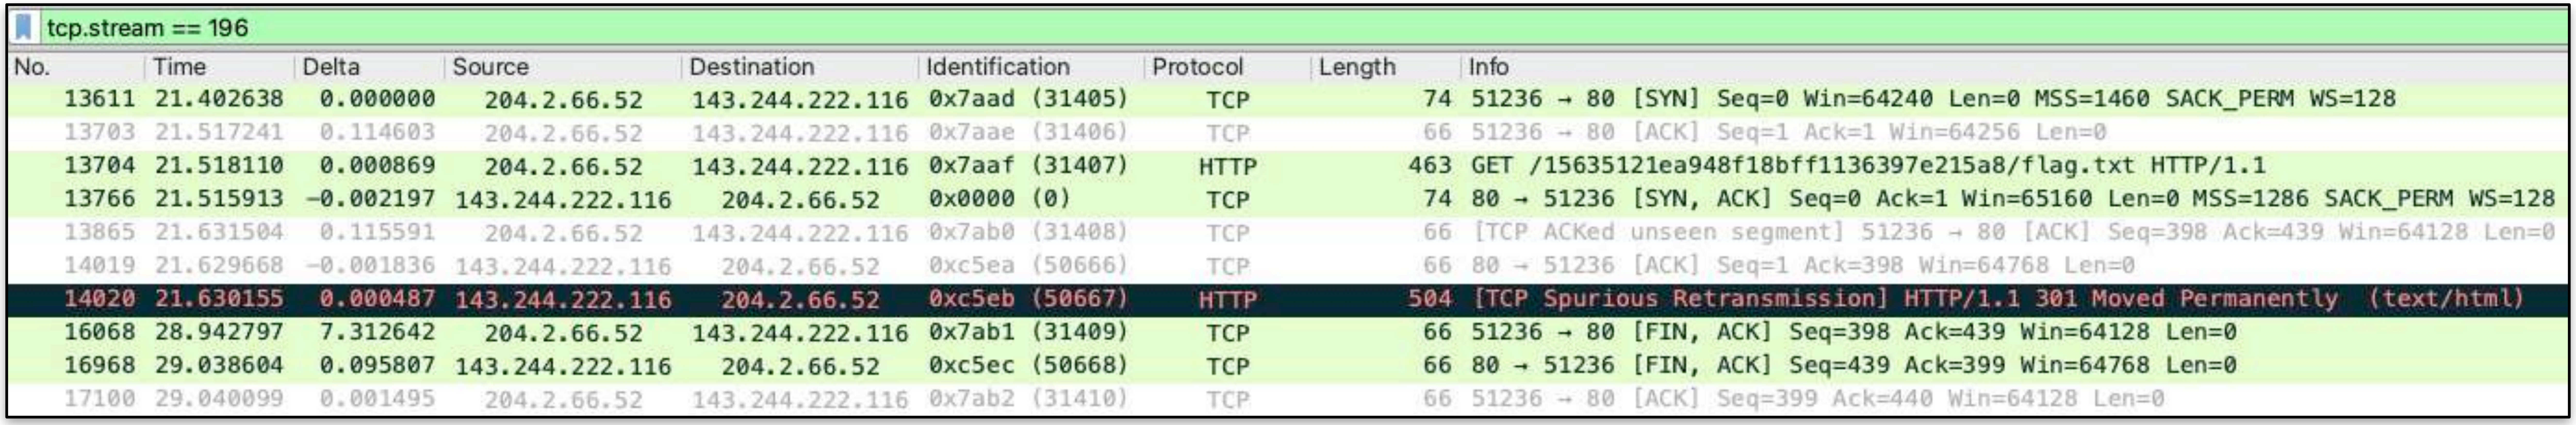

tcp.stream == 196

No.	Time	Delta	Source	Destination	Identification	Protocol	Length	Info
13611	21.402638	0.000000	204.2.66.52	143.244.222.116	0x7aad (31405)	TCP	74	51236 → 80 [SYN] Seq=0 Win=64240 Len=0 MSS=1460 SACK_PERM WS=128
13703	21.517241	0.114603	204.2.66.52	143.244.222.116	0x7aae (31406)	TCP	66	51236 → 80 [ACK] Seq=1 Ack=1 Win=64256 Len=0
13704	21.518110	0.000869	204.2.66.52	143.244.222.116	0x7aaf (31407)	HTTP	463	GET /15635121ea948f18bff1136397e215a8/flag.txt HTTP/1.1
13766	21.515913	-0.002197	143.244.222.116	204.2.66.52	0x0000 (0)	TCP	74	80 → 51236 [SYN, ACK] Seq=0 Ack=1 Win=65160 Len=0 MSS=1286 SACK_PERM WS=128
13865	21.631504	0.115591	204.2.66.52	143.244.222.116	0x7ab0 (31408)	TCP	66	[TCP ACKed unseen segment] 51236 → 80 [ACK] Seq=398 Ack=439 Win=64128 Len=0
14019	21.629668	-0.001836	143.244.222.116	204.2.66.52	0xc5ea (50666)	TCP	66	80 → 51236 [ACK] Seq=1 Ack=398 Win=64768 Len=0
14020	21.630155	0.000487	143.244.222.116	204.2.66.52	0xc5eb (50667)	HTTP	504	[TCP Spurious Retransmission] HTTP/1.1 301 Moved Permanently (text/html)
16068	28.942797	7.312642	204.2.66.52	143.244.222.116	0x7ab1 (31409)	TCP	66	51236 → 80 [FIN, ACK] Seq=398 Ack=439 Win=64128 Len=0
16968	29.038604	0.095807	143.244.222.116	204.2.66.52	0xc5ec (50668)	TCP	66	80 → 51236 [FIN, ACK] Seq=439 Ack=399 Win=64768 Len=0
17100	29.040099	0.001495	204.2.66.52	143.244.222.116	0x7ab2 (31410)	TCP	66	51236 → 80 [ACK] Seq=399 Ack=440 Win=64128 Len=0

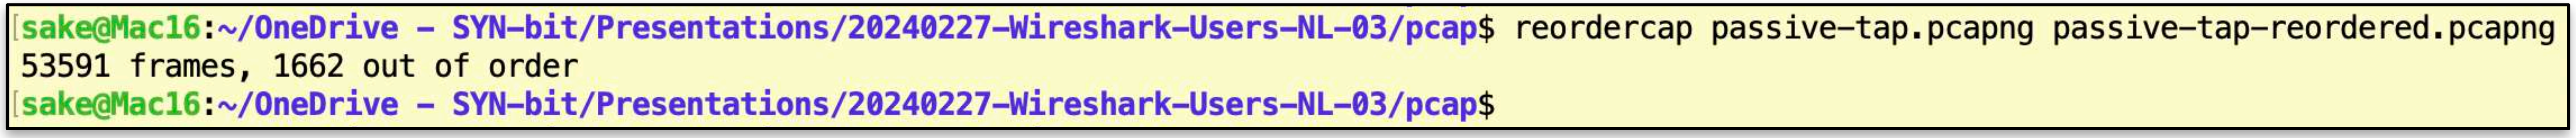

```
sake@Mac16:~/OneDrive - SYN-bit/Presentations/20240227-Wireshark-Users-NL-03/pcap$ reordercap passive-tap.pcapng passive-tap-reordered.pcapng
53591 frames, 1662 out of order
sake@Mac16:~/OneDrive - SYN-bit/Presentations/20240227-Wireshark-Users-NL-03/pcap$
```

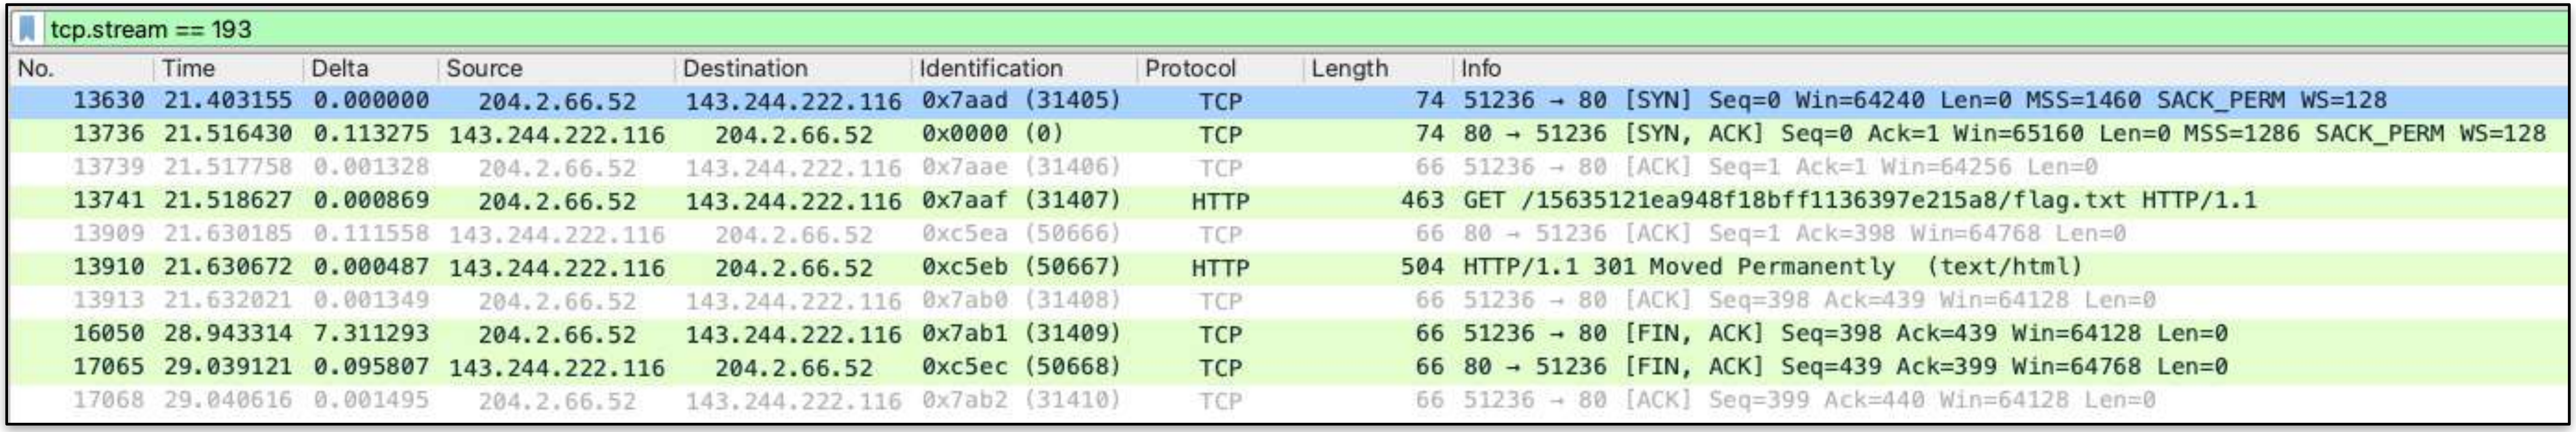

tcp.stream == 193

No.	Time	Delta	Source	Destination	Identification	Protocol	Length	Info
13630	21.403155	0.000000	204.2.66.52	143.244.222.116	0x7aad (31405)	TCP	74	51236 → 80 [SYN] Seq=0 Win=64240 Len=0 MSS=1460 SACK_PERM WS=128
13736	21.516430	0.113275	143.244.222.116	204.2.66.52	0x0000 (0)	TCP	74	80 → 51236 [SYN, ACK] Seq=0 Ack=1 Win=65160 Len=0 MSS=1286 SACK_PERM WS=128
13739	21.517758	0.001328	204.2.66.52	143.244.222.116	0x7aae (31406)	TCP	66	51236 → 80 [ACK] Seq=1 Ack=1 Win=64256 Len=0
13741	21.518627	0.000869	204.2.66.52	143.244.222.116	0x7aaf (31407)	HTTP	463	GET /15635121ea948f18bff1136397e215a8/flag.txt HTTP/1.1
13909	21.630185	0.111558	143.244.222.116	204.2.66.52	0xc5ea (50666)	TCP	66	80 → 51236 [ACK] Seq=1 Ack=398 Win=64768 Len=0
13910	21.630672	0.000487	143.244.222.116	204.2.66.52	0xc5eb (50667)	HTTP	504	HTTP/1.1 301 Moved Permanently (text/html)
13913	21.632021	0.001349	204.2.66.52	143.244.222.116	0x7ab0 (31408)	TCP	66	51236 → 80 [ACK] Seq=398 Ack=439 Win=64128 Len=0
16050	28.943314	7.311293	204.2.66.52	143.244.222.116	0x7ab1 (31409)	TCP	66	51236 → 80 [FIN, ACK] Seq=398 Ack=439 Win=64128 Len=0
17065	29.039121	0.095807	143.244.222.116	204.2.66.52	0xc5ec (50668)	TCP	66	80 → 51236 [FIN, ACK] Seq=439 Ack=399 Win=64768 Len=0
17068	29.040616	0.001495	204.2.66.52	143.244.222.116	0x7ab2 (31410)	TCP	66	51236 → 80 [ACK] Seq=399 Ack=440 Win=64128 Len=0

Insert a "real" network TAP

- Pros
  - Transparant
  - Can be left inline (especially with optical TAPs)
  - (sometimes) Specialised hardware (FPGA)
    - Forward frames, regardless of size, error etc
    - accurate timestamping
    - port information
- Cons
  - Can be Expensive
  - Inserting/Removing causes interruptions

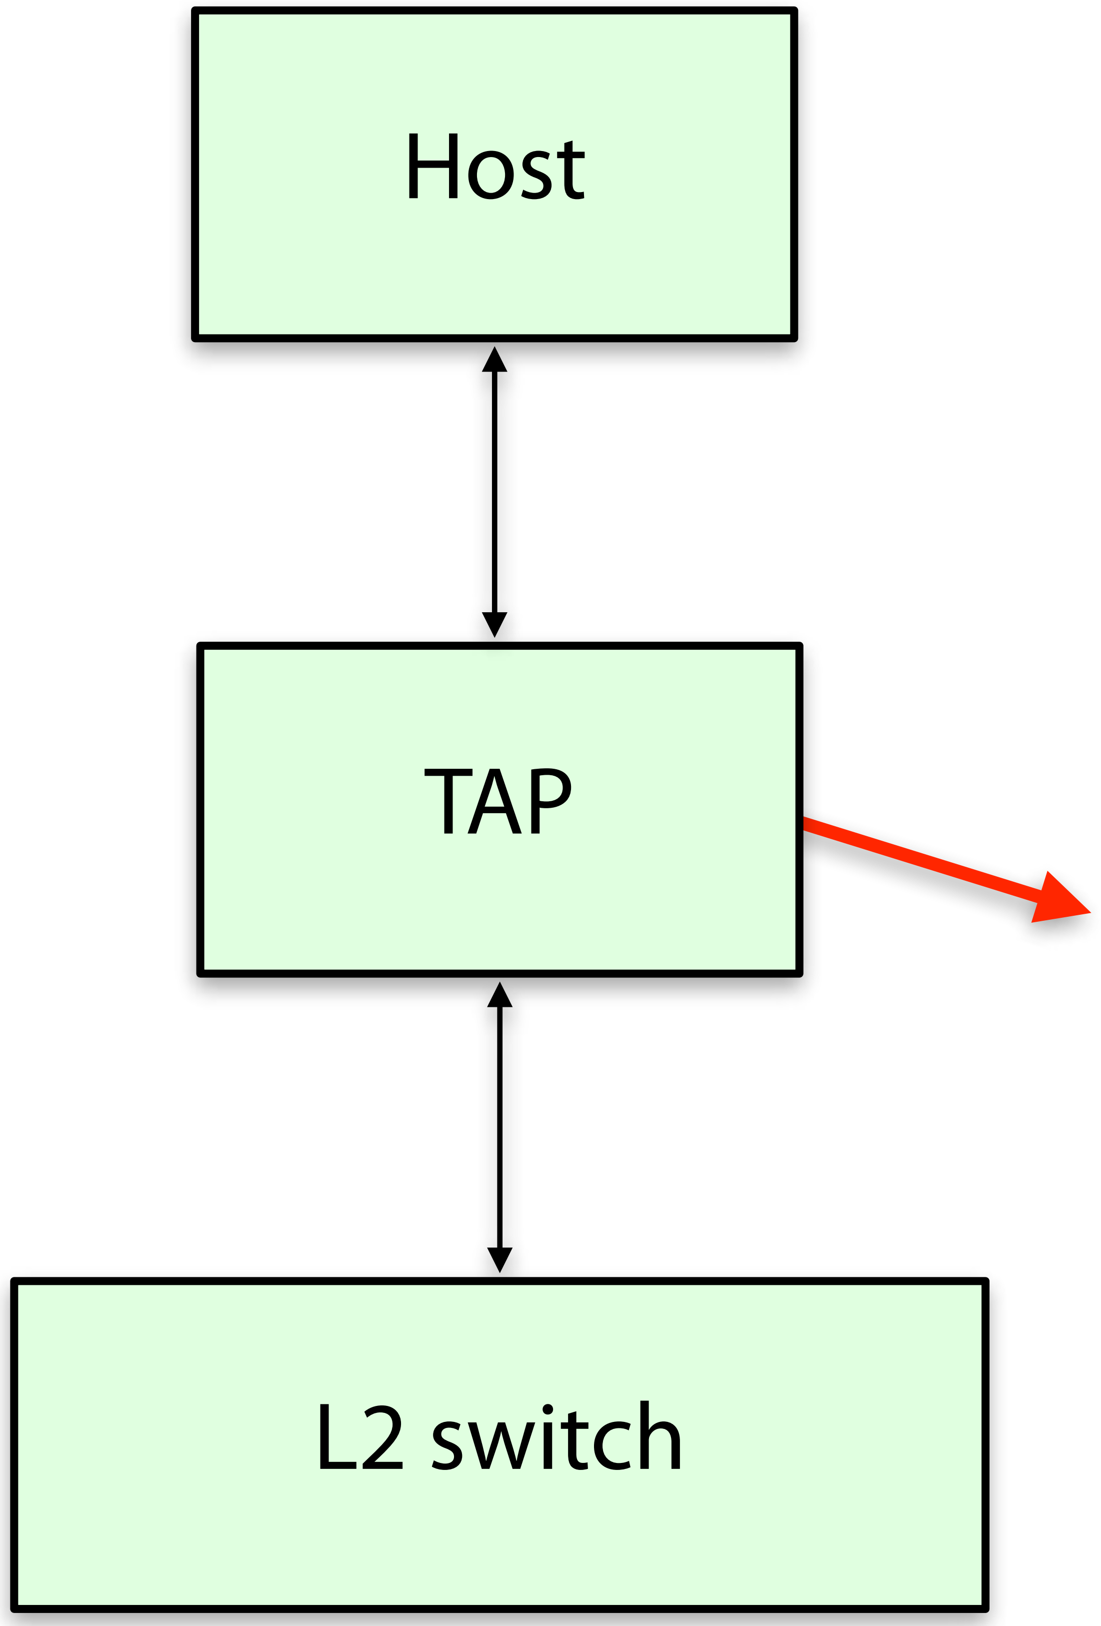

We want a *portable* troubleshooting TAP!

- Many TAP models, mostly rackmount
  - 1, 10, 40, 100 Gbps / fiber or copper / bypass / etc
  - One to many ports

- Requirements for a TAP in your laptop bag:
  - 10/100/1000 copper ethernet
  - POE forwarding
  - Aggregating and/or breakout
  - Forwarding small/large/bad frames
  - Windows / MacOS / Linux compatible
  - Preferably USB powered (so no heavy power adapter)

Models found

- ETAP-2003 (Dualcomm)
- PacketRaven PRP-SCC-1GA (NEOX networks)
- *P1GCCAS (Garland Technology)*
- LANProbe (Qlinx / RTNsystems)
- SharkTap - multiple versions (midBitTech)
- *USR4524-MINI (US Robotics)*
- ProfiShark 1G / 1G+ / 10G / 10G+ (Profitap)

ETAP-2003 (Dualcomm)

- Network ports: **2x 1000baseT**
- Monitor port: **1x 1000baseT**
- TAP mode: **aggregating**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **~ €200**

- Tested: **YES**

PacketRaven PRP-SCC-1GA (NEOX networks)

- Network ports: **2x 1000baseT**
- Monitor ports: **2x 1000baseT**
- TAP modes: **aggregating/breakout/regenerating**
- Powered by: **(dual) power adapter**
- POE forwarding: **yes**

- Listprice: **on request**

- Tested: **YES**

P1GCCAS (Garland Technology)

- Network ports: **2x 1000baseT**
- Monitor ports: **2x 1000baseT**
- TAP modes: **aggregating/breakout/regenerating**
- Powered by: **(single) power adapter**
- POE forwarding: **yes**

- Listprice: **on request**

- Tested: **NO**

LANProbe (Q-linx / RTNSystems)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x 1000baseT + 1x USB 3.0 (1 Gbps)**
- TAP modes: **aggregating**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **$199**

- Tested: **YES**

SharkTap (midBitTech)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x 1000baseT**
- TAP modes: **aggregating**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **$219,95**

- Tested: **NO**

SharkTapUSB (midBitTech)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x 1000baseT + 1x USB 3.0 (1 Gbps)**
- TAP modes: **aggregating**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **$269,95**

- Tested: **YES**

SharkTapBYP (midBitTech)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x 1000baseT + 1x USB 3.0 (1 Gbps)**
- TAP modes: **aggregating**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **$329,95**

- Tested: **NO**

SharkTapHUB (midBitTech)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x 1000baseT**
- TAP modes: **"Full Duplex Gigabit HUB"**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **$229,95**

- Tested: **NO**

USR4524-MINI (US Robotics)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x USB 3.0 (2x 1 Gbps?)**
- TAP modes: **breakout? aggregating?**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **on request**

- Tested: **NO**

ProfiShark 1G (Profitap)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x USB 3.0 (2 Gbps)**
- TAP modes: **aggregating (inline or span mode)**
- Powered by: **USB**
- POE forwarding: **yes**
- Listprice: **on request**
- Tested: **YES**

ProfiShark 1G+ (Profitap)

- Network ports: **2x 1000baseT**
- Monitor ports: **1x USB 3.0 (2 Gbps)**
- TAP modes: **aggregating (inline or span mode)**
- Powered by: **USB**
- POE forwarding: **yes**

- Listprice: **on request**

- Tested: **NO**

All TAPs are created equal, but...

- From € to €€€
- Switch-ASIC, RGMII carbon copy, FPGA, etc
- Features for permanent deployment
- Accuracy of aggregation
- Timestamping in hardware
- Compliancy

- Let's test and see which features (that are important to us) are supported!

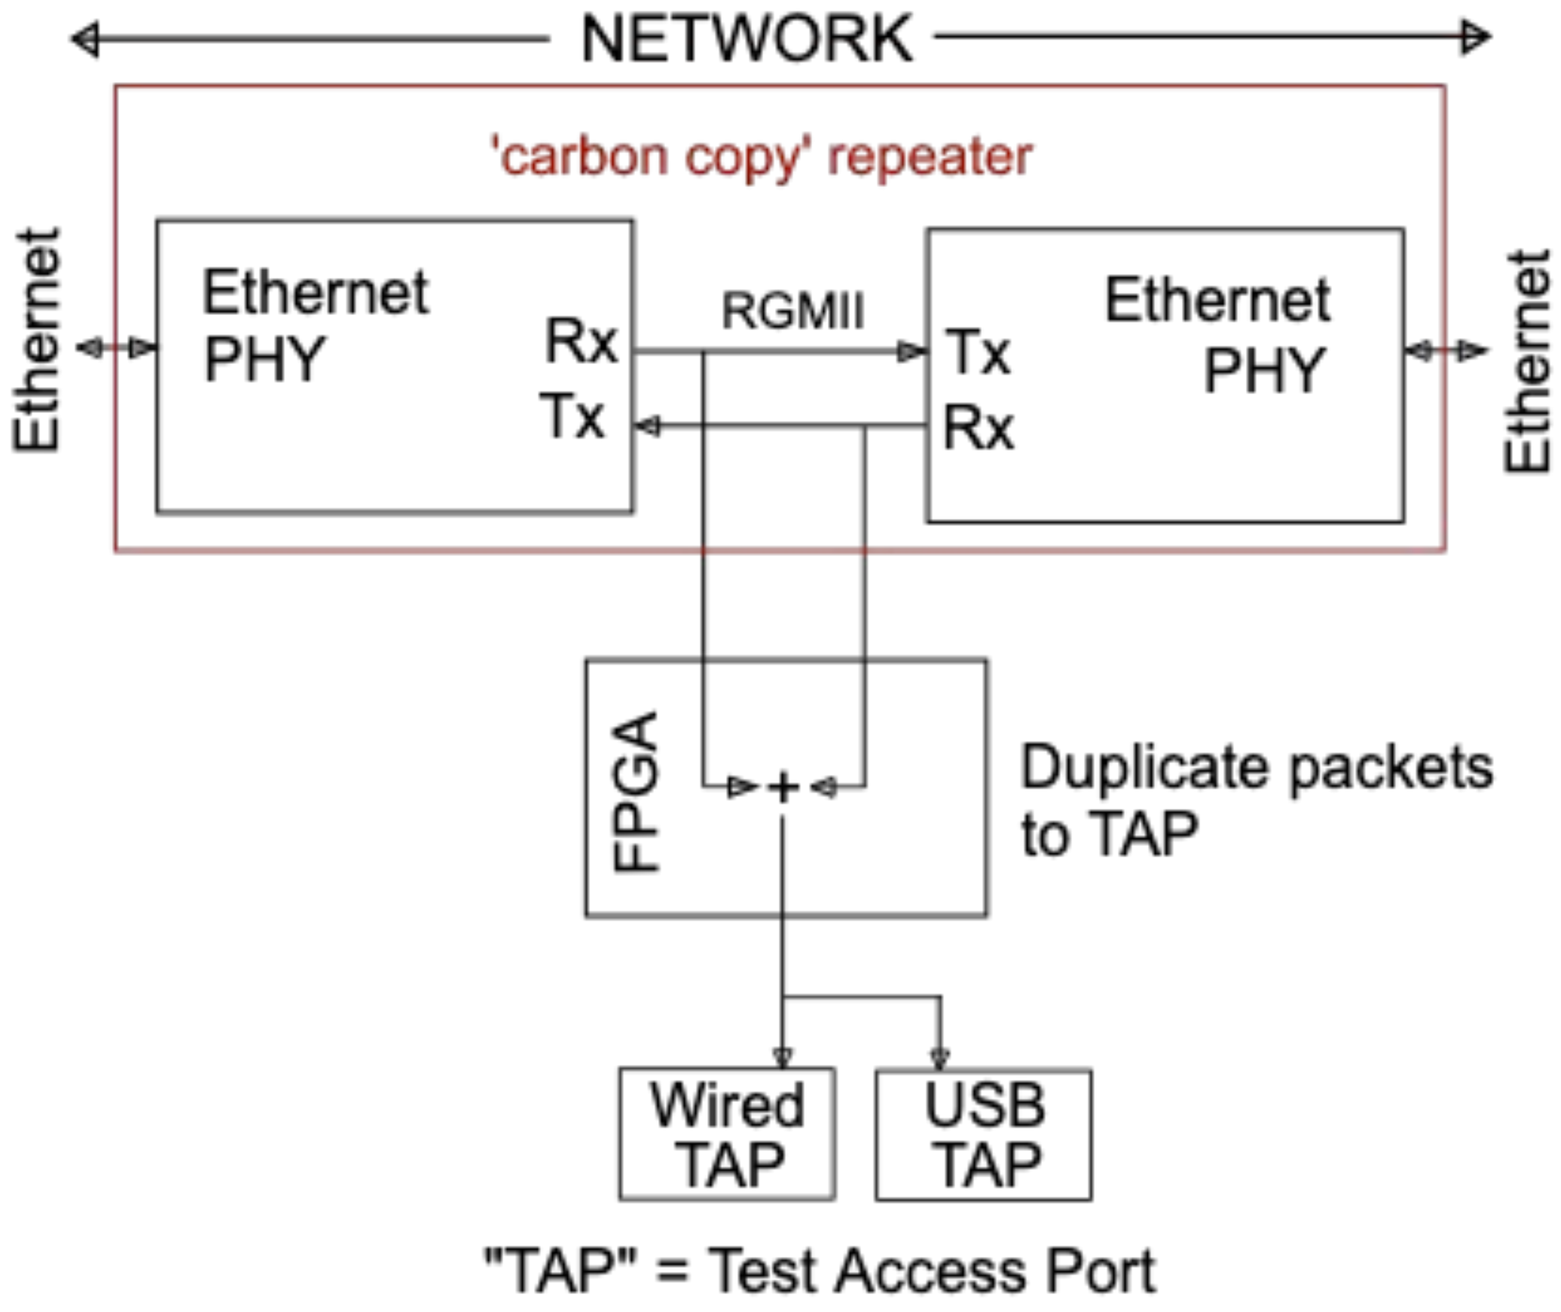

Test equipment

- 5 TAP "devices under test" (DUT)
- A Cisco switch with POE (WS-C2960C-12PC-L)
- A VoIP phone (Avaya 9650)
- FMADIO FMAD20p3 - 20 Gbps capture device
  - Uses FPGAs and can use 1Gbps and 10Gbps SFPs
  - Packet generating functionality
- Lenovo Thinkpad X270
  - Ubuntu 22.04.3 (6.2.0-39-generic)
  - On board Intel I219-V 10/100/1000 ethernet
- An untwisted CAT5e cable and a lighter!

Physical characteristics

TAP	Monitor Ports	USB Powered	Power (cable)	Size (LxBxH) Weight
DUALCOMM ETAP-2003	1x RJ45 (aggregate 1Gbps)	✅	USB-A power cable	94x70x26 mm 184 gram
NEOX PRP-SCC-1GA	2x RJ45 (breakout/aggregate/regenerate)	❌	Power Adapter or POE	170x106x35 mm 470 + 110 gram
LANProbe	1x RJ45 + 1x USB 3.0 eth (aggregate 1Gbps)	✅	USB 3.0 (USB-A to USB-B)	123x66x28 mm 173 gram
SharkTapUSB	1x RJ45 + 1x USB 3.0 eth (aggregate 1Gbps)	✅	USB 3.0 (USB-A to USB-B)	130x70x28 mm 140 gram
ProfiShark 1G	1x USB 3.0 (aggregate 2Gbps)	✅	USB 3.0 (USB-A to USB-B)	124x69x24 mm 174 gram

Testing Network Port Features

- POE forwarding
  - Will the VoIP phone power on over the TAP?
- Link Negotiation Forwarding
  - Connect and change settings on one side, will the other side follow? If not, speed mismatches can occur
- Link Failure Propagation
  - Disconnect the phone, does the switch port go down?
- Bypass on power failure
  - Will the TAP forward packets when not on power

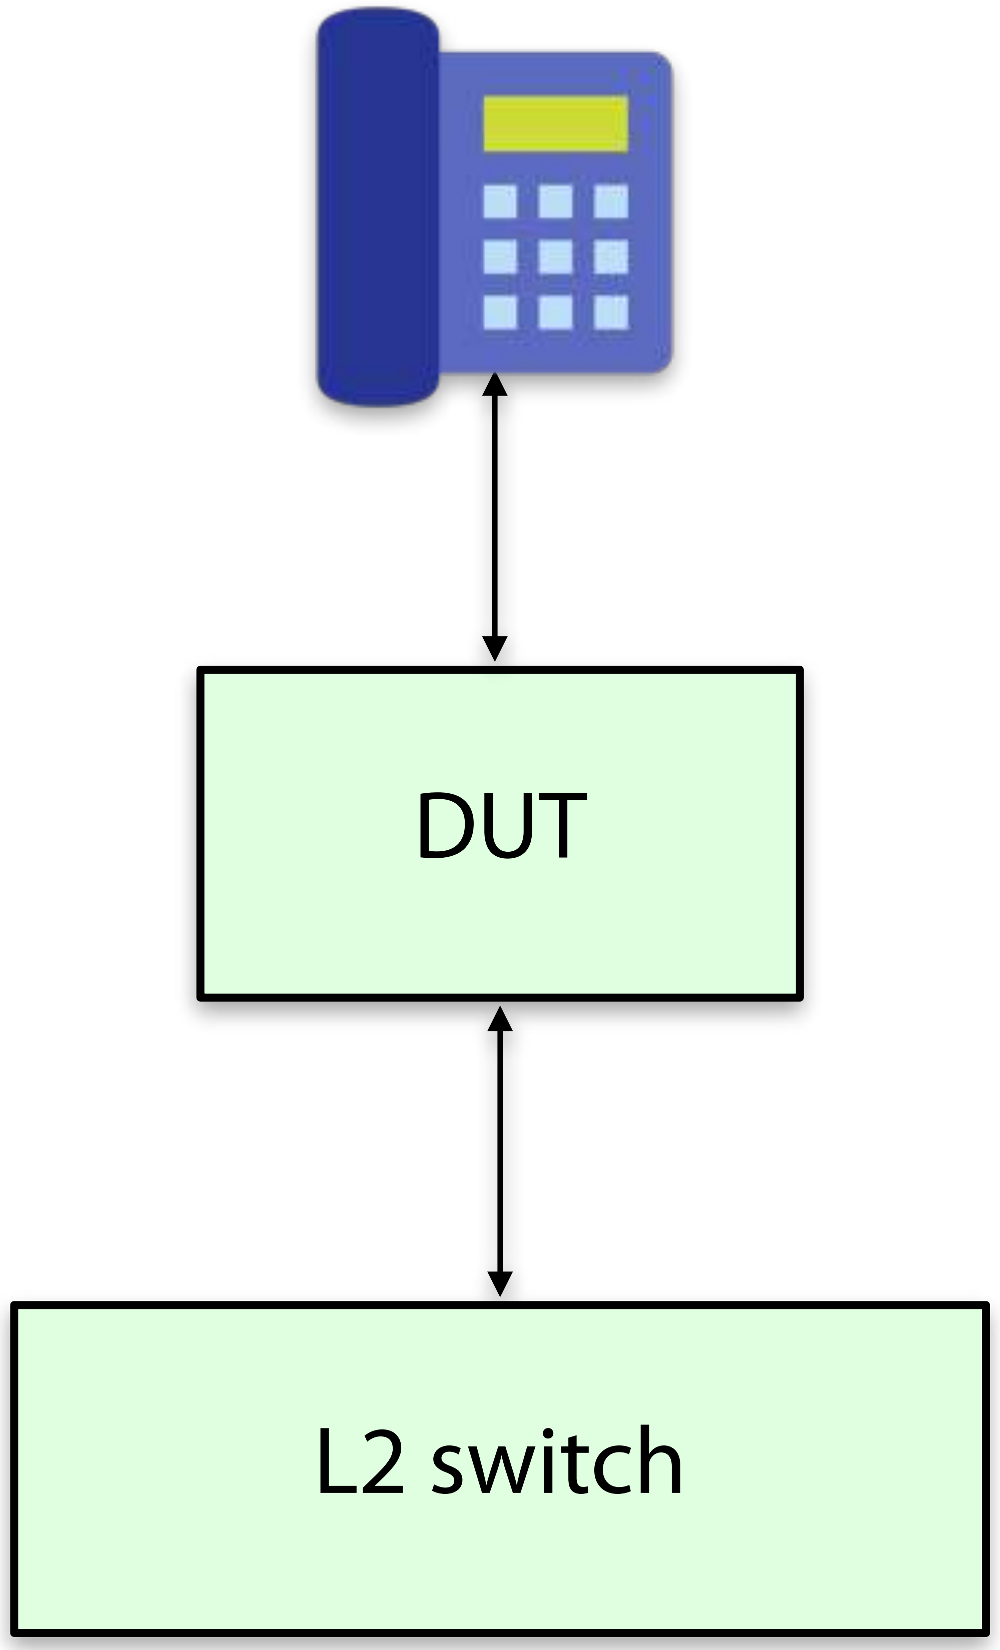

Network port features

TAP	POE Forward	Link Negotiation Forwarding	Link Failure Propagation	BYPASS (on power failure)
DUALCOMM ETAP-2003	✅ (? W)	❌	❌	❌
NEOX PRP-SCC-1GA	✅ (12,95 W)	❌ (configurable port speed)	✅ (configurable)	✅ (configurable)
LANProbe	✅ (0,75A max at 57V)	✅	❌	✅
SharkTapUSB	✅ (350mA per pair at 57V)	✅	❌	❌ (✅ With SharkTabBYP)
ProfiShark 1G	✅ (POE+)	✅ (plus manual override)	✅	✅

Capturing bad and "special" packets

- Minimum/maximum frame size
  - Will it forward runts (<64 bytes)
  - Will it forward jumbo frames (up to 9022 bytes)
  - Will it forward giants (>9022 bytes)

- FCS errors
  - Will it forward packets with bad FCS

- 01-80-C2-00-00-00 to 01-80-C2-00-00-0F
  - IEEE 802.1D MAC Bridge Filtered MAC Group Addresses
    - https://standards.ieee.org/products-programs/regauth/grpmac/public/
    - https://interestingtraffic.nl/2017/11/21/an-oddly-specific-post-about-group_fwd_mask/
  - STP, flow control, LACP, 802.1x, LLDP, etc

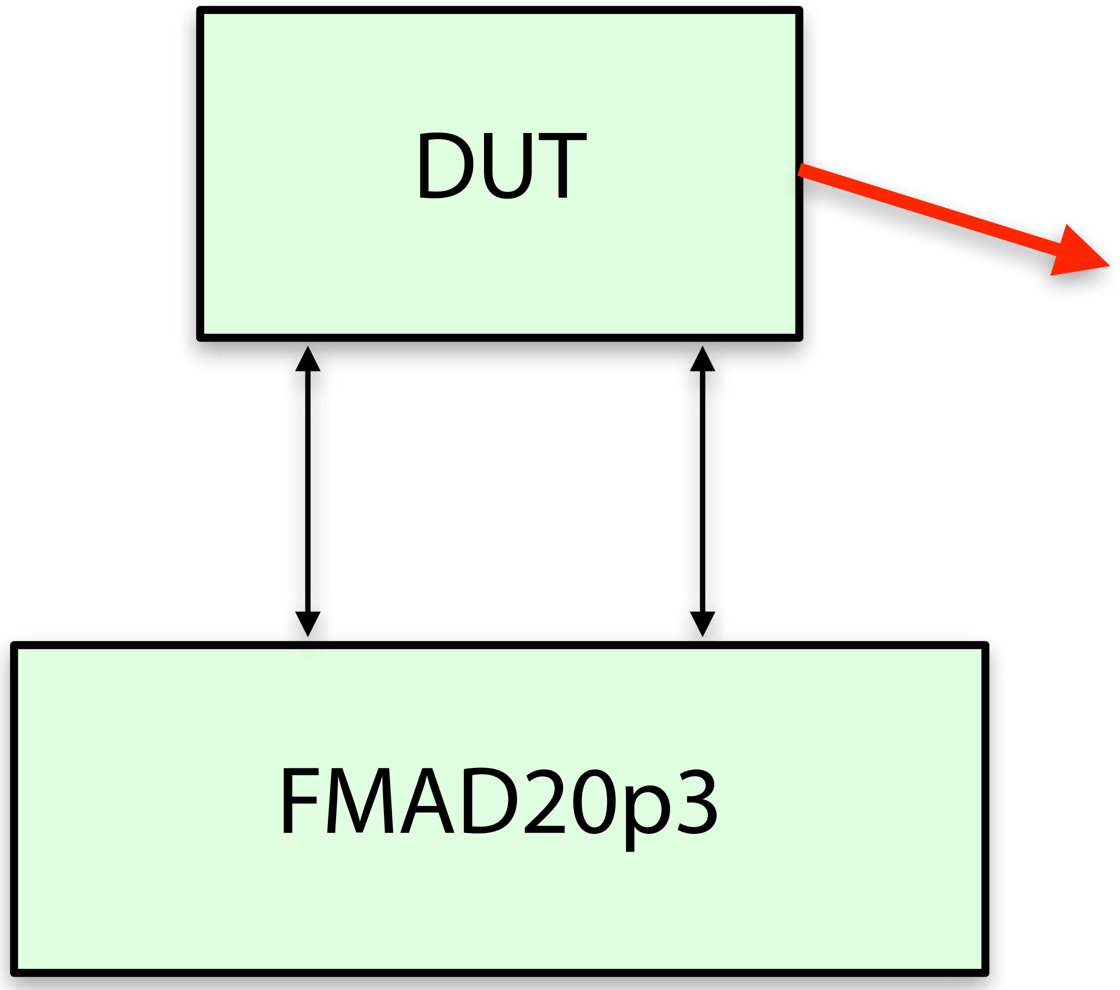

Can we capture (bad) FCS

- All but one TAP use ethernet or USB-ethernet monitor port
  - So capture dependent on NIC / Driver / OS-kernel
- FCS usually stripped by NIC
  - Override on linux: `ethtool -K <interface> rx-fcs on`
  - Out of luck on Windows / MacOS
- Bad packets usually dropped by NIC
  - Override on linux: `ethtool -K <interface> rx-all on`
  - Out of luck on Windows / MacOS

https://www.flickr.com/photos/editor/192671597

FMADIO - FMAD20p3

- Portable 2x 10Gbps sniffer
- FPGA based
- Can generate line rate traffic
  - ... and capture at the same time
- Internal path is 10Gbps, so had to conquer a few challenges for the 1Gbps tests
- Feature and new fw version created very fast!

Test script for capabilities

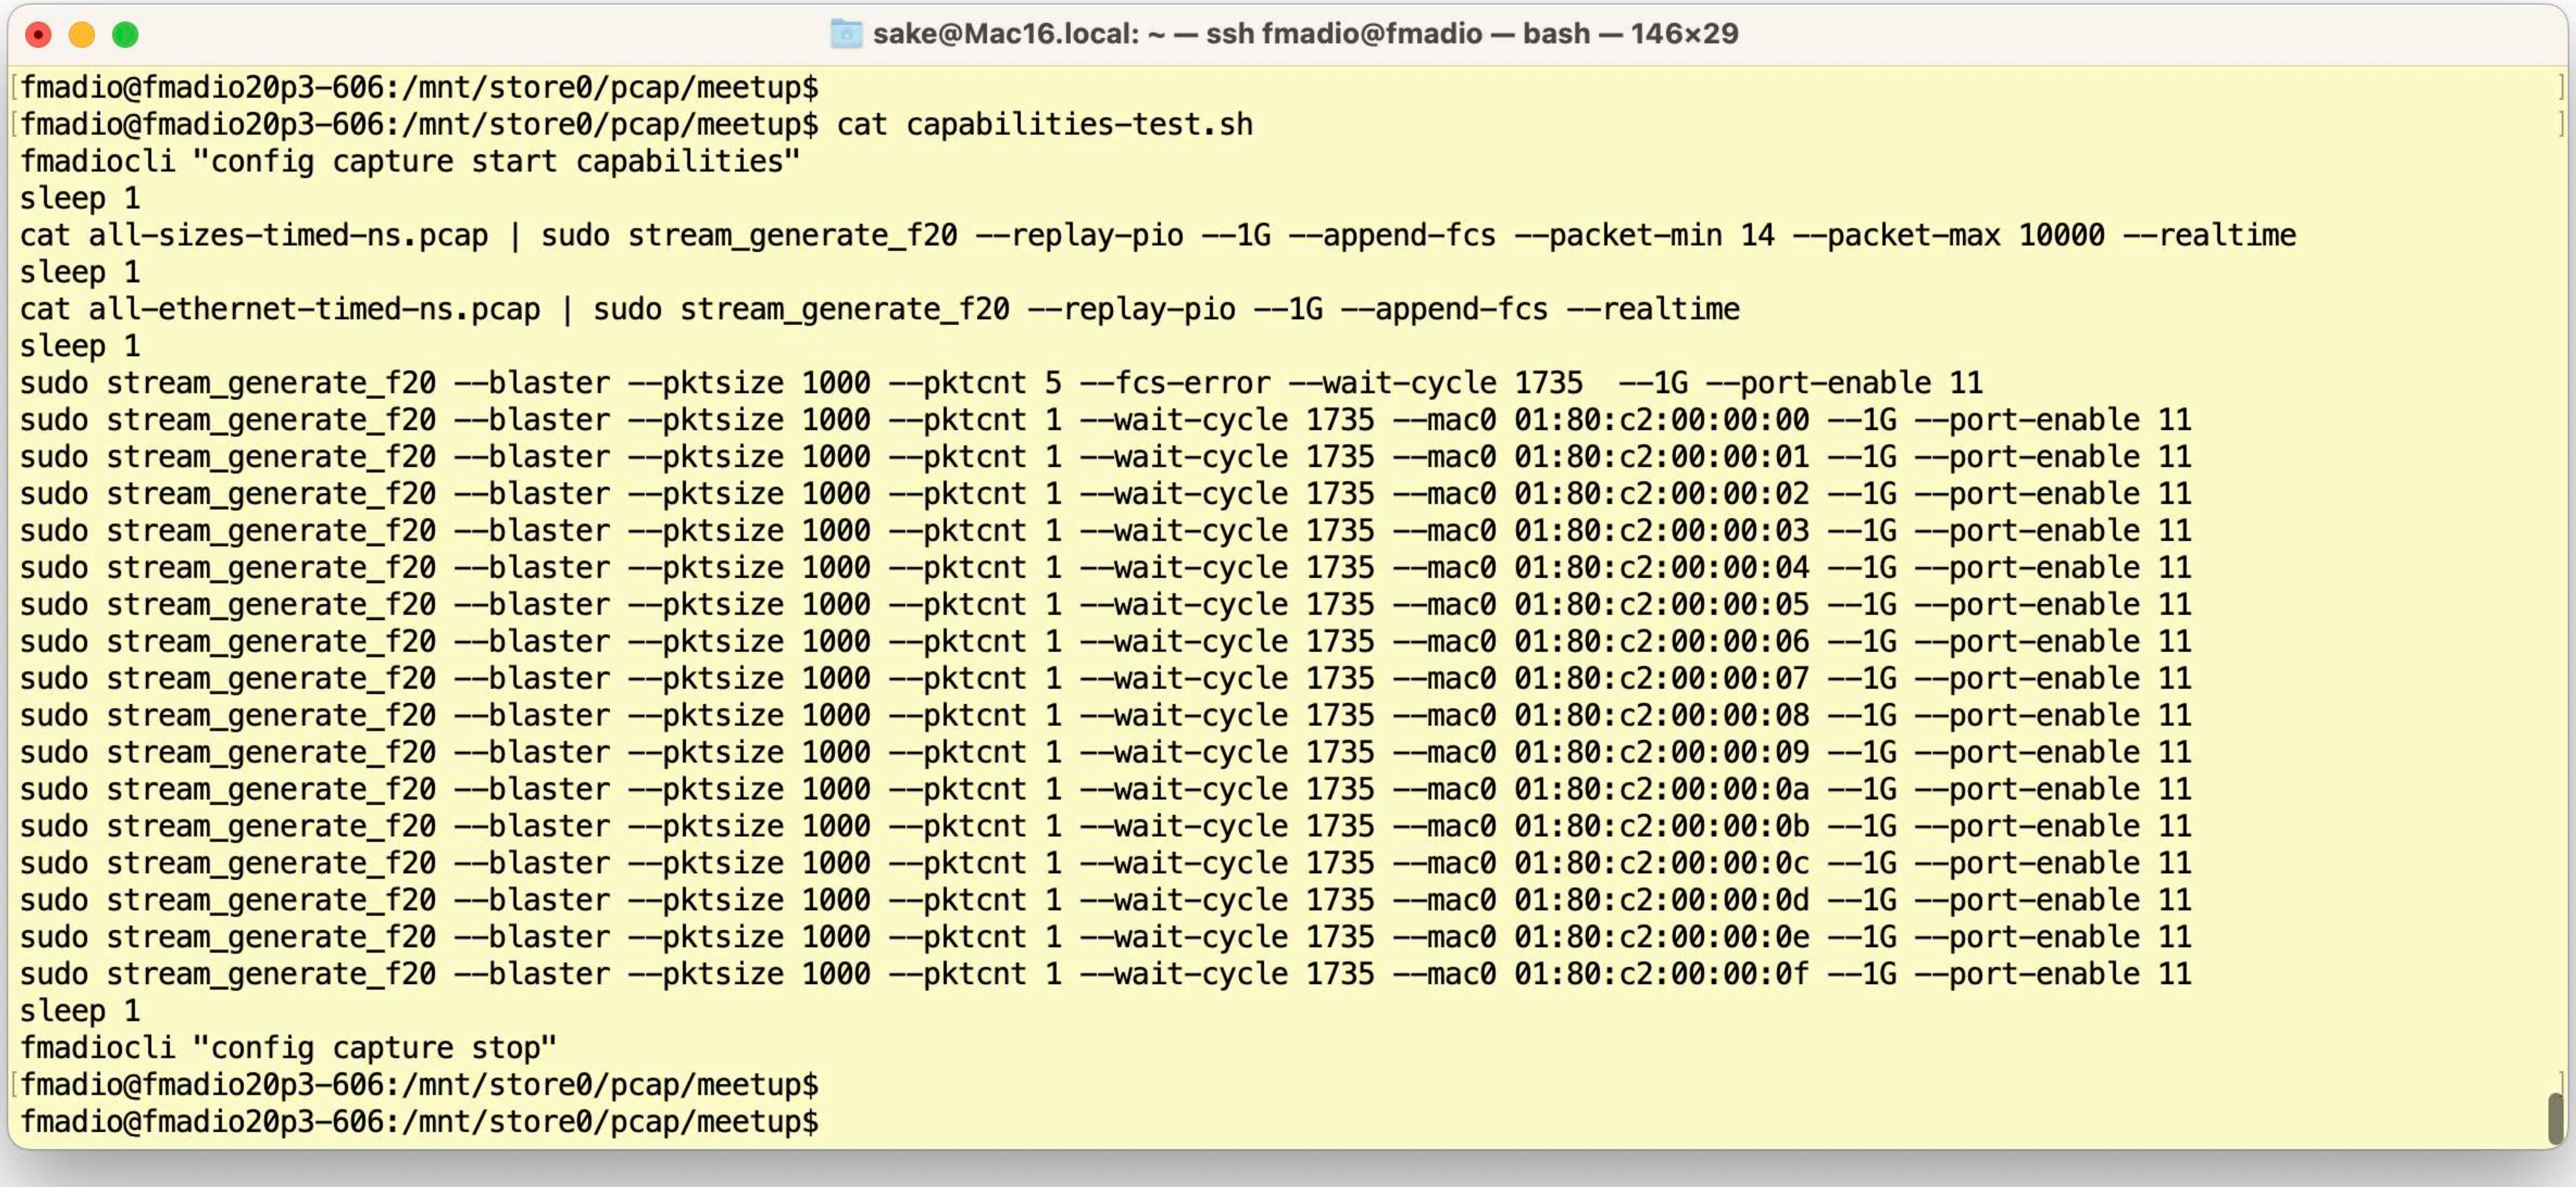

```
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$ cat capabilities-test.sh
fmadiocli "config capture start capabilities"
sleep 1
cat all-sizes-timed-ns.pcap | sudo stream_generate_f20 --replay-pio --1G --append-fcs --packet-min 14 --packet-max 10000 --realtime
sleep 1
cat all-ethernet-timed-ns.pcap | sudo stream_generate_f20 --replay-pio --1G --append-fcs --realtime
sleep 1
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 5 --fcs-error --wait-cycle 1735  --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:00 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:01 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:02 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:03 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:04 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:05 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:06 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:07 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:08 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:09 --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:0a --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:0b --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:0c --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:0d --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:0e --1G --port-enable 11
sudo stream_generate_f20 --blaster --pktsize 1000 --pktcnt 1 --wait-cycle 1735 --mac0 01:80:c2:00:00:0f --1G --port-enable 11
sleep 1
fmadiocli "config capture stop"
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$
```

FMADIO
10G 40G 100G PACKET CAPTURE
DASHBOARD
CAPTURE
REPLAY
SEARCH
FILES
CONFIG
TOOLS
ABOUT
English

PACKET BROWSER

/capture/

PCAP2	PScp	File	Bytes	GBytes	Packet Count	Description	Del
📄 GZ	🔍	test_20240227_1004	2.836.135.936 B	2 GB		Tue . 10:04:42 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0937	2.836.135.936 B	2 GB		Tue . 09:38:14 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0145	2.836.135.936 B	2 GB		Tue . 01:45:36 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0120	2.836.135.936 B	2 GB		Tue . 01:20:57 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0101	2.837.184.512 B	2 GB		Tue . 01:02:13 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0050	2.837.184.512 B	2 GB		Tue . 00:51:00 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0038	2.837.184.512 B	2 GB		Tue . 00:39:40 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0019	2.836.135.936 B	2 GB		Tue . 00:20:18 . 27-02-2024	⊗
📄 GZ	🔍	test_20240227_0008	2.836.135.936 B	2 GB		Tue . 00:09:23 . 27-02-2024	⊗

Let's test the capture laptop

- Change the interface settings:
  - "Jumbo Frames not supported on this device when CRC stripping is disabled."
  - rx-all and rx-fcs "fixed" (adapter and kernel dependent)
- MTU=1500:
  - frame sizes 18-1664 are captured
  - frames >=1665 crash the kernel (OOPS!)
- MTU=9000:
  - all 18-10000 byte frames are captured
- Packets with bad FCS are captured (rx-all on)
- All 01-80-C2-00-00-0x addresses are captured

```
echo "All enabled settings:"
sudo ethtool -k $1 | grep ": on"
echo ""
echo "Disabling offload features:"
sudo ethtool -K $1 rx off tx off sg off \
                   tso off lro off \
                   gro off gso off \
                   rxvlan off txvlan off \
                   ntuple off rxhash off
echo "Enabling rx-all and rx-fcs:"
sudo ethtool -K $1 rx-all on rx-fcs on
echo ""
echo "All enabled settings after disabling all:"
sudo ethtool -k $1 | grep ": on"
echo ""
echo "Changing MTU size to 9000"
sudo ifconfig $1
sudo ifconfig $1 down
sudo ifconfig $1 mtu 9000
sudo ifconfig $1 up
sudo ifconfig $1
```

Network port forwarding

TAP	Minimum Frame Size	Maximum Frame Size	FCS Errors	Blocked Low Level Bridge Protocols
DUALCOMM ETAP-2003	64	1522	❌	FlowControl, LACP 01-80-C2-00-00-01, -02
NEOX PRP-SCC-1GA	18	10000	✅	none
LANProbe	64	9000	❌	none[1]
SharkTapUSB	18	10000	✅	none
ProfiShark 1G (both save and eth mode)	18	10000	✅	none

[1] Ports are bridged, so if packets are sourced from the same mac address on both ports, they get dropped

Monitor port forwarding

TAP	Minimum Frame Size	Maximum Frame Size	FCS Errors	Blocked Low Level Bridge Protocols
DUALCOMM ETAP-2003	64	1522	❌	FlowControl, LACP 01-80-C2-00-00-01, -02
NEOX PRP-SCC-1GA	18 (= min tested)	10000 (= max tested)	✅	none
LANProbe - RJ45 LANProbe - USB	18 (= min tested) 64 (60)	9000 6148 (6144)	✅ ❌	none none
SharkTapUSB - RJ45 SharkTapUSB - USB	18 (= min tested) 64 (60)	10000 (= max tested) 9022 (9018)	✅ ❌	none eth.type==0x8100 (802.1Q)
ProfiShark 1G (both save and eth mode)	18 (= min tested)	10000 (= max tested)	✅	none

fast, faster, fastest

	max pps & bps @ 64 bytes		**max pps & bps @ 1518 bytes**	
	Bytes	**Bits**	**Bytes**	**Bits**
Preamble	8	64	8	64
Framesize	64	512	1518	12144
Interpacket Gap	12	96	12	96
Total	**84**	**672**	**1538**	**12304**
Half Duplex	**1,488 Mpps** **0,762 Gbps (goodput)** **1 Gbps (throughput)**		**0,081 Mpps** **0,987 Gbps (goodput)** **1 Gbps (throughput)**	
Full Duplex	**2,976 Mpps** **1,524 Gbps (goodput)** **2 Gbps (throughput)**		**0,163 Mpps** **1,974 Gbps (goodput)** **2 Gbps (throughput)**	

Speed test scripts

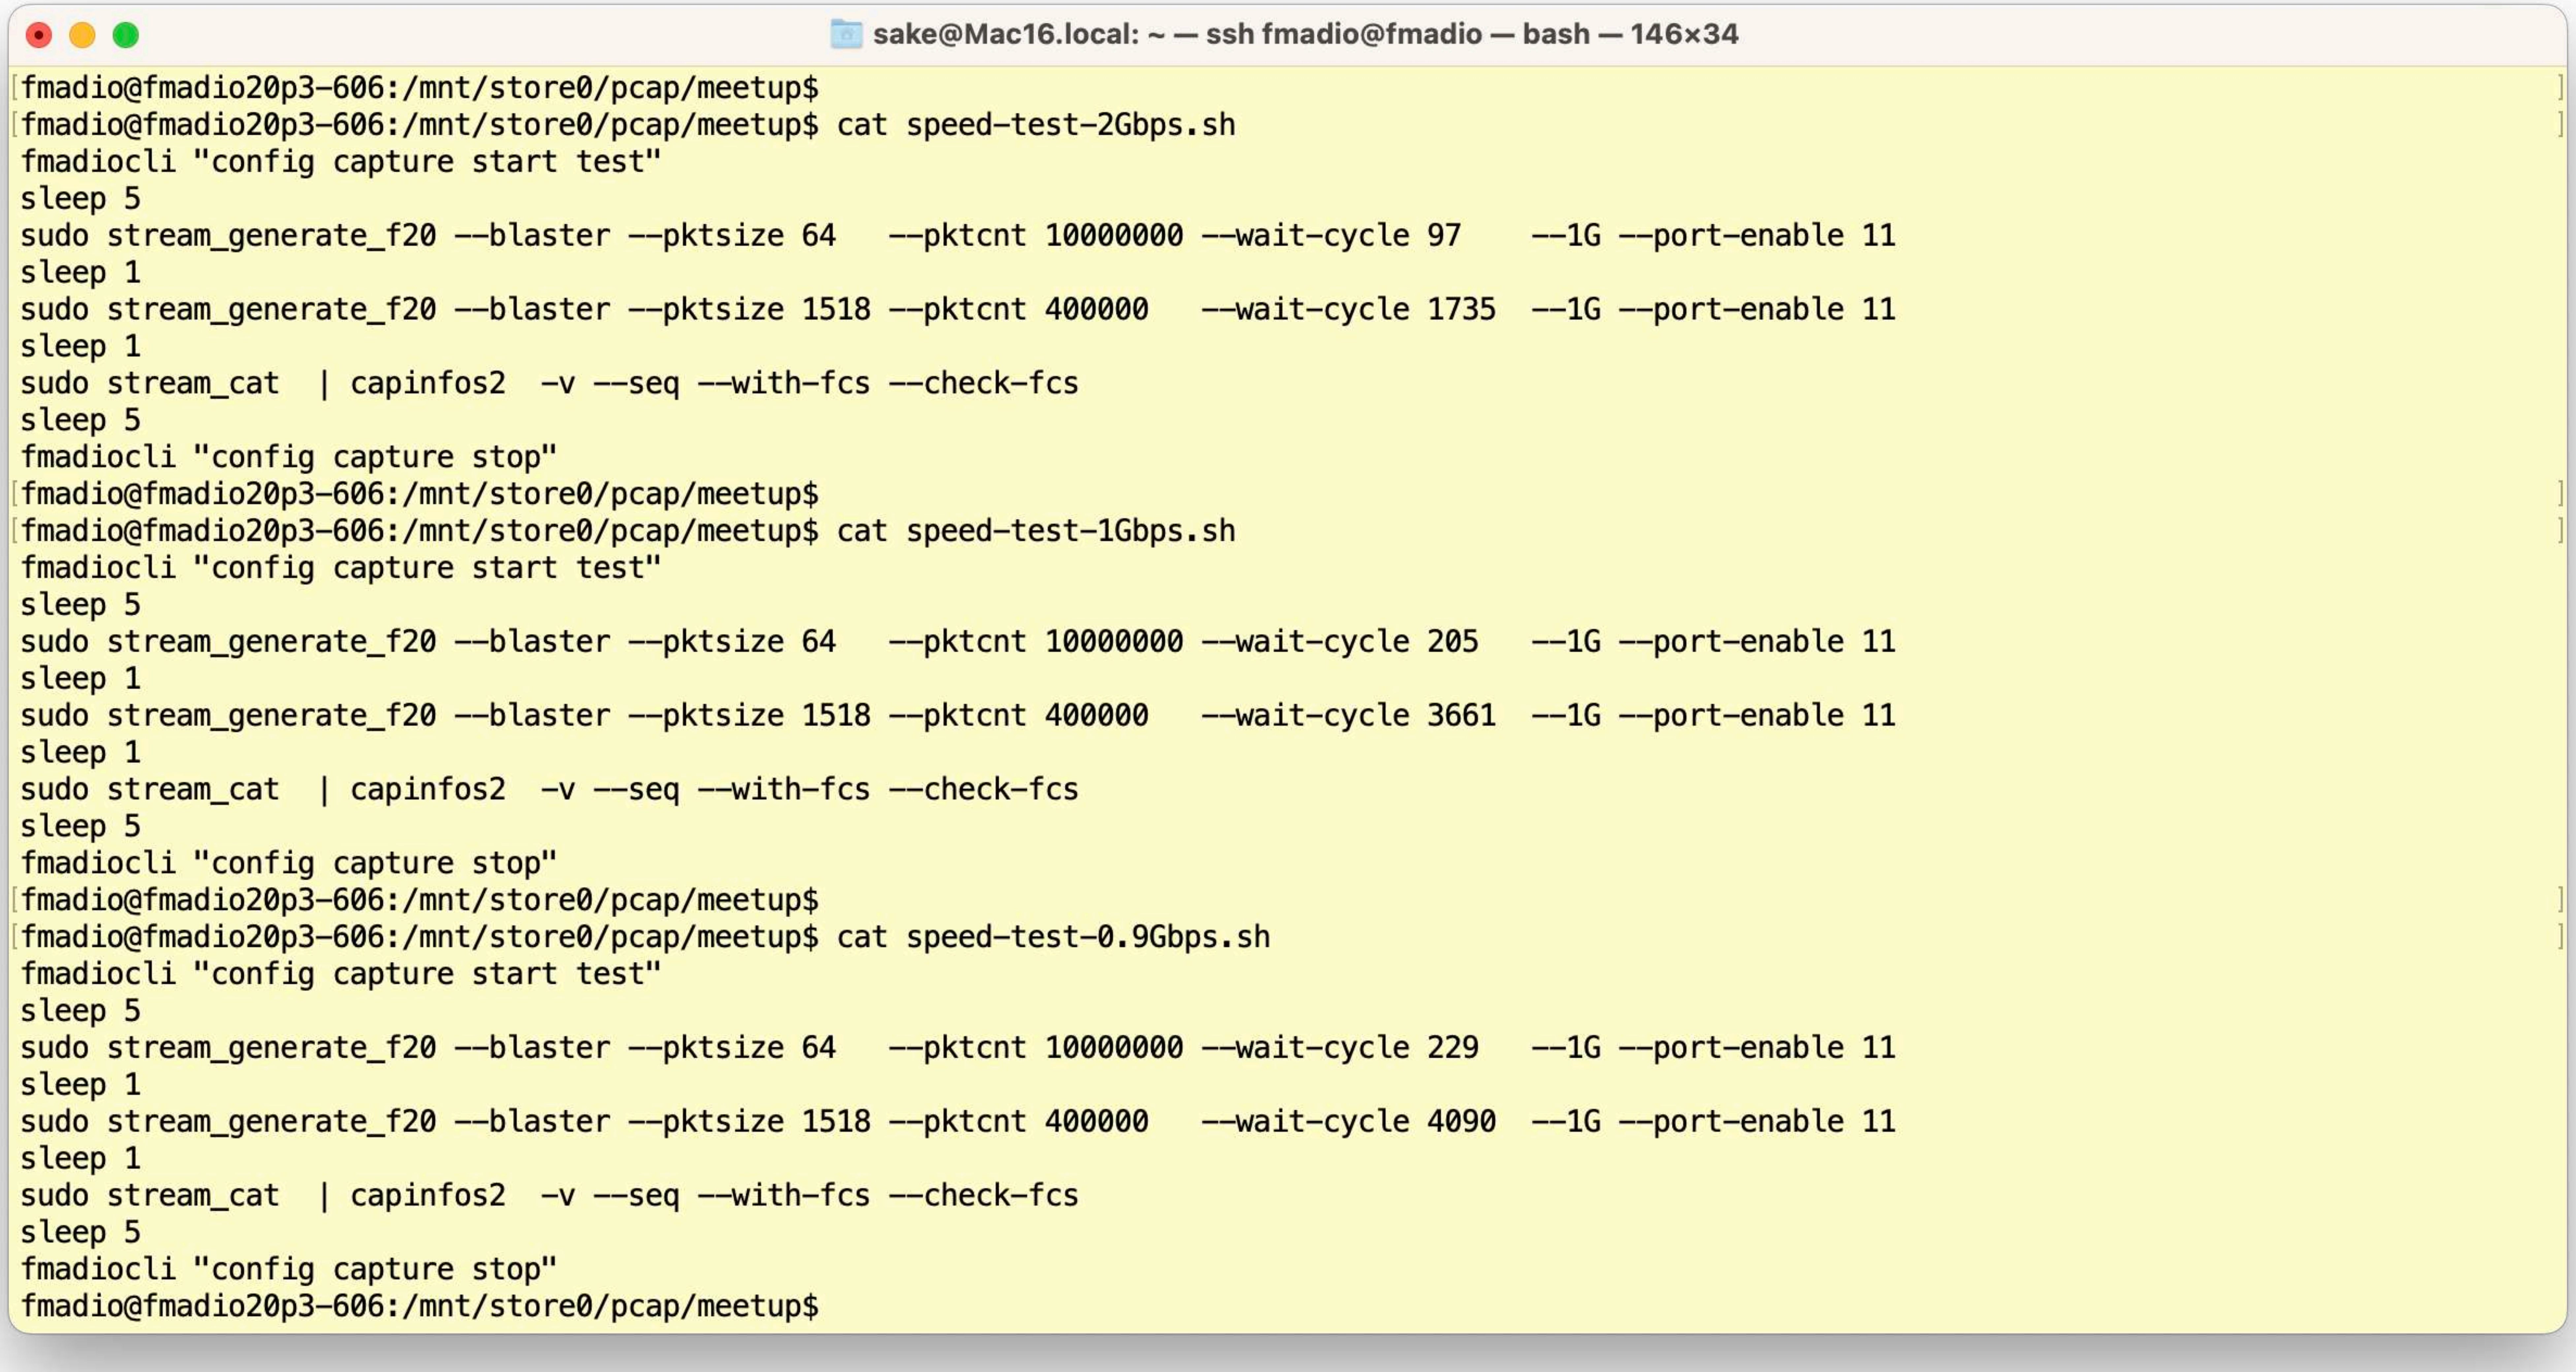

```
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$ cat speed-test-2Gbps.sh
fmadiocli "config capture start test"
sleep 5
sudo stream_generate_f20 --blaster --pktsize 64   --pktcnt 10000000 --wait-cycle 97    --1G --port-enable 11
sleep 1
sudo stream_generate_f20 --blaster --pktsize 1518 --pktcnt 400000   --wait-cycle 1735  --1G --port-enable 11
sleep 1
sudo stream_cat  | capinfos2  -v --seq --with-fcs --check-fcs
sleep 5
fmadiocli "config capture stop"
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$ cat speed-test-1Gbps.sh
fmadiocli "config capture start test"
sleep 5
sudo stream_generate_f20 --blaster --pktsize 64   --pktcnt 10000000 --wait-cycle 205   --1G --port-enable 11
sleep 1
sudo stream_generate_f20 --blaster --pktsize 1518 --pktcnt 400000   --wait-cycle 3661  --1G --port-enable 11
sleep 1
sudo stream_cat  | capinfos2  -v --seq --with-fcs --check-fcs
sleep 5
fmadiocli "config capture stop"
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$ cat speed-test-0.9Gbps.sh
fmadiocli "config capture start test"
sleep 5
sudo stream_generate_f20 --blaster --pktsize 64   --pktcnt 10000000 --wait-cycle 229   --1G --port-enable 11
sleep 1
sudo stream_generate_f20 --blaster --pktsize 1518 --pktcnt 400000   --wait-cycle 4090  --1G --port-enable 11
sleep 1
sudo stream_cat  | capinfos2  -v --seq --with-fcs --check-fcs
sleep 5
fmadiocli "config capture stop"
fmadio@fmadio20p3-606:/mnt/store0/pcap/meetup$
```

Network port performance (FD)

TAP	64 bytes @ 2.894MPps (2 Gbps)		1518 bytes @ 0.162MPps (2 Gbps)	
	Packetloss	Sequence Errors	Packetloss	Sequence Errors
DUALCOMM ETAP-2003	0	0	0	0
NEOX PRP-SCC-1GA	0	0	0	0
LANProbe	0	0	0	0
SharkTapUSB	0	0	0	0
ProfiShark 1G	0	0	0	0

Monitor port performance

TAP	64 bytes			1518 bytes		
	Packetloss	Sequence Errors	Mpps/Gbps	Packetloss	Sequence Errors	Mpps/Gbps
DUALCOMM ETAP-2003	0	0	@ 1.447MPps (1 Gbps)	0	0	@ 0.081MPps (1 Gbps)
NEOX PRP-SCC-1GA (aggregation mode)	0	0	@ 1.447MPps (1 Gbps)	0	0	@ 0.081MPps (1 Gbps)
NEOX PRP-SCC-1GA (breakout mode)	0	0	@ 2.894MPps (2 Gbps)	0	0	@ 0.162MPps (2 Gbps)
LANProbe (RJ45)	0	0	@ 1.447MPps (1 Gbps)	0	0	@ 0.081MPps (1 Gbps)
SharkTapUSB (RJ45)	*0*	*0*	***@ 1.302MPps (875 Mbps)***	*0*	*0*	***@ 0.073MPps (898 Mbps)***
ProfiShark 1G (save mode)	0	0	@ 2.894MPps (2 Gbps)	0	0	@ 0.162MPps (2 Gbps)

Packets jumping the queue!

No.	Time	Delta	Source	Destination	Identification	Protocol	Length	Info
744	0.000513520	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
745	0.000514192	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
746	0.000514864	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
747	0.000515536	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
748	0.000516208	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
749	0.000516880	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
750	0.000517552	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
751	0.000518416	0.000000864	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
752	0.000519088	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
753	0.000519760	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
754	0.000520432	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
755	0.000521104	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
756	0.000521776	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
757	0.000522448	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
758	0.000523120	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
759	0.000523792	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
760	0.000524464	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
761	0.000525328	0.000000864	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II
762	0.000526000	0.000000672	00:af:2a:03:02:00	00:af:2a:03:01:00		0x0000	64	Ethernet II
763	0.000526672	0.000000672	00:af:2a:03:01:00	00:af:2a:03:02:00		0x0000	64	Ethernet II

Aggregation Reordering

TAP	64 bytes			1518 bytes		
	Jumps	%	Mpps/Gbps	Jumps	%	Mpps/Gbps
DUALCOMM ETAP-2003	3469836	17,3%	@ 1.447MPps (0.694Gbps)	140743	17,6%	@ 0.081MPps (0.982Gbps)
NEOX PRP-SCC-1GA (aggregation mode)	0	0%	@ 1.447MPps (0.694Gbps)	0	0%	@ 0.081MPps (0.982Gbps)
LANProbe (RJ45)	142334	0,7%	@ 1.447MPps (0.694Gbps)	13242	1,7%	@ 0.081MPps (0.982Gbps)
SharkTapUSB (RJ45)	174554	0,9%	***@ 1.302MPps (0.625Gbps)***	34613	4,3%	***@ 0.073MPps (0.884Gbps)***
ProfiShark 1G (save mode)	0	0%	@ 2.894MPps (1.389Gbps)	0	0%	@ 0.162MPps (1.964Gbps)

Review

ETAP-2003 (Dualcomm)

- Pros
  - Affordable
  - 1000baseT monitor output, no drivers needed

- Cons
  - No link negotiation sync/forwarding
  - No jumbo frame support
  - Does not forward/mirror error packets
  - Filters some bridge mac-addresses
  - High amount of aggregation reorderings

PacketRaven PRP-SCC-1GA (NEOX networks)

- Pros
  - Fully transparant
  - Exact capture with no loss or aggregation reorderings
  - Many security certifications

- Cons
  - A bit on the heavy side for the laptop bag
  - Enterprise level pricing
  - Fixed port speed configuration
  - Mainly geared at permanent deployments in high security environments (is also a Pro of course!)

LANProbe (Qlinx / RTNsystems)

- Pros
  - Affordable
  - 1000baseT monitor output, no drivers needed
  - USB monitor output if you're short on ports

- Cons
  - Not fully transparant as it bridges packets
    - Won't be a problem under normal circumstances

SharkTapUSB (midBitTech)

- Pros
  - Affordable
  - 1000baseT monitor output, no drivers needed
  - USB monitor output if you're short on ports
  - Fully transparant on network ports
  - Forwards all frames to 1000baseT monitor port

- Cons
  - Does not handle full 1Gbps monitoring load (0,9 Gbps is fine)
  - Linux USB monitor port driver drops vlan tagged frames
  - No USB nic driver for MacOS

ProfiShark 1G (Profitap)

- Pros
  - Fully transparant
  - Exact (full 2Gbps) capture with no loss or aggregation reorderings
  - Timestamping on the FPGA (8 ns accuracy)
  - Inline and SPAN mode available
  - Not depending on OS and driver capabilities

- Cons
  - Enterprise level pricing
  - Driver & management software needed

Summary

- All taps are fine for medium level general traffic
  - Normal IP traffic like one workstation or VoIP phone, etc
- Capturing network errors can be a challenge
- There is a justification for enterprise level taps

- It was fun to test these TAPs!!!

- **Big thanks to: Dualcomm, Neox, RTNSystems, midBitTech, Profitap and FMADIO**

https://www.flickr.com/photos/157270154@N05/38494483572/in/album-72157689436445124/

Q&A
You have
Questions
We have
Answers

FIN/ACK/FIN/ACK

Still questions?

sake.blok@SYN-bit.nl

SYN-bit

deep traffic analysis