EtherOops

Exploring practical methods to exploit Ethernet Packet-in-Packet attacks

Ben Seri, VP of Research
Gregory Vishnepolsky, Researcher

ARMIS™

Who we are

- Prior work includes:
  - BlueBorne Bluetooth vulns
  - Urgent/11 VxWorks vulns
  - CDPwn Cisco vulns

- Researchers at Armis since 2016

- Armis is an IoT security company that allows enterprises to better identify the devices on their networks and what they’re doing

Motivations for bypassing NAT/FW

- The majority of zero-click “remote” code execution vulnerabilities require network adjacency (EternalBlue, BlueKeep)
- CDPwn / Urgent11 – single packet layer2 RCEs – how to turn them into true remote attacks?

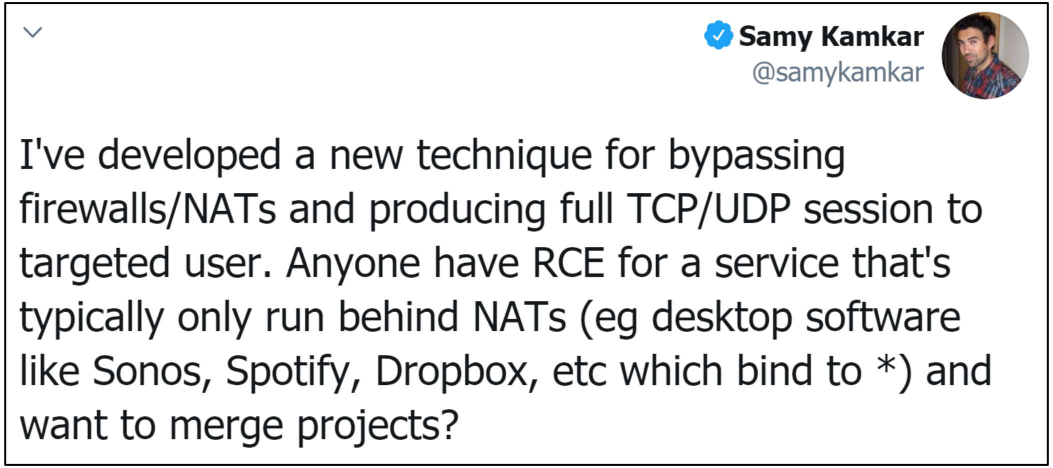
Samy Kamkar
@samykamkar

I've developed a new technique for bypassing firewalls/NATs and producing full TCP/UDP session to targeted user. Anyone have RCE for a service that's typically only run behind NATs (eg desktop software like Sonos, Spotify, Dropbox, etc which bind to *) and want to merge projects?

Attack target: Inject layer 2 packets from the Internet

- Attacker is behind a FW/NAT, needs to inject packets into the LAN
- NAT allows RELATED/ESTABLISHED connections
- Attacker can send some TCP/UDP packets that are allowed through the FW, but not anything malicious

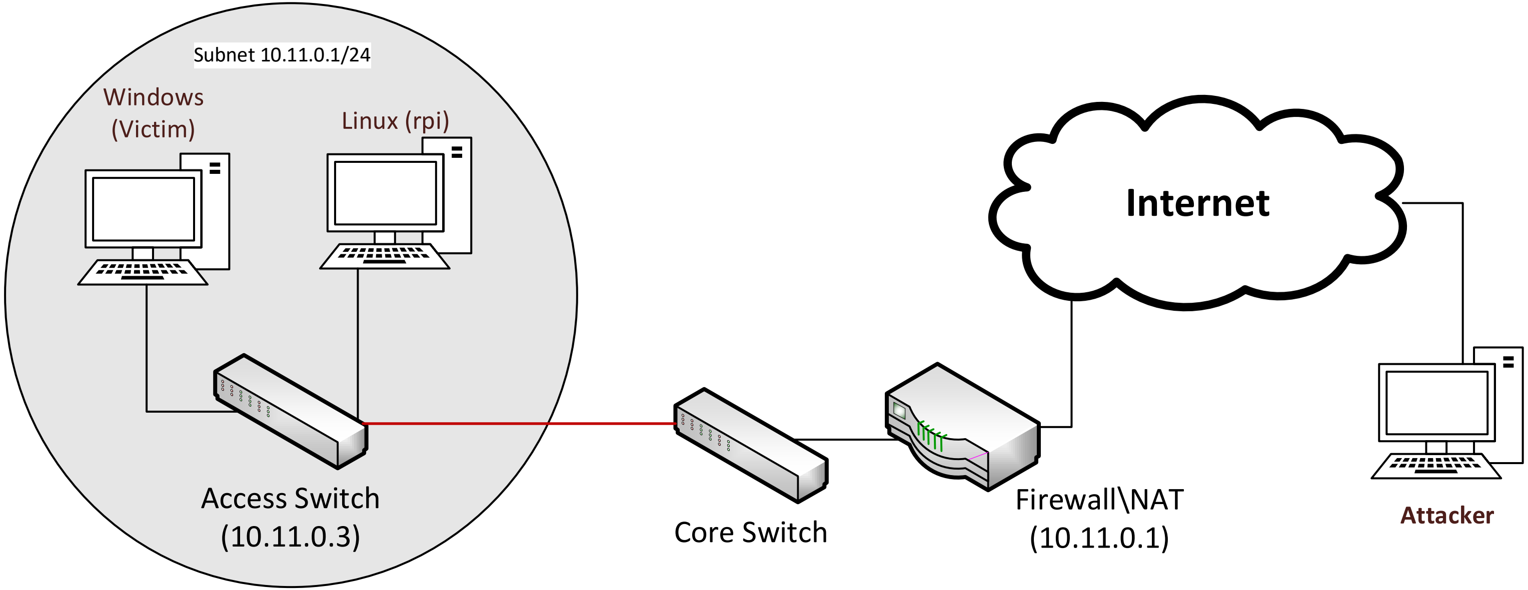

Packet-in-Packet in Ethernet???

- Travis Goodspeed – "802.11 Packets in Packets (2011, 28c3)"
  - Possibly coined the term "Packet in Packet"

Preamble	Sync	Payload
`00 00 00 00`	`a7`	`0f ...`
`00 00 00 00`	`a^`	`0f ... 00 00 00 00 a7 ...`

802.15.4 Packet-in-Packet!

- "Injection Attacks on 802.11n MAC Frame Aggregation (2015)"
  - Very nice practical tool on Github
- A significant amount of other wireless protocols, like ZigBee (802.15.4) are vulnerable to this
- But in wired protocols???

Ethernet PHY Encoding

- FastEthernet (100 Mbps) and GigabitEthernet have different PHY encodings
- In FastEthernet, 4B5B encoding is used --
  5 bit **symbol** for every 4 bits of data. Special symbols also exist:

Symbol	4B5B code	Description
H	00100	Halt
I	11111	Idle
J	11000	Start #1
...		
T	01101	End

- There is **no** error detection at this layer, except for detecting invalid symbols

The GMII/RGMII hardware interface

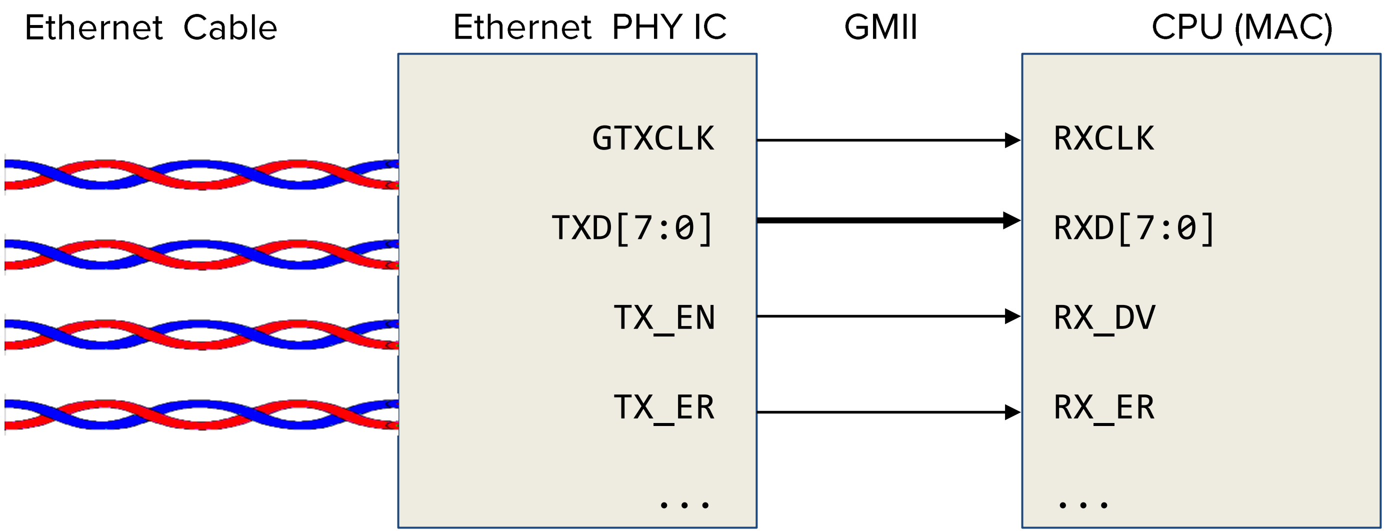

- The PHY IC translates Ethernet **symbols** to our familiar 8-bit bytes on the parallel RXD port.
- On the PHY layer, there are Start & End framing symbols
- On the GMII side, these indications are partially **in band** on the RXD port

PHY / GMII / MAC data flow

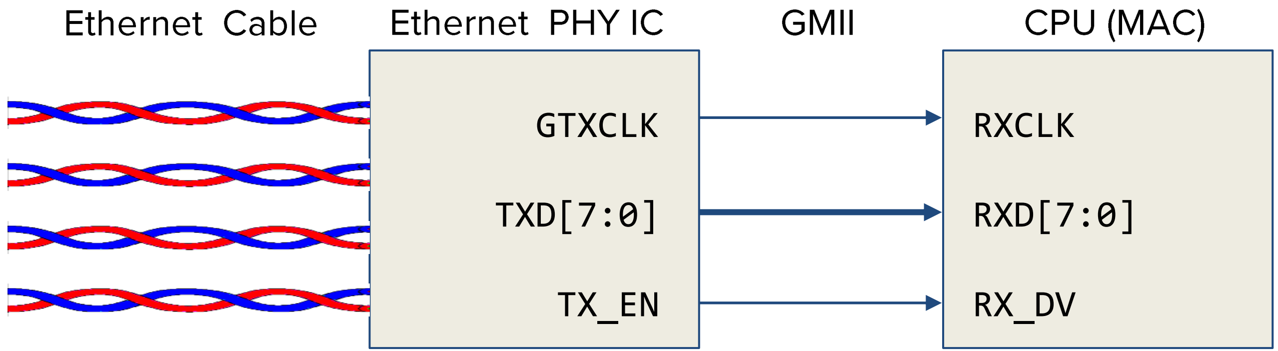

PHY	Start #1	55 55 ...	D5	...	FCS	End
GMII	RX_DV high	55 55 ...	D5	...	FCS	RX_DV low
CPU		Preamble	SFD	Payload	FCS	

<Ethernet-Header> | <IP-Header> | <TCP/UDP-Header> | <Payload>

Ethernet Packet-in-Packet data flow

<Ethernet-Header> | <IP-Header> | <TCP/UDP-Header> | <Payload>

PHY	Start #1	55 …	**D4**	…	D5	…	FCS	End
GMII	RX_DV high	55 …	**D4**	…	**D5**	…	FCS	RX_DV low
CPU		Preamble			**SFD**	PiP	FCS	

<Ethernet-Header> | <IP-Header> | <TCP/UDP-Header>

- The corrupted symbol has to be a valid data symbol (50% chance for FastEthernet, 41% for GBE)

Ethernet Packet-in-Packet – explained

Ethernet Frame

SF^ | <Ethernet-Header> | <IP-Header> | <TCP/UDP-Header> | <Payload> | <FCS>

Ethernet Packet-in-Packet – explained

Ethernet Frame

SF^ | <…> | **<** Payload: SFD | <Inner Packet> **>** <FCS >

Ethernet Packet-in-Packet – explained

Ethernet Frame

SF^ | <…> | **<** Payload: SFD | <Inner Packet> **>** <FCS>

<Ethernet-Header> | <IP-Header> | <TCP/UDP-Header> | <Payload>

The 32-bit CRC (FCS) must match **both** inner and outer packets, thus requiring the attacker to know the source\destination MAC addresses, and the internal IPs

Ethernet Packet-in-Packet – CRC32 collisions

- The CRC32 of the outer packet (the one allowed through the FW) must match the CRC32 of the inner packet (the one we want to inject).

- Therefore, a 4 byte complement is needed inside the outer packet, before the inner packet:

Outer packet header	*Complement*	*Inner packet*	CRC32

The entire *outer packet*

- CRC32(A + X + B) == CRC32(B)
  - Trivial for any A, B as long as X is 4 bytes long

Ethernet Packet-in-Packet - Prior work & background

- BH 2013, “Fully arbitrary 802.3 packet injection”, detailed the packet-in-packet scenario in Ethernet!
  However, it was deemed impractical.
  - “...though the reliability and extremely low error rate of wired cables make it unrealistic.”

- In reality, the industry standard for IEEE 802.3ab (GBE) specifies an acceptable BER of $1/10^{10}$
  - This means that one bit-flip would occur for every 10Gb of data
  - On a 1Gb/s Ethernet cable, this means a bit-flip would occur every 10 seconds!

Ethernet Cables - Survey

- At Armis, our product has access to the network infrastructure of many large enterprises in order to improve their network visibility. Additionally, it allows us to collect anonymized data.
- We added rules to extract information about Symbol Errors from all managed switches, such as using the following commands on Cisco switches:

```
#show controllers ethernet-controller | inc Sym
     0 Excessive collisions              15704 Symbol error frames
     0 Excessive collisions                  0 Symbol error frames
```

- This information is also available via SNMP, at OID 1.3.6.1.2.1.10.7.2.1.18 “dot3StatsSymbolErrors”, along with counters of all valid packets

Ethernet Cables - Survey

- The results we got from 2 large enterprise networks:

Number of active ports	Number of ports with BER of 1e-10 or more	Number of ports with BER of 1e-08 or more
71920	997 (1.3%)	230 (0.3%)
20774	298 (1.4%)	53 (0.25%)

- **When BER is 1e-08 or more, a packet-in-packet condition can occur within minutes!**
  (assuming the attacker can send packets at full line throughput)
- Each switch port above counts the errors on the series-combination of cables, connectors and sockets that lead to it.
- From this data, it's impossible to know what's faulty exactly. But the attack will still work...

Ethernet Cables - CAT 5 & 6

- In practice, the BER of Ethernet cables varies greatly
  - Short cables pretty much never experience bit flips
  - Very long cables will likely experience the standard acceptable BER (defined for a 90m max length)
  - Faulty cables might experience orders of magnitude greater BER!
- There are multiple parameters for cables
  - CAT 5/5e/6/6e/6a
  - UTP/FTP/STP
  - Length

- Any of these cables can be just as faulty as any other

Ethernet Cables - Twisted pairs

- Ethernet cables consist of 4 tightly twisted pairs of wire

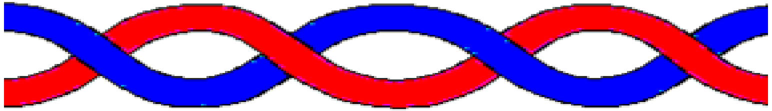

- Simplistically, in each pair, one wire will always be set to the opposite voltage of the other. The signal is the **difference** between the 2 wires in the pair.

- STP and FTP cables have additional shielding, to further prevent the interference from noise, as the twisted pairs are imperfect (and can interfere with other pairs)

- “Common mode” interference can also be a problem for receivers

Ethernet Cables - Long cables / Not shielded

CAT 5e

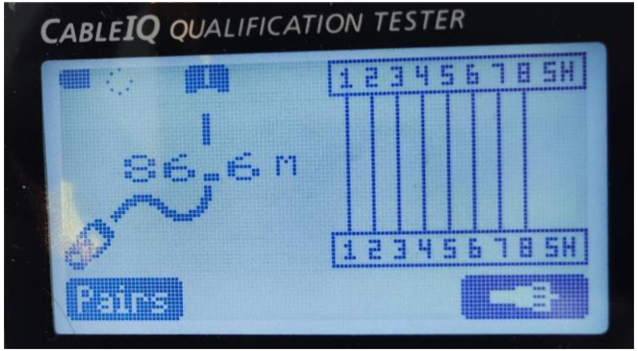

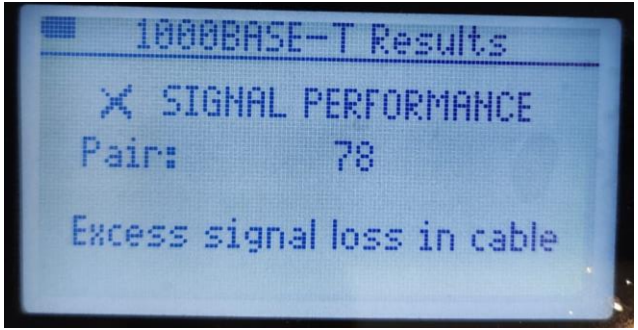

- Almost 90m long cable
- Shield not connected

Ethernet Cables - Internal short

CAT 5e FTP

- One of the pairs is shorted to another!
- Fluke calls this “bridge tap”
- The cable still appears to work at 100mb/s (which needs only 2 pairs)
- Has BER of about $1/10^7$…

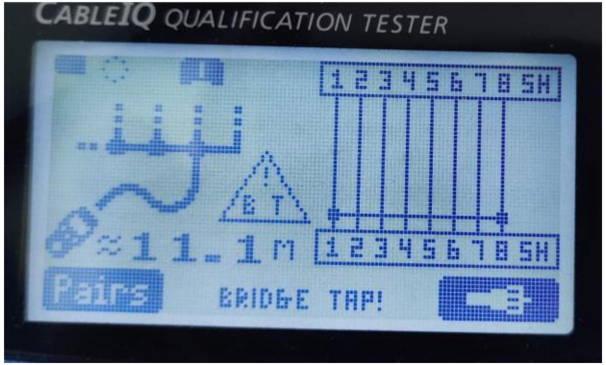

Reproduce your own faulty cables - Crosstalk

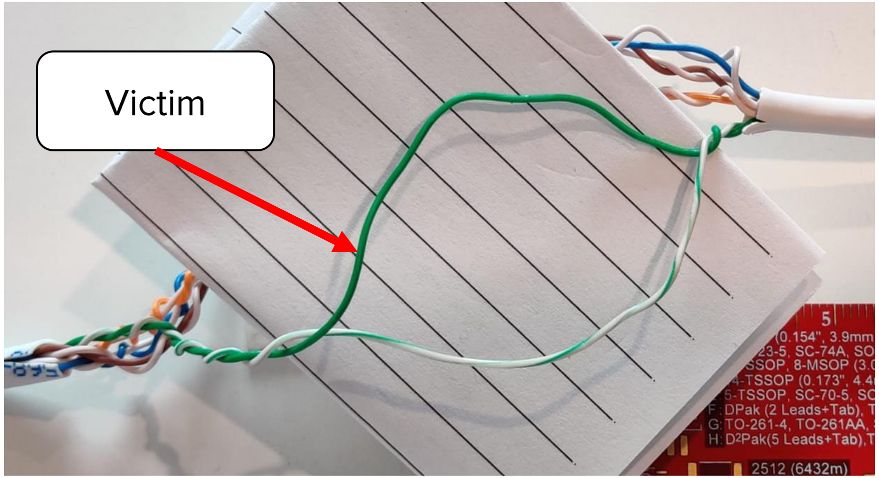

“Faulty” pair made into a loop

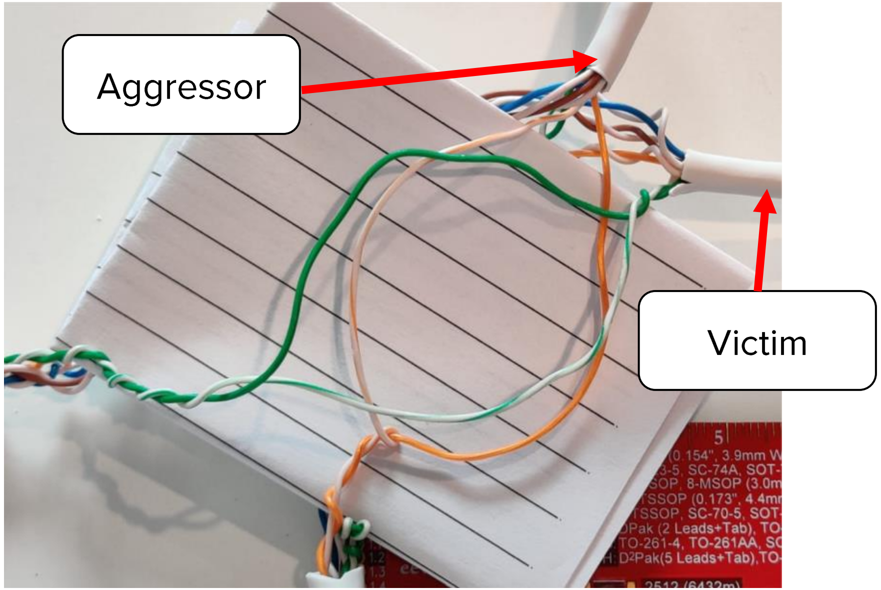

- 2 adjacent “faulty” 1GB Ethernet cables mutually experience “alien crosstalk”
- This is a **highly exaggerated** “fault”! Meant to produce 10s of bit flips per second
- 100Mb/s Ethernet can have internal crosstalk between its own TX and RX pairs

Reproduce your own faulty cables - Crosstalk

About 20-25dB of coupling between aggressor and victim from previous slide

Reproduce your own faulty cables - Short to shield

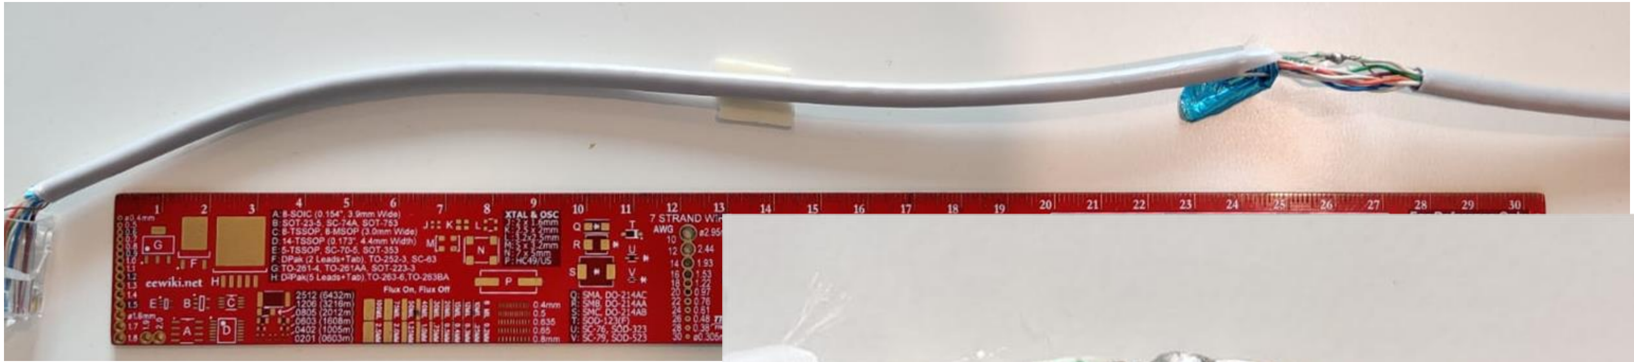

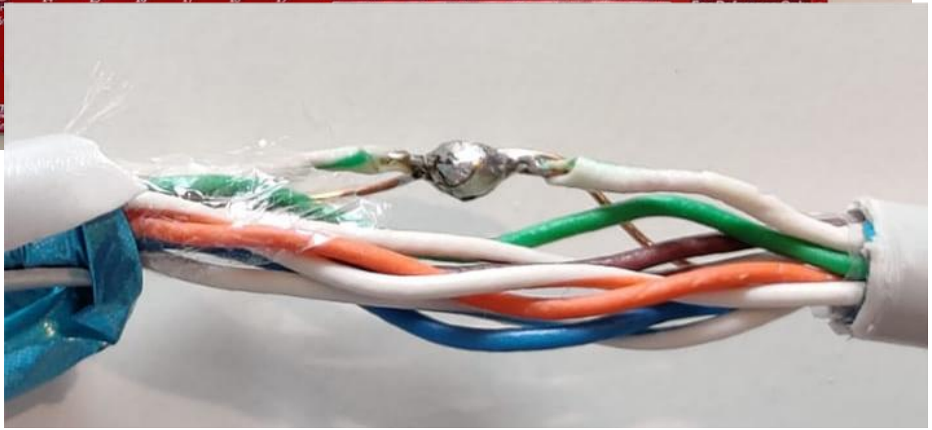

- Exaggerated fault, but does happen in reality
- A wire is shorted to the shield (the copper strand is connected to the blue aluminum shield)

- This cable is 2m long. If the shield is not connected, it’s now a 2m long antenna

Reproduce your own faulty cables - Short to shield

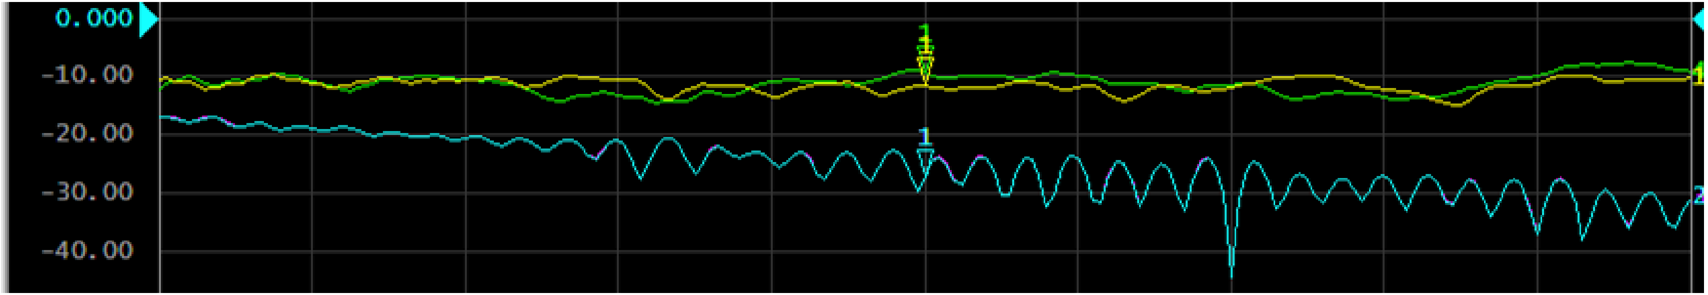

Original S-Params for non-faulty 60m cable

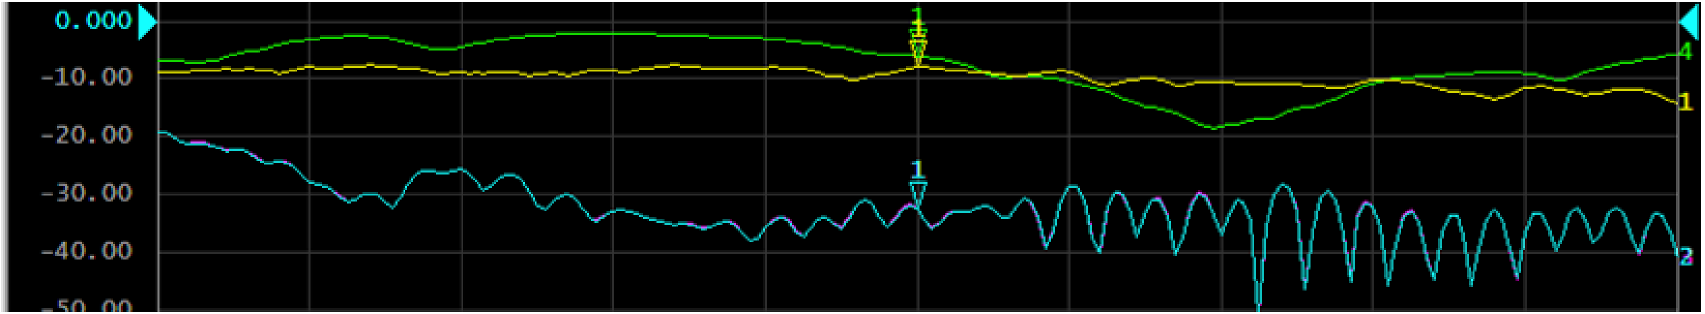

60m cable **in series** with the 2m shield-shorted cable (10dB difference!)

Ethernet Cables - Model scenario

- A model scenario for actual cables:

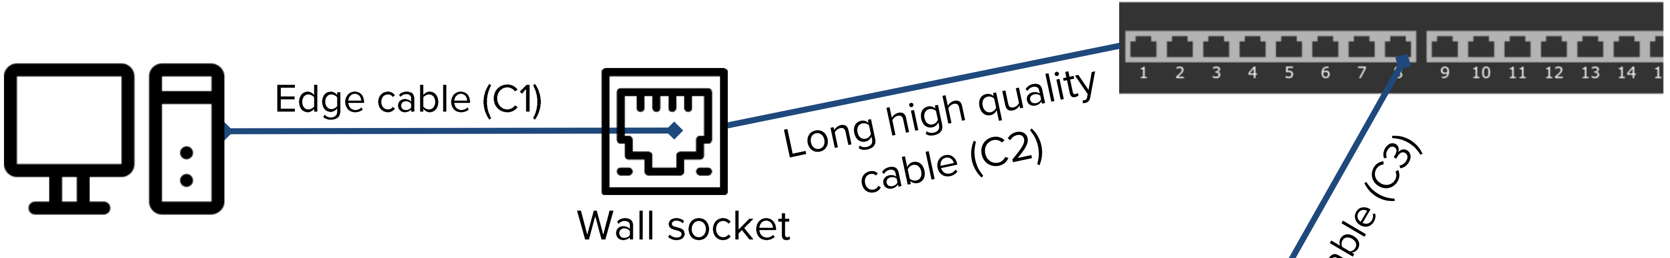

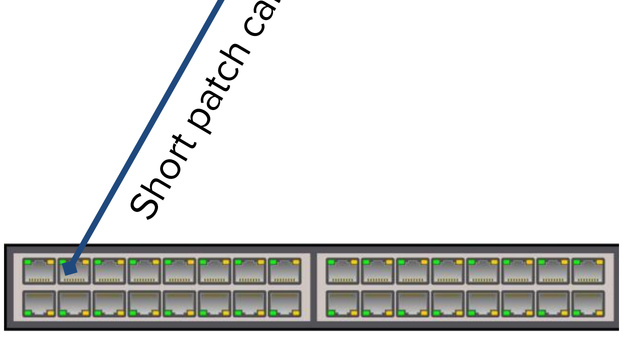

- Electrically, the “cable” consists of all these 3 cables, jacks, sockets and connectors between the device and the switch combined in series.

- It’s common for C2 to be high quality, maybe CAT 7 or above (in new deployments)

- C2 will also be a significantly longer cable

Ethernet Packet Injection – Single packet attacks?

- To clarify, an attack consists of sending lots of packets (preferably at line rate), that encapsulate our PiP payload, over the faulty cable.
- The attacker then hopes that a bit-flip will occur on the SFD, the odds of which are decreased according to the number of bytes in every packet.

- This means that an attacker can reasonably hope to inject **one** packet during an attack that may take hours. So what single packet can do the most damage?

  - 1-packet RCE attacks (CDPwn, Urgent/11)
  - Apple ICMP of death (CVE-2018-4407) (affected all Apple products)
  - IPv6 Router Advertisement
    - Allows an attacker to set DNS servers and even WPAD on Windows!

Ethernet Packet Injection – Example

Ethernet Frame

SF^ | <…> | **<** Payload: SFD | <Fake-Packet> **>** <FCS>

<Ethernet-Header> | <IP-Header> | <TCP/UDP-Header> | <Payload>

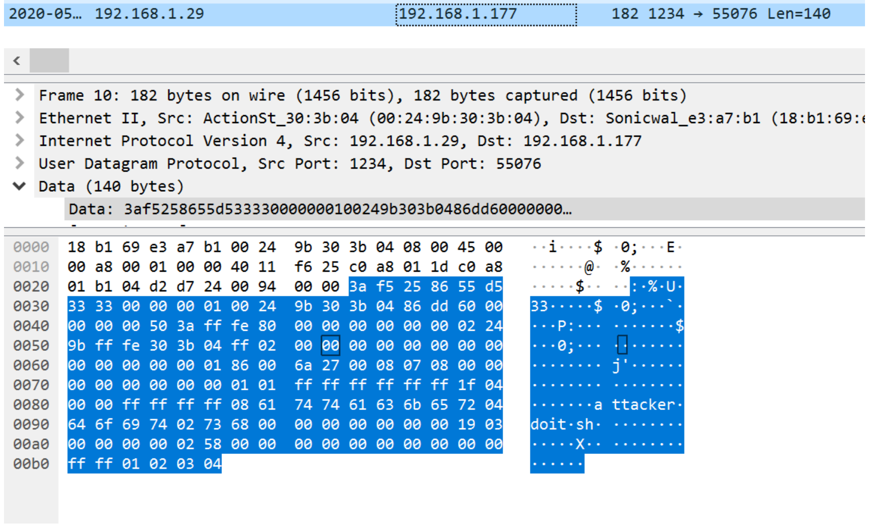

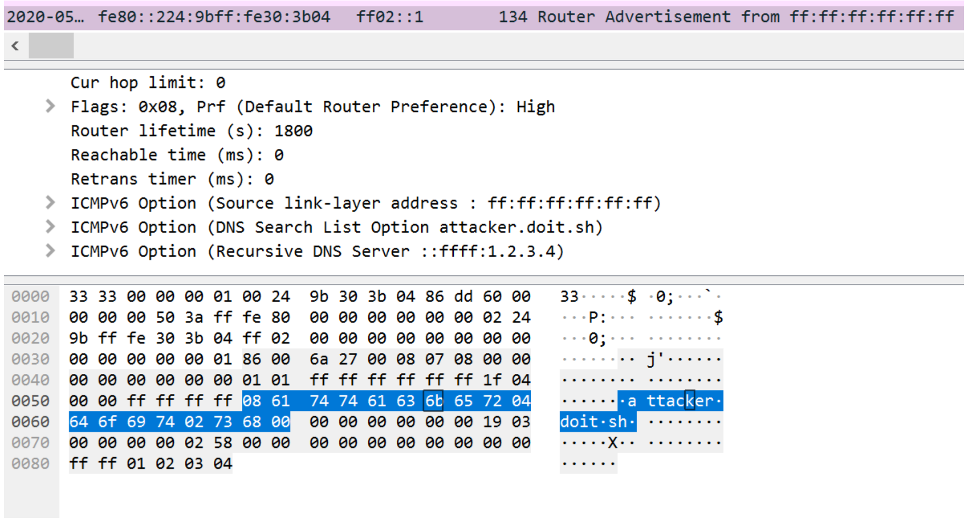

IPv6 Router Advertisement

- IPv6 is enabled by default on all interfaces in all modern OS’s
- Unlike DHCP in IPv4, an IPv6 RA can arrive unsolicited. It’s more like the ancient RARP

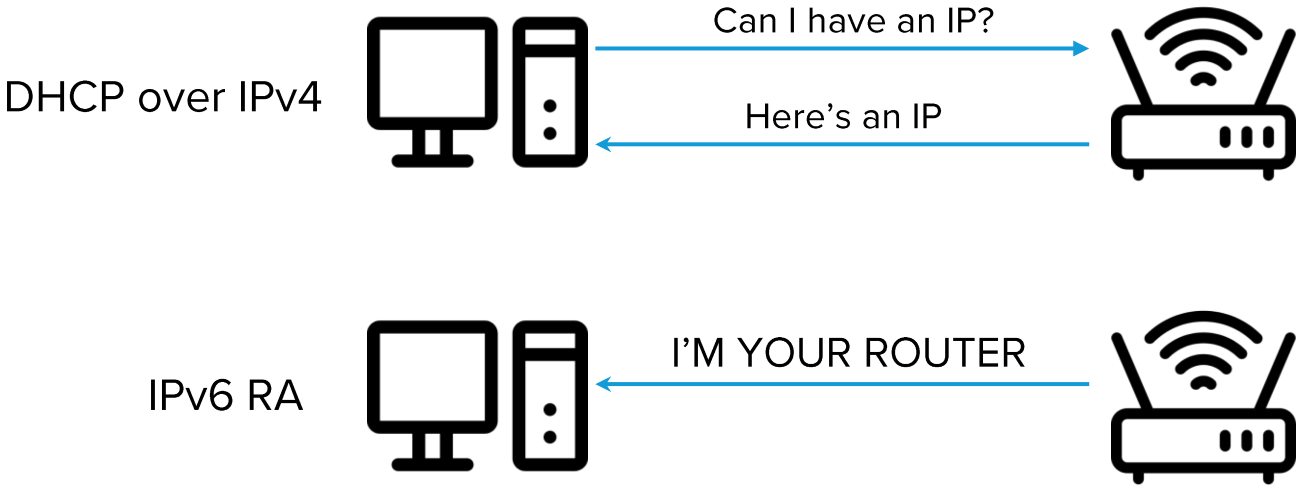

IPv6 Router Advertisement

```
2020-05…  fe80::224:9bff:fe30:3b04   ff02::1        134 Router Advertisement

    Cur hop limit: 0
  > Flags: 0x08, Prf (Default Router Preference): High
    Router lifetime (s): 1800
    Reachable time (ms): 0
    Retrans timer (ms): 0
  > ICMPv6 Option (Source link-layer address : ff:ff:ff:ff:ff:ff)
  > ICMPv6 Option (DNS Search List Option attacker.doit.sh)
  > ICMPv6 Option (Recursive DNS Server ::ffff:1.2.3.4)
```

- A working IPv6 network is not required. An attacker can add DNS servers that'll work over IPv4 using "IPv6 mapped IPv4 addresses", of the form `::ffff:X.X.X.X`
- Setting the "search domain" will force Windows machines to look for WPAD on `wpad.attacker-domain` (that too is enabled by default)

1-click Attack Scenario (+Demo)

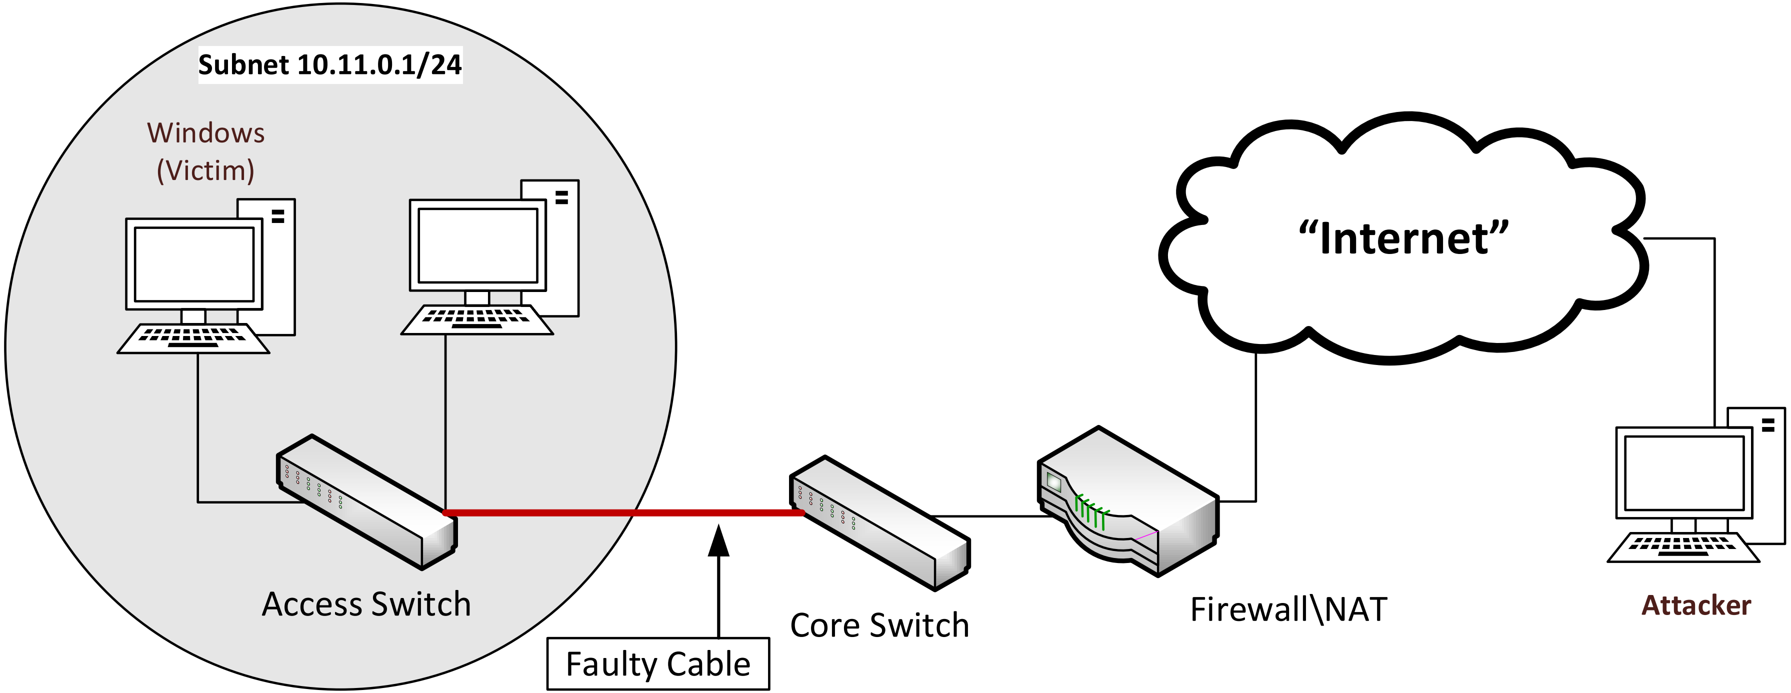

Finding out the MAC addresses

- Knowing the MACs behind the faulty cable is a requirement of the attack
  - However, MACs are not a secret!
  - FWs will have adjacent MACs for their physical ports. An attacker in a DMZ one hop from the firewall (not over the internet) will see one of them.

- WiFi exposes MAC addresses over the air
  - WPA2 encrypted traffic still has the MACs appear in clear-text
  - The exposed MACs are **the same ones** as on the wired LAN behind the AP (the AP is bridged to the LAN)
  - MACs never change. Visiting a site once, prior to the attack, is enough

MAC addresses from WiFi monitor mode

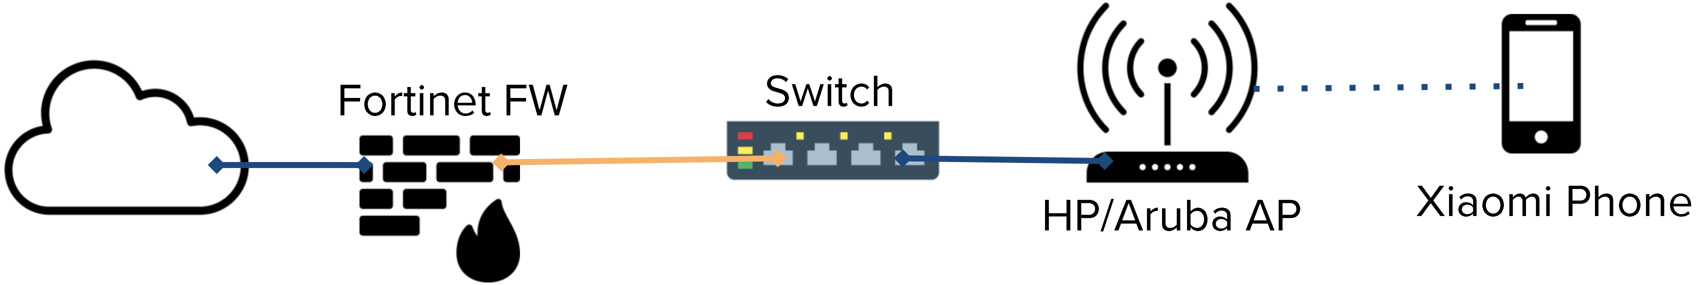


```
Fortinet_80:[redacted]   XiaomiCo_b2:33:37   QoS Data, SN=2110, FN=0, Flags=.p....F.C   802.11
Fortinet_80:[redacted]   XiaomiCo_b2:33:37   QoS Data, SN=2111, FN=0, Flags=.p....F.C   802.11
```

```
Receiver address: XiaomiCo_b2:33:37 (48:2c:a0:b2:33:37)
Transmitter address: HewlettP_cd:c9:e0 (44:48:c1:cd:c9:e0)
Destination address: XiaomiCo_b2:33:37 (48:2c:a0:b2:33:37)
Source address: Fortinet_80:[redacted]
BSS Id: HewlettP_cd:c9:e0 (44:48:c1:cd:c9:e0)
STA address: XiaomiCo_b2:33:37 (48:2c:a0:b2:33:37)
```

The Receiver & Source addresses in the 802.11 header are straight from the wired LAN

Proximity attacks

- “Faulty cables” are cables that are susceptible to normal, reasonable background EMI noise.
- But what about unreasonable noise?

- An unshielded cable, carrying an already attenuated signal, may become susceptible at higher EMI levels.

- EMP weapons are a thing.
    - These commonly use wideband pulses between 100MHz - 2GHz to interfere with any cabling longer than 5cm or so

Public research into EMP “simulation” components

- Public Research:

  [1] A Peaking Switch to Generate a High Voltage Pulse of Sub-nanosecond Rise Time [2012]
  [2] Self contained source based on an innovating resonant transformer and an oil peaking switch [2011]
  [3] An oil peaking switch to drive a dipole antenna for wideband applications
  [4] Generation of sub-nanosecond pulses using peaking capacitor [2016]
  [5] Impulse Electromagnetic Interference Generator [2004]
  [6] A 500Kv pulser with fast risetime for EMP simulation [2013]
  [7] Analysis of half TEM horn antenna for high power UWB system [2017]

- The above research describes the following:
  - Charge a capacitor to a very high voltage
  - Discharge it through a fast spark-gap in parallel to an antenna
  - Created pulse acts as a powerful ultra wide band signal

Poor man's EMP

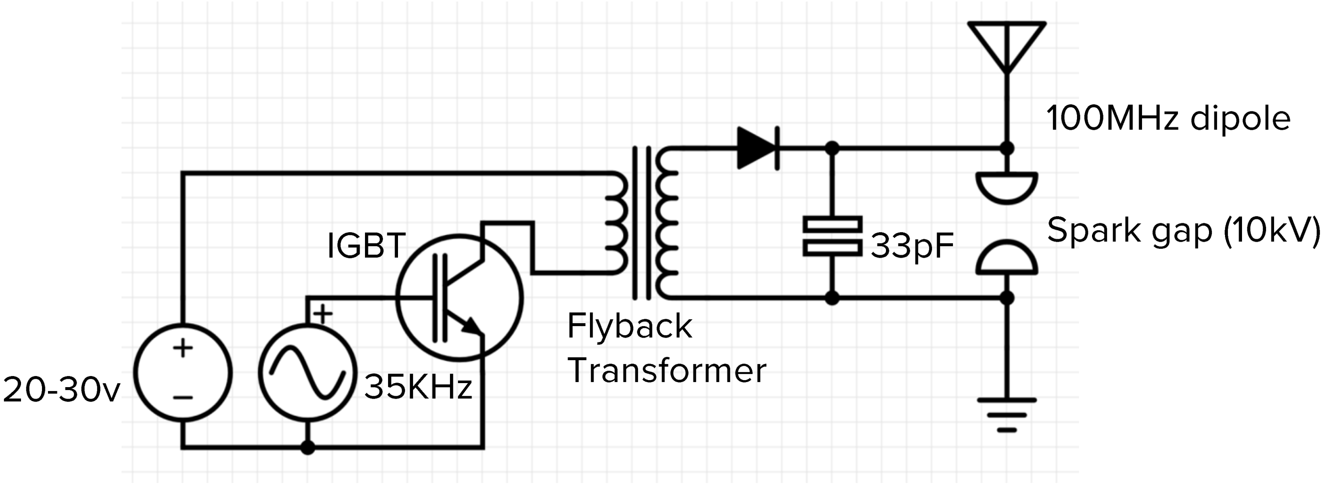

- A spark gap radio! The first kind of radio transmitter. Transmits wideband pulses at around 100MHz. Very short pulses (5-10ns) at high power.
- The discharges happen at a rate of 1-2KHz or so.
- Please don't make this, it can kill you.

Poor man's EMP

Transmitter on

Background

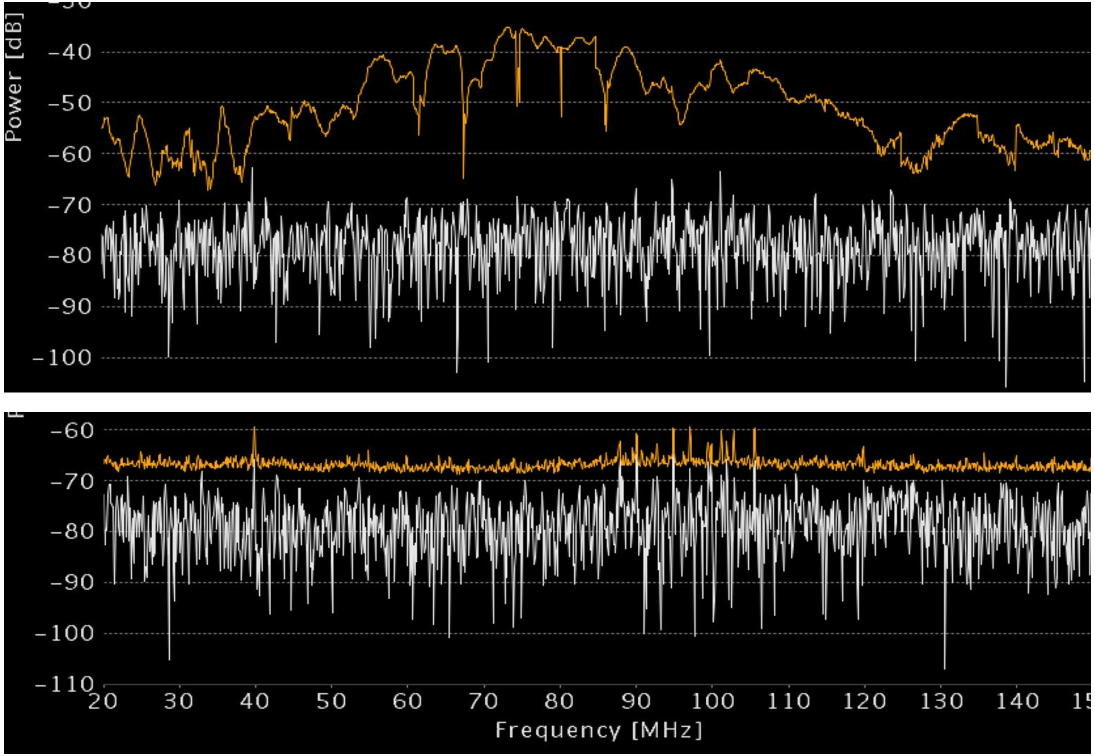

Spectrum analyzer view, 4 meters away

Poor man's EMP

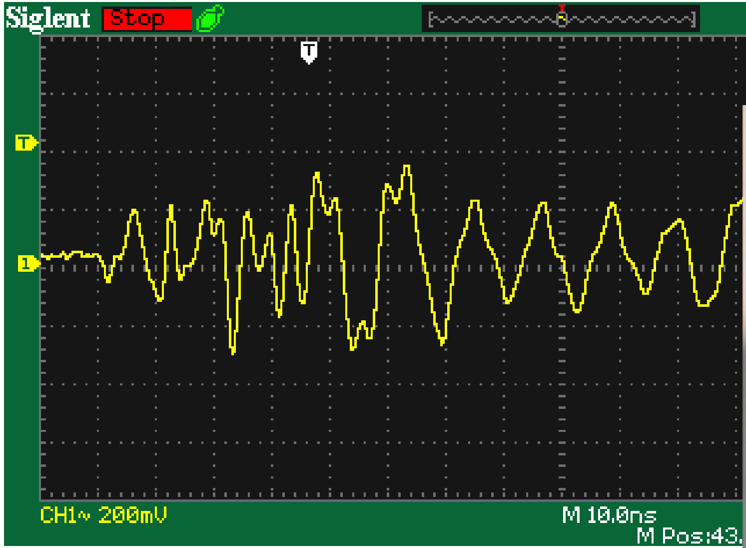

200mV / 10ns per div

- 600 millivolt peak-to-peak pulse on scope probe loop at a distance of 2.5m
- Main frequency around 80MHz

- Attenuated ethernet pairs have voltage differences in the range of 100-200 millivolts...

- The previously cited papers describe far, far more powerful setups

Poor man's EMP

200mV / 10ns per div

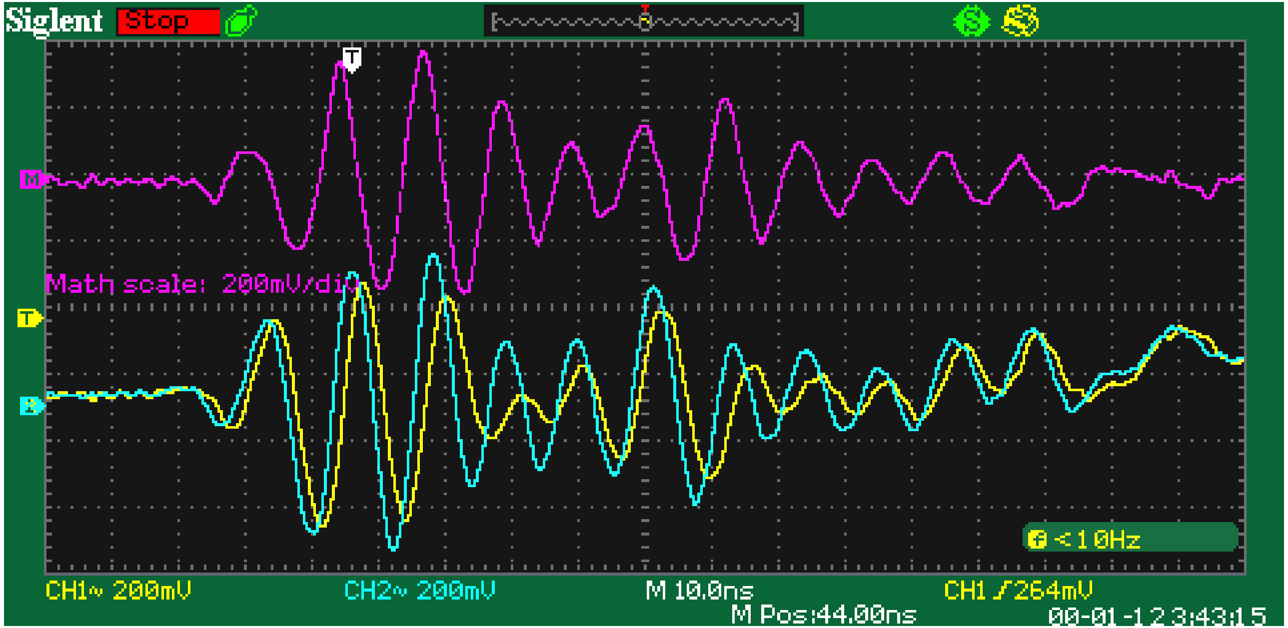

Blue and yellow are the induced voltage in 2 wires of an Ethernet twisted pair, 10m long. Purple is the differential signal.

Proximity attack scenario (+Demo)

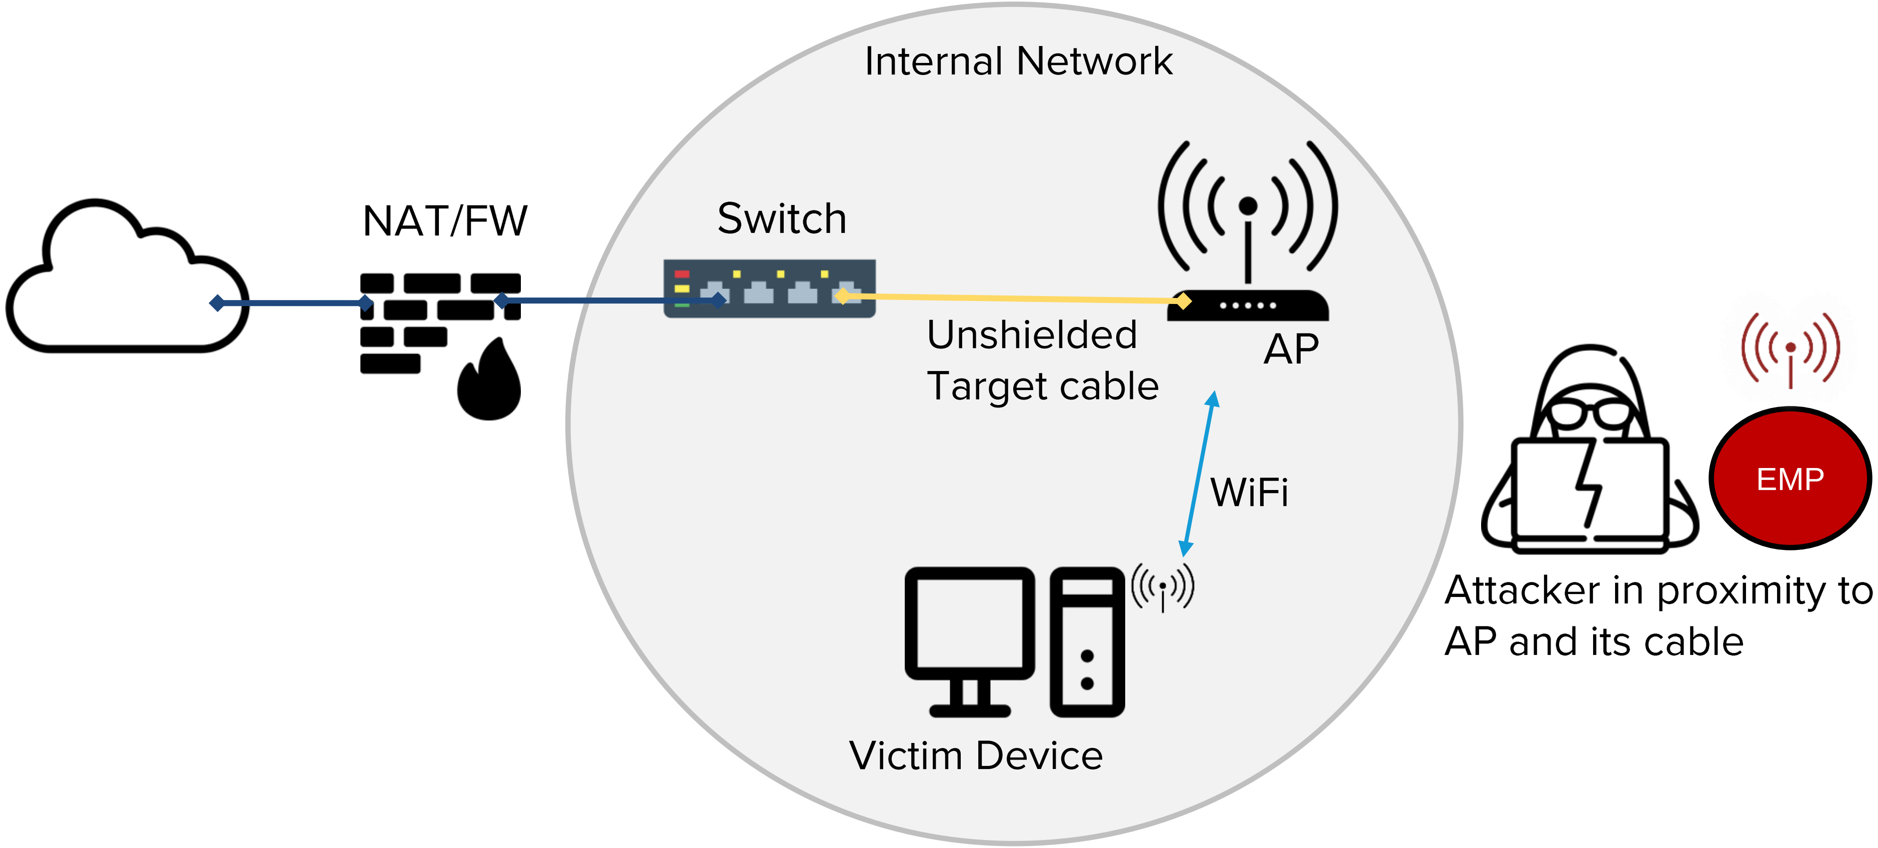

Proximity attack scenario (+Demo)

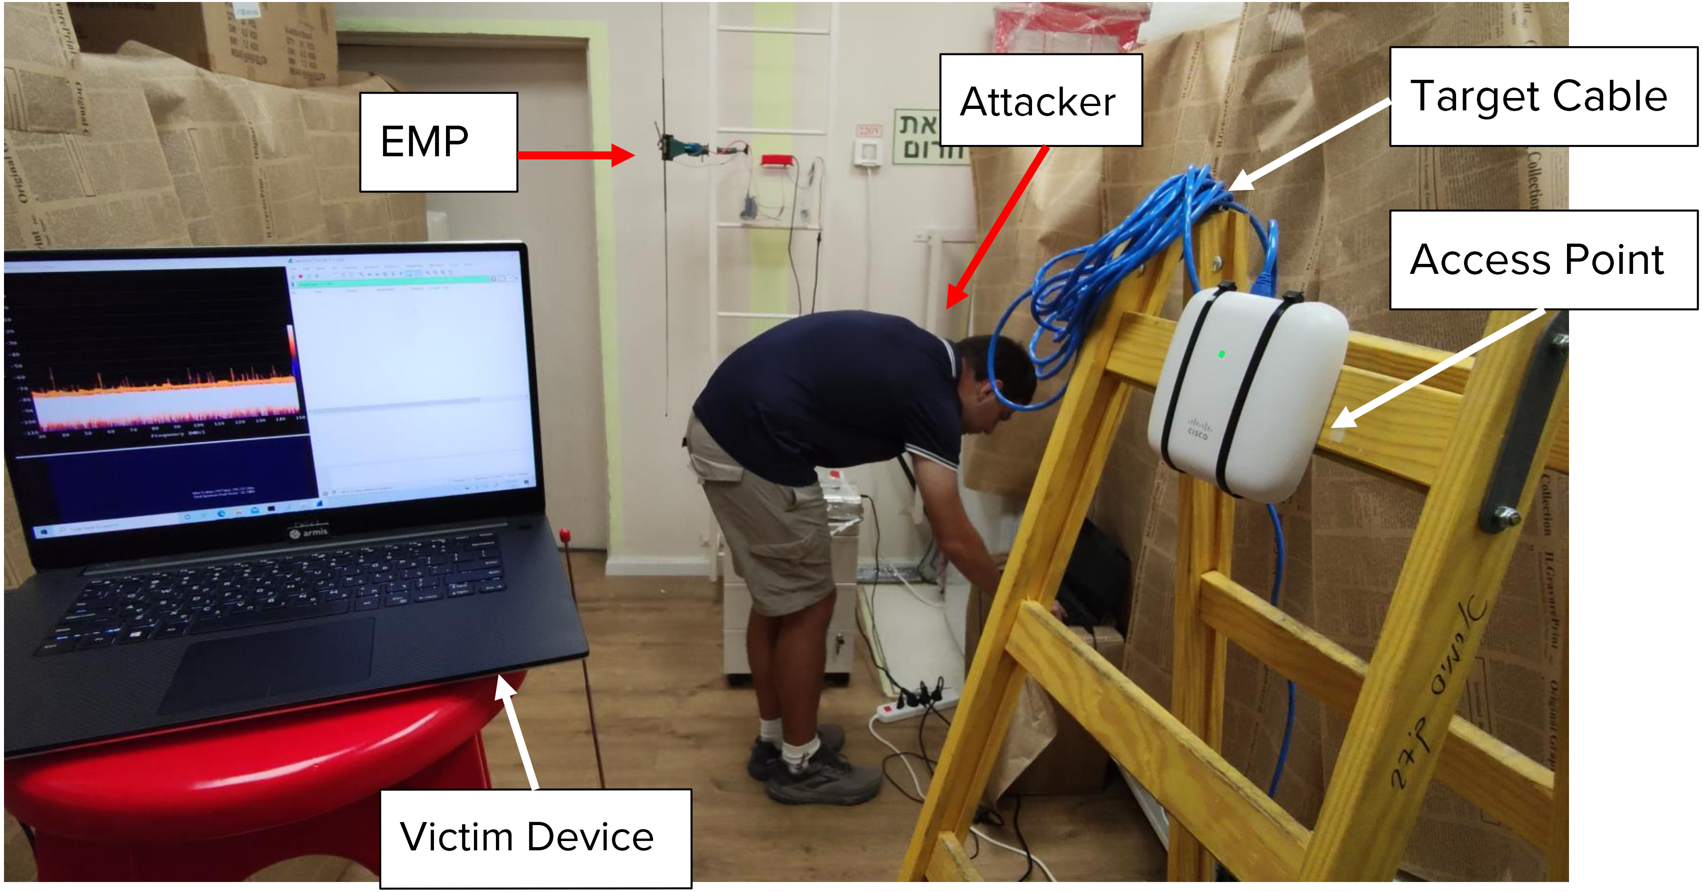

Proximity attack scenario (+Demo)

EMP

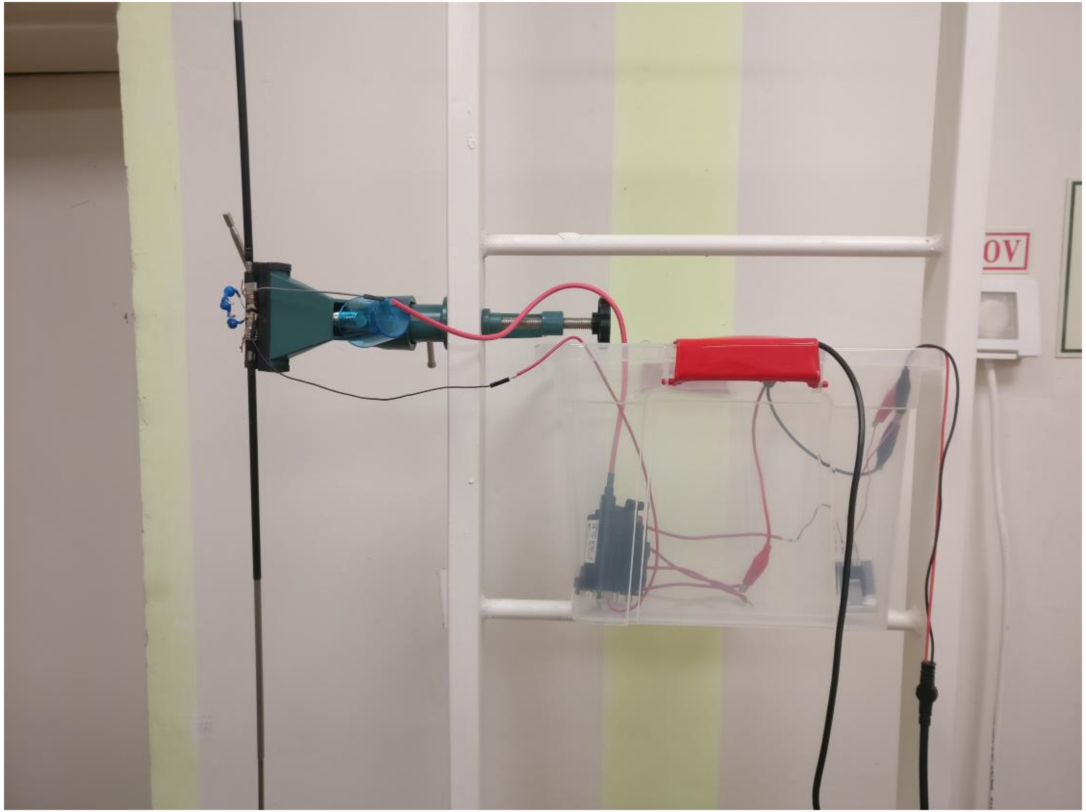

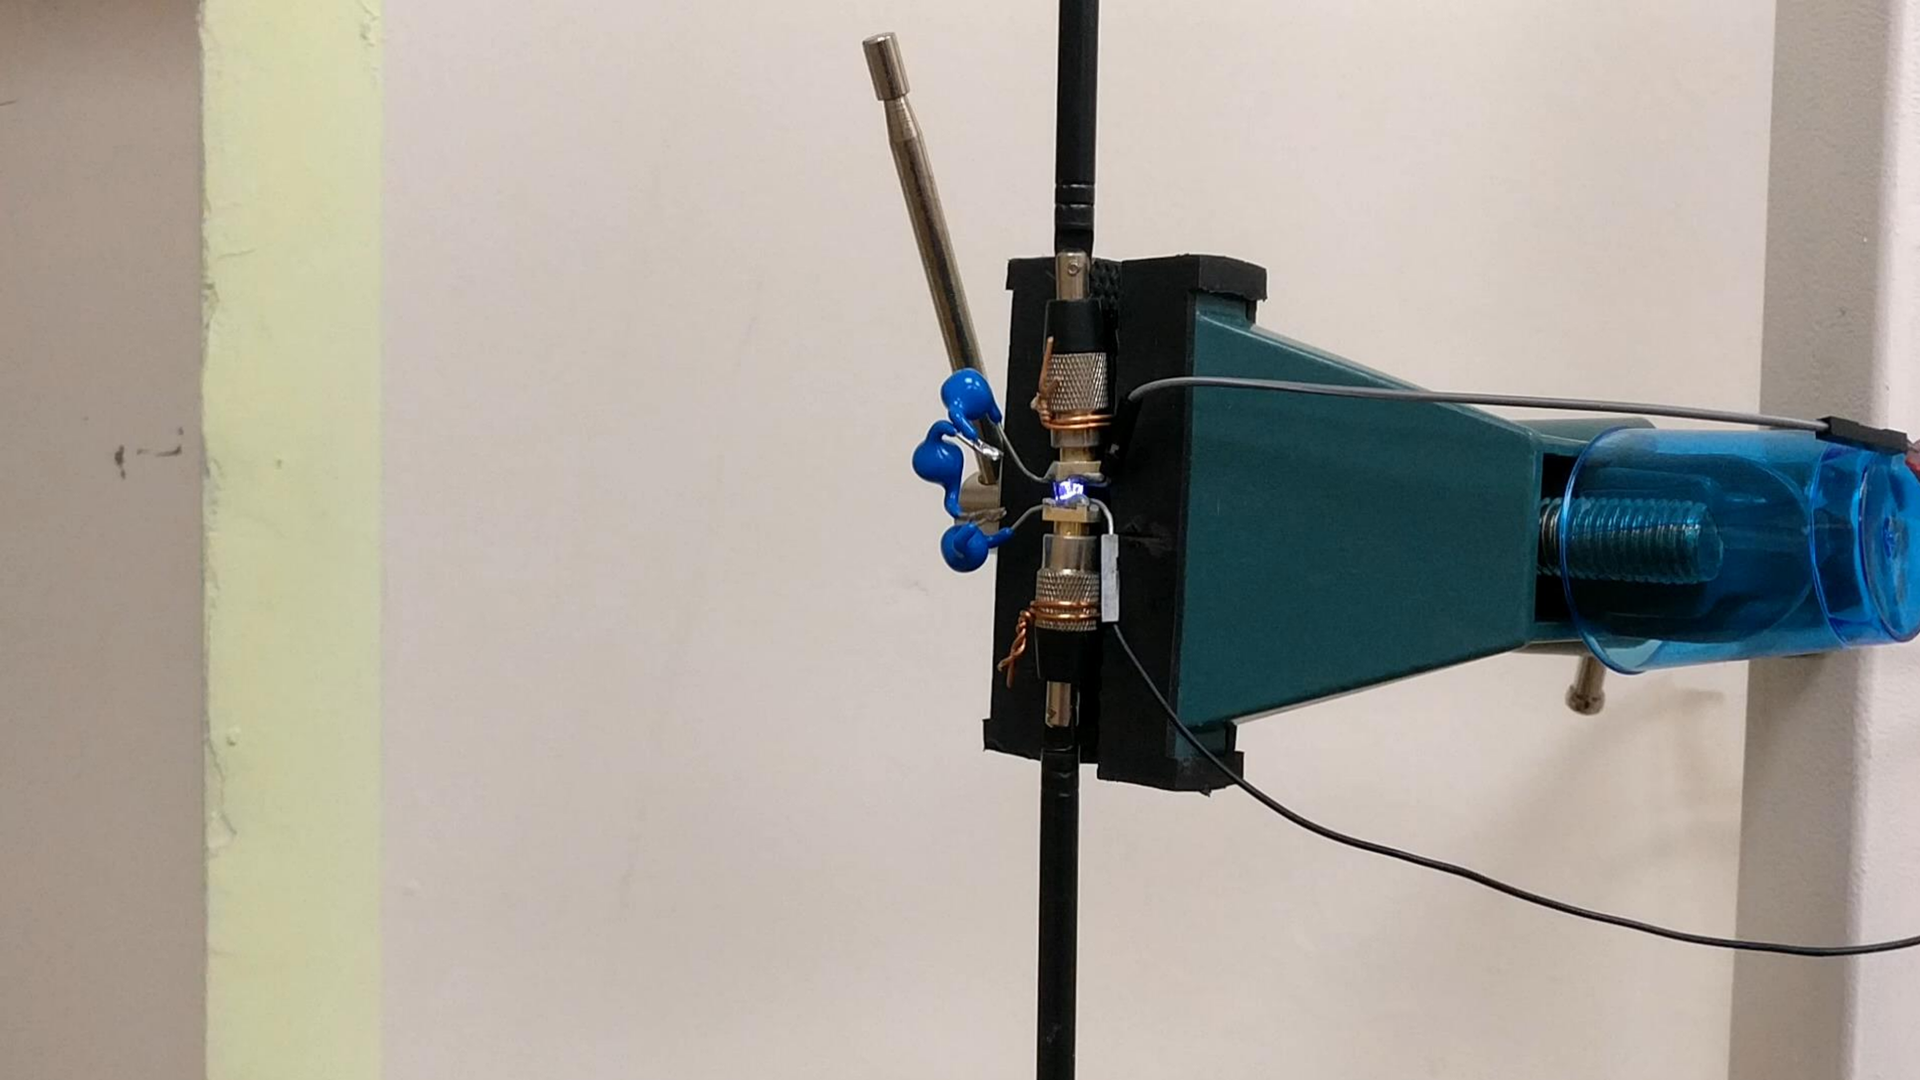


יציאה
יציאת
חרום
220V
armis
CISCO

Final notes

- Ethernet Packet-in-Packet attacks are complicated, but possible!

- Things to do about it:
  - Develop mitigations in network infrastructure
  - Monitor the condition of Ethernet cables in networks

- Further research is required:
  - Getting a better understanding of how EMI attacks can impact Ethernet cables
  - Defining the exact parameters and quality of Ethernet cables that are at risk

- More info: https://armis.com/EtherOops

ARMIS™