Alexey Petrenko

Detecting physical layer attacks on Ethernet networks

Helsinki Metropolia University of Applied Sciences
Bachelor of Engineering
Degree Programme in Information Technology
Thesis
8 October 2015

Helsinki Metropolia University of Applied Sciences

Author(s) Title Number of Pages Date	Alexey Petrenko Detecting physical layer attacks on Ethernet networks 58 pages + 3 appendices 8 October 2015
Degree	Bachelor of Engineering
Degree Programme	Degree Programme in Information Technology
Specialisation option	Networking Technology
Instructor(s)	Matti Puska, Principal Lecturer

The goal of this project was to find out if it is possible to select a set of metrics available from networking equipment which could be used to detect known physical layer attacks on Ethernet networks.

The project was carried out in a laboratory environment on a testing topology which was designed specifically for the testing purposes of this project.

Known physical layer attacks on Ethernet networks were described in detail. A set of metrics which might be used for attack detection was suggested. All metric values were gathered on each link in a topology in a normal state and under each of the attacks. Effectiveness of the suggested metrics was analysed.

The project showed that it is possible to use metrics obtained from networking devices to detect known physical layer attacks on Ethernet networks.

Keywords	network attacks, Ethernet, security

Contents

Abbreviations

1 Introduction 1

2 Theoretical background 2

2.1 General attacks classification 2
2.1.1 Confidentiality, Integrity, Availability 2
2.1.2 Disclosure, Alteration, Destruction 3
2.1.3 Attacks on Confidentiality 4
2.1.4 Attacks on Integrity 4
2.1.5 Attacks on Availability 5
2.1.6 Summary 6
2.2 OSI network model 6
2.3 TCP/IP network model 7
2.4 Ethernet networks 8
2.4.1 Ethernet on the OSI model 8
2.4.2 Media specifications 9
2.4.3 Power over Ethernet 11
2.5 Monitoring network traffic 12
2.5.1 Port mirroring 13
2.5.2 Taps 13
2.6 Physical layer attacks 14
2.6.1 Hub 15
2.6.2 Switch with port mirroring 16
2.6.3 Electrically active Ethernet tap 16
2.6.4 Electrically passive Ethernet tap 17
2.6.5 PC bridge man in the middle 19

3 Methods and materials 20

3.1 Selected attacks 20
3.2 Metrics 21
3.2.1 Link state 21
3.2.2 Link speed 22
3.2.3 Link duplex 23
3.2.4 Power over Ethernet 24
3.2.5 Observed cable length 24
3.2.6 Observed cable wiring 24
3.2.7 Error counters 24
3.2.8 Latency 25
3.2.9 Link quality 25
3.2.10 Throughput 25
3.3 Environment 26
3.3.1 Topology 26
3.3.2 Equipment 28
3.4 Gathering metrics 29
3.4.1 Link state as a passive metric 29
3.4.2 Link state as an active metric 30
3.4.3 Link speed and duplex as passive metrics 30
3.4.4 Link speed and duplex as active metrics 31
3.4.5 Power over Ethernet 32
3.4.6 Observed cable length and wiring 32
3.4.7 Error counters 32
3.4.8 Latency and link quality 33
3.4.9 Throughput 33

4 Results 35

4.1 Baseline 35
4.2 PC bridge man in the middle 37
4.3 Switch with port mirroring 40
4.4 Throwing Star LAN Tap 43
4.5 Electrically passive Ethernet tap with direct connection to the wires 45
4.6 Net Optics Fast Ethernet tap 47
4.7 Net Optics Gigabit tap 49
4.8 Combined results 52

5 Discussion 54

6 Conclusions 55

References 56

Appendices

Appendix 1 Detailed information about the SW1

Appendix 2 Detailed information about the SW2

Appendix 3 Configuration of the switch with port mirrorring

Abbreviations

CDP	Cisco Discovery Protocol
CPU	Central Processing Unit
DDoS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DoD	Department of Defence
DoS	Denial of Service
FCS	Frame Check Sequence
HTTP	Hypertext Transfer Protocol
HTTPS	HTTP Secure
ICMP	Internet Control Message Protocol
IOS	Internetwork Operating System
IP	Internet Protocol
LAN	Local Area Network
LLC	Logical Link Control
MAC	Media Access Control
MDI	Medium Dependent Interface
MDI-X	Medium Dependent Interface Crossover
MitM	Man in the Middle
NIC	Network Interface Controller
OSI model	Open Systems Interconnection model
PC	Personal Computer
PD	Powered Device
PSE	Power Sourcing Equipment
PoE	Power over Ethernet
RTT	Round-Trip Time
SFP	Small Form-factor Pluggable
SPAN	Catalyst Switched Port Analyzer
SSH	Secure Shell
SSL	Secure Socket Layer
STP	Shielded Twisted Pair
STP	Spanning Tree Protocol
TDR	Time-Domain Reflectometer
UTP	Unshielded Twisted Pair

1 Introduction

Nowadays security issues are becoming more and more important. The amount of sensitive data transfered over networks is growing every day. More people get affected by major breaches [1,5]. Backups with private photos are getting stolen [1,64]. Governments spy on one another and on their own citizens [2,3]. In this situation it is important to be aware of the weaknesses in the technologies that are used every day.

What are the consequences of someone inserting a tiny box between an unsuspecting user and the Internet and this box saves a copy of all traffic passing by, including photos, messages and other sensitive information? Or someone putting one of these boxes before a network-enabled printer/scanner and getting a copy of every scanned and printed document. Unfortunately, it is not a fantasy and such devices exist. It is easy to buy one, or build it [3]. What is worse: no antivirus or firewall on the end device can detect such attacks. During these attacks system settings are not changed, network settings are not changed, and visible network traffic is not changed.

This project aims to research known attacks on the networks and find ways to detect them. More specifically, physical layer attacks on the Ethernet networks are covered. The focus is on Ethernet networks, due to the fact that Ethernet is the most widely used network standard nowadays [4,6]. Physical layer attacks are selected because they are the hardest to detect. Unlike other network attacks they can function without introducing, or modifying any network traffic.

The main feature of this research is that no specific networking equipment is used to detect attacks. The idea is to try to detect attacks using only features, built in in a common networking equipment, such as switches, computers and access points. In this way the results of this research may be adopted by anyone and integrated into existing networks without having to add any devices to their existing infrastructure.

Different network equipment has different sets of features and it is impossible to test each device. Some of the methods used in this project may not be available on every networking device. In this case methods described here should be adopted accordingly.

Summarizing all of the above, this research aims to answer the following question: Is it possible to reliably detect all known physical layer attacks in Ethernet networks by observing information available to the network devices? Additionally, it will be interesting to see if there is a single metric which could be used to detect each physical layer attack in Ethernet networks.

2 Theoretical background

2.1 General attacks classification

In order to get a picture of what the main classes of attacks are, it is important to understand what the main security principles are, what security aims to protect and what security professionals care about. The Confidentiality, Integrity and Availability (CIA) triad, which is often considered to be the cornerstone of security, could be used to answer these questions [5,11].

2.1.1 Confidentiality, Integrity, Availability

The CIA triad or, simply, CIA is the most popular security model. It defines three main principles which have to be guaranteed if a system is to be considered secure: Confidentiality, Integrity and Availability [6,63].

Figure 1 shows how Confidentiality, Integrity and Availability form a secure system.

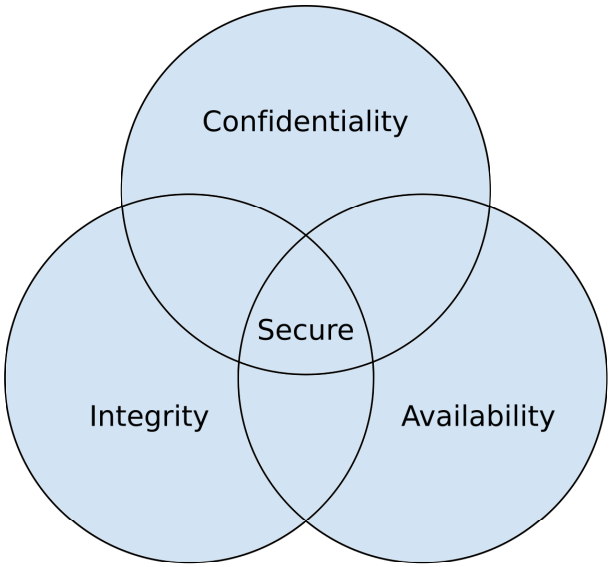

Figure 1: Confidentiality, Integrity and Availability relationship (Redrawn from [5,11]).

These three principles are described in detail as follows:

Confidentiality

Means that protected data should be accessible only by authorized people or systems [5,10].

Integrity

Means that protected data should not be subject to unauthorized or undesirable changes [7,6].

Availability

Means that protected data should be accessible whenever needed [7,6].

Security professionals agree that it is important to maintain a balance between Confidentiality, Integrity and Availability of a system, despite the fact that sometimes it may be challenging [5,10;8,12].

2.1.2 Disclosure, Alteration, Destruction

Some sources prefer to invert the CIA model and use it as DAD Model (see figure 2). DAD stands for Disclosure, Alteration, Destruction [8,11].

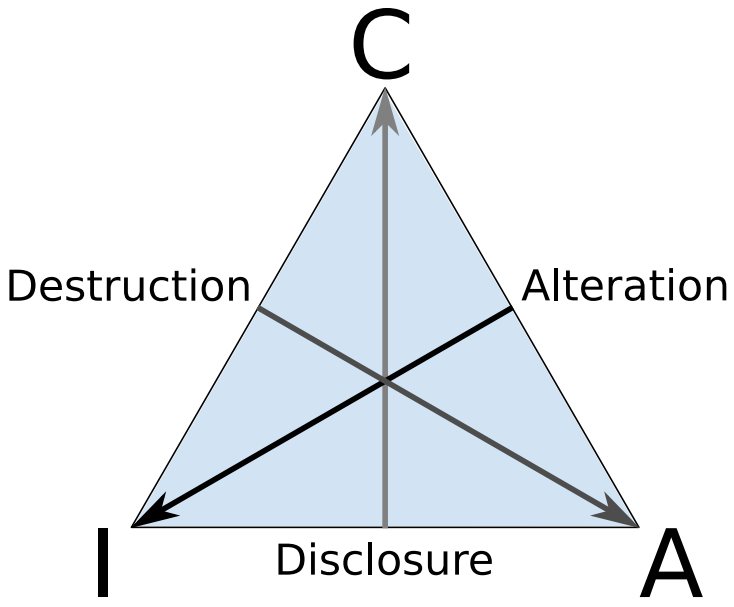

Figure 2: Disclosure, Alteration and Destruction (Redrawn from [8,11]).

As can be seen, DAD is CIA the other way around. Disclosure is opposite to Confidentiality, Alteration is opposite to Integrity and Destruction is opposite to Availability. The DAD model serves the same purpose as the CIA model, but it uses an opposite relation to the information. CIA determines principles which should be provided by a system, whereas DAD determines principles which should not be allowed by a system. [7,6.]

These key security principles lead to the main classes of the attacks. As illustrated in figure 3, it is possible to classify attacks based on what they are targeting. [9,178;10,187-194.]

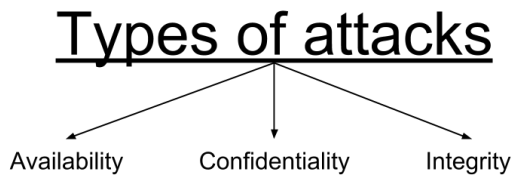

Figure 3: An overview of different attack types (Modified from [10,188]).

In this way attacks could be classified as attacks against Confidentiality, attacks against Integrity and attacks against Availability.

2.1.3 Attacks on Confidentiality

The most common attack on confidentiality in networks is eavesdropping. Sometimes it is also called sniffing. The essence of this attack is that the attacker who listens to network traffic could potentially gain access to the passwords, emails, personal documents or other sensitive information transfered over the network [11,541]. The amount of data available to the attacker depends on the position in the network and on the portion of traffic which is unencrypted.

Eavesdropping is not the only network attack against confidentiality. Man in the Middle attack, covered in section 2.1.4 is also often violates confidentiality of a communication. It allows the attacker to actively interfere with a traffic in order to gain access to the information which was meant to be secret.

2.1.4 Attacks on Integrity

An example of the attacks on integrity are Man in the Middle attacks or MitM. MitM attacks assume that an attacker gains access to a position in a network where all or part of the traffic between a victim and the target server or network passes through him [12,257]. An example of this attack is illustrated in figure 4.

Figure 4: An attacker performing MitM attack on a victim.

Remaining in this position an attacker can manipulate traffic in any desired way. For example, it is possible to inject malicious code to web pages visited by the victim, modify unencrypted emails send by the victim, or, simply, block access to certain resources.

Here, MitM attacks are described as attacks against integrity. However, most of the time they are attacks against the confidentiality as well. If the attacker circumvents encryption, or deals with plaintext traffic passing by, he has access to the information, which is transferred, and therefore it is an attack against confidentiality as well.

Having a possibility to modify the traffic, the attacker can sometimes trick the victim into an establishing unencrypted or significantly weakened connection to a server instead of a properly encrypted connection [13,173;14,277-278]. Alternatively, if there is no proper authentication in place, the attacker may trick the victim into an establishing encrypted connection to a different endpoint, which will then establish encrypted connection to a real server and acts as a proxy, forwarding the traffic between both ends [15,52]. If any of these attacks are successful, the attacker will be able to see or manipulate the traffic which was meant to be secret.

Another way to use MitM attack is to forward all traffic without any changes. In this case MitM will efficiently become eavesdropping [16,1-3].

It is important to understand that MitM attacks require a special position in a network. The attacker needs to have a possibility to receive traffic, modify it and send it on, while original not-modified traffic should be discarded. This adds some limitations to which network attacks can be used for MitM.

2.1.5 Attacks on Availability

The main class of the attacks on the availability is Denial of Service or DoS. The idea behind a DoS attack is to render a system or some part of it unusable for legitimate users [9,178].

DoS attacks can be divided into the following categories:

- exhaustion of resources
- crash or misbehaviour [9,178].

Exhaustion DoS attacks try to exceed the amount of data a host or a network can handle and render it unusable for legitimate users, whereas other DoS attacks rely on the system improperly handling certain type of data and crashing. A distinctive feature of the latter is that often human intervention is needed to bring the system back online. [9,178-179.]

There is also a variety of DoS attacks called Distributed Denial of Service or DDoS. Essentially they are the same as DoS attacks. The difference is that DDoS attacks are performed from multiple points at the same time while DoS attacks are performed from a

single point. Having an attack originating from multiple points at once makes it harder to mitigate it. [17,17;9,179-180;18,117.]

2.1.6 Summary

To summarize all of the above: DoS and DDoS are attacks on the Availability, eavesdropping is an attack on Confidentiality and MitM is an attack on Integrity and Confidentiality. After adding this information to figure 3, figure 5 was constructed. It shows examples of attacks against the main security principles.

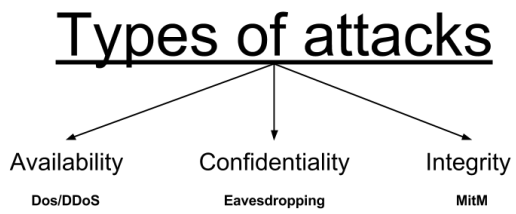

Figure 5: An overview of different attack types with examples (Modified from [10,188]).

This project focuses only on the following two attack classes: Eavesdropping and MitM. Attacks on the availability (DoS) are not covered, as the whole nature of these attacks makes them trivial to notice. Successful attacks on the availability of a network make the network or a part of it non-functional, which is really easy to detect. Unsuccessful attacks, on the other hand, do not cause network outages and it is harder to notice them. But if they are unsuccessful, there will be no reasons to look for them. In contrast to that, both Eavesdropping and MitM network attacks tend not to cause network outages and may remain unnoticed for very long periods of time.

2.2 OSI network model

The Open Systems Interconnection (OSI) model is a model designed for the classification of network protocols [19,18-19]. The OSI model includes seven layers, which are application, presentation, session, transport, network, data link, physical. Each layer is responsible for certain functions.

The functions of OSI layers from lowest to highest could be described as follows:

Physical layer (L1)

Defines functional electrical and mechanical control of data circuits physical media is connected to.

Data link layer (L2)

Standardizes communication between stations on a single link.

Network layer (L3)
Standardizes communication between stations across the network.
Transport layer (L4)
Is responsible for flow control and error-recovery.
Session layer (L5)
Is responsible for connections between applications.
Presentation layer (L6)
Is responsible for data representation for applications.
Application layer (L7)
Provides interface for end user applications.
[20,11-12]

Any network protocol could be placed on the OSI model. Some protocols are operating on a single layer, while others are operating on multiple layers at the same time.

The OSI model incorporates the peer communication concept. Communications within OSI model layers are independent of one another. Protocols interact with protocols on the same OSI level, without having to be aware of the protocols being used on the other layers. [21,6.]

2.3 TCP/IP network model

The OSI network model is not the only model used for network protocols classification. There is another model which could be used for this purpose. It is called Transmission Control Protocol / Internet Protocol (TCP/IP) model, or DoD Model. DoD abbreviation comes from the US Department of Defence, which was responsible for its development.

TCP/IP model has only four layers: application, transport, internet and a network interface layer. These layers can be mapped on the OSI model in the way shown in table 1. [19,24.]

Table 1: TCP/IP model mapped onto OSI model (Copied from [19,24]).

OSI model	**TCP/IP model**
Application layer Presentation layer Session layer	Application layer
Transport layer	Transport layer
Network layer	Internet layer
Data-Link layer Physical layer	Network interface layer

To cover the same functionality with a fewer number of layers, some of the TCP/IP model layers perform functions of several OSI model layers at once. As mentioned in chapter 1, attacks covered in this project are located on the OSI physical layer, or network interface layer of the TCP/IP model.

2.4 Ethernet networks

Ethernet is a set of standards describing Local Area Network (LAN) technologies. It resides on OSI physical (L1) and data link (L2) layers . Ethernet includes standards for different media and speeds. It defines media specifications, physical signaling, frame structure and addressing on L2. [4,11-13.]

2.4.1 Ethernet on the OSI model

As it is already said in the beginning of this section, Ethernet is located on the OSI layers L1 and L2. However, the Ethernet standard IEEE 802.3 defines more precise division to a sublayers, then the OSI model. There are four sublayers: Logical Link Control (LLC), Media Access Control (MAC), physical signaling and media specifications. These sublayers are mapped to the OSI model in the way shown in figure 6. [20,12.]

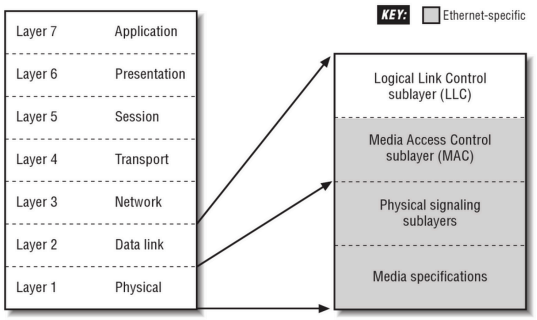

Figure 6: Ethernet sublayers (Copied from [20,13]).

The LLC sublayer is independent of Ethernet and standardized in a separate standard: IEEE LLC [20,12]. MAC layer is the same for all versions of Ethernet, whereas physical layer standards differ between Ethernet varieties. [20,13.] This research focuses on physical layer attacks on the networks and out of these four sublayers the main interest is in the Ethernet media specifications and physical signaling.

2.4.2 Media specifications

Ethernet networks can function over a wide variety of cable types: coaxial cables, Unshielded Twisted Pair (UTP) cables, Shielded Twisted Pair (STP) cables and fiber optic cables. There are multiple different Ethernet signaling standards for each of these media types. Different standards provide different speeds and have different requirements to a media. [22,44.]

For example, Gigabit Ethernet (1000BASE-T), Fast Ethernet (100BASE-TX) and a legacy 10BASE-T and all are able to function on the same category 5 cable, providing different speeds. At the same time newer and faster standards, such as 10 Gigabit Ethernet (10GBASE-T) would not work over the same category 5 cables and require at least category 6 cables. [23,703.]

Out of this available medias this project focuses on the twisted pair cables. Coaxial cables are not covered because they are outdated and not used in modern networks anymore. Fiber optical cables are not as widely used for horizontal cabling as twisted pair cables and they require expensive equipment in order to manipulate with them. Therefore, they are less likely to become a target for an attacker. That are the reasons behind a decision to focus on the twisted pair cables. Even though this project covers only Ethernet networks with twisted pair media, some of the methods described in it may be applied to other Ethernet media with minor modifications

Twisted pair cables are divided into several categories based on a TIA/EIA 568 standard [20,211]. TIA/EIA 568 is a vendor-independent standard for cabling specifications developed by Telecommunication Industries Associations (TIA) and Electronic Industries Association (EIA) [20,207].

Cable categories are numbered from one onwards, one being the oldest. The bare minimum in modern networks is to use category 5 cables. As mentioned before, category five cables are suitable for use with Ethernet speeds up to a Gigabit. 10BASE-T is the only standard which can work over cables with a lower category: category 3 and 4 cables.

UTP cables usually have four pairs of wires, but some of them may remain unused, depending on the used speed.

10BASE-T
: Uses two pairs of category 3 or better cables

100BASE-TX
: Uses two pairs of category 5 or better cables

1000BASE-T
: Uses four pairs of category 5 or better cables

10GBASE-T

Uses four pairs of category 6 or better cables

For convenience wires are colored in pairs: green and white-green, orange and white-orange, blue and white-blue, and brown and white-brown. Ethernet over twisted pair uses RJ-45 connectors and sockets to connect cables to a devices. The most widespread wiring schemes used for twisted pair cables are T568A and T568B [24,130]. Figure 7 shows wiring schemes both connector types.

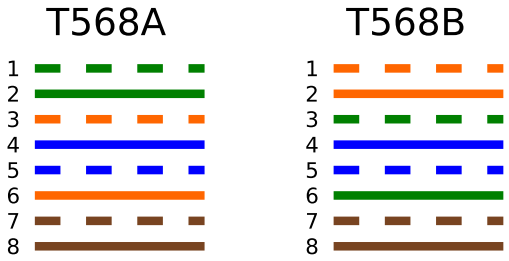

Figure 7: T568A and T568B connectors wiring.

As can be seen, T568B is essentially the T568A with white and orange pairs swapped [24,131].

There are two types of Ethernet interfaces on the network equipment: medium dependent interface (MDI) and medium dependent interface crossed (MDI-X) [25,110]. These interfaces differ in the choice wire pairs, which they use to transfer and receive data. On 100BASE-TX and 10BASE-T MDI interfaces use pins 1 and 2 to transmit and pins 3 and 4 to receive. MDI-X does use these pairs in the opposite way. This raises a question about interconnecting devices with the same or different interface types. Different cables are required to connect different interface types.

To connect MDI and MDI-X interfaces a straight-through cable should be used. If both ends of a cable use connectors with the same wiring, the cable is straight-through. Diagram of a straight-through cable is shown in figure 8.

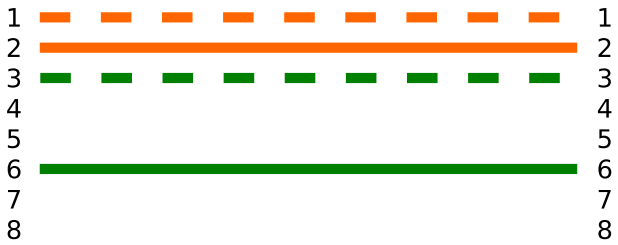

Figure 8: Wiring diagram of a four-wire straight-through cable.

Straight-through cables connect pin 1 on the one end to a pin 1 on the other end, pin 2 to the pin 2 and so on. In this way transmitting pins of MDI interface are connected to the receiving pins of a MDI-X interface and vice versa.

In order to connect two interfaces of the same type (MDI or MDI-X) a crossover cable is required. Crossover cable uses one T568B connector and one T568A connector. Figure 9 shows a diagram of a crossover cable.

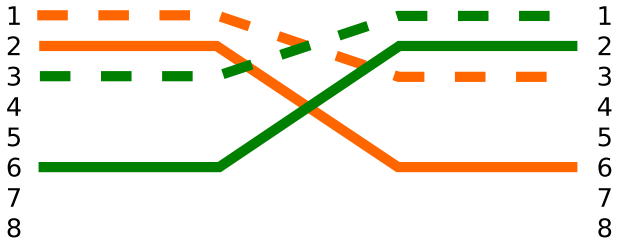

Figure 9: Wiring diagram of a four-wire crossover cable.

A crossover cable connects pins 1 and 2 on one end to the pins 3 and 4 on the other end and pins 3 and 4 on the first end to pins 1 and 2 on the second end. In this way transmitting pins of one interface are connected to a receiving pins of another interface and vice versa.

To eliminate the need of using different cable types for different interfaces, auto MDI-X was developed. It allows automatic detection of a cable type and selection of appropriate pairs for data transfer and reception. [4,280.]

Unlike 10BASE-T and 100BASE-TX, 1000BASE-T and newer standards use all four pairs to transmit and receive data at the same time. These standards all have auto MDI-X built in and enabled by default. [4,280.]

2.4.3 Power over Ethernet

It is important to mention Power over Ethernet (PoE) as it would be used in a testing topology. PoE is a technology of providing a power to an end device over an Ethernet twisted pair cable. It is usually used to power devices such as IP phones, security cameras or wireless access points.

The most widely used PoE is defied in IEEE standard 802.3af [4,89]. There is an extension to that standard - IEEE 802.3at, which allows devices to provide more power. This standard defines two possible device roles: Powered Device (PD) and Power Sourcing Equipment (PSE). PSE is a device which provides power. PSE may be provided by a networking device such as a switch, or by a separate power injector. PD is a device which is powered by PSE. [4,92.]

The standards define methods for detecting if a PoE capable device (PD) is connected to a port. When the PD is detected, PSE and PD negotiate amount of power, provided on a cable. The amount of power is determined by a class of a PD. [4,94-97.]

Table 2 shows available PD classes.

Table 2: PoE power classes (Copied from [4,96]).

Class	Power at PSE output (watts)	Description
0	15.4	Unclassified
1	4.0	Very low power
2	7.0	Low power
3	15.4	Mid power
4	36.0	High power

Class zero is the default class for devices which do not support power negotiation. In this case full power is provided to a device (15.4 watt). [4,95.]

PoE uses two pairs of wires to provide the power. One pair is used for negative current and another for positive. There are two alternatives to provide power over Ethernet: over data pairs or over spare pairs. The pairs used for the current are negotiated between PSE and PD. [4,98.]

There are multiple other versions and standards of PoE, some of which are vendor-specific [4,105-106]. They are not covered in this section.

2.5 Monitoring network traffic

This section is focused on the available methods and tools for network traffic monitoring, such as port mirroring and network taps. Even though these are meant for legitimate use by network engineers and system administrators to troubleshoot network-related problems, they could be used by malicious actors to perform eavesdropping. That is why these technologies are covered in this project.

There are software sniffers available for all major operating systems [26,35]. Sniffer software switches network card in a promiscuous mode and captures all network traffic which reaches the network card [26,18]. They have a wide variety of features, such as packet dissection, stream assembly, powerful filters and much more [26,74-76]. So, why are any additional methods or devices needed?

Software sniffers are good for analyzing the traffic seen on a cable. The problem arises when there is a need to monitor a traffic which does not originate from or is not meant for the PC running a sniffer. During the operation switches learn where the devices are located and use this information to forward non-broadcast traffic directly to the correct port. In this way, non-broadcast traffic is only forwarded to a port with the destination device, meaning that devices behind other ports are not able to see this traffic, regardless of the

mode their network card operates in [26,20]. Moreover, sometimes there is a need to monitor network traffic from a different LAN, which is separated from a sniffer by a router, and this traffic is not forwarded to a network with a sniffer at all. This is exactly the reason the monitoring tools and methods mentioned in the beginning of this section are used for.

2.5.1 Port mirroring

Port mirroring is a feature of a switch, which allows a switch to be configured so that it copies all incoming, outgoing or both directions of a traffic from a selected port to a designated monitoring port [26,35].

There are also variations of port mirroring, which allow to transfer traffic of a selected port across multiple switches and L3 networks respectively [27,405;27,408]. Depending on a vendor names and details of this feature may vary. For example, port mirroring is called Catalyst Switched Port Analyzer (SPAN) in Cisco terminology [27,400].

2.5.2 Taps

Network taps are separate network devices, designed to copy all the traffic which passes through it without modifying the original traffic in any way[26,24]. Network taps are installed in the middle of the existing network link.

Network taps have some extra features compared to switches with port mirroring functionality, such as:

Failsafe
: When a tap loses power, it will connect the wires of the cables directly so that the traffic continues to flow.

Link failure propagation
: If a link goes down on one port, tap will turn off a link on the other port as well.

Speed synchronization
: A tap negotiates the same speed and duplex on both ports.

PoE passthrough
: A tap allows power over Ethernet to pass through it.

Built-in battery
: When a tap loses power, it will continue to work on a battery for some time.

There is an opinion that Fast Ethernet taps and Gigabit Ethernet taps rely on different principles to be able to function. Particularly, Fast Ethernet taps use high impedance repeater equipment, while Gigabit Ethernet taps terminate an Ethernet link. [28]

2.6 Physical layer attacks

This section describes known attacks on Ethernet over twisted pair networks. The only exception are DoS attacks, which are not covered in this research. The reasoning behind that decision was explained in section 2.1.6.

Figure 10 shows an example segment of a network, which is used to illustrate possible attack points.

Figure 10: Example network segment.

The targets of attacks covered in this research are twisted pair cables between network devices, not the devices themselves. Figure 11 shows the target scope of possible attacks.

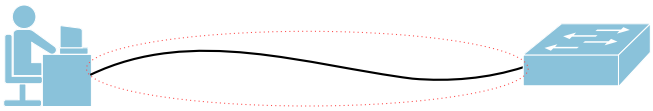

Figure 11: Target scope of the attacks.

Attacks can occur anywhere along the cable, or on either of the cable ends.

Some of the attacks rely on an attacker connecting additional network devices to a network, such as a network tap or a switch. These devices use standard RJ-45 sockets for connections. In this case, the attacker can unplug the cable from one of the end devices, connect this cable to his equipment and use extra network cable to connect another port of his device to the end device. This scenario is shown in figure 12.

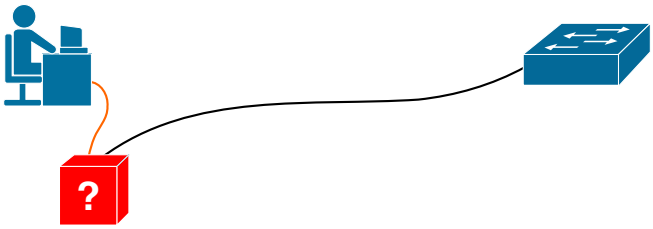

Figure 12: Example of an attack on the cable end.

Attacks described above are happening on the end of a cable. However, it does not always have to happen there. It is possible to perform such an attack even if the attacker has no access to any of the cable ends.

Twisted pair cables are made of relatively thin wires and isolation. And they can be easily cut with a knife or a scissors. Later it is possible to put new RJ-45 connectors on it with a help of inexpensive crimping tool. [4,271].

After reflecting this on the example network segment, figure 13 will be constructed.

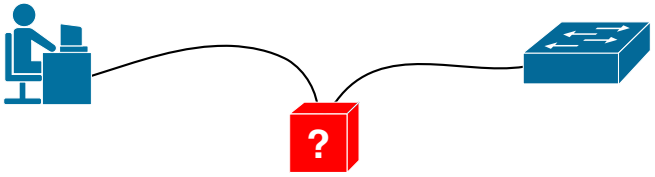

Figure 13: Example of an attack along the cable.

The process of cutting a cable and attaching new connectors to it, depending on the skills, will take a few minutes. If attack is performed out of work hours, link downtime may remain unnoticed.

2.6.1 Hub

Hub interconnects all devices connected to it, retransmitting every received message to the each port (except of the port the message was received on). In this way everyone, gets a copy of all the traffic transmitted over a hub. However, there is a limitation to it, due to a nature of a hub, it can carry only one signal at a time. Therefore, devices connected to a hub cannot receive and transmit at the same time. This mode of operation is called half-duplex. [29,57.]

Using any of the methods described above it is possible to perform eavesdropping attack by inserting a hub between the devices under attack. Figure 14 shows how a hub can be used for an attack.

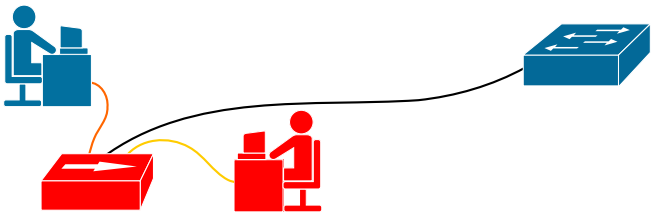

Figure 14: Example of an attack with a hub.

To perform an attack the attacker needs to connect attacked devices through a hub and connect a sniffer to the hub.

Besides the fact that hubs are outdated and do not support speeds faster than 100BASE-TX, they work only in half duplex mode [29,56]. Therefore, link under attack would degrade to a half duplex when hub is connected. Such behaviour could be easily detected. Due to these facts, no experiments with a hub would be performed.

2.6.2 Switch with port mirroring

The attacker may replace a hub with a switch which supports port mirroring to perform another variant of eavesdropping attack. This is reflected in figure 15.

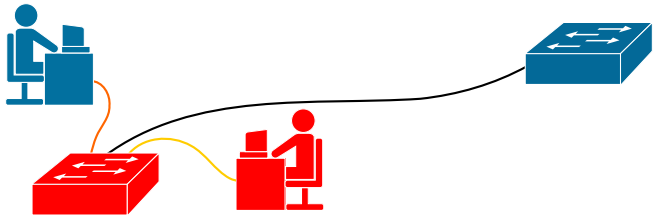

Figure 15: Example of an attack with a switch with port mirroring functionality.

Using a switch instead of a hub allows the attacker to stay less noticeable, as switches support full duplex mode, unlike hubs [29,65]. Therefore, they can establish connections in a full duplex mode, without downgrading a link duplex.

However, depending on speeds supported by devices under attack and a switch, the speed of a link may change. When both switch and a device connected to it support speed negotiation, they establish a link on a higher speed supported by both devices [4,67]. Depending on a speeds supported by a switch, the speed of a link between a switch and an end device may be lower or higher than the speed of a link before an attack was performed. If the maximum speed supported by end devices is different, they may negotiate links with different speeds, which may be used to detect an attack. Test results regarding these assumptions may be found in section 4.3.

2.6.3 Electrically active Ethernet tap

The attacker can use an Ethernet tap instead of a switch with port mirroring and achieve the same result. Figure 16 shows how a tap can be used for this attack.

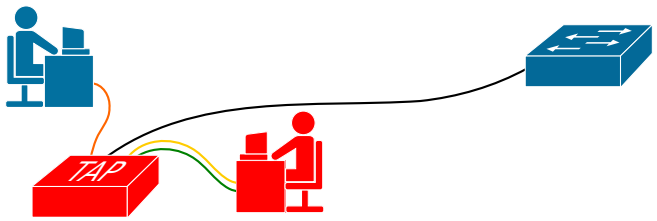

Figure 16: Example of an attack with an electrically active tap.

Additional features built in taps may make a tap less noticeable than a switch. For example, if device power goes down, a tap should reconnect wires directly, so that the link will continue working. Another example is that taps maintain the same speed and duplex on both ports. Switches do not usually have this functionality. [30]

2.6.4 Electrically passive Ethernet tap

Unlike the taps described before, the tap covered in this section does not require any external power source to function. That is why it is called electrically passive.

As mentioned in section 2.4.2, 100BASE-TX and 10BASE-T Ethernet standards use only two pairs of wires in a cable. One pair of wires is used to transfer the data in one direction and a second pair of wires is used to transfer the data in the opposite direction.

First Ethernet networks used a single shared coaxial cable to interconnect multiple network devices. To achieve this, straight-through bus topology was used. Even though Ethernet over twisted pair cables uses dedicated cable to interconnect each pair of network devices, 100BASE-TX and 10BASE-T standards are still able to work over a bus topology. In a fact, Ethernet hub incorporates a bus topology inside itself. So, in any network segment with a hub network devices share the same bus. [29,56.]

The fact that both 100BASE-TX and 10BASE-T are able to function over a shared medium, results in an interesting attack vector. The attack consist of an attacker connecting the listening pins of his Network Interface Controller (NIC) in parallel to a pair of wires in a cable, which is being attacked. After that the original link continues to work as before and the attacker gets a copy of a traffic flowing on this pair. If the attacker connects a second NIC in the same way to a second pair of wires, he gains access to a copy of the traffic flowing in another direction. [28]

Figure 17 shows an example of this attack. Monitoring cables for the attacker are marked TAP1 and TAP2.

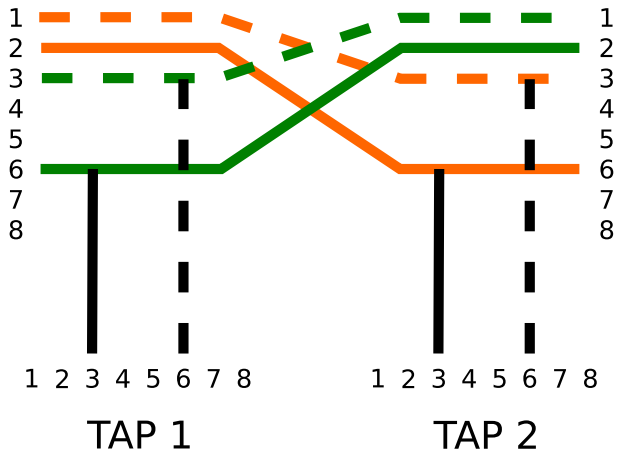

Figure 17: Wiring diagram of an electrically passive Ethernet tap.

This attack does not work on a faster Ethernet networks, such as 1000BASE-T and 10GBASE-T, as they use all four pairs of wire for data transfer and they were not designed to work over the bus.

There are different practical approaches towards performing this attack. It is possible to use a separate circuit with RJ-45 sockets, which connects monitoring interfaces in parallel to the wires of a cable. An example of this approach is a Throwing Star LAN Tap by Great Scott Gadgets [31]. It is shown in figure 18.

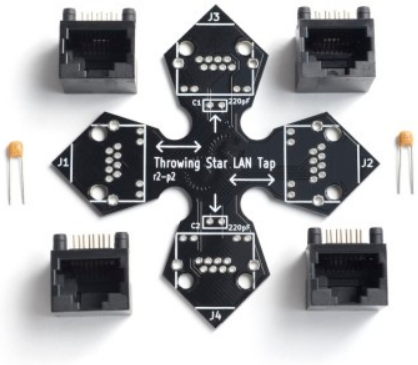

Figure 18: Throwing Star LAN Tap (Copied from [31]).

Using a separate device with a standard RJ-45 connectors leaves the cable intact. However, it requires the attacker to reconnect a cable, when performing an attack. This usually causes a downtime, which may be undesirable.

If downtime is highly undesirable, a second approach may be used. This approach involves connecting listening wires directly to a cable which is attacked. After isolation is removed, listening cables may be soldered, or connected in any other way to the transmitting wires.

One clever and easy way of connecting listening cables is to use a standard RJ-45 sockets. After that standard Ethernet cables are used to connect a tap to the listening network cards. Figure 19 shows a photo of such a tap.

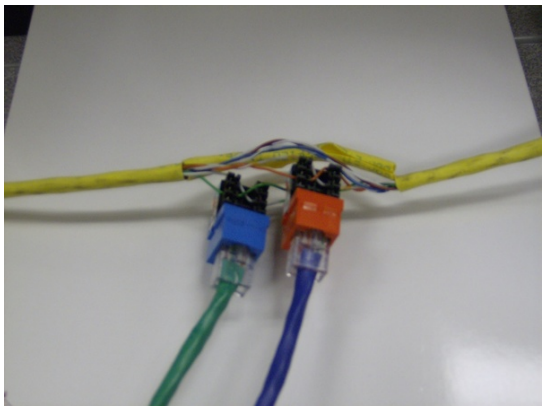

Figure 19: Ethernet tap example (Copied from [32]).

Even though this attack only works on 100BASE-TX and 10BASE-T networks, it is possible to perform it on faster networks as well. If the attacker cuts open the remaining wires in a cable, devices would not be able to negotiate any speed which requires all four

pairs and they will fall back to a Fast Ethernet connection. [4,79] Throwing Star LAN Tap has capacitors connected in line with remaining wires to achieve the same effect. These capacitors degrade a link quality, so that 1000BASE-T does not work on this link. [31]

2.6.5 PC bridge man in the middle

During the MitM attack network traffic flow has to be modified so that traffic will go through an attacker's computer or appliance. This project covers physical layer attacks, so only physical ways of redirecting the traffic would be considered. All the victim's traffic should go through the attacker's PC and only through it, so that an unmodified version of traffic would not be visible in the network. This condition leaves the only option: an existing Ethernet link should be terminated at the attacker's PC and a new Ethernet link should go from the attacker's PC further.

Figure 20: An example of a PC bridge MitM attack.

The attacker has to have at least two NICs in his PC, so that he/she can connect to both Ethernet cables at the same time.

3 Methods and materials

This chapter provides a list of attacks which are selected for testing, followed by a set of metrics, which are available from network devices and could be used to detect physical layer attacks. A general description of each metric is provided as well as technical details on how these metrics are extracted in a test environment. In the end the testing environment is described.

3.1 Selected attacks

To have a wide coverage of various attacks, the following attacks are selected for testing:

PC bridge MitM
: Attack is described in section 2.6.5

Switch with port mirroring
: Attack is described in section 2.6.2

Electrically passive tap (Throwing Star LAN Tap)
: Attack is described in section 2.6.4

Electrically passive tap with direct connection to the wires
: Attack is described in section 2.6.4

Electrically active 100BASE-TX capable tap
: Attack is described in section 2.6.3

Electrically active 1000BASE-T capable tap
: Attack is described in section 2.6.3.

As one may notice, one attack, which was covered in section 2.6, is not included in this list. This is an attack with an Ethernet hub. Section 2.6.1 provides an attack description, as well as the reasoning why it is not tested in this project.

Another interesting fact about the list of selected attacks is that there are two different attacks involving electrically active taps. The first attack uses Net Optics TP-CU tap, which does not support Gigabit Ethernet, and the second uses Net Optics TP-CU3-ZD tap, which supports Gigabit Ethernet [33;30]. The reason for testing attacks with two different taps is that, Gigabit Ethernet capable taps are believed to be based on different principles then Fast Ethernet taps, as mentioned in section 2.5.2. The difference in the way these taps work may make one of them easier to detect than another.

Generally speaking, an attack is considered successful, if after performing the setup phase, the link is functional and the end device is reachable. Short downtimes during the attack setup phase are tolerated. Moreover, they could be used as an indicator of attacks as well. However, if the end device is not reachable after the attack setup, then the attack is considered unsuccessful and no measurements are performed on that link.

3.2 Metrics

As mentioned in chapter 1, this projects aims to detect physical layer attacks on Ethernet networks, using only the commonly available networking equipment, which is used in a network infrastructures. In other words, most of the attack detection methods described here should be applicable to the arbitrary network without a need for any additional equipment.

This section focuses on determining the set of metrics which are available on the standard networking devices and could be used to detect the selected attacks. Here, metric is a property of a connection between two network devices. Some of the metrics will change their value during attacks, which would be an indicator of compromise.

It is important to mention that even though the goal of a research is to detect attacks anywhere in a network, actual testing is done on each link separately. A network is treated as a composition of individual links between devices. If there is a way to detect an attack on a single link, then it is just a matter of scaling up the same technique in order to detect the same attack anywhere in a network. Later on all of the discussed metrics are applicable to a single link only.

Measurements are carried out to gain metric values. Measurements could be divided into two groups: passive and active. A measurement would be called passive, if it does not change any network parameters of a link. In contrast, a measurement would be called active, if it introduces any changes to a link.

Following subsections cover in detail selected metrics, which potentially may be used for detecting physical layer network attacks.

3.2.1 Link state

A link state shows whether the link is up and functioning or down and non functioning. The terms on and off could be used in the same meaning. Although it is possible to have links in a topology which are constantly disabled, we would not be focusing on them, as

they usually pose no interest to attackers. Thus, an expected default state of the link is up.

Link state as a passive metric

A link state can be used as a passive metric. In this case it is enough to observe the link state over time and detect changes in it. Normal behaviour of the link state depends on an environment. Links to the end devices could change their state as end device is restarted, powered off or disconnected, however links between network devices are usually expected to be constantly up.

It is expected that, if the link state goes down and back up again, it can be an indicator of attack. It may happen when Ethernet cable is disconnected from one device and connected to another, or cable is cut, crimped and connected to a networking device. However, links could go down for various non-malicious reasons, such as problems with a cable, or device on the other end being powered off or restarted. This metric is expected to be prone to a false positives.

Link state as an active metric

It may be possible to use a link state as an active metric. During active measurements it is necessary to change the link state on one end of a cable and check a link state on the other end of a cable. It is expected that the link state of another end of a cable goes down and then up during. If the link state on the other end of a cable does not match the link state on the first end of a cable, it may be an indicator that the devices are not connected directly to one another. If there is no way to check a link state on the other end of a link during a test, indirect indicators could be used, such as logs on a device or Dynamic Host Configuration Protocol (DHCP) renewal request when the link goes up again.

3.2.2 Link speed

Link speed as a passive metric

Passively observing a link speed over time and detecting any changes in it should be possibly to detect changes in a network. It is expected that the link speed may downgrade during some of the attacks. For example, electrically passive Ethernet taps do not support speeds higher than 100 Mb/s [31;28]. A link speed may fluctuate for various non-malicious

reasons. However, in a well-maintained environment a link speed fluctuation should be investigated, as it could be an indicator of other problems, such as physical media or hardware problems. [20,93;20,359]

Link speed as an active metric

A link speed may be used as an active metric as well. By changing a link speed on one end of a cable and checking a link speed on the other end of a cable, it may be possible to detect if devices are not connected directly to each other. It is expected that during the normal conditions, depending on whether or not the link speed auto negotiation is enabled on a device on the other end of a cable, the link would be renegotiated on the same speed as set on the first device, or it will stop working. If the link speed on the other end does not change during a test, it may be an indicator of an attack.

3.2.3 Link duplex

Link duplex as a passive metric

Observing a link duplex values over time and detecting changes in it, may be used as a metric as well. It is expected that during the normal operation the link duplex stays the same all the time. If the link duplex changes its value, it may be an indicator of an attack. A link duplex may fluctuate for various non-malicious reasons. However, in a well-maintained environment a link duplex fluctuation should be investigated, as it could be an indicator of other problems, such as physical media or hardware problems.

Link duplex as an active metric

To use a link duplex as an active metric it is necessary to change the link duplex on one end of a cable and then check a link duplex on the other end of a cable. It is expected that depending on whether or not the link duplex auto negotiation is enabled on a device on the other end of a cable, the link would be renegotiated on the same duplex as set on the first device, or it will stop working. If the link duplex does not change during a test, it is an indicator that devices are not directly connected to each other.

3.2.4 Power over Ethernet

Switch ports capable of being a PSE, know if device on the other end is a PD and for each PD they know it's class and a requested power [4,92-97]. These values can be used as a metrics as well. It is expected that they stay the same all the time, as PD do not change its properties over time.

3.2.5 Observed cable length

Some of the network devices support Time Domain Reflectometry (TDR) on their network ports. TDR allows a device to measure a cable length with certain accuracy and detect cable faults, if any. [34;35,160.]

It is expected that the observed cable length stays close to the same value over the time, as well as the status of the cable pairs stays the same. Some fluctuations in a cable length are expected, as tests are performed on a live link and some variations in results may happen. However, if observed cable length changes dramatically in either direction or cable-fault indicator appears, it is an indicator that something on a link have changed. For example, cable was replaced with a cable of a different length, or extra equipment was connected to a cable.

This metric would react to a non-malicious cable changes as well. For example, cable replaced due to a network improvement may cause change in the metric value. However, if company implements proper change control, authorized cable change should be reflected in the documentation [36,725].

3.2.6 Observed cable wiring

Some of the cable-testing functions provide information about the cable wiring [34]. This information also could be observed. If cable wiring suddenly changes, it is an indicator that the cable was replaced with a cable of a different type, or tampered with, which requires further investigation.

3.2.7 Error counters

Network devices usually provide multiple counters on per-interface basis. Some of these counters could be used in conjunction with specially generated traffic to form active metrics. For example, one device generates multiple frames with incorrect Frame Check Se-

quence (FCS), jumbo frames or runt frames. Normal behaviour for a switch is to discard such frames and increment corresponding error counters. It is expected that the corresponding error counters would increase on the device on the other end of a cable [37,23]. If it does not happen, it is an indicator that the device have not received these frames. Most likely the reason for such behaviour is that devices are not connected directly and there is a device between, which filtered these frames.

3.2.8 Latency

Latency shows how long it takes for a frame to reach the other device. There is no direct metric on a switch corresponding to a link latency. However, if switch has an IP address, it may be possible to use Internet Control Message Protocol (ICMP) ping, or other protocols to get Round-Trip Time (RTT) to the device on the other end of a cable. [38,6.]

It is expected that the latency, being the characteristic of a link, will remain approximately the same all the time. However, this test utilizes Central Processing Unit (CPU) resources, therefore its results could be inaccurate depending on the CPU load [39,245]. Another factor which may affect the test results is the link congestion, caused by a non test network traffic passing the link during the test. This traffic may interfere with the test traffic and cause higher latency values [38,6]. But, if latency increases significantly, it may be a sign that a slow device is inserted in between the devices.

3.2.9 Link quality

ICMP ping test described in the previous chapter provides one more valuable output: amount of lost packets. It can be used as a metric as well. This counter shows how many ICMP packets were lost during the test. This value closely correlate with a link quality and may be influenced by an additional equipment connected during attacks.

3.2.10 Throughput

Throughput is a characteristic of a network, which shows how much data per second can be transfered over it. The same way as with latency, there are no direct metric available from a switch which shows throughput of a link. Active tests are required to measure throughput values.

Common way of measuring throughput is to transfer as much data as possible through a tested network over a given period of time. There may be difficulties in such measurement,

due to a normal non-test network traffic taking a part of available throughput, which lowers results showed by a throughput test. Another important aspect is that such measurements may be harmful for production networks, as they temporary place a link under high load. Maybe such tests should be performed only out of business hours. It is expected that the throughput of a link remains the same over time If it drops significantly, it should be investigated, as some other device may have been added to a network, or other link parameters may have degraded.

3.3 Environment

Designing a test topology for this project is a challenging task, as there are some controversial requirements. For the quality of results it is beneficial to have multitude of links in a system, each with a different combination of parameters and capabilities, that would provide granular information about the way each metric behave depending on the link properties. However, each link should be thoroughly tested in a normal conditions and during each of the attacks. That leads to a huge number of measurements. For example, having 6 attacks and 10 metrics would require 140 measurements per link in a topology. That is one measurement per metric for each end of a link for each attack plus baseline measurements.

3.3.1 Topology

The topology used for tests is a compromise between these two requirements. There is a small number of links in total with wide diversity of a link parameters and features. This topology is shown in figure 21.

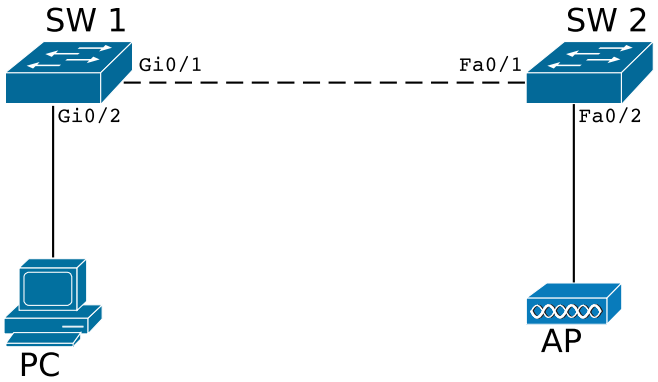

Figure 21: Topology used for experiment.

As one can see there are only three links in a topology, which keeps the overall amount of required measurements small enough. At the same time topology provides desired diversity by including connections with different speeds, cable types, cable lengths and

features.

For people familiar with network designing process this topology may seem absurd, as it does not have logical structure and it uses two switches, where one would have been enough, also it has a 1000BASE-T capable port connected to a port, which supports only 100BASE-TX. Please, keep in mind that this topology is designed for testing and the main focus is to have a wide variety of link parameters.

As mentioned in section 3.1, six attacks are selected for testing. section 2.6 provided attacks description. As can be seen, five out of six attacks use standard RJ-45 connectors and sockets for connecting additional equipment to a cable. The remaining attack is performed along the cable and does not require an Ethernet connector.

Even though, the first five attacks require RJ-45 connector, in order to connect attack equipment, they do not necessary have to happen at the end of a cable. It is possible to perform such attacks anywhere along the cable, as well. Ethernet UTP and STP cables could be easily crimped. Which makes it possible for an attacker to cut the cable at any point, crimp both ends and use them to connect his equipment.

When attack is happening on the either end of a cable, extra cable is required to connect the attacker's equipment to the existing network equipment. Cable used for that may be straight through or crossover. When attack happens along the cable, cable is cut and crimped. In a process of crimping different wiring schemes could be used for each of the connectors (T568A or T568B). All of the above means that for each of these five attacks additional network equipment can be connected to a link in eight different ways:

1. One end of a cable
    a) Straight through cable
    b) Crossover cable
2. Somewhere along the cable
    a) T568A and T568A
    b) T568A and T568B
    c) T568B and T568A
    d) T568B and T568B
3. Another end of a cable
    a) Straight through cable
    b) Crossover cable

Thus, in order to cover all possible ways of performing these attacks, it is necessary to perform each of these attacks eight times on the each link. Which makes an overall amount of measurements required for this project too big. To make the amount of measurements manageable, such attacks would be tested only once per link. The attack point and a connection cable would be defined for each link separately.

The only attack which does not involve connecting to a cable with Ethernet connectors is electrically passive Ethernet tap with direct connection to the wires, which was described in section 2.6.4. It relies on stripping the cable isolation and connecting listening wires directly to the transmitting and receiving pairs of the cable. This attack damages a cable, therefore it is unacceptable to use this cable for any other tests after this attack is tested. In order to keep the topology identical for all the attacks and main network cables intact, this attack would be performed on a short cables connected directly to the end of each topology link using a coupler. The topology updated with consideration of the short cables is shown in figure 22. Short cables are marked in red.

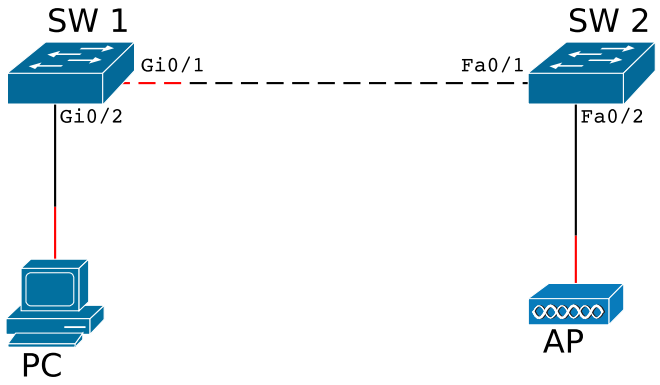

Figure 22: Topology used for experiment with actual cables displayed.

These short cables will remain in the topology constantly, so that it is the same for all the attacks. Cables damaged during the attacks would be replaced with a similar new ones after the attack is over and before the next one. The same cables would be used to connect the attacker's network equipment for other attacks. The attacker's equipment would be connected instead of couplers. Attack points for the electrically passive Ethernet tap with direct connection to the wires are in the middle of short cables.

3.3.2 Equipment

Topology includes two Cisco switches with and without PoE support, as well as full or limited TDR support, PC and an wireless access point. These network devices were selected with a goal of using standard common network equipment with no special features.

SW1 is a Cisco Catalyst 2960 switch. It has 24 Fast Ethernet ports and two Gigabit ports. TDR is supported on all the ports. Fast Ethernet ports only provide information about two pairs of wires. Gigabit ports provide information about all four pairs of wires. PoE is not supported by this switch. [40,26.41,19] SW2 is a Cisco Catalyst 3560v2 switch. It has 24 Fast Ethernet ports and two Small Form-factor Pluggable (SFP) slots. TDR is supported on the Fast Ethernet ports. Fast Ethernet ports only provide information about two pairs of wires. PoE is supported by all Fast Ethernet ports on this switch. [42,4;42,17.] Detailed information about switches SW1 and SW2 is available in appendices 1 and 2 respectfully.

Access point is UniFi UAP-PRO. It is capable of being a PD and receiving a power over the Ethernet cable. [43,6.]

3.4 Gathering metrics

In order to have precise and reliable results, it is important to have information about the processes happening in the network devices at all the time while performing experiments. There are multiple ways to manage most of the devices. All testing devices are in the same location, therefore standard management methods can be used. Switches have serial console connection for management, PC could be managed via keyboard and mouse. Access point do not have a separate management interface and has to be managed via network.

Following sections provide technical details on how to obtain metrics results from a Cisco switch and a Linux PC. UniFi UAP-PRO access point is Linux based and has Secure Shell (SSH) access. Therefore, Linux PC commands should work. However, if there are any differences, it would be separately noted.

3.4.1 Link state as a passive metric

Commands illustrated in listing 1 and listing 2 could be used to obtain current link state on a Cisco switch and a Linux PC respectfully.

```
SW1#show ip interface brief Gi0/1
Interface              IP-Address  OK?  Method  Status  Protocol
GigabitEthernet0/1     unassigned  YES  unset   up      up
```

Listing 1: Example output of `show ip interface brief` command on a switch (SW1).

Cisco Internetwork Operating System (IOS) `show ip interface brief` command provides an information about the interface status and IP address assigned it [44,66].

```
# ip link show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP
    mode DEFAULT group default qlen 1000
     link/ether 78:e7:d1:86:a1:c9 brd ff:ff:ff:ff:ff:ff
```

Listing 2: Example output of `ip link show` command on a Linux PC (PC).

`ip link show` command is a command from the Linux iproute2 toolset, which provides detailed information about a link and its parameters [45,203].

3.4.2 Link state as an active metric

As mentioned in section 3.2.1, an active link state metric assumes that the link state is changed on one device and observed on the second device. Above it was described how to check a link state on a Cisco switch and on a Linux PC. To disable an interface on a Cisco switch `shutdown` command in an interface configuration section may be used. Iproute2 `ip link set dev eth0 down` command disables an interface on a Linux system. To bring interfaces back Cisco `no shutdown` command and Linux `ip link set dev eth0 up` commands are used.

For some unknown reason neither `ip link set dev eth0 down` nor `ip link set dev eth0 up` works on the AP. Corresponding `ifconfig` commands do not work as well. Therefore, active link speed tests were not performed from the AP.

3.4.3 Link speed and duplex as passive metrics

Commands to get current speed and duplex of an interface on a Cisco switch and on a Linux PC are provided in listing 3 and listing 4 respectfully.

```
SW1#show interface Gi 0/1 status
Port   Name   Status      Vlan Duplex   Speed   Type
Gi0/1         connected   1    a-full   a-100   10/100/1000BaseTX
```

Listing 3: Example output of `show interface status` command on a switch (SW1).

Cisco IOS `show interface status` command provides an information about the interface status, speed and duplex [36,251].

```
# ethtool eth0
Settings for eth0:
    Supported ports: [ TP ]
    Supported link modes:   10baseT/Half 10baseT/Full
                            100baseT/Half 100baseT/Full
                            1000baseT/Half 1000baseT/Full
    Supported pause frame use: No
    Supports auto-negotiation: Yes
    Advertised link modes:  10baseT/Half 10baseT/Full
                            100baseT/Half 100baseT/Full
                            1000baseT/Half 1000baseT/Full
    Advertised pause frame use: Symmetric
    Advertised auto-negotiation: Yes
    Link partner advertised link modes:  10baseT/Half 10baseT/Full
                                         100baseT/Half 100baseT/Full
                                         1000baseT/Half 1000baseT/Full
    Link partner advertised pause frame use: No
    Link partner advertised auto-negotiation: Yes
```

```
Speed: 1000Mb/s
Duplex: Full
Port: Twisted Pair
PHYAD: 1
Transceiver: internal
Auto-negotiation: on
MDI-X: on
Supports Wake-on: g
Wake-on: g
Current message level: 0x000000ff (255)
               drv probe link timer ifdown ifup rx_err tx_err
Link detected: yes
```

Listing 4: Example output of `ethtool` command on a Linux PC (PC).

Linux `ethtool` command provides detailed information about a speed and duplex modes of an interface [45,205]. Alternatively `mii-tool` or log messages may be used to determine a link speed and duplex on a Linux system [46,302-303].

Example of a log messages produced on a PC when the link is established is shown in listing 5.

```
eth0: Link is up at 1000 Mbps, full duplex
```

Listing 5: Example of a log message generated on a Linux system when a network connection is established.

Unfortunately, neither of the methods described above work on the AP. It has no `ethtool` or `mii-tool` binaries installed and no log entries are generated on a link status change. Therefore, no active link speed or duplex tests would be performed to or from it.

3.4.4 Link speed and duplex as active metrics

Following the steps described in section 3.2.2 and in section 3.2.3, link speed and duplex should be changed on one device and observed on the other device to form active link speed and active link duplex metrics. Previous section showed how link speed and duplex values may be obtained on a Cisco switch and on a Linux PC. To change speed of a link on a Cisco switch `speed` command in an interface configuration section is used. Duplex of a link is changed with a `duplex` command in the same configuration section. Linux uses `ethtool` or `mii-tool` commands to achieve the same result.

3.4.5 Power over Ethernet

To obtain power over Ethernet related information from a Cisco switch `show power inline` command is used. Example output of this command is shown in listing 6.

```
SW2#show power inline Fa 0/2
Interface Admin  Oper       Power   Device              Class Max
                            (Watts)
--------- ------ ---------- ------- ------------------- ----- ----
Fa0/2     auto   on         15.4    Ieee PD             4     15.4
```

Listing 6: Output of a `show power inline` command for a port (Fa0/2) of a switch (SW2).

As can be seen, output of a command includes PoE operational mode, provided power, type and a class of a PD connected to a port, which may all be used as metrics.

3.4.6 Observed cable length and wiring

Cisco IOS `test cable-diagnostic tdr` and `show cable-diagnostic tdr` commands are used to perform a TDR cable test and to display its results. Example output of such test is shown in listing 7.

```
SW1#test cable-diagnostics tdr interface Gi 0/2
...
SW1#show cable-diagnostics tdr interface Gi 0/2
Interface Speed Local pair Pair length        Remote pair Pair status
--------- ----- ---------- ------------------ ----------- -----------
Gi0/2     1000M Pair A     8    +/- 10 meters Pair B      Normal
                Pair B     8    +/- 10 meters Pair A      Normal
                Pair C     8    +/- 10 meters Pair D      Normal
                Pair D     8    +/- 10 meters Pair C      Normal
```

Listing 7: Example of cable-diagnostic command output

This output includes per cable pair information about the corresponding remote pair, pair length and pair status. These values are used as metrics as well.

3.4.7 Error counters

As described in section 3.2.7, the original idea was to use available network card error counters in combination with the bad traffic sent over the wire, but no cross platform and uniform way to generate such traffic was found. Therefore, this test is not implemented.

3.4.8 Latency and link quality

ICMP RTT and lost packets counter are shown after the corresponding ping test is over. A shortened example of such a test is shown in listing 8.

```
SW1#ping 10.0.0.3 repeat 1000
Type escape sequence to abort.
Sending 1000, 100-byte ICMP Echos to 10.0.0.3, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
...
!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (1000/1000), roundp min/avg/max = 1/2/9 ms
```

Listing 8: Results of ping test performed from a switch (SW1) to a PC (PC).

Test output includes minimum, average and maximum RTT times as well as the amount of requests sent and responses received. In this example 1000 packets was sent and 1000 was received, meaning that an ICMP packets lost counter equals to zero.

The test is performed over a big number of packets and an average RTT value is used as a metric to compensate for the possible variations in response times due to an uncontrolled network events

3.4.9 Throughput

The first idea was to use `ttcp` tool for throughput testing. It is truly cross platform: It is included in Cisco IOS on switches and routers and there are binaries for Windows and Linux system.

However, the test `ttcp` run on a link between switches SW1 and SW2 produced results showed in listing 9.

```
ttcp-t: 67108864 bytes in 33051 ms (33.051 real seconds) (~1982 kB/s)
```

Listing 9: Switch (SW2) `ttcp` speed test results for a link between switches (SW1) and (SW2).

Data transfer of 67108864 bytes in 33051 milliseconds corresponds to a 2030463.9 bytes per second, or 1.95 Mb/s. Result provided by ttcp is a lot lower than the theoretical 100 Mb/s throughput of a link.

An observation was made that switch CPU load increases dramatically during the test. Listing 10 provides CPU load statistics for a switch SW2, which was measured right after the ttcp test finished.

```
CPU utilization for five seconds: 96%/5%; one minute: 42%; five minutes
    : 15%
```

Listing 10: CPU utilization statistic for a switch (SW2).

As can be seen, CPU load reached very high values during a test. Most likely ttcp test results are low due to a fact that CPU cannot process bigger amounts of data and not due to a limited link throughput.

To confirm this assumption a different speed measurement tool, called `iperf`, was used to measure bandwidth of the same link. Results showed by `iperf` over the same link are multitude times higher than the results provided by `ttcp`.

Unfortunately, `iperf` runs only on PCs and cannot run on a switches. Therefore, it cannot be used to test throughput of each individual link separately. For testing purposes it is possible to connect a separate PC to each device in the topology to perform throughput measurements with `iperf`. But, it would go against initial goal of not using any additional hardware. Therefore, throughput tests would not be performed on this topology.

4 Results

This chapter provides results of the experiments. First, the baseline results are provided. Next, the results of the measurements for each link under each of the attacks are provided. In the end, the results are combined and analyzed.

4.1 Baseline

This section provides a results of the measurements in a topology without any attacks. Later on these results would be used as a baseline for a comparison with a results received from measurements when the network is under each of the attacks. For each Ethernet port in a topology the results were extracted according to the methods described in section 3.4.

During these measurements it was found out that even though AP is running Linux, it does not have necessary tools or log messages to check the link parameters. Therefore, link speed and duplex metrics are not available for the AP, whereas ink state is deduced from the fact that the AP was accessible via network, therefore it had to be up. This means that active link tests cannot be performed on the corresponding switch port, as there are no information about the link state on the other end of a link. These metrics are marked as n/a in the results table 3.

TDR test on a Fa 0/2 port of SW2 stuck in 'TDR test is in progress' for half an hour. After that port had to be turned off and on to bring a port back to the operational state. However, it was impossible to repeat TDR test attempt until the switch was restarted.

Surprisingly, link state active test from a PC does not work as expected. When interface is shut down on the PC, switch detects that the line protocol goes down and than back up immediately. The same behaviour is observed for both Windows and Linux on the PC.

The obtained results are provided in table 3.

Table 3: Results of baseline measurements.

	PC *eth0*	SW1 *Gi 0/2*	SW1 *Gi 0/1*	SW2 *Fa 0/1*	SW2 *Fa 0/2*	AP *Main*
Link						
State	Up	Up	Up	Up	Up	Up
Speed (Mb/s)	1000	1000	100	100	100	n/a
Duplex	Full	Full	Full	Full	Full	n/a
PoE	n/a	n/a	n/a	PSE	PSE	PD
Client detected	-	-	-	No	Yes	-
Client class	-	-	-	-	4	-
Requested power	-	-	-	-	15.44	-
TDR	n/a			Wrong	Stuck	n/a
Pair A						
Length (m)	-	8 +/- 10	N/A	0 +/- 2	-	-
Remote pair	-	B	B	A	-	-
Status	-	Normal	Normal	Short	-	-
Pair B						
Length (m)	-	8 +/- 10	N/A	0 +/- 2	-	-
Remote pair	-	A	A	A	-	-
Status	-	Normal	Normal	Short	-	-
Pair C						
Length (m)	-	8 +/- 10	21 +/- 2	N/A	-	-
Remote pair	-	D	D	A	-	-
Status	-	Normal	Short	N/A	-	-
Pair D						
Length (m)	-	8 +/- 10	20 +/- 2	N/A	-	-
Remote pair	-	C	C	A	-	-
Status	-	Normal	Short	N/A	-	-
ICMP						
RTT	1.9	2	3	4	2	2.1
Packets lost	0	0	0	0	0	0
Link - active						
State	Down/Up[1]	Match	Match	Match	Match [2]	n/a
Speed	Match	Match	Match	Match	n/a	n/a
Duplex	Match	Match	Match	Match	n/a	n/a

1 Link went down and then immediately up
2 When link is disabled, PoE goes down as well.

When testing TDR on a Fa 0/1 port of SW2, link went down for more than a minute. During this test PoE related log messages were generated. These messages are shown in listing 11.

```
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1,
    changed state to down
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to down
%ILPOWER-7-DETECT: Interface Fa0/1: Power Device detected: Cisco PD
%ILPOWER-5-POWER_GRANTED: Interface Fa0/1: Power granted
%ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/1: PD removed
%ILPOWER-7-DETECT: Interface Fa0/1: Power Device detected: Cisco PD
%ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/1: PD removed
%ILPOWER-7-DETECT: Interface Fa0/1: Power Device detected: Cisco PD
%ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/1: PD removed
%ILPOWER-7-DETECT: Interface Fa0/1: Power Device detected: Cisco PD
%ILPOWER-5-POWER_GRANTED: Interface Fa0/1: Power granted
%ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/1: PD removed
%ILPOWER-7-DETECT: Interface Fa0/1: Power Device detected: Cisco PD
%ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/1: PD removed
%ILPOWER-7-DETECT: Interface Fa0/1: Power Device detected: Cisco PD
%ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/1: PD removed
%ILPOWER-7-DETECT: Interface Fa0/1: Power Device detected: Cisco PD
%ILPOWER-5-IEEE_DISCONNECT: Interface Fa0/1: PD removed
%LINK-3-UPDOWN: Interface FastEthernet0/1, changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/1,
    changed state to up
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/2,
    changed state to down
%LINEPROTO-5-UPDOWN: Line protocol on Interface FastEthernet0/2,
    changed state to up
```

Listing 11: Log messages generated while performing TDR tests on a SW1 port Fa 0/1.

After manually disabling PoE on a port, test completed successfully without PoE log messages. However, results were the same as in the first test.

Unfortunately, these TDR results do not seem reasonable or reliable in any way. These results showed that each cable pair is connected to a remote pair A, which was not the case, as cable was fully functioning. At the same time TDR determined incorrect cable length for two wire pairs. Regardless of this fact, in a following measurements TDR tests on this port are performed. But, their results are analyzed very thoroughly.

4.2 PC bridge man in the middle

Events generated on the network devices during the installation of a MitM PC are shown in table 4. The table includes shortened log messages. Even though AP is running Linux,

no network related events are logged on it. Therefore, this cell of a table left empty.

Table 4: Events logged on a devices during the installation.

Device	Port	Logged events
PC	eth0	Link is down Link is up at 1000 Mbps, full duplex
SW1	Gi 0/2	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW1	Gi 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/2	Line protocol changed state to down PD removed Interface changed state to down Interface changed state to up Line protocol changed state to up
AP	Main	-

Events generated on the devices are usual log events which happen when a link goes down and then up. These events may be used as an indicator of abnormal behaviour, which needs to be investigated. However, it may be hard to use this indicator, depending on the environment. It is important to know what is expected behaviour for each link in a network before reacting to such events. For some links in a network it may be absolutely normal to go down and up. For others it may be not.

The only odd event in these log messages is the `PD removed` message on SW2 Fa 0/2 without a matching `Power Device detected` message. From this it can be concluded that the device behind the port was replaced. This finding is confirmed by a PoE metric as shown in table 5.

Results of all measurements are shown in table 5.

Table 5: Results of measurements after MitM PC is installed on each of the links.

	PC *eth0*	SW1 *Gi 0/2*	SW1 *Gi 0/1*	SW2 *Fa 0/1*	SW2 *Fa 0/2*	AP *Main*
Link						
State	Up	Up	Up	Up	Up	Down [1]
Speed (mb/s)	1000	1000	1000	100	100	-
Duplex	Full	Full	Full	Full	Full	-
POE	n/a	n/a	n/a	PSE	PSE	-
Client detected	-	-	-	No	No [1]	-
Client class	-	-	-	-	-	-
Requested power	-	-	-	-	-	-
TDR	n/a			Wrong	-	-
Pair A						
Length (m)	-	0 +/- 10	24 +/- 10	N/A	-	-
Remote pair	-	B	B	A	-	-
Status	-	Normal	Normal	Normal	-	-
Pair B						
Length (m)	-	0 +/- 10	24 +/- 10	N/A	-	-
Remote pair	-	A	A	A	-	-
Status	-	Normal	Normal	Normal	-	-
Pair C						
Length (m)	-	0 +/- 10	20 +/- 2	N/A	-	-
Remote pair	-	D	D	A	-	-
Status	-	Normal	Open	N/A	-	-
Pair D						
Length (m)	-	0 +/- 10	19 +/- 2	N/A	-	-
Remote pair	-	C	C	A	-	-
Status	-	Normal	Open	N/A	-	-
ICMP						
RTT	2.5	3	5	5	-	-
Packets lost	0	0	0	0	-	-
Link - active						
State	No changes	No changes	No changes	No changes	-	-
Speed	No changes	No changes	No changes	No changes	-	-
Duplex	No changes	No changes	No changes [2]	No changes [2]	-	-

1 No PoE was provided to the AP. Attack considered unsuccessful on this link.
2 CDP duplex mismatch events were logged.

When the MitM PC was inserted in a link between SW2 Fa 0/2 and AP, no PoE was provided to the AP. Attack considered unsuccessful on this link. Therefore, no measurements were done on the corresponding switch port.

As one can see from table 5, Gigabit Ethernet port Gi 0/1 on a switch SW1 negotiated a link on a 1000 Mb/s, whereas the Fa 0/1 port on the SW2, which is on the other end of a cable does not support speeds above 100 Mb/s. This speed mismatch is a strong indicator that the devices are not connected directly.

TDR results showed radical cable length shortage for a SW1 port Gi 0/2. On the cable connected to this port, attack point is right next to a SW1. And TDR accurately reflected the cable length change.

For a Gi 0/1 port of the SW1 TDR reflected pair status change for pairs C and D, and a slight variation in a cable length. These changes together are a really strong indicator that the device on the end of a cable was changed.

TDR test on a SW2 was performed only on a port Fa 0/1. Cable on a port Fa 0/2 was not tested because an attack on the AP is considered unsuccessful. Unfortunately, TDR results for SW2 Fa 0/2 port does not seem reasonable again, like it was in the baseline measurements. The difference is that this time TDR was unable to determine a cable length for any of the pairs.

RTT values increased slightly, however, this metric alone does not seem sufficient to reliably indicate that there is additional device along the cable.

All active link tests showed the same results: Devices on the opposite ends were not seeing any link parameter changes during the tests, but the links were still working. This indicates that the devices are not connected directly.

When the link-active duplex test were performed on a cable between switches, multiple Cisco Discovery Protocol (CDP) duplex mismatch error messages were generated. Example of such message is shown in listing 12.

```
%CDP-4-DUPLEX_MISMATCH: duplex mismatch discovered on GigabitEthernet0
   /1 (not full duplex), with SW2 FastEthernet0/1 (full duplex).
```

Listing 12: Example of a message generated on a SW1 during the link-active duplex test.

These CDP messages could be used as an attack indicator as well.

4.3 Switch with port mirroring

A switch used to perform this attack is a Cisco Catalyst 2960 switch, similar to SW1. Switch configuration is shown in appendix 3.

Events generated during the attack setup process are shown in table 6.

Table 6: Events logged on a devices during the installation.

Device	Port	Logged events
PC	eth0	Link is down Link is up at 1000 Mbps, full duplex
SW1	Gi 0/2	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW1	Gi 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/2	PD removed Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
AP	Main	-

As can be seen, generated log messages are similar to the ones generated during the PC MitM attack. The only difference is the order of `PD removed` and `Line protocol changed state to down` messages on a SW2 port Fa 0/2.

Exactly as in the previous attack, no power was provided to the AP and an attack on this link was considered unsuccessful.

Link state, speed and duplex metrics are identical to the ones observed in the previous attack. PoE metrics are the same as well.

TDR results for the SW1 port Gi 0/1 are different. This time different wiring is detected and all four pairs have the same status and length. TDR results for other ports are identical to the results from the previous attack.

This time RTT values are identical to the baseline values.

Measurements results for this attack are shown in table 7.

Table 7: Results of measurements after a switch with port mirroring is installed.

	PC *eth0*	SW1 *Gi 0/2*	SW1 *Gi 0/1*	SW2 *Fa 0/1*	SW2 *Fa 0/2*	AP *Main*
Link						
State	Up	Up	Up	Up	Up	Down [1]
Speed (mb/s)	1000	1000	1000	100	100	-
Duplex	Full	Full	Full	Full	Full	-
POE	n/a	n/a	n/a	PSE	PSE	-
Client detected	-	-	-	No	No [1]	-
Client class	-	-	-	-	-	-
Requested power	-	-	-	-	-	-
TDR	n/a			Wrong	-	-
Pair A						
Length (m)	-	0 +/- 10	24 +/- 10	0 +/- 2	-	-
Remote pair	-	B	A	A	-	-
Status	-	Normal	Normal	Short	-	-
Pair B						
Length (m)	-	0 +/- 10	24 +/- 10	N/A	-	-
Remote pair	-	A	B	A	-	-
Status	-	Normal	Normal	Normal	-	-
Pair C						
Length (m)	-	0 +/- 10	24 +/- 2	N/A	-	-
Remote pair	-	D	C	A	-	-
Status	-	Normal	Normal	N/A	-	-
Pair D						
Length (m)	-	0 +/- 10	24 +/- 2	N/A	-	-
Remote pair	-	C	D	A	-	-
Status	-	Normal	Normal	N/A	-	-
ICMP						
RTT	1.9	2	3	4	-	-
Packets lost	0	0	0	0	-	-
Link - active						
State	No changes	No changes	No changes	No changes	-	-
Speed	No changes	No changes	No changes	No changes	-	-
Duplex	No changes	No changes	No changes	No changes	-	-

1 No PoE was provided to the AP. Attack considered unsuccessful on this link.

Active link tests have the same results. The only difference is that CDP was not complaining about the duplex of a link between switches. To investigate this behaviour, some additional checks were performed. Apparently, CDP messages do not pass an attacker's switch. Thus, switches have no CDP information about one another.

4.4 Throwing Star LAN Tap

Log events generated during the installation are shown in table 8.

Table 8: Events logged on a devices during the installation.

Device	Port	Logged events
PC	eth0	Link is down Link is up at 100 Mbps, full duplex
SW1	Gi 0/2	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW1	Gi 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/2	Line protocol changed state to down PD removed Interface changed state to down Power Device detected: IEEE PD Power granted Interface changed state to up Line protocol changed state to up Line protocol changed state to down Line protocol changed state to up
AP	Main	-

These events differ from the ones seen in the previous attacks. This time PC link come up on a lower speed. Another difference is that the tap does not interrupt a PoE operation and AP receives power. The same observations are confirmed in a results table 9.

Measurements results for this attack are shown in table 9

Table 9: Results of measurements after Throwing Star LAN Tap is installed and listening PC is connected to it.

	PC *eth0*	SW1 *Gi 0/2*	SW1 *Gi 0/1*	SW2 *Fa 0/1*	SW2 *Fa 0/2*	AP *Main*
Link						
State	Up	Up	Up	Up	Up	Up
Speed (mb/s)	100	100	100	100	100	n/a
Duplex	Full	Full	Full	Full	Full	n/a
POE	n/a	n/a	n/a	PSE	PSE	PD
Client detected	-	-	-	No	Yes	-
Client class	-	-	-	-	4	-
Requested power	-	-	-	-	15.44	-
TDR	n/a			Wrong	Stuck	n/a
Pair A						
Length (m)	-	0 +/- 2	19 +/- 2	0 +/- 2	-	-
Remote pair	-	A	B	A	-	-
Status	-	Impedance mismatch	Impedance mismatch	Normal	-	-
Pair B						
Length (m)	-	N/A	20 +/- 2	0 +/- 2	-	-
Remote pair	-	B	A	A	-	-
Status	-	Normal	Impedance mismatch	Short	-	-
Pair C						
Length (m)	-	0 +/- 2	21 +/- 2	N/A	-	-
Remote pair	-	C	D	A	-	-
Status	-	Short	Short	N/A	-	-
Pair D						
Length (m)	-	3 +/- 2	19 +/- 2	N/A	-	-
Remote pair	-	D	C	A	-	-
Status	-	Short	Short	N/A	-	-
ICMP						
RTT	1.9	3	4	3	2	2.0
Packets lost	0	0	0	0	0	0
Link - active						
State	Down/Up[1]	Match	Match	Match	Match [2]	n/a
Speed	Match	Match	Match	Match	n/a	n/a
Duplex	Match	Match	Match	Match	n/a	n/a

1 Link went down and then immediately up
2 When link is disabled, PoE goes down as well.

From the first line of the table, the results differ from the ones seen in the previous attacks. Now both SW1 Gi 0/1 and SW2 Fa 0/2 share the same speed. However, link between

PC and SW1 Gi 0/2 degraded to a 100 Mb/s. Both of these observations are perfectly in line with the theory behind an attack.

Another major difference is that a TDR test detected impedance mismatch on multiple links. This is a really strong indicator that the link was tampered with. Cable wiring and cable length changes are detected as well. Beside of that, TDR reported some short pairs, where the normal pairs originally were. This is due to a fact that a Throwing Star LAN Tap degrades a quality of two pairs, which are not used for the data transfer by 100BASE-TX and 10BASE-T.

RTT values are a little bit higher than the baseline values. Unlike all the previous attacks, all link active tests showed the same results as during the baseline measurements.

4.5 Electrically passive Ethernet tap with direct connection to the wires

Log events generated during an installation phase of this attack are shown in table 10.

Table 10: Events logged on a devices during the installation.

Device	Port	Logged events
PC	eth0	Link is down Link is up at 100 Mbps, full duplex
SW1	Gi 0/2	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW1	Gi 0/1	-
SW2	Fa 0/1	-
SW2	Fa 0/2	-
AP	Main	-

As can be seen, log messages were generated on PC and SW1, when the link was renegotiated to a 100 Mb/s speed. There were no log messages generated on links which worked at 100 Mb/s.

Results of measurements for the topology under this attack are shown in table 11

Table 11: Results of measurements after electrically passive Ethernet tap is connected directly to wires and a listening PC is connected to it.

	PC *eth0*	SW1 *Gi 0/2*	SW1 *Gi 0/1*	SW2 *Fa 0/1*	SW2 *Fa 0/2*	AP *Main*
Link						
State	Up	Up	Up	Up	Up	Up
Speed (mb/s)	100	100	100	100	100	n/a
Duplex	Full	Full	Full	Full	Full	n/a
POE	n/a	n/a	n/a	PSE	PSE	PD
Client detected	-	-	-	No	Yes	-
Client class	-	-	-	-	4	-
Requested power	-	-	-	-	15.44	-
TDR	n/a			Wrong	Stuck	n/a
Pair A						
Length (m)	-	0 +/- 2	19 +/- 2	0 +/- 2	-	-
Remote pair	-	B	B	A	-	-
Status	-	Impedance mismatch	Impedance mismatch	Normal	-	-
Pair B						
Length (m)	-	N/A	20 +/- 2	0 +/- 2	-	-
Remote pair	-	A	A	A	-	-
Status	-	Normal	Impedance mismatch	Short	-	-
Pair C						
Length (m)	-	0 +/- 2	21 +/- 2	N/A	-	-
Remote pair	-	D	D	A	-	-
Status	-	Short	Short	N/A	-	-
Pair D						
Length (m)	-	3 +/- 2	19 +/- 2	N/A	-	-
Remote pair	-	C	C	A	-	-
Status	-	Short	Short	N/A	-	-
ICMP						
RTT	1.9	3	4	3	2	2.0
Packets lost	0	0	0	0	0	0
Link - active						
State	Down/Up[1]	Match	Match	Match	Match [2]	n/a
Speed	Match	Match	Match	Match	n/a	n/a
Duplex	Match	Match	Match	Match	n/a	n/a

1 Link went down and then immediately up
2 When link is disabled, PoE goes down as well.

The results of the measurements for this tap are generally the same as for the Throwing Star LAN Tap. The main difference is that this tap does not change wiring and does not

disrupt connectivity when installed on a 100BASE-TX or 10BASE-T link.

4.6 Net Optics Fast Ethernet tap

The log messages generated on devices while tap was installed are shown in table 12.

Table 12: Events logged on a devices during the installation.

Device	Port	Logged events
PC	eth0	Link is down Link is up at 100 Mbps, full duplex
SW1	Gi 0/2	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW1	Gi 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/2	PD removed Line protocol changed state to down Interface changed state to down Power Device detected: IEEE PD Power granted Interface changed state to up Line protocol changed state to up Line protocol changed state to down Line protocol changed state to up
AP	Main	-

Log messages match with the messages generated during the attack with a Throwing Star LAN Tap.

As can be seen from the log messages and from the results table 13, after installation link between the PC and the Gi 0/2 port on SW1 degraded to a 100 Mb/s. This is expected, as the tap used in this attack does not support speeds higher than 100BASE-TX.

Results of measurements are shown in table 13.

Table 13: Results of measurements after Net Optics TP-CU tap is installed and listening PC is connected to it.

	PC *eth0*	SW1 *Gi 0/2*	SW1 *Gi 0/1*	SW2 *Fa 0/1*	SW2 *Fa 0/2*	AP *Main*
Link						
State	Up	Up	Up	Up	Up	Up
Speed (mb/s)	100	100	100	100	100	n/a
Duplex	Full	Full	Full	Full	Full	n/a
POE	n/a	n/a	n/a	PSE	PSE	PD
Client detected	-	-	-	No	Yes	-
Client class	-	-	-	-	4	-
Requested power	-	-	-	-	15.44	-
TDR	n/a			Wrong	Stuck	n/a
Pair A						
Length (m)	-	N/A	N/A	0 +/- 2	-	-
Remote pair	-	B	A	A	-	-
Status	-	Normal	Normal	Short	-	-
Pair B						
Length (m)	-	N/A	21 +/- 2	0 +/- 2	-	-
Remote pair	-	A	B	A	-	-
Status	-	Normal	Impedance mismatch	Short	-	-
Pair C						
Length (m)	-	1 +/- 2	19 +/- 2	N/A	-	-
Remote pair	-	D	C	A	-	-
Status	-	Short	Short	N/A	-	-
Pair D						
Length (m)	-	0 +/- 2	19 +/- 2	N/A	-	-
Remote pair	-	C	D	A	-	-
Status	-	Short	Short	N/A	-	-
ICMP						
RTT	1.9	2	4	4	2	2.0
Packets lost	0	0	0	0	0	0
Link - active						
State	Down/Up[1]	Match	Match	Match	Match [2]	n/a
Speed	Match	Match	Match	Match	n/a	n/a
Duplex	Match	Match	Match	Match	n/a	n/a

1 Link went down and then immediately up
2 When link is disabled, PoE goes down as well.

As expected, the results of the measurements are similar to the results observed in the attacks with electrically passive tap. However, it was highly unexpected to see an impedance mismatch with this tap. As one may think, a network device manufactured by a big vendor

should not allow this.

4.7 Net Optics Gigabit tap

Log messages observed during the tap installation are shown in table 14.

Table 14: Events logged on a devices during the installation.

Device	Port	Logged events
PC	eth0	Link is down Link is up at 1000 Mbps, full duplex
SW1	Gi 0/2	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW1	Gi 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/1	Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
SW2	Fa 0/2	Line protocol changed state to down PD removed Interface changed state to down Power Device detected: IEEE PD Power granted Interface changed state to up Line protocol changed state to up Line protocol changed state to down Interface changed state to down Interface changed state to up Line protocol changed state to up
AP	Main	-

Events shown in logs on a SW2 Fa 0/1 and SW2 Fa 0/2 differ from the usual log entries for link going down and up. For some reason during the tap installation link on a switch SW2 port Fa 0/1 went down and up twice.

Unusual behaviour is also observed when performing link-active tests from a switch. As can be seen from listing 13, the link is negotiated at 1000 Mb/s before the link-active test.

```
eth0: Link is up at 1000 Mbps, full duplex
```

Listing 13: The last log message on the PC before a link active speed test is started.

When switch port is set to 100 Mb/s, log messages from listing 14 are logged on a PC.

```
eth0: Link is down
eth0: Link is up at 1000 Mbps, full duplex
eth0: Link is down
eth0: Link is up at 100 Mbps, full duplex
```

Listing 14: Log messages generated on the PC after switch (SW1) port (Fa 0/2) speed is changed to 100 Mb/s.

Setting a duplex to half duplex on a switch produces the output, which is shown in listing 15.

```
eth0: Link is down
eth0: Link is up at 1000 Mbps, full duplex
eth0: Link is down
eth0: Link is up at 100 Mbps, half duplex
```

Listing 15: Log messages generated on the PC after switch (SW1) port (Fa 0/2) duplex is changed to half duplex.

As can be seen, the link between the PC and the tap is first negotiated at 1000 Mb/s, full duplex and then renegotiated to a speed supported by the SW2 Fa 0/2 port. It looks like the tap which was used to perform this attack negotiates maximum available speed on each end and then drops it to a higher speed supported by both ends, if needed. Such double log messages potentially could be used to detect this or similar tap installed in a network.

Results of all the measurements for this attack are shown in table 15.

Table 15: Results of measurements after Net Optics TP-CU3-ZD tap is installed.

	PC *eth0*	SW1 *Gi 0/2*	SW1 *Gi 0/1*	SW2 *Fa 0/1*	SW2 *Fa 0/2*	AP *Main*
Link						
State	Up	Up	Up	Up	Up	Up
Speed (mb/s)	1000	1000	100	100	100	n/a
Duplex	Full	Full	Full	Full	Full	n/a
POE	n/a	n/a	n/a	PSE	PSE	PD
Client detected	-	-	-	No	Yes	-
Client class	-	-	-	-	4	-
Requested power	-	-	-	-	15.44	-
TDR	n/a			Wrong	Stuck	n/a
Pair A						
Length (m)	-	4 +/- 10	N/A	0 +/- 2	-	-
Remote pair	-	A	B	A	-	-
Status	-	Normal	Normal	Short	-	-
Pair B						
Length (m)	-	4 +/- 10	N/A	0 +/- 2	-	-
Remote pair	-	B	A	A	-	-
Status	-	Normal	Normal	Short	-	-
Pair C						
Length (m)	-	4 +/- 10	N/A	N/A	-	-
Remote pair	-	C	D	A	-	-
Status	-	Normal	Normal	N/A	-	-
Pair D						
Length (m)	-	4 +/- 10	N/A	N/A	-	-
Remote pair	-	D	C	A	-	-
Status	-	Normal	Normal	N/A	-	-
ICMP						
RTT	1.9	2	4	3	2	2.0
Packets lost	0	0	0	0	0	0
Link - active						
State	Down/Up[1]	Match	Match	Match	Match [2]	n/a
Speed	Match	Match	Match	Match	n/a	n/a
Duplex	Match	Match	Match	Match	n/a	n/a

1 Link went down and then immediately up
2 When link is disabled, PoE goes down as well.

Besides the odd log messages, TDR is the only metric which detected this attack.

4.8 Combined results

Returning back to the goals stated in the introduction, this project aims to answer following research questions:

1. Is it possible to reliably detect all known physical layer attacks in Ethernet networks by observing information available to the network devices?
2. Is there a single metric which could be used to detect each physical layer attack in Ethernet networks?

Results of experiments which are provided in sections 4.1 to 4.7 can be used to answer the first research question. To answer the second research question, table 16 is constructed.

Table 16: Effectiveness of metrics based on the attack type.

	PC MitM	Switch with port mirroring	Throwing Star LAN Tap	Electrically passive tap	Fast Ethernet tap	Gigabit tap
Link						
State	Partial	Partial	None	None	None	None
Speed	Partial	Partial	Partial	Partial	Partial	None
Duplex	None	None	None	None	None	None
POE						
Client detected	Partial	Partial	None	None	None	None
Client class	Partial	Partial	None	None	None	None
Requested power	Partial	Partial	None	None	None	None
TDR						
Length	Full	Full	Full	Full	Full	Full
Remote pair	None	Partial	Partial	None	Partial	Partial
Status	Partial	Partial	Full	Full	Full	Partial
ICMP						
RTT	None	None	None	None	None	None
Packets lost	None	None	None	None	None	None
Link - active						
State	Full	Full	None	None	None	None
Speed	Full	Full	None	None	None	None
Duplex	Full	Full	None	None	None	None
Log events	Full	Full	Full	Partial	Full	Full

Table 16 groups together experiment results and shows how effective each metric is in detecting each of the selected attacks. Effectiveness of a metric is marked as full, partial or none, where full corresponds to a metric, which detected an attack on all links in a topology, partial corresponds to a metric which detected an attack on at least one link and none corresponds to a metric which did not detected an attack.

TDR metrics results for each pair are merged in one. In this way, if at least one pair showed a value which differs from the baseline value, combined metric is considered different from the baseline value. TDR results for SW2 switch are excluded, as they are not reliable.

5 Discussion

The main benefit of the approach covered in this project is that it allows to detect physical layer attacks on Ethernet networks without having to purchase any additional equipment. It makes this approach highly scalable. However, there are some drawbacks connected to it. The biggest drawback is that multiple tests used in this project are disruptive. During these tests link downtime varies from a couple of a second break for a link parameters renegotiation up to an undefined amount of time due to a complete test failure where the device has to be restarted. Some tests cause a device power loss, which leads to a device shutting down, and this may cause a data loss. These facts may be sufficient for some not to implement such tests in their environment.

A couple of the proposed metrics were not tested due to limitations posed by network devices selected for testing. It is possible that these metrics may be successfully used on the other models of network equipment.

As shown in table 16 some of the metrics show no effectiveness at all. These metrics are link duplex, ICMP RTT and packet lost counter. The fact that these metrics did not react to any of the attacks does not mean that they are useless for attack detection. They may react to attacks in different conditions. For example, link duplex degrades to a half duplex if a hub is used for attack. ICMP RTT and packet lost counter may react to an attack with a device which is not able to hold the load, such as a slow PC acting as a bridge.

6 Conclusions

The goal of this project was to determine if it is possible to reliably detect known physical layer attacks on Ethernet networks using only common network equipment, such as switches and PCs. Results covered in chapter 4 show that it is possible to detect all tested attacks in the example topology using selected metrics.

The most efficient method for detecting physical layer attacks happened to be time domain reflectometry, which is, unfortunately, not widely supported by various network equipment. According to the test results, the second most efficient method of detecting such attacks is to monitor log events. Luckily, almost every device can be configured to generate logs, which then can be aggregated and analyzed. The problem is that, in order to efficiently use log events for attacks detection, it is important to know which log events are normal for a given network and which are not.

During the tests some network equipment demonstrated highly unexpected behaviour and some equipment provided unreliable test results. The most important conclusion that can be made is that it is important to know the network and network equipment really well to be able to implement attack detection efficiently.

References

1. LeClair J et al. Cybersecurity in Our Digital Lives. Albany, NY: Hudson Whitman/Excelsior College Press; 2015.
2. Cole E. Advanced Persistent Threat: Understanding the Danger and How to Protect Your Organization. Waltham, MA: Elsevier Science; 2012.
3. Pwnie Express. Pwn Plug Elite [online]. Boston, MA: Pwnie Express; 2014. URL: https://www.pwnieexpress.com/product/pwn-plug-elite/. Accessed Mar. 14, 2015.
4. Spurgeon C, Zimmerman J. Ethernet: The Definitive Guide. Sebastopol, CA: O'Reilly Media; 2014.
5. Pfleeger C, Pfleeger S. Security in Computing. Upper Saddle River, NJ: Prentice Hall PTR; 2003.
6. Quigley M. Encyclopedia of Information Ethics and Security. Hershey, PA: Information Science Reference; 2007.
7. Andress J. The Basics of Information Security: Understanding the Fundamentals of InfoSec in Theory and Practice. Waltham, MA: Elsevier Science; 2014.
8. Conrad E, Misenar S, Feldman J. CISSP Study Guide. Waltham, MA: Elsevier Science; 2012.
9. Cole E. Hackers Beware: Defending Your Network from the Wily Hacker. Indianapolis, IN: New Riders; 2002.
10. Sloan R, Warner R. Unauthorized Access: The Crisis in Online Privacy and Security. Boca Raton, FL: Taylor & Francis; 2013.
11. Annual Review of Communications. Volume. 59. Chicago, IL: International Engineering Consortium; 2007.
12. Toxen B. Real World Linux Security: Intrusion Prevention, Detection, and Recovery. Upper Saddle River, NJ: Prentice Hall PTR; 2003.
13. Weidman G. Penetration Testing. San Francisco, CA: No Starch Press; 2014.
14. Alcorn W, Frichot C, Orru M. The Browser Hacker's Handbook. Indianapolis, IN: Wiley; 2014.
15. Barrett D, Silverman R, Byrnes R. SSH, The Secure Shell: The Definitive Guide. Sebastopol, CA: O'Reilly Media; 2005.
16. Helba S, ed. Computer Forensics: Investigating Wireless Networks and Devices. Clifton Park, NY: Cengage Learning; 2009.
17. Sud R, Edelman K. SECUR Exam Cram 2. Indianapolis, IN: Que; 2003.
18. Douligeris C, Serpanos D. Network Security: Current Status and Future Directions. Alameda, CA: Wiley; 2007.
19. Blank A. TCP/IP Foundations. Alameda, CA: Wiley; 2006.
20. Spurgeon C. Ethernet: The Definitive Guide: The Definitive Guide. Sebastopol, CA: O'Reilly Media; 2000.
21. Cannon K, Caudle K, Chiarella A. CCNA Guide to Cisco Networking Fundamentals. Boston, MA: Course Technology; 2008.

22. Reynders D, Wright E. Practical TCP/IP and Ethernet Networking. Burlington, MA: Newnes; 2003.

23. West J, Dean T, Andrews J. Network+ Guide to Networks. Boston, MA: Cengage Learning; 2015.

24. Semenov A, Strizhakov S, Suncheley I. Structured Cable Systems. Heidelberg: Springer; 2002.

25. Bird D, Harwood M. Network+ Training Guide. Indianapolis, IN: Que; 2002.

26. Sanders C. Practical Packet Analysis: Using Wireshark to Solve Real-world Network Problems. San Francisco, CA: No Starch Press; 2011.

27. Froom R, Sivasubramanian B, Frahim E. Implementing Cisco IP Switched Networks (SWITCH) Foundation Learning Guide. Indianapolis, IN: Cisco Press; 2010.

28. Marget C. Ethernet Taps - Don't Get Me Started [online]. New Hampshire: Chris Marget; Jan. 19, 2012.
URL: http://www.fragmentationneeded.net/2012/01/ethernet-taps-dont-get-me-started.html. Accessed Sept. 24, 2016.

29. Harrington J. Ethernet Networking for the Small Office and Professional Home Office. Burlington, MA: Morgan Kaufmann; 2010.

30. Net Optics. Installation Guide for Gig Zero DelayTM Tap and 10/100/1000BaseT Tap [online]. Santa Clara, CA: Net Optics; 2010.
URL: http://www.netoptics.com/sites/default/files/PUBTPCU3ZDU.pdf. Accessed Sept. 28, 2015.

31. Ossmann M. Throwing Star LAN Tap [online]. Evergreen, CO: Great Scott Gadgets; July 2011.
URL: http://greatscottgadgets.com/throwingstar/. Accessed Mar. 14, 2015.

32. Karunaratne J. The Passive Splice Network Tap [online]. Austin, TX: Janitha Karunaratne; Dec. 24, 2009.
URL: http://janitha.com/articles/passive-splice-network-tap/. Accessed Mar. 14, 2015.

33. Installation Guide for 10/100BaseT Tap. Sunnyvale, CA: Net Optics; 2006.

34. Cisco. Catalyst 2960 and 2960 -S Switch Command Reference, Release 12.2(58)SE - Catalyst 2960 Switch Cisco IOS Commands - rmon collection through show vtp [online]. San Jose, CA: Cisco; Oct. 14, 2013.
URL: http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_58_se/command/reference/2960cr/cli2.html. Accessed Sept. 23, 2015.

35. Reynolds H, Marschke D. JUNOS Enterprise Switching. Sebastopol, CA: O'Reilly Media; 2009.

36. Donahue G. Network Warrior. Sebastopol, CA: O'Reilly Media; 2011.

37. Lammle T. Todd Lammle's CCNA/CCENT IOS Commands Survival Guide. Indianapolis, IN: Wiley; 2014.

38. Schroder C. Linux Networking Cookbook. Sebastopol, CA: O'Reilly Media; 2007.

39. Wallace K. CCNP TSHOOT 642-832 Official Cert Guide. Indianapolis, IN: Cisco Press; 2010.

40. Cisco. Cisco Catalyst 2960-S and 2960 Series Switches with LAN Base Software [online]. San Jose, CA: Cisco; Mar. 2013.
URL: http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-2960-series-switches/product_data_sheet0900aecd80322c0c.pdf. Accessed Sept. 28, 2015.

41. Cisco. Cisco Catalyst Switch Guide: Scalable, intelligent LAN switching for campus, branch, and data center networks of all sizes [online]. San Jose, CA: Cisco; 2006. URL: http://www.cisco.com/web/AT/unified_partners/smb/vertriebliche-positionierung/-switching/downloads/guide_switching_e.pdf. Accessed Sept. 28, 2015.

42. Cisco. Cisco Catalyst 3560 v2 Series Switches [online]. San Jose, CA: Cisco; Mar. 2013.
URL: http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-3560-series-switches/data_sheet_c78-530976.pdf. Accessed Sept. 28, 2015.

43. Ubiquiti Networks. UniFi AP Datasheet. [online]. San Jose, CA: Ubiquiti Networks; 2015.
URL: http://dl.ubnt.com/datasheets/unifi/UniFi_AP_DS.pdf. Accessed Sept. 28, 2015.

44. Harrington D. CCNP Practical Studies: Troubleshooting. Indianapolis, IN: Cisco Press; 2003.

45. Turnbull J, Matotek D, Lieverdink P. Pro Linux System Administration. New York, NY: Apress; 2009.

46. Nemeth E, Snyder G, Hein T. Linux Administration Handbook. Upper Saddle River, NJ: Pearson Education; 2006.

1 Detailed information about the SW1

```
Cisco IOS Software, C2960 Software (C2960-LANBASEK9-M), Version 15.0(2)
   SE8, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2015 by Cisco Systems, Inc.
Compiled Thu 14-May-15 02:39 by prod_rel_team

ROM: Bootstrap program is C2960 boot loader
BOOTLDR: C2960 Boot Loader (C2960-HBOOT-M) Version 12.2(44)SE5, RELEASE
     SOFTWARE (fc1)

SW1 uptime is 2 minutes
System returned to ROM by power-on
System image file is "flash:c2960-lanbasek9-mz.150-2.SE8.bin"


This product contains cryptographic features and is subject to United
States and local country laws governing import, export, transfer and
use. Delivery of Cisco cryptographic products does not imply
third-party authority to import, export, distribute or use encryption.
Importers, exporters, distributors and users are responsible for
compliance with U.S. and local country laws. By using this product you
agree to comply with applicable laws and regulations. If you are unable
to comply with U.S. and local laws, return this product immediately.

A summary of U.S. laws governing Cisco cryptographic products may be
   found at:
http://www.cisco.com/wwl/export/crypto/tool/stqrg.html

If you require further assistance please contact us by sending email to
export@cisco.com.

cisco WS-C2960-24TT-L (PowerPC405) processor (revision C0) with 65536K
   bytes of memory.
Processor board ID FOC1121ZEVZ
Last reset from power-on
1 Virtual Ethernet interface
24 FastEthernet interfaces
2 Gigabit Ethernet interfaces
The password-recovery mechanism is enabled.

64K bytes of flash-simulated non-volatile configuration memory.
Base ethernet MAC Address       : 00:1C:57:6B:F5:00
Motherboard assembly number     : 73-10390-03
Power supply part number        : 341-0097-02
Motherboard serial number       : FOC112131ZQ
Power supply serial number      : AZS111704SS
Model revision number           : C0
Motherboard revision number     : C0
Model number                    : WS-C2960-24TT-L
System serial number            : FOC1121ZEVZ
```

```
Top Assembly Part Number         : 800-27221-02
Top Assembly Revision Number     : D0
Version ID                       : V02
CLEI Code Number                 : COM3L00BRA
Hardware Board Revision Number   : 0x01


Switch Ports Model              SW Version            SW Image
------ ----- -----              ----------            ----------
*    1 26    WS-C2960-24TT-L    15.0(2)SE8            C2960-LANBASEK9-M


Configuration register is 0xF
```

Listing 16: 'show version' command output for a SW1.

2 Detailed information about the SW2

```
Cisco IOS Software, C3560 Software (C3560-IPSERVICESK9-M), Version
   15.0(2)SE8, RELEASE SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2015 by Cisco Systems, Inc.
Compiled Wed 13-May-15 23:32 by prod_rel_team

ROM: Bootstrap program is C3560 boot loader
BOOTLDR: C3560 Boot Loader (C3560-HBOOT-M) Version 12.2(50r)SE, RELEASE
    SOFTWARE (fc1)

Switch uptime is 1 minute
System returned to ROM by power-on
System image file is "flash:c3560-ipservicesk9-mz.150-2.SE8.bin"


This product contains cryptographic features and is subject to United
States and local country laws governing import, export, transfer and
use. Delivery of Cisco cryptographic products does not imply
third-party authority to import, export, distribute or use encryption.
Importers, exporters, distributors and users are responsible for
compliance with U.S. and local country laws. By using this product you
agree to comply with applicable laws and regulations. If you are unable
to comply with U.S. and local laws, return this product immediately.

A summary of U.S. laws governing Cisco cryptographic products may be
   found at:
http://www.cisco.com/wwl/export/crypto/tool/stqrg.html

If you require further assistance please contact us by sending email to
export@cisco.com.

cisco WS-C3560V2-24PS (PowerPC405) processor (revision G0) with 131072K
    bytes of memory.
Processor board ID FDO1433X0AU
Last reset from power-on
1 Virtual Ethernet interface
24 FastEthernet interfaces
2 Gigabit Ethernet interfaces
The password-recovery mechanism is enabled.

512K bytes of flash-simulated non-volatile configuration memory.
Base ethernet MAC Address       : 68:BD:AB:20:58:00
Motherboard assembly number     : 73-11706-12
Power supply part number        : 341-0266-02
Motherboard serial number       : FDO14330HHM
Power supply serial number      : LIT14291K0L
Model revision number           : G0
Motherboard revision number     : A0
Model number                    : WS-C3560V2-24PS-E
System serial number            : FDO1433X0AU
```

```
Top Assembly Part Number         : 800-32964-04
Top Assembly Revision Number     : A0
Version ID                       : V04
CLEI Code Number                 : COMKV00DRA
Hardware Board Revision Number   : 0x03


Switch Ports Model              SW Version            SW Image
------ ----- -----              ----------            ----------
*    1 26    WS-C3560V2-24PS    15.0(2)SE8            C3560-
   IPSERVICESK9-M


Configuration register is 0xF
```

Listing 17: 'show version' command output for a SW2.

3 Configuration of the switch with port mirrorring

Following is the configuration, which was used for the switch during the attack.

```
SW-Mirror#sh run
Building configuration...

Current configuration : 2466 bytes
!
! Last configuration change at 00:26:45 UTC Mon Mar 1 1993
!
version 15.0
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname SW-Mirror
!
boot-start-marker
boot-end-marker
!
!
no aaa new-model
system mtu routing 1500
!
!
no ip domain-lookup
!
!
!
!
!
!
!
!
!
spanning-tree mode pvst
spanning-tree extend system-id
no spanning-tree vlan 1
!
vlan internal allocation policy ascending
!
!
!
!
!
!
interface FastEthernet0/1
 switchport mode access
 no keepalive
!
```

```
interface FastEthernet0/2
 switchport mode access
 no keepalive
!
interface FastEthernet0/3
 switchport mode access
 no keepalive
!
interface FastEthernet0/4
 switchport mode access
 no keepalive
!
interface FastEthernet0/5
 switchport mode access
 no keepalive
!
interface FastEthernet0/6
 switchport mode access
 no keepalive
!
interface FastEthernet0/7
 switchport mode access
 no keepalive
!
interface FastEthernet0/8
 switchport mode access
 no keepalive
!
interface FastEthernet0/9
 switchport mode access
 no keepalive
!
interface FastEthernet0/10
 switchport mode access
 no keepalive
!
interface FastEthernet0/11
 switchport mode access
 no keepalive
!
interface FastEthernet0/12
 switchport mode access
 no keepalive
!
interface FastEthernet0/13
 switchport mode access
 no keepalive
!
interface FastEthernet0/14
 switchport mode access
 no keepalive
!
interface FastEthernet0/15
 switchport mode access
 no keepalive
!
interface FastEthernet0/16
```

```
 switchport mode access
 no keepalive
!
interface FastEthernet0/17
 switchport mode access
 no keepalive
!
interface FastEthernet0/18
 switchport mode access
 no keepalive
!
interface FastEthernet0/19
 switchport mode access
 no keepalive
!
interface FastEthernet0/20
 switchport mode access
 no keepalive
!
interface FastEthernet0/21
 switchport mode access
 no keepalive
!
interface FastEthernet0/22
 switchport mode access
 no keepalive
!
interface FastEthernet0/23
 switchport mode access
 no keepalive
!
interface FastEthernet0/24
 switchport mode access
 no keepalive
!
interface GigabitEthernet0/1
 switchport mode access
!
interface GigabitEthernet0/2
 switchport mode access
!
interface Vlan1
 no ip address
 shutdown
!
ip http server
ip http secure-server
no cdp run
!
!
line con 0
 exec-timeout 0 0
 logging synchronous
line vty 0 4
 login
line vty 5 15
 login
```

```
!
!
monitor session 1 source interface Gi0/1
monitor session 1 destination interface Fa0/1
end
```

Listing 18: Configuration of a switch used during the attack.