Information Security

Security Analysis and Improvements of Time-Triggered Ethernet

OSI Layer 2 Network Security

Diploma Thesis

For attainment of the academic degree of

Diplom-Ingenieur/in

submitted by

Andreas Hagl

is211821

in the

University Course Information Security at St. Pölten University of Applied Sciences

The interior of this work has been composed in LaTeX.

Supervision

Advisor: Dipl.-Ing. Gabor Österreicher, BSc

St. Pölten, February 5, 2024

(Signature author) (Signature advisor)

Ehrenwörtliche Erklärung

Ich versichere, dass

- ich diese Diplomarbeit selbständig verfasst, andere als die angegebenen Quellen und Hilfsmittel nicht benutzt und mich sonst keiner unerlaubten Hilfe bedient habe.
- ich dieses Diplomarbeitsthema bisher weder im Inland noch im Ausland einem Begutachter/einer Begutachterin zur Beurteilung oder in irgendeiner Form als Prüfungsarbeit vorgelegt habe.
- diese Arbeit mit der vom Begutachter/von der Begutachterin beurteilten Arbeit übereinstimmt.

Der Studierende/Absolvent räumt der FH St. Pölten das Recht ein, die Diplomarbeit für Lehre- und Forschungstätigkeiten zu verwenden und damit zu werben (z.B. bei der Projektevernissage, in Publikationen, auf der Homepage), wobei der Absolvent als Urheber zu nennen ist. Jegliche kommerzielle Verwertung/Nutzung bedarf einer weiteren Vereinbarung zwischen dem Studierenden/Absolventen und der FH St. Pölten.

Ort, Datum *Unterschrift*

Kurzfassung

Um die Integrationsdichte zu erhöhen als auch Gewicht und Kosten einzusparen, werden in der Luft- und Raumfahrt verstärkt Netzwerksysteme eingesetzt, welche alle Anforderungen mit nur einem gemeinsamen physischen Netzwerk erfüllen können. Eines dieser Netzwerksysteme ist Time-Triggered Ethernet (TTEthernet). TTEthernet erlaubt es, zeitgesteuerte Kommunikation nach dem Standard SAE AS6802 gemeinsam mit Kommunikation nach ARINC 664 und gewöhnlichem Ethernet nach IEEE 802.3 auf derselben physischen Hardware zu verwenden. Jedes dieser drei Netzwerkprotokolle hat seine Berechtigung, da sie unterschiedliche Stärken und Schwächen haben. Sicherheitskritischer und unkritischer Datenverkehr, der früher durch verschiedene Netzwerke physisch getrennt war, wird jetzt über dasselbe Netzwerk abgewickelt; dies hat negative Auswirkungen auf die IT-Sicherheit. Erschwerend kommt hinzu, dass bei der Entwicklung dieser Netzwerksysteme Anforderungen an die IT-Sicherheit noch keine Rolle gespielt haben.

Das Ziel dieser Arbeit war es, die Sicherheitsstärken und -schwächen dieser Protokolle zu ermitteln und gegenüberzustellen. Anhand dieser Erkenntnisse wurden technische Maßnahmen entwickelt, um deren Sicherheit zu verbessern.

In dieser Arbeit wurde STRIDE als Analysemethode verwendet, um die Sicherheitsbedrohungen und Schwachstellen zu ermitteln. Zur Modellierung der Systeme wurden Datenflussdiagramme für jedes der drei Netzwerkprotokolle und die für die zeitgesteuerte Kommunikation erforderliche Uhrensynchronisation erstellt.

Insgesamt konnten 45 Bedrohungen identifiziert werden: 33 Bedrohungen für Standard IEEE 802.3 Ethernet-Netzwerke, 26 Bedrohungen für ARINC 664 Netzwerke und 37 Bedrohungen für zeitgesteuerte Netzwerke. 27 der identifizierten Sicherheitsbedrohungen und Schwachstellen können durch die elf Sicherheitsmaßnahmen, die in dieser Arbeit vorgeschlagen wurden, entschärft werden. Die Sicherheitsmaßnahme *Keyless Remote Attestation*, die in dieser Arbeit entwickelt wurde, ist ein neuartiger Ansatz zur Attestierung, der keine kryptografischen Funktionen und Schlüssel erfordert. Die vorgeschlagenen Sicherheitsmaßnahmen betreffen alle drei Verkehrsklassen und wirken sich auf die Konfiguration des Netzwerkes, die Netzwerktopologie, die Netzwerkprotokolle, sowie auf das Design und die Implementierung der Netzwerkgeräte aus. Die Ergebnisse und die in dieser Arbeit verwendete Analysemethode können als Grundlage für den Vergleich von TTEthernet mit weiteren OSI-Schicht 2 Netzwerkprotokollen verwendet werden.

Index: Time-Triggered Ethernet, Netzwerk Security, ARINC 664, AS6802, Threat Modeling, STRIDE, Security Maßnahmen, Keyless Remote Attestation

Abstract

In order to increase the integration density as well as to save weight and costs, mixed-critical network systems are increasingly being used in the aerospace industry which can fulfill all requirements with just one shared physical network. One of these network systems is Time-Triggered Ethernet (TTEthernet). TTEthernet makes it possible to use time-controlled communication according to the SAE AS6802 standard together with communication according to ARINC 664 and ordinary Ethernet according to IEEE 802.3 on the same physical layer. Each of these three network protocols has its use case, as they have different strengths and weaknesses. Security-critical and non-critical data traffic that used to be physically separated by different networks is now handled via the same network, which has a negative impact on security. To make matters worse, IT security requirements have not played a role in the development of these network systems.

The aim of this work was to determine and compare the security strengths and weaknesses of these protocols. Based on these findings, technical measures were developed to improve their security.

In this thesis, STRIDE was used as a threat modeling method to identify the security threats and weaknesses. Data flow diagrams were created for each of the three network protocols and for the clock synchronization required for time-triggered communication to model the system.

A total of 45 threats have been identified, these can be broken down into 33 threats affecting standard IEEE 802.3 Ethernet networks, 26 threats for ARINC 664 networks and 37 threats for time-trigged AS6802A networks. 27 of the identified security threats and weaknesses can be mitigated by the eleven security measures that have been proposed in this thesis. The security measure *Keyless Remote Attestation* developed in this thesis is a novel approach to attestation that does not require cryptographic functions and keys. The proposed security measures address all three traffic classes and have an impact on the configuration of the network, the network topology, the network protocols and the design and implementation of the network devices. The findings and the analysis method used in this thesis can be used as a baseline to compare TTEthernet with other OSI layer 2 network protocols.

Keywords: Time-Triggered Ethernet, Network Security, ARINC 664, AS6802, Threat Modeling, STRIDE, Security Measure, Keyless Remote Attestation

Contents

1 Introduction **1**

1.1 Motivation and problem statement 2

1.2 Objective of this Thesis 3

1.3 Contribution 3

1.4 Thesis Outline 4

2 Prerequisites **5**

2.1 ARINC 664 5

2.1.1 History 5

2.1.2 Operational priniciple 6

2.2 Time-Triggered Ethernet 7

2.2.1 Clock synchronization 8

2.2.2 Fault tolerance 10

3 Related Work **11**

4 Approach **13**

4.1 Scope 13

4.2 Methodology 14

4.3 Security Environment 15

5 Threat Modeling **17**

5.1 Identifying Assets 17

5.2 Data Flow Diagram 18

5.2.1 Best-effort communication 18

5.2.2 Rate-constrained communication 19

5.2.3 Time triggered communication 20

5.2.4 Clock synchronization mechanism 21
5.3 Identifying Security Threats 24
5.3.1 Best-Effort Communication 25
5.3.2 Rate-Constrained Communication 33
5.3.3 Time-Triggered Communication 39
5.3.4 Clock Synchronization 48
5.4 Analysis Results 63
6 Security Measures 69
6.1 M.1 - Static Configuration 69
6.2 M.2 - MACsec 70
6.3 M.3 - Keyless Remote Attestation 72
6.4 M.4 - Remote Attestation 74
6.5 M.5 - Authenticated PCFs 75
6.6 M.6 - Cryptographically Signed Configuration 76
6.7 M.7 - Optimize SM placement and selection 76
6.8 M.8 - Synchronize PHY and MAC 77
6.9 M.9 - Vote on redundant frames 77
6.10 M.10 - Restrict access to the frame memory 79
6.11 M.11 - Perform traffic policing on ES ingress ports 79
6.12 Out of Scope Security Measures 79
6.13 Residual Threats 80
7 Conclusion 83
7.1 Limitations 86
7.2 Future Work 86
A Detailed description of specific threats 87
List of Figures 94
List of Tables 97
Glossary 99

Bibliography **103**

1. Introduction

In safety and reliability critical systems a shift from segmented single-purpose networks to mixed-criticality networks is ongoing [1]–[6]. Mixed-criticality networks are systems that allow several traffic classes with different requirements, such as latency, jitter, bandwidth or fault tolerance, to coexist on the same physical network without interfering with the requirements of the other traffic classes. Many of these new network systems are based on Ethernet (IEEE 802.3) [7]–[9]. However, regular Ethernet does not provide the necessary performance and reliability guarantees like deterministic latency and redundancy that are necessary for safety critical applications [10], [11] therefore the new network systems extend Ethernet to provide these capabilities. Examples for the aerospace industry are ARINC 664 [7] or Time-Sensitive Networking (TSN) [8] for the automotive [12], industrial [13] and aerospace market [14]. These new network systems can reduce the size, weight, power consumption, system complexity and therefore cost and time-to-market, by replacing several legacy networks with just one combined network [15]. For example, in a spacecraft application, one network can be shared between vehicle command and control systems and onboard science payload [16], [17].

Another network system that supports mixed-critical traffic is Time-Triggered Ethernet (TTEthernet). TTEthernet was developed by the company TTTech with joint research done at the Vienna University of Technology [15], [18]. TTEthernet is a deterministic, fault-tolerant and time-critical Layer 2 network protocol that extends IEEE 802.3 Ethernet with Quality of Service (QoS) capabilities so it can be used as a reliable data network for safety-critical applications [9].

TTEthernet is used as the main network backbone in the following real-world applications:

- **Space:**
  - NASA's Orion Multi-Purpose Crew Vehicle (MPCV) capsule [18]–[21]
  - ESA's Ariane 6 launcher [22], [23]
  - NASA's Lunar Gateway space station [24]
- **Aerospace:**
  - Sikorsky's S-97 Raider helicopter [25]

- **Industrial:**
  - Wind turbines made by Vestas [26]

TTEthernet is also part of the International Avionics System Interoperability Standards (IASIS) for spacecraft [27]. This and additional standards are part of an international partnership by the National Aeronautics and Space Administration (NASA), the European Space Agency (ESA), the Canadian Space Agency (CSA), Roscosmos and the Japan Aerospace Exploration Agency (JAXA) to establish interoperable systems and interfaces.

1.1. Motivation and problem statement

TTEthernet was first developed in 2004 [15]. It was developed with safety in mind, there were no security requirements at the time when the network protocol was designed [28].

Since then, security has become more and more important in the industries where TTEthernet is used, with the primary goal of ensuring safety. The first security standard for avionic systems (DO-326) was released in 2010 [29]. The first subset of security standards for industrial automation and control systems (IEC 62443) was released in 2009 [30] and for the automotive market, it was in 2016 (SAE J3061) [31].

Attacks or proof of concept attacks on systems in these industries show that these systems are of interest to attackers. For example, Stuxnet was discovered in 2010 [32]. At the Black Hat USA 2018 security conference, a security analysis of satellite communication systems that provide in-flight Internet connectivity for crew and passengers was presented. Due to insufficient network segmentation, it would have been possible to gain direct access from the public Internet to the network used by passengers on the aircraft [33]. Another proof of concept attack was presented at the same conference one year later. By chaining multiple vulnerabilities in the network components of a Boeing 787, it could be possible to gain access to the avionics network through an unprivileged network segment, such as the in-flight entertainment system or the Internet access used by the passengers [34].

In 2022 an attack called PCspooF that can break the isolation guarantees of TTEthernet was discovered [35]. The attack demonstrated that it is possible to force a TTEthernet network to resynchronize by having a malicious best-effort end system perform a packet-in-packet attack using a well-dosed electromagnetic interference (EMI).

These rather new security standards and newly discovered security threats result in increased security needs for safety-critical applications and therefore will lead to additional security requirements for safety-critical systems. TTEthernet has some characteristics that can have the potential to improve the security of mixed-criticality systems. For example, the time-triggered nature of a TTEthernet network can render denial-of-service attacks ineffective.

1.2. Objective of this Thesis

TTEthernet devices support (at least) the following three traffic classes: Best-effort traffic (BE) is standard IEEE 802.3 Ethernet, rate-constrained traffic (RC) as defined in ARINC 664-part-7 and time-triggered traffic (TT). The clock synchronization mechanism that is necessary for time-triggered traffic is defined in AS6802A. All of these traffic classes have their strength and weaknesses and are used together for mixed critical networks, like the aforementioned real-world applications. It is likely that these three traffic classes also have different security strengths and weaknesses. This would be valuable information for the manufacturers, integrators and operators to improve the security of their network systems.

Therefore, the goal of this thesis is to answer the following research questions:

1. What are the security threats and weaknesses of time-trigged AS6802A traffic in comparison to standard IEEE 802.3 Ethernet and ARINC 664-part-7?
2. What technical security measures can be implemented to mitigate the identified security threats and weaknesses to improve the security properties of a TTEthernet network for all three traffic classes?

To answer the first question it is necessary to identify and analyze the security threats of the three traffic classes and compare them with each other. After that, security measures can be developed based on the analysis results to answer the second question. Technical security measures may include changes to the respective network protocol, the network configuration settings and requirements for switch and end system implementation.

1.3. Contribution

The STRIDE method was used for threat modeling to answer the first research question. A data flow diagram was created for every traffic class and the clock synchronization mechanism to model the network system. This has not been done before for TTEthernet and its clock synchronization mechanism. A total of 45 threats have been identified, these can be broken down into 33 threats for standard IEEE 802.3 Ethernet networks, 26 threats for ARINC 664-part-7 networks and 37 threats for time-trigged AS6802A networks. These findings and the analysis method used in this thesis can be used as a baseline to compare TTEthernet with other industrial OSI layer 2 network protocols in the future, such as TSN. 27 of the identified security threats and weaknesses can be mitigated by the eleven security measures that have been proposed in this thesis. Mitigations for the remaining threats are either well beyond OSI layer 2 or not a technical measure and therefore out of scope. The security measure *Keyless Remote Attestation* developed in this thesis is a novel approach to attestation that does not require cryptographic functions and keys. The proposed security measures address all three traffic classes and affect the configuration of the network, the network topology, the network protocols and suggest changes in the design and implementation of the network devices.

1.4. Thesis Outline

The overall structure of this document is as follows:

- **Introduction**: Chapter 1 introduces the topic, describes the motivation, defines the problem and outlines the research goal of this thesis.
- **Prerequisites**: A description of TTEthernet and ARINC 664 is provided in chapter 2, based on information from standards and research papers.
- **Related Work**: The related work section in chapter 3 presents existing research and literature that pertains to this study.
- **Approach**: Chapter 4 defines the scope and context for the security analysis of TTEthernet. It contains the identified assets and security properties and describes the system and environment to be analyzed.
- **Threat Modeling**: A security analysis of Ethernet, ARINC 664 and TTEthernet is performed in chapter 5 and lists the identified security weaknesses and strengths.
- **Security Measures**: The developed security measures can be found in chapter 6.
- **Conclusion**: Finally, chapter 7 draws a conclusion, summarizes the key findings of this thesis, and gives an outlook on possible further work.

2. Prerequisites

This chapter explains what Time-Triggered Ethernet (TTEthernet) is and how it works. This is necessary information for the subsequent chapters. TTEthernet was developed for mixed critical systems and supports multiple traffic classes that can operate alongside each other on the same single physical network. TTEthernet equipment that is available on the market supports three different traffic classes. The first traffic class is regular Ethernet as specified in IEEE 802.3 [36] which is used for best-effort (BE) communication. The BE traffic class is typically used for uncritical low-priority data. IEEE 802.3 Ethernet is a well-known network standard and is therefore not further explained in this chapter. The next traffic class is rate-constrained (RC) traffic. The traffic class is implemented as specified in the ARINC 664 standard. RC traffic can be used for critical asynchronous communication. Time-triggered (TT) traffic is the third traffic class and the one that is specific to TTEthernet. TT traffic provides synchronous communication and is best suited for critical applications that require deterministic latency and minimum jitter. TT traffic has the highest priority in the network. RC traffic has a lower priority than TT traffic and BE is the traffic class with the lowest priority. The name of the traffic class and its associated standard are used interchangeably throughout this document.

2.1. ARINC 664

ARINC 664, also known as AFDX® (Avionics Full DupleX Switched Ethernet), is an OSI layer 2 network protocol based on IEEE 802.3 Ethernet with additional Quality of Service capabilities that was developed for critical data communication in avionic systems. The standard ARINC 664-part-7 defines the requirements for an ARINC 664 network and its network switches [7]. AFDX® is an implementation of the ARINC 664 standard and patented by Airbus. [37]

2.1.1. History

ARINC 664 operates quite differently from regular IEEE 802.3 Ethernet. To understand why this is the case, it is helpful to have a look at its predecessors. ARINC 429 was the primary field bus used in commercial aircraft and is still used today. The data is transmitted digital, serial, unidirectional and differential with a data rate of up to 100kbit/s. As shown in Figure 2.1, one differential pair is required per transmitter to which up to a maximum of 19 receivers can be connected. This 1-to-n connection in the form of a bus topology is called a link and is only unidirectional. To send data back, for example from Line-Replaceable Unit (LRU) 2 to LRU 1, an additional differential pair is required [38].

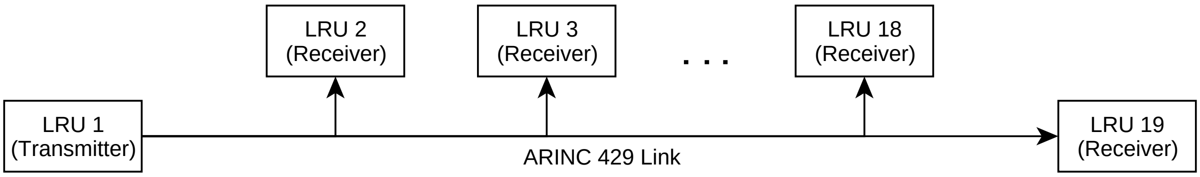

Figure 2.1.: Network topology of ARINC 429 [6]

The advantage of this setup is that no collisions of data packets can occur on the bus and the maximum bandwidth is always available with constant latency. These three factors are key characteristics necessary for a real-time safety-critical communication system. The disadvantage of this setup is an increased cabling effort, which has a negative effect on weight.

2.1.2. Operational priniciple

ARINC 664 was developed to save weight and costs and to enable higher data rates. The IEEE 802.3 Ethernet physical layer is used as a basis. Although ARINC 664 uses a star architecture and ARINC 429 uses a bus architecture, the safety features of ARINC 429 could be preserved by the applied changes to the OSI layer 2 protocol of Ethernet. As with Ethernet, the star node is implemented via an OSI layer 2 switch and the network devices connected to it are referred to as end systems (ES).

The physical communication to the end systems is bidirectional, as is the case with standard Ethernet, but the physical unidirectional links known from ARINC 429 have been replaced by logical links called virtual links (VLs). A VL defines a unidirectional path from a sending end system to one or more receiving end systems (1-to-n relationship). Each VL has a unique VL identifier (ID). The destination MAC address field is used to encode the VL ID. One part of the destination MAC address field consists of the CT marker (Critical Traffic) and the last two bytes of the MAC address are used for the VL ID. All VLs have the same CT marker. Instead of requiring one cable pair per outgoing connection, as necessary with ARINC 429, there is only one cable pair per ES, and on top of that are multiple logical (virtual) connections. This mechanism provides a logical segmentation of the network.

To prevent the VLs from influencing each other due to the multiple assignment of several VLs on one network cable, a maximum bandwidth limit can be configured for each VL. The configured bandwidth limit per VL combined must not exceed the maximum bandwidth of the physical connection. Although this prevents the available bandwidth from getting overbooked, it does not ensure a fixed latency for every transmission. If two data packets arrive at the switch via two separate physical ports but are forwarded to one and the same port, one of the two data packets must be buffered in the switch memory and therefore arrives later at the end system. Therefore an ARINC 664 network switch requires more memory space than a time-triggered only switch. Nevertheless, the system operates deterministic, because the maximum latency difference (jitter) can be calculated.

The bandwidth limit of a virtual link is defined by the following parameters:

- Maximum frame size
- Minimum time between two frames, the so-called BAG (Bandwith Allocation Gap)
- Bandwidth limit per unit of time

The User Datagram Protocol (UDP) is primarily used at the transport layer because virtual links are unidirectional by design and real-time capabilities are important in an avionics network. It is possible to use Transmission Control Protocol (TCP) by configuring two VLs, one in each direction.

Redundancy is used to minimize the probability of failure of the communication system. The network switches and cabling are redundant and the end systems are connected to all these redundant networks. The redundant networks are called channel A, channel B and so on. An end system transmits its data in parallel on both channels. Every ARINC 664 compliant RC network frame contains a frame counter. The ES uses this frame counter to identify redundant frames. Frame deduplication works as follows. The first frame that arrives at the receiver and is also valid is forwarded to the host, all other frames with the same frame counter get discarded. The frame counter value can be reused after a configurable time period called *SkewMax* [7].

2.2. Time-Triggered Ethernet

TTEthernet is an OSI layer 2 network protocol that enhances IEEE 802.3 Ethernet with additional Quality of Service capabilities. TTEthernet supports synchronous time-triggered communication. A TTEthernet network consists of time-triggered capable network switches and time-triggered end systems (in IEEE 802.3 known as *Data Terminal Equipment*). TTEthernet also uses VLs to segment the time-triggered traffic.

TTEthernet capable switches can segment the network not only by space (physical ports, virtual links) but also by time. Every time-triggered capable device stores a pre-calculated timetable, called a network schedule, locally. This schedule is created in the design phase of a TTEthernet network. The schedule contains information on when time-triggered frames are planned to be sent and received. The TTEthernet devices synchronize their local times over the network to establish a common time and align their schedule to this time. The standard AS6802A specifies a fault-tolerant synchronization protocol needed for time-triggered messaging and partitioning of mixed critical traffic on the same physical network [9].

2.2.1. Clock synchronization

The TTEthernet clock synchronization service provides a mechanism to establish a common time between all TTEthernet network devices, called global time. The global time is only a relative time and not an absolute time such as the Coordinated Universal Time (UTC). Synchronization messages are used to establish and maintain this global time. The synchronization messages are called Protocol Control Frames (PCFs). The synchronization is necessary because, without a common global time, the TTEthernet network devices would not be able to align the dispatch time of their messages with the predefined schedule.

A TTEthernet network contains three different types of synchronization devices:

- **Synchronization Master** - An end system that periodically sends PCFs with its local time to the compression master. One TTEthernet network contains several synchronization masters (SM), this is necessary to provide redundancy against faults.
- **Crompression Master** - A network switch that receives the PCFs from the synchronization masters and calculates a fault-tolerant average time. This calculated time is then transmitted via PCF to all TTEthernet network devices to establish the global time. A least one compression master (CM) is needed in the network.
- **Synchronization Client** - All other time-triggered network devices that are not SMs and CMs are synchronization clients (SC). These devices will never create PCFs. They only receive PCFs from the compression master to synchronize their local time with the received global time. A network switch that acts as a synchronization client will relay PCFs between synchronization masters, compression masters and synchronization clients.

The payload of a PCF contains the information shown in Figure 2.2.

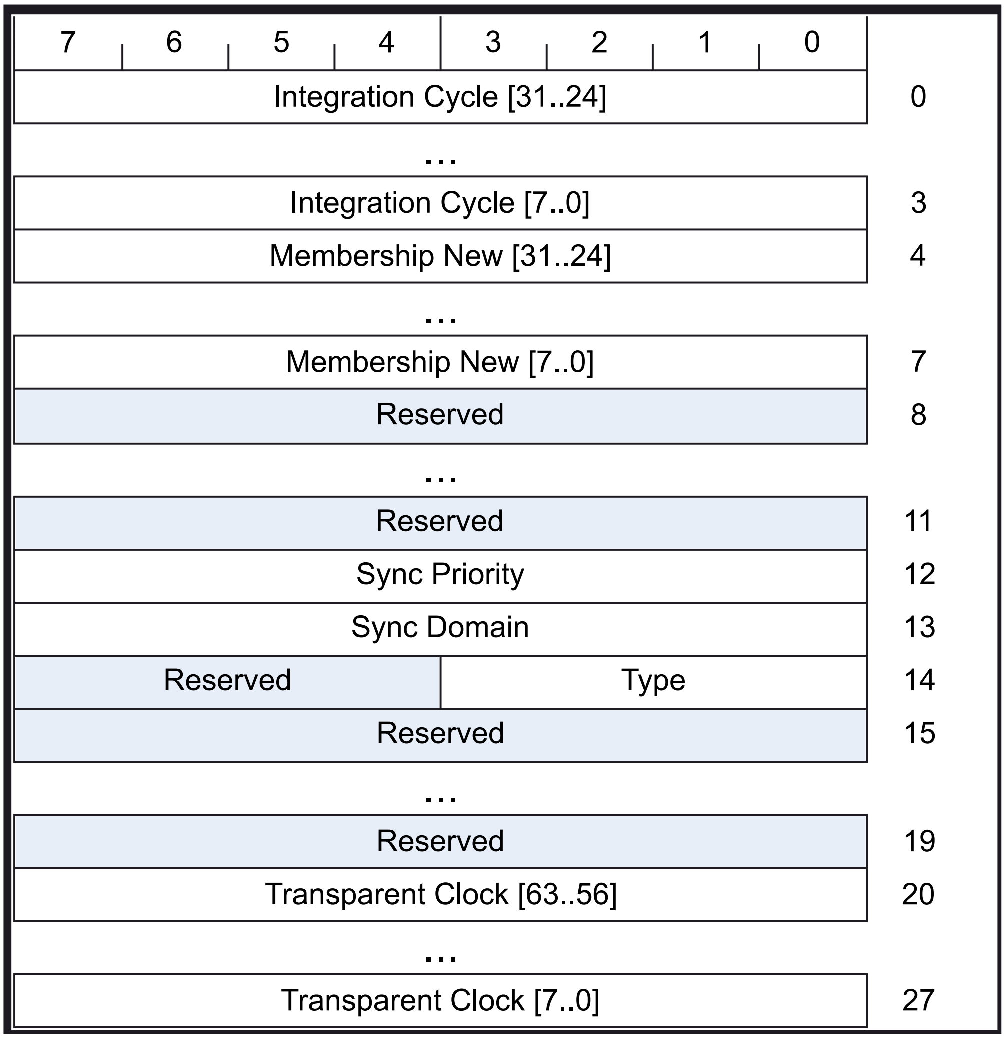

7	6	5	4	3	2	1	0	
Integration Cycle [31..24]								0
...								
Integration Cycle [7..0]								3
Membership New [31..24]								4
...								
Membership New [7..0]								7
Reserved								8
...								
Reserved								11
Sync Priority								12
Sync Domain								13
Reserved				Type				14
Reserved								15
...								
Reserved								19
Transparent Clock [63..56]								20
...								
Transparent Clock [7..0]								27

Figure 2.2.: PCF format [9]

There are three different frame types of PCFs indicated by the *Type* field in the payload:

1. **Coldstart PCF** - is sent by an SM, if the device is in the unsynchronized state and no synchronized network was detected.
2. **Coldstart acknowledgment (ACK) PCF** - is sent by a CM as a response to a received Coldstart PCF from an SM.
3. **Integration (IN) PCF** - is sent periodically by an SM that is in the synchronized state to the CM. The CM compresses the PCF IN frame and distributes the compressed PCF IN frame to all time-triggered devices in the network.

A synchronized clock between time-triggered devices is a necessity for TT traffic, PCFs therefore have a higher priority in the network than TT frames. A TTEthernet network distinguishes the following four different traffic types that are numbered by their priority:

1. **PCF** - Protocol Contol Frame messages
2. **TT** - Time-Triggered traffic
3. **RC** - Rate-Constrained traffic, see section 2.1
4. **BE** - Best-Effort traffic

2.2.2. Fault tolerance

TTEthernet uses redundant networks, called channels, to increase fault tolerance. Each network channel has its own switches and cabling. The end systems are connected to all channels. The clock synchronization mechanism establishes one global time over all channels. The AS6802A standard establishes two failure hypotheses that define how many synchronization devices are needed. The so-called *single-failure hypothesis* requires at least two CMs operating on redundant channels and four SMs. In this scenario, the clock synchronization mechanism can tolerate one rogue SM or the loss of one CM. A single rogue SM can send arbitrary PCFs to the CM without affecting the clock synchronization but a network switch must never send arbitrary PCFs. The second hypothesis is called *dual-failure hypothesis* and requires at least three CMs operating on redundant channels and five SMs, but this scenario requires high-integrity synchronization masters. Market research with leading suppliers, such as TTTech Computertechnik AG [39], has shown that unfortunately no commercial-of-the-shelf (COTS), high-integrity SMs are currently available on the market.

3. Related Work

Plenty of scientific research exists on the security of regular Ethernet networks and their well-known weaknesses. Kiravuo *et al.* [40] provide a very complete survey about Ethernet Local Area Network (LAN) security. It correctly identifies that Ethernet's greatest strength of zero-configurability is also its greatest security weakness because the self-configuration features can be misused easily. Lang [41] analyses security threats of Ethernet systems used in automotive applications and lists potential security measures. The main focus of the thesis is to analyze the impact of integrity and authenticity security measures on network latency in order to find a balance between security and safety. Considerably less research can be found on the security of rate-constrained and time-triggered communication.

A clock synchronization protocol frequently used in Ethernet networks is IEEE 1588 [42]. Existing research on the security of IEEE 1588, such as [43], may provide valuable input for the security analysis of TTEthernet. Mizrahi describes how IPsec [45] and MACsec [45] can improve the security of IEEE 1588 in [44]. In [46] more than 30 potential security issues and threats of IEEE 802.1 TSN protocols have been identified and categorized with STRIDE. TSN is a competing standard for time-triggered communication. Although the clock synchronization mechanism of TSN is different from the one used by TTEthernet, some of the identified threats are relevant for all systems with time-triggered communication.

Steiner [28] analyzed security threats to the TTEthernet clock synchronization protocol and discussed possible mitigations but only in the context of the failure hypothesis defined by AS6802. Loveless *et al.* [35] published the first attack on the clock synchronization that breaks the isolation guarantees of TTEthernet. In [47] a platform-independent security measure was developed on top of the existing clock synchronization algorithm of TTEthernet.

No research exists that covers the scope of this thesis of comparing the security of Ethernet and ARINC 664 with TTEthernet but several papers are related to individual security aspects of TTEthernet.

4. Approach

4.1. Scope

TTEthernet is a layer 2 network protocol of the Open Systems Interconnection model (OSI). Therefore, the scope of this analysis is threats on OSI layer 2 that can have an effect on TTEthernet. Although the main focus is on OSI layer 2, it has to be noted that upper OSI layers can be reliant on the OSI layer below. In this case, the layer below is the physical layer. Attacks on weaknesses of the TTEthernet protocol itself that are facilitated by attacks on other OSI layers are in the scope of this analysis. Weaknesses in OSI layer protocols and services like the absence of secure authentication in SNMPv1 can have an indirect effect on TTEthernet network devices. However, this is not a deficiency of the OSI layer 2 protocol that is analyzed and is therefore out of scope. Attacks on the logical level like an attack on weaknesses of the protocol are in scope. Attacks on the physical level are in scope but attacks on the social level like social engineering attacks are out of the scope of this analysis. Other competing layer 2 network protocols like TSN are out of scope. TTEthernet has several failure hypotheses that require different equipment, in this thesis only the single-failure hypothesis is part of the scope. Identifying the risk of the threat is out of scope because the risk is very dependent on the specific system and its application.

The assessment only covers security threats of the operational phase and not the design, manufacturing or decommission phase because these threats are tied to a secure development lifecycle process and are not TTEthernet specific.

Security measures can be improvements of the TTEthernet protocol itself but also measures implemented on the switches and end systems of the network as well as choosing suitable configuration settings and parameters. Implementing or testing the suggested security improvements listed in chapter 6 is out of the scope of this work.

4.2. Methodology

This work is split into two parts. In the first part a security analysis of TTEthernet is performed, see chapter 5. The goal of the security analysis is to identify threats and their impact on TTEthernet. TTEthernet not only supports time-triggered traffic but can be used alongside other traffic classes like best-effort traffic (regular Ethernet traffic) and rate-constraint traffic. The security analysis will identify the security strengths and weaknesses of time-triggered traffic in comparison to best-effort and rate-constraint traffic. To identify possible threats and threat scenarios literature research on known layer 2 threats and thread modeling for TTEthernet specific threats will be performed. Attack libraries like MITRE's Common Attack Pattern Enumeration and Classification (CAPEC) [48] and MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) [49] are good sources for well-known threats and attack techniques. The Space Attack Research & Tactic Analysis (SPARTA) created by the Aerospace Corporation is based on the MITRE ATT&CK matrix and provides a list of Tactic, Techniques, and Procedures (TTP) and threats that are specific to spacecrafts [50].

A structured approach to threat modeling will be used to identify threats that are unknown and specific to TTEthernet. The system that is analyzed is modeled with the help of a Data Flow Diagram (DFD). The DFD uses the following elements: *External Entities*, *Processes*, *Data Flows* and *Data Stores*. In addition, the diagram shows the *Trust Boundaries* in the system. Modeling the targeted system is essential to be able to analyze it. Figure 4.1 shows a depiction of the symbols used in the DFD. DFDs are well suited to model network systems and as an input for threat modeling [51].

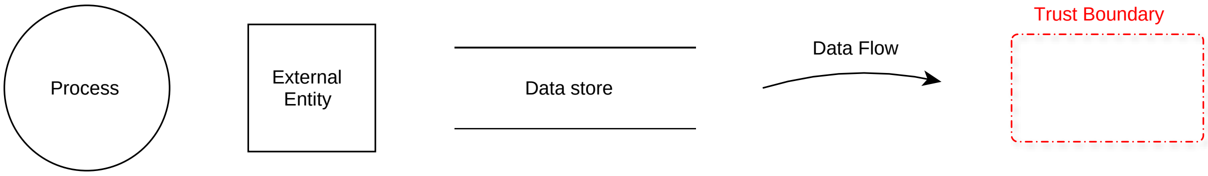

Figure 4.1.: Key to the symbols used in DFD

For this analysis, the STRIDE model is used to help find security threats within TTEthernet. Although STRIDE was developed at Microsoft with IT systems and web-based applications in mind, research showed that it can be effective for industrial control systems (ICS) [52] and embedded systems [53], [54] as well. The STRIDE model organizes threats into the following five categories: spoofing (S), tampering (T), repudiation (R), information disclosure (I), denial-of-service (D) and elevation-of-privilege (E). Not all categories apply to all the DFD elements equally. The STRIDE-per-Element method accounts for that by mapping STRIDE categories to the DFD elements. The mapping that was used for this analysis is shown in Table 4.1. This mapping is from [51].

DFD Element Type	**DFD Item ID**	**S**	**T**	**R**	**I**	**D**	**E**
External Entity	EE.x	X		X			
Process	P.x	X	X	X	X	X	X
Data Flow	DF.x		X		X	X	
Data Store	DS.x		X	[1]	X	X	

Table 4.1.: STRIDE-per-Element allocation [51]

The purpose of the second part in chapter 6 is to come up with security improvements for TTEthernet. Based on the security threats and weaknesses of TTEthernet that were identified in the security analysis, potential security measures to mitigate these weaknesses should be developed and described. Advantages and disadvantages as well as the threats that it addresses should be described for every security measure.

4.3. Security Environment

The ECSS standard [55] describes several network topologies: Single Channel network topology, Dual Channel network topology, Triple Channel network topology and Mixed network topology.

The difference between Single-, Dual-, and Triple-Channel is the redundance that the system provides. A Dual Channel network contains twice as many switches and electrical wiring as a Single Channel network. All channels are always active and transport frames to increase the redundancy. If one channel fails in a Dual Channel network, then there is still network traffic over the remaining channel. Every extra network channel adds an additional network to the network system. The amount of redundancy needed for a particular application depends on the reliability and safety requirements of that application. Due to the high effort (needs lots of resources, more cabling and more switches) multiple channels are only used to connect systems where redundancy is required. In real-world applications it is therefore common to have a Mixed network topology where parts of the network are built with redundancy and have multiple channels in parallel and other parts of the network system are only a Single Channel.

The Lunar Gateway uses a Mixed network topology and has no dedicated best-effort network. Best-effort network devices are connected to the same network as the devices that require high redundancy, but these

[1] In the literature repudiation is also defined as a threat to data stores if it is used to store logs. None of the data stores in a TTEthernet system is used for log data and therefore this threat category is not relevant in this case.

devices are connected to only one channel. Notice that not all best-effort devices must be connected to the same channel [17]. The network topology described in the International Avionics System Interoperability Standards for spacecraft is also based on a Mixed topology [27].

To reflect network topologies used in real-world applications, all analyses that are performed in this document will be based on the Mixed network topology that is defined in the ECSS standard. The network shown in Figure 4.2 consists of 3 network switches and 8 end systems. Two of the switches are configured as compression masters and the end systems ES1 to ES4 as synchronization masters. The third switch and the end systems ES5 to ES8 act as synchronization clients. The network consists of two channels, channel A and channel B. The end systems ES1 to ES5 are connected to both channels, therefore they have a redundant connection. To satisfy the independence requirement, there is no connection between the switches of different channels.

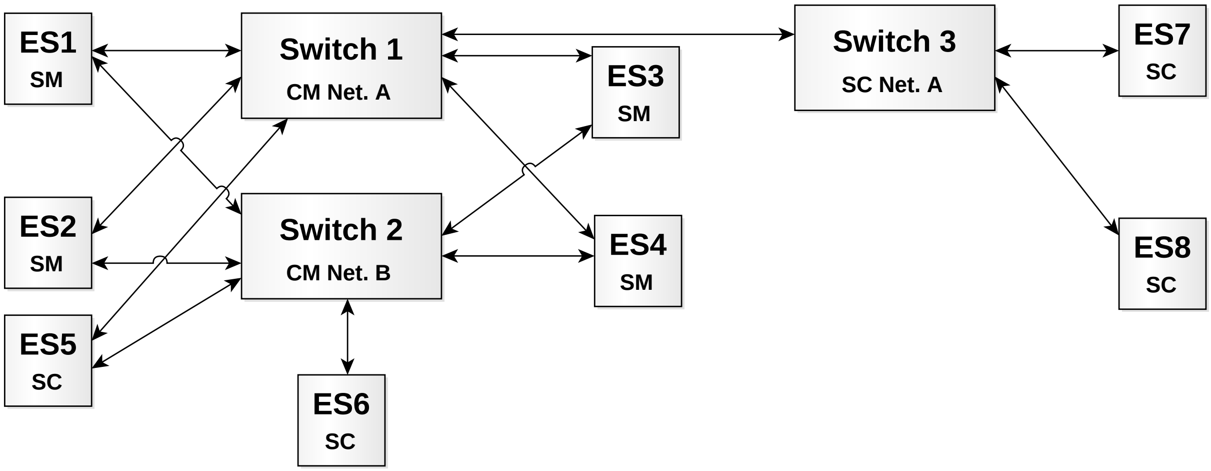

Figure 4.2.: Mixed network topology [55]

5. Threat Modeling

5.1. Identifying Assets

To be able to proceed with the security analysis and come up with security improvements, it is necessary to identify the security assets first. Security assets are the network functions that are critical for the operation of the network.

The primary asset of every computer network is to provide data communication within the network (A1). Critical security attributes for this function are the availability of the network to transmit and receive network frames and the integrity of the transmitted data. Note that in the information security context data integrity does not only refer to accidental changes of data, but also deliberate changes by an unauthorized entity. Confidentiality of the transmitted data can also be a security requirement, but this depends on the application. It is included in this analysis for the sake of completeness, but it has to be noted that confidentiality can also be addressed on upper layers, for example with IPsec. The asset A1 is applicable for all traffic classes. Enforcing traffic policing and frame filtering functionality (A1.1) is an asset of Time-Triggered (TT) and Rate-Constrained (RC) traffic. This asset is dependent on A1. Applicable security attributes for this asset are availability and integrity. A time-triggered network is reliant on a mechanism to synchronize the clocks of the network devices (A1.2). This asset is dependent on A1 because TTEthernet uses special data packets (PCFs) for clock synchronization that are transmitted over the same network as all the other network data. Table 5.1 lists the identified assets of a TTEthernet network.

ID	Typ	Asset	Traffic Class	Security Attribute
A1	Primary	Transmit network frames	BE/RC/TT	Availability, Integrity, Confidentiality
A1.1	Secondary	Enforce traffic policing and frame filtering	RC/TT	Availability, Integrity
A1.2	Secondary	Clock synchronization	TT	Availability, Integrity

Table 5.1.: Identified Security Assets

5.2. Data Flow Diagram

For threat modeling it is necessary to create a model of the system that is analyzed. In this work data flow diagrams (DFD) are used to model the TTEthernet network. The DFD is broken up into several diagrams that model the different use cases and assets that have been identified in section 5.1.

5.2.1. Best-effort communication

Figure 5.1 shows the communication with best-effort traffic. The diagram only shows one switch and one end system, but this is enough to display all interactions of the modeled system.

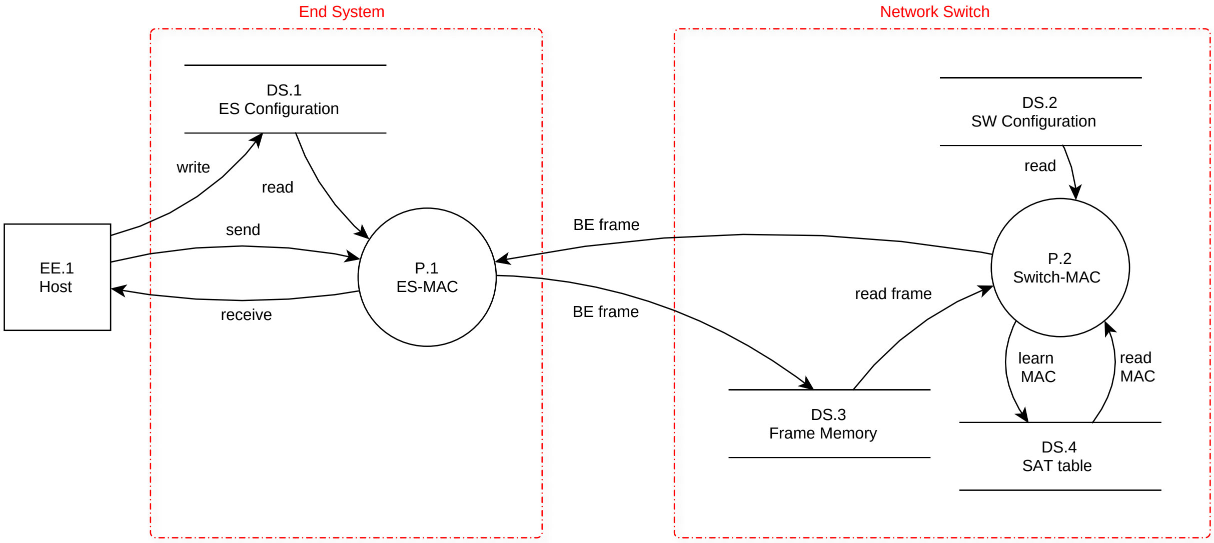

Figure 5.1.: DFD of best-effort traffic

The network switches and end systems (aka network interface cards) behave as in any regular Ethernet network. The *ES-MAC* transmits BE network frames that are sent by the *Host* to the *End System* immediately into the Ethernet network. The MAC address of the end system is stored in the *ES Configuration*. Every BE network frame that is received by the *ES-MAC* is forwarded to the *Host*. A *Host* typically has the capability to change the *ES Configuration* to be able to install updates of the configuration.

The *Switch-MAC* of the network switch forwards every valid Ethernet frame it receives to other network devices according to its *Source Address Table (SAT)*. The SAT contains a list of MAC addresses together with the corresponding switch port. These entries can be either static or dynamic. Static entries must be pre-configured and are stored in the *SW Configuration*. Every MAC address can only exist once in the table. If dynamic address learning is active, then the switch will learn and update MAC Addresses (plus switch port) that are not statically configured and store them in the SAT. Every network switch has a limited amount of frame memory. Received network frames are dropped if the *Frame Memory* is full.

5.2.2. Rate-constrained communication

Figure 5.2 shows the communication with rate-constrained traffic. The diagram only shows one channel with one switch and one end system, but this is enough to display all interactions of the modeled system.

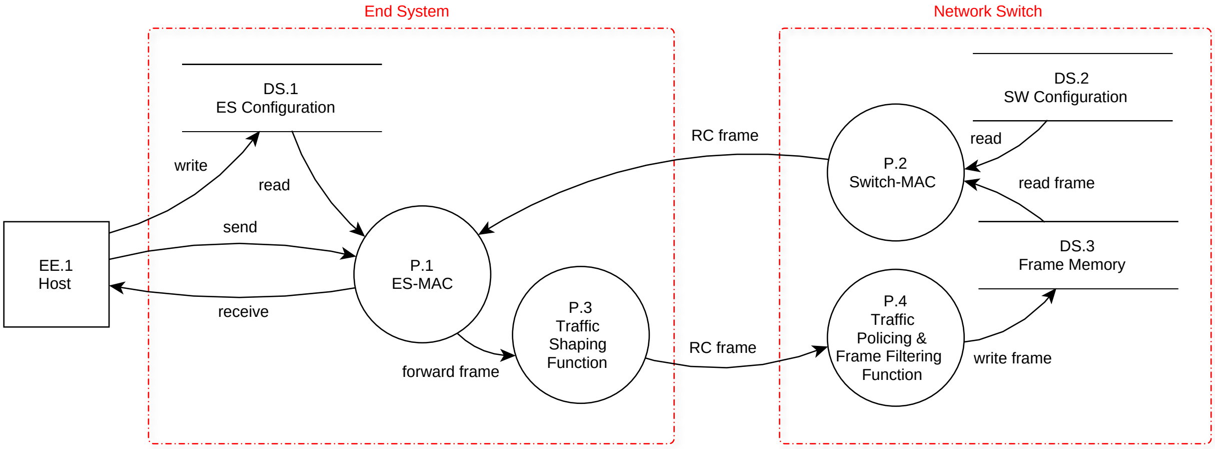

Figure 5.2.: DFD of rate-constrained traffic

The *ES-MAC* forwards RC network frames that are sent by the *Host* to the *Traffic Shaping Function*. The End System has a cache memory that can buffer a certain number of frames. This is important if the host sends a burst of frames in a short period. The purpose of this function is to ensure that the rate at which data is sent into the network does not exceed a configured maximum bandwidth per VL ID. This information is pre-defined in the design phase of the system and stored in the *ES Configuration*. A *Host* typically has the capability to change the *ES Configuration* to be able to install updates of the configuration. Every RC frame received by an RC capable network switch is verified by the *Traffic Policing and Frame Filtering Function* first. The *Traffic Policing and Frame Filtering Function* checks if the semantics and syntax of the received frames are correct. The function also ensures that only frames with a VL ID that are allowed at the switch port where the frames were received are forwarded to the *Switch-MAC*. RC frames that do not pass these checks are dropped. Network frames are also dropped if the received data rate exceeds the configured bandwidth limit. This limit is set per VL. The *Switch-MAC* uses a pre-configured table of all VLs to forward an RC frame based on its VL. This table and the configuration data needed for traffic policing and filtering are stored in the *SW Configuration*.

5.2.3. Time triggered communication

Figure 5.3 shows the communication with time-triggered traffic. The diagram only shows one channel with one switch and one end system, but this is enough to display all interactions of the modeled system.

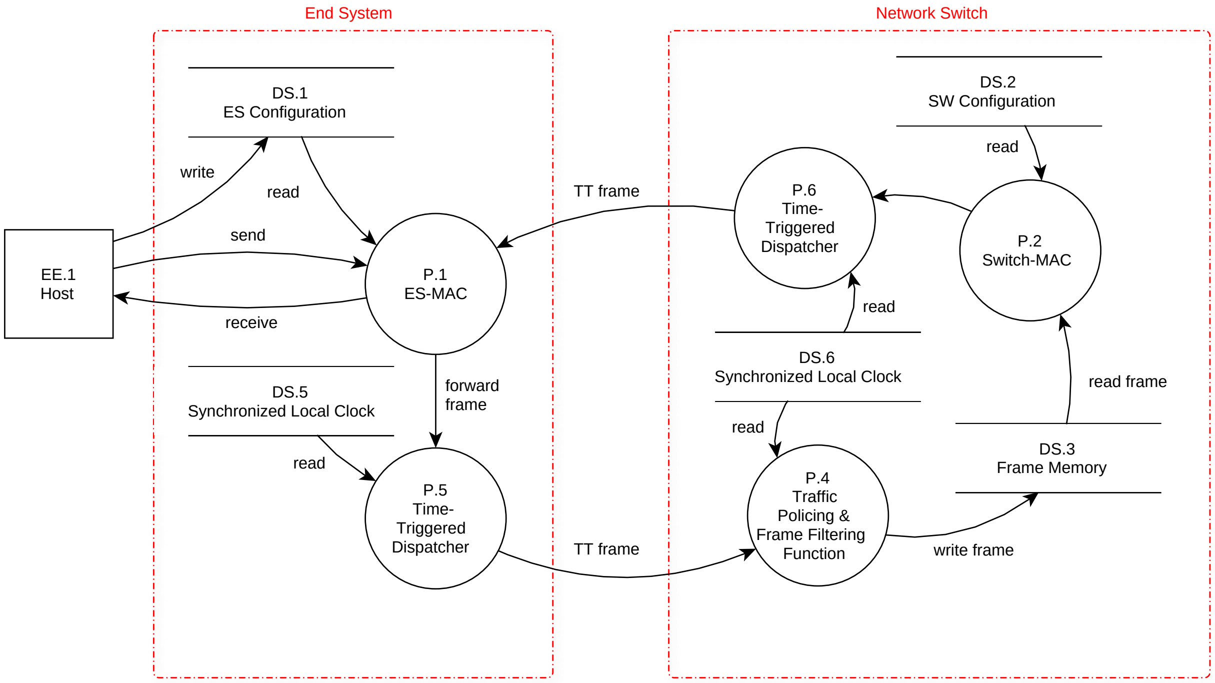

Figure 5.3.: DFD of time-triggered traffic

The *ES-MAC* forwards TT network frames that are sent by the *Host* to the *Time-Triggered Dispatcher*. This function dispatches the TT frames in time based on the *Synchronized Local Clock* and a pre-configured network schedule. Like with RC frames, VL IDs are also used for routing TT frames. The network schedule and VL IDs are stored in the *ES Configuration*. A *Host* typically has the capability to change the *ES Configuration* to be able to install updates of the configuration. Every TT frame that is received by the network switch gets processed by the *Traffic Policing and Frame Filtering Function*. This function enforces the time-triggered communication based on the network schedule and the *Synchronized Local Clock*. A TT frame that is received out of schedule gets dropped. The *Traffic Policing and Frame Filtering Function* also checks if the VL IDs, semantics and syntax of the received frames are correct. TT frames that are passing these checks are forwarded to the *Switch-MAC*, otherwise they are dropped. The *Switch-MAC* uses a pre-configured table of all VLs to forward a TT frame based on its VL. This table and the configuration data needed for traffic policing and filtering are stored in the *SW Configuration*. The *Time-Triggered Dispatcher* dispatches the TT frame in time, that was forwarded from the *Switch-MAC*. This works the same as in a time-triggered capable end system. All TT traffic is dependent on the clock synchronization function to establish a common time base in the network and to synchronize the local clocks of the network devices.

5.2.4. Clock synchronization mechanism

A TTEthernet network can be either in a synchronized or un-synchronized state. Because a data flow diagram is not a state diagram, both states must be modeled in their own diagram. Figure 5.4 shows the clock synchronization mechanism of a synchronized TTEthernet network. The local clocks of the synchronized network devices will drift over time. The purpose of the clock synchronization mechanism is to align the local clocks in the network and keep them aligned to provide a common time base within the network.

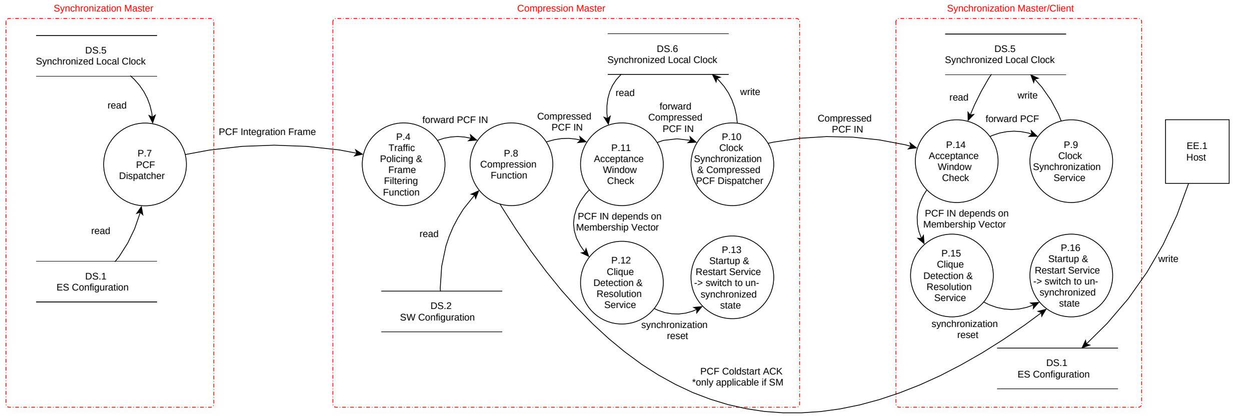

Figure 5.4.: DFD of clock synchronization mechanisms in synchronized operation

The end systems that act as synchronization masters periodically send PCF integration frames (IN) over the network to the switches that act as compression masters. The information of its content, when and how often a PCF IN frame is sent, is stored in the *ES Configuration* on the ES. The *Synchronized Local Clock* provides the point in time when the PCF must be dispatched by the *PCF Dispatcher*. Between the synchronization master and compression master can be several switches but this is not shown in the diagram. In addition, the diagram only shows one synchronization master (or synchronization client, respectively) and one channel of the network.

The PCF IN frame is received and processed by the *Traffic Policing & Frame Filtering Function* of the compression master. The *Traffic Policing & Frame Filtering Function* only accepts PCF IN frames on network ports and VLs respectively from a synchronization master. The information about who is a synchronization master is stored in the *SW Configuration*. PCFs that are not integration frames will be dropped. The *Compression Function* collects all the received PCF IN frames it received from the synchronization masters within a certain time period, but only one PCF per synchronization master, and calculates a fault-tolerant average value that is then used to establish a common time in the network. Static delays like wire delays or propagation delays from senders and receivers are also considered by the calculation. All static and pre-calculated parameters, such as the wire delay, are stored in the configuration data of the devices. The *ES Configuration* and *SW Configuration* are created by the network integrator in the design phase of the system. This calculated average value is then sent as a so-called compressed PCF IN frame to the next process. The compressed PCF also contains the membership vector with information about which and how many synchronization masters have contributed to the calculation.

The *Acceptance Window Check* ensures that only Compressed PCF IN frames that arrive within a certain time frame that is aligned with the *Synchronized Local Clock* can contribute to the clock synchronization process. The *Clock Synchronization Service and Compressed PCF Dispatcher* updates the value of the *Synchronized Local Clock* to keep the local clock synchronized with the common time base, and the compressed PCF IN frame is then forwarded to all synchronization masters and synchronization clients in the network. If compressed PCF IN frames with a membership vector that exceeds the preconfigured limit are received outside of this time frame, then the *Clique Detection and Resolution Service* will be triggered. This mechanism is called asynchronous clique detection. There is also a synchronous clique detection mechanism that is triggered when the number of synchronization masters that are aligned with the device falls below a preconfigured threshold. The purpose of the *Clique Detection and Resolution Service* is to reestablish synchronization on a network level. The compression master will switch into its un-synchronized state if a clique is detected.

Like the compression master, the synchronization master and synchronization client will use the Compressed PCF IN frame to re-synchronize its *Synchronized Local Clock*. The *Host* of the End System can also use the *Synchronized Local Clock* to synchronize its applications with the common time base of the network. A *Host* typically has also the capability to change the *ES Configuration*. Their *Acceptance Window Check*, *Clock Synchronization Service* and *Clique Detection and Resolution Service* work analogous to those in the compression master but with some differences in the synchronization master. One difference is that a synchronization master will switch to the un-synchronized state if it receives a PCF Coldstart Acknowledgement frame. The *Clique Detection and Resolution Service* of a synchronization master operates a third clique detection mechanism, called relative clique detection. It compares if the number of synchronized synchronization masters is equal to or below the number of unsynchronized synchronization masters in the network. If this is the case, it will switch into the un-synchronized state.

Figure 5.5 shows the clock synchronization mechanism of an un-synchronized TTEthernet network.

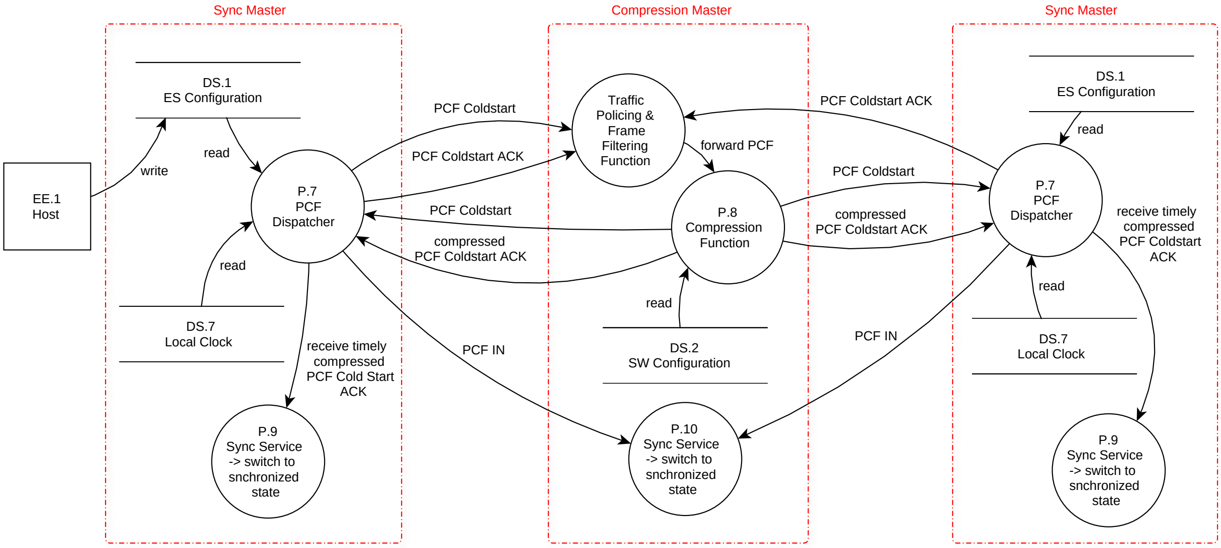

Figure 5.5.: DFD of clock synchronization mechanisms in un-synchronized operation

A synchronization master that is in the un-synchronized state (this is also its state after powering up the device) will listen for compressed PCF IN frames to integrate itself into an already synchronized network. If this is the case, it will switch to the synchronized state. Otherwise, the synchronization master will start with the so-called fault-tolerant handshake. The *PCF Dispatcher* sends a PCF Coldstart frame that contains its *Local Clock* to the compression master. The data necessary to create the PCFs is stored in the *ES Configuration*.

A compression master that is in its un-synchronized state will also listen for compressed PCF IN frames to integrate itself into an already synchronized network. All PCFs are received and processed by the *Traffic Policing & Frame Filtering Function* of the compression master. The *Traffic Policing & Frame Filtering Function* only accepts PCFs on network ports and VLs respectively from a synchronization master. The information about who is a synchronization master is stored in the *SW Configuration*. If there is none, it will wait for a PCF Coldstart frame transmitted by a synchronization master. Information about which synchronization masters should be present on the network is stored in the *SW Configuration*. When the *Compression Function* receives a PCF Coldstart frame it relays the frame back and forwards it to all other synchronization masters in the network.

A synchronization master that receives a PCF Coldstart frame while in its un-synchronized state responds by sending a PCF Coldstart Acknowledgement frame. A compression master that receives a PCF Coldstart Acknowledgement frame (while in its un-synchronized state) will compress the frame and send the compressed PCF Coldstart Acknowledgement frame to all synchronization masters. To compress the frame means that the *Compression Function* collects all PCFs it receives within a certain time span and generates one compressed frame. This completes the fault-tolerant handshake. After the synchronization master receives a compressed PCF Coldstart Acknowledgement frame it switches to the synchronized state by sending the first PCF IN frame. A compression master switches to the synchronized state when it receives a PCF IN frame with a sufficiently high number of bits set in the membership vector.

Note that synchronization clients are not shown in the DFD because a synchronization client that is in its unsynchronized state only waits for the reception of a compressed PCF IN frame. The synchronization client switches into the synchronized state if the received compressed PCF IN frame has a membership vector that exceeds the preconfigured limit.

5.3. Identifying Security Threats

For every security threat that was identified a table with the following information is provided:

- **ID** is a unique identifier in the form of: T. + a letter for the STRIDE category + a consecutive number
- **Preconditions** that have to be fulfilled for the threat to materialize.
- **Threat Description** provides a brief summary of the threat.
- **Impact on the Element** is the local effect on the DFD element that is analyzed.
- **Impact on Asset** that is affected by this threat.

The following list provides a definition for the attributes that have been used to further categorize and describe the effect on an element:

- **Temporary**: The effect on the element stops with the attack or the system can recover automatically.
- **Persistent**: The effect on the element is persistent even after the attack has stopped.
- **Destructive**: The element is physically destroyed by the attack.
- **Recoverable**: Normal operation can be restored but extra steps are needed, like performing a reset.
- **Non-recoverable**: Maintenance work or physical access is necessary to restore operation.

In general, an attacker with full physical access could have a direct impact on the availability of every element of the system, such as denial-of-service (DoS). The effect could be persistent, or even destructive to the element and might therefore be non-recoverable. This cannot be mitigated within the system but requires organizational security measures, like for example physical access control. This type of threat and security measure is out of scope and was therefore not analyzed. Neither the Ethernet standard nor the ARINC 664 standard nor the TTEthernet standard requires the system to store logs. In theory, this could be a threat in the category of repudiation, but no case could be found where an attack could exploit this.

5.3.1. Best-Effort Communication

Threats on best-effort communication affect the asset A1.

5.3.1.1. DS.1 - ES Configuration

Table 5.2 lists threats of the data store DS.1 that stores the ES configuration. An attacker who is in control of the host can reconfigure the ES by writing its configuration into the registers of the ES (T.T1). This is an attack on the integrity of the ES configuration. An attacker can also erase the content of the ES configuration registers (T.D1). The ES configuration is used by the ES-MAC to set for example the network speed of the Ethernet physical layer (PHY). Without this configuration or an incorrect configuration present (for example a wrong Ethernet speed), the ES is not operational and is not able to send or receive network frames. Therefore, the threats T.T1 and T.D1 can have an effect on the availability of P.1 and then in turn on the availability of asset A1. These attacks are persistent until a correct ES configuration is written to recover the content of the ES registers. The ES configuration does not contain confidential information, therefore no threat could be identified in the category of information disclosure.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T1	Host access	An attacker with control of the host can change the ES configuration	Integrity, persistent, recoverable	Availability
T.D1	Host access	An attacker with control of the host can erase the ES configuration	Availability, persistent, recoverable	Availability

Table 5.2.: Identified threats of data store DS.1 ES configuration for BE communication

5.3.1.2. DS.2 - SW Configuration

Table 5.3 lists threats of the data store DS.2 that stores the SW configuration. An attacker with management access to the network switch can change the SW configuration (T.T2). This is a threat to the integrity of the SW configuration. The SW configuration contains for example the static a priori configured entries of the SAT. This can be used to intercept and re-route network traffic and affects the Switch-MAC (P.2) and then in turn on the confidentiality, integrity and availability of asset A1. An attacker can also erase the content of the SW configuration registers (T.D2). Without this configuration or an incorrect configuration present, the network switch is not operational and affects the availability of the switch. These attacks are persistent until a correct SW configuration is written to recover the content of the SW registers, but the network cannot be used to upload the configuration because the networking part of the switch is not operational in this scenario. Another threat like T.D2 but with a severe impact on recoverability can occur if the storage medium has a finite number of erase cycles, such as flash memory. An attacker might try to wear out the flash memory by performing excessive erase cycles (T.D3). This is a destructive attack that might be executed via remote access and is not recoverable without replacing the damaged component. Both threats (T.D2 and T.D3) can affect the availability of a network switch and are therefore a threat to the availability of asset A1. The SW configuration does not contain confidential information, therefore no threat could be identified in the category of information disclosure.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T2	SW management access	An attacker with management access might be able to change the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.D2	SW management access	An attacker with management access can delete the SW configuration	Availability, persistent, non-recoverable	Availability
T.D3	SW management access, data stored on flash memory	An attacker with management access can wear out and physically destroy the flash memory	Availability, persistent, non-recoverable, destructive	Availability

Table 5.3.: Identified threats of data store DS.2 SW configuration for BE communication

5.3.1.3. DS.3 - Frame Memory

Table 5.4 lists threats of the data store DS.3 that stores the network frames. With management access to the network switch, an attacker might be able to read (T.I1) or manipulate (T.T3) the content of the frame memory. But this is dependent on the actual implementation of the switch. It is likely that the management Central Processing Unit (CPU) inside the switch does not have direct access to the frame memory because this access is not necessary for the operation of a network switch. Although switch ASICs exist that have this from a security point of view undesired capability [56]. With management access to the application CPU easier attacks to manipulate, delete and read network frames exist, such as changing the SAT, as described in section 5.3.1.4, to forward or drop incoming network frames.

Another threat to the frame memory is the exhaustion of memory space via a denial-of-service attack (T.D4). This attack requires that the attacker has access to the network and can send a large amount of network frames within a short period of time. Depending on the network topology, the bandwidth of frame memory and network ports, it might be possible for an attacker to fill up the frame memory with frames from the attacker. This could cause increased delays, jitter and frame drops for genuine network traffic. All threats listed in this paragraph have a temporary effect on the assets and they will stop when the attack stops. These threats can have a direct impact on integrity (T.T3), confidentiality (T.I1) or availability (T.D4) of network frames and therefore on asset A1.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T3	SW management access	An attacker with management access might be able to manipulate the frame memory	Integrity, temporary	Integrity
T.I1	SW management access	An attacker with management access might be able to read the frame memory	Confidentiality, temporary	Confidentiality
T.D4	Network access	An attacker with network access can try to exhaust the memory space	Availability, temporary	Availability

Table 5.4.: Identified threats of data store DS.3 Frame Memory for BE communication

5.3.1.4. DS.4 - SAT

Table 5.5 lists threats of the data store DS.4 that stores the source address table (SAT). An attacker can manipulate the entries in the SAT by sending network frames with spoofed source MAC addresses (T.T4). This is a well-known threat for network switches with dynamic MAC address learning enabled. This is a threat to the integrity of the SAT and can be used to intercept and re-route network traffic to for example launch an adversary-in-the-middle attack. A detailed description of MAC address spoofing can be found in Table A.1. Another possible threat that also exploits dynamic MAC address learning is MAC flooding (T.D5). This attack is described in detail in Table A.2. The effect of an SAT that overflows with entries depends on the actual implementation. In a typical implementation, a network switch falls back into hub mode. This means every frame with a new MAC received by a switch with a full SAT is broadcasted on all ports. This can have a negative effect on network bandwidth, availability and confidentiality of network frames and therefore asset A1 because the frames are forwarded to all network devices. Both threats can only have a temporary effect because dynamically learned MAC addresses are deleted after the MAC addresses age out over time if the attacker stops sending frames that spoof or flood the SAT.

An attacker with management access to the network switch can manipulate the entries in the SAT (T.T5) by changing the switch configuration as shown in threat T.T2. This attack is persistent until a working SAT is loaded, but a recovery can be difficult or not possible if it is performed via the network. Threats T.T4 and T.T5 can be used to intercept and re-route network traffic and affect the confidentiality, integrity and availability of asset A1. The SAT does not contain confidential information, therefore no threat could be identified in the category of information disclosure.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T4	Network access, dynamic MAC address learning	An attacker can perform a MAC address spoofing attack to add or change entries in the SAT	Integrity, temporary	Confidentiality, Integrity, Availability
T.T5	SW management access	An attacker with management access can add and change entries in the SAT	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.D5	Network access, dynamic MAC address learning	An attacker can perform a MAC-Flooding attack to overflow the SAT	Availability, temporary	Confidentiality, Availability

Table 5.5.: Identified threats of data store DS.4 SAT for BE communication

5.3.1.5. P.1 - ES-MAC

Table 5.6 lists threats of the process P.1 ES-MAC that provides the main functionality of the end system. An attacker can send network frames with a spoofed source MAC address to the end system (T.S1). This is possible because Ethernet does not offer any authenticity guarantees for its network traffic. This attack can for example be used to spoof the Address Resolution Protocol (ARP). The direct effect of this threat is temporary, but it can be used to trigger further threats with longer lasting consequences, therefore it can have an impact on the confidentiality, integrity and availability of asset A1. The process ES-MAC is dependent on settings stored in the ES configuration. Without or with incorrect configuration, the ES is not ready for operation. Tampering (T.T6) or deleting (T.D7) the ES configuration (see T.T1 and T.D1) can be a threat to P.1 and can affect for example the integrity and availability of the ES-MAC. This threat is persistent until the ES-MAC is correctly reconfigured.

An attacker with network access might be able to overwhelm the network input of the host (T.D6). This is only a threat if the host does not have enough resources to process the incoming network data. This is not an unrealistic concern for embedded systems, because a simple network sensor or actuator does not need the processing power to send data at line speed in order to perform its task. This can lead to dropped network frames if the host is not fast enough to empty the receive buffer. Another threat that can have an effect on the availability is if the host holds the ES-MAC in reset (T.D8). The effect of the last two threats is only temporary and will stop when the attack stops. All these threats can interrupt the process and affect the ability to send and receive network frames and therefore the availability of asset A1.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.S1	Network access	Attacker spoofs source MAC address of received network frame	Authenticity, temporary	Confidentiality, Integrity, Availability
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Availability
T.D6	Network access	An attacker might be able to overwhelm the input of the host	Availability, temporary	Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability
T.D8	Host access	An attacker with host access can hold the ES in reset	Availability, temporary	Availability

Table 5.6.: Identified threats of process P.1 ES-MAC for BE communication

5.3.1.6. P.2 - Switch-MAC

Table 5.7 lists threats of the process P.2 Switch-MAC that provides the main functionality of the switch. If spanning tree is used to automatically configure the network topology, an attacker can send Bridge Protocol Data Units (BPDU) messages to spoof the spanning tree mechanism inside the SW (T.S2). This threat can then be used by an attacker to force constant topology changes leading to a denial-of-service (DoS) attack on the network communication. This threat might also be used to perform an adversary-in-the-middle attack on network frames by establishing the device controlled by the attacker as the new root bridge. Therefore, this threat can then in turn have a negative effect on the availability, confidentiality and integrity of network frames and asset A1. The effect of threat T.S2 is temporary and will stop when the attack stops.

The process Switch-MAC is dependent on settings stored in the SW configuration. Without or with incorrect configuration, the SW is not ready for operation. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can be a threat to P.2 and can affect the integrity and availability of the Switch-MAC. A tampered process P.2 can be used to perform an adversary-in-the-middle attack by rerouting the network frames. This would compromise the confidentiality (T.I2) and integrity (T.T7) of network frames and asset A1. These three threats are persistent until the Switch-MAC is correctly reconfigured. Recovering a switch via the network with a defective or missing configuration can be difficult or is not possible. A non-working switch affects the availability of asset A1.

Another threat can be a lack of network segmentation because virtual local area networks (VLANs) are not used or the VLAN has too many peers (T.E1). In this case, an attacker only needs network access to be able to send arbitrary network frames to any network device. This includes even network devices that do not need to communicate with each other to perform their tasks. This violates core security principles, such as least privilege or minimizing the attack surface. This threat is temporary and will stop when the attack stops. A lack of effective network segmentation can impact the integrity and availability of asset A1.

The bypassing of network segmentation is another threat (T.E2). An attacker can try to inject network frames into privileged network segments to gain unauthorized access. This might be possible via a packet-in-packet attack if the switch is susceptible to this type of attack. A detailed description of packet-in-packet attacks, like EtherOops, can be found in Table A.3. This can enable the attacker to inject network frames into other VLANs or parts of the network that are otherwise not accessible by the attacker because for example due to static SAT. This threat is temporary and will stop when the attack stops. Injecting network frames impacts the integrity of asset A1.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.S2	Network access, spanning tree	An attacker can send BPDU messages to spoof the spanning tree protocol	Authenticity, temporary	Confidentiality, Integrity, Availability
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.I2	SW management access	An attacker can perform an adversary-in-the-middle attack by rerouting the network frames	Confidentiality, persistent, recoverable	Confidentiality, Integrity
T.D9	SW management access	An attacker with SW management access can disable the switch	Availability, persistent, recoverable	Availability
T.E1	Network access, no VLANs	Attacker can send arbitrary network frames to any network device	Authorization, temporary	Integrity, Availability
T.E2	Network access, EMI	Attacker can perform a packet-in-packet attack to inject network frames	Authorization, temporary	Integrity

Table 5.7.: Identified threats of process P.2 Switch-MAC for BE communication

5.3.1.7. EE.1 - Host

Table 5.8 lists threats of the external entity EE.1 that acts as a host for the ES. An attacker can send ARP messages to poison the ARP cache of the host. This is possible if ARP is used to discover the MAC address of unknown hosts because this protocol does not provide any authentication. A detailed description of ARP spoofing can be found in Table A.4. This can be used to intercept and re-route network traffic to affect the confidentiality, integrity and availability of asset A1. This threat is temporary and will stop when the attack stops.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.S3	Network access, ARP is used	An attacker can send ARP messages to spoof the ARP cache of the host	Authenticity, temporary	Confidentiality, Integrity, Availability

Table 5.8.: Identified threats of external entity EE.1 Host for BE communication

5.3.1.8. DF.1 - BE Frame

Table 5.9 lists threats of data flow DF.1 BE frame from P.1 ⇒ DS.3 and P.2 ⇒ P.1 that represents the network traffic (DF.1) sent between hosts and switches in the network. Note that the analysis in this section and all analyses of the data flow element in the following sections only assess direct threats to this connection. This means that in this case only threats to BE frames that originate at the connection between end systems and switches will be analyzed. BE frames can of course be manipulated or sniffed before or after they are transmitted, but this is already analyzed in the respective elements.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T8	Physical access to network wiring	An attacker can modify network frames	Integrity, temporary	Integrity
T.T9	Physical access to network wiring	An attacker can create new network frames	Integrity, temporary	Integrity
T.I3	Physical access to network wiring	An attacker can sniff network frames	Confidentiality, temporary	Confidentiality
T.D10	Physical access to network wiring	An attacker can drop network frames	Availability, temporary	Availability

Table 5.9.: Identified threats of data flow DF.1 BE frame for BE communication

An attacker with physical access to the cable between switches and end systems can modify (T.T8), create new (T.T9), sniff (T.I3) or drop (T.D10) network frames and therefore affect confidentiality, integrity and availability of asset A1. These threats are temporary and will stop if the attacker or the system of the attacker loses physical access to the network wiring.

5.3.1.9. DF.2 - write ES Configuration

Table 5.10 lists threats of data flow DF.2 from EE.1 ⇒ DS.1. The host writes the configuration of the end system into the ES configuration memory via the host interface after a reset. An attacker with physical access to the host interface might be able to modify (T.T10) or drop (T.D11) part or all configuration data that is transmitted to the end system. These threats have a direct effect on the ES configuration and impact the availability of asset A1 as described in section 5.3.1.1 DS.1. The ES configuration does not contain confidential information, therefore no threat could be identified in the category of information disclosure. These threats are temporary and will stop if the attacker loses physical access to the host interface.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T10	Physical access to host interface	An attacker can modify the ES configuration data on the fly	Integrity, temporary	Availability
T.D11	Physical access to host interface	An attacker can drop ES configuration data	Availability, temporary	Availability

Table 5.10.: Identified threats of data flow DF.2 write ES Configuration for BE communication

5.3.1.10. DF.3 - send and receive BE Frames to host

Table 5.11 lists threats of data flow DF.3 BE frame from EE.1 ⇒ DS.3 and back to the host. The data contains the network traffic that is received by the end system and forwarded to the host or that is sent from the host to the network. An attacker with physical access to this interface could be able to (T.T11), create new (T.T12), sniff (T.I4) or drop (T.D12) some or all network frames that are processed by this end system and therefore affect confidentiality, integrity and availability of asset A1. These threats are temporary and will stop if the attacker loses physical access to the host interface.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T11	Physical access to host interface	An attacker can modify network frames	Integrity, temporary	Integrity
T.T12	Physical access to host interface	An attacker can create new network frames	Integrity, temporary	Integrity
T.I4	Physical access to host interface	An attacker can sniff network frames	Confidentiality, temporary	Confidentiality
T.D12	Physical access to host interface	An attacker can drop network frames	Availability, temporary	Availability

Table 5.11.: Identified threats of data flow DF.3 for BE communication

5.3.2. Rate-Constrained Communication

5.3.2.1. DS.1 - ES Configuration

Table 5.12 lists threats of the data store DS.1 that stores the ES configuration. The identified threats and impacts are identical to the threats of DS.1 for BE communication described in section 5.3.1.1.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T1	Host access	An attacker with control of the host can change the ES configuration	Integrity, persistent, recoverable	Availability
T.D1	Host access	An attacker with control of the host can erase the ES configuration	Availability, persistent, recoverable	Availability

Table 5.12.: Identified threats of data store DS.1 ES configuration for RC communication

5.3.2.2. DS.2 - SW Configuration

Table 5.13 lists threats of the data store DS.2 that stores the SW configuration. The identified threats and impacts are identical to the threats of DS.2 for BE communication described in section 5.3.1.2.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T2	SW management access	An attacker with management access might be able to change the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.D2	SW management access	An attacker with management access can delete the SW configuration	Availability, persistent, non-recoverable	Availability
T.D3	SW management access, data stored on flash memory	An attacker with management access can wear out and physically destroy the flash memory	Availability, persistent, non-recoverable, destructive	Availability

Table 5.13.: Identified threats of data store DS.2 SW configuration for RC communication

5.3.2.3. DS.3 - Frame Memory

Table 5.14 lists threats of the data store DS.3 that stores the ES configuration. The identified threats T.T3 and T.I1 and their impacts are identical to the threats of DS.3 for BE communication described in section 5.3.1.3. It should be noted, that the threat T.D4 exhaust memory space from BE communication does not apply for RC traffic if the bandwidth limit that is needed to constrain the frame rate is correctly calculated, configured and executed by process P.4. The frame memory is therefore not susceptible to threats of the category denial-of-service.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T3	SW management access	An attacker with management access might be able to manipulate the frame memory	Integrity, temporary	Integrity
T.I1	SW management access	An attacker with management access might be able to read the frame memory	Confidentiality, temporary	Confidentiality

Table 5.14.: Identified threats of data store DS.3 Frame Memory for RC communication

5.3.2.4. P.1 - ES-MAC

Table 5.15 lists threats of the process P.1 ES-MAC that provides the main functionality of the end system. The identified threats T.T6, T.D6, T.D7 and T.D8 and their impacts are identical to the threats of P.1 for BE communication described in section 5.3.1.5. In the case of threat T.D6 the attack surface is smaller for RC traffic because an attack is only effective if the system that performs the denial-of-service attack is directly connected to the end system without an ARINC 664 switch in between. Otherwise, the switch would drop all network frames that exceed the configured rate limit. However, traffic policing is only performed at the ingress ports of the network switches and not by end systems.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Availability
T.T18	Adversary-in-the-middle on one channel	An attacker can perform a race attack by exploiting the redundancy management to ensure that injected frames win over legitimate frames	Integrity, temporary	Integrity, Availability
T.D6	Network access	An attacker might be able to overwhelm the input of the host	Availability, temporary	Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability
T.D8	Host access	An attacker with host access can hold the ES in reset	Availability, temporary	Availability

Table 5.15.: Identified threats of process P.1 ES-MAC for RC communication

An ARINC 664 compliant RC frame adds a sequence number that is incremented with every frame. The sequence number is used in networks with redundant channels to de-duplicate frames in the ES before the frames are forwarded to the host. The redundancy management inside the ES-MAC uses a “first valid frame wins” policy, when a second frame is received with the same sequence number, it is simply discarded. An attacker can try to exploit this behavior, by performing a race attack (T.T18). At first, an attacker has to be able to perform an adversary-in-the-middle attack on the VL the attacker wants to inject network frames. The big danger is that even in a multi-channel network an attacker only needs access to one channel to perform this attack. The redundancy management sends only the first frame with a particular sequence number to the host. If an attacker injects network frames with the correct sequence number before the genuine frame arrives at the receiver, then the injected frame always wins. To maximize the success rate, the injected frame should be sent immediately after the previous frame. This attack is a threat to the integrity of P.1 and asset A1. In addition, this threat can also affect the availability of asset A1 because an attacker can use it to ensure that the genuine frame gets dropped.

An attacker can send network frames with a spoofed source MAC address to the end system, but this does not have an effect on an ARINC 664 system because the virtual links (VLs) are statically configured and unidirectional. In ARINC 664 all virtual links have one transmitter and one or more receivers. The destination MAC address is used as the VL ID. This means that in an ARINC 664 network, the source MAC address is not necessary to identify the transmitter. Therefore, no threat could be identified in the category of spoofing.

5.3.2.5. P.2 - Switch-MAC

Table 5.16 lists threats of the process P.2 Switch-MAC that provides the main functionality of the switch. The identified threats T.T7, T.I2, T.D9 and T.E1 and their impacts are identical to the threats of P.2 for BE communication described in section 5.3.1.6. Note that ARINC 664 does not support dynamic network topology changes. A mechanism like spanning tree is therefore not necessary and hence cannot be spoofed. As a result, no threat in the category of spoofing could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.I2	SW management access	An attacker can perform an adversary-in-the-middle attack by rerouting the network frames	Confidentiality, persistent, recoverable	Confidentiality, Integrity
T.D9	SW management access	An attacker with SW management access can disable the switch	Availability, persistent, recoverable	Availability
T.E2	Network access, EMI	Attacker can perform a packet-in-packet attack to inject network frames	Authorization, temporary	Integrity

Table 5.16.: Identified threats of process P.2 Switch-MAC for RC communication

5.3.2.6. P.3 - Traffic Shaping Function

Table 5.17 lists threats of the process P.3 Traffic Shaping Function. This process ensures that the end system sends its frames in the configured data transfer rate. This configuration is stored in the ES configuration. The process RC Traffic Shaping Function is therefore dependent on settings stored in the ES configuration. Without or with incorrect configuration, the ES is not ready for operation. Tampering (T.T6) or deleting (T.D7) the ES configuration (see T.T1 and T.D1) can be a threat to P.3 and can affect for example the integrity and availability of the Traffic Shaping Function. The identified threats T.T6 and T.D7 and their impacts are identical to the threats of P.1 for BE communication described in section 5.3.1.5. The process only controls the data transfer rate and does not process the content of the network frame, therefore no threat in the category of information disclosure could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability

Table 5.17.: Identified threats of process P.3 Traffic Shaping Function for RC communication

5.3.2.7. P.4 - Traffic Policing & Frame Filtering Function

Table 5.18 lists threats of the process P.4 Traffic Policing & Frame Filtering Function. This process enforces the bandwidth limit and the network segmentation per VL. The configuration for this function is stored in the SW configuration. The process RC Traffic Shaping Function is therefore dependent on settings stored in the SW configuration. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can therefore be a threat to P.4 and can affect the integrity and availability of the Traffic Policing & Frame Filtering Function and asset A1.1. A tampered process P.4 can be used to disable the network segmentation. This could compromise the confidentiality and integrity of network frames and asset A1. Recovering a switch via the network with a defective or missing configuration can be difficult or impossible. A non-working switch affects the availability of asset A1. These threats are persistent until the function is correctly reconfigured.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.1
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.18.: Identified threats of process P.4 Traffic Policing & Frame Filtering Function for RC traffic

5.3.2.8. EE.1 - Host

The external entity EE.1 acts as a host for the ES. In an ARINC 664 network, the VL ID corresponds with the network device that transmits a frame. In ARINC 664 a virtual link only has one transmitter and one or more receivers. The destination MAC address defines the VL ID. A VL ID defines a static route that the network frame takes through the network. The source MAC address is not important for an ARINC 664 system because the VL ID already defines the transmitter. Spoofing the source MAC address has no effect, hence no threats in the category of spoofing could be identified. In addition, VL IDs and therefore MAC Addresses are statically configured, so ARP is not needed and therefore ARP spoofing is not possible in an ARINC 664 network.

5.3.2.9. DF.1 - RC Frame

Table 5.19 lists threats of data flow DF.1 RC frame from P.3 ⇒ P.4 and P.2 ⇒ P.1 that represents the network traffic (DF.1) between hosts and switches in the network. The identified threats and impacts are identical to the threats of DF.1 for BE communication described in section 5.3.1.8.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T8	Physical access to network wiring	An attacker can modify network frames	Integrity, temporary	Integrity
T.T9	Physical access to network wiring	An attacker can create new network frames	Integrity, temporary	Integrity
T.I3	Physical access to network wiring	An attacker can sniff network frames	Confidentiality, temporary	Confidentiality
T.D10	Physical access to network wiring	An attacker can drop network frames	Availability, temporary	Availability

Table 5.19.: Identified threats of data flow DF.1 RC frame for RC communication

5.3.2.10. DF.2 - write ES Configuration

Table 5.20 lists threats of data flow DF.2 from EE.1 ⇒ DS.1. The host writes the configuration of the end system into the ES configuration memory via the host interface after a reset. The identified threats and impacts are identical to the threats of DF.2 for BE communication described in section 5.3.1.9.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T10	Physical access to host interface	An attacker can modify ES configuration data on the fly	Integrity, temporary	Availability
T.D11	Physical access to host interface	An attacker can drop ES configuration data	Availability, temporary	Availability

Table 5.20.: Identified threats of data flow DF.2 write ES Configuration for RC communication

5.3.2.11. DF.3 - send and receive RC Frames to host

Table 5.21 lists threats of data flow DF.3 RC frame from EE.1 ⇒ DS.3 and back to the host. The data contains the network traffic that is received by the end system and forwarded to the host or that is sent from the host to the network. The identified threats and impacts are identical to the threats of DF.3 for BE communication described in section 5.3.1.10.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T11	Physical access to host interface	An attacker can modify network frames	Integrity, temporary	Integrity
T.T12	Physical access to host interface	An attacker can create new network frames	Integrity, temporary	Integrity
T.I4	Physical access to host interface	An attacker can sniff network frames	Confidentiality, temporary	Confidentiality
T.D12	Physical access to host interface	An attacker can drop network frames	Availability, temporary	Availability

Table 5.21.: Identified threats of data flow DF.3 for RC communication

5.3.3. Time-Triggered Communication

5.3.3.1. DS.1 - ES Configuration

Table 5.22 lists threats of the data store DS.1 that stores the ES configuration. The identified threats and impacts are identical to the threats of DS.1 for BE communication described in section 5.3.1.1.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T1	Host access	An attacker with control of the host can change the ES configuration	Integrity, persistent, recoverable	Availability
T.D1	Host access	An attacker with control of the host can erase the ES configuration	Availability, persistent, recoverable	Availability

Table 5.22.: Identified threats of data store DS.1 ES configuration for TT communication

5.3.3.2. DS.2 - SW Configuration

Table 5.23 lists threats of the data store DS.2 that stores the SW configuration. The identified threats and impacts are identical to the threats of DS.2 for BE communication described in section 5.3.1.2.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T2	SW management access	An attacker with management access might be able to change the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.D2	SW management access	An attacker with management access can delete the SW configuration	Availability, persistent, non-recoverable	Availability
T.D3	SW management access, data stored on flash memory	An attacker with management access can wear out and physically destroy the flash memory	Availability, persistent, non-recoverable, destructive	Availability

Table 5.23.: Identified threats of data store DS.2 SW configuration for TT communication

5.3.3.3. DS.3 - Frame Memory

Table 5.24 lists threats of the data store DS.3 that stores the ES configuration. The identified threats T.T3 and T.I1 and their impacts are identical to the threats of DS.3 for BE communication described in section 5.3.1.3. However, the threat T.D4 exhaust memory space from BE communication does not apply to TT traffic because process P.4 will drop all frames that arrive out of schedule. The frame memory is therefore not susceptible to threats of the category of denial-of-service.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T3	SW management access	An attacker with management access might be able to manipulate the frame memory	Integrity, temporary	Integrity
T.I1	SW management access	An attacker with management access might be able to read the frame memory	Confidentiality, temporary	Confidentiality

Table 5.24.: Identified threats of data store DS.3 Frame Memory for TT communication

5.3.3.4. DS.5 - Synchronized Local Clock ES

Table 5.25 lists threats of the data store DS.5 that stores the local time that is synchronized with the global time by the clock synchronization mechanism. This value is therefore dependent on the clock synchronization mechanism. Threats to the clock synchronization mechanism are analyzed in detail in section 5.3.4. An attacker could compromise the clock synchronization of the network to adversely affect the integrity of the local clock (T.T13). An incorrect time affects the time-triggered dispatcher P.5 and the availability of the asset A1 because if time-triggered frames are sent at the wrong point in time, they might be dropped by the switch. This threat is temporary until the local clock is synchronized with the global clock again. The global time is non-confidential information, therefore no threats of that category could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T13	Compromised clock sync.	A successful attack of the clock snychronization mechanism will affect the local time	Integrity, temporary	Availability

Table 5.25.: Identified threats of data store DS.5 Synchronized Local Clock for TT communication

5.3.3.5. DS.6 - Synchronized Local Clock SW

Table 5.26 lists threats of the data store DS.6 that stores the local time that is synchronized with the global time by the clock synchronization mechanism. This value is therefore dependent on the clock synchronization mechanism. Threats to the clock synchronization mechanism are analyzed in detail in section 5.3.4. An attacker could compromise the clock synchronization of the network to adversely affect the integrity of the local clock (T.T13). An incorrect time affects the acceptance window check of P.4 and the time-triggered dispatcher P.5. If the acceptance window operates with the wrong time, it can drop all time-triggered frames that arrive at the correct time affecting the availability of the asset A1. This threat is temporary until the local clock is synchronized with the global clock again. The global time is non-confidential information therefore no threats of that category could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T13	Compromised clock sync.	A successful attack of the clock snychronization mechanism will affect the local time	Integrity, temporary	Availability

Table 5.26.: Identified threats of data store DS.6 Synchronized Local Clock for TT communication

5.3.3.6. P.1 - ES-MAC

Table 5.27 lists threats of the process P.1 ES-MAC that provides the main functionality of the end system. The identified threats T.T6, T.D6, T.D7 and T.D8 and their impacts are identical to the threats of P.1 for BE communication described in section 5.3.1.5. In the case of threat T.D6, the attack surface is smaller for TT traffic because an attack is only effective if the system that performs the denial-of-service attack is directly connected to the end system without a TTEthernet switch in between. Otherwise, the switch would drop all network frames that are received out of schedule. End systems do not perform traffic policing on their ingress ports.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Availability
T.T18	Adversary-in-the-middle on one channel	An attacker can perform a race attack by exploiting the redundancy management to ensure that injected frames win over legitimate frames	Integrity, temporary	Integrity, Availability
T.D6	Network access	An attacker might be able to overwhelm the input of the host	Availability, temporary	Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability
T.D8	Host access	An attacker with host access can hold the ES in reset	Availability, temporary	Availability

Table 5.27.: Identified threats of process P.1 ES-MAC for TT communication

TT frames also have a sequence number that is incremented with every frame for frame de-duplication like ARINC 664. Therefore the threat T.T18 can also affect TT traffic in a similar way as described in section 5.3.2.4. The difference is that an attacker must not inject a frame immediately after the previous frame, but right after the acceptance window opens at the ES. Therefore an attacker either needs to know the schedule or needs to observe the timing and pattern of the VL to extract the correct point in time to inject a frame. This makes it more difficult but in general, the attack should still be possible because the acceptance window must be larger than the frame to account for jitter in the transmission.

An attacker can send network frames with a spoofed source MAC address to the end system, but this does not have an effect on a TTEthernet system because the virtual links (VLs) are statically configured and unidirectional. In TTEthernet all virtual links have one transmitter and one or more receivers. The destination MAC address is used as the VL ID. This means that in a TTEthernet network, the source MAC address is not necessary to identify the transmitter. Therefore, no threat could be identified in the category of spoofing.

5.3.3.7. P.2 - Switch-MAC

Table 5.28 lists threats of the process P.2 Switch-MAC that provides the main functionality of the switch. The identified threats T.T7, T.I2, T.D9 and T.E1 and their impacts are identical to the threats of P.2 for BE communication described in section 5.3.1.6. TTEthernet does not support dynamic network topology changes. A mechanism like spanning tree is therefore not necessary and hence cannot be spoofed. As a result, no threat in the category of spoofing could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.I2	SW management access	An attacker can perform an adversary-in-the-middle attack by rerouting the network frames	Confidentiality, persistent, recoverable	Confidentiality, Integrity
T.D9	SW management access	An attacker with SW management access can disable the switch	Availability, persistent, recoverable	Availability
T.E2	Network access, EMI	Attacker can perform a packet-in-packet attack to inject network frames	Authorization, temporary	Integrity

Table 5.28.: Identified threats of process P.2 Switch-MAC for TT communication

5.3.3.8. P.4 - Traffic Policing & Frame Filtering Function

Table 5.29 lists threats of the process P.4 Traffic Policing & Frame Filtering Function. This process enforces the schedule of the time-triggered frames and the network segmentation per VL. To be able to enforce the network schedule correctly, the function needs the correct schedule stored in the SW configuration and the correct global time. At the right time, the function opens an acceptance window at the input port to allow time-triggered frames that arrive in-schedule to pass. If time-triggered frames arrive out of schedule they will be dropped. This process reads the global time from the synchronized local clock DS.6. If an attacker has managed to manipulate the local clock (see T.T13), the acceptance window check could wrongly drop time-triggered frames that arrive actually at the correct point in time (T.T14). This threat is temporary until the local clock is synchronized with the global clock again. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can be a threat to P.4 because it needs the schedule and can affect the integrity and availability of the Traffic Policing & Frame Filtering Function and asset A1.1. A tampered process P.4 can be used to disable the network segmentation. This could compromise the confidentiality and integrity of network frames and asset A1. Recovering a switch via the network with a defective or missing configuration can be difficult or impossible. A non-working switch affects the availability of asset A1. The threats T.T7 and T.D9 are persistent until the function is correctly reconfigured.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Confidentiality, Integrity, Availability
T.T14	Manipulated local clock	The acceptance window check drops correct frames because of the local clock is out of sync	Integrity, temporary	Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.29.: Identified threats of process P.4 Traffic Policing & Frame Filtering Function for TT traffic

5.3.3.9. P.5 - Time-Triggered Dispatcher ES

Table 5.30 lists threats of the process P.5 Time-Triggered Dispatcher. This process ensures that the end system sends its frames at the right point in time into the network. Every TTEthernet device needs a predefined schedule that defines at what point in time (relative to the global time) what frames are sent and received. This configuration is stored in the ES configuration. The Time-Triggered Dispatcher is therefore dependent on the ES configuration DS.1 and the Synchronized Local Clock DS.5. Without or with incorrect configuration, the ES is not ready for operation. Tampering (T.T6) or deleting (T.D7) the ES configuration (see T.T1 and T.D1) can be a threat and can affect the integrity and availability of the Time-Triggered Dispatcher.

These threats are persistent until a correct ES configuration is restored.

An attacker can also try to manipulate the local clock (see T.T13) to change the point in time when the time-triggered frames are transmitted (T.T15). This affects the integrity of the process. This threat is temporary until the local clock is synchronized with the global clock again. All these threats can interrupt the process and affect the ability to send frames at the right point in time and impact the availability of asset A1 because frames can then be dropped by the acceptance window check in P.4. The process P.5 only controls the point in time when a frame is transmitted and does not process the content of the network frame, therefore no threat in the category of information disclosure could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Availability
T.T15	Manipulated local clock	The time-triggered frames are sent at the wrong point in time because the local clock is out of sync	Integrity, temporary	Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability

Table 5.30.: Identified threats of process P.5 Time-Triggered Dispatcher ES for TT communication

5.3.3.10. P.6 - Time-Triggered Dispatcher SW

Table 5.31 lists threats of the process P.6 Time-Triggered Dispatcher. This process ensures that the switch sends its frames at the right point in time into the network. Every TTEthernet device needs a predefined schedule that defines which frames are sent and received at which time, relative to the global time. This configuration is stored in the SW configuration. The Time-Triggered Dispatcher is therefore dependent on the SW configuration DS.2 and the Synchronized Local Clock DS.6. Without or with incorrect configuration, the SW is not ready for operation. Tampering (T.T6) or deleting (T.D7) the SW configuration (see T.T2, T.D2 and T.D3) can be a threat and can affect the integrity and availability of the Time-Triggered Dispatcher. These threats are persistent until a correct SW configuration is restored.

An attacker can also try to manipulate the local clock (see T.T13) to change the point in time when the time-triggered frames are transmitted (T.T15). This affects the integrity of the process. This threat is temporary until the local clock is synchronized with the global clock again. All these threats can interrupt the process, affect the ability to send frames at the right point in time and impact the availability of asset A1 because frames can then be dropped by the acceptance window check in P.4. The process P.6 only controls the point in time when a frame is transmitted and does not process the content of the network frame, therefore no threat in the category of information disclosure could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T6	SW management access, load SW configuration	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, recoverable	Availability
T.T15	Manipulated local clock	The time-triggered frames are sent at the wrong point in time because the local clock is out of sync	Integrity, temporary	Availability
T.D7	SW management access, delete SW configuration	An attacker can disrupt the process by deleting the SW configuration	Availability, persistent, recoverable	Availability

Table 5.31.: Identified threats of process P.6 Time-Triggered Dispatcher SW for TT communication

5.3.3.11. EE.1 - Host

The external entity EE.1 acts as a host for the ES. Like an ARINC 664 network, TTEthernet also uses VLs. The VL ID corresponds with the network device that transmits a frame and the destination MAC address defines the VL ID. The source MAC address is not important for TTEthernet because the VL ID already defines the transmitter. Spoofing the source MAC address has no effect, hence no threats in the category of spoofing could be identified. In addition, VL IDs and therefore MAC Addresses are statically configured, so ARP is not needed and ARP spoofing is not possible.

5.3.3.12. DF.1 - TT Frame

Table 5.32 lists threats of data flow DF.1 TT frame from P.5 ⇒ P.4 and P.6 ⇒ P.1 that represents the network traffic (DF.1) sent between hosts and switches in the network. The identified threats and impacts are identical to the threats of DF.1 for BE communication described in section 5.3.1.8. Note that for time-triggered communication, it is not necessary to drop a frame to impact the availability, delaying the frame to arrive out of schedule can be enough to achieve the same effect (T.D13).

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T8	Physical access to network wiring	An attacker can modify network frames	Integrity, temporary	Integrity
T.T9	Physical access to network wiring	An attacker can create new network frames	Integrity, temporary	Integrity
T.I3	Physical access to network wiring	An attacker can sniff network frames	Confidentiality, temporary	Confidentiality
T.D10	Physical access to network wiring	An attacker can drop network frames	Availability, temporary	Availability
T.D13	Physical access to network wiring	An attacker can delay network frames	Availability, temporary	Availability

Table 5.32.: Identified threats of data flow DF.1 TT frame for TT communication

5.3.3.13. DF.2 - write ES Configuration

Table 5.33 lists threats of data flow DF.2 from EE.1 ⇒ DS.1. The host writes the configuration of the end system into the ES configuration memory via the host interface after a reset. The identified threats and impacts are identical to the threats of DF.2 for BE communication described in section 5.3.1.9.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T10	Physical access to host interface	An attacker can modify ES configuration data on the fly	Integrity, temporary	Availability
T.D11	Physical access to host interface	An attacker can drop ES configuration data	Availability, temporary	Availability

Table 5.33.: Identified threats of data flow DF.2 write ES Configuration for TT communication

5.3.3.14. DF.3 - send and receive TT Frames to host

Table 5.34 lists threats of data flow DF.3 TT frame from EE.1 ⇒ DS.3 and back to the host. The data contains the network traffic that is received by the end system and forwarded to the host or that is sent from the host to the network. The identified threats and impacts are identical to the threats of DF.3 for BE communication described in section 5.3.1.10.

ID	Precondition	Threat Description	Impact on Element	Impact on A1
T.T11	Physical access to host interface	An attacker can modify network frames	Integrity, temporary	Integrity
T.T12	Physical access to host interface	An attacker can create new network frames	Integrity, temporary	Integrity
T.I4	Physical access to host interface	An attacker can sniff network frames	Confidentiality, temporary	Confidentiality
T.D12	Physical access to host interface	An attacker can drop network frames	Availability, temporary	Availability

Table 5.34.: Identified threats of data flow DF.3 for TT communication

5.3.4. Clock Synchronization

The task of the clock synchronization service is to establish a common time between the time-triggered network devices, this time is called global time. This service does not rely on confidential data for its security. Therefore, no security threat in the category of information disclosure could be identified in this section.

5.3.4.1. DS.1 - ES Configuration

Table 5.35 lists threats of the data store DS.1 that stores the ES configuration. The identified threats and impacts are identical to the threats of DS.1 for BE communication described in section 5.3.1.1.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T1	Host access	An attacker with control of the host can change the ES configuration	Integrity, persistent, recoverable	Integrity, Availability
T.D1	Host access	An attacker with control of the host can erase the ES configuration	Availability, persistent, recoverable	Availability

Table 5.35.: Identified threats of data store DS.1 ES configuration for clock synchronization

5.3.4.2. DS.2 - SW Configuration

Table 5.36 lists threats of the data store DS.2 that stores the SW configuration. The identified threats and impacts are identical to the threats of DS.2 for BE communication described in section 5.3.1.2.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T2	SW management access	An attacker with management access might be able to change the SW configuration	Integrity, persistent, non-recoverable	Integrity, Availability
T.D2	SW management access	An attacker with management access can delete the SW configuration	Availability, persistent, non-recoverable	Availability
T.D3	SW management access, data stored on flash memory	An attacker with management access can wear out and physically destroy the flash memory	Availability, persistent, non-recoverable, destructive	Availability

Table 5.36.: Identified threats of data store DS.2 SW configuration for clock synchronization

5.3.4.3. DS.5 - Synchronized Local Clock ES

Table 5.37 lists threats of the data store DS.5 stores the local time that is synchronized with the global time by the clock synchronization mechanism. This value is therefore dependent on the clock synchronization mechanism. An attacker could compromise the clock synchronization of the network to adversely affect the integrity of the synchronized local clock (T.T13). An incorrect time affects the PCF dispatcher P.7 and the availability of the asset A1.2 because if PCFs are sent at the wrong point in time they might be dropped by the switch. This threat is temporary until the local clock is synchronized with the global clock again. The global time is not confidential information, therefore no threats of that category could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T13	Compromised clock sync.	A successful attack of the clock snychronization mechanism will affect the local time	Integrity, temporary	Availability

Table 5.37.: Identified threats of data store DS.5 Synchronized Local Clock for clock synchronization

5.3.4.4. DS.6 - Synchronized Local Clock SW

Table 5.38 lists threats of the data store DS.6 that stores the local time that is synchronized with the global time by the clock synchronization mechanism. This value is therefore dependent on the clock synchronization mechanism. An attacker could compromise the clock synchronization of the network to adversely affect the integrity of the synchronized local clock (T.T13). An incorrect time affects the acceptance window check of P.11 and the time-triggered dispatcher P.12. If the acceptance window operates with the wrong time, it can drop all PCFs that arrive at the correct time affecting the availability of the asset A1.2. This threat is temporary until the local clock is synchronized with the global clock again. The synchronized local clock is not confidential information, therefore no threats of that category could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T13	Compromised clock sync.	A successful attack of the clock snychronization mechanism will affect the local time	Integrity, temporary	Availability

Table 5.38.: Identified threats of data store DS.6 Synchronized Local Clock for clock synchronization

5.3.4.5. DS.7 - Local Clock

Table 5.39 lists threats of the data store DS.7 that stores the local clock. The local clock is the local time source and reference of a time-triggered device. While in the unsynchronized state, the SM sends the local time to the CM in the form of a Coldstart PCF. Threats to the local clock are very dependent on the actual implementation. The local clock may be implemented as a fixed logic that cannot be changed or affected. For this analysis, it is assumed that in some form or another, it might be possible for an attacker that is in control of the host to manipulate the value of the local clock (T.T16) or at least to perform a denial-of-service attack on it (T.D14). The local clock can for example run too fast or too slow or maybe also stop or jump in its value. Threats to the local clock affect the integrity and availability of DS.7 and asset A1.2. This threat is temporary until the manipulation stops and the clock is progressing uninterrupted again. The local clock does not contain any confidential information.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T16	Host access	An attacker with control of the host can manipulate the local clock	Integrity, temporary	Integrity, Availability
T.D14	Host access	An attacker with control of the host can stop the local clock	Availability, temporary	Availability

Table 5.39.: Identified threats of data store DS.7 Local Clock for clock synchronization

5.3.4.6. P.7 - PCF Dispatcher ES

Table 5.40 lists threats of the process P.7 PCF Dispatcher. This process ensures that the sync master (SM) sends the IN (integration) PCFs at the right point in time into the network. The SMs are providing a reference used to create the global time. Note that for this analysis, it is assumed that the sync masters are located in the end systems, as shown in Figure 4.2. The configuration for this is stored in the ES configuration. The

PCF Dispatcher is therefore dependent on the ES configuration DS.1 and the Synchronized Local Clock DS.5. Without or with incorrect configuration, the ES is not ready for operation. Tampering (T.T6) or deleting (T.D7) the ES configuration (see T.T1 and T.D1) can be a threat and can affect the integrity and availability of the PCF Dispatcher. These threats are persistent until a correct ES configuration is restored. An attacker can also try to manipulate the local clock (see T.T13) that is synchronized with the global time to change the point in time when PCF IN frames are transmitted (T.T15). This affects the integrity of the process. This threat is temporary until the local clock is synchronized with the global clock again. Another threat that can have an effect on the availability is if the host holds the ES-MAC in reset (T.D8). The threats T.D7 and T.D8 have only an impact on the clock synchronization of the entire network (asset 1.2) if multiple SMs are affected. Note that the minimum number of sync masters that are required to operate is configurable. The default value is three. The threats T.T6 and T.T15 can affect the integrity and availability of asset A1.2 if an attacker controls multiple SMs in the network.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Integrity, Availability
T.T15	Manipulated local clock	The PCFs are sent at the wrong point in time because of a manipulated local clock	Integrity, temporary	Integrity, Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability
T.D8	Host access	An attacker with host access can hold the ES in reset	Availability, temporary	Availability

Table 5.40.: Identified threats of process P.7 PCF Dispatcher ES for clock synchronization

5.3.4.7. P.4 - Traffic Policing & Frame Filtering Function

Table 5.41 lists threats of the process P.4 Traffic Policing & Frame Filtering Function. This process polices incoming PCFs dependent on the sync state and VL IDs, and is executed by every compression master (CM) in the network. All PCFs that are received on ports and from VL IDs that are not connected to synchronization masters will be dropped. This should ensure that only PCF integration (IN) frames from an SM can contribute to the clock synchronization. Coldstart PCFs and Coldstart acknowledgment PCFs are dropped by the frame filtering function while the CM is in the synchronized state. The necessary information to perform this functionality is stored in the SW configuration. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can be a threat to P.4 because malicious configuration can affect the integrity and availability of the Traffic Policing & Frame Filtering Function and therefore asset A1.2. The threats T.T7 and T.D9 are persistent until the function is correctly reconfigured.

The next threat is dependent on the network topology. If the SMs are not connected directly to the CM but there is for example one network switch between an SM and the CM, then the attacker can try to perform an adversary-in-the-middle attack via the switch in between to spoof the origin of the PCF to bypass the traffic policing function (T.S4). PCFs are not protected with a message authentication code (MAC) therefore an adversary-in-the-middle can arbitrarily change the data in a PCF. Note that this attack is only effective if an attacker can spoof PCFs from multiple SMs because malicious PCFs from a single SM cannot impact the integrity and availability of asset A1.2. The threat T.S4 has a temporary effect on the process and goes away if the attack stops.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.S4	Adversary-in-the-middle	An attacker can spoof PCF IN frames by faking its origin	Authenticity, temporary	Integrity, Availability
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Integrity, Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.41.: Identified threats of process P.4 Traffic Policing & Frame Filtering Function for clock sync.

5.3.4.8. P.8 - Compression Function

Table 5.42 lists threats of the process P.8 Compression Function. The compression function periodically collects multiple PCF IN frames within a certain time frame (observation window) and calculates a fault-tolerant average time over all the input values. Only one PCF per compression master (CM) is used for the calculation to avoid the time of a single CM having more impact than others. This time is then later used to establish the global time in the network. Note that this process is performed on all synchronization masters (SM) independently, but there is only one global time over all channels established by the clock synchronization protocol. The compression function runs unsynchronized to the synchronized local clock and is therefore not dependent on it. The configurable settings of this function (like the observation window) are stored in the SW configuration. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can be a threat to P.8 because malicious configuration can affect the integrity and availability of the Compression Function. This can also impact asset A1.2 because misconfigured values between channels can in the worst case lead to different calculation times between the channels. The threats T.T7 and T.D9 are persistent until the function is correctly reconfigured.

The next threat is dependent on the network topology. If the SMs are not connected directly to the CM but there is for example one network switch between a SM and the CM, then the attacker can try to perform an adversary-in-the-middle attack via the switch in between to spoof the PCF (T.S4). PCFs are not protected with a message authentication code therefore an adversary-in-the-middle can arbitrarily change the data in a PCF. This can be used to trick the compression algorithm to use PCFs from just one sender for the time calculation. Note that this attack is only effective if an attacker is able to spoof PCFs from multiple SMs in the network because malicious PCFs from a single SM cannot impact the integrity and availability of asset A1.2. The threat T.S4 has a temporary effect on the process and goes away if the attack stops. The same is true if malicious PCFs are sent directly from multiple SMs (T.T19). An attacker can either misconfigure these SMs or generate and send malicious PCFs from their host to perform this attack.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.S4	Adversary-in-the-middle	An attacker can spoof PCFs so that several PCFs are used from one sender	Authenticity, temporary	Integrity, Availability
T.T7	SW management access	An attacker can tamper the process by changing its settings stored in the SW configuration	Integrity, persistent, non-recoverable	Integrity, Availability
T.T19	Host access	An attacker controls multiple SMs and sends malicious PCFs	Integrity, temporary	Integrity, Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.42.: Identified threats of process P.8 Compression Function for clock synchronization

5.3.4.9. P.11 - Acceptance Window Check SW

Table 5.43 lists threats of the process P.11 Acceptance window check. A compression master that is in the synchronized state uses an acceptance window to check if the PCFs produced by the compression function are aligned to the already synchronized local clock. PCFs that arrive outside of the acceptance window will be dropped and do not contribute to the clock synchronization process. The process needs the synchronized local clock (that is aligned with the global clock) and the SW configuration to get the correct parameters for the window size and the point in time when the window should open and close. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can be a threat to P.11 because a malicious configuration can affect the integrity and availability of the Acceptance window check. A misconfigured acceptance window can lead to a loss of the availability of the clock synchronization (asset A1.2). The threats T.T7 and T.D9 are persistent until the function is correctly reconfigured.

This function is dependent on the synchronized local clock. A manipulated synchronized local clock can cause valid PCFs to be dropped and malicious PCFs to pass the acceptance window check (T.T17). The PCFs that pass the check are used to calculate the new global time. This threat affects the integrity of P.11 and can cause a loss of availability of asset A1.2. The threat T.T17 is temporary until the local clock is synchronized with the global clock again. Note that for networks with redundant channels, one manipulated acceptance window check is not enough to cause a sync loss on all channels.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Integrity, Availability
T.T17	Manipulated local clock	Valid PCFs are dropped and malicious PCFs can pass the check because of a manipulated local clock	Integrity, temporary	Integrity, Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.43.: Identified threats of process P.11 Acceptance window check SW for clock synchronization

5.3.4.10. P.10 - Clock Synchronization & Compressed PCF Dispatcher

Table 5.44 lists threats of the process P.10 Clock Synchronization & Compressed PCF Dispatcher. This process receives the compressed PCFs that successfully passed the acceptance window check and uses the time information from the PCF to readjust its own synchronized local clock with the global time. The process also forwards the compressed PCF to all time-triggered devices in the network and there the compressed PCF is also used for local clock synchronization. This function is dependent on the compressed PCFs it receives and on the SW configuration. The SW configuration for example defines the network devices that should receive the compressed PCFs. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can be a threat to P.10 because a malicious configuration can affect the integrity and availability of the Clock Synchronization & Compressed PCF Dispatcher. Time-triggered devices that do not receive compressed PCFs cannot synchronize to a global clock and are not operational, affecting the availability of asset A1.2. It could reduce the integrity of asset A1.2 if compressed PCFs are only forwarded to a subset of the time-triggered devices because some devices are now able to synchronize and some devices are not. The threats T.T7 and T.D9 are persistent until the function is correctly reconfigured.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Integrity, Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.44.: Identified threats of process P.10 Clock Synchronization & Compressed PCF Dispatcher

5.3.4.11. P.12 - Clique Detection & Resolution Service SW

Table 5.45 lists threats of the process P.12 Clique Detection & Resolution Service SW. The purpose of this process is to monitor the compressed PCFs and detect if a clique has formed that is not in sync with the time of the rest of the network. The CM performs two different types of clique detection checks: synchronous clique detection and asynchronous clique detection. A PCF contains a data field called membership vector that lists all SMs that have been used to create this frame and therefore contributed to its time. The synchronous clique detection is triggered if the PCF is received in-schedule (within the acceptance window), but the number of bits set in the membership vector is below a predefined threshold. This can happen if there are some SMs that are synchronized with the CM but there are too few of them. The asynchronous clique detection is triggered if the PCF was received out-of-schedule (outside the acceptance window), but the number of bits set in the membership vector is equal to or above a predefined threshold. This can happen if several SMs are not synchronized with the CM.

The predefined thresholds are stored in the SW configuration. Tampering (T.T7) or deleting (T.D9) the SW configuration (see T.T2 and T.D2) can be a threat to P.12 because a malicious configuration can affect the integrity and availability of the Clique Detection & Resolution Service. A misconfigured clique detection can cause the function to trigger a resynchronization of the network even if everything is correct or it does not resynchronize even if it should. This can affect the availability of asset A1.2. While the network is unsynchronized or in the process of resynchronization, TT frames cannot be sent. This process is also dependent on the output of the Acceptance Window Check because it determines if the PCF was received in-schedule or out-of-schedule.

An attacker can try to trigger the clique detection by sending malicious PCFs from multiple SMs in the network (T.T19). This affects the integrity of P.12 and a reset of CMs also affects the availability of the entire network. The next threat is dependent on the network topology. If the SMs are not connected directly to the CM but there is for example one network switch between a SM and the CM, then the attacker can try to perform an adversary-in-the-middle attack via the switch in between to spoof the PCF (T.S4). This can be used to trick the clique detection into using PCFs from just one sender to determine if a clique has formed. Note that this attack is only effective if an attacker is able to spoof PCFs from multiple SMs in the network because malicious PCFs from a single SM cannot impact the integrity and availability of asset A1.2. The threats T.S4 and T.T19 have a temporary effect on the process and go away if the attack stops.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.S4	Adversary-in-the-middle	An attacker can spoof PCFs so that several PCFs are used from one sender	Authenticity, temporary	Integrity, Availability
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Integrity, Availability
T.T19	Host access	An attacker controls multiple SMs and sends malicious PCFs	Integrity, temporary	Integrity, Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.45.: Identified threats of process P.12 Clique Detection & Resolution Service SW

5.3.4.12. P.13 - Startup & Restart Service SW

Table 5.46 lists threats of the process P.13 Startup & Restart Service SW. This process handles the state transition from the synchronized or power-up into the un-synchronized state. This service includes procedures for cold starts and the integration or reintegration of a time-triggered network device. If the Clique Detection process P.12 detects a clique it triggers a restart of the clock synchronization mechanism that is executed by P.13. This process is therefore dependent on P.12.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T7	SW management access	An attacker can tamper the process by changing its setting stored in the SW configuration	Integrity, persistent, non-recoverable	Integrity, Availability
T.D9	SW management access	An attacker with SW management access can disable the function	Availability, persistent, recoverable	Availability

Table 5.46.: Identified threats of process P.13 Startup & Restart Service SW for clock synchronization

5.3.4.13. P.14 - Acceptance Window Check ES

Table 5.47 lists threats of the process P.14 Acceptance window check inside the ES. A synchronization master or synchronization client that is in the synchronized state uses an acceptance window to check if the PCFs sent by the CM are aligned to the already synchronized local clock. PCFs that arrive outside of the acceptance window will be dropped and do not contribute to the clock synchronization process. The process needs the synchronized local clock (that is aligned with the global clock) and the ES configuration to get the correct parameters for the window size and the point in time when the window should open and close. Tampering (T.T6) or deleting (T.D7) the ES configuration (see T.T1 and T.D1) can be a threat to P.14 because a malicious configuration can affect the integrity and availability of the Acceptance Window Check. A misconfigured acceptance window can lead to a loss of the availability of the clock synchronization (asset A1.2). The threats T.T7 and T.D9 are persistent until the function is correctly reconfigured.

This function is dependent on the synchronized local clock. A manipulated synchronized local clock can cause valid PCFs to be dropped and malicious PCFs to pass the acceptance window check (T.T17). The PCFs that pass the check are used to calculate the new global time. This threat affects the integrity of P.14 and can cause a loss of availability of asset A1.2. The threat T.T17 is temporary until the local clock is synchronized with the global clock again. Note that one synchronization master with a manipulated acceptance window check is not enough to cause a sync loss and will not trigger a resynchronization of the entire network. A network with redundant channels will not change the threat because end systems are connected to all channels.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Integrity, Availability
T.T17	Manipulated local clock	Valid PCFs are dropped and malicious PCFs can pass the check because of a manipulated local clock	Integrity, temporary	Integrity, Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability

Table 5.47.: Identified threats of process P.14 Acceptance window check ES for clock synchronization

5.3.4.14. P.9 - Clock Synchronization Service

Table 5.48 lists threats of the process P.9 Clock Synchronization Service. This process receives the compressed PCFs that successfully passed the acceptance window check and uses the time information from the PCF to readjust its own synchronized local clock with the global time. This function is dependent on the compressed PCFs it receives from the CM and on the ES configuration. To readjust the local clock, it is necessary to read the time information from the PCF and add the time it took to reach the ES. This information, called propagation delay, is calculated a priori and stored in the ES configuration.

Tampering (T.T6) or deleting (T.D7) the ES configuration (see T.T1 and T.D1) can be a threat to P.9 because a malicious configuration can affect the integrity and availability of the Clock Synchronization Service. Time-triggered devices that do not receive compressed PCFs cannot synchronize to a global clock and are not operational affecting the availability of asset A1.2. It could reduce the integrity of asset A1.2 if compressed PCFs are only forwarded to a subset of the time-triggered devices because some devices are now able to synchronize and some devices are not. It can also happen that a device can synchronize but is not synchronized correctly with the rest of the network because of a for example incorrect propagation delay. The threats T.T7 and T.D9 are persistent until the function is correctly reconfigured.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Integrity, Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability

Table 5.48.: Identified threats of process P.9 Clock Synchronization Service

5.3.4.15. P.15 - Clique Detection & Resolution Service ES

Table 5.49 lists threats of the process P.15 Clique Detection & Resolution Service of synchronization master and client. The job of this process is to monitor the compressed PCFs and detect if a clique has formed that is not in sync with the time of the rest of the network. An SM performs three different types of clique detection checks: synchronous clique detection, asynchronous clique detection and relative clique detection. An SC performs two different types of clique detection checks: synchronous clique detection and asynchronous clique detection. A PCF contains a data field called membership vector that lists all SMs that have been used to create this frame and therefore contributed to its time. The synchronous clique detection is triggered if the PCF was received in-schedule (within the acceptance window), but the number of bits set in the membership vector is below a predefined threshold. This can happen if there are some SMs that are synchronized with the CM but there are too few of them. The asynchronous clique detection is triggered if the PCF was received out-of-schedule (outside the acceptance window), but the number of bits set in the membership vector is equal to or above a predefined threshold. This can happen if several SMs are not synchronized with the CM. Relative clique detection is triggered if the number of synchronized SMs is equal to or below the number of unsynchronized SMs.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.S5	Adversary-in-the-middle	An attacker can spoof PCFs to trigger the clique detection mechanism	Authenticity, temporary	Integrity, Availability
T.S6	Network access, EMI	An attacker performs a PCspooF attack to trigger the clique detection mechanism	Authenticity, temporary	Integrity, Availability
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Integrity, Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability

Table 5.49.: Identified threats of process P.15 Clique Detection & Resolution Service ES

The predefined thresholds are stored in the ES configuration. Tampering (T.T6) or deleting (T.D7) the ES configuration (see T.T1 and T.D1) can be a threat to P.15 because a malicious configuration can affect the integrity and availability of the Clique Detection & Resolution Service. A misconfigured clique detection can cause the function to trigger a resynchronization of the network, even if everything is correct or it does not resynchronize even if it should. This can affect the availability of asset A1.2. While the network is unsynchronized or in the process of resynchronization TT frames cannot be sent. This process is also dependent on the output of the Acceptance Window Check because it determines if the PCF was received in-schedule or out-of-schedule.

PCFs do not provide a mechanism to verify if they are authentic. If an attacker can perform an adversary-in-the-middle attack or can inject a PCF that is then received by the time-triggered network device, this device can not verify if the PCF and its content are genuine. The PCspooF attack exploits this vulnerability by performing a packet-in-packet attack to inject a malformed PCF that triggers the clique detection mechanism (T.S6). This attack can be performed by an unprivileged network device that is connected directly to a port of a compression master. The attacker has to be able to send hand-crafted network frames and the device needs a hardware implant that can send an EMI pulse to a CM. If an attacker manages to send a network frame that contains a PCF with a membership vector that is set to trigger the clique detection, it could cause a resynchronization of the network device that received this PCF. This attack is described in detail in Table A.5. If this attack is successful on multiple SMs at the same time, then the effect is not limited to these SMs but to the entire network.

Another way to achieve the same result as with PCspooF is if the attacker can inject arbitrary PCFs by an adversary-in-the-middle attack (T.S5). To do so, an attacker can for example be in control of a network switch connected to SM and SC. The effects of threats T.S5 and T.S6 are temporary and will stop when the attack stops. These threats can affect the availability of asset A1.2.

5.3.4.16. P.16 - Startup & Restart Service ES

Table 5.50 lists threats of the process P.16 Startup & Restart Service ES. This process handles the state transition from the synchronized or power-up into the un-synchronized state. This service includes procedures for cold starts and the integration or reintegration of a time-triggered network device. If the Clique Detection process P.15 detects a clique, it triggers a restart of the clock synchronization mechanism that is executed by P.16. This process is therefore dependent on P.15. An SM that receives a Coldstart Acknowledgement PCF will immediately trigger the Startup & Restart Service. TT frames cannot be sent while the network is unsynchronized or in the process of resynchronization. An attacker can inject a Coldstart ACK PCF via the PCspooF attack (T.S6) or an adversary-in-the-middle attack (T.S5). The PCspooF attack is described in detail in Table A.5. The effects of threats T.S5 and T.S6 are temporary and will stop when the attack stops. These threats can affect the availability of asset A1.2.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.S5	Adversary-in-the-middle	An attacker can spoof PCFs to inject a Coldstart ACK PCF	Authenticity, temporary	Integrity, Availability
T.S6	Network access, EMI	An attacker performs a PCspooF attack to inject a Coldstart ACK PCF	Authenticity, temporary	Integrity, Availability
T.T6	Host access, load ES configuration	An attacker can tamper the process by changing its setting stored in the ES configuration	Integrity, persistent, recoverable	Integrity, Availability
T.D7	Host access, delete ES configuration	An attacker can disrupt the process by deleting the ES configuration	Availability, persistent, recoverable	Availability

Table 5.50.: Identified threats of process P.16 Startup & Restart Service ES

5.3.4.17. EE.1 - Host

The external entity EE.1 acts as a host for the ES. It can write the ES configuration, besides that the host is not involved in the clock synchronization process. Therefore, no threats to the host could be identified within the scope of this analysis. Note that the EE.1 can be a threat to the ES configuration but not the other way around.

5.3.4.18. DF.1 - PCF

Table 5.51 lists threats of data flow DF.1 PCF from SM ⇒ CM and CM ⇒ SM and SC. This represents the exchange of PCFs between the network devices. Note that for time-triggered communication, it is not necessary to drop a frame to impact the availability, delaying the frame to arrive out of schedule can be enough to achieve the same effect. PCFs don't contain any confidential information therefore no threat in the category of information disclosure could be identified.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T8	Physical access to network wiring	An attacker can modify network frames	Integrity, temporary	Integrity
T.T9	Physical access to network wiring	An attacker can create new network frames	Integrity, temporary	Integrity
T.D10	Physical access to network wiring	An attacker can drop network frames	Availability, temporary	Availability
T.D13	Physical access to network wiring	An attacker can delay network frames	Availability, temporary	Availability

Table 5.51.: Identified threats of data flow DF.1 PCF for clock synchronization

5.3.4.19. DF.2 - write ES Configuration

Table 5.52 lists threats of data flow DF.2 from EE.1 ⇒ DS.1. The host writes the configuration of the end system into the ES configuration memory via the host interface after a reset. An attacker with physical access to the host interface might be able to modify (T.T10) or drop (T.D11) part or all of the configuration data that is transmitted to the end system. These threats have a direct effect on the ES configuration and impact the availability of asset A1.2 as described in section 5.3.4.1 DS.1. The ES configuration does not contain confidential information, therefore no threat could be identified in the category of information disclosure. These threats are temporary and will stop if the attacker loses physical access to the host interface.

ID	Precondition	Threat Description	Impact on Element	Impact on A1.2
T.T10	Physical access to host interface	An attacker can modify the ES configuration data on the fly	Integrity, temporary	Availability
T.D11	Physical access to host interface	An attacker can drop the ES configuration data	Availability, temporary	Availability

Table 5.52.: Identified threats of data flow DF.2 write ES Configuration for clock synchronization

5.4. Analysis Results

Comparing only the number of threats identified per traffic class with each other, as shown in Table 5.53, is not a good indicator to find out if one traffic class has superior security performance over another. It is also necessary to take the type and impact of the threat into consideration.

Traffic class	**Number of threats**	**S**	**T**	**R**	**I**	**D**	**E**
BE	33	3	12	0	4	12	2
RC	26	0	11	0	4	10	1
TT[1]	37	3	17	0	4	12	1

Table 5.53.: Analysis results per traffic class

It is therefore better to compare the traffic classes on a threat-by-threat basis. The absence of a threat is also important to consider. Table 5.54 lists all threats in the category of spoofing that have been identified.

Threat	**Element**	**BE**	**RC**	**TT**	**CS**
T.S1	P.1	X			
T.S2	P.2	X			
T.S3	EE.1	X			
T.S4	P.4, P.8, P.12				X
T.S5	P.15, P.16				X
T.S6	P.15, P.16				X

Table 5.54.: Analysis results for spoofing threats

Although both BE and TT traffic are vulnerable to three threats in this category, they are not the same threats. Several spoofing threats of a BE network are caused by the fact that a typical Ethernet network has built-in mechanisms to dynamically configure the network. These functions make it easier to set up the network, but this plug-and-play convenience not only helps a network administrator but is also plug-and-play for an attacker. An attacker can spoof these mechanisms to exploit their functionality. For example, threat T.S2 is a spoofing attack on the spanning tree protocol that allows a network to automatically configure its network topology. Threats T.S1 and T.S3 are both attacks that spoof ARP messages. This functionality allows the

[1] The number of threats of traffic class TT includes also the threats of the clock sync mechanism because TT traffic is dependent on it.

end systems to automatically learn the MAC addresses of new network devices. An RC and TT network does not have this functionality because the whole network is configured a priori in the design phase of the network. This greatly reduces this type of threat. The three spoofing threats (T.S4, T.S5, T.S6) that can affect a TT network are different ways to spoof PCFs that are used for the clock synchronization mechanism. These three threats have in common that they are only effective if PCFs from multiple SMs are spoofed at the same time, otherwise, the fault-tolerant synchronization mechanism ensures that the attack does not affect the whole TT network.

Threat	**Element**	**BE**	**RC**	**TT**	**CS**
T.T1	DS.1	X	X	X	X
T.T2	DS.2	X	X	X	X
T.T3	DS.3	X	X	X	
T.T4	DS.4	X			
T.T5	DS.4	X			
T.T6	P.1, P.3, P.5, P.6, P.7, P.9, P.14, P.15, P.16	X	X	X	X
T.T7	P.2, P.4, P.8, P.10, P.11, P.12, P.13	X	X	X	X
T.T8	DF.1	X	X	X	X
T.T9	DF.1	X	X	X	X
T.T10	DF.2	X	X	X	X
T.T11	DF.3	X	X	X	
T.T12	DF.3	X	X	X	
T.T13	DS.5, DS.6			X	X
T.T14	P.4			X	
T.T15	P.5, P.6, P.7			X	X
T.T16	DS.7				X
T.T17	P.11, P.14				X
T.T18	P.1		X	X	
T.T19	P.8, P.12				X

Table 5.55.: Analysis results for tampering threats

Table 5.55 lists all identified threats in the category tampering. Most of the threats that have been identified belong to this category. Many of these threats are independent of the traffic class and are inherent to a network system. Good examples of that are threats T.T1, T.T2, T.T6 and T.T7 that affect the configuration of the network devices. Manipulated configurations can affect many elements and have therefore a big impact on the security of the network. Threat T.T3 is implementation-dependent and only possible if the frame memory is accessible, for example via the management CPU. Threats that tamper with the SAT entries (T.T4 and T.T5) only affect BE communication if dynamic MAC address learning is enabled in the network switches. This is a plug-and-play feature of a typical Ethernet network that is not present in RC and TT networks. Tampering of data in transit is another type of threat that is independent of the traffic class (T.T8, T.T9, T.T10, T.T11, T.T12). An attacker with physical access to connections and cables can tamper with the data when the message is not protected by authentication. A TT network is dependent on a working clock synchronization mechanism. This mechanism needs additional functionality that increases the attack surface and introduces new threats. Every TT device stores a local clock that is synchronized with the global time and stored as synchronized local clock. If an attacker manages to tamper with this data store, the network device can lose synchronization with the network (T.T13, T.T16). The synchronized clock is used by several processes, for example by the acceptance window check inside the network switches and by the frame dispatcher. TT and PCFs will be dropped if they are sent out of schedule (T.15). If the clock of the acceptance window check is out of sync, the switch might drop all TT (T.T14) or PCF (T.T17) frames. Depending on which TT network device and how many network devices are tampered with, the impact on the network is limited to only these devices. This applies in particular to TT devices that have the SM role. The clock synchronization can be disrupted if multiple SMs are sending malicious PCFs (T.19). Threat T.18 is only applicable for RC or TT networks with redundant channels. There, an attacker can inject frames that overrule genuine frames on redundant channels even if the attacker controls only one channel. Note that TT traffic is only affected if the ES uses a “first valid frame wins” policy instead of a voting policy for frame deduplication. BE networks are not affected because of their lack of redundancy in the first place. Therefore, injecting BE frames is possible without the need to perform T.18. This threat is a notable example that redundancy is an appropriate measure to increase the safety of a system, but not its security.

Threat	**Element**	**BE**	**RC**	**TT**	**CS**
T.I1	DS.3	X	X	X	
T.I2	P.2	X	X	X	
T.I3	DF.1	X	X	X	
T.I4	DF.3	X	X	X	

Table 5.56.: Analysis results for Information Disclosure threats

Table 5.56 lists all identified threats in the category of information disclosure. No network system of the three traffic classes is dependent on or requires confidential information. All threats (T.I1, T.I2, T.I3, T.I4) are related to potential disclosure of confidential information in the payload. If the payload data that is sent over the network is not confidential or is protected on an upper OSI layer, then none of these threats have an impact on confidentiality. The threat T.I1 is only applicable for certain implementations, for example, if the management CPU has read access to the frame memory.

Threat	Element	BE	RC	TT	CS
T.D1	DS.1	X	X	X	X
T.D2	DS.2	X	X	X	X
T.D3	DS.2	X	X	X	X
T.D4	DS.3	X			
T.D5	DS.4	X			
T.D6	P.1	X	X	X	
T.D7	P.1, P.3, P.5; P.6, P.7, P.9, P.14, P.15, P.16	X	X	X	X
T.D8	P.1, P.7	X	X	X	X
T.D9	P.2, P.4, P.8, P.10, P.11, P.12, P.13	X	X	X	X
T.D10	DF.1	X	X	X	X
T.D11	DF.2	X	X	X	X
T.D12	DF.3	X	X	X	
T.D13	DF.1				X
T.D14	DS.7				X

Table 5.57.: Analysis results for Denial-of-Service threats

Table 5.57 lists all identified threats in the category denial-of-service. DoS is the Achilles heel of BE traffic due to the lack of QoS in Ethernet. RC and TT systems add QoS capabilities to Ethernet to provide performance guarantees for throughput and latency. This is the reason why threat T.D4 is only applicable to BE networks. RC networks have the functionality to limit the data rate per VL to mitigate the threat effectively if the rate limit has been calculated correctly. The DoS protection of a TT network is even stronger because it is an implicit functionality of how time-triggered systems work. A network schedule defines when a frame is allowed to pass the input filter. Out-of-schedule frames are dropped. This is enforced by the switches in both (RC and TT) cases. Another threat that only affects BE communication is flooding the SAT with MAC addresses (T.D5) if dynamic MAC address learning is enabled in the network switches. This attack can lead to an information disclosure threat because a typical network switch would fall back from a switch to a bridge operating mode. In this case, all network traffic would be forwarded to the attacker. In RC and TT networks the whole network configuration must be done a priori in the design phase of the network, therefore dynamic MAC address learning is not supported. An attacker can try to affect the availability of the host input by sending data at full line speed to it (T.D6). Some embedded systems, such as sensors and actuators, can lack the computational resources to handle high data load and there is the potential that this degrades its capability to perform tasks. This threat is relevant for all three traffic classes but for BE traffic it is easier to achieve due to its lack of QoS capabilities. In RC and TT networks the maximum bandwidth can be limited to a data rate the host can handle but this bandwidth limit is enforced in the switches of the network. The ES does not have a built-in bandwidth limiter, therefore the host input can be flooded if an attacker gains direct network access to the ES without a protective SW in between. This threat can also affect networks with redundancy if the bandwidth of the host input is the limiting factor. An attacker with access to only one channel of the network can try to send as many frames as the network allows to saturate the connection between ES and host to affect the communication of the other channels. This case shows that it is important to also consider the maximum throughput of the host interface when the bandwidth limit is calculated. Deleting the configuration of a network device (T.D1, T.D2, T.D7, T.D9) has a negative effect on the availability of the network device and has therefore a major impact on the network. This is independent of the traffic class that is used. Another threat applicable to all traffic classes is if the host holds the end system in reset (T.D8). This threat is less relevant for data communication because, with the access required to perform this attack, an attacker can probably also stop the host itself from sending and receiving frames. However, it can be relevant if it affects an SM because a TT network needs a minimum number of SM for clock synchronization. For TT ES an attacker with host access might also be able to disable or erase the value of the local clock (DS.7). This is implementation-dependent and only possible if the ES design allows it. A special denial-of-service attack is threat T.D3. In this threat scenario, an attacker can physically damage a network switch via a remote attack by exploiting the fact that flash memory has a finite number of erase cycles. An attacker with physical access to connections and cables can drop all or only specific data in transit (T.D10, T.D11, T.D12). This is a threat to any network system and can range from simple attacks like cutting network cables to attacks that need extra equipment with the capability to selectively drop frames. Redundancy can help to make this attack less effective. An attacker with physical access and extra equipment can also delay network frames (T.D13). This threat can impact especially time-triggered networks because their use case is to provide communication with guaranteed minimum latency.

Delayed TT traffic will most likely be dropped. Delaying PCFs on multiple routes can lead to a loss of synchronization in the network, in a worst-case scenario. BE and RC traffic can also be delayed but it depends on the network congestion if frames are dropped.

Threat	**Element**	**BE**	**RC**	**TT**	**CS**
T.E1	P.2	X			
T.E2	P.2	X	X	X	

Table 5.58.: Analysis results for elevation-of-privilege threats

Table 5.58 lists all identified threats in the category elevation-of-privilege. A typical Ethernet network has either no network segmentation or VLANs are used to segment the network. However, it can happen that more network devices as necessary have access to one VLAN. Therefore an attacker might gain access to privileged segments of the network (T.E1). RC and TT networks use VLs for network segmentation. Every VL consists of one sender and only the receivers that need to receive the network frame for their operation. This means that for network segmentation, the important security principle need-to-know is an intrinsic property of how RC and TT networks work. Threat T.E2 is another attack to bypass network segmentation. In this threat, an attacker exploits a weakness of the physical layer to perform a packet-in-packet attack to inject arbitrary network frames into privileged network segments. This is a threat to any Ethernet network with a physical layer that is susceptible to a packet-in-packet attack like EtherOops.

The analysis also shows that an attacker does not care about safety or security guarantees or the failure hypothesis a network system was designed for. An attacker will attack multiple switches and end systems if possible. It is therefore not possible to define an upper bound of infected components.

6. Security Measures

This chapter describes potential security measures that could mitigate or at least reduce potential security threats and weaknesses identified in chapter 5 Threat Modeling. The measures address all three traffic classes. The measure could be a change in the configuration of the network system, a change of the network topology, a protocol change or a change in the design and implementation of the network devices. For every proposed security measure, a list of the threats it aims to mitigate is provided.

6.1. M.1 - Static Configuration

Data networks used in a particular aerospace, industrial or automotive use case are designed once and are then deployed to many identical systems. For example, for one variant of a car model, one and the same network configuration can be used for millions of cars. These quantities are not achieved in the aerospace industry, but due to the nature of this industry, the network topology and what data packages have to be sent between devices are known years in advance before the first flight takes place. In addition, these networks are not changing or expanding in the field without proper planning done ahead of time. This is in stark contrast to company networks where no company network is like another, and company networks are constantly evolving. A typical enterprise Ethernet network has several mechanisms to decrease the workload of network administrators. STP can for example automatically establish a logical loop-free network topology that can automatically react to topology changes. The capability that STP offers is not needed and is even undesirable for networks used in safety-critical applications. Building a network system that can automatically cause a topology change literally on the fly is a bad idea from a safety point of view, in the context of an aerospace use case. But these plug-and-play convenience features to automate network configuration can also affect the security of a network system. Various attacks on STP exist that can be exploited by an attacker. An effective security measure to mitigate all of these threats is to statically configure the network. If the network topology does or should not change, STP is not needed. Instead, the network topology can be statically configured when the network is designed. The same is true for functionalities that enable the network to automatically learn MAC addresses of new devices, like ARP and network switches that can learn MAC addresses. There is no need for ARP and dynamic MAC address tables, if all network devices and their MAC addresses are known in advance. A static configuration can enable capabilities that a typical Ethernet network does not have. A statically configured network can for example use all redundant paths and their bandwidth loop-free. Most automated mechanisms would deactivate redundant paths and hence lose bandwidth. It has to be noted that protections for these automatic functionalities exist, but they have their own disadvantages. Removing these capabilities decreases the complexity and attack surface. For example, no software and open network ports are needed anymore to handle BPDU messages for STP. This security measure is only applicable for BE traffic because RC and TT systems only support a static configuration.

Using static configuration is an effective security measure to mitigate the following threats:
T.S1, T.S2, T.S3, T.T4, T.D5

6.2. M.2 - MACsec

MACsec is a security standard that adds confidentiality, integrity and authenticity to OSI layer 2 communication. MACsec is standardized in IEEE 802.1AE [57]. With MACsec, it is possible to encrypt a network frame and check if it is authentic. The standard defines GCM-AES-128 as the default cipher suite. This means the Advanced Encryption Standard (AES) cryptographic cipher with a key length of 128-bit is used in Galois/Counter Mode (GCM) for encryption and message authentication. GCM needs one symmetric key and a unique initialization vector (IV) that must never be reused [58]. The IV comprises a packet number that is incremented with every network frame that gets encrypted. The packet number counter is 32-bit long and re-keying needs to be performed every time before this counter overflows. Issuing a new key is also necessary after powering up or resetting to avoid the re-use of a key and IV combination. The re-keying mechanism is not specified in IEEE 802.1AE. It is therefore not enough to implement only IEEE 802.1AE for MACsec, but a key exchange mechanism is needed as well. IEEE 802.1X [59] specifies the MACsec Key Agreement protocol (MKA) that is used as the de facto standard in combination with MACsec. Every MACsec device needs a shared secret that is used by MKA to derive a new key for the re-keying mechanism. Several established ways to provision MACsec devices with this shared secret exist, but for the type of applications where TTEthernet is used, a pre-shared key (PSK) is most likely the best option. Key management and distribution via a centralized service like Remote Authentication Dial-In User Service (RADIUS) is not desirable because it adds complexity, increases dependency, requires more resources, increases the time after power-up or reset to become operational and lacks the necessary redundancy. MACsec has the feature to “bypass” traffic based on EtherType or VLAN-ID. This allows to mix unprotected and MACsec protected communication on the same network. To increase the security, only MACsec traffic should be allowed to prevent an attacker with physical access to inject plain text network frames, i.e. all non-MACsec traffic shall be dropped by the network devices. The use of MACsec comes with several disadvantages and limitations. MACsec increases the power consumption for every network device. This might be an issue for applications with a limited power budget like in the aerospace sector. TTEthernet, ARINC 664 and standard Ethernet operate without the need for any secrets. Cryptographic functions need key material and with them comes the need to manage these keys. Keys are a supporting asset for MACsec and they need to be protected from information disclosure threats. The initial key exchange performed by the MKA that is needed after every power-up or reset increases the time before the first network frame can be sent. This increased delay might be unacceptable for certain applications. The document AUTOSAR CP R22-11 [60] defines a statically configured MKA with only a subset of features suitable for automotive applications. IEEE 802.1X requires the system to provide a random number generator (RNG) for the key derivation function (KDF) that is used to derive new keys. A TTEthernet network is a very deterministic system, hence implementing a deterministic RNG is probably the only option, but even a deterministic RNG needs sufficient entropy to create a good quality seed [61]. A weak RNG can therefore weaken the security of MACsec. This can increase the implementation and testing effort for MACsec in low entropy embedded systems or increase the hardware cost if a non-deterministic hardware RNG is used instead. MACsec (in its basic form) does not offer end-to-end encryption and only provides protection between network ports. This means that the network frames are protected on the wire but are unprotected inside the network devices.

MACsec offers effective protection against attacks on OSI layer 1, like an attacker with physical access to the network cables. The encryption protects the confidentiality of network frames against information disclosure threats and message authentication protects the integrity and authenticity of network frames against tampering threats. MACsec is also an effective measure against packet-in-packet attacks like EtherOops and PCspooF because these threats enable the manipulation of data on the wire and not inside the network device. The physical interface between the host and ES is not protected by MACsec because the ES would typically perform the data encryption and decryption. MACsec cannot protect against DoS threats like dropping or delaying network frames. It is important to mention that although MACsec offers data delay protection, it is not suitable to protect traffic with tight latency requirements.

MACsec is an effective security measure to mitigate the following threats:
T.S6, T.T8, T.T9, T.I3, T.E2

6.3. M.3 - Keyless Remote Attestation

Remote Attestation allows a third party to verify the configuration of a system, such as settings, software, and hardware. The security analysis has identified many threats that can occur if the configuration of network devices is manipulated. This is especially true for the configuration of the Synchronization Masters (SM) because they provide the time that is needed to synchronize the TTEthernet network. The synchronization mechanism can tolerate only a certain number of rogue SMs that depend on the total number of SMs in the network system. A TTEthernet network typically has four SMs and this is also the case for the example topology displayed in Figure 4.2, that was used for threat modeling. If two of the four SMs fail silently, that means they stop sending PCFs, it would not affect the clock synchronization for the rest of the network. A rogue SM that is controlled by an attacker can also fail arbitrarily by sending malicious PCFs. In this case, the clock synchronization mechanism of the example network can only tolerate one rogue SM. To generate malicious PCFs, an attacker can either send PCFs from the host or have them generated by the ES itself by changing its configuration. The network's resistance to rogue SMs could therefore be increased if the Compression Master (CM) can detect that the ES configuration has been manipulated. Remote attestation can be used for exactly this purpose.

Yet attestation requires provisioning and managing of key material and implementing cryptographic functions. Both are not needed for TTEthernet. However, under certain conditions, it is possible to implement remote attestation without the need for keys and asymmetric cryptography. In order to achieve the low latency and jitter performance that TTEthernet offers, most or all of the critical functions are implemented in hardware and not in software. The functionality of TTEthernet end systems and switches is built with digital logic gates inside an FPGA or ASIC. An ASIC is an application-specific integrated circuit that was built to perform a specific function, for example, a TTEthernet switch, and it cannot be changed. The functionality of a field-programmable gate array (FPGA) can be changed, but this typically requires physical access to re-flash the FPGA bitstream. Some FPGAs can be locked to a specific bitstream to improve their integrity. In practice, this means the logic inside FPGA or ASIC cannot be changed via a remote attack. This property can be used to implement the keyless remote attestation functionality because the hardware can be trusted to perform the necessary operation as designed. The end systems that act as SMs append their current ES configuration to every PCF they send. To save bandwidth it is not necessary to send the whole ES configuration. The ES can either calculate a hash of the ES configuration and send the digest or only append the bits of the configuration that are actually security relevant instead. To save power, it is also not necessary to re-calculate the attestation data for every PCF but only when the ES configuration has changed. If the

FPGA or ASIC has a test or debug mode, then this information shall also be part of the attestation data. The payload of a PCF has a size of 28 bytes but 9 bytes plus 4 bits of them are reserved for future use. The minimum payload length of Ethernet is 46 bytes. If the payload is shorter than that the remaining bytes are filled with padding. This means a PCF has 27 bytes plus 4 bits, in total 220 bits, that are unused at the moment. These 220 bits can be used for the attestation data without affecting the throughput. The message digest of a suitable hash function like SHA-265 is 256 bit long [62] but it can be truncated to a length that provides enough security strength for this use case. NIST document SP 800-107 [63] provides guidance and recommendations on how to truncate a message digest. If a hash is used for attestation, it must be a cryptographic hash function with the property that it is not feasible to find another input that matches the hash value of the ES configuration within the lifetime of the devices.

A practical implementation of the keyless remote attestation may look as follows. The SMs shall have an unmodifiable logic that calculates the attestation data and appends it to every PCF that is sent into the network. The attestation data can be different for every SM, but the calculation is deterministic and does not change if the ES configuration does not change. The attestation data can therefore be pre-calculated in the design phase. The pre-calculated attestation data shall be stored in the SW configuration together with the network port the associated SM is connected to. The SW configuration shall be able to store two attestation values per SM. This is necessary to enable ES configuration updates in the field. Before an ES configuration can be updated, the configuration of the network switches must be updated first with the attestation data of the current ES configuration and the attestation data of the new ES configuration. With their traffic policing and frame filtering functions, network switches already play an important role in enforcing the security attributes of a TTEthernet network. Every network switch shall verify the attestation data inside the PCF at its ingress ports. If the attestation data from the PCF does not match one of the two associated attestation values stored in the SW configuration the PCF shall be dropped. Additional requirements must be fulfilled to prevent an attacker from bypassing the attestation procedure. PCFs are generated by the PCF Dispatcher function on the end system and not by the host. A host shall therefore not be able to send a valid PCF into the network. This can for example be implemented by ensuring that only the PCF Dispatcher can send frames with EtherType "0x891D" that is used for PCFs. It shall not be possible for an attacker with only remote access to manipulate the calculation of the attestation data. This can for example be achieved if the calculation logic is implemented in hardware. Many Ethernet PHYs have built-in frame generators for testing purposes [64]–[67]. Some Ethernet PHYs even allow the generation and sending of valid network frames with arbitrary data [67]. If this is the case, a host shall not be able to generate a valid PCF via the Ethernet PHY.

The advantage of this novel approach is that it does not require the implementation of cryptographic functions and there is no need to manage keys. It can protect against remote attacks on the clock synchronization where an attacker was able to gain access to several SMs, but not more than the total number of SMs minus the configured minimum number of SMs required.

The disadvantage of this approach is that it does not help against adversary-in-the-middle attacks, where an attacker is able to inject PCFs between SM and CM, because the attestation data is static and a switch cannot verify the integrity and authenticity of a received PCF. It also does not protect against physical threats if an attacker can manipulate the attestation calculation or bypass it completely. However, an attacker needs to manipulate PCFs from at least half the SMs to impact the clock synchronization of the network.

Keyless Remote Attestation is a security measure that can improve the security of TT traffic against the threats **T.T13, T.T15** and is an effective measure to mitigate: **T.T16, T.T19**

6.4. M.4 - Remote Attestation

Remote attestation allows a third party to verify the configuration of a system. It relies on asymmetric cryptography and a private key. The private key needs to be stored on the attested device and must be kept secret. This security measure can be used to verify the ES configuration of the SMs in a similar way as the keyless remote attestation described in section 6.3. The big advantage of using public-key cryptography over the previously mentioned keyless method is that the attestation data can be protected by a signature. The Edwards-curve Digital Signature Algorithm (EdDSA) [68], like Ed25519, is a good choice for this application because it does not require any randomness to create a signature [69]. A verifier can use the public key to confirm the integrity and authenticity of the attestation data. The signature itself is not enough to provide a significant security improvement over the keyless method because the signed data is still static, which makes it vulnerable to replay attacks. To mitigate this the attestation data must be protected while in transit with additional security improvements like authenticated PCFs described in section 6.5.

A disadvantage of this security measure is that it is necessary to implement cryptographic functions and the private keys must be protected on the devices.

Keyless Remote Attestation is a security measure that can improve the security of TT traffic against the following threats: **T.T8, T.T9, T.T13, T.T15** and is an effective measure to mitigate: **T.T16, T.T19**

6.5. M.5 - Authenticated PCFs

TTEthernet devices exchange PCFs to synchronize their clocks, this is necessary for TT communication. This makes PCFs a primary target for attacks. Authenticated PCFs can improve security because the receiver can verify if the frame was manipulated and which device has generated the frame.

MACsec as described in section 6.2 provides these capabilities but has some limitations for this use case. MACsec was created to protect network frames in transit between devices, but not for end-to-end communication. In this case, network frames are decrypted and unprotected inside all network switches on the path between transmitter and receiver. Therefore, the receiver cannot verify the origin of the PCF. Another disadvantage is that MACsec uses symmetric keys. This means every device that needs the key is also able to generate a valid message authentication code. This defeats the proof of data origin as well. A MACsec unaware network switch can still forward MACsec frames because MACsec does not decrypt the header of the MAC frame. This can be used as a workaround to solve the proof of data origin problem. To achieve that, only the sender and the receiver need to share the same key but not the hops in between. The receiver still needs to know the key, but the receiver is also the verifier. This increases the effort to manage the keys because now a different key is needed for every connection.

Another way to implement this security measure is to cryptographically sign every PCF that is sent. To implement this, asymmetric cryptography and a private key is needed on every network device that sends PCFs. The signature algorithm EdDSA is a good choice for the reasons stated in section 6.4. Calculating the signature takes time and can vary with every calculation depending on the signature algorithm. TTEthernet uses precalculated values that are stored in the network configuration for the propagation delay of its network components like wires, switches and end systems. The additional network delay caused by the cryptographic operation is therefore not a problem as long as it is constant. The PCF Dispatcher shall always wait a fixed time period before it sends the signed PCF. This time period has to be long enough to calculate the signature in the worst-case scenario. The signature can be placed at the end of the payload and shall protect the entire network frame and not only the payload. The devices that receive the PCF can verify its origin with the associated public key of the transmitter and can account for the additional latency. To mitigate replay attacks it is also necessary to implement a challenge-response mechanism that provides a nonce from the verifier, a persistent counter, or a mixture of both. With asymmetric cryptography a private key is only needed in the device that calculates the signature, hence it is less exposed and easier to keep secret than symmetric keys. A disadvantage of this security measure is that it is necessary to implement cryptographic functions and the private keys must be protected on the devices. It also increases the complexity and time to calculate the signature as well as power consumption because a new signature must be calculated for every PCF that is sent.

Authenticated PCFs can improve the security against **T.T8, T.T9** and are an effective security measure to mitigate the following threats: **T.S4, T.S5, T.S6**

6.6. M.6 - Cryptographically Signed Configuration

The security analysis has shown that manipulated configurations of network devices can have a major impact on the security of a network system. It is therefore imperative that the correct untampered configuration is applied. Network switches and end systems shall verify the cryptographic signature of their configuration before they load it and only proceed if the signature is valid. It is recommended that the signature check is performed in hardware to achieve the best protection. A disadvantage of this security measure is that it is necessary to implement cryptographic functions and to handle the keys. The network device only needs to store the public key that is used to verify the signature and not a private key, thus it is not necessary to protect the public key from information disclosure threats. Therefore, no additional procedures are needed for the decommissioning of a device. Another disadvantage is that it would not be possible to directly write to individual registers storing the configuration in the SW or ES anymore.

Cryptographically Signed Configuration is an effective security measure to mitigate the following threats: **T.T1, T.T2, T.T6, T.T7, T.T10, T.T16, T.I2, T.D14**

6.7. M.7 - Optimize SM placement and selection

The chosen network topology can have an impact on the security of a network. This security measure provides recommendations as to which ES should be promoted to the SM role to improve the security of the clock synchronization. To mitigate adversary-in-the-middle attacks the SM shall be directly connected to the CM without other network devices in between. This does not protect against an attacker with physical access to the network wiring but against remote attacks. In addition, preferably only ESs with minimal attack surface should be promoted to SMs. Network routers or gateways are examples of devices with large attack surfaces because these devices typically connect network segments with different privileges and interact with many devices in the network. Devices that provide many network services and have a lot of open ports also have larger attack surfaces. The ideal SM is a rather simple network device that has only outgoing VLs, such as a sensor. A simpler device typically has a lower failure rate, this is also advantageous from a safety point of view. Promoting devices that are dissimilar can be an advantage to avoid common vulnerabilities between SMs.

An optimized placement and selection of SMs in the network can be effective in mitigating the following threat: **T.S4**

6.8. M.8 - Synchronize PHY and MAC

Packet-in-packet attacks like EtherOops and PCSpooF exploit the fact that some Ethernet PHYs will immediately forward every data they receive on their media-independent interface (MII) to the media-dependent interface (MDI) even if the PHY establishes a link on its MDI side in the middle of a current MII transmission.

This type of attack works because the MAC is unaware of the link state of the PHY, which allows frame fragments to be transmitted in case a link is re-established on the PHY during an ongoing frame transmission of the MAC. An attacker can cause a link loss for example with EMI. This requires physical access or at least physical proximity to place an electronic device, that can generate a strong enough electromagnetic pulse (EMP) to cause a link loss on the targeted PHY.

The deficiency of the PHY can be mitigated by the MAC. Ethernet PHYs have a management interface like Management Data Input/Output (MDIO) or Inter-Integrated Circuit (I2C) that is used to configure and read its status information [36], [70]. The current link state is also readable from the PHY registers. The MAC shall read this register periodically and disable the MII transmission to the PHY as long as the link is down. Polling the PHYs needs to be faster than the PHY specific minimum link build-up time to ensure that the MAC can disable its transmission fast enough.

Synchronize PHY and MAC is an effective security measure to mitigate the following threats:
T.E2, T.S6

6.9. M.9 - Vote on redundant frames

Every ARINC 664 compliant RC network frame contains a frame counter. Redundant frames are frames that share the same counter value that is then used by the receiver to match the frames for redundancy management. The ARINC 664 standard defines that the first valid frame that arrives at a receiver wins over a redundant frame that is received at a later point in time. All redundant frames that are received after the first frame, within a configurable time period called *SkewMax*, on redundant channels are dropped. If a redundant frame arrives very late, after the *SkewMax* period has ended, it counts as a new frame. Therefore, even with access to a single channel, a malicious frame sent by an attacker can always win over genuine frames sent on the other redundant channels if the malicious frame is sent first. TT traffic uses synchronous redundancy management as defined in the AS6802A standard. A synchronized receiver knows the worst-case point in time when redundant frames will arrive and stores all redundant frames that are received within this worst-case reception window, but the standard does not define which of the redundant frames shall be forwarded to the host.

A mitigation for this threat is to wait for the time period *SkewMax* after a frame was received, then compare all the redundant frames that have been received within the time period and count the number of identical frames. Identical frames mean identical copies, if a single bit or more differs they are not identical. For TT traffic, the worst-case reception window shall be used instead of *SkewMax*. The receiver shall perform a vote and the frame with the highest count wins. All frames with equal or lower counts are dropped. If several frames have been received on the same channel, only the first frame is counted. This prevents an attack scenario where an attacker sends the malicious frames several times to win the vote. If only one channel received a valid frame, this frame wins the vote as well. This increases the availability but decreases the integrity. A disadvantage is that this approach increases the latency by the time *SkewMax*. If for example two different frames with the same frame counter are received within *SkewMax* on a dual-channel network, both frames will be dropped, because neither frame can win the vote.

Another possibility to avoid the problem is to protect the traffic with a message authentication code. With that functionality, the receiver could detect if the integrity and origin of a frame are authentic and drop it if it is not. MACsec can be used for that but for maximum effectiveness, the authentication must happen end-to-end, so only the sender and receiver share the secret key. Even better would be the use of asymmetric cryptography for message authentication because then only the sender needs to possess the secret key. Using a signature or message authentication code has the disadvantage that it is necessary to implement, cryptographic functions and to manage the keys.

Vote on redundant frames is an effective security measure to mitigate the following threat:
T.T18

6.10. M.10 - Restrict access to the frame memory

The frame memory stores the received network frames on a network switch until they are forwarded to the network. The only component inside a network switch that needs read and write access to this frame memory to perform its operation is the switch MAC. However, some MACs allow other components inside the switch, like a management CPU, to access the frame memory [56]. This is a bad security design because it infringes on the principle of least privilege. The frame memory shall only be accessible by the switch MAC and the MAC shall not provide frame memory access to other components to minimize the attack surface area.

Restricting access to the frame memory of a network switch is an effective security measure to mitigate the following threats: **T.T3, T.I1**

6.11. M.11 - Perform traffic policing on ES ingress ports

RC ESs do not enforce a data rate limit on the network traffic they receive. The same is true for TT ESs, which do not perform an acceptance window check, because traffic policing is implemented on the network switches only. Although the MAC of the ES can operate at full line speed on all of its redundant ports, this might not be the case for the host interface or the host itself. Especially small embedded systems like sensors or actuators do not need the full bandwidth of 100BASE-TX Ethernet to perform their tasks. An attacker can try to saturate the host interface by sending as many network frames as possible to the ES. This can also affect network traffic from redundant channels if they share the same host interface. An ES shall therefore be able to limit the data rate of ingress traffic to protect its host interface from saturation and resource exhaustion of the host. This feature enables the network designer to limit the total combined bandwidth from the network ports to not exceed the maximum bandwidth of the host interface.

Performing traffic policing on ES ingress ports is an effective security measure to mitigate the following threats: **T.D6**

6.12. Out of Scope Security Measures

The mitigation of some of the identified threats requires changes to components that are out of the scope of this thesis. For example, many threats require an attacker to gain access to the management interface of a network switch to be of concern. Discussing the protection of management interfaces is well outside the scope of OSI Layer 2 security improvements. The security of a host and its network services are out of scope as well. The same is true for software and firmware that is running on the network devices.

The following threats require security measures on components that are out of the scope of this work:
T.T5, T.D2, T.D3, T.D9

6.13. Residual Threats

Table 6.1 provides an overview of all the identified security threats and measures to mitigate them. A threat can be mitigated by one of several security measures, this means one measure is enough to be effective. Measures that are marked with an "X" indicate that these measures cover all aspects of a specific threat. The "~" indicates that a security measure is only effective to mitigate a specific case. For example, the security measure M.5 implements authentication for PCFs and this ensures that an attacker with physical access to the cabling can not tamper the frame unnoticed, but this protects only PCFs and not all traffic.
Residual threats are threats where no technical mitigations can be identified and are therefore out of scope. Table 6.1 shows that there are 14 residual threats, these threats are: T.T11, T.T12, T.T14, T.T17, T.I4, T.D1, T.D4, T.D7, T.D8, T.D10, T.D11, T.D12, T.D13, TE.1.

Threat	**Affected Element**	**M.1**	**M.2**	**M.3**	**M.4**	**M.5**	**M.6**	**M.7**	**M.8**	**M.9**	**M.10**	**M.11**	**Out of Scope Measures**
T.S1	P.1	X											
T.S2	P.2	X											
T.S3	EE.1	X											
T.S4	P.4, P.8, P.12					X		X					
T.S5	P.15, P.16					X							
T.S6	P.15, P.16		X			X			X				
T.T1	DS.1						X						
T.T2	DS.2						X						
T.T3	DS.3										X		
T.T4	DS.4	X											
T.T5	DS.4												X
T.T6	P.1, P.3, P.5, P.6, P.7, P.9, P.14, P.15, P.16						X						
T.T7	P.2, P.4, P.8, P.10, P.11, P.12, P.13						X						
T.T8	DF.1		X		~	~							
T.T9	DF.1		X		~	~							
T.T10	DF.2						X						
T.T11	DF.3												
T.T12	DF.3												
T.T13	DS.5, DS.6			~	~								
T.T14	P.4												
T.T15	P.5, P.6, P.7			~	~								
T.T16	DS.7			X	X		X						
T.T17	P.11, P.14												
T.T18	P.1									X			
T.T19	P.8, P.12			X	X								
T.I1	DS.3										X		
T.I2	P.2						X						
T.I3	DF.1		X										
T.I4	DF.3												
T.D1	DS.1												
T.D2	DS.2												X
T.D3	DS.2												X
T.D4	DS.3												
T.D5	DS.4	X											
T.D6	P.1											X	
T.D7	P.1, P.3, P.5; P.6, P.7, P.9, P.14, P.15, P.16												
T.D8	P.1, P.7												
T.D9	P.2, P.4, P.8, P.10, P.11, P.12, P.13												X
T.D10	DF.1												
T.D11	DF.2												
T.D12	DF.3												
T.D13	DF.1												
T.D14	DS.7						X						
T.E1	P.2												
T.E2	P.2		X						X				

Table 6.1.: Matrix of identified security threats and suggested security measures

7. Conclusion

The security analysis carried out in chapter 5 has identified 33 threats for BE networks, 26 threats for RC networks and 37 threats for TT networks. However, comparing only these absolute figures of threats per traffic class with each other can lead to a misleading conclusion, because the threats are not always the same and the severity of the impact is also different from threat to threat. Similar threats could also be treated as one threat or individually, which can affect and skew the number of threats. To draw a conclusion, it is therefore necessary to focus on the individual threats and their impacts. This section deals with the key findings, but a detailed evaluation of the analysis results can be found in section 5.4.

A major outcome of the security analysis is that protecting the configuration of the network devices is key to the security of the network system. This is true for all traffic classes. The best mitigation that was identified to counter threats targeting the integrity of the configuration is to have configuration data that is cryptographically signed and that a network device verifies the signature before applying the configuration (M.6). Another weakness that applies to all traffic classes is denial-of-service threats where an attacker gains physical access to the network. Even low-skill attacks, like cutting the network wiring, are difficult to defend with only technical measures. Organizational security measures would be needed, but they are outside the scope of this thesis. MACsec (M.2) was identified as an appropriate security measure against an attacker with physical access to the network wiring, for all traffic classes to protect the integrity, authenticity and confidentiality of network frames but it is not effective against remote attacks.

BE traffic has turned out to be the most insecure of the analyzed traffic classes. A typical IEEE 802.3 Ethernet network uses many plug-and-play convenience features, like ARP and STP that help to automatically configure the network, but these features increase the attack surface because this plug-and-play functionality can also be abused by an attacker. The easiest and most effective security measure to mitigate this category of threats is to disable these functionalities and configure the network statically (M.1). RC and TT traffic do not suffer from these threats because they only use static configuration. Effective network segmentation is another area where the security of RC and TT traffic is sublime over BE traffic. RC and TT networks are segmented with the help of VLs and additionally, there is only one sender per VL and only the receivers that are configured to receive the frame. The principle of least privilege is therefore an intrinsic security property of TTEthernet and ARINC 664. VLANs are an established way to segment Ethernet networks but the segmentation cannot be performed as granular and strictly as with RC and TT networks because the latter only support unidirectional communication while VLANs typically allow bidirectional communication [71]. The biggest weakness of BE networks and the reason why ARINC 664 and TTEthernet were developed, is the lack of Quality of Service. The safety feature of RC traffic to constrain the data rate per VL and to limit when TT traffic can be sent also improves the security of ARINC 664 and TTEthernet and is an effective and robust security feature against DoS attacks. This is not the case for regular Ethernet.

RC traffic has the lowest number of identified threats because it has the lowest attack surface of the analyzed traffic classes due to its design. The switches of RC and TT networks perform traffic policing and frame filtering and are therefore critical for the security of the network. In contrast to TT networks, RC switches are not dependent on other systems to perform these important security functions, because they do not rely on external functions for traffic policing, such as the global time, as TTEthernet does. But ARINC 664 also has its weaknesses. For example, the way how deduplication of redundant frames is designed can be exploited by an attacker. ARINC 664 uses a "first frame wins" policy for deduplication. This means an attacker who can inject traffic to only one channel will always overrule genuine traffic on the other channels, as long as traffic from the attacker is sent first. A voting mechanism (M.9) can be used to mitigate this threat, but this mitigation will increase network latency. This is not the case for TTEthernet because TT traffic is already synchronized between channels to perform the voting without adding additional latency.

As TT traffic is dependent on a functioning clock synchronization, additional functionalities are required that in turn lead to additional threats. The higher amount of threats in comparison to ARINC 664 is a result of the necessary clock synchronization mechanism; on the other hand, the resulting time-triggered nature of TTEthernet traffic also offers additional security advantages. For instance, TTEthernet can segment the network not only in space but also in time. The threat analysis also showed that it is not sufficient to only consider attacks within the failure hypothesis because any vulnerability may be exploited whether it is inside or outside the failure hypothesis.

Although redundancy can significantly improve the safety of a network system, the analysis shows that this only contributes little to security, because the same or similar redundant devices also share the same vulnerabilities. Even if an attacker only has access to a single channel, threats exist for ARINC 664 and TTEthernet that target the redundancy of channels. Nevertheless, redundancy in the clock synchronization system can improve the security of the network. Multiple redundant clock synchronization masters (SM) can limit the impact of malicious PCFs. This can be further improved if malicious PCFs can be detected by the TT devices. One method to confirm the authenticity of a PCF is to sign it (M.5). Security can be improved further if the compression masters (CM) can detect a manipulated SM. Remote attestation can be used by the CM as a security measure (M.4) to detect if the configuration of an SM has changed. The proposed solution is to append the attestation data to every PCF. This design would change the content of a PCF but is backward compatible because there is enough unused space left in the frame for the attestation data. A novel approach to attestation called *Keyless Remote Attestation* (SM.3) was developed in this thesis that does not require cryptographic functions or keys. Several security measures have been suggested to improve the resilience of the clock synchronization function but even if all measures would be applied, there are still limits in defending all conceivable attack scenarios. For example, if all the SMs in the network are compromised and are sending malicious or no PCFs at all, TT communication is made impossible.

A detailed description of all the additional security measures developed and proposed in this thesis can be found in chapter 6. The descriptions of the individual security measures provide concrete requirements, denoted with the term "shall", and guidance on how to implement them. These critical network systems should work highly predictable and deterministic, but this property will yield only little entropy, therefore cryptographic algorithms that do not need random numbers should be used where possible. In addition, the advantages and disadvantages of a measure are described as well. For example, for some of the additional security measures it is necessary to handle key material which can increase threats in the category of information disclosure. It was also discovered that some of the well-known security measures like MACsec are not well suited for critical applications. One obstacle is the additional time before the first frame can be sent. It will increase due to design decisions that have been made for MACsec and its key exchange mechanism.

Table 7.1 lists the total number of identified threats per traffic class and the remaining threats if all the suggested mitigations are implemented. This includes the mitigations that are out of the scope of this analysis as well. The table also shows the number of residual threats per STRIDE category.

Traffic class	**Identified threats**	**Residual threats**[1]	**S**	**T**	**R**	**I**	**D**	**E**
BE	33	11	0	2	0	1	7	1
RC	26	9	0	2	0	1	6	0
TT[2]	37	12	0	4	0	1	7	0

Table 7.1.: Number of threats per traffic class with and without mitigations

[1] Threats for which no technical mitigation was found within the scope of this thesis.

[2] This also includes threats to the clock synchronization mechanism.

7.1. Limitations

The conducted security analysis is only valid under the scope defined in section 4.1 and the assumptions listed in section 4.3. Although a structured approach to threat modeling was used to model the system and to identify the threats, there is no guarantee that all threats have been found. No threat analysis method can assure that because new attack vectors will likely be discovered in the future. In addition, the chosen detail of the system model can change the analysis result. Models are intentional abstractions from reality and should represent reality well enough for its purpose. The AS6802A standard also defines a second operating mode for the clock synchronization mechanism called *dual-failure hypothesis* that requires high-integrity SMs. However, in this thesis, only a *single-failure hypothesis* network was analyzed. This thesis only analyses the network protocols and systems as they are defined in the various standards and does not take bugs and implementation differences to the standards into account. The identified threats and weaknesses have not been tested and verified on real hardware.

7.2. Future Work

This thesis focuses on the security of TTEthernet, ARINC 664 and regular Ethernet because these are the protocols that are supported by the TTEthernet devices currently on the market. A comparison with other competing network protocols for mixed critical networks, such as TSN, would be of additional interest.

The proposed security measures only exist as a theoretical concept at the moment. Future work is needed to build and implement these measures to verify their effectiveness and viability in real-world systems.

A. Detailed description of specific threats

Threat ID: T.T4 **Name:** MAC address spoofing

Description:

For every given MAC Address only one entry is allowed in the SAT. If a network switch receives a network frame with an already known MAC address, the entry gets updated. Dynamically learned entries have an aging time. If the time has elapsed since the entry was created or last updated, then the entry gets automatically removed from the SAT. [72]

In order to perform this attack, the attacker needs to know the MAC address of its target. With this knowledge, the attacker can then send an Ethernet frame with this MAC address set as the source address. The switch will learn that the MAC address is now connected to the network port that is connected to the device controlled by the attacker. Network frames that are sent to the target device will get forwarded to the attacker.

BE:

This threat is an attack on the integrity of the SAT inside the network switch. This attack can affect confidentiality (network sniffing) and can be a precursor for an adversary-in-the-middle attack.

RC:

The switches in an ARINC 664 network are using a predefined static SAT and are therefore not susceptible to this attack.

TT:

The switches in a TTEthernet network are using a predefined static SAT and are therefore not susceptible to this attack.

Threat ID: T.D5 **Name:** MAC flooding (SAT overflow)

Description:

The difference between a network switch and a network bridge is that a network switch forwards unicast network frames not to all ports but only to the port that is connected to the network device with the corresponding destination MAC address. Therefore, a network switch needs to know the MAC addresses that are connected to it per port. An Ethernet switch can dynamically learn new MAC addresses by inspecting the source MAC address of an incoming Ethernet frame and storing this data in the Source Address Table (SAT). An entry in the SAT consists of a MAC address and a port number. The SAT has a finite number of entries. If the switch receives an Ethernet frame with a destination MAC address that is not yet present in the SAT, then the network switch acts like a bridge and the Ethernet frame is forwarded to all ports. This also happens if the number of MAC addresses exceeds the maximum number of entries in the SAT. [72]

This behavior can be exploited by an attacker by generating many Ethernet frames with various source MAC addresses thus flooding the SAT of the Ethernet switch. This effectively downgrades the switch to a bridge and the attacker can now sniff all network traffic.

BE:

This threat is an attack on the integrity of the SAT inside the network switch. This attack can affect confidentiality (network sniffing) and can be a precursor for an adversary-in-the-middle attack. Because all traffic is now forwarded to all switch ports, the traffic on the network will increase, and this can lead to packet loss and high latency and therefore affect the availability.

RC:

The switches in an ARINC 664 network are using a predefined static SAT and are therefore not susceptible to this attack.

TT:

The switches in a TTEthernet network are using a predefined static SAT and are therefore not susceptible to this attack.

Threat ID: T.E1	**Name:** Packet-in-packet attack (EtherOops)

Description:

This threat can be used to inject valid network frames into privileged network segments that are otherwise not accessible by an attacker. It only works in one direction, the exfiltration of network frames from the privileged network segment is not possible with the packet-in-packet attack on its own.

To perform this attack, it is necessary to send a manually crafted network frame with a payload that is controlled by the attacker to the network switch that enforces this network segmentation, and the network frame must be forwarded to network ports connected to a segmented network. The payload contains a preamble, the start frame delimiter (SFD), another MAC header and a payload. This means the network frames or network packet contains another network packet in its payload, hence the name of this attack. To get a valid inner frame it is important that the Cyclic Redundancy Check (CRC) of the outer network frame must be the same as the CRC of the inner frame because the tail of the network frame is shared by both frames. An Ethernet broadcast frame as the outer frame is a good choice as a carrier for the inner frame because it is typically forwarded to all switch ports. In addition, the attacker has to be able to inject electromagnetic interference into the network cable connected to or in close proximity to this network switch. This can for example be accomplished by installing a hardware implant in the form of a high voltage generator into a network device that is directly connected to the network switch. This high-voltage pulse must be strong enough to disrupt the operation of the network switch. It is important that the EMI does not cause permanent physical damage to the network devices, because they need to work after the EMI.

There are at least two different methods to cause a packet-in-packet attack, but their success is dependent on the actual implementation of the network device. The first method is called the runaway preamble attack. In this case, the EMI is used to interfere with the SFD of the fabricated network frame. The EMI has to change the SFD of the outer MAC header into a different bit pattern after the network frame is forwarded by the network switch. The attack is successful if the receiver that is part of the privileged network segment treats not the outer SFD but the inner SFD as the actual start of the frame. This attack does not work if the MAC of the receiver does not allow random data in the preamble before the SFD [73]. Independent tests showed that none of the tested TTEthernet devices are susceptible to this attack because they drop frames if non-preamble data is in the preamble [35].

The second method is called a link reset attack. At the start of the attack, the EMI is used to cause a link loss on Ethernet PHYs at several Ethernet ports of the network switch before the fabricated network frame is sent to the network switch.

It is important that the link of the switch port connected to the device that sends the fabricated network frame gets reestablished first, before the other ports of the switch ports connected to the privileged network segments. The fabricated network frame is sent immediately after the link is re-established to the network switch. The switch MAC receives the frame and forwards it to the other switch ports. The attack is successful if the PHYs on the other switch ports reestablishes their links while the preamble in the payload is sent to the PHYs. In this case, the outer frame is cut off and the PHYs only send the inner frame to the receiver. This attack is only successful if the Ethernet PHY is susceptible to this attack. Tests showed that many PHYs will start to send immediately after their link is reestablished even during an ongoing transmission from the MII side [35].
BE: This threat can enable an attacker to gain unauthorized access to a privileged part of an Ethernet network. VLANs are typically used to segment Ethernet networks. An attacker can use the packet-in-packet attack to inject arbitrary frames into a VLAN by setting the VLAN ID of the inner header to the VLAN ID of the targeted VLAN.
RC: The threat can be used to bypass the network segmentation of an ARINC 664 network by using different VL IDs for the outer and inner headers of the injected frame. It is therefore possible to hop from one VL to the targeted privileged VL.
TT: The threat can be used to bypass the network segmentation of a TTEthernet network. Like ARINC 664, TTEthernet also uses VL IDs to segment the network.

<table>
<tr><td><b>Threat ID:</b> T.S3</td><td><b>Name:</b> ARP spoofing (also known as ARP cache poisoning)</td></tr>
<tr><td colspan="2"><b>Description:</b>
The Address Resolution Protocol (ARP) is an OSI layer 2 network protocol that is used to resolve the MAC address (layer 2/link layer address) for a given IPv4 address (layer 3/internet layer address). Before a network device sends a network packet, it checks if the destination MAC address of the receiver is known and present in the MAC table. If this is not the case, the network device uses ARP to get the correct MAC address for a specific IPv4 address. In this case, the network device sends an ARP request as a broadcast to all devices on the same network. The network device that matches the IPv4 address in the ARP request will send an ARP reply message. This information is added as a new entry in the ARP table. The authenticity of ARP messages is not protected, therefore an attacker on the same network can spoof any ARP message with the intention of the attacker to change the ARP table of the attacked network devices [74].
An ARP-Spoofing attack can be performed as follows. The ARP protocol allows a network device to send an ARP message without being asked. This gratuitous ARP message can be used by the attacker to change the ARP tables of the other devices on the network. An attacker can use this attack to redirect traffic to the target device by announcing that the IPv4 address of the target device is now associated with the MAC address of the attacker device. Every network packet for the target device will now be redirected to the network device that is controlled by the attacker. The attacker can read, change, or drop any received message and send it to the target device to complete the adversary-in-the-middle attack [75].</td></tr>
<tr><td colspan="2"><b>BE:</b>
This attack is a serious threat to an Ethernet network that relies on ARP. The attack can affect the confidentiality (network sniffing), availability (block message) and integrity (transmitted data manipulation) of the network data.</td></tr>
<tr><td colspan="2"><b>RC:</b>
In an ARINC 664 network the destination MAC address is used to address a specific network path that can contain one or several receivers. These network paths are defined when the network is designed and are therefore static. ARP is not used and has no use in an ARIC 664 network.</td></tr>
<tr><td colspan="2"><b>TT:</b>
ARP is not used and has no use in a TTEthernet network for the same reason as for ARINC 664.</td></tr>
</table>

Threat ID: T.S6	**Name:** PCspooF
Description: TTEthernet relies on a common time base that is distributed to all network devices that are part of the time-triggered communication. This is accomplished by the clock synchronization protocol defined in AS6802A [9] and described in section 2.2. TTEthernet is reliant on PCFs, because a single PCF, like a cold start acknowledgment frame, can trigger the resynchronization of a synchronization master or client. The resynchronization affects all channels. Unsynchronized devices cannot communicate with each other. This threat uses the packet-in-packet attack described in Table A.3 to inject malicious PCFs via a link reset attack. To perform this attack, it is necessary to send a network frame with a payload that is controlled by the attacker to the TTEthernet switch that acts as a compression master and the TTEthernet switch has to forward this frame to time-triggered network devices. A hardware implant that can generate EMI needs to be directly connected to the compression master to trigger the link reset. The EMI source can, for example, be integrated into a network device that is connected to the CM. This EMI has to be strong enough to cause a link loss on several Ethernet ports of the TTEthernet switch, but must not cause permanent damage or disruption to the TTEthernet switch, because a working switch is necessary to perform (and repeat) this attack. To execute the PCspooF attack, the attacker has for example managed to manipulate ES6 with the necessary EMI hardware implant and is also able to send network frames with a malicious PCF in its payload from the end system. As shown in Figure 4.2, ES6 is only connected to the compression master of channel B. At first, the manipulated end system hardware transmits a high voltage pulse that produces enough EMI (Electro Magnetic Interference) inside the TTEthernet switch to cause a link loss of its PHYs. The more ports of the switch experience a link loss the more effective the attack can be. To cause resynchronization of the entire TTEthernet network more than half of the compression masters need to receive the malicious PCF. The malformed network frame must be sent immediately after the PHY the switch reestablishes a link with ES6. The link of the switch port connected to the device controlled by the attack must get reestablished first, before the other ports of the switch. This adds some uncertainty to the attack. The attack can be repeated to increase the success rate and to increase the duration of the synchronization loss. This attack shows that the individual OSI layers are not independent of each other. The lower layers can affect the security of the upper layers. The PCspooF attack exploits the fact that, although with the redundancy of multiple network channels, the time base is shared over all network channels. This means that the network channels are not fully independent. [35]	

BE:

Ethernet has no clock synchronization mechanism and is therefore not susceptible to this threat.

RC:

ARINC 664 has no clock synchronization mechanism and is therefore not susceptible to this threat.

TT:

Time-triggered communication is not possible without a working and reliable clock synchronization. This threat will lead to a loss of availability.

List of Figures

2.1 Network topology of ARINC 429 [6] 6

2.2 PCF format [9] 9

4.1 Key to the symbols used in DFD 14

4.2 Mixed network topology [55] 16

5.1 DFD of best-effort traffic 18

5.2 DFD of rate-constrained traffic 19

5.3 DFD of time-triggered traffic 20

5.4 DFD of clock synchronization mechanisms in synchronized operation 21

5.5 DFD of clock synchronization mechanisms in un-synchronized operation 23

List of Tables

4.1 STRIDE-per-Element allocation [51] 15

5.1 Identified Security Assets 17

5.2 Identified threats of data store DS.1 ES configuration for BE communication 25

5.3 Identified threats of data store DS.2 SW configuration for BE communication 26

5.4 Identified threats of data store DS.3 Frame Memory for BE communication 27

5.5 Identified threats of data store DS.4 SAT for BE communication 28

5.6 Identified threats of process P.1 ES-MAC for BE communication 29

5.7 Identified threats of process P.2 Switch-MAC for BE communication 30

5.8 Identified threats of external entity EE.1 Host for BE communication 31

5.9 Identified threats of data flow DF.1 BE frame for BE communication 31

5.10 Identified threats of data flow DF.2 write ES Configuration for BE communication 32

5.11 Identified threats of data flow DF.3 for BE communication 33

5.12 Identified threats of data store DS.1 ES configuration for RC communication 33

5.13 Identified threats of data store DS.2 SW configuration for RC communication 34

5.14 Identified threats of data store DS.3 Frame Memory for RC communication 34

5.15 Identified threats of process P.1 ES-MAC for RC communication 35

5.16 Identified threats of process P.2 Switch-MAC for RC communication 36

5.17 Identified threats of process P.3 Traffic Shaping Function for RC communication 37

5.18 Identified threats of process P.4 Traffic Policing & Frame Filtering Function for RC traffic . . 37

5.19 Identified threats of data flow DF.1 RC frame for RC communication 38

5.20 Identified threats of data flow DF.2 write ES Configuration for RC communication 39

5.21 Identified threats of data flow DF.3 for RC communication 39

5.22 Identified threats of data store DS.1 ES configuration for TT communication 40

5.23 Identified threats of data store DS.2 SW configuration for TT communication 40

5.24 Identified threats of data store DS.3 Frame Memory for TT communication 41

5.25 Identified threats of data store DS.5 Synchronized Local Clock for TT communication . . . 41
5.26 Identified threats of data store DS.6 Synchronized Local Clock for TT communication . . . 42
5.27 Identified threats of process P.1 ES-MAC for TT communication 42
5.28 Identified threats of process P.2 Switch-MAC for TT communication 43
5.29 Identified threats of process P.4 Traffic Policing & Frame Filtering Function for TT traffic . . 44
5.30 Identified threats of process P.5 Time-Triggered Dispatcher ES for TT communication 45
5.31 Identified threats of process P.6 Time-Triggered Dispatcher SW for TT communication . . . 46
5.32 Identified threats of data flow DF.1 TT frame for TT communication 47
5.33 Identified threats of data flow DF.2 write ES Configuration for TT communication 47
5.34 Identified threats of data flow DF.3 for TT communication 48
5.35 Identified threats of data store DS.1 ES configuration for clock synchronization 48
5.36 Identified threats of data store DS.2 SW configuration for clock synchronization 49
5.37 Identified threats of data store DS.5 Synchronized Local Clock for clock synchronization . . 49
5.38 Identified threats of data store DS.6 Synchronized Local Clock for clock synchronization . . 50
5.39 Identified threats of data store DS.7 Local Clock for clock synchronization 50
5.40 Identified threats of process P.7 PCF Dispatcher ES for clock synchronization 51
5.41 Identified threats of process P.4 Traffic Policing & Frame Filtering Function for clock sync. . 52
5.42 Identified threats of process P.8 Compression Function for clock synchronization 53
5.43 Identified threats of process P.11 Acceptance window check SW for clock synchronization . 54
5.44 Identified threats of process P.10 Clock Synchronization & Compressed PCF Dispatcher . . 55
5.45 Identified threats of process P.12 Clique Detection & Resolution Service SW 56
5.46 Identified threats of process P.13 Startup & Restart Service SW for clock synchronization . . 56
5.47 Identified threats of process P.14 Acceptance window check ES for clock synchronization . . 57
5.48 Identified threats of process P.9 Clock Synchronization Service 58
5.49 Identified threats of process P.15 Clique Detection & Resolution Service ES 59
5.50 Identified threats of process P.16 Startup & Restart Service ES 60
5.51 Identified threats of data flow DF.1 PCF for clock synchronization 61
5.52 Identified threats of data flow DF.2 write ES Configuration for clock synchronization 62
5.53 Analysis results per traffic class 63
5.54 Analysis results for spoofing threats 63
5.55 Analysis results for tampering threats 64
5.56 Analysis results for Information Disclosure threats 65

5.57 Analysis results for Denial-of-Service threats 66
5.58 Analysis results for elevation-of-privilege threats 68

6.1 Matrix of identified security threats and suggested security measures 81

7.1 Number of threats per traffic class with and without mitigations 85

Glossary

A	Asset
ACK	Acknowledgment
AES	Advanced Encryption Standard
AFDX	Avionics Full-Duplex Switched Ethernet
ARINC	Aeronautical Radio, Incorporated
ARP	Address Resolution Protocol
ASIC	Application-Specific Integrated Circuit
ATT&CK	Adversarial Tactics, Techniques, and Common Knowledge

BAG	Bandwith Allocation Gap
BE	Best-Effort traffic class
BPDU	Bridge Protocol Data Units

CAPEC	Common Attack Pattern Enumeration and Classification
CM	Compression Master
COTS	Commercial-Of-The-Shelf
CPU	Central Processing Unit
CRC	Cyclic Redundancy Check
CSA	Canadian Space Agency
CT	Critical Traffic

DFD	Data Flow Diagram
DoS	Denial of Service

ECSS	European Cooperation for Space Standardization
EdDSA	Edwards-Curve Digital Signature Algorithm
EMI	Electromagnetic Interference
EMP	Electromagnetic Pulse
ES	End System
ESA	European Space Agency
FPGA	Field-Programmable Gate Array
GCM	Galois/Counter Mode
I2C	Inter-Integrated Circuit
IASIS	International Avionics System Interoperability Standards
ICS	Industrial Control System
ID	Identifier
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IN	Integration
IV	Initialization Vector
JAXA	Japan Aerospace Exploration Agency
KDF	Key Derivation Function
LAN	Local Area Network
LRU	Line-Replaceable Unit
M	Security Measure

MAC	Media Access Control
MDI	Media-Dependent Interface
MDIO	Management Data Input/Output
MII	Media-Independent Interface
MKA	MACsec Key Agreement protocol
NASA	National Aeronautics and Space Administration
NIST	National Institute of Standards and Technology
OSI	Open Systems Interconnection
PCF	Protocol Control Frame
PHY	Physical layer
PSK	Pre-Shared Key
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
RC	Rate-Constrained traffic class
RNG	Random Number Generator
SAE	Society of Automotive Engineers
SAT	Source Address Table
SC	Synchronization Client
SDF	Start Frame Delimiter
SHA	Secure Hash Algorithm
SM	Synchronization Master
SPARTA	Space Attack Research & Tactic Analysis
STP	Spanning Tree Protocol

STRIDE	Spoofing Tampering Repudiation Information disclosure Denial of service Elevation of privilege
SW	Switch
TCP	Transmission Control Protocol
TSN	Time-Sensitive Networking
TT	Time-Triggered traffic class
TTEthernet	Time-Triggered Ethernet
TTP	Tactic, Techniques, and Procedures
UDP	User Datagram Protocol
UTC	Coordinated Universal Time
VL	Virtual Link
VLAN	Virtual Local Area Network

Bibliography

[1] H. Kopetz and G. Bauer, "The time-triggered architecture," *Proceedings of the IEEE*, vol. 91, no. 1, pp. 112–126, 2003. DOI: `10.1109/JPROC.2002.805821`.

[2] Pierre Bieber, Frédééric Boniol, Marc Boyer, Eric Noulard, and Claire Pagetti, "New Challenges for Future Avionic Architectures," *Aeropsacelab Journal*, vol. 04, May 2012. [Online]. Available: `https://www.researchgate.net/publication/244484691_New_Challenges_for_Future_Avionic_Architectures`.

[3] Till Steinbach, Hyung-Taek Lim, Franz Korf, Thomas Schmidt, Daniel Herrscher, and Adam Wolisz, "Tomorrow's In-Car Interconnect? A Competitive Evaluation of IEEE 802.1 AVB and Time-Triggered Ethernet (AS6802)," Sep. 2012. DOI: `10.1109/VTCFall.2012.6398932`.

[4] José Simó, Patricia Balbastre, Juan Francisco Blanes, José-Luis Poza-Luján, and Ana Guasque, "The Role of Mixed Criticality Technology in Industry 4.0," *Electronics*, vol. 10, no. 3, 2021, ISSN: 2079-9292. DOI: `10.3390/electronics10030226`. [Online]. Available: `https://www.mdpi.com/2079-9292/10/3/226`.

[5] Mitch Fletcher, "Progression of an Open Architecture from Orion to Altair and LSS," *Honeywell International*, May 2009. [Online]. Available: `https://www.academia.edu/20896448/Open_systems_architecture_Both_boon_and_bane`.

[6] Christian Fuchs, "The Evolution of Avionics Networks From ARINC 429 to AFDX," *Innovative Internet Technologies and Mobile Communications and Aerospace Networks (IITM & AN)*, vol. 65, Jan. 2012. [Online]. Available: `https://www.net.in.tum.de/fileadmin/TUM/NET/NET-2012-08-1/NET-2012-08-1_09.pdf`.

[7] *ARINC 664P7-1 Aircraft Data Network, Part 7, Avionics Full-Duplex Switched Ethernet Network*, en, Standard, ARINC. [Online]. Available: `https://www.aviation-ia.com/products/664p7-1-aircraft-data-network-part-7-avionics-full-duplex-switched-ethernet-network-2`.

[8] (2022). Time-Sensitive Networking (TSN) Task Group. Accessed: 2023-01-04, IEEE, [Online]. Available: https://1.ieee802.org/tsn.

[9] *AS6802A - Time-Triggered Ethernet*, en, Standard, SAE, Nov. 2023. [Online]. Available: https://www.sae.org/standards/content/as6802a.

[10] E. Webb, "Ethernet for space flight applications," vol. 4, Feb. 2002, pp. 4–1927, ISBN: 0-7803-7231-X. DOI: 10.1109/AERO.2002.1036905. [Online]. Available: https://ntrs.nasa.gov/citations/20020039726.

[11] (2018). Time-Triggered Ethernet Slims Down Critical Data Systems. Accessed: 2023-01-04, NASA, [Online]. Available: https://spinoff.nasa.gov/Spinoff2018/t_4.html.

[12] (2022). P802.1DG – TSN Profile for Automotive In-Vehicle Ethernet Communications. Accessed: 2023-01-04, IEEE, [Online]. Available: https://1.ieee802.org/tsn/802-1dg/.

[13] (2022). IEC/IEEE 60802 TSN Profile for Industrial Automation. Accessed: 2023-01-04, IEEE, [Online]. Available: https://1.ieee802.org/tsn/iec-ieee-60802.

[14] (2022). P802.1DP – TSN for Aerospace Onboard Ethernet Communications. Accessed: 2023-01-04, IEEE, [Online]. Available: https://1.ieee802.org/tsn/802-1dp.

[15] Petr Grillinger, "Time-Triggered Ethernet History, Present, and Evolution," *ČESKO-SLOVENSKÝ SEMINÁŘ*, p. 9, [Online]. Available: https://www.kiv.zcu.cz/pad2013/template/sbornikpad2013.pdf#page=17.

[16] Andrew Loveless, "On TTEthernet for Integrated Fault-Tolerant Spacecraft Networks," *AIAA SPACE Forum and Exposition*, Aug. 2015. [Online]. Available: https://ntrs.nasa.gov/citations/20150014489.

[17] ——, "On Time-Triggered Ethernet in NASA's Lunar Gateway," *NASA Avionics Architectures Community of Practice, Houston, TX*, Jul. 2020. [Online]. Available: https://ntrs.nasa.gov/citations/20205005104.

[18] Wilfried Steiner, Guenther Bauer, Brendan Hall, and Michael Paulitsch, "Time-Triggered Ethernet," in. Jan. 2011, pp. 181–220. [Online]. Available: https://www.researchgate.net/publication/330310738_Time-triggered_ethernet.

[19] (2022). Time-Triggered Ethernet. Accessed: 2023-01-04, ESA, [Online]. Available: https://www.esa.int/Enabling_Support/Space_Engineering_Technology/Onboard_Computers_and_Data_Handling/Time-Triggered_Ethernet.

[20] Timothy Cichan, Kerry Timmons, Kathleen Coderre, and Willian D. Pratt, "The Orion Spacecraft as a Key Element in a Deep Space Gateway," Jul. 2017, Accessed: 2023-01-04. [Online]. Available: https://www.lockheedmartin.com/content/dam/lockheed-martin/eo/photo/webt/Orion-Spacecraft-as-a-Key-Element-to-Deep-Space.pdf.

[21] (2022). NASA's Orion Spacecraft - Systems Integration with Deterministic Ethernet. Accessed: 2023-01-04, TTTech, [Online]. Available: https://www.tttech.com/sites/default/files/documents/TTTech-NASA-Casestudy-Orion.pdf.

[22] (2016). Time-Triggered Ethernet. Accessed: 2023-01-04, ESA, [Online]. Available: https://www.esa.int/ESA_Multimedia/Images/2016/11/Time_triggered_Ethernet.

[23] (2022). European launcher Ariane 6 - Creating a single, redundant avionics network for all data traffic with TTEthernet®. Accessed: 2023-01-04, TTTech, [Online]. Available: https://www.tttech.com/sites/default/files/documents/202212-TTTech-Aerospace-casestudy-Ariane6.pdf.

[24] Christopher Adamek, "Gateway System Requirements," Jun. 2019, Accessed: 2023-05-15. [Online]. Available: https://ntrs.nasa.gov/citations/20190029153.

[25] Woodrow Bellamy III. (2015). TTEthernet Avionics Backbone a Technology Breakthrough for S-97 Raider. Accessed: 2023-01-04, Access Intelligence, LLC, [Online]. Available: https://www.aviationtoday.com/2015/07/20/ttethernet-avionics-backbone-a-technology-breakthrough-for-s-97-raider.

[26] Nick Flaherty. (2022). Vestas taps TTTech for world's first 15 MW wind turbine. Accessed: 2023-01-04, European Business Press SA, [Online]. Available: https://www.eenewseurope.com/en/vestas-taps-tttech-for-worlds-first-15-mw-wind-turbine.

[27] *International Avionics System Interoperability Standards (IASIS)*, en, Standard, National Aeronautics and Space Administration, Mar. 2019. [Online]. Available: https://nasasitebuilder.nasawestprime.com/wp-content/uploads/sites/45/2019/09/avionics_baseline_final_3-2019.pdf.

[28] Wilfried Steiner, "Candidate security solutions for TTEthernet," in *2013 IEEE/AIAA 32nd Digital Avionics Systems Conference (DASC)*, 2013, 4A5-1-4A5-10. DOI: 10.1109/DASC.2013.6712579.

[29] *DO-326 - Airworthiness Security Process Specification*, en, Standard, RTCA, Aug. 2010. [Online]. Available: `https://my.rtca.org/productdetails?id=a1B36000001IcfwEAC`.

[30] *IEC TS 62443-1-1:2009 - Industrial communication networks - Network and system security - Part 1-1: Terminology, concepts and models*, en, Standard, IEC, Jul. 2009. [Online]. Available: `https://webstore.iec.ch/publication/7029`.

[31] *SAE J3061 - Cybersecurity Guidebook for Cyber-Physical Vehicle Systems*, en, Standard, SAE, Jan. 2016. [Online]. Available: `https://doi.org/10.4271/J3061_201601`.

[32] Nicolas Falliere, Liam O Murchu, and Eric Chien, "W32. stuxnet dossier," *White paper, symantec corp., security response*, vol. 5, no. 6, p. 29, 2011. [Online]. Available: `https://pax0r.com/hh/stuxnet/Symantec-Stuxnet-Update-Feb-2011.pdf`.

[33] Ruben Santamarta, "Last Call for SATCOM Security," *Black Hat USA 2018*, Aug. 2018, Accessed: 2023-01-04. [Online]. Available: `https://ioactive.com/wp-content/uploads/2018/08/us-18-Santamarta-Last-Call-For-Satcom-Security-wp.pdf`.

[34] ——, "Arm IDA and Cross Check: Reversing the Boeing 787's Core Network," *Black Hat USA 2019*, Aug. 2019, Accessed: 2023-01-04. [Online]. Available: `https://ioactive.com/arm-ida-and-cross-check-reversing-the-787s-core-network/`.

[35] A. Loveless, L. T. X. Phan, R. Dreslinski, and B. Kasikci, "PCspooF: Compromising the Safety of Time-Triggered Ethernet," *IEEE Symposium on Security and Privacy (S&P)*, 2023. [Online]. Available: `https://web.eecs.umich.edu/~barisk/public/pcspoof.pdf`.

[36] "IEEE Standard for Ethernet," *IEEE Std 802.3-2022 (Revision of IEEE Std 802.3-2018)*, pp. 1–7025, 2022. DOI: `10.1109/IEEESTD.2022.9844436`.

[37] MOREAUX JEAN-PAUL, "Data transmission system for aircraft," US6925088B1, 2005. [Online]. Available: `https://worldwide.espacenet.com/patent/search/family/007928755/publication/US6925088B1?q=pn%3DUS6925088`.

[38] *ARINC 429P1-19 Digital Information Transfer System (DITS), Part 1, Functional Description, Electrical Interfaces, Label Assignments and Word Formats*, en, Standard, ARINC. [Online]. Available: `https://www.aviation-ia.com/products/429p1-19-digital-information-transfer-system-dits-part-1-functional-description-electrical`.

[39] TTTech. (). Aerospace Products List. Accessed: 2024-01-31, TTTech, [Online]. Available: `https://www.tttech.com/aerospace/products`.

[40] Timo Kiravuo, Mikko Sarela, and Jukka Manner, "A Survey of Ethernet LAN Security," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1477–1491, 2013. DOI: `10.1109/SURV.2012.121112.00190`.

[41] Martin Lang, *Secure Automotive Ethernet : Balancing Security and Safety in Time Sensitive Systems*, 2019. [Online]. Available: `http://bth.diva-portal.org/smash/record.jsf?pid=diva2%3A1332281`.

[42] "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," *IEEE Std 1588-2019 (Revision ofIEEE Std 1588-2008)*, pp. 1–499, 2020. DOI: `10.1109/IEEESTD.2020.9120376`.

[43] Elena Lisova, "Securing Clock Synchronization in Industrial Heterogeneous Networks," PhD thesis, Mälardalen University, 2016. [Online]. Available: `https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A916720`.

[44] Tal Mizrahi, "Time synchronization security using IPsec and MACsec," in *2011 IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication*, 2011, pp. 38–43. DOI: `10.1109/ISPCS.2011.6070153`. [Online]. Available: `https://www.researchgate.net/publication/236123600_Time_synchronization_security_using_IPsec_and_MACsec`.

[45] Karen Seo and Stephen Kent, *Security Architecture for the Internet Protocol*, RFC 4301, Dec. 2005. DOI: `10.17487/RFC4301`. [Online]. Available: `https://www.rfc-editor.org/info/rfc4301`.

[46] Doğanalp Ergenç, Cornelia Brülhart, Jens Neumann, Leo Krüger, and Mathias Fischer, "On the Security of IEEE 802.1 Time-Sensitive Networking," in *2021 IEEE International Conference on Communications Workshops (ICC Workshops)*, 2021, pp. 1–6. DOI: `10.1109/ICCWorkshops50388.2021.9473542`.

[47] Haris Isakovic, "A secure global time base for time triggered systems," PhD thesis, 2011. [Online]. Available: `https://resolver.obvsg.at/urn:nbn:at:at-ubtuw:1-44562`.

[48] MITRE. (). Common Attack Pattern Enumerations and Classifications. Accessed: 2023-05-15, MITRE, [Online]. Available: `https://capec.mitre.org/`.

[49] ——, (). Adversarial Tactics, Techniques, and Common Knowledge. Accessed: 2023-05-15, MITRE, [Online]. Available: `https://attack.mitre.org/`.

[50] SPARTA. (). Space Attack Research & Tactic Analysis. Accessed: 2023-05-15, Aerospace Corporation, [Online]. Available: `https://sparta.aerospace.org/`.

[51] Adam Shostack, *Threat modeling: Designing for security*. John Wiley & Sons, 2014, ISBN: 978-1-118-80999-0.

[52] Mounir Kamal, "ICS Layered Threat Modeling," *SANS Institute*, Jan. 2019. [Online]. Available: `https://www.sans.org/white-papers/38770/`.

[53] Amit Kumar Gupta, *Security risk analysis of automotive ethernet networks*, Oct. 2017. [Online]. Available: `http://essay.utwente.nl/73894/`.

[54] Anton Bretting and Mei Ha, "Vehicle Control Unit Security using Open Source AUTOSAR," 2015. [Online]. Available: `https://hdl.handle.net/20.500.12380/219822`.

[55] *ECSS-E-ST-50-16C - Time-Triggered Ethernet*, en, Standard, European Cooperation for Space Standardization, Sep. 2021. [Online]. Available: `https://ecss.nl/get_attachment.php?file=2021/09/ECSS-E-ST-50-16C(30September2021).pdf`.

[56] *TTE Switch Controller Space*, TT6802-3-SW-C, TTTech. [Online]. Available: `https://www.tttech.com/sites/default/files/documents/Flyer_TTE-SW%20Controller%20Space_D-TTE-ASIC-E-04-004_V4.1.pdf`.

[57] "IEEE Standard for Local and Metropolitan Area Networks-Media Access Control (MAC) Security," *IEEE Std 802.1AE-2018 (Revision of IEEE Std 802.1AE-2006)*, pp. 1–239, 2018. DOI: `10.1109/IEEESTD.2018.8585421`.

[58] *NIST SP 800-38D - Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC*, en, Standard, NIST, Nov. 2007. [Online]. Available: `https://doi.org/10.6028/NIST.SP.800-38D`.

[59] "IEEE Standard for Local and Metropolitan Area Networks–Port-Based Network Access Control," *IEEE Std 802.1X-2020 (Revision of IEEE Std 802.1X-2010 Incorporating IEEE Std 802.1Xbx-2014 and IEEE Std 802.1Xck-2018)*, pp. 1–289, 2020. DOI: `10.1109/IEEESTD.2020.9018454`.

[60] "Specification of MACsec Key Agreement," *AUTOSAR CP R22-11*, pp. 1–76, 2022. [Online]. Available: `https://www.autosar.org/fileadmin/standards/R22-11/CP/AUTOSAR_SWS_MACsecKeyAgreement.pdf`.

[61] *NIST SP 800-108 Rev. 1 - Recommendation for Key Derivation Using Pseudorandom Functions*, en, Standard, NIST, Aug. 2022. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-108r1.

[62] *FIPS 180-4 - Secure Hash Standard (SHS)*, en, Standard, NIST, Aug. 2015. [Online]. Available: https://doi.org/10.6028/NIST.FIPS.180-4.

[63] *NIST SP 800-107 Rev. 1 - Recommendation for Applications Using Approved Hash Algorithms*, en, Standard, NIST, Aug. 2012. [Online]. Available: https://doi.org/10.6028/NIST.SP.800-107r1.

[64] *Integrated 10/100/1000 Mbps Energy Efficient Ethernet Transceiver*, 88E1514, Rev. G, Marvell, Jan. 2023. [Online]. Available: https://www.marvell.com/content/dam/marvell/en/public-collateral/phys-transceivers/marvell-ethernet-phys-alaska-88e151x-datasheet.pdf.

[65] *1000BASE-T1 Automotive Ethernet PHY with SGMII and RGMII*, DP83TG720S, Rev. E, Texas Instuments, Nov. 2022. [Online]. Available: https://www.ti.com/lit/ds/symlink/dp83tg720s-q1.pdf.

[66] *Robust, Industrial, Low Latency and Low Power 10 Mbps, 100 Mbps, and 1 Gbps Ethernet PHY*, ADIN1300, Rev. A, Analog Devices, Aug. 2019. [Online]. Available: https://www.analog.com/media/en/technical-documentation/data-sheets/adin1300.pdf.

[67] *Single Port Gigabit Ethernet Copper PHY with RGMII/RMII Interfaces*, VSC8531XMW-05, Rev. 4.1, Microsemi, Feb. 2019. [Online]. Available: https://ww1.microchip.com/downloads/en/DeviceDoc/VMDS-10514.pdf.

[68] *FIPS 186-5 - Digital Signature Standard (DSS)*, en, Standard, NIST, Feb. 2023. [Online]. Available: https://doi.org/10.6028/NIST.FIPS.186-5.

[69] Daniel Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang, "High-Speed High-Security Signatures," vol. 2, Sep. 2011, pp. 124–142, ISBN: 978-3-642-23950-2. DOI: 10.1007/978-3-642-23951-9_9.

[70] *INF-8074i Specification for SFP (Small Formfactor Pluggable) Transceiver*, en, Standard, SFF Committee, May 2001. [Online]. Available: https://members.snia.org/document/dl/26184.

[71] "IEEE Standard for Local and Metropolitan Area Networks–Bridges and Bridged Networks," *IEEE Std 802.1Q-2022*, pp. 1–2163, 2022. DOI: 10.1109/IEEESTD.2022.10004498.

[72] "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges," *IEEE Std 802.1D-2004 (Revision of IEEE Std 802.1D-1998)*, pp. 1–281, 2004. DOI: 10.1109/IEEESTD.2004.94569.

[73] Ben Seri, Gregory Vishnepolsky, and Yevgeny Yusepovsky, "EtherOops Bypassing Firewalls and NATs By Exploiting Packet-in-Packet Attacks in Ethernet," Aug. 2020. [Online]. Available: https://info.armis.com/rs/645-PDC-047/images/Armis-EtherOops-TWP-20200805-1.pdf.

[74] *An Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware*, RFC 826, Nov. 1982. DOI: 10.17487/RFC0826. [Online]. Available: https://www.rfc-editor.org/info/rfc826.

[75] Sean Whalen, "An introduction to arp spoofing," *Node99 [Online Document]*, vol. 563, 2001. [Online]. Available: https://api.semanticscholar.org/CorpusID:59638215.