UTRECHT UNIVERSITY

MASTER'S THESIS

Authentication and Authorization for the Internet of Things for Health

Author:
Jarno BREDENOORD

Supervisors:
Drs. Lennart HERLAAR
Dr. Fabiano DALPIAZ

A thesis submitted in fulfillment of the requirements for the degree of Master of Science

in

Information Science
Department of Information and Computing Sciences

April 24, 2019

Utrecht University

UTRECHT UNIVERSITY

Abstract

Faculty of Science
Department of Information and Computing Sciences

Master of Science

Authentication and Authorization for the Internet of Things for Health

by Jarno BREDENOORD

Background: The Internet of Things (IoT) may transform health and other sectors but this requires (technological) solutions that ensure adequate security and privacy. Proper authentication and authorization are essential to achieve these goals, but may also harm these if the solutions for authentication and authorization are designed or configured improperly. **Goal:** The goal of this thesis project is to find out how authentication and authorization can help ensuring security and privacy for IoT devices in the health sector. **Method:** Current solutions for authentication and authorization in the IoT are analyzed. Using semi-structured interviews, requirements for authentication and authorization solutions for the IoT in health are presented. **Results:** The analysis of available solutions for authentication and authorization shows that there are many different models, architectures, and mechanisms for authentication and authorization, each having their own advantages and disadvantages. The results of the interviews show that the main objectives of authentication and authorization are related to privacy, confidentiality, and integrity of data. The most important challenges to achieve these objectives are heterogeneity and a lack of standardization, as well as problems related to managing (large amounts of) data. To achieve a desired level of security and privacy, authentication and authorization must offer transparency, anonymity / pseudonymity, unlinkability, unobservability, confidentiality, integrity, availability, usability, accountability, auditability, trustworthiness, and non-repuditation. A general set of guidelines for secure and privacy preserving authentication and authorization is proposed and validated. **Conclusion:** In health care, organizations are vulnerable to security and privacy threats. In some cases there is a trade-off between some security and privacy objectives. There is an orientation towards centralized IoT solutions. Potential negative effects for privacy are avoided through legal and organizational measures. Current trends such as virtualization of networks may affect the way authentication and authorization is carried out.

Contents

List of Abbreviations

ABAC	Attribute-Based Access Control
ADF	Access Control Decision Facility
AEF	Access Control Enforcement Facility
ALFA	Abbreviated Language For Authorization
CoAP	Constrained Application Protocol
DDoS	Distributed Denial of Service
DoS	Denial of Service
DTLS	Datagram Transport Layer Security
ECC	Elliptic Curve Cryptography
HTML	Hyper Text Markup Language
HTTP	Hypertext Transfer Protocol
ISO	The International Organization for Standardization
IoT	Internet of Things
IP	Internet Protocol
JSON	JavaScript Object Notation
M2M	Machine to Machine (communication)
PDP	Policy Decision Point
PEP	Policy Evaluation Point
PII	Personally Identifiable Information
PKC	Public Key Cryptography
PKI	Public Key Infrastructure
RBAC	Role-Based Access Control
REST	Representational State Transfer
RFID	Radio-Frequency Identification
RSA	Rivest–Shamir–Adleman
SDN	Software Defined Networking
SKC	Symmetric Key Cryptography
SSL	Secure Sockets Layer
TLS	Transport Layer Security
VNF	Virtual Network Function
VPN	Virtual Private Network
WoT	Web of Things
WSN	Wireless Sensor Network
XACML	eXtensible Access Control Markup Language
XML	eXtensible Markup Language

Chapter 1

Introduction

1.1 The Internet of Things

Even though there is no generally accepted definition of the Internet of Things (IoT) (Bojanova and Voas, 2017), it is often described as a paradigm which refers to a network of software, sensors, actuators, and micro controllers embedded in physical objects, or "things" (Bertino, 2016; Kouicem, Bouabdallah, and Lakhlef, 2018). A large number of different connected devices can be included in the IoT, such as smart devices, automobiles, fridges, smart phones, and health care or industrial appliances (Alaba et al., 2017). The number of connected things is growing; according to Ericsson, around 29 billion connected devices will be in use in 2022, of which 18 billion will be related to the IoT (Telefonaktiebolaget LM Ericsson, 2018).

The goal of the IoT is to connect physical objects using a network infrastructure with inter-operable communication protocols and software (Alaba et al., 2017). These devices then exchange data with (centralized) systems or other connected devices (Bertino, 2016). The IoT is expected to bring change to society in many different ways. In logistics, the IoT may be used to track temperature, locations, and movement of products and vehicles (Atzori, Iera, and Morabito, 2010). Connected devices may enable the power grid to become a "smart grid", i.e., to monitor and control the entire power grid chain from production to consumption (Bekara, 2014). IoT technologies may be applied on an urban level with the aim of supporting added-value services for the administration of cities and their citizens, for example in public spaces (Zanella et al., 2014). IoT technologies may also have large impact on health care. IoT devices may for example be used to monitor patients at home (AL-mawee, 2015).

The IoT is a relatively new paradigm that differs from its predecessors such as traditional internet, mobile internet, sensor networks and M2M networks in a few different ways. For example, because of its focus on widely available services and universal access on top of heterogeneous network architectures (Hussein, Bertin, and Frey, 2017). The IoT includes a large number of heterogeneous devices. As most of these devices can be connected to the internet, these devices often support common web technologies, such as HTTP, JSON, XML etc. An advantage of using these technologies is that they are well supported so they can be integrated with existing (non-IoT) infrastructure. There are also newer protocols, specifically tuned to constrained IoT environments such as the Constrained Application Protocol (CoAP) as alternative for HTTP.

1.1.1 IoT reference models

In the IoT, terminology is not standardized, and the domain which encompasses the IoT is sometimes vague (Rose, Eldridge, and Chapin, 2015). The following subsections aim to overcome this ambiguity and provide a further clarification of a simplified IoT environment. A multi-layered IoT system is common in both literature and practice (Rahmani et al., 2018). For example, sensors may gather a patient's heart activity and transmit data to so-called *gateway* devices. These gateways and sensors are physically close to each other, for example in the case of a smart routers with more computational resources in the same physical room as the patient. These gateways transfer the data to cloud databases for storage and a user application for analytics (Baker, Xiang, and Atkinson, 2017).

It may be useful to think about how elements in the IoT interact with each other. To do this, CISCO suggested a seven-layer generic IoT reference model. The goal of this model is to provide definitions and descriptions that can be applied accurately to elements and functions of IoT systems and applications. (CISCO, 2014). The lowest level *(edge nodes)* is the level that consists of computing nodes such as RFID tags and reader, smart controllers, and sensors. The second layer *(communication)* consists of all the components that enable transmission of information or commands between or within layers. On the third layer *(edge computing)* simple data processing is carried out locally to provide a fast response and reduce computational load. The fourth layer *(data accumulation)* enables conversion of data in motion to data at rest, for example for storage in databases. The fifth layer *(data abstraction)* provides the opportunity to render and store data such that future processing becomes more efficient or simpler. The sixth layer *(applications)* provides information interpretations. The last layer *(users and centers)* consists of users that make use of applications and their data (CISCO, 2014). The flow in the model is usually bidirectional, but the dominant direction of data flow depends on the application.

Mosenia and Jha (2017) present a three-layer simplification of CISCO's generic IoT reference model. The *edge nodes*, *communication*, and *edge computing* layers are combined into the *edge-side layer*. The *data accumulation* and *data abstraction* layers are combined into the *server / cloud-side layer*. Finally, the *applications* and *users and centers* layers are combined into the *user-side layer* (Mosenia and Jha, 2017). This model is used for the remainder of this work. The relationship between the CISCO model and the Mosenia and Jha model can be found in Figure 1.1.

1.1.2 IoT communication patterns

Besides studying the IoT based on the layers of which it is made up, it can also be studied from the communication that takes place within or between layers. To do this, and to speed up the design of internet-connected smart objects, the Internet Architecture Board (IAB) proposes four smart object communication patterns for the IoT (Tschofenig et al., 2015). These communication patterns are: Device-To-Device, Device-To-Cloud, Device-To-Gateway, and Back-End Data-Sharing. These communication patterns are explained below.

In the Device-To-Device pattern two or more devices connect and communicate directly with each other. This pattern is commonly used for applications which typically use small data packets and communicate between devices with low data rate requirements. For security, this usually means that these devices often have a direct

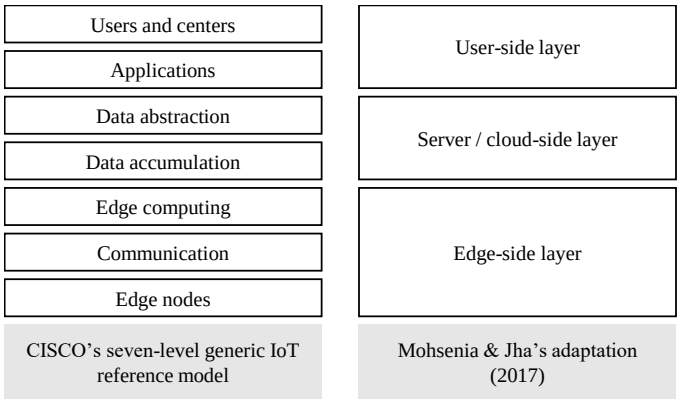

FIGURE 1.1: CISCO (2014) generic IoT reference model and the Mosenia and Jha, 2017 adaptation

relationship and have built-in security and trust mechanisms, but also use device-specific data models and potentially a limited number of supported communication protocols, which makes interoperability a challenge (Rose, Eldridge, and Chapin, 2015).

In the Device-To-Cloud pattern, devices at the edge of the network collect data and transmit these to a cloud server. This approach is usually based around existing IP-based communications mechanisms, such as wired Ethernet or WiFi (Rose, Eldridge, and Chapin, 2015). A potential downside of this pattern is vendor lock in because of specific communication protocols that are used on top of more generic protocols (e.g., IP).

In the Device-To-Gateway model, the "edge" device, or the device that is located at the edge of the network, connects through a so-called gateway service. This gateway acts as an intermediary between the device and the cloud or other devices and provides security or other functionaries. Devices that carry out these tasks are sometimes said to be in the "fog", because they are in between the edge devices and the cloud. Examples of gateway device can be smart phones or "hub" home devices or controllers. Modern smart-gateways have more computational resources and may carry out some tasks previously carried out by the sensor or cloud such as authentication, authorization, local storage, real-time local data processing, embedded data mining or data or protocol translation (e.g., if the smart objects require interoperability with non-IP devices) (Rahmani et al., 2018). This model is common for consumer devices, but also for medical devices which often do not have enough resources to send data to the cloud or make use of different protocols then cloud servers (Rahmani et al., 2018). In the Back-End Data-Sharing model, users export and analyze smart object data from a cloud service in combination with data from other sources. This architecture supports granting access to third parties.

1.1.3 The IoT in health care

There are many different opportunities for the IoT in health care. Combining IoT with other (digital) technologies, creates a new form of health care called *smart home health care* (Bennett, Rokas, and Chen, 2017), *E-health* (Scarpato et al., 2017), or *smart health* (Kang et al., 2018), or Internet of Things for Medical Devices (IoT-MD) (Solapurkar, 2016). These forms of health care refer to the provision of health care services at any time or place using different forms of information technology and health care

(Kang et al., 2018). For example, IoT devices can be used to monitor non-critical patients at home instead of in hospitals, reducing strain on resources while monitoring people continuously and enabling them to stay at home longer (Baker, Xiang, and Atkinson, 2017). Besides at home, medical IoT networks can also be applied in hospitals, doctor's offices, and even on a city level (Scarpato et al., 2017). The devices that are used can include (among others) wearables, diabetes sensors, heart rate monitoring sensors, oxygen saturation sensors, and pulmonary disease sensors (Scarpato et al., 2017). For the remainder of this work, the term *health care* refers to the provision of health services regardless of the type of health services and location of where these services are provided. The use of consumer electronics used for monitoring for fun or sports is not regarded as a health care service.

AL-mawee (2015) mentions that IoT can be used for a wide range of applications in the health domain, from preventing diseases, to managing chronic diseases and disabilities. This includes remote monitoring for people staying at home, early prevention to avoid future diseases or disabilities, and medical treatment of the institutionalized disabled who benefit from continuous monitoring and are living in nursing homes, long-term care facilities, or hospitals. Data that results from these or other sources can be used for gaining insights (analytics) and personalized medicine and precision medicine (Scarpato et al., 2017). Health care IoT-devices may collect, transmit and process private or confidential data, potentially without the user noticing (Conti et al., 2018). The increased adoption of these IoT-related devices and the confidential nature of the data collected therefore increases the need for appropriate (technological) security and privacy measures.

However, despite its importance, security and privacy remain two of the main challenges that the IoT faces today (Alaba et al., 2017; Conti et al., 2018). In fact, security is seen as one of the most important factors that obstructs further adoption of the IoT (Trnka, Cerny, and Stickney, 2018). There are a few key differences between traditional IT systems and IoT systems that make securing the IoT different (Bojanova and Voas, 2017). For example, IoT nodes generally have limited computational power and storage capacity, which makes public key encryption harder. Traditional security and privacy solutions can therefore not be easily applied within the IoT. IoT devices may have limited computational resources, which asks for scalable solutions that can provide security and are able to operate using limited computational resources (Conti et al., 2018). Numerous specialized tools, techniques, frameworks, certificates, algorithms and procedures specifically optimized to secure IoT networks have therefore been suggested (Alaba et al., 2017; Kouicem, Bouabdallah, and Lakhlef, 2018), such as lightweight encryption algorithms to secure IoT devices (Alaba et al., 2017). More lightweight solutions however, may be more easily compromised by attackers (Porambage et al., 2016) and still require proper implementation. As a result, security and privacy therefore remain serious challenges within the IoT.

1.2 Problem statement

In this work, the focus is on authentication (similarly known as identification) and authorization (sometimes also called access control) for the IoT for health care. Authentication is the process of verifying the identity of an entity. Authorization is the process of granting permission on specific actions to certain entities (Trnka, Cerny, and Stickney, 2018). Because the decision for allowing or denying access to a resource can only be made if the identity of an entity is confirmed, authentication is

often seen as an important prerequisite for authorization. These concepts play a key role in security and privacy as limiting access to resource to authorized entities is an important aspect of providing security and privacy. Due to the differences between traditional IT-solutions and the IoT, traditional solutions for authentication and authorization cannot easily be applied within the IoT. Authentication and authorization problems therefore remain a serious challenge in securing the IoT (Conti et al., 2018; Ouaddah et al., 2017; Trnka, Cerny, and Stickney, 2018; Alaba et al., 2017; Kouicem, Bouabdallah, and Lakhlef, 2018).

Several design challenges are mentioned in the literature. Some of these challenges include the question to what extent the solution should be centralized (i.e. where the decision for allowing or denying access is made) (Ouaddah et al., 2017), what model or which mechanisms, technologies, and protocols to use. Which architecture or solution is most secure or privacy preserving? How to deal with limited computational resources of IoT devices and find smart trade-offs between privacy, security, and technical feasibility? Existing solutions have their own advantages and disadvantages, but it is unclear which matches the needs of the health care sector the most. Ouaddah et al., 2017 mention the dilemma between adapting existing solutions or creating new ones with IoT specific requirements in mind as one of the main open issues for authentication and authorization within the IoT.

In order to answer these questions, an assessment of current solutions for authentication and authorization, as well as an analysis of security and privacy objectives and more specific requirements is needed. A detailed assessment of solutions for authentication and authorization can be made by comparing these solutions with each other (Bertino, 2016). Not all challenges are purely technical though. Needs related to security and privacy in health care are not specified. These objectives and related requirements need to be identified so current solutions for authentication and authorization can be assessed from a health care perspective, which results in guidance for secure and privacy preserving authentication and authorization in health. The resulting guidelines do not just fill theoretical gaps but may also enable practitioners to find ways to optimize the security or integrity of their data in the IoT. In order to achieve these goals, the following research questions are presented:

Main research question

How can authentication and authorization be managed in order to ensure security and privacy of the IoT in health care? *Explanation:* Within health care, specific requirements are placed on security and privacy of personal information (Ponemon Insfitute, 2016). Traditional access control solutions that are applied in non-IoT environments may not meet these needs sufficiently (AL-mawee, 2015). There are many different architectures and models for authentication and authorization within the IoT, but it is not clear what the relationship is between these architectures and models, and security and privacy in a health environment.

Sub-questions

1. What are the general security and privacy objectives that are applicable to authentication and authorization solutions?

Explanation: the first step in understanding the importance of authentication and authorization and its mechanisms is to understand the general security and privacy objectives that exist for digital environments. A clear view of these different objectives and motivations must exist before current authentication and authorization

systems can be compared. This question is answered based on an analysis of relevant literature in Chapter 2.

2. Which requirements do solutions for authentication and authorization have to fulfill in health care?
Explanation: there are specific requirements for authentication and authorization in health. Because of the nature of the authentication and authorization problem, it is important to not only focus on theoretical requirements, but also expert opinions for practical requirements. This research question is answered in Chapter 3. Available literature is studied and a preliminary answer is given. A number of experts is interviewed in a semi-structured manner in order to define practical requirements.

3. What are the characteristics of currently available IoT authentication and authorization solutions?
Explanation: there are many different solutions for authentication and authorization. This raises several questions about the way these solutions differ. Questions that can now be answered are for example: which architectures are used, which communication protocols are used, how mature are current systems? This question is answered in Chapter 4 based on an analysis of currently available solutions, technologies, and concepts.

4. To what extent do current solutions for authentication and authorization meet the general security and privacy objectives and heath care requirements?
Explanation: now the different security objectives and health requirements of authentication and authorization solutions are clear, it is possible to compare the identified authentication and authorization solutions. These systems can now be evaluated to reveal their strengths and weaknesses. This question is answered in Chapter 5 based on expert interviews and an analysis of available literature.

5. How may current trends influence authentication and authorization for the IoT in health care?
Explanation: this holistic question aims to explore current trends and opportunities for future research. For example, what are challenges, opportunities, risks, threats of authentication and authorization? How will they function in the future, and what will their influences be on security, privacy, and the adoption of the IoT for health care devices? What is the future of promising solutions that are not mature yet? This question will be answered in Chapter 7 based on expert interviews and an analysis of current trends in theory and practice.

1.3 Outline of the research method

There are many security and privacy mechanisms for the IoT. In practice however, security and privacy remain some of the most important challenges in the IoT. The abundance of security compromises and privacy breaches in the IoT in practice proves that security and privacy are still important challenges for the IoT. In this work, semi-structured interviews were chosen because of its flexible, accessible, intelligible nature, and its capability of disclosing information that often remains hidden as part of human or organizational behavior (Qu and Dumay, 2011).

Sub-question 1 and 3 are answered with an analysis of relevant literature. To answer sub-question 2, 4, and 5, semi-structured interviews with IT experts that are active in the field of health care are conducted. The relevant interviewees' professional descriptions include IT auditor, cyber security advisor, data analyst, and health care professional. Interviews are transcribed, and analyzed in qualitative data analysis software NVivo. For a complete overview of the research plan, see Section 3.4.

1.3.1 Approach

Hevner's Design Science method functions as the basis of this research (Hevner, 2007). Within this design paradigm, three inherent cycles contribute to the creation of new artifacts (see Figure 1.2). The rigor cycle connects design science activities with the knowledge base of scientific theories and methods, experience & expertise, and meta-artifacts. The relevance cycle combines design science activities with the application domain, such as people, organizational systems, technical systems, and problems & opportunities. The central design cycle iterates between the two other activities (Hevner, 2007).

To complete the rigor cycle, literature is reviewed to analyze current scientific knowledge concerning authentication and authorization and to analyze the different solutions that have been suggested for authentication and authorization. The aim of this review is to answer sub-question 3. More about this literature review can be found in Chapters 2, 3, and 4.

The application domain is the whole of people, organizational systems, and technical systems interact with each other to achieve a certain goal. To analyze the application domain and in order to complete the relevance cycle, semi-structured interviews with experts are conducted. More about these interviews can be found in Section 3.4. As design science projects aim to identify opportunities and solve problems, the first goal is to define the problem scope of authentication and authorization in these interviews. Besides sketching scope, acceptance criteria for the ultimate evaluation of the research results is also part of this cycle (Hevner, 2007). A secondary goal of these interviews is therefore to define acceptance criteria for authentication and authorization for IoT devices within health care. In practice, this means that requirements of authentication and authorization systems are determined using semi-structured interviews, requirements. This translates to sub-question 2.

Within the design cycle, the requirements are input from the relevance cycle and design and evaluation theories and methods are drawn from the rigor cycle (Hevner, 2007). This cycle consists of building design artifacts & processes by making the artifact. When designing an artifact, such as a set of guidelines, the artifact should later be validated to justify that the artifact would contribute to stakeholder goals. This validation is done before implementation of an artifact within its original problem context (Wieringa, 2014, p. 31). In this case, validation is done using semi-structured expert-interviews.

To complete the relevance cycle, and with that all design science research cycles, the output of the design science research should be returned into the environment for study and evaluation within the application domain (Hevner, 2007). The results of this "field testing" is to evaluate whether additional iterations of the design cycle are necessary. Implementation and evaluation however not in scope of this thesis because of time constraints.

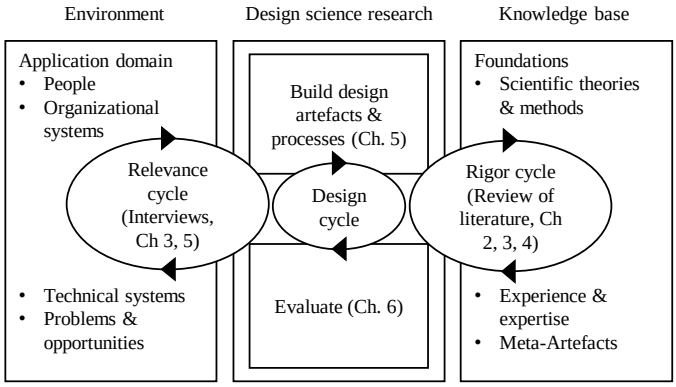

FIGURE 1.2: Design Science Research Cycles by Hevner (2007)

1.4 How this work is organized

The remainder of this work is structured as follows: Chapter 2 presents general security and privacy objectives that are applicable to authentication and authorization solutions, followed by Chapter 3 on the requirements that solutions for authentication and authorization have to fulfill in health care. Chapter 4 presents an overview of current solutions for authentication and authorization. In Chapter 5, these solutions for authentication and authorization are compared to the general security and privacy objectives and health care requirements. Based on this comparison, a set of guidelines for secure and privacy preserving authentication and authorization is created, which is validated using expert interviews in Chapter 6. Chapter 7 finalizes this work with a discussion, expectations, limitations, conclusions, and opportunities for future research.

1.5 Summary of this chapter

There is no generally accepted definition of the IoT but its goal is to connect physical objects using a network infrastructure with inter-operable communication protocols and software. The IoT differs from its predecessors such as traditional internet, mobile internet, sensor networks and M2M networks in a few different ways. Terminology is not standardized and the domain which encompasses the IoT is sometimes vague. Therefore, it is useful to think about how elements in the IoT interact with each other. To do this, generic IoT reference models are introduced. Besides looking at an IoT system from a multi-layered perspective, it can also be viewed from the communication that takes place within or between layers.

The IoT can be used for a wide range of applications in the health care domain and creates new forms of health care services. Health care IoT-devices may collect, transmit and process private or confidential data, which increases need for security and privacy measures. However, despite its importance, security and privacy remain two of the main challenges that the IoT faces. Authentication and authorization are important to ensure security and privacy of the IoT but remain a serious challenge. What is necessary is an assessment of current solutions and a comparison to objectives and requirements. In order to achieve these goals, the following main research question is presented: How can authentication and authorization be managed in order to ensure security and privacy of the IoT in health care? This research question is supported by the following sub-questions: What are the general security and

privacy objectives that are applicable to authentication and authorization solutions? Which requirements do solutions for authentication and authorization have to fulfill in health care? What are the characteristics of currently available IoT authentication and authorization solutions? To what extent do current solutions for authentication and authorization meet the general security and privacy objectives and heath care requirements? How may current trends influence authentication and authorization for the IoT in health care? Because of the nature of the problem, semi-structured interviews with relevant experts are chosen in combination with Hevner's design science method.

Chapter 2

Security and privacy

2.1 About this chapter

Assuming a digital environment, this chapter defines security and privacy and explains their importance. Based on a review of scientific literature, this chapter aims to answer the first sub-question: *"What are the general security and privacy objectives that are applicable to authentication and authorization solutions?"* In order to answer this question, Section 2.2 defines security and privacy. Section 2.3 explains the relationship between security and privacy on the one hand and authentication and authorization on the other. Section 2.4 presents requirements for solutions for authentication and authorization. Section 2.5 concludes the chapter.

2.2 Security & privacy defined

2.2.1 Security

Security is commonly described as the interplay of confidentiality, integrity, and availability, or the CIA triad (AL-mawee, 2015). Confidentiality is defined by ISO as the property that information is not made available or disclosed to unauthorized individuals, entities, or processes. Integrity, or data integrity, is the property that data has not been altered or destroyed in an unauthorized manner. Availability it the property of being accessible and useable upon demand by an authorized entity (ISO, 2000).

The CIA triad is sometimes supplemented with additional concepts to provide further depth, for example *accountability*, or the ability of a system to hold users responsible for their actions, *auditability*, or the ability of a system to conduct persistent monitoring of all actions, *trustworthiness*, or the ability of a system to verify identity and establish trust in a third party, and *non-repudiation*, or the ability of a system to confirm occurrence or non-occurrence of an action (Mosenia and Jha, 2017).

There are different types of security threats that make achieving these security objectives complicated. These threats may include physical attacks, in which an attacker tamperers with physical elements directly, or cyber attacks, which are deployed through malicious software (i.e., malware) or by gaining access to elements of communication networks (Giraldo et al., 2017). The risk that these threats pose may differ for each system, application, or solution. This work mainly focuses on attacks that are focused on gaining unauthorized access.

Security issues and breaches in the digital world are unceasingly common. What started with spam in the 1970s transitioned into viruses and malware, and became increasingly sophisticated and coordinated (Kruse et al., 2017a). Today it is one of the main IT challenges. Organizations may take numerous measures to ensure security of their information. As each system has its own risks, and risk levels may change

over time, it is important that organizations are aware of cyber security trends and threats as they emerge (Kruse et al., 2017a).

2.2.2 Privacy

rmally defined as *"The right of individuals to determine for themselves when, how, and to what extent information about them is collected, processed, and communicated; that includes individuals having the right to determine these aspects within their area of control explicitly; trust that the right above is respected when control is not possible."* (Alpár et al., 2016). Privacy issues may lead to economic, social, and other forms of discriminatory treatment, which may feel as invasive, unexpected, or unwelcome (Wachter, 2018). Ensuring a user's privacy is important for organizations in order to comply with (international) privacy regulation, to ensure a user's trust, and reduce the chance on privacy breaches which may lead to financial loss or reputational damage.

According to (Loukil et al., 2017), complying with eleven privacy safeguarding requirements set by ISO is a good way to respect data privacy laws. These principles are: consent and choice; purpose legitimacy and specification; collection limitation; data minimization; use, retention and disclosure limitation; accuracy and quality; openness, transparency and notice; individual participation and access; accountability; information security and privacy compliance (Loukil et al., 2017). Within the EU's General Data Protection Regulation (GDPR) that came into effect in 2018, the concept of "informed consent" plays an important role. This concept means that the user of a technology fully understands how and why their data will be utilized. In the absence of this consent, collecting, processing, or transmitting of personal data is illegal (Bäumer, Oelffen, and Keil, 2017). As a result, data controllers and processors are expected to be proactive in addressing privacy implications of any system (O'Connor et al., 2017).

Combining and processing seemingly innocent data may introduce privacy problems for users of digital technologies such as the IoT. For example, combining and analyzing data from smart phone sensors can be used to infer a user's mood, stress level, personality type, bipolar disorder, demographics (e.g., gender, marital status, job status, age), smoking habits, overall well-being, sleep pattern, happiness, levels of exercise and types of physical activity or movement (Peppet, 2014). In a PKI-based signature scheme, an entity signs a message using its private key, and can be read using its public key. This public key can be used as unique identifier for an entity and may enable tracing (Peppet, 2014).

Especially in the IoT, privacy problems may arise. This is because of the increasing presence of sensors and devices and the increase of data that these devices collect. As a result, the tendency of collecting more information than necessary for the delivery of certain services is on the rise (Alpár et al., 2016). Collecting seemingly innocent data or combining data can lead to the identification of individuals. For example, Porambage et al. present an example situation in which a user buys an RFID-tagged object. In some cases, the user's personal information could be automatically linked to the object and be known to (cloud) service providers or other entities that process the user's data. Such user information leakage may lead to privacy threats in terms of tracking, localizing, and personalization. If the user possesses a number of objects that can be linked together, malicious entities may be able to estimate the ownership of those objects too, allowing user profiling and tracking.

The larger concept of privacy may have different implications that can be grouped into several privacy preserving "objectives" (Ouaddah et al., 2017), which articulate what should be achieved for someone's privacy to be preserved. These objectives, as

defined by Ouaddah et al. (2017) are: *transparency*, which is a state of affairs in which people understand who knows what about them, how their data will be used, with whom it is shared, and how long it is held; *user-driven*, when users are master of their own data, and have full and granular access control over the data that they share in the network or cloud; *anonymity*, when IoT applications are required to not disclose the identity of their users, for example by allowing anonymous communication, such as hiding location, identity, time, frequency, and volume details, as well as communication context; *pseudonymity*, in which actors are linked with a pseudonym, or random identifier, rather than an identity; *inlinkabily*, when specific actions of a person cannot not be linked together; *unobservability*, when a user may use a resource or service without others being able to observe that the resource or service is being used; *decentralization*, when every node in a network shares its data with other nodes directly, without intervention of a third or trusted party.

Privacy is sometimes seen as completely different from security (e.g., Fabiano, 2017), as two related concepts (e.g., Ouaddah et al., 2017), or as part of security (e.g., Kouicem, Bouabdallah, and Lakhlef, 2018). One could argue that the boundaries between privacy and security are not always clear. Especially the confidentiality aspect of security seems to be partly overlapping with the concept of privacy (Mendez Mena, Papapanagiotou, and Yang, 2018; AL-mawee, 2015). What both privacy and security have in common is that both concepts are concerned with protecting data and restricting access. In this work, privacy and security are seen as two different, but closely related, constructs. This has two main reasons. First, even though not all systems that are perfectly secure (i.e., free from malicious entry) safeguard a user's privacy, nor compliance with privacy legislation, privacy cannot be guaranteed without any form of security in place. Second, both privacy and security are supported or affected by the combination of authentication and authorization.

2.3 Security, privacy, authentication, and authorization

In order to study the relationship between security and privacy mechanisms and the objectives of these mechanisms, Sandhu's OM-AM framework can be used (Sandhu, 2000). This four-layered framework consists of Objective, Model, Architecture, and Mechanism (in that sequence). The Objective and Model (OM) layers express what security objectives, and tradeoffs (priorities, and requirements) are. The Architecture and Mechanism (AM) layers address how these requirements are met. The framework is not aimed at building one abstraction on top of another, but deals with very different kinds of concepts at each layer. Each layer requires its own tools, notations and abstractions. The advantage of this framework is that it clearly establishes the issues at each layer. There are many-to-many relationships between successive layers, meaning that, for example a mechanism may support multiple objectives and objectives may be supported by multiple mechanisms. In this case, we see transparency, user-driven, anonymity, pseudonymity, unlinkability, unobservability, decentralization, confidentiality, integrity, availability, usability, accountability, auditability, trustworthiness, and non-repuditation as objectives, and different aspects of authentication and authorization as model, architecture, or mechanism.

Limiting access to resources to authorized entities is an important aspect of providing security and privacy, for example to ensure confidentiality of information. Because authentication and authorization aim to achieve exactly this, authentication and authorization are essential to provide security and privacy. When designing or using solutions for authentication and authorization however, one should ensure

security and privacy, for example due to misconfiguration or inappropriate mechanisms. The aim of this section is to explain the role of authentication and authorization for security and privacy.

2.3.1 Security

Security problems related to for example incorrect or misconfiguration of solutions for authentication may lead to exposure of personal or device identifiers (e.g., MAC, IP, username) and related tokens (e.g., passwords or cryptographic proofs). Insecure authorization may lead to exposure of confidential information or stolen access tokens.

To prevent such security problems, cryptographic measures can be takenEncryption is important to establish a secure connection between two parties to prevent eavesdropping, and to cryptographically sign messages to prove the authenticity of a message. The mathematical mechanisms behind the cryptography are far beyond scope of this work. The two main forms of cryptography that are discussed here are Public Key Cryptography (PKC) and Symmetric Key Cryptography (SKC). Both forms of cryptography are used in practice to secure authentication and authorization.

PKC, also known as asymmetric cryptography is a form of cryptography which uses two keys, a public key, which may be distributed widely, and a private key, which is not to be shared. PKC can be used for encrypting messages which can only be read by the holder of the private key. Another use is for digital signatures. A sender of a message can use its private key to generate digital signatures. The receiver can use the sender's public key to verify the signature. A downside of PKC is that it is computationally relatively resource intensive. Therefore, lightweight public key encryption algorithms have been proposed, such as Ecliptic Curve Cryptography (ECC). ECC uses smaller key sizes for achieving the same level of security (Ye et al., 2014), resulting in lower computing and memory requirements (Hernández-Ramos et al., 2016).

In SKC, the same key is used for encrypting and decrypting information. It is computationally less resource intensive than PKC but the problem is that the key must be exchanged with the other party over a potentially insecure channel that might be eavesdropped. A symmetric key can be established over an insecure channel using a so-called key exchange or handshake, such as Diffie-Hellman key exchange.

2.3.2 Privacy

Even though authentication and authorization are essential to achieve confidentiality and therefore privacy, they may also harm a user's privacy. For example, when communicating over the internet, it is almost always necessary to disclose information that can be used to identify devices, users, or natural persons, such as IP-addresses of communicating parties. Messages sent or actions taken by entities may be linked to the same entity. That may make users vulnerable to undesired or malicious tracking (Krasnova, 2017). But even when the communicating parties are successful in concealing these properties for third parties, privacy issues may occur as a result of authentication. This is because authentication generally consists of an identity (such as a user name), and a token (such as a password or a cryptographic proof). Therefore, there must always be a party that controls a database with identities. The entity that is responsible for maintaining this list of identities is therefore (at

least technically) able to track users in a system, creating potential privacy problems (Krasnova, 2017).

To preserve the privacy of users that authenticate and protect personal data in the IoT, several privacy preserving approaches can be used, such as data minimization, anonymization, pseudonymization, group signatures, and attribute-based authentication. What these approaches have in common is that they do not rely on an identity that can easily be linked to a natural person. In this section, a brief outline of these approaches is given. A more detailed overview of how these approaches are used in practice for authentication and authorization can be found in Chapter 4.

One approach to preserve the privacy of users could be to minimize the amount of data gathered or processed by IoT devices, known as data minimization. Data minimization is an important area of concern in governmental policies (Wachter, 2018), but hard to accomplish in IoT scenarios as devices may even gather information without the user being aware of this (Conti et al., 2018). Data minimization generally refers to a minimization of the amount of data that devices gather (for example a home surveillance camera that only transmits signals when the home owner is not at home) but can also be used to minimize the data that is necessary for authentication.

Zhou et al. (2017) discusses how authentication can play a role in preserving user's privacy. When identifying themselves, users may for example adopt a pseudonym rather than an identity, known as pseudonymization. A pseudonym may protect a user's privacy because the pseudonym cannot be linked directly to the user. However, it may be necessary to update pseudonyms periodically, which can be hard (Zhou et al., 2017). Besides that, each pseudonym is unique which potentially makes entities identifiable if additional information is available.

Another approach is by using group signatures, which are a form of digital signatures in which any member of a group can sign on behalf of the others. The identity of a member is therefore not necessarily disclosed (Khader, 2007). However, because this signature is the same for an entire group, this may not allow for fine-grained access control which is often necessary in IoT scenarios.

It is also possible to authenticate users via attributes, which are properties of users such as environmental conditions like time and locations (Alpár et al., 2016). Some argue that using attributes instead of identities for access control improves user's privacy, because decisions for granting or denying access can be made solely on those qualities that are seen as essential (Alpár et al., 2016). Attribute-based authentication allows users to control their personal information, by data minimization, and limitation of goal (Krasnova, 2017).

2.4 Requirements for solutions for authentication and authorization

The European Commission, United Nations' authorities, and other worldwide law enforcement organizations are trying to find a common ground for addressing IoT privacy issues while empowering the existing legal framework. According to Porambage et al. (2016) such an IoT privacy framework should include authentication (in the form of identity privacy) and access control (fine tune the granularity). Privacy frameworks for IoT health care applications must be open and transparent to patients, specify the reasons for collecting necessary health information, maintain accurate and real-time information, and ensure the protection of patient records (Porambage et al., 2016)

This work looks at the relationship between different aspects of privacy and security as objectives compared to authentication and authorization models, architectures, and mechanisms. The sub-concepts of privacy and security are presented as objectives. To study the relationship between these objectives and solutions for authentication and authorization, a few high-level (theoretical) requirements that are based on these objectives are presented in Table 2.1. This list of objectives and their requirements is based on the literature that is presented in this chapter. The list is validated in Chapter 3 to examine correctness, completeness, and relevance for authentication and authorization for the IoT in health care. Additionally, potential areas to consider are identified. In Chapter 5, guidelines for achieving the objectives are presented.

TABLE 2.1: An overview of the objectives and requirements that were identified in this chapter

Objective	**Requirement**
Transparency	The system must allow people to understand who knows what about them, how their data will be used, with whom it is shared and how long it is held.
User-driven	The system must allow users to have full and granular access control over the data they share in the network or in the cloud.
Anonymity	The system must not disclose the identity of their users.
Pseudonymity	The system must link actions of a person with a pseudonym rather than an identity; trades off anonymity with accountability.
Unlinkability	The system must not link specific actions of the same person should together unless necessary.
Unobservability	The system must not allow users and / or subjects to determine whether an operation is being performed by another user.
Decentralization	Each node in the network shares its data with others nodes directly, without intervention of any third or trusted entity.
Integrity	The system must prevent unauthorized modications of resources.
Availability	The system must support a high readiness for usage. Offline mode and short- and long-time availability.
Confidentiality	The system must prevent unauthorized disclosure of resources through granular (fine-grained), revocable, delegatable access control.
Usability	The system must allow access control to be easily managed, expressed and modied.
Accountability	The system must be able to hold users responsible for their actions.
Auditability	The system must be able to conduct persistent monitoring of all actions.
Trustworthiness	The system must be able to verify identity and establish trust in a third party.
Non-repudiation	The system must be able to conrm occurrence or non-occurrence of an action.

2.5 Conclusion

This chapter starts with the question: *"What are the general security and privacy objectives that are applicable to authentication and authorization solutions?"*. Security is commonly described as the interplay of confidentiality, integrity, and availability, potentially extended with additional constructs such as accountability, auditability, trustworthiness, or non-repudiation. Privacy issues may lead to different kinds of negative effects for individuals. Also from a legal point of view it is important to ensure privacy of personal data. Combining and processing seemingly innocent data may introduce privacy problems for IoT-users. Especially in the IoT, privacy problems may arise. Privacy can be grouped into several "objectives". In this work, privacy and security are seen as two different, but closely related, constructs. The OM-AM model can be used to describe the relationship from objective to mechanism. A list of objectives and requirements for privacy preserving authentication and authorization is presented. This list includes transparency, user-driven, anonymity, pseudonymity, unlinkability, unobservability, decentralization, confidentiality, integrity, availability, usability, accountability, auditability, trustworthiness, and non-repuditation.

Chapter 3

The IoT in health care

3.1 About this chapter

This chapter explains the importance of security and privacy of health-related data, introduces the IoT in health care, and the role of security and privacy in IoT. The chapter aims to give an answer to sub-question 2: *"Which requirements do solutions for authentication and authorization have to fulfill in health care?"* Section 3.2 explains the importance of security and privacy for heath care data, followed by Section 3.3 which explains the most important threats for the IoT in health care. Section 3.4 presents the most important challenges and objectives of security and privacy for the IoT in health care based on expert interviews. Section 3.5 presents requirements and areas to consider for achieving secure and privacy preserving authentication and authorization for the IoT in health care. Section 3.6 concludes this chapter.

3.2 The importance of security and privacy for health care data

The previous chapter explained the importance of security and privacy of digital information. This threat is even larger for health care organizations, as they are more vulnerable to modern trends and threats to security and privacy. This is because they experience difficulties adopting cyber security measures and lag behind in security (Kruse et al., 2017b). According to Kruse et al. (2017b), there are two primary drivers that result in an increase of cyber threats for health care organizations. Their US-oriented study found that one of them is a result of new US federal policy initiatives, which takes a lot of health care organization's IT budget, allowing them to spend less than 5% of their IT budget on security. The other one is the quickly changing technological landscape. New technologies are implemented faster than the security systems can be created or updated to protect these devices. Previously stand-alone systems are becoming integrated within IT systems and are no longer immune to traditional cyber attacks (Kruse et al., 2017a).

Health care as a sector is a prime target for medical information theft. This information may contain sensitive personal and financial information. (Kruse et al., 2017b). One US study found that that 90 percent of health care organizations had a data breach in the past two years, and almost half of them had more than five data breaches in the same period. Criminal attacks are the main cause of health data security breaches in health care (Ponemon Insfitute, 2016). Personal health information are more valuable than most other types of information (Ponemon Insfitute, 2016). Stolen patient health records can be sold for up to 60 USD per record, which is about 10 to 20 times more than credit card information (Freeze, 2019). Medical information

allows criminals to commit identity theft, medical fraud, extortion, and the ability to illegally obtain controlled substances (Kruse et al., 2017a).

Besides the cyber risks, the increasing accessibility and availability of personal health records accessed via the internet can also lead to significant privacy issues (Porambage et al., 2016). As personal information is gathered, transmitted, and processed, these personal data are at risk by definition. If the privacy of patients is not preserved, these data may eventually be visible for health insurers, governments, unauthorized medical personal, or the general public.

3.3 Threats for the IoT in health care

IoT devices in health care are not free from privacy and security risks. These devices have the capability to collect, process, and transmit large amounts of personal data, potentially without the user knowing or directly able to change this. This causes the problem that personal data is flowing out of sight or control of users. This has two main downsides. First, data that often can be linked to individuals may be stored or processed by potentially insecure devices which may by subject to malicious attacks, causing the risks that these data are leaked. Second, personal data may be collected, transported, or processed by legitimate entities, but for reasons that go beyond the primary reason why these data are gathered, which causes privacy issues.

User's privacy may be at risk as a result of IoT technologies. For example, in the US, insurers have started giving customers discounts on their insurance if they are willing to share data from wearables such as Fitbits or Apple Watches (Ingraham, 2018). Apple Watches already have the capability to accurately measure heart rate (Abt, Bray, and Benson, 2018) and other data such as location. Data that the insurer already has (e.g., zip code, gender, historical declarations), may provide an accurate image about someone's historical and present health state. Combining these date data with those gathered by wearables may allow the insurer to make future predictions. For the user, this may lead to unwanted profiling, possibly resulting in economic, social, or other forms of discriminatory treatment (Wachter, 2018). These negative effects could be avoided by (international) privacy legislation or informed consent. However, legislation may only be effective combined with adequate law enforcement, and consumers may not understand what they agree with. Because of these and other reasons, the concept of Privacy by Design (PbD) (or Privacy by Default) may offer a solution. PbD is defined as a philosophy and methodology of *embedding privacy into the design specifications of information technologies, business practices, and networked infrastructures as a core functionality* (Cavoukian, 2011). PbD is an important aspect of the new European General Data and Privacy Regulation (GDPR), which came into effect in May 2018, and shifts organization's the attention from traditional approaches to protecting privacy reactively, to a more proactive approach (Cavoukian, 2011). Because of the increased focus on a system's design, clarity about what secure and privacy preserving authentication and authorization for the IoT entails is necessary. The remainder of this chapter aims to do that.

Especially for health care, the level of adoption of IoT devices is expected to depend on the level of the user's perceived confidentiality and integrity (Ouaddah et al., 2017). It is therefore important that these data are protected and only available by those who have legitimate reasons to access these data, such as the patient and relevant medical personnel.Adoption of the IoT in health care is not going as fast as expected. Koop et al. (2008) argued in 2008 that by 2018, health care capabilities would be more or less evenly distributed in hospital, community and home, and

by 2028, the majority of health care should be moved out of the hospital and into the community and home. To reap the benefits of IoT in health care, it is therefore important that this environment is made secure. There are some challenges however that have to be overcome specifically for the IoT.

Kouicem, Bouabdallah, and Lakhlef (2018) mentions a few challenges for the IoT that potentially compromise privacy and security of the IoT. These challenges are the amount of (cyber) attacks, data management, quality-of-service (QoS) constraints, scalability, heterogeneity, lack of standardization, mobility, and resource constraints. Based on these challenges, requirements for solutions for authentication and authorization for the IoT in health care can be identified. In order to do so, several experts are interviewed.

3.4 Responses from interviews

Based on the available literature, several general objectives of security solutions have been suggested, which are introduced in Chapter 2, such as scalability, transparency, context-awareness etc.. Users of IoT-technologies may however have different levels of privacy awareness and concerns and therefore have different needs. In order to further define health care-specific areas to consider for authentication and authorization solutions for the IoT and to qualitatively evaluate current solutions (see Section 3.4.1), several interviews are carried out. During the interviews, two topics have been mainly relevant. First, the biggest *challenges* for security and privacy for the IoT in health care (see Section 3.4.1) and second, the most relevant *objectives* for security and privacy solutions in health care (see Section 3.4.2). For both topics, a pre-defined list of challenges and objectives is available that contains the most important challenges and objectives that are identified in Chapter 2. Interviewees are also asked to mention additional challenges or objectives if the lists seems incomplete to them and to identify specific *areas to consider*, which are the most important areas that one should consider to achieve an objective.

The number of interviewees needed for this study is determined based on saturation sampling (Wohlin, Höst, and Henningsson, 2003), meaning that interviews were conducted until a consensus had been reached, indicating that multiple participants have given similar answers and new insights from additional interviews are not expected anymore. In order to find out whether the level of saturation is reached, all interviews were coded using qualitative data analysis computer software package NVivo. The aim of the coding is to produce a systematic recording of the themes and issues addressed in the interviews by tagging textual segments for later comparison. The main categories of interest were identified first. Table 3.1 shows that of these categories of interest were discussed in the interviews at least once, indicating that saturation has been achieved and additional interviews would not lead to new topics.

All interviewees are active in the field of health care, and are employed by health organizations directly (e.g., a hospital), or by an organization that delivers products or services to health organizations, such as manufacturers of IoT solution, or consultants. The interviewees' job description include: cyber security expert (interview no. 1), IT auditor (interview no. 2), manager IoT manufacturer (interview no. 3), cyber security expert (interview no. 4), IT infrastructure manager at a hospital (interview no. 5), IT auditor (interview no. 6), IT auditor (interview no. 7), cyber security expert (interview no. 8), IT auditor (interview no. 9), IT auditor (interview no. 10), and advisor E-health (interview no. 11).

TABLE 3.1: New concepts in interviews; indicating that no new concepts within the pre-defined topics were introduced by the interviewees after the 9th interview

	Interview no.										
Category of interest	1	2	3	4	5	6	7	8	9	10	11
Challenges for achieving privacy and security in the IoT	■	■	■	■		■					
The IoT in health	■	■									
Environmental conditions		■	■								
Privacy and perception of privacy	■	■	■	■					■		
Security and security measures	■	■	■	■							
Data minimization	■		■								
The IoT in the future / current trends	■		■								
Identities and their roll	■	■									
Regulation			■								

■ = A new concept within the topic was introduced by the interviewee.
□ = No new concept within the topic was introduced by the interviewee.

The data collection technique is of a first-degree level, which means direct interaction with the subjects and recording of the data in real time (Wohlin, Höst, and Henningsson, 2003). When using semi-structured interviews, themes are planned ahead, but additional questions can be asked if deemed necessary. A simple slide deck was used for visual assistance. Before the interview, the interviewee was told about the confidentiality of the interview and their rights related to stopping the interview at any time for any or no reason. A brief introduction of the IoT in health care is given. Because of the level of experience of the interviewees, it is expected that they have an understanding of IoT in health care. A brief introduction has been given anyway to ensure that the setting and scope of the problems are clear. The first set of interviews contribute to sub-question 2 *(Which requirements do solutions for authentication and authorization have to fulll in health care?)*.

3.4.1 Topic 1: main security challenges for the IoT

In the first part of the interview, a few typical IoT health care scenarios were given. The question *"What are the main challenges for security and privacy for the IoT in health?"* was asked, with which the interviewee was asked to mention the most significant privacy and security risks. Interviewees could choose from a pre-defined list of challenges that were identified in Chapter 2 that included *safety*, *amount of attacks*, *data management*, *QoS constraints*, *scalability*, *heterogeneity / lack of standardization*, *mobility*, and *resource constraints*. The interviewee is asked to specify what they believe are the main challenge for security and privacy for the IoT in health care, and which are less relevant. The interviewee is asked to explain their reasoning and to provide other significant challenges that they feel are missing on the list.

The reason this question was added to the interview was to specify the main challenges for the IoT. Kouicem, Bouabdallah, and Lakhlef (2018), mention these challenges as the main challenges for the IoT in general. The authors also recognize however that that the significance of these challenges is different per application

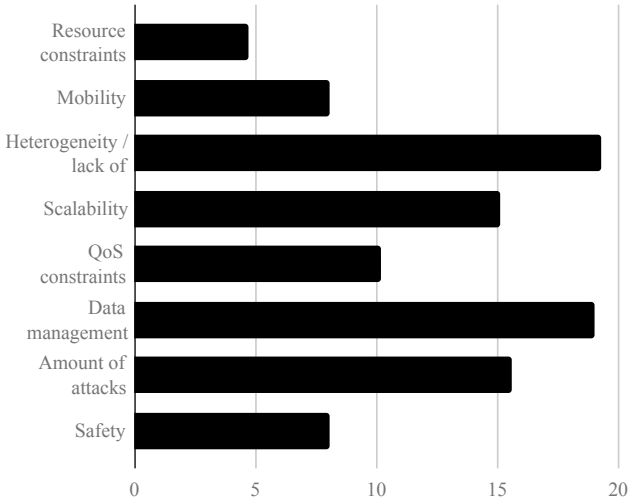

FIGURE 3.1: Main security challenges according to the interviewees

domain. For health care, the authors specify resource constraints as the main challenge for achieving security and privacy for the IoT. Mobility, heterogeneity / lack of standardization, scalabilty, QoS constraints, and safety are mentioned as moderately large challenges. Data management and amount of attacks as mentioned as smaller challenges. The authors do not specify the reasoning for the ranking of these challenges. This question therefore aims to find out what really are the main security challenges for IoT in health care and why these are the main challenges. Figure 3.1 shows the level of importance that the interviewees consider these constructs to be. To prioritize each of the constructs, each interviewee was given 100 prioritization points, which he of she was allowed to distribute across all items on the vertical axis of the figure. The length of the horizontal bar represents the average number of points that the interviewees awarded to the item across all interviews. Interestingly, these outcomes differ considerably from those of Kouicem, Bouabdallah, and Lakhlef (2018). The following sub-sections summarize the interviewee's reasoning for mentioning a challenge as relevant or not relevant.

Safety

Just like the concept of cyber security, safety deals with the concept of risk. The difference however is that safety is related to accidental risk, such natural disasters, damages, and human errors (Kouicem, Bouabdallah, and Lakhlef, 2018). Put differently, safety relates to whether a system may harm its environment, whereas security consist of protecting the system from attacks that come from the environment. As Figure 3.1 shows, safety is seen as a moderately important security challenge. An interviewee noted:

> *Safety relates to whether a lot is already optimized using IoT. You can say that they cannot operate without it anymore. I have seen that about 20 beds can be monitored by one person. An alarm goes off when someone gets out of bed, if that (monitoring) goes wrong, it can have life-changing consequences.*

Amount of attacks

In the interviews, the amount cyber attacks were seen as a real risk for the condentiality of user's health care data, which is in line with recent ndings, that health care organizations consider cyber attackers and malicious insiders an important security threat (Ponemon Insfitute, 2016). One interviewee noted when asked what is a big security problem:

> *Hackers, people who are consciously looking for data (...) previously, this was about data which you could use to make money with, like credit card information and emails. Now, people are looking for context data to put people under pressure. This is much more sensitive than for example credit card information, which you can block.*

Data management

Data management challenges are related to the management of the huge amount of data generated by smart devices. Questions related to locating data, controlling access to the data and preserving integrity and privacy are related to data management (Kouicem, Bouabdallah, and Lakhlef, 2018). Kouicem, Bouabdallah, and Lakhlef do not consider data management an important security challenge for the IoT in health care. Interestingly, the experts that are interviewed for this study regarded data management to be an important challenge. Patient information is often stored on-premise and secured within the organization's own environment. Information comes from different systems and has to be integrated somehow. How to do this securely remains a challenge for health care organizations. The interviewee's answers mainly focused on ensuring confidentiality of information and preventing data leaks, which is difficult due to the amount and nature of data gathered. For example one interviewee noted:

> *Data management is important because the gathered data is very personal. If you can combine data, you can make a large number of predictive models. It is therefore important that the data stays where it should be, to prevent data leaks.*

QoS constraints

QoS (Quality-of-Service) is seen as a moderately important problem. A 100% available connection cannot be achieved in practice, especially outside of hospital environments. It is therefore only possible to monitor non-critical patients outside of hospital environments and IoT solutions should always take into account that connections may be lost and that a decent level of robustness is required. One interviewee said:

> *This will always remain an issue, which is hard to overcome.*

Scalability

Scalability is seen as a relatively large problem for privacy and security of information in the case of the IoT in health care. Especially for ensuring robustness and availability of information. One interviewee noted:

> *Scalability is mainly an issue for availability of information. In terms of privacy and security, it is mainly a functional problem.*

Heterogeneity / lack of standardization

Heterogeneity and a lack of standardization are seen as a significant problem for security and privacy for the IoT in health care. There are many different devices, applications, and tools are not able to transmit or process data in a standardized way. There are standards for data protection, and also standards for storing and sending electronic health records and standards for medical devices. There are however, no standards for security and privacy of medical IoT devices. As a result, privacy impact analyses are performed regularly on an ad hoc basis. For security, there is a reliance on other standards for information security such as ISO27001.

> *I think the ability to link data to other systems is one of the main issues. (...) How to handle that decently and securely, that is hard.*

Mobility

Mobility is a fairly large problem, but at the same time, mobility is one of the advantages of IoT technologies and one of the main opportunities that IoT technologies offer relative to other types of digital technologies.

> *The fact that you do not know where your devices are or where they can go makes ensuring privacy and security harder, especially security. On the other hand however, the fact that devices can be taken anywhere is one of the main advantages of the IoT, because you can be monitored all day without being in a hospital all day for example.*

Resource constraints

Digital devices may have vulnerabilities. Within the IoT, smaller devices may have problems securing tokens and credentials. IoT device components are often low-cost and do not have sophisticated means to protect code, data, and tokens. Due to power and hardware constraints they may only be able to process tokens with low complexity, which makes it easier to compromise the token. Over the years, computing power has increased, even for smaller devices, but, due to heterogeneity of devices and the number of devices, it cannot always be expected that non-light-weight security solutions will scale well in the IoT. The trustworthiness of devices may degrade over time due to changes. As a result, it must be possible to dynamically withdraw certain authorizations from devices or users (Hu, 2016, p. 240).

Resource constraints related challenges are mentioned as one of the main challenges for authentication and authorization in the IoT in health care in literature. The practitioners interviewed however, did not entirely agree with this. Devices are expected to at least be able to support encryption, and have hardware that allows protection of cryptographic keys and tokens, but if this is achieved, a device's resource constraints are no longer seen as problematic for a device's security or privacy. When asked if smaller devices have enough capability for data protection mechanisms such as encryption one interviewee agreed and said:

> *Resource constraints is becoming less of a problem. For example, the capacity of batteries and of power of microprocessors is increasing. This, in the end, allows for more opportunities for securing the device.*

Please note that *resource constraints* is seen as a security challenge by Kouicem, Bouabdallah, and Lakhlef, similarly to the concept of *technology constraints* as defined

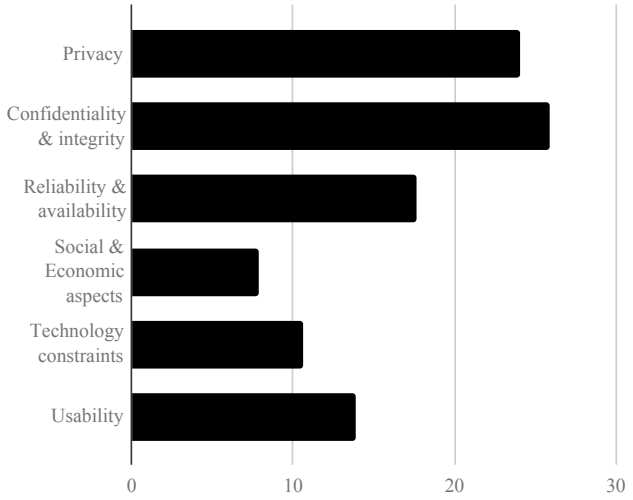

FIGURE 3.2: Main security and privacy preserving objectives

as an objective by Ouaddah et al. In this work, the concept of *resource constraints* is seen as a challenge, so as an obstruction to achieve an objective, rather than an objective itself.

3.4.2 Topic 2: main privacy and security objectives for the IoT in health care

With the question: *"Which objectives of privacy and security are most relevant for IoT solutions in health?"*, the interviewee was asked to mention the most important objectives for achieving security and privacy as well as what they believe are less important. Just like for the first question, the interviewee was be asked to explain its reasoning and to provide other significant objectives that they feel are missing on the list.

Just like for the first question, the significance of these objectives is said to differ per application domain (Ouaddah et al., 2017). However, the authors do not specify the relative importance of these objectives for health care. This question therefore aims to find out what the main security and privacy goals for IoT in health care are. Figure 3.2 shows which security and privacy preserving objectives are considered most important by the interviewees. The prioritization in this figure is done similarly to the one in Figure 3.1 in which the length of the horizontal bar represents the average number of prioritization points that the interviewees awarded to the item across all interviews. To explain the interviewee's reasoning, a summary of given answers including fragments of interview answers is given here for every challenge.

Privacy

Privacy as a general umbrella term for different sub-objectives of privacy is seen as one of the most important objectives. To make the ambiguous construct of privacy more specific, it is divided among sub-objectives (transparency, user-driven, anonymity, pseudonymity, unlinkability, unobservability, decentralization). Sub-objectives such as transparency or unlinkability are seen as essential. In the opinion of the interviewees, individuals should be able to access and understand who may

access their personal data to ensure that they trust the IoT technology. Other objectives that are seen as important are user-drivenness, as users may not agree with default authorization policies, anonymity or pseudonymity as identity attributes may be exposed to third parties when interacting with the IoT, making users directly or indirectly identifiable, and unobservability, as devices may collect, process, or store more identifying attributes than the minimum amount necessary for the correct working of the system. In one case mentioned as example in one of the interviews, collected data comes in a secured in-house portal (on-premise) at the health care provider. The interviewee was asked if suppliers of IoT devices are able to store data themselves, and what the consequences may be for privacy for the patient. Its response was:

> *Interviewee 1: That does not happen (yet). Interviewee 2: if we would store data decentrally, we carry out a Privacy Impact Assessment (PIA) first. Then we will evaluate the product. If the product does not adhere to our standards, it will not be used.*

Decentralization, and delegation are not found to be an important sub-objective, and are therefore omitted. This is probably because currently, most IoT systems in health care are based around a centralized architecture.

Confidentiality and integrity

Confidentiality and integrity are seen as one the most important objectives, as there is a chance that unauthorized entities may access or modify information. One interviewee mentions:

> *The most important is confidentiality and integrity (...) especially with IoT in combination with health. If confidentiality is ensured, privacy is protected partly as well. (...) If you have not taken care of this, you will never know for sure if the data that you have are integer. I think this is the most important.*

Reliability and availability

In summary, reliability and availability are seen as important. The extent to which reliability and availability are hard to achieve differs per scenario. Interviewees mentioned that this is easier in hospitals or formal health care institutions but harder in to achieve a home environment or outdoor. This is mainly due to environmental conditions that may cause imperfections in connection. One of the interviewees notes:

> *There is a need not just for privacy and security, but definitely also reliability. From what I know, especially in care, they are trying to optimize processes using IoT appliances.*

Social & economic aspects

Social and economic aspects include interoperability, cooperativity and collaboration, and context awareness. These concepts may allow sharing of IoT authorization policies between hospitals for example. These objectives are seen as useful, but not as essential for access control, mostly because in health care, data that is gathered by an entity is usually the primary user of those data. Only after processing and analyzing these data, for example into a structured format or patient record, it becomes

important to share these data between entities. As sharing of patient records is a topic on its own, it is not included in this work.

Usability

Access control should be easily managed, expressed and modied. Even though it is essential that this can be done, it does not have to be the patient who has control over authorization decisions. In case usability cannot be guaranteed, users of the IoT environment may not be the appropriate, because of a failure to manage access policies well.

Accountability, auditability, trustworthiness, and non-repudiation

The concepts accountability, auditability, trustworthiness, and non-repudiation are not in the list of objectives as defined by Ouaddah et al. (2017). Based on the comments of several interviewees however, these are added as objectives. Interviewees mentioned the problem that organizations may face a lack of visibility of unauthorized actions. There is the possibility that unauthorized actions, entities, attackers, or malicious nodes cannot be identified. This may result in the inability to take actions against these actions. As changes in device's or third-party service's perceived level of security may occur, inadequate estimates of trust or reputation of devices may exist.

It is seen as hard to monitor IoT devices, as these devices often do not save log files themselves. As a solution, monitoring devices on a network level is seen as an effective and mature solution to provide security in a network. This requires however that devices within this network can be monitored, allowing establishing trustworthiness of a device in a network. Virtualization technologies are seen as useful to do this.

3.5 Requirements for authentication and authorization for the IoT in health care

Based on the discussion above, several requirements for secure and privacy-preserving authentication and authorization for the IoT are presented in this chapter. A summary of the objectives, requirements, and areas to consider is found in Table 3.2. The *objectives* and *requirements* in the first and second column are introduced in Chapter 2. The *area to consider* column is based on the answers given by the interviewees that are identified in this chapter. Based on this list, mitigation strategies are presented in the form of "guidelines" in Chapter 5.

TABLE 3.2: An overview of identified objectives, requirements, and areas to consider identified in this and previous chapter based on literature and interviews

Objective	Requirement	Area to consider
Transparency	The system must allow people to understand who knows what about them, how their data will be used, with whom it is shared and how long it is held	Individuals may not be able to access or understand access policies, and may, as a result not trust the IoT technology.
User-driven	The system must allow users to have full and granular access control over the data they share in the network or in the cloud.	The user may not agree with default authorization policies.
Anonymity	The system must not disclose the identity of their users.	Identity attributes may be visible to third parties when interacting with the IoT.
Pseudonymity	The system must link actions of a person with a pseudonym rather than an identity; trades off anonymity with accountability.	Individuals may be directly identifiable based on the identifiers used for authentication.
Unlinkability	The system must not link specific actions of the same person should together unless necessary	Users may be subject to undesired linking of data by service providers or other third parties.
Unobservability	The system must not allow users and / or subjects to determine whether an operation is being performed by another user.	IoT devices may collect, process, or store more identifying attributes then necessary for the correct working of the system.
Integrity	The system must prevent unauthorized modications of resources.	Unauthorized entities may modify data.
Availability	the system must support a high readiness for usage.	A lack of availability may lead to connectivity or reliability issues.
Confidentiality	The system must prevent unauthorized disclosure of resources through granular (fine-grained), revocable, delegatable access control.	Unauthorized entities may access personal data.
Usability	The system must allow access control to be easily managed, expressed and modied	Users of the IoT environment may not be the intended users because of a failure to manage access policies well.
Accountability	the system must be able to hold users responsible for their actions (e.g. misuse of information.	Unauthorized entities, attackers, or malicious nodes may not be identified.

Objective	Requirement	Area to consider
Auditability	The system must be able to conduct persistent monitoring of all actions	The organization may face a lack of visibility of unauthorized actions.
Trustworthiness	The system must be able to verify identity and establish trust in a third party	Inadequate estimates of trust or reputation of devices may occur due to changes in device's or service's perceived level of security
Non-repuditation	The system must be able to conrm occurrence or non-occurrence of an action	Actions performed by malicious entities cannot be made undone.

3.6 Conclusion

"This chapter started with the question: *"Which requirements do solutions for authentication and authorization have to fulll in health care?"*. In health care, organizations are vulnerable to cyber threats and lag behind in security. Health care data is an important target for cyber attacks. Increasing digitization of health records may also lead to privacy issues. The IoT in health care is not free from privacy and security risks and its users may be therefore be at risk. The adoption rate of IoT technologies in health care however, depends on the perceived security and privacy of these technologies. Interviews are carried out to define health care-specific requirements for security and privacy in this domain. The main challenges for achieving privacy and security are related to data management, heterogeneity, and a lack of standardization. Confidentiality, integrity, transparency, and unlinkability are the most important security and privacy related objectives for the IoT in health care. Based on literature, requirements are defined for each of the objectives. For each objective, important areas to consider are defined based on interviews.

Chapter 4

Authentication & authorization for the IoT

4.1 About this chapter

In order to allow or restrict access to a resource, a system has to identify (a) who or what the requesting entity is (authentication), and (b) evaluate if the entity has rights to access the resource (authorization). There are different approaches how this can be done in the IoT. The aim of this chapter is to give an answer to the sub-question: *"What are the characteristics of currently available IoT authentication and authorization solutions?"*. This chapter presents an overview of the current literature on authentication and authorization for the IoT. Section 4.2 presents the approach taken to find relevant papers/ Section 4.3 explains a generic process for authentication and authorization. Different implementations of solutions for authentication and authorization are presented in section 4.4 and compared in section 4.5. Section 4.6 concludes this chapter.

4.2 Approach

This chapter presents an overview of the current literature on authentication and authorization for IoT devices. The focus of this review is on the "higher levels" of authentication and authorization systems such as security objectives, frameworks, projects, schemes, and architectures, and not on lower levels of the network stack such as hardware tools or encryption algorithms.

In this chapter, the focus is on authentication and authorization specifically for the IoT, as the IoT differs from traditional IT systems in a few ways. First, due to the size of some smaller IoT devices, some of these devices may suffer from resource constraints, such as a limited battery life, limited memory, or limited computational power. Second, the IoT consists of a large number of heterogeneous devices, that may use different protocols or standards. Third, the IoT generally consists of a larger number of devices. Fourth, IoT devices may adopt different roles. Fifth, IoT devices may not be able to rely on traditional user name-password-style authentication for different reasons, for example as some devices may not have keyboards or screens to enter credentials, or because there is no single user that "owns" the device, for example in case of security cameras.

The following approach is be taken to find and analyze the current literature: The following indexing sites have been used: IEEE Xplore, ACM Digital Library, Web of Science, SpringerLink and ScienceDirect using the query *("Internet of Things" OR IoT) AND (authorization OR authorisation OR "access control")*. Please note that the exact query syntax differed slightly per indexing site. Therefore, the query is

adapted slightly in some cases. Backward snowballing, or using the reference list to identify new papers to include (Wohlin, 2014), is also used to identify papers that are not found in the initial search.

Papers from 2012 and older, papers that have four or less pages (excluding references), papers that are not (a) articles, (b) conference papers or (c) book chapters, and duplicates of papers have been deleted. As a final filter, papers that still seem out of scope based on their abstract have been deleted.

4.3 Restricting or allowing access

There is a complex interplay of many different identity components in IoT platforms, devices at the edge of the platform communicate to gateways, which requires trust between edge device and gateway, attributes of users and devices have to be stored in a registry, and applications for data analysis need to communicate to databases through APIs. This creates a complicated interplay of identities, trust, and authorizations. This section aims to explain the process for allowing access to a resource to a device that an entity previously did not have a connection to. In practice, the entire process differs per authorization solution. An overview of the general (and traditional) process is explained here.

In general, the first step is securing the connection between the resource owner and future resource requester to prevent eavesdropping. Second, the devices must establish an identity, so that (third) access permission can be granted to these identities. Fourth, when requesting access to a resource, a device must prove its identity (authentication), after which the resource owner can (fifth) grant access (authorization). In the rest of this chapter, this process is explained in more detail.

Establishing a secure connection in the IoT

In order to prevent eavesdropping, the transport layer between two devices must be secured. There are different mechanisms to do this. In case the RESTful architectural style is applied, HTTP or CoAP are common protocols which can be secured using Transport Layer Security (TLS) (its predecessor Secure Sockets Layer (SSL) is no longer considered secure). Datagram Transport Layer Security (DTLS) is a common implementation of TLS.

Identities in the IoT

An identity is a set of properties that makes an entity (e.g., a user or device) unique. This identity can be established in multiple ways. For human users, this can be done using for example a user name. There are several attributes that can be used for identification of devices, for example a laptop has its manufacturer model number, product key of operating system, and an IP or Media Access Control (MAC) address.

Granting of access in the IoT

Resource holders can grant access to resources by updating authorization policies. How this works in practice differs per solution for authentication and authorization.

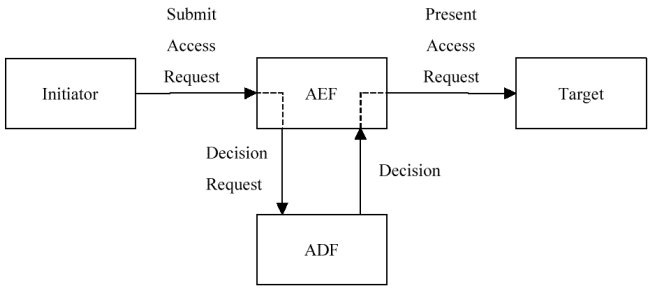

FIGURE 4.1: ISO 10181-3 Access Control Framework by The Open Group (2000)

Authentication in the IoT

Controlling which objects can access a resource is essential to ensure privacy and security in the IoT. Confirming an object's identity in the IoT is a critical prerequisite for access control, because a decision whether to grant access to a user or non-manipulated device can only be made if the identity of the object can be confirmed. Authentication is the process of identifying users and objects in networks and is therefore an important step in ensuring security and privacy. When identifying users, authentication can consist of multiple "factors". It can be based on *something the user knows*, usually a user name and password, *something the user has*, such as a physical device like a hardware token, or *something the user is*, for example on biometric information, such as a finger print (Trnka, Cerny, and Stickney, 2018). Traditionally, authentication relied on user names and passwords, but this requires user interaction, which limits its use in the IoT (Atwady and Hammoudeh, 2017).

Authorization in the IoT

Many architectures for traditional authorization systems are based on ISO/IEC 10181-3 (Ouaddah et al., 2017), an ISO standard which presents an architectural framework for authorization of internet resources and requirements for authorization protocols. Figure 4.1 shows an example of this standard. In the architecture shown in the picture, the system consists of a Policy Enforcement Point (PEP, also called Access Enforcement Facility (AEF)), where the policy decisions are enforced, and a policy decision point (PDP, also called Access Decision Facility (ADF)), where the decision for access is made. PDP's decision making can be considered as a part of authentication, while the AC policy enforcing of the PEP can be considered authorization Yang, 2016.

4.4 Implementations of solutions for authentication and authorization in the IoT

Based on the available literature, five high-level architectures are identified. The centralized architecture has a centralized PDP and PEP. The capability-based architecture has a central PDP that issues a token, which proves the holder authorization to a resource. The enforcement of this decision is done locally. With the third architecture, locally-centralized, globally-distributed, devices at the edge of the network belong to a local trust domain and access control decisions are made and enforced by

less resource-constrained devices in a trust domain. In a decentralized architecture, access decisions are made and enforced locally. And finally, with blockchain-based authentication and authorization architectures, decisions are made based on a distributed blockchain-style ledger.

4.4.1 Centralized architectures

The centralized architecture has a centralized PDP and PEP, meaning that access control policies are stored and enforced at a central point in the network, for example by a cloud server. Based on the available literature, three common models are found to use make use of this architecture. These are Role-Based Access Control (RBAC), Attribute-Based Access Control (ABAC), and XACML (eXtensible Access Contol Markup Language)

RBAC

In the early 1990s, Role-Based Access Control (RBAC) was introduced. RBAC groups users into groups known as roles. An entity may be assigned to multiple roles. Access rules are assigned to these roles, but not specific entities (Trnka, Cerny, and Stickney, 2018). Barka, Mathew, and Atif (2015) propose an architecture for prescribers of WoT (Web of Things) services to control authorization based on the RBAC model. The solution supports web-enabled (IP-based) "things". A similar solution specific to the health care sector was introduced by Chen et al. (2018). Zhang and Tian (2010) proposed a service-oriented approach which extends the RBAC model with contextual information. This context can be anything that can be used to characterize the situation of an entity, such as time or space.

ABAC

Attribute Based Access Control (ABAC), which bases access rules on attributes rather then roles, was later introduced to make authorization more context aware and flexible (Trnka, Cerny, and Stickney, 2018). In the ABAC model, access rights can be defined based on characteristics that are relevant for the system's security, known as attributes. ABAC allows for more fine-grained access control than RBAC (Ye et al., 2014). Every data consumer is in possession of a specific set of attributes (e.g., age, responsibilities, or role). An enforcement policy is a set of rules that determine whether a certain set of attributes suffice to allow data access. The subject and object described through attributes that are associated with their characteristics. The associated attributes can be defined according to the system's needs. The user is granted access to a system according to their attributes when they initiate an access request (Ye et al., 2014).

Ye et al. (2014) present a scheme for mutual authentication and authorization using ABAC. The authors propose to use ECC for establishing a secure session key. In the initialization phase, a Base Station (BS) generates necessary keys and parameter information (identity, private key, public key, hash function, elliptic curve and its parameters). After the initialization phase, mutual authentication and key establishment take place. When a user enters the network, they need to register at the gateway node. When a user initiates a request to access node resources, it needs to submit their own attribute certificate. Nodes then determine themselves whether the user is authorized to access the data. If the attributes that are presented by the user match those stored in the node, access is granted.

Huang, Wang, and Yang (2018) present DECENT, a fine-grained authorization scheme for constrained devices and the cloud based on hierarchical attribute-based encryption. This form of encryption reduces key management by introducing hierarchical attribute authorities which are responsible for key management. To reduce local computation burden, the authors propose an outsourced encryption construction, in which gateway and cloud carry out the first part of en- and decrypting data.

XACML

XACML is an XML-based language for authorization, and may be seen as an implementation of ABAC. It was standardized by the Organization for the Advancement of Structured Information Standards (OASIS). The language describes both an authorization policy language and an authorization decision language (El-Aziz and Kannan, 2013). Tasali, Chowdhury, and Vasserman (2017) present a solution that is based on both roles and attributes. The authors make use of ALFA (Abbreviated Language For Authorization), a simplified XML-based language.

4.4.2 Capability-based authorization

A capability (also known as a key, ticket, or token) is a communicable, unforgeable token of authority. The capability refers to a value that uniquely references an object along with an associated set of access rights and gives the possessor permission to access an entity or object. Put simply, the capability token grants the device access. The resource holder only has to verify the capability, but not the identity of the requester (Gusmeroli, Piccione, and Rotondi, 2013). An advantage is that the capability is sent together with the request. Therefore, the number of interactions is reduced and the authorization is simplified (Hernández-Ramos et al., 2016).

Lee, Huang, and Yang (2017) present TBAS, a token-based authorization service architecture for the IoT. The token is issued by a third-party authentication center. The framework is based on the OpenID protocol for authentication, an existing protocol that was originally not designed with the IoT in mind.

Another example comes from Seitz, Selander, and Gehrmann (2013), who adapted XACML to be specifically used for IoT devices. The author's model is built upon the observation that evaluating XACML policies is too heavyweight for constrained devices. Most of the actual authorization process is therefore externalized and capability tokens are issued to entities seeking access to a resource. The device only performs authorization enforcement. The solution is independent of transport protocol, but CoAP is suggested. As alternative HTTP(S) is mentioned, in case devices are not resource constrained.

Seitz, Selander, and Gehrmann (2013) also present a framework for authorization for the IoT based on XACML. The authors propose a framework that is independent of authentication mechanism, key management, or secure transport protocol. The authors make use of SAML (Security Assertion Markup Language) authorization decision assertions as alternative for OAuth access tokens (see Section 4.4.2). Assertions are digitally signed data objects containing asserted information.

The model by Gusmeroli, Piccione, and Rotondi (2013) supports delegation, which means that the subject cannot just grant access rights to another subject, but also grant the right to further delegate rights to others. The depth of delegation can be controlled at each stage. Capabilities can be revoked and a level of details for rights can specified. The PDP manages resource access request validation and decision. It deals with validation of access rights granted in the capability and local policies, as

well as checking the revocation status of the capability. The resource manager checks the acceptability of the capability token and acts as a PEP. This system anonymizes entity identities by hiding identity details from the service provider and guarantee the entity's identity by the trustworthiness of the identity management service itself.

OAuth

OAuth is a protocol for authorization based on the representational state transfer (REST) web architecture and can be seen as a subset of Capability-based Authorization. It is a generic authorization protocol on which implementations can be based. OAuth tokens have a limited life time and are attached to the client, so the server does not have to go back to the authorization server to check for validity. OAuth provides an authorization layer on top of a secure transport layer such as HTTPS. In standard OAuth, the user grants the client an access token, which contains both the user's and server's identity.

Limited computational power of smart objects may not be sufficient to perform cryptographic primitives required for message authentication, integrity checks, and digital signatures. Also, if access permissions for the services reside on the smart object itself, it can be hard to dynamically update them. For this reason, Cirani et al. (2015) present an OAuth-based framework for enforcement of access policies. Delegating the authorization functionality benefits from lower processing load with respect to solutions where access control is implemented on the smart object, fine-grained (remote) customization of access policies, and scalability, without the need to operate directly on the device.

Solapurkar (2016) proposes a scheme based on OAuth 2.0 extended with JSON web tokens (JWT). The scheme uses JSON web signing instead of OAuth 2.0 SSL / TLS. Within the experimental scenario mentioned in the paper, the proposed scheme was more resource efficient than the existing OAuth 2.0 scheme.

Fremantle, Kopeckỳ, and Aziz (2015) base their solution on OAuth and do not just focus on access control, but also related activities such as metadata publishing, key management, and monitoring. A centralized architecture is proposed, in which all authorization logic is externalized to a more powerful entity than the IoT device itself. Scalability and flexibility are ensured by using OAuth tokens.

4.4.3 Locally-centralized, globally-distributed authorization architectures

In the locally-centralized, globally-distributed authorization architecture, parts of the authentication and authorization process is carried out by local gateway devices. Resource constrained devices at the edge of the network reside within a "trust" domain or "bubble", in which a less constrained device is also included. The less constrained devices communicate, authenticate, authorize, and set up a secure transport layer. Constrained devices make use of derived trust.

An example of such a system is Auth, a network architecture that uses local authentication and authorization entities (Kim et al., 2016). This solution serves (locally) centralized trusted entities for local IoT nodes, but as gateway for authorization through interaction with other networks based on globally distributed trust (Kim and Lee, 2017).

Moosavi et al. (2015) present a similar authentication and authorization solution for IoT-based health care. Authentication and authorization of an end-user is done by distributed smart e-health gateways. The solution uses a certificate-based DTLS handshake protocol. The authors argue that the distributed architecture reduces the

impact of DoS attacks and that the architecture uses a more secure key management scheme between sensor nodes and smart gateways.

Gerdes, Bergmann, and Bormann (2014) propose Delegated CoAP Authentication and Authorization Framework (DCAF). In this framework, a node can delegate authentication of communication peers and management of authorization information to a trusted less resource-constrained host. The framework uses a secure channel using DTLS to secure CoAP. The system uses symmetric cryptography and keys of third parties are shared between trusted peers.

4.4.4 Decentralized authentication and authorization architectures

In distributed or decentralized authorization architectures, the decision and enforcement of authorization policies is not made at one central point in the system, but at the edge of the network, physically close to the device itself. In a distributed architecture, all access control logic is embedded into end-devices. These are enabled to obtain, process, and transmit information to other entities directly (Hernández-Ramos et al., 2016).

Mahalle et al. (2013b) use a distributed capability-based access control model. In such a capability-based model, an issuer issues a capability token to the subject. This token proves access rights to other resources. This solution makes IoT devices aware of the trust of other devices to adapt their authorization decisions accordingly. Hernández-Ramos et al. (2016) also uses a decentralized capability-based access control (DCapBAC) model which is directly implemented on resource-constrained devices themselves to ensure scalability. An authenticated key exchange is performed first to compute a session key. Then, the key is used to establish a secure channel for the second state in which the capability token is used to get access to a specific resource. Because this model is distributed, edge devices require quite some computational power. To deal with this, lightweight protocols are used such as an optimized version of ECC.

4.4.5 Blockchain-based authentication and authorization architectures

Existing distributed authentication and authorization technologies can be (partially) substituted by (immature) emerging technologies. Some of these are based on blockchain technology, created by pseudonymous Satoshi Nakamoto and famous for its application in the Bitcoin (Nakamoto, 2008). Nakamoto came up with the technology to create this electronic cash system that combines peer-to-peer data sharing with public key cryptography. Using this Bitcoin system, transactions are exchanged without the need of a third (trusted) party, such as a bank or financial institution. Blockchain functions as a distributed ledger, in which transactions are approved by majority votes by peers. Blockchains are seen as a secure technique for saving and sending data, and by some even as a promising and functional solution to some data problems in the IoT health care sector (Kang et al., 2018). In this work the focus is on blockchain solutions for authentication and authorization in the IoT and neglect blockchain solutions for storing and sharing data.

The type of encryption that is used in Bitcoin's blockchain technology is asymmetric. The public key is used to identify the owner of the coin, and to encrypt a transaction, which can only be decrypted by the holder of the private key. A blockchain is a chain of time stamped blocks that are linked by cryptographic hashes. (Fernández-Caramés and Fraga-Lamas, 2018). Peers validate a node's transaction.

The peers vote and when a majority consensus is agreed upon, the transaction is validated. A validated transaction is added to the blockchain, including a time stamp.

A few practical examples of blockchain-based IoT authentication or authorization solutions are discussed here. The advantage of using blockchain technology for authorization is that it does not depend on a central point, which mitigates the risk of a single-point-of-failure, while ensuring the integrity of the token and preventing double spending. One of the first examples of the use of blockchain technology for authorization is presented by (Zyskind and Nathan, 2015). A blockchain functions as an automated access-control manager. It is used for carrying instructions, such as storing, querying and sharing data. In the proposed solution, non-authorization related data is also saved in the blockchain, but it can also be combined with off-blockchain storage to construct a personal data management platform focusing on privacy. One of the disadvantages of this system is that it requires quite some computational resources, because besides using a blockchain for authorization, data is also saved in a blockchain. It is questionable however, if these resources are available on constrained devices. Shafagh et al. (2017) propose a similar system for storing and sharing IoT data. The system splits the data plane and control plane. Data is saved in the cloud and nodes determine which nodes have access to the data via a blockchain-based control plane.

Ouaddah, Abou Elkalam, and Ait Ouahman (2016) propose a framework for access control in the IoT based on the UTXO blockchain model and for the second-generation account model in Ouaddah, Elkalam, and Ouahman (2017). In the author's second work, two levels of access control are distinguished. The first level is concerned with the management of access policies over operations between cooperative organizations. The second level with management of access control within an organization. In the proposed system, the fist level is decentralized, and second centralized (Ouaddah et al., 2017). The authors suggest a distributed blockchain-based solution for access control aimed at providing decentralized pseudonymous and privacy preserving access control. The blockchain is adapted not to be used as crypto currency, to support new types of transactions that are used to grant, get, delegate, and revoke access. Instead of transferring coins, cryptographically signed tokens are transferred. These tokens are used to request access. The blockchain is considered as a policy retrieval point, where authorization policies for each pair (resource, requester) are stored in form of transactions. The device must send the request to the Authorization Management Point (AMP, also wallet), which also acts as Policy Enforcement Point (PEP). The PEP formulates the request to a GetAccess transaction. The PEP broadcasts this transaction to the whole network until it reaches miners. These miners act as Policy Decision Point (PDP) and evaluate the transaction by checking the request with the defined policy. This is done by comparing the unlocking script of this transaction to the locking script of the GrantAccess transaction that proceeded this transaction. The decision for access is made by executing a SmartContract. If it is allowed, the SmartContract sends the user a token to their address. This token can then be used by the requester to prove access to the client.

There are more solutions that offer solutions for authorization or authentication using blockchain for the IoT, including (Alphand et al., 2018; Hammi et al., 2018; Di Pietro et al., 2018; Kinkelin et al., 2018; Wu et al., 2018; Xu et al., 2018; Pinno, Gregio, and De Bona, 2017). These solutions are inspired on the solution proposed by Ouaddah, Abou Elkalam, and Ait Ouahman (2016) and are generally not as mature as that solution.

4.5 Comparison of solutions for authentication and authorization

The section above shows that there are many different architectures for authentication and authorization for the IoT. A comparison of the mechanisms that are used to implement these architectures can be found in Table 4.1. The reviewed solutions for authentication and authorization differ in several ways. In this section, their main differences are explained. The *"Application domain"* column refers to the application domain of the proposed solution as specifically suggested by its author(s).

TABLE 4.1: Comparison of mechanisms used for different solutions for authentication and authorization

Citation	**Implementation name**	**Framework / model**	**Transport Protocol**	**Format for saving policies or tokens**	**Cryptographic mechanisms**	**Application domain**	**Context awareness**	**Implementation (e.g. prototype, experiment or simulation)**
Barka, Mathew, and Atif, 2015	*Not specified*	RBAC	*Not specified*	Access Contol Policies are saved as XML file	Symmetric key, algorythm not specified	Any	No	No
Chen et al., 2018	*Not specified*	RBAC	*Not specified*	*Not specified*	Symmetric key, Diffie–Hellman key exchange	Any	No	Yes (simulation)
Sicari et al., 2017	IoTPlatform	ABAC	HTTPS	XML	Combination of symmetric (for data) and public key (authentication) cryptography	Health	No	No
Hemdi and Deters, 2016	*Not specified*	ABAC	CoAP	JSON attributes	No mention of encryption of messages for authentication	Any	No	Yes
Tasali, Chowdhury, and Vasserman, 2017	MDCF	XACML	*Not specified*	Uses ALFA, an abbriviated version of XML.	*Not specified*	Health	Yes	Yes (simulation)
Huang, Wang, and Yang, 2018	DECENT	Other	*Not specified*	*Not specified*	Combination of Advanced Encryption Standard (AES) (symmetric) and ECC (public key)	Any	No	Yes
Hussein, Bertin, and Frey, 2017	COBAC	Other	*Not specified*	*Not specified*	*Not specified*	Multiple	No	No
Ouaddah et al., 2015	SmartOrBAC	OrBac	RESTful (exact protocol not mentioned)	*Not specified*	*Not specified*	Any	Yes	No

Citation	Implementation name	Framework / model	Transport Protocol	Format for saving policies or tokens	Cryptographic mechanisms	Application domain	Context awareness	Implementation (e.g. prototype, experiment or simulation)
Gerdes, Bergmann, and Bormann, 2014	Delegated CoAP Authentication and Authorization Framework (DCAF)	Other	CoAP	*Not specified*	Only symmetric encryption on constrained nodes	Any	No	Yes
Kim et al., 2016	Auth	Other	HTTPS	*Not specified*	Both symmetric and private key cryptography	Any	Yes	Yes
Seitz, Selander, and Gehrmann, 2013	*Not specified*	XACML	Independent of transport protocol, suggests CoAP. As alternative HTTP(S) if devices are not resource constrained	JSON	HMAC-SHA256 for signing, IETF JSON Web Encryption for wrapping assertion and payload	Any	No	Yes
Seitz et al., 2016	ACE	OAuth	CoAP (or others)	CBOR (or others)	(D)TLS or COSE_Encrypted Wrappers	Any	No	Yes
Fremantle et al., 2014	*Not specified*	OAuth	MQTT instead of HTTP or CoAP	Any	No	Any	No	No
Cirani et al., 2015	IoT-OAS	OAuth	CoAP, but also HTTP	*Not specified*	HMAC-SHA1 (digital signature for signing tokens)	Any	No	Yes
Lee, Huang, and Yang, 2017	TBAS	OAuth	HTTP(S)	JSON	(D)TLS over HTTP and SHA-2 or SHA-3 hashing for signing tokens	Any	No	Yes
Gusmeroli, Piccione, and Rotondi, 2013	*Not specified*	Capability-based	HTTP	SAML / XACML based tokens	RSA for signing capabilities	Any	No	No

Citation	Implementation name	Framework / model	Transport Protocol	Format for saving policies or tokens	Cryptographic mechanisms	Application domain	Context awareness	Implementation (e.g. prototype, experiment or simulation)
Bandara et al., 2016	*Not specified*	Capability	HTTP	Tokens are saved in JSON, policies in XACML	RSA for signing tokens	Any	No	Yes
Mahalle et al., 2013b	IACAC	Capability	Multiple	*Not specified*	ECC-Diffie Hellman (ECCDH)	Any	No	Yes
Hernández-Ramos et al., 2016	DCapBAC	Distributed Capability based	CoAP	JSON tokens	Uses asymmetric EEC	Any	Yes	Yes
Mahalle et al., 2013a	FTBAC	Fuzzy	*Not specified*	*Not specified*	*Not specified*	Any	Yes	No
Bernabe, Ramos, and Gomez, 2016	TACIoT	Other (trust-based)	CoAP	Policies saved in XACML	ECC for end-to-end security via digital signatures	Any	Yes (but only trust)	Yes
Ouaddah, Abou Elkalam, and Ait Ouahman, 2016	Fairaccess	Blockchain	*Not specified*	In a Blockchain	PKC, algorithm not specified	Any	No	Yes
Ouaddah, Elkalam, and Ouahman, 2017	Fairaccess 2.0	Blockchain	*Not specified*	In a Blockchain	PKC, algorithm not specified	Any	Yes	No
Pinno, Gregio, and De Bona, 2017	ControlChain	Blockchain	*Not specified*	In a Blockchain	PKC, algorithm not specified	Any	No	No

4.5.1 Centralized vs. decentralized

For centralized architectures, there is only one trusted party which verifies entities' access requests. The process of verification happens via a server-side application where access policies are saved. It is easier to manage policies in these centralized architectures, but both server and client must completely trust the central entity (Roman, Zhou, and Lopez, 2013). This centralized architecture works well for traditional IT-systems in which there is one server and multiple clients. In the IoT however, this paradigm is reversed by having many devices (which function as servers) and possibly many clients (Ouaddah, Abou Elkalam, and Ait Ouahman, 2016). The size of some networks makes it hard to create efficient centralized authentication systems. Decentralized solutions aim to solve this problem. For these solutions, there is no centralized trusted party. Instead, the participants coordinate autonomously to build further trust (Kim and Lee, 2017). In this way, the IoT benefits from connectivity that facilitates collaboration among nodes and takes advantage of emerging edge computing technologies (Kim and Lee, 2017). A downside of these distributed architectures however, is that they are more complex and harder to manage (Ouaddah et al., 2017).

Some argue that the centralized approach is not scalable and does not support future growth of the IoT. While others (e.g., Barka, Mathew, and Atif, 2015) claim that RBAC is a scalable solution, because it reduces the number of tasks that the devices at the edge of the network have to carry out. This makes the finding the equilibrium between end device autonomy to control access over produced information and computing efforts requested by authorization mechanisms an open issue (Ouaddah et al., 2017).

4.5.2 Virtualization of network functions and Software Defined Networking

Until now, a system's architecture has been regarded as a static design decision, that cannot change dynamically. Technologies like Software Defined Networking (SDN) Network Function Virtualization (NFV) however, are capable of dynamically changing design and operation of functions by delivering virtual appliances at the edge of networks. This allows the creation of virtualized services to existing architecture and counter measures such as firewall rules or DDoS mitigation (Zarca et al., 2019).

NFV delivers virtual appliances at the edge of networks and allows the creation of virtualized services to existing architecture and counter measures such as firewall rules or DDoS mitigation . For these solutions, authorization and data planes are split. One of them is done centrally, while the other is done locally (Kim and Lee, 2017).

A problem with IoT is oversight: there are lots of devices, which makes them hard to manage, they are far from the one who controls them, are static (once in place they stay there) and are therefore hard to monitor. SDN and NFV also offer centralized approaches for identity, authentication, and authorization management.

4.5.3 Attributes or identities

Most solutions for authentication and authorization rely on identities. Some solutions however, do not require the explicit identity of users or devices, but attributes for authentication. Some argue that using attributes instead of identities preserves the privacy of users better, because users do not have to give away their identity in case it is not essential (Alpár et al., 2016).

FIGURE 4.2: An example of a high-level overview of a networked system with centralized (yellow), hybrid (green), or decentralized (blue) architectures.

4.5.4 Transport protocol

RESTful systems usually communicate through HTTP(S) (Hypertext Transfer Protocol). In the case of the IoT, the constrained application protocol (CoAP) is often used. HTTP(S) is a very common, mature and widely supported protocol. It is however relatively old and not originally designed with the IoT in mind. To bridge the gap between the internet and the physical world and to allow smart objects into the internet, effort has been made to adapt existing protocols to interoperable, but more efficient used for constrained environments (Hernández-Ramos et al., 2016). Seitz, Selander, and Gehrmann (2013) suggest the use of CoAP instead of HTTP(S) because CoAP is specifically designed for constrained devices and features a low overhead compared to HTTP.

4.5.5 Policy or token storage format

JSON and XML are standards that can be used for the storage of access control policies, but also for expressing access tokens or capabilities. Hernández-Ramos et al. (2016) argue that JSON is more suitable than XML in IoT scenarios because it is more lightweight. Seitz, Selander, and Gehrmann (2013) mention that the use of JSON instead of XML for assertions reduces the size of the assertion by roughly a factor of ten. Another approach has been taken by Tasali, Chowdhury, and Vasserman (2017), who use the Abbreviated Language for Authorization (ALFA), a simplified version of XACML. On average, the policies that the authors wrote contained 528 non-whitespace characters in ALFA, compared to 3903 in XACML.

4.5.6 Cryptographic measures

Both PKC and SKC are used for authentication and authorization in the IoT. SKC is computationally more lightweight, but keys must be shared securely and keys must be stored for every pair of communicating devices. A full handshake protocol may be computationally intensive for a constrained IoT device. The initial handshake may therefore be delegated to the owner of an object or a device with more computational resources, such as smart gateways (Atwady and Hammoudeh, 2017). It also limits scalability in case a device is connected to a very large number of devices (Miettinen et al., 2018). Hussein, Bertin, and Frey (2017) deal with this problem by only sharing keys between devices that are physically close to each other (edge devices and the fog). Another alternative is by using PKC to overcome these challenges (Ye et al., 2014). The problem with PKC however, is that it is computationally harder, so lightweight encryption algorithms are often used for this, such as ECC (Hernández-Ramos et al., 2016).

4.5.7 Maturity of solutions

The solutions that are compared here differ significantly in maturity and prevalence. In order to prove the feasibility of the approaches compared here, some authors provide a proof of concept, prototype, experiment or simulation. Other solutions, such as those built around OAuth are more mature solutions have already been implemented in real life situations. This shows that some of these approaches are more mature than others.

4.5.8 Blockchain solutions

Within a blockchain, all users are identified by their public key or hash (Fernández-Caramés and Fraga-Lamas, 2018). Therefore, anonymity is not guaranteed by default. For Bitcoin for example, a transaction is a cryptographically signed transfer of funds from one public key to another. So while payer and payee are not explicitly matched to a real-world entity, all transactions are transparent and can even be traced in some cases (Meiklejohn et al., 2013). This reduces privacy, and this may become even worse if a private key is compromised. As Bitcoin creator Satoshi Nakamoto says: *'If the owner of a key is revealed, linking could reveal other transactions that belonged to the same owner'* (Nakamoto, 2008). There are some solutions to this problem proposed in the literature. For example, (Fernández-Caramés and Fraga-Lamas, 2018) mentions the idea to create a unique address for every transaction or for every counter party, but they do not offer any suggestions on what such a solution could look like in practice. Probably, a more lightweight approach is needed because of the high overheads that are related to this. (Dorri, Kanhere, and Jurdak, 2017) uses a private immutable ledger that is managed centrally to make sure that the ledger is kept private. Storing the ledger centrally however, introduces possible risks that are related to having a single point of failure.

4.6 Conclusion

This chapter started with the question: *"What are the characteristics of currently available IoT authentication and authorization solutions?"*. An overview of current literature on authentication and authorization for the IoT is presented, focusing on the "higher levels" of authentication and authorization systems. To restrict or allow access to resources, a secure connection must be established, users and devices must be assigned with identities and granted access to resources, after which they must authenticate themselves before accessing resources. How this takes place in practice differs per implementation. An important difference is where the decisions for allowing or restricting access are made or enforced (i.e., where the PEP and PDP are placed) Different frameworks, protocols and tools for authentication and authorization exist, five different high-level architectural designs for authentication and authorization are introduced, including centralized architectures that are based around RBAC, ABAC, XACML, Capability-based Authentication, which relies on an issuer of capability tokens, such as the OAuth protocol, Locally-centralized / globally-distributed, or decentralized architectures, where access control decisions are not made at a central point in the network, but closer to the edge of the network.

This chapter shows that there is not one single form in which authentication and authorization takes place, but rather a large number of different possibilities. This reflects the large number of different contexts in which different types of IoT devices operate. Every sector, from health care to transportation, to utilities and

consumer electronics have their own IoT solutions. As Ouaddah et al. (2017) already noted, there is no single approach (always) better than another. All of them have advantages and disadvantages. Different design approaches seem complement each other, rather than compete with each other. Because of the diversity of these design approaches, it does not seem easy to present guidance to when to use a which design approach.

Chapter 5

Authentication, authorization, security, and privacy

5.1 About this chapter

Based on the findings in the previous chapters, this chapter aims to explain the effect of solutions for authentication and authorization on privacy and security in the IoT in health care, answering the sub-question: *"To what extent do current solutions for authentication and authorization meet the general security and privacy objectives and heath care requirements?"*. Section 5.2 compares objectives and requirements for secure and privacy preserving authentication and authorization to current solutions. Several guidelines for secure and privacy preserving authentication and authorization are presented in Section 5.3. Section 5.4 concludes the chapter.

5.2 Objectives and requirements compared to current solutions

5.2.1 Objectives and challenges for security and privacy in IoT

The previous chapters show that the most important security and privacy related objectives for IoT devices in health care are related to privacy, confidentiality, and integrity of information. As shown in Chapter 3, achieving these objectives is complicated by cyber attacks, scalability issues, data management problems, and a lack of standardization.

To achieve these objectives, access control plays a key role. Based on the interview results, we can make several observations on the implementation, design, or operating of authentication and authorization in practice. First, authentication is commonly based around unique identities, and limited use is made of alternative methods for authentication that are described in this chapter. Second, current solutions for authentication and authorization are focused on centralized architectures. The most common way to authenticate individuals is by using federated identity management, in which an identity provider intervenes between a user and a service provider to authenticate a user. The (potential) effects of these solutions on privacy and security are discussed in this chapter.

Health care organizations however, seem to take organizational and technical measures to make sure privacy of users is preserved. In general however, always-identity authentication is common and their negative privacy effect of always-identity for authentication or centralized architectures is not experienced. When aiming to preserve privacy, the focus is usually on what (and how much) data is gathered, stored, or processed, not on how users or devices authenticate themselves.

5.2.2 Alternative approaches for authentication

In Chapter 4 it is explained that authentication and authorization are important to ensure a minimum level of privacy, as access to personal information should be limited to authorized entities. Evaluating the proof of a requester's identity is traditionally considered an important prerequisite for authorization. A set of data with respect to a particular individual is called the individual's identity. Usually, at least one of the pieces of available information acts as an identifier, which is a direct link between the individual and the system (Alpár et al., 2016). An example could be a user name. Authentication protocols often reveal the full identity of a user, which may be more information than necessary for secure authorization, as a proof of permission to a resource may suffice (Alpár et al., 2016). These types of always-identity mechanisms for authorization often rely on a central identity management authority to assess the identity of access requesters before making an authorization decision (Hussein, Bertin, and Frey, 2017). These always-identity solutions for authorization have several downsides.

First, relying on identities for authenticating a requester introduces more complexity in IoT scenarios where IoT devices' identities are hard to maintain and assert (Hussein, Bertin, and Frey, 2017). Second, user's privacy might be threatened because transactions that are linked to the same identifier are traceable, which introduces privacy risks (Alpár et al., 2016). Therefore, authentication and authorization may hurt privacy. Generally, the more we know of an entity, the more secure we can design a solution for authentication and authorization, but privacy may be harmed because every action of a user can be tracked, potentially without the user knowing this. Some suggest that privacy is damaged in centralized (cloud) IoT architectures and with always-identity authentication scenarios (Krasnova, 2017). In Chapter 4 it is explained that authentication and authorization are important concepts to ensure a minimum level of privacy. Authentication and authorization may however be designed in a way which hurts privacy.

There are multiple ways to preserve privacy by not always having to prove a full identity every time requests for access to a resource or service is made. For example, a system can make use of a capability-based approach, in which one computational entity evaluates access policies and issues a capability token which allows direct access to a resource. Assuming that the token does not contain the requester's identity, the client does not know who accessed the resource. A downside however, is that there is one centralized entity (the authorization server) that still has to verify the user's identity, and therefore may record all actions the user takes.

Another approach is by adopting a pseudonym instead of an identity. An advantage is that this pseudonym is not directly identifiable to a human user. This does not mean however that these pseudonyms cannot be tracked, and does not guarantee that the user is not identifiable by using additional information from other sources.

Another solution is to use attributes for authentication instead of identities. Attribute-based authentication relies on attributes, or characteristics of users or devices embedded in cryptographic containers rather than identities (e.g., current location, owner, or manufacturer). This makes a unique identity become a variable in a list of attributes needed to evaluate an access request. Attribute-based authentication may be a solution, in which only the essential information is verified. According to Krasnova (2017, p. 33), attributes such as the device's brand, user nationality can be anonymous, allowing authentication can be achieved by using the minimum amount of information required to successfully access a resource.

Context-based authentication solutions are for mutual authentication of devices that belong to the same trust domain (e.g., have the same owner). These solutions make use of contextual information that the devices' sensors observe in their physical environment. This context is used to mutually authenticate devices that are closely located (i.e., share the same context). A context fingerprint is derived from the physical features of the context by the device. Another device that shares the same context (e.g., because it is in the same room) has the same context. This shared context is used as a shared secret used for mutual authentication (Miettinen et al., 2018). Kalamandeen et al. (2010) propose a solution to pair devices that share the same received signal strength with a third, observing device. The system that is proposed determines if two devices are in close proximity by comparing the strength of the signal between the two devices and a third observing device. Miettinen et al., 2018 compared several different solutions for context-based authentication. The authors concluded that it is possible for two devices to establish a shared secret, as long as sufficient time is given to analyze the environment. For contexts where a complete contextual separation from the larger environment cannot be made, the authentication process should be repeated several times.

In practice, current solutions for authentication and authorization only make very limited use of these privacy-preserving methods. Alternative forms of authentication such group signatures, tokens, pseudonyms, attributes etc. are available, but, except for pseudonymization, not desired by the interviewed experts. This has several reasons. For example, devices are only used for one purpose; usually measuring or monitoring. In those cases, it does not make sense to hide an identity, as the information must be linked to an individual in order to be useful.

5.2.3 Decentralized approaches for authentication and authorization

Based on the interviews that are carried out, it can be concluded that in practice, there is a reliance on centralized architectures for IoT systems. There may be a preference for on-premise centralized solutions instead of using cloud solutions because of privacy reasons, compliance with legislation, and best-practices. A centralized approach is said to be easier to secure. According to the interviewees especially accountability, auditability, trustworthiness, and non-repudiation are said to be achieved more easily in centralized solutions.

IoT solution architects have generally preferred centralized authentication and authorization due to their convenience for users. They allow for flexible authentication, but the identity provider is involved in all transaction, so it may trace users connecting to services. According to Wachter (2018), they face challenges of scalability, cross-border governance harmonization, and pose privacy risks to users as they allow greater exchange and linkage of potentially sensitive personal data between service providers. According to Krasnova (2017, p. 32), decentralized IoT architectures are preferred over centralized architectures in terms of privacy.

Technologies, such as SDN and NVF which virtualize networks, shift centralized computation to the edge of networks, delivering better throughput, and more context-specific functions. These technologies can be used for carrying out identity, authentication, authorization management and offer more flexibility (Zarca et al., 2019). Identity, authentication, authorization, and accountability may be embedded in the design of a platform, offering opportunities for security (Zarca et al., 2019). In terms of authentication, these virtual technologies may allow key management,

reducing strain on IoT devices and increasing scalability. In terms of accountability the technologies may increase network visibility, which has security benefits and may enable new flexible and powerful network solutions.

Even though the security related benefits of SDN / NVF can be large, there is very little known about their potential privacy effects. Even though functions that make use of the advantages of SDN / NVF are more decentralized, with more computing and decision making at the edge of networks, closer to users themselves, it is still questionable whether linking and tracing of users is really less likely in these scenarios, due to central monitoring and control. The real and practical effects of these technologies on privacy are largely unknown. Identifying potential risks and conducting privacy impact assessments may be therefore be necessary in case these technologies are used by organizations to ensurea user's privacy is not harmed and to ensure compliance with legislation can be guaranteed.

5.3 Guidelines

Based on the findings in the literature and on the general consensus of the interviewed experts, some general guidelines for secure and privacy-preserving authentication and authorization for the IoT in health care are mentioned. The guidelines aim to find a balance that satisfies the security and privacy objectives. The exact reasoning behind this table is found below. The guidelines aim to find a balance that satisfies the security and privacy objectives that are introduced in Chapter 1. A set of guidelines per requirement is presented, as can be seen in Table 5.1. The model consists of three layers. The first layer or the edge-side layer consists of edge nodes, communication, and edge communication. The server / cloud-side layer, which consists of data accumulation and data abstraction. This layer is generally responsible for centralized decision making and enforcement. The user / application layer is responsible for information interpretations, and users (note that this does not always has to be the patient). For each of the 14 identified objectives, an explanation of these guidelines is given.

TABLE 5.1: Guidelines for secure and privacy-preserving authentication and authorization for 14 objectives

	Guidelines		
Area to consider	**User-side layer**	**Server / cloud-side layer**	**Edge-side layer**
Objective: Transparency; the system must allow people to understand who knows what about them, how their data will be used, with whom it is shared and how long it is held.			
Individuals may not be able to access or understand access policies, and may, as a result not trust the IoT technology.	Ensure transparency concerning authorization policies for end-users.	Identify where personal information is stored. Monitor and review who accessed personal information.	Allow user to understand when a device collects data and ensure understanding of what, how, and where personal information is stored or processed on the device.
Objective: User-driven; the system must allow users to have full and granular access control over the data they share in the network or in the cloud.			
The user may not agree with default authorization policies.	Allow user to influence authorization policies (in)directly.	Maintain fine-grained and dynamically updatable access control and ensure access is limited to authorized users, devices, and processes. Formally manage assets throughout removal, transfers, and disposal.	Allow users to determine when a device gathers data.
Objective: Anonymity; IoT applications must not disclose the identity of their users.			
Identity attributes may be visible to third parties when interacting with the IoT.	Allow anonymization of data (e.g. for research purposes).	Use privacy-enhancing cryptographic techniques such as group signatures or attribute-based authentication to make identity attributes less visible to third parties.	Ensure that there is no loss of identity attributes or personal data after disposal or loss a of device.
Objective: Pseudonymity; the system must link actions of a person with a pseudonym rather than an identity; trades off anonymity with accountability.			
Individuals may be directly identifiable based on the identifiers used for authentication.	Use pseudonyms by default.	Use pseudonyms by default.	Ensure that there is no loss of linkable identity attributes or personal data after disposal or loss a of device.

	Guidelines		
Area to consider	**User-side layer**	**Server / cloud-side layer**	**Edge-side layer**
Objective: Unlinkability; the system must not link specific actions of the same person should together unless necessary.			
Users may be subject to undesired linking of data by service providers or other third parties.	Achieve data minimization by limiting the amount of data gathered. Reduce the use of unique data per entity.	Limit the amount of identification attributes gathered.	Use methods to limit tracing of tags, such as pseudonymous or anonymous tagging, physical isolation, kill / sleep commands, blocking, or personal firewalls.
Objective: Unobservability; the system must not allow users and / or subjects to determine whether an operation is being performed by another user.			
IoT devices may collect, process, or store more identifying attributes then necessary for the correct working of the system.	Allow data anonymization	Limit data retention and keep in mind that event logs may include personal information.	Use privacy-enhancing cryptographic techniques or alternative authentication measures (non-identity) or multi-factor authentication.
Objective: Integrity; the system must prevent unauthorized modications of resources.			
Unauthorized entities may modify data.	Prevent unauthorized modification of data in front-end applications.	Ensure logical and physical access limitation of data storage to authorized users and encryption of logging.	Make use of sufficiently strong secured channels for exchange of information; make use of at least transport-layer security. Ensure that the device has built-in physical security controls to protect it from tampering to ensure the integrity of device identifiers and tokens. Ensure that hardware allows protection of for example source code and firmware.
Objective: Availability; the system must support a high readiness for usage. Offline mode and short- and long-time availability.			
A lack of availability may lead to connectivity or reliability issues.	Monitor and review historic availability.	Ensure correct working of fail-over systems.	Make use of protocols that allow limitations of connected devices (e.g. CoAP instead of HTTPS when necessary).

	Guidelines		
Area to consider	**User-side layer**	**Server / cloud-side layer**	**Edge-side layer**
Objective: Confidentiality; the system must prevent unauthorized disclosure of resources through granular (fine-grained), revocable, delegatable access control.			
Unauthorized entities may access personal data.	Frequently review access control policies, monitoring, and alerting.	Maintain fine-grained access control to ensure access is limited to authorized users, devices, and processes. Formally manage assets throughout removal, transfers, and disposal.	Ensure devices restrict each user, device, and process to the minimum access privileges necessary and prevent unauthorized access to all sensitive data stored on the device. Ensure data that is stored on the device can be encrypted and sanitized.
Objective: Usability; the system must allow access control to be easily managed, expressed and modied.			
Users of the IoT environment may not be the intended users because of a failure to manage access policies well.	Ensure a clear and accurate overview of access policies.	Allow fine-grained access control.	Allow users to determine when a device may interact with other devices.
Objective: Accountability; the system must be able to hold users responsible for their actions (e.g. misuse of information).			
Unauthorized entities, attackers, or malicious nodes may not be identified.	Identify malicious entities using incident analysis activities.	Ensure vulnerability management tasks, such as vulnerability scanning, are carried out. Allow removal of malicious devices from networks. Ensure identities and credentials are issued, managed, verified, revoked, and audited for authorized devices, users and processes and that the organization's (asset) management system can access or understand devices' and users' identities.	Allow devices to interact with enterprise asset management systems. Ensure the device can uniquely identify and authenticate each user, device, and process attempting to logically access it and that the device can thwart attempts to gain unauthorized access.

	Guidelines		
Area to consider	**User-side layer**	**Server / cloud-side layer**	**Edge-side layer**
Objective: Auditability; the system must be able to conduct persistent monitoring of all actions.			
The organization may face a lack of visibility of unauthorized actions.	Achieve network visibility, monitoring, and alerting.	Ensure identities and credentials are issued, managed, verified, revoked, and audited for authorized devices, users and processes. Ensure that authentication of users, devices, and other assets is in line with the risk of the transaction. Achieve network visibility and dynamically updatable trust and reputation of devices on the network. Perform regular vulnerability scans.	Make sure the device has an unique identifier. Allow devices to interact with existing enterprise log management systems. Ensure that the device supports the use of vulnerability scanners or provides built-in vulnerability identification and reporting capabilities.
Objective: Trustworthiness; the system must be able to verify identity and establish trust in a third party.			
Inadequate estimates of trust or reputation of devices may occur due to changes in devices' or services' perceived level of security.	Allow network visibility using monitoring, alerting, and base actions based on this visibility. Inventarize physical devices, software platforms, assets, and systems.	Reveal dependencies on external services and information systems. Monitor external party's trustworthy by reviewing SLAs, certification, or auditing. Monitor whether the manufacturer provides patches or upgrades for all software and firmware throughout each devices' lifespan.	Ensure that hardware allows protection of source code and firmware. If the device does not have secure built-in patch, upgrade, and configuration management capabilities, make sure it can interface with enterprise vulnerability management systems with such capabilities. Ensure the device is capable of having its software patched or updated. Ensure awareness of all external software and services the device uses, such as software running on or dynamically downloaded from the cloud.
Objective: Non-repudiation; the system must be able to conrm occurrence or non-occurrence of an action.			
Actions performed by malicious entities cannot be made undone.	Have monitoring and alerting in place based on event logs.	Enable logging of critical processes and actions. Ensure integrity and availability of these logs, e.g. by encrypted logging, and backup and restore processes	Ensure the device is able to log its operational and security events in sufficient detail.

Transparency

A system is said to be transparent if it allows people to understand who knows what about them, how their data will be used, with whom it is shared and how long it is held (Ouaddah et al., 2017). Not being able to achieve transparency may cause that a user is not aware of access policies, is unaware of consequences of use of IoT-technology, and may, as a result not trust it. Security and privacy breaches may undermine trust among IoT users, objects, and device or service providers. From the interviews it became clear that users should at least have insights whom have rights to access their personal information. Additionally, users may have different expectations, levels of awareness, and desires regarding privacy. On a user-side layer there should be transparency about authorization policies for patients and other end-users so it becomes clear who may have access to their personal information. As a first step in protecting personal information, organizations should identify all locations where it is stored. Monitoring who accessed or had access to personal or confidential information on a server / cloud side layer may also help ensuring transparency. This can be achieved by periodically reviewing historical logging, taking into account who needed access based on a role within an organization. At the edge of a network, on the device itself, it must be clear an IoT device gathers data. Furthermore, there must be understanding of what, how, and where personal information is stored or processed on the device.

User-driven

A system is said to be user-driven if users have full and granular access control over the data they share in the network or in the cloud (Ouaddah et al., 2017). A potential risk is that the patient or user does not agree with default authorization policies, which may lead to a lack of trust or disagreement. On the user-side layer, users should be able to influence authorization policies, either by controlling them themselves or through another party. On a server / cloud-side layer, an IoT system must support fine-grained and dynamically updatable access control to ensure access is limited to authorized users, devices, and processes. Assets should be formally managed throughout removal, transfers, and disposal of devices. On the edge-side layer, users must be able to determine for themselves when personal data is gathered.

Anonymity

Anonymity is the quality that IoT applications do not disclose the identity of their users (Ouaddah et al., 2017). A failure to do so may lead to an (undesired) exposure of identity attributes to third parties when interacting with the IoT. On the highest level, the user-side layer, it should be possible to anonymize data for secondary purposes, such as research purposes. Preferably, this should be done in a way that guarantees that the individuals who are the subjects of the data cannot be re-identified based on the available data. Protection models such as those proposed by Sweeney (2002) may be used for this. On the server / cloud-side layer, privacy enhancing cryptographic techniques may make identifiers that are used for authentication less visible to transmitting parties. Examples of these techniques are discussed in Chapter 2 and include group signatures and some instances of attribute-based authentication. On the edge-side layer, assure that there is no loss of identity or personal data on the device itself after disposal or loss of device. This means that personal or other

confidential data must either not be stored on the device itself (on a longer-term basis), or be securely encrypted in a way that cannot easily be broken, now or in the foreseeable future.

Pseudonymity

Pseudonymity is said to be achieved in case a system must link actions of individuals with a pseudonym rather than an identity; it therefore trades off anonymity with accountability (Ouaddah et al., 2017). In case pseudonymity cannot be achieved, individuals may be directly identifiable based on the identifiers used for authentication. On all the three layers suggested by Mosenia and Jha (2017), pseudonyms should be applied by default. It must be noted however that pseudonymity does not protect against linking and tracing of pseudonyms and pseudonymization alone does therefore not ensure user's privacy. On the edge-side layer, pseudonymity may help ensuring there is no loss of directly linkable identity attributes or personal data in the device after disposal or loss of the device.

Unlinkability

Unlinkability is achieved if specific actions of the same person will not be linked together unless necessary (Ouaddah et al., 2017). Failing to achieve unlinkability may lead to undesired or unwelcome linking, combining, or enhancing of personal data by service providers or other third parties. On the user-side layer, the use of unique data per entity should be reduced to make sure individuals (who are the subjects of the data) cannot be re-identified based on the data itself. On the server / cloud-side layer unlinkability can be accomplished by data minimization, or by monitoring service providers or other third-parties, for example via review of documentation, Service Level Agreements (SLAs), auditing, audit statements, or certifications. On the edge-side layer, different kinds of devices may be linked or traced. Solutions could be pseudonymous tagging of RFID tags or other device, in which the identity of tags is protected using mapping algorithms or encryption, anonymous tagging, in which the used pseudonym is re-issued frequently, physical isolation, kill / sleep blocking, in which the tag is temporarily or permanently blocked to prevent unauthorized access, or personal firewalls for less constrained devices (Mosenia and Jha, 2017). These (technical) measures ensure that there is no information stored on the device itself that can be related to individuals.

Unobservability

Unobservability is achieved if users or subjects cannot determine whether an operation is being performed by another user (Ouaddah et al., 2017). Failing to achieve unobservability may lead to privacy threats due to logging or monitoring more then necessary for the correct functioning of the system. Other privacy risks include user tracking and localizing, which permit the creation and misuse of detailed user profiles (Mosenia and Jha, 2017). On the server / cloud-side layer, anonymization of information or actions may help achieving unobservability. On this layer, where data is stored, gathering of information should be limited to only as much data as required to successfully accomplish a given task, known as data minimization. Furthermore, limiting data retention may also have positive effects for unobservability. Organizations should be aware that event logs may also contain personal information. On the edge-side layer, unobservability can be realized party by using

alternative authentication measures that do not make use of directly recognizable identities on a device level such as multi-factor authentication, group signatures, or attribute-based authentication. As mitigation, RFID systems must therefore provide anonymity, even when the state of a tag has been disclosed (Mosenia and Jha, 2017). This may be accomplished using cryptographic techniques.

Integrity

Integrity is the property of accuracy and completeness (ISO, 2000). A potential risk in case integrity cannot be guaranteed is modification of data by unauthorized entities. On the user-side layer, unauthorized modification of data using front-end applications must be prevented. Both logical and physical access to data storage should be limited. Besides that, logging should be secured to make sure attacks or direct data changes cannot be made invisible. Data should be transmitted over sufficiently strong secured channels. There are many cryptographic ways of achieving this, each having their own advantages and disadvantages. On a very basic level, the channel is secured using for example transport-layer security. In case more or more confidential information is transmitted, more advanced methods can be applied that deliver end-to-end security, such as virtually or physically isolated connections. The interviewees in this study agreed that information should not be transmitted over insecure channels. Besides using other cryptographic measures, using a VPN for transmitting health care information over the internet is seen as a practical, yet attainable and necessary solution that should be used in many cases. Besides that, organizations should ensure the integrity of device identifiers and tokens. This can be achieved by having built-in physical security controls on the devices' hardware that allows protection of source code, firmware to protect it from tampering

Availability

Availability is the property of being accessible and usable on demand by an authorized entity (ISO, 2000), indicating that the system must support a high readiness for usage. Potential risks include connectivity or reliability issues. Monitoring and review of historic availability on a user-side layer may be useful as a first step in ensuring availability. Fail-over systems on the server / cloud-side layer should be in place and their correct working should be ensured. To ensure availability, organizations should make use of protocols that suit limitations of connected devices. Using CoAP (Constrained Application Protocol) for example, may be preferred over HTTP(S) in some situations.

Confidentiality

Confidentiality is the property that information is not made available or disclosed to unauthorized individuals, entities, or processes (ISO, 2000). Confidentiality should be attained through granular (fine-grained), revocable, delegatable access control. A potential risk in case confidentiality cannot be ensured is unauthorized access of personal data. Frequent review of access control policies supported by monitoring, and alerting for unauthorized or suspicious actions on the user-side layer helps achieving confidentiality. On the server / cloud-side layer, fine-grained access control, or granular, flexible, dynamically updatable, scalable, available, context-aware, heterogeneous, collaborative policies can contribute to confidentiality of information by ensuring access is limited to authorized, users, devices, and processes. Assets should

be formally managed throughout removal, transfers, and disposition. On the edge-side layer, unauthorized modification of data in front-end applications must be prevented. Both logical and physical access to data storage should be limited. Besides that, logging should be encrypted to make sure attacks or direct data changes cannot be made invisible. Data should be transmitted over sufficiently strong secured channels.

Usability

Usable access control models are those which reduce user effort in system administration and facilitate more autonomous establishment of security context (Ouaddah et al., 2017). This means that access control should be easily managed, expressed, and modied. If this is not achieved, there is a risk that desired access control policies do not reflect actual policies because of a failure to manage access policies well. On the user-side layer, a clear, timely and accurate overview of access policies should be available. On the server / cloud-side layer fine-grained access control should be in place. This means that rights can be assigned to users or devices individually, not having to rely on group or role-based access control. On the edge-side layer, users should be able to determine when a device may interact with other devices.

Accountability

Accountability is the ability of a system to hold users responsible for their actions (Mosenia and Jha, 2017). Accountability is necessary to ensure that proper steps to prevent, resolve, or mitigate attacks can be taken, when unauthorized entities, attackers, or malicious nodes are identified. On the user-side layer, it should be possible to identify malicious entities. Incident analysis activities and centralized network visibility, monitoring, and alerting may help achieving this. On the server / cloud-side layer, achieving network visibility and dynamically updating trust and reputation of devices on the network is necessary to achieve accountability. This means that vulnerability management tasks such as vulnerability scanning are carried out. Based on continuous monitoring, malicious devices must be able to be removed from networks. Identities and credentials should be issued, managed, verified, revoked, and audited for authorized devices, users and processes. Organization's (asset) management system should be able to access or understand identities and credentials for devices as well as individuals (Boeckl et al., 2018). On the edge-side layer, devices should have unique identifiers. Devices that operate on this level may be black boxes that provide little or no information on hardware, software and firmware and cannot provide sufficient visibility into characteristics of the device. Logging may not be available on edge device themselves due to the constrained nature of some devices. Therefore, interaction with available log management systems should be allowed. The device should support the use of vulnerability scanners or provides built-in vulnerability identification and reporting capabilities. On the edge-side layer, the used hardware should allow protection of source code, firmware, or processor. Devices should be able to interact with enterprise asset management systems and should be able to uniquely identify and authenticate each user, device, and process attempting to logically access it and thwart attempts to gain unauthorized access (e.g., a lock-out after several unsuccessful log on attempts).

Trustworthiness

Trustworthiness is the ability of a system to verify identity and establish trust in a third party (Mosenia and Jha, 2017). Perceived security of devices may change over time. Therefore, inadequate estimates of trust or reputation of devices may occur due to changes in a device's or service's perceived level of security. On the user-side layer, network visibility, monitoring, and alerting should be in place and actions should be based on this visibility. Physical devices, software platforms, assets, and systems should be registered. On the server / cloud-side layer, dependencies on external services and information systems should be revealed. The trustworthiness of these external parties should be classified by means of certification, or auditing. It should be clear whether the manufacturer provides patches or upgrades for all software and firmware throughout each devices' lifespan. In case the devices cannot be trusted, the ability to remove malicious devices from a network should exist. On the edge-side layer, hardware should allow protection of source code, firmware, and processor in order to be trustworthy. Devices should be able to interface with enterprise vulnerability management systems if the device does not have secure built-in patch, upgrade, and configuration management capabilities itself. It should be ensured that the device is capable of having its software patched or up-to-date awareness of all external software and services the device uses, such as software running on or dynamically downloaded from the cloud.

Non-repudiation

Non-repudiation is the ability of a system to confirm occurrence or non-occurrence of an action (Mosenia and Jha, 2017). A potential threat that is that actions performed by malicious entities cannot be revoked. To achieve non-repudiation there should be monitoring and alerting in place on the user-side layer. Furthermore, there should be event logging of critical processes and actions available on the server / cloud-side layer. Integrity and availability of these logs should be guaranteed by for example encrypted logging, and backup / and restore processes. On the edge-side layer, devices should be able to log operational and security events in sufficient detail. Note that logging on the device itself may or may not be possible.

Decentralization

Decentralization is said to be achieved if each node in the network shares its data with other nodes directly, without intervention of any third or trusted entity (Ouaddah et al., 2017). In most cases however, centralized architectures are used in which nodes mostly communicate with cloud servers (or on-premise centralized servers). In the opinion of the interviewees under study, decentralization would not lead to differences in privacy or security, as long as organizational or legal measures are taken to prevent privacy problems that may be associated with centralized approaches. Centralized architectures are said to be easier to secure and monitor for malicious software or entry. Because of these reasons, decentralization is not included as an objective by itself for authentication and authorization. This does not mean however, that increased decentralization does not have privacy benefits. The potential effect of decentralization on a user's privacy is further discussed in Section 7.2.

5.4 Conclusion

"This chapter started with the question: *"To what extent do current solutions for authentication and authorization meet the general security and privacy objectives and heath care requirements?"*. The previous chapter shows that the most important objectives for IoT devices in health care are related to privacy, confidentiality, and integrity of information. This chapter shows that authentication and authorization play a key role in achieving these goals. Authentication and authorization are essential to provide security and privacy. Insecure authentication however may lead to exposure of identifiers, and insecure authorization may lead to exposure of confidential information or stolen access tokens. Cryptographic measures may be taken to ensure security. To preserve the privacy of users and protect personal data in the IoT, several privacy preserving approaches can be used, such as data minimization, pseudonymization, group signatures, or attribute-based authentication. In practice, centralized architectures and an always-identity approach are common and the potentially negative privacy effect of always-identity for authentication or centralized architectures is not experienced, as technical, legal, and organizational measures are taken to preserve privacy. Alternative forms of authentication and authorization are available, but not desired. Edge and fog technologies such as SND and NVF shift centralized computation to the edge of networks. Virtualizing networks offers the capability to have full control over networks, which potentially has a positive effect on security. However, as the control that these entities have increases, and data is gathered about information that may be linked to human entities of a platform. Finding a balance between all objectives may be hard. Based on literature and interviews, a set of guidelines is proposed to evaluate the extent to which security and privacy objectives are satisfied.

Chapter 6

Validation

6.1 Approach

In this chapter, the guidelines presented in the previous chapter are validated. Validation refers to the process of collecting validity evidence to evaluate the appropriateness of the interpretations, uses, and decisions (Cook and Hatala, 2016) of an artifact. There are different ways to validate an artifact. In this work, validation interviews are carried out. In these interviews, relevant domain experts are asked their professional opinion about the artifact under study.

These interviews are carried out in a similar fashion as the first set of interviews, however the sample size is smaller (5). and included both interviewees that were already interviewed, as well as interviewees that were not interviewed previously. Therefore, some of the validating interviewees were the same as that were consulted during the first set of interviews which ultimately gave input for the guidelines themselves. New interviewees were asked to discover whether new interviewees could also understand the process.

During the validation interviews, several questions relating to the appropriateness of the artifact have been asked asked. To examine the usefulness when applying the guidelines in practice the question *"Would applying the guidelines improve security and privacy within the domain?"* is asked. To assess the correctness of the guidelines the questions *"Are the guidelines proper transformations of the areas to consider? Are the areas to consider proper transformations of the guidelines?"* are asked. To measure the extent whether the domain of the concept is clear and the measures fully represent the domain the question *"Do the objectives reflect the concepts security and privacy accurately and completely?"* is asked.

6.2 Result

Based on the validation interviews, interviewees generally agreed with the questions asked but had some remarks. For the question *"Would applying the guidelines improve security and privacy within the domain?"* The interviewees generally agreed. The guidelines are intended to guide the design of new IoT solutions or to review existing solutions. A proposed recommendation is that there should be a process in place to use the guidelines, in which adherence to the guidelines is be reviewed periodically. The guidelines should have a role within organization's managerial processes. What this role is has to be specified further by organizations that seek to implement the guidelines.

Even though interviewees generally agreed with the questions *"Are the guidelines proper transformations of the areas to consider? Are the areas to consider proper transformations of the guidelines?"*, one piece of criticism is that the guidelines were initially not

formulated in the same way. For example, the guidelines were sometimes formulated using “may be”, sometimes “must” “should”, implying that some guidelines were more important than others. Uniformity in wording is seen as important to take away ambiguity. This issue is solved by making small changes in the formulation of some guidelines to achieve uniform in the choice of words and formulation of the guidelines. “should”

“Do the objectives reflect the concepts security and privacy accurately and completely?” Even though the interviewees generally agreed with this question in terms of completion, they also noted that not all objectives are equally important in each situation. For example, in some situations, transparency may be essential to achieve user acceptance of IoT technology, whereas in other situations, integrity of data may be essential to ensure a correct diagnosis. Also, some objectives may be hard to combine at the same time (e.g., visibility and anonymity). The relative importance of each objective therefore has to be assessed before using the guidelines.

The main takeaway of the validation interviews is that the guidelines are seen as valid, but for practical application there has to be a process in place in which the relative importance of each objective is taking into account. It is suggested that organizations can do this by performing an initial assessment to discover the relative importance of each objective.

Chapter 7

Discussion

7.1 About this Chapter

The aim of this chapter is to answer the last sub-question: *"How may current trends influence authentication and authorization for the IoT in health care?"* and to conclude by answering the main research question *"How can authentication and authorization be managed in order to ensure security and privacy of the IoT in health care?"* Section 7.2 presents an interpretation of the findings of this work. Section 7.3 explains how current trends may affect authentication and authorization for the IoT in health care. Section 7.4 presents the main limitations of this work. Section 7.5 presents an overview of the conclusions of this work by providing an answer to all research questions. The work is finalized in Section 7.6 which presents opportunities for future research.

7.2 Interpretation of findings

The results of the interviews show that privacy, confidentiality, and integrity of data are the main objectives related to authentication and authorization. There seems to be a trade off between some of these security and privacy objectives, especially with auditability, visibility, and accountability on the one hand, and anonymity, unlinkability, and unobservability, and decentralization on the other. In health, due to a centralized orientation of many IoT solutions, the focus seems to be on auditability, visibility, and accountability to achieve security. In practice, potential negative effects for privacy are avoided through legal and organizational measures, such as medical confidentiality promises, confidentiality of EHRs, legislation, privacy impact assessments, data minimization, or avoidance of cloud- or other third party services.

The most important challenges to achieve privacy and security are heterogeneity, a lack of standardization, and problems related to managing (large amounts of) data. In the opinion of the interviewees under study, decentralizing authentication and authorization would not lead to increased privacy or security, as long as organizational or legal measures are taken. Centralized architectures are said to be easier to secure and monitor. Because of these reasons, decentralization is not included as a guideline or best practice for secure and privacy preserving authentication and authorization. The main takeaway of the validation interviews is that the guidelines that are presented in Chapter 5 are seen as valid, but for practical application there has to be a process in place in which the relative importance of each objective is taken into account. Organizations can do this by performing an initial assessment to discover the relative importance of each of these objectives.

7.3 Expectations for the future

7.3.1 IoT maturity, and expectations from the past

The IoT was originally envisioned as an ubiquitous network comprising everyday things that can sense, change, and process information about its environment. In this vision, devices would interact and share data which each other based on machine-to-machine communication. This state has however not been realized (Krasnova, 2017). Devices seem to be less smart than they were once envisioned, and generally only serve one goal (e.g., temperature monitoring). Similarly, an excess of devices, services, and connections does not seem to exist yet either in health care, where the focus is on centralized solutions. Another important difference is its heavy reliance on a single network (the internet) as alternative for direct machine-to-machine communication (Krasnova, 2017).

This explains and justifies a centralized approach, as health care organizations must have and want overview of all information in a system. Central control also simplifies authentication and authorization (Ouaddah et al., 2017; Wachter, 2018). Similarly, the negative impact of authentication and always-identity paradigms for privacy does not seem as large as suggested by some. If this situation may change in the future, for example in case of more machine-to-machine communication, alternative technological approaches such as group signatures, attribute / token based authentication will have more value in the IoT in health care. Changing authentication and authorization from always-identity to these alternatives may be hard due to path dependencies and embeddidness of authentication and authorization within an IoT solution, making it hard to adopt new architectures for authentication and authorization for the IoT. Therefore, a paradigm shift in design of IoT systems, not simply a change in protocols that are used for authentication and authorization may occur. Virtualization technologies such as SDN / NFV may help achieving this flexibility and allow for better monitoring and security of IoT networks. Virtualization may lead to exible provisioning, deployment and management of networks and decouples the control plane from the data plane, moving the control logic from the edge to a central controller (Kobo, Abu-Mahfouz, and Hancke, 2017).

The findings of this work suggest however, that it is still hard to implement such systems for home situations that can be controlled or monitored remotely. Even if these challenges can be overcome, and monitoring of the network can be simplified, this would however increase complexity by adding another layer. At the same time however, due to increasing visibility of users within networks, and the increase of the amount of gathered meta-data, virtualization may also have negative consequences for privacy. Pseudonymization is currently commonly applied to split between authentication and identity of user and device. It is questionable however if these may be enough to entirely preserve privacy of users in virtual networks. These potential threats should be taken into account when investing into these technologies, potentially in the form of Privacy Impact Assessments. At the same time however, due to increasing visibility of users within networks, virtualization may also have negative consequences for privacy. This should be taken into account when investing into these technologies, potentially in the form of Privacy Impact Assessments.

7.4 Limitations

This work has mostly focused on the effect of authentication and authorization on security and privacy related aspects relative to the application of the IoT in health care, and described these issues as main problems that limit further adoption of application of IoT technologies. There are also other challenges however that have to be overcome in order to reap the benefit that IoT technologies offer. These would require their own analysis. For example, one could ask themselves legal, organizational, or possibly even philosophical questions related to the application of IoT technologies in health care. For example, what is the role of "third parties" that transmit, process, or store personal health related data such as (cloud) service providers, internet carriers, cloud owners, internet providers and others related to the medical confidentiality? Or how can health care organizations effectively achieve data minimization, and set and monitor adequate storage and purpose limitations? Besides this, there are also other limitations which are explained in this section.

7.4.1 Security and capabilities of devices

Goal of this work is providing guidance on an architectural level on how authentication and authorization can help ensuring security and privacy of the IoT for health devices, and looked at authentication and authorization on an architectural level in the IoT for health care. Hardware, RFID tags, or the mathematics behind cryptographic algorithms are outside of the scope of this work, but may be important in providing security of information. For example, the OAuth protocol works with cryptographic keys which have to remain confidential to ensure that the device that holds the key is not compromised. However, the OAuth protocol does not specify how this key is protected. Security on a hardware level is therefore important, and the relevant key must be stored on the hardware in a way that the key remains confidential. Many (cheaper) devices however, are not capable of doing this. Even though security on a hardware level is outside of the scope of this work, it is important to provide secure authentication. Some other security the threats that are outside the scope of this work include Trojans, Side-channel attacks (D)DoS physical attacks, eavesdropping, side-channel attacks, fraudulent packet injection, routing attacks (Mosenia and Jha, 2017).

7.4.2 Health and IoT

It is likely that application of IoT technologies in health care will increase in the coming years. A white paper from a European consortium from 2016 predicts that advanced sensing, computing, and communicating technologies will enable personalized and preventive medicine. Continuous monitoring of relevant parameters throughout life to will allow for personalized health care and improvement in the way drugs are developed (Lehrach, Ionescu, and Benhabiles, 2016). It is very unlikely that the IoT in health care will not change significantly over time. As said previously, adoption, device's capabilities, and the place of IoT in health are expected to change. It seems reasonable to assume that parallel to these changes, challenges and requirements for security and privacy will also change over time. Similarly, for security, technological changes will change security needs. Devices that are sufficiently secure at time that they left the factory may become insecure before the end of their lifespan, which may involve recalls, additional maintenance, or firmware updates.

This work regarded health care as a sector and only looked at formal health care. Given the growth of health appliances for personal use however, the line between consumer appliances and health care appliances seems to blur. Apple watches for example, are already capable of capturing accurate ECGs (Abt, Bray, and Benson, 2018). As the boundaries between institutionalized health, home care, or fun seem to blur, the rise of health data may allow companies to create new business models, but at the same time the risk on privacy problems may increase. This may be because of health care data is no longer purely available by formal health care organizations, but also to a very small number of large technology companies who have a history of combining so much data that it leads to privacy issues and consumer concerns.

7.4.3 Change of privacy through digitalizalization

In this work, privacy is regarded as a more or less static and definable concept. In practice however, privacy may also mean different things for different people. People may not be aware, or have a lack of interest in their privacy. People's perception of privacy, and therefore the requirements for privacy may change over time and according to the problem of *value dynamism*, also through the use of technology. Value dynamism in the ethics of technology tells us that technologies often change the value frameworks we use to evaluate them (Kudina and Verbeek, 2018). As a result, a user's perceived privacy depends on how the technology is used, and to what extent the technology is useful for the user.

Over time, this may lead to a redefinition of people's perception of privacy, which is hard to predict in advance, as the application of technologies mediate this perception (Kudina and Verbeek, 2018). The application of IoT technologies may therefore change the perception of privacy even further, in ways that cannot be foreseen right now. Influencing these technological developments is easy when its applications are not yet manifest, yet, once the implications are known, they are difficult to change Kudina and Verbeek, 2018. As a result, limits of what we deem acceptable may blur as the adoption of these technologies increases.

Regardless of the perceived usefulness or emotional value that individuals place on these types of devices, potential threats to privacy should be evaluated. How this should happen and what value we place on privacy may change over time however. Legal frameworks however may be more static, even though is likely to follow up on these societal changes as well in the long term. Health care organizations have limited resources to ensure security of IoT devices and to preserve the privacy of users. So, organizations must make sure that they comply with the law, even though they must be aware that these legal frameworks may not be equal to the (ethical) expectations of users. For example, The European Group on Ethics noted that implantation of digital devices in human bodies should only be governed if the objective of the device is important, the implant is necessary to achieve this objective, and there is no less invasive or cost-effective method available that achieves the same objective (Kumar, 2007).

7.5 Conclusion

Chapter 1 of this document started with the main research questions, and five sub-questions. This section aims to provide a summarized answer to these sub-questions. Together, these questions answer the main research question *"How can authentication*

and authorization be managed in order to ensure security and privacy of the IoT in health care?".

1. What are the general security and privacy objectives that are applicable to authentication and authorization solutions? Security is commonly described as the interplay of confidentiality, integrity, and availability, potentially extended with additional constructs such as accountability, auditability, trustworthiness, or non-repudiation. Privacy issues may lead to different kinds of negative effects for individuals. Also from a legal point of view it is important to ensure privacy of personal data. Combining and processing seemingly innocent data may introduce privacy problems for IoT-users. Especially in the IoT, privacy problems may arise. Privacy can be grouped into several "objectives". In this work, privacy and security are seen as two different, but closely related, constructs. The OM-AM model can be used to describe the relationship from objective to mechanism. Authentication and authorization are essential to provide security and privacy. Insecure authentication however may lead to exposure of identifiers, and insecure authorization may lead to exposure of confidential information or stolen access tokens. Cryptographic measures may be taken to ensure security. To preserve the privacy of users and protect personal data in the IoT, several privacy preserving approaches can be used, such as data minimization, pseudonymization, group signatures, or attribute-based authentication. Objectives and requirements for privacy preserving authentication and authorization include transparency, user-driven, anonymity, pseudonymity, unlinkability, unobservability, decentralization, confidentiality, integrity, availability, usability, accountability, auditability, trustworthiness, and non-repuditation.

2. Which requirements do solutions for authentication and authorization have to fulfill in health care? In health care, organizations are vulnerable to cyber threats and lag behind in security. Health care data is an important target for cyber attacks. Increasing digitization of health records may also lead to privacy issues. The IoT in health care is not free from privacy and security risks and its users may be therefore be at risk. The adoption rate of IoT technologies in health care however, depends on the perceived security and privacy of these technologies. Interviews are carried out to define health care-specific requirements for security and privacy in this domain. The main challenges for achieving privacy and security are related to data management, heterogeneity, and a lack of standardization. Confidentiality, integrity, transparency, and unlinkability are the most important security and privacy related objectives for the IoT in health care. Based on literature, requirements are defined for each of the objectives. For each objective, important areas to consider are defined based on interviews.

3. What are the characteristics of currently available IoT authentication and authorization solutions? An overview of current literature on authentication and authorization for the IoT is presented, focusing on the “higher levels“ of authentication and authorization systems. To restrict or allow access to resources, a secure connection must be established, users and devices must be assigned with identities and granted access to resources, after which they must authenticate themselves before accessing resources. How this takes place in practice differs per implementation. An important difference is where the decisions for allowing or restricting access are made or enforced. Different frameworks, protocols and tools for authentication and authorization exist, five different high-level architectural designs for authentication and authorization are introduced, including centralized architectures that are based around RBAC, ABAC, XACML, Capability-based Authentication, which relies on an issuer of capability tokens, such as the OAuth protocol, Locally-centralized / globally-distributed, or decentralized architectures, where access control decisions are not made at a central

point in the network, but closer to the edge of the network.

This shows that there is not one single form in which authentication and authorization takes place, but rather a large number of different possibilities. This reflects the large number of different contexts in which different types of IoT devices operate. Every sector, from health care to transportation, to utilities and consumer electronics have their own IoT solutions. As Ouaddah et al. (2017) already noted, there is no single approach (always) better than another. All of them have advantages and disadvantages. Different design approaches seem complement each other, rather than compete with each other. Because of the diversity of these design approaches, it does not seem easy to present guidance to when to use a which design approach.

4. To what extent do current solutions for authentication and authorization meet the general security and privacy objectives and heath care requirements? Chapter 3 shows that the most important objectives for IoT devices in the health sector are related to privacy, confidentiality, and integrity of information. To achieve these goals, access control may play a key role. In general, centralized architectures and an always-identity approach are common and the potentially negative privacy effect of always-identity for authentication or centralized architectures is not experienced, as technical, legal, and organizational measures are taken to preserve privacy. Alternative forms of authentication and authorization are available, but not desired. Edge and fog technologies such as SND and NFV shift centralized computation to the edge of networks. Virtualizing networks offers the capability to have full control over networks, which potentially has a positive effect on security. However, as the oversight that these entities have increases, and data is gathered about information that may be linked to human entities of a platform. Finding a balance between all objectives may be hard. Based on literature and interviews, a set of guidelines is proposed to evaluate the extent to which security and privacy objectives are satisfied.

This chapter started with the fifth sub-question: *"How may current trends influence authentication and authorization for the IoT in health care?"* Currently, there seems to be a trade-off between some security and privacy objectives. In health there is an orientation centralized IoT solutions. Potential negative effects for privacy are avoided through legal and organizational measures. The IoT was envisioned as a network of devices connected with each other, vision has not been realized however (Krasnova, 2017). Devices are less smart than envisioned and only serve one goal. An abundance of devices, services, and connections does not seem to exist yet, and there is a heavy reliance on a single network (the internet). This explains and justifies a central approach, as health care organizations must and want overview of all information in a system. This central control also simplifies authentication and authorization. Similarly, the impact of authentication and always-identity paradigms is not as large as imagined. If this changes in the future, alternative approaches to authentication or authorization are worth considering. Virtualization technologies may be applied here, even though applying these may still be challenging in practice.

The main research question of this work is: *"How can authentication and authorization be managed in order to ensure security and privacy of the IoT in health care?"*. There are different reasons why security and privacy are important, creating several different requirements for providing security and privacy. In health care, due to the importance and the confidential nature of information the need for protecting information becomes even larger. The IoT may transform health care, but its differences with traditional IT systems pose new security and privacy threats. Authentication

and authorization are basic mechanisms to ensure the flow of information is controlled and ensure security and privacy for the IoT in health care. There are different architectures, models, and mechanisms for authentication and authorization but there is not one de-facto standard for the IoT in health care. In practice, authentication is often linked the identity of the natural user and the use of centralized IoT architectures is common. Organizations must take measures to guarantee security and privacy of information. In this work, several guidelines are proposed (and validated) that aim to manage authentication and authorization in a way that ensures security and privacy for the IoT in health care.

7.6 Future research

This work presented several guidelines for secure and privacy preserving authentication and authorization. At the moment however, these guidelines are focused on one moment in time, do not see security and privacy as dynamically changing concepts and can therefore only be used to provide guidance on one specific moment. In order for organizations to implement these guidelines more easily, it might be necessary to view the guidelines as step-by-step plan that organizations can use to implement these guidelines. Future research may focus for example on how to determine what the main security and privacy objectives are in a specific situation, how to implement the guidelines, and how to evaluate their effectiveness. Many of the respondents in the interviews prefer centralized, closed, IoT systems because of the ease of updating access policies, software, and monitoring. What would be interesting, is a system with decentralized PEP and PDP, which is dynamically updatable, monitorable and transparent to users of devices, and allows software updates, but at the same time privacy preserving. What such a system would look like is hard to say. Advances have been made with SDN / NFV-based authentication and authorization. Very little is known however about these IoT technology's impact on privacy of users. For blockchain-based solutions, access control polices are saved in a distributed ledger, privacy issues may arise as everybody may have an overview of a user's or device's rights. distributed, alternative methods for authentication that separate authentication from a natural user's personal identity such as group signatures, attribute-based authentication, or alternative authentication measures (non-identity) such as multi-factor authentication may make identity attributes less visible to transmitting parties. Future research may focus on these areas, to find a balance between auditability, visibility, and transparency on the one hand and user's privacy and unlinkability on the other.

Bibliography

Abt, Grant, James Bray, and Amanda Clare Benson (2018). "The validity and inter-device variability of the Apple Watch™ for measuring maximal heart rate". In: *Journal of sports sciences* 36.13, pp. 1447–1452.

AL-mawee, Wassnaa et al. (2015). "Privacy and security issues in IoT healthcare applications for the disabled users a survey". In:

Alaba, Fadele Ayotunde et al. (2017). "Internet of things security: A survey". In: *Journal of Network and Computer Applications* 88, pp. 10–28.

Alpár, Gergely et al. (2016). "New directions in IoT privacy using attribute-based authentication". In: *Proceedings of the ACM International Conference on Computing Frontiers*. ACM, pp. 461–466.

Alphand, Olivier et al. (2018). "IoTChain: A blockchain security architecture for the Internet of Things". In: *Wireless Communications and Networking Conference (WCNC), 2018 IEEE*. IEEE, pp. 1–6.

Atwady, Yahya and Mohammed Hammoudeh (2017). "A survey on authentication techniques for the internet of things". In: *Proceedings of the International Conference on Future Networks and Distributed Systems*. ACM, p. 8.

Atzori, Luigi, Antonio Iera, and Giacomo Morabito (2010). "The internet of things: A survey". In: *Computer networks* 54.15, pp. 2787–2805.

Baker, Stephanie B, Wei Xiang, and Ian Atkinson (2017). "Internet of Things for Smart Healthcare: Technologies, Challenges, and Opportunities". In: *IEEE Access* 5, pp. 26521–26544.

Bandara, Syafril et al. (2016). "Access control framework for API-enabled devices in smart buildings". In: *Communications (APCC), 2016 22nd Asia-Pacific Conference on*. IEEE, pp. 210–217.

Barka, Ezedine, Sujith Samuel Mathew, and Yacine Atif (2015). "Securing the web of things with role-based access control". In: *International Conference on Codes, Cryptology, and Information Security*. Springer, pp. 14–26.

Bäumer, Ulrich, Sabine von Oelffen, and Miriam Keil (2017). "Internet of Things: Legal Implications for Every Business". In: *The Palgrave Handbook of Managing Continuous Business Transformation*. Springer, pp. 435–458.

Bekara, Chakib (2014). "Security issues and challenges for the IoT-based smart grid". In: *Procedia Computer Science* 34, pp. 532–537.

Bennett, Jamie, Osvaldas Rokas, and Liming Chen (2017). "Healthcare in the Smart Home: A Study of Past, Present and Future". In: *Sustainability* 9.5, p. 840.

Bernabe, Jorge Bernal, Jose Luis Hernandez Ramos, and Antonio F Skarmeta Gomez (2016). "TACIoT: multidimensional trust-aware access control system for the Internet of Things". In: *Soft Computing* 20.5, pp. 1763–1779.

Bertino, Elisa (2016). "Data Security and Privacy in the IoT." In: *EDBT*. Vol. 2016, pp. 1–3.

Boeckl, Kaitlin et al. (2018). *Considerations for Managing Internet of Things (IoT) Cybersecurity and Privacy Risks*. Tech. rep. National Institute of Standards and Technology.

Bojanova, Irena and Jeffrey Voas (2017). "Trusting the internet of things". In: *IT Professional* 5, pp. 16–19.

Cavoukian, Ann (2011). "Privacy by design in law, policy and practice". In: *A white paper for regulators, decision-makers and policy-makers*.

Chen, Fulong et al. (2018). "An infrastructure framework for privacy protection of community medical internet of things". In: *World Wide Web* 21.1, pp. 33–57.

Cirani, Simone et al. (2015). "Iot-oas: An oauth-based authorization service architecture for secure services in iot scenarios". In: *IEEE sensors journal* 15.2, pp. 1224–1234.

CISCO (2014). *The Internet of Things Reference Model*.

Conti, Mauro et al. (2018). "Internet of Things security and forensics: Challenges and opportunities". In: *Future Generation Computer Systems* 78, pp. 544–546.

Cook, David A and Rose Hatala (2016). "Validation of educational assessments: a primer for simulation and beyond". In: *Advances in Simulation* 1.1, p. 31.

Di Pietro, Roberto et al. (2018). "A blockchain-based Trust System for the Internet of Things". In: *Proceedings of the 23nd ACM on Symposium on Access Control Models and Technologies*. ACM, pp. 77–83.

Dorri, Ali, Salil S Kanhere, and Raja Jurdak (2017). "Towards an optimized blockchain for IoT". In: *Proceedings of the Second International Conference on Internet-of-Things Design and Implementation*. ACM, pp. 173–178.

El-Aziz, AA Abd and A Kannan (2013). "A comprehensive presentation to XACML". In:

Fabiano, Nicola (2017). "The Internet of Things ecosystem: The blockchain and privacy issues. The challenge for a global privacy standard". In: *Internet of Things for the Global Community (IoTGC), 2017 International Conference on*. IEEE, pp. 1–7.

Fernández-Caramés, Tiago M and Paula Fraga-Lamas (2018). "A Review on the Use of Blockchain for the Internet of Things". In: *IEEE Access*.

Freeze, Di (2019). *2019 Cybersecurity Almanac: 100 Facts, Figures, Predictions and Statistics*. URL: `https://cybersecurityventures.com/cybersecurity-almanac-2019/`.

Fremantle, Paul, Jacek Kopeckỳ, and Benjamin Aziz (2015). "Web api management meets the internet of things". In: *International Semantic Web Conference*. Springer, pp. 367–375.

Fremantle, Paul et al. (2014). "Federated identity and access management for the internet of things". In: *Secure Internet of Things (SIoT), 2014 International Workshop on*. IEEE, pp. 10–17.

Gerdes, Stefanie, Olaf Bergmann, and Carsten Bormann (2014). "Delegated authenticated authorization for constrained environments". In: *Network Protocols (ICNP), 2014 IEEE 22nd International Conference on*. IEEE, pp. 654–659.

Giraldo, Jairo et al. (2017). "Security and privacy in cyber-physical systems: A survey of surveys". In: *IEEE Design & Test* 34.4, pp. 7–17.

Gusmeroli, Sergio, Salvatore Piccione, and Domenico Rotondi (2013). "A capability-based security approach to manage access control in the internet of things". In: *Mathematical and Computer Modelling* 58.5-6, pp. 1189–1205.

Hammi, Mohamed Tahar et al. (2018). "Bubbles of Trust: A decentralized blockchain-based authentication system for IoT". In: *Computers & Security* 78, pp. 126–142.

Hemdi, Marwah and Ralph Deters (2016). "Using REST based protocol to enable ABAC within IoT systems". In: *Information Technology, Electronics and Mobile Communication Conference (IEMCON), 2016 IEEE 7th Annual*. IEEE, pp. 1–7.

Hernández-Ramos, José L et al. (2016). "DCapBAC: embedding authorization logic into smart things through ECC optimizations". In: *International Journal of Computer Mathematics* 93.2, pp. 345–366.

Hevner, Alan R (2007). "A three cycle view of design science research". In: *Scandinavian journal of information systems* 19.2, p. 4.

Hu, Fei (2016). "Security and privacy in Internet of Things (IoTs): models, algorithms, and implementations". In:

Huang, Qinlong, Licheng Wang, and Yixian Yang (2018). "DECENT: Secure and fine-grained data access control with policy updating for constrained IoT devices". In: *World Wide Web* 21.1, pp. 151–167.

Hussein, Dina, Emmanuel Bertin, and Vincent Frey (2017). "Access control in IoT: From requirements to a candidate vision". In: *Innovations in Clouds, Internet and Networks (ICIN), 2017 20th Conference on*. IEEE, pp. 328–330.

Ingraham, Christopher (2018). *An insurance company wants you to hand over your Fitbit data so it can make more money. Should you?* URL: `https://www.washingtonpost.com/business/2018/09/25/an-insurance-company-wants-you-hand-over-your-fitbit-data-so-they-can-make-more-money-should-you/?noredirect=on&utm_term=.75b8349b8cbf`.

ISO (2000). *nformation processing systems — Open Systems Interconnection — Basic Reference Model — Part 2: Security Architecture*. Standard. Geneva, CH: International Organization for Standardization.

Kalamandeen, Andre et al. (2010). "Ensemble: cooperative proximity-based authentication". In: *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, pp. 331–344.

Kang, Minhee et al. (2018). "Recent Patient Health Monitoring Platforms Incorporating Internet of Things-Enabled Smart Devices". In: *International neurourology journal* 22.Suppl 2, S76.

Khader, Dalia (2007). "Attribute Based Group Signatures." In:

Kim, Hokeun and Edward A Lee (2017). "Authentication and Authorization for the Internet of Things". In: *IT Professional* 19.5, pp. 27–33.

Kim, Hokeun et al. (2016). "A secure network architecture for the internet of Things based on local authorization entities". In: *Future Internet of Things and Cloud (FiCloud), 2016 IEEE 4th International Conference on*. IEEE, pp. 114–122.

Kinkelin, Holger et al. (2018). "Trustworthy Configuration Management for Networked Devices using Distributed Ledgers". In: *arXiv preprint arXiv:1804.04798*.

Kobo, Hlabishi I, Adnan M Abu-Mahfouz, and Gerhard P Hancke (2017). "A Survey on Software-Defined Wireless Sensor Networks: Challenges and Design Requirements." In: *IEEE Access* 5.1, pp. 1872–1899.

Koop, C Everett et al. (2008). "Future delivery of health care: Cybercare". In: *IEEE Engineering in Medicine and Biology Magazine* 27.6.

Kouicem, Djamel Eddine, Abdelmadjid Bouabdallah, and Hicham Lakhlef (2018). "Internet of Things Security: a top-down survey". In: *Computer Networks*.

Krasnova, A (2017). "Smart invaders of private matters: Privacy of communication on the Internet and in the Internet of Things (IoT)". PhD thesis. Radboud University Nijmegen.

Kruse, Clemens Scott et al. (2017a). "Cybersecurity in healthcare: A systematic review of modern threats and trends". In: *Technology and Health Care* 25.1, pp. 1–10.

Kruse, Clemens Scott et al. (2017b). "Security techniques for the electronic health records". In: *Journal of medical systems* 41.8, p. 127.

Kudina, Olya and Peter-Paul Verbeek (2018). "Ethics from within: Google Glass, the Collingridge dilemma, and the mediated value of privacy". In: *Science, Technology, & Human Values*, p. 0162243918793711.

Kumar, Vikas (2007). "Implantable RFID Chips-Security versus Ethics." In: *FIDIS*, pp. 151–157.

Lee, Shih-Hsiung, Ko-Wei Huang, and Chu-Sing Yang (2017). "TBAS: Token-based authorization service architecture in Internet of things scenarios". In: *International Journal of Distributed Sensor Networks* 13.7, p. 1550147717718496.

Lehrach, H, A Ionescu, and N Benhabiles (2016). "The Future of Health Care: deep data, smart sensors, virtual patients and the Internet-of-Humans". In: *Future health manifesto*.

Loukil, Faiza et al. (2017). "Privacy-Aware in the IoT Applications: A Systematic Literature Review". In: *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems"*. Springer, pp. 552–569.

Mahalle, Parikshit N et al. (2013a). "A fuzzy approach to trust based access control in internet of things". In: *Wireless Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems (VITAE), 2013 3rd International Conference on*. IEEE, pp. 1–5.

Mahalle, Parikshit N et al. (2013b). "Identity authentication and capability based access control (iacac) for the internet of things". In: *Journal of Cyber Security and Mobility* 1.4, pp. 309–348.

Meiklejohn, Sarah et al. (2013). "A fistful of bitcoins: characterizing payments among men with no names". In: *Proceedings of the 2013 conference on Internet measurement conference*. ACM, pp. 127–140.

Mendez Mena, Diego, Ioannis Papapanagiotou, and Baijian Yang (2018). "Internet of things: Survey on security". In: *Information Security Journal: A Global Perspective* 27.3, pp. 162–182.

Miettinen, Markus et al. (2018). "Revisiting context-based authentication in IoT". In: *2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC)*. IEEE, pp. 1–6.

Moosavi, Sanaz Rahimi et al. (2015). "SEA: a secure and efficient authentication and authorization architecture for IoT-based healthcare using smart gateways". In: *Procedia Computer Science* 52, pp. 452–459.

Mosenia, Arsalan and Niraj K Jha (2017). "A comprehensive study of security of internet-of-things". In: *IEEE Transactions on Emerging Topics in Computing* 5.4, pp. 586–602.

Nakamoto, Satoshi (2008). "Bitcoin: A peer-to-peer electronic cash system". In:

Ouaddah, Aafaf, Anas Abou Elkalam, and Abdellah Ait Ouahman (2016). "FairAccess: a new Blockchain-based access control framework for the Internet of Things". In: *Security and Communication Networks* 9.18, pp. 5943–5964.

Ouaddah, Aafaf, Anas Abou Elkalam, and Abdellah Ait Ouahman (2017). "Harnessing the power of blockchain technology to solve IoT security & privacy issues". In: *Proceedings of the Second International Conference on Internet of things and Cloud Computing*. ACM, p. 7.

Ouaddah, Aafaf et al. (2015). "Security analysis and proposal of new access control model in the Internet of Thing". In: *Electrical and Information Technologies (ICEIT), 2015 International Conference on*. IEEE, pp. 30–35.

Ouaddah, Aafaf et al. (2017). "Access control in the Internet of Things: Big challenges and new opportunities". In: *Computer Networks* 112, pp. 237–262.

O'Connor, Yvonne et al. (2017). "Privacy by Design: Informed Consent and Internet of Things for Smart Health". In: *Procedia Computer Science* 113, pp. 653–658.

Peppet, Scott R (2014). "Regulating the internet of things: first steps toward managing discrimination, privacy, security and consent". In: *Tex. L. Rev.* 93, p. 85.

Pinno, Otto Julio Ahlert, Andre Ricardo Abed Gregio, and Luis CE De Bona (2017). "ControlChain: Blockchain as a Central Enabler for Access Control Authorizations in the IoT". In: *GLOBECOM 2017-2017 IEEE Global Communications Conference*. IEEE, pp. 1–6.

Ponemon Insfitute, LLC (2016). *Sixth annual benchmark study on privacy & security of healthcare data*.

Porambage, Pawani et al. (2016). "The quest for privacy in the internet of things". In: *IEEE Cloud Computing* 2, pp. 36–45.

Qu, Sandy Q and John Dumay (2011). "The qualitative research interview". In: *Qualitative research in accounting & management* 8.3, pp. 238–264.

Rahmani, Amir M et al. (2018). "Exploiting smart e-Health gateways at the edge of healthcare Internet-of-Things: A fog computing approach". In: *Future Generation Computer Systems* 78, pp. 641–658.

Roman, Rodrigo, Jianying Zhou, and Javier Lopez (2013). "On the features and challenges of security and privacy in distributed internet of things". In: *Computer Networks* 57.10, pp. 2266–2279.

Rose, Karen, Scott Eldridge, and Lyman Chapin (2015). "The internet of things: An overview". In:

Sandhu, Ravi (2000). "Engineering authority and trust in cyberspace: The OM-AM and RBAC way". In: *Proceedings of the fifth ACM workshop on Role-based access control*. ACM, pp. 111–119.

Scarpato, Noemi et al. (2017). "E-health-IoT Universe: A Review". In: *International Journal on Advanced Science, Engineering and Information Technology* 7.6, pp. 2328–2336.

Seitz, L et al. (2016). "Authorization for the Internet of Things using OAuth 2.0". In: *Internet Engineering Task Force (IETF): Fremont, CA, USA*.

Seitz, Ludwig, Göran Selander, and Christian Gehrmann (2013). "Authorization framework for the internet-of-things". In: *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2013 IEEE 14th International Symposium and Workshops on a*. IEEE, pp. 1–6.

Shafagh, Hossein et al. (2017). "Towards blockchain-based auditable storage and sharing of iot data". In: *Proceedings of the 2017 on Cloud Computing Security Workshop*. ACM, pp. 45–50.

Sicari, S et al. (2017). "A policy enforcement framework for Internet of Things applications in the smart health". In: *Smart Health* 3, pp. 39–74.

Solapurkar, Prajakta (2016). "Building secure healthcare services using OAuth 2.0 and JSON web token in IOT cloud scenario". In: *Contemporary Computing and Informatics (IC3I), 2016 2nd International Conference on*. IEEE, pp. 99–104.

Sweeney, Latanya (2002). "k-anonymity: A model for protecting privacy". In: *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10.05, pp. 557–570.

Tasali, Qais, Chandan Chowdhury, and Eugene Y Vasserman (2017). "A flexible authorization architecture for systems of interoperable medical devices". In: *Proceedings of the 22nd ACM on Symposium on Access Control Models and Technologies*. ACM, pp. 9–20.

Telefonaktiebolaget LM Ericsson, Ericsson (2018). *Internet of Things forecast – Ericsson*. URL: `https://www.ericsson.com/en/mobility-report/internet-of-things-forecast`.

The Open Group (2000). *AUTHORIZATION (AZN) API*. URL: https://publications.opengroup.org/c908.

Trnka, Michal, Tomas Cerny, and Nathaniel Stickney (2018). "Survey of Authentication and Authorization for the Internet of Things". In: *Security and Communication Networks* 2018.

Tschofenig, H et al. (2015). *Architectural considerations in smart object networking*. Tech. rep.

Wachter, Sandra (2018). "Normative challenges of identification in the Internet of Things: Privacy, profiling, discrimination, and the GDPR". In: *Computer law & security review* 34.3, pp. 436–449.

Wieringa, Roel J (2014). *Design science methodology for information systems and software engineering*. Springer.

Wohlin, Claes (2014). "Guidelines for snowballing in systematic literature studies and a replication in software engineering". In: *Proceedings of the 18th international conference on evaluation and assessment in software engineering*. Citeseer, p. 38.

Wohlin, Claes, Martin Höst, and Kennet Henningsson (2003). "Empirical research methods in software engineering". In: *Empirical methods and studies in software engineering*. Springer, pp. 7–23.

Wu, Longfei et al. (2018). "An out-of-band authentication scheme for internet of things using blockchain technology". In: *2018 International Conference on Computing, Networking and Communications (ICNC)*. IEEE, pp. 769–773.

Xu, Quanqing et al. (2018). "Blockchain-based decentralized content trust for docker images". In: *Multimedia Tools and Applications* 77.14, pp. 18223–18248.

Yang, Huihui (2016). "Cryptographic enforcement of attribute-based authentication". In:

Ye, Ning et al. (2014). "An efficient authentication and access control scheme for perception layer of internet of things". In: *Applied Mathematics & Information Sciences* 8.4, p. 1617.

Zanella, Andrea et al. (2014). "Internet of things for smart cities". In: *IEEE Internet of Things journal* 1.1, pp. 22–32.

Zarca, Alejandro Molina et al. (2019). "Enabling Virtual AAA Management in SDN-Based IoT Networks". In: *Sensors* 19.2, p. 295.

Zhang, Guoping and Jiazheng Tian (2010). "An extended role based access control model for the Internet of Things". In: *Information Networking and Automation (ICINA), 2010 International Conference on*. Vol. 1. IEEE, pp. V1–319.

Zhou, Jun et al. (2017). "Security and privacy for cloud-based IoT: challenges". In: *IEEE Communications Magazine* 55.1, pp. 26–33.

Zyskind, Guy, Oz Nathan, et al. (2015). "Decentralizing privacy: Using blockchain to protect personal data". In: *Security and Privacy Workshops (SPW), 2015 IEEE*. IEEE, pp. 180–184.