CHALMERS
UNIVERSITY OF TECHNOLOGY

UNIVERSITY OF GOTHENBURG

Architecture Framework for Blockchain Implementation

A Design Science Study

Master's thesis in Computer science and engineering

Felix Lissåker
Johan Sjöberg

Department of Computer Science and Engineering
Chalmers University of Technology
University of Gothenburg
Gothenburg, Sweden 2019

Master's Thesis 2019

Architecture Framework for Blockchain Implementation

A Design Science Study

Felix Lissåker
Johan Sjöberg

UNIVERSITY OF GOTHENBURG

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Computer Science and Engineering
Chalmers University of Technology
University of Gothenburg
Gothenburg, Sweden 2019

Architecture Framework for Blockchain Implementation
A Design Science Study

Felix Lissåker
Johan Sjöberg

Supervisor: Eric Knauss, Department of Computer Science and Engineering
Advisor: Ulf Liljensten, Biswise
Examiner: Riccardo Scandariato, Department of Computer Science and Engineering

Master's Thesis 2019
Department of Computer Science and Engineering
Chalmers University of Technology and University of Gothenburg
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Typset in LaTeX
Gothenburg, Sweden 2019

Architecture Framework for Blockchain Implementation
A Design Science Study

Felix Lissåker, Johan Sjöberg
Department of Computer Science and Engineering
Chalmers University of Technology and University of Gothenburg

Abstract

Background: The release of Nakamoto's whitepaper, describing the inner workings of Bitcoin, triggered a vast interest in blockchain's core. In resent years, an elevated technological fascination is evident in the large number of new blockchain platforms, tokens and technological solutions aiming to solve the technology's shortcomings. An enabler of high security, transparency, and auditability, blockchain has been envisioned as the missing piece of the foreseen IoT breakthrough. Industries outside cryptocurrencies are starting to show an elevated interest in the technology.

Objective: To facilitate future blockchain implementations, an architecture framework is derived with appropriate viewpoints, models and corresponding guidelines that software architects are recommended to consider.

Methods: The framework is empirically constructed following a design science research methodology over three cycles. Challenging scenarios and guidelines are gathered from a workshop with five case company participants, 14 interviews, and a structured literature review including 229 primary studies of which 17 are thoroughly perused. Two industry evaluation sessions with case company representatives assess the framework. An interviewee survey sent out to study participants and a SAPSA conference talk with subsequent feedback from impartial industry representatives further evaluate the work.

Results: Beyond well known viewpoints traditionally used in architecture frameworks, five new viewpoints can be derived from the findings. The *applicability* viewpoint addresses the pronounced difficulty of finding an appropriate use case and using blockchain properly. The *ecosystem* viewpoint addresses the current scarcity of blockchain expertise and the means by which developers can be supported in their smart contract development. The *infrastructure* viewpoint highlights how the system's topology and platform parameters affect its performance, scalability and transaction throughput. The *legal* viewpoint helps architects taking informed decisions in the light of current legal developments with an accent on network liability and private data management. Lastly, the *end user* viewpoint attend to the identified prioritization of technological aspects over customer value.

Conclusions: The framework need further refinement, e.g. through the addition of extra viewpoints from other frameworks. Nevertheless, the ambition to establish initial implementation practices is satisfied. The viewpoints and their respective guidelines are anticipated to be of industrial value, based on evaluations provided by the case company and other industry representatives.

Acknowledgements

This thesis has been performed within the Department of Computer Science and Engineering, for the division of Software Engineering, at Chalmers University of Technology. The work has involved close correspondence with industry representatives who have significantly conduced to the result.

We would like to express our sincere appreciation to our industry supervisor Ulf Liljensten for his constructive engagement throughout the process of our research. Furthermore, many thanks to Kristoffer Syversen for generously sharing his time and knowledge in blockchain systems and the specific use case. We would also like to give a special thanks to our academic supervisor Eric Knauss for his unabating support and constructive feedback.

Finally, a last thank you to all industry participants who have shared their thoughts and expertise for the benefit of our conclusions.

Felix Lissåker and Johan Sjöberg, Gothenburg, June 2019

Contents

1 Introduction **2**

1.1 Statement of the problem . . . 2
1.2 Purpose of the study . . . 3
1.3 Case company . . . 4
1.4 Research questions . . . 4

2 Background **5**

2.1 Blockchain fundamentals . . . 5
2.2 Consensus algorithms . . . 8
2.3 An outlook on blockchain and IoT . . . 9
2.4 Blockchain applications and challenges . . . 11

3 Methodology **13**

3.1 Design science research . . . 13
3.1.1 Problem type . . . 13
3.1.2 Artifact . . . 14
3.1.3 Design science research components . . . 15
3.2 Literature review . . . 18
3.3 Interviews and workshop . . . 18

4 Findings **21**

4.1 Research question 1 . . . 21
4.1.1 Stakeholders . . . 21
4.1.2 Concerns . . . 23
4.1.3 Challenging scenarios . . . 26
4.2 Research question 2 . . . 34
4.2.1 Applicability viewpoint . . . 35
4.2.2 Ecosystem viewpoint . . . 42
4.2.3 Infrastructure viewpoint . . . 45
4.2.4 Legal viewpoint . . . 51
4.2.5 End user viewpoint . . . 60
4.3 Research question 3 . . . 63
4.3.1 Workshop evaluation . . . 64
4.3.2 Survey evaluation . . . 66
4.3.3 The SAPSA conference evaluation . . . 67

5 Artifact **68**

5.1 Architecture framework 68
5.2 Industry relevance 74

6 Discussion **75**
6.1 Implications to research 76
6.2 Validity concerns 77

7 Conclusion and Outlook **79**

A Interview Template **91**

B Survey Template **92**

1. Introduction

Blockchain has the potential to considerably change the world economy by providing opportunities to enable decentralized, traceable, and secure ways of managing sensitive data (Tapscott, Tapscott, & Kirkland, 2016). The technology is not limited to financial transactions but available to all forms of interactions and data exchanges (Conoscenti, Vetrò, & Martin, 2016). It holds the promise to reduce the inevitable trust gap in most exchanges by eliminating misgivings about data validity. In the absence of such a solution, current interactions are contingent upon reliable intermediaries such as financial institutions and governments alike (Rodrigues, Bocek, & Stiller, 2018).

The infancy of blockchain is indubitable, given the rather narrow range of applications and their individual success (Halaburda, 2018). Many hurdles, such as performance implications and legal obstacles need to be solved before blockchain usage can be fully motivated (Kumar & Mallick, 2018). Unlike the Internet, which has a sophisticated governance ecosystem, the current discipline of blockchain applications is still in need of more robust governance frameworks. There are multiple questions and important aspects of blockchain technologies that require further research, including performance, security, and applicability (Reyna, Martín, Chen, Soler, & Díaz, 2018).

Blockchain's potential goes hand in hand with the evolution of IoT (Internet of Things). These devices are anticipated to grow from 11BN units in 2018 to 125BN by 2030 (Mittal, Tam, & Ko, 2018). This surge will demand a reliable infrastructure for vast amounts of private user data. The increasing adoption of IoT devices and the large amount of accessible data will change the way individuals and companies manage their businesses and conduct transactions. Privacy, tractability, and data validity are but a few of many aspects needing consideration. Blockchain's aspiration to facilitate transparency and trust is therefore a strong candidate for such use cases (Tapscott et al., 2016).

1.1 Statement of the problem

Dieterich et al. (2017) highlights several application areas that could potentially benefit from applying some sort of blockchain technology. Blockchain is not constrained to manage intangible assets but can also be used in the trading of tangible goods. Blockchain's provenance and most well known application, cryptocurrencies, has paradoxically compromised the reputation of the technology. Many are, however, still positive about its future. Mittal et al. (2018) predicts a rapid increase of IoT applications suggesting global adoption by 2030, realizing 125BN devices. As a result, data security will become even more prominent. It is suggested by both Dieterich et al. (2017) and Mittal et al. (2018) that blockchain has the potential to safely store user data and mitigate the trust discrepancies between involved parties. Different use cases will require divergent properties and therefore request different blockchain technologies. Software engineering guidelines and frameworks are needed to facilitate this potential and to make the technology more accessible.

As reported by Halaburda (2018), blockchain is still a novel technology that arouses confusion. Its integral components – composed of smart contracts, encryption, and a decentralized ledger – are not strictly contingent upon each other but can be implemented and used independently. The unique composition of blockchain requires an equally unique set of use case requirements to be valid. Nevertheless, as application areas supporting IoT rapidly grow, the potential utilization of blockchain follow suite. Blockchain's prospects within these areas are considered high but the lack of sufficient standards hamper future progress. Generally, numerous ideas regarding blockchain and its merge with domains outside the cryptocurrency sphere remain largely hypothetical and suffer from formal structures. The architectural considerations are many and necessitate careful investigations. The results of these investigations need to be compiled and formalized in a framework to facilitate future blockchain implementations. The following problem statement has therefore been formulated:

> ***Current blockchain research has to be complemented with proper guidelines, aiding decision making with respect to implementation strategies.***

These guidelines will be derived with influence from a specific use case. The use case is obtained from GreenStar Marine's ambition to develop a blockchain platform for battery verification. In the initial steps of their implementation efforts, the architecture framework is aspired to provide directions in their decision-making.

1.2 Purpose of the study

The overall purpose of this research goes beyond aiding the specific case company. It aspires to compile and structure the knowledge base surrounding blockchain to support other blockchain implementations as well. Hence, the purpose of this study is as follows:

> ***The purpose of this study is to derive an architecture framework for blockchain implementation that provides useful guidelines and support for decision-making.***

This study aspires to discern the major challenges of blockchain present in literature and among industry practitioners. The architecture framework is later established around a set of viewpoints derived from these challenges. To facilitate blockchain implementation efforts, each viewpoint will be complemented with model suggestions and a set of architectural guidelines addressing the identified challenges. These tools will assist system architect in modelling system specific views, following recommended practices.

The framework is based on a structured literature review of 100 primary studies obtained by searching two academic databases using the keywords *blockchain*, *IoT*, and *challenges*. This investigation is complemented by an interview review involving 14 industry participants so to supplement the academic findings with practical aspects present in industry. Except for two researchers, the interviewees are working professionals with varying blockchain focus.

Researchers in the field of blockchain and practitioners implementing blockchain applications are anticipated to benefit from this work. Use cases other than battery tracking and secondary marketplaces will benefit by utilizing the architecture framework derived from this work. The particular use case have characteristics shared by many IoT networks, including multiple participants, privacy concerns, performance challenges and ample amounts of accessible data, making the findings more generalizable.

1.3 Case company

The thesis is carried out in collaboration with GreenStar Marine and Biswise, two companies situated at Lindholmen, Göteborg. GreenStar Marine constructs electric propulsion systems. They are currently developing a tracking software for battery systems and consulted by Biswise. More specifically, a secondary marketplace for batteries are under development. Secondary marketplaces shares the problem of mistrust between buyers and sellers. GreenStar Marine has been developing a centralized tracking platform, but their current ambition is to upgrade it with blockchain support. This would enable more accurate and verifiable battery assessments, encouraging secure and efficient trading. This makes marketplaces for secondary items an interesting field of study with respect to blockchain. Furthermore, the use case showcases the potential use of the technology for environmental purposes.

1.4 Research questions

The study objective is divided into three research questions that revolve around three aspects: challenges of blockchain implementations (RQ1), architectural guidelines (RQ2), and framework evaluation (RQ3). These questions are aligned with the elements of the standard architecture framework definition (see ISO/IEC/IEEE (2011)) and follow the iterative approach of the DSR methodology.

RQ1: *Based on stakeholders and their concerns, which challenging scenarios are important to consider when assessing the use of blockchain?*

RQ2: *What viewpoints are necessary to consider in blockchain implementations and how can different architectural principles address their respective challenges?*

RQ3: *By evaluating the framework derived from the preceding two questions, to what extent does the framework aid practitioners when reasoning about blockchain?*

2. Background

The academic literature surrounding the purpose of this study is scarce and this study aims to fill part of this void; in particular by consolidating information from resent research and interviews. The results will be grouped and organized into an architecture framework for blockchain implementation.

Given the novelty of blockchain and the ambiguity surrounding the general concept, this chapter is dedicated to elaborate on some of its important fundamentals. Understanding these concepts are prerequisites for following the reasoning in subsequent sections. Furthermore, the research underpinning this chapter has facilitated the creation of the architecture framework. Beyond the technical fundamentals of blockchain, an outlook for IoT is given to elaborate on the entwinement of these two technologies. The future prospects of IoT is often linked to the development of blockchain. The chapter concludes with related work on blockchain challenges.

2.1 Blockchain fundamentals

In early 2009, Bitcoin brought blockchain to the world. Though the initial focus has been on developing and defining new cryptocurrencies, in later years, the interest has shifted to the underlying technology: blockchain. At large, blockchain enables support for autonomous networks that leverage on a distributed architecture. A blockchain can be described as a digital ledger that is shared and synchronously updated between all participating nodes delimiting the risk of central data ownership. The technology was developed with security in mind and designed to support immutability, transparency, and operational robustness (Panarello, Tapas, Merlino, Longo, & Puliafito, 2018). The topics of security and privacy is still highly relevant to discuss as the technology is still in its development phase. The exclusion of a central authority transfers the trust to the infrastructure of the network itself that must work appropriately so to not erode its validity.

Different blockchain technologies have been developed in the wake of Bitcoin to address some of its problems. All technologies are, however, constructed around a set of shared features. Panarello et al. (2018) describes four of the most prominent features: (1) each technology institutes trust between network participants that leverage on the distributed network, removing the need for a trusted central authority; (2) at its core, the blockchain provides a distributed ledger that is maintained by all network participants; (3) the ledger is synchronized between all nodes based on protocols ensuring transaction distribution, validation and inclusions; (4) network agreement is reached through a consensus protocol that warrant consistent ledger updates, despite network delays and dissonance.

(Reyna et al., 2018) describe blockchains as chained blocks containing validated transactions. The chain links the blocks together by ensuring that the present block stores the hash of the proceeding one, creating an immutable storage system. As new transactions are generated, the network needs to agree on the transaction ordering considering network delays and local latency. Each node responsible

for updating the ledger generates and proposes a new block containing a set of transactions. The network then assigns the task of updating the ledger to a specific node following a consensus algorithm that differs in terms of security and performance.

Many of today's blockchain platforms have extended their functionality beyond simple transactions to support logical procedures. Logic is added onto the blockchain by the inclusion of smart contracts. Christidis and Devetsikiotis (2016) describe smart contracts as program based routines that reside on the blockchain, each with its own address. A transaction sent to the address of a smart contract results in the execution of the contract on each network node. This requires a virtual machine to be present. The Ethereum blockchain, for example, can be described as a distributed virtual machine on top of which developers can add their defined smart contracts.

Some blockchains liberates centralized databases from inefficiency, high costs, and vulnerability. Instead it promises to introduce *consensus*, *provenance*, *immutability*, and *finality* (Cole & Gorman, 2019). These aspects are defined in the context of blockchain as follows:

- **Consensus:** Determining how to agree on a certain state.
- **Provenance:** Assures historical tracking of data.
- **Immutability:** Makes data impracticable to tamper with.
- **Finality:** Makes information doubtlessly truthful.

Consensus among distributed ledgers are essential to determine the data stored on a blockchain. Different platforms host different consensus algorithms which poise the benefits and drawbacks of the chain. Consensus algorithms will be covered in more detail in a subsequent section in this chapter. *Provenance* is a general feat of blockchain, common to most platforms, but does still vary depending on solution. Certain implementation strategies, where parts of the blockchain functionality is used, provenance can be de-prioritized in favor for storage capacity or some other stressed feature. Established blockchains such as the Bitcoin, Ethereum, Hyperledger all store historical transaction data indefinitely. *Immutability* and *finality* are both the original drivers of blockchains. That is, they are databases that are impracticable to tamper with and ensures the verified transactions. These claims are, as previously mentioned, bred into the fabrics of blockchain. However, events proving the contrary – hacked cryptocurrency exchanges and loss of valuable data – have surfaced in the early phases of the technology.

The following sections will augment the notion of blockchain by presenting the most anticipated blockchain platforms found by this study. These platforms are at their respective ends of the blockchain spectrum in terms of openness and privacy, which will be elaborated in the subsequent sections. Furthermore, the public-permissioned blockchain is a combination of the two which has proliferated in recent years and are described below.

Public-permissionless blockchains: Ethereum and Bitcoin

Most centralized data bases can create, read, update and delete data. The fundamental rules of blockchain dictate that data only can be created and read (Richardson, 2017). Ethereum is the

pioneer of *public* and *permissionless* chains together with the Bitcoin blockchain (Mohanty, 2018). This means that everyone can read and write to the chains with the same permissions and ability to join (Christidis & Devetsikiotis, 2016). This information is limited to transaction information, the identity of each party is still anonymous. The openness of these chains allows anyone to join and leads to much larger networks. The keepers of ledgers exchange their storage capacity and computational power with token remittance. Resulting trust differences on an open network require sophisticated consensus algorithms. Regarding Ethereum and the Bitcoin blockchain, the current use of the Proof-of-Work (PoW) consensus protocol satisfy security requirements but is vastly expensive in terms of computational power, energy requirements, and update frequency (Reyna et al., 2018).

Different methods to compromise blockchains have surfaced. The most comprehensible attack requires 51% of the entire computational power of the network, which is becoming ever more impracticable given its increasing size (Fernández-Carmés & Fraga-Lamas, 2018). As the number of transactions grow, the computational power required to create a block (the hashing power) increases. As it does, less conventional nodes will be able to participate on the chain incentivizing initiatives such as mining pools and computational farms. This trend centralizes computational power and could again increase the chances of 51% attacks (Zheng, Xie, Dai, Chen, & Wang, 2017). In summary, this sophisticated consensus algorithm tries to ensure security on blockchains but is threatened by centralization of computational power Reyna et al. (2018).

Private-permissioned blockchains: Hyperledger

Hyperledger being *private/closed* and *permissioned* is on the opposite side of the spectrum (Androulaki et al., 2018). Users outside the network are incapable of reading or writing to the blockchain. All peers that want to join the network need to be permitted by a central governance system before doing so. This facilitate a much smaller, more trustworthy community where transactions and participants are transparent and entrusted to each other. The individual dedication also diminish the issue of storage capacity and computational power since the participants are more dedicated to the cause of the chain. Ambitious consensus algorithms are much less needed which rules out PoW. Instead selective endorsement is applied where authorities decided upon validation, allowing a subset of nodes confirm transactions (*Hyperledger Fabric*, 2019). However, the central governance system drastically undermine the core essence of a blockchain. In these cases, a single point of authority distributes permissions to participants in the network, which is the fundamental flaw that blockchain initially sought to resolve (Reyna et al., 2018)

Public-permissioned blockchains

As described by Meng, Tischhauser, Wang, Wang, and Han (2018) public-permissioned blockchains share characteristics from both its private and public counterparts. Although these blockchains have a central authority exerting some network control, everyone is still allowed to read the blockchain state. It is the *writing access* that is regulated. The rational behind leaving *reading access* completely open is to encourage miners to take part in the network. Allowing everyone to partake in the verification process incorporates the same strength in numbers as on public blockchains.

2.2 Consensus algorithms

Consensus protocols manages a common, unambiguous ordering of transactions on the blockchain. Despite geographical spread of participating nodes, the protocol ensures integrity and consistency (Baliga, 2017). It does so, however, in very different ways that introduces different challenges. It is, therefore, important to consider the consensus algorithm in blockchain implementations which also makes it relevant for the architecture framework.

The original consensus mechanism was based on the sophisticated Proof-of-Work (PoW) algorithm. Drawbacks of this algorithm, in combination with an increasing interest around blockchain, has produced numerous other algorithms, fit for different applications. Instead of suffocating the disadvantages of early consensus algorithms one can chose blockchain platform based on its associated consensus algorithm to fit particular use cases. The following sections will elaborate the basics of the two most popular consensus algorithms found by this study.

Proof of Work

The PoW algorithm is essential to every early version of blockchain, including Ethereum and the Bitcoin blockchain. Actions on the chain requires processing of the algorithm on all participating nodes on the network, thus, requiring vast computational power (Panarello et al., 2018). A transaction T_1 constitute a challenge (c) distributed to all mining nodes on the network. All miners aim their computational power to find a particular ten digit number or proof (p) to c. Both c and varying attempts of p are run through the hash function, SHA265. This function takes a file or string and produces a strict, unambiguous, 265 digit binary number. The first miner that finds a p that produces a SHA265 output with an array of initial consecutive zeros wins the challenge and earns a reward. The number of zeroes required determines the difficulty to solve c and are proportional to the number of historical transactions stored on the chain. This development ensures that computational power to solve a certain c increases as the number of historical transactions increases. More computational power is, thus, needed today compared to an early transaction on, for example, the Bitcoin blockchain. Once the correct p is found the winner distributes p to all other participating nodes. Since all other nodes competed in the quest to find p solving c, they all have access to c and can therefore test the distributed p by running both p and c through the hash function to confirm T_1. Since p is already provided, only one run through the hash function is needed to verify the truthfulness of p which requires very little resources. The demanding part of the algorithm is to randomly test p until it solves c. This does in essence make it difficult to provide truth but very easy to validate a transaction, which is exactly what is required of a consensus algorithm.

The requirement on miners to store all historical transaction data and to participate in all computational efforts is problematic only in relation to the size of the network (Zheng, Xie, Dai, Chen, & Wang, 2018). In the beginning, when transactions and participants were few, this brilliant algorithm made total sense. Now, however, one c needs to be processed by thousands of sophisticated computers whereas only one p from one computer gets distributed as truth. This results in an unnecessary waste of power and resources. Furthermore, hashing difficulty is set to increase as a result of widespread

adoption of blockchain which results in greater computational requirements and extended difficulty to defend the PoW algorithm's excessive use of resources.

Proof of Stake

PoS utilizes the brilliance of PoW but reduces the computational cost by letting only one miner find p (Lisk, 2019). The entrusted miner is chosen randomly and is constantly changing which makes it impracticable for outsiders to pinpoint the miner and compromise its process. However, the miner itself can tamper with both c and p which incentivizes a "stake". This stake is a deposit that each miner has to contribute with in order to be considered in the mining process. The chance of getting "randomly" picked is proportional to the size of the stake, meaning that a larger deposit increases the chance of reward. If the miner chooses to break the trust by meddling with c or p it will loose its deposit. Since the deposit needs to be larger than the block reward little incentive lies in tampering with the process.

By using this consensus mechanism powerful computational assets are no longer preferable since no advantage is given to quick resolution of p (Gaži, Kiayias, & Russell, 2018). This brings back the decentralized nature of blockchain, at least to some extent. Capital intensive peers can generate much larger stakes than smaller actors which again creates unbalance. This unbalance is less severe than for PoW, however, since computational power can be bought in bulk, utilizing economies of scale. In the case of Bitcoin a computational majority is held by some large mining pools at times. This centralized power could be used with malicious intent. PoS can be used advantageously in similar manners by owning a majority of the total stake on the network. In the case of Bitcoin, however, that would mean an incomprehensible monetary effort impracticable compared to owning a 51% share of all computational power since it could be accessible through crowd sourcing. In contrast to PoW where a faster computer always, on average, solves a particular issue faster a node with lower stakes can still be entrusted to do the computational work. This inclusion of all participants on the network is a promising aspect of the PoS algorithm because it facilitates the fundamental goal of blockchain: decentralization.

2.3 An outlook on blockchain and IoT

The demand for IoT enabled devices is set to surge during the 21th century. By 2030 IoT enabled devices are estimated to reach 125BN units, half of which are customer related (Mittal et al., 2018). As the line between device and user blurs more connected devices will result in more sensitive data. Given 2018's many conundrums on data privacy (e.g. Cadwalladr and Graham-Harrison (2018) on Facebook's Cambridge Analytica scandal) measures will be taken to find new ways of storing data more securely. Blockchain's promise to decentralize and verify data storage could therefore be a substantial driver of the IoT success. In this particularly promising application of blockchain, numerous benefits are identified by Mittal et al. (2018):

- Sensor data can be tracked to prevent duplication of malicious data.
- Each ledger can be provided with unique identification, authentication, and seamless data

transfer.

- Direct communication between devices can delimit the need for third parties, thus reducing costs.
- Historical records of connected devices can be kept for troubleshooting.

In general, blockchain-based IoT solutions can be used to emphasize the growing issues of data privacy as the use of IoT devices increases. All in all, Mittal et al. (2018) concludes that the rise of IoT is inevitable; its steepest period of growth is assumed to be in the upcoming decades. Consequently, more data of sensitive nature will be generated which will further fuel the trust issues between the parties in a transaction. The increasing adoption of IoT within many industries will make the consideration and development of blockchain solutions relevant when establishing new corporations.

According to Sagirlar, Carminati, Ferrari, Sheehan, and Ragnoli (2018), current IoT platforms are managed through centralized Cloud infrastructures. This contravene the IoT networks' intrinsically distributed nature. Moreover, the centralized architecture imposes threats in terms of high maintenance costs, insufficient performance, and security breaches. The researchers identify blockchain as a strong candidate for overcoming such disadvantages, while supporting distributed consensus and verification of data. However, the architectural design decisions are challenging and will require further research. As an example, Danzi, Kalør, and Popovski (2018) acknowledge the limited capacity of IoT devices. Performance and memory constrains make them ill-suited for storing copies of the complete ledger and performing consensus algorithms. Hypothetical designs have been proposed, involving periodic updates, and intermediate network nodes, connecting the IoT devices to the separate blockchain network. However, other aspects (e.g. connectivity and security) require more consideration and further research.

According to Juels (2006) the amount if RFID tags will proliferate into the billions in the upcoming decade. These IoT enabled devices can silently be tracked within its short radius of range. In sufficient amounts, networks of RFID tags will be able to generate personal data. Unlike phones, which have similar capabilities, RFID tags usually have no batteries and are therefore not subject to being switched off. Furthermore, mobile transmitted signals are only receivable by specialized telecommunication equipment whereas RFID tags are scannable by commodity readers. The microchip itself can be as small as a grain of sand, some 0.4 mm^2. This makes it impractical for the consumer to control access, unlike a phone which hosts sophisticated encryption software. Given the impossibility to locally ensure privacy on a RFID tag, the network must instead facilitate that trust. Blockchains are high contenders in these use cases, as will be touched upon in subsequent chapters.

The oncoming paradigm of measuring and communicating devices have furthermore resulted in a general public unease on what the future holds. Famed Italian clothing company, Benetton, received threats of boycott when proposing implementation of IoT enabled RFID tags into some of its clothing lines (Violino, 2003). The purpose being to only track progress through production lines and evaluate store stock and therein disable the tag upon purchase. Yet, customers were afraid of technology misuse and breach of privacy assertions. This implies that security and privacy will be of uttermost importance if the IoT shall succeed.

2.4 Blockchain applications and challenges

Dieterich et al. (2017) have conducted a literature review of potential blockchain use cases and assessed their potential. To better understand the prospects of the technology, a survey has been conducted and analyzed together with discussions from seven interviews with scientific experts and blockchain practitioners. The authors draw the conclusion that the technology's potential is currently being investigated, and customized solutions are being developed for new application areas. The report presents data suggesting that 10% of global GDP will be harvested on blockchain networks by 2027. A report published by Markets and Markets (2019) suggests that the global blockchain IoT market will grow at a compound annual growth rate close to 93%. This prediction is based on the significant increase in IoT compatible devices and the growing appreciation for the ensuing security issues. Blockchain is anticipated to play an important role in addressing these concerns as well as enhancing business operations leveraging on smart contracts and increased transparency. The report further states that the major threats opposing these promising projections include legal complications, performance limitation, technology incomprehension and the absence of standards.

Focusing on the manufacturing industry, Dieterich et al. (2017) present tested and ongoing blockchain projects aiming to increase trust and traceability within supply chain networks, while reducing paperwork and costs. The technology's potential to support and manage the rapid increase in IoT applications are also mentioned. However, it is emphasized that technological problems such as scalability and transaction throughput currently prevent sufficient adoption. Though a few attempts to circumvent these limitations of the technology have been recorded, many remain hypothetical and unvalidated.

According to Dieterich et al. (2017), certain conditions required for efficient use of the technology includes the need of overcoming trust issues and efficient management of data shared by multiple actors. Though this would suggest high potential in shaping new business models around blockchain solutions, the benefits over existing technologies are not very clear and require careful consideration (Rodrigues et al., 2018). The technical limitations in terms of scalability and verification speed, together with legal aspects must be studied in more detail according to Kumar and Mallick (2018). Dieterich et al. (2017) conclude the report with a discussion related to the Gartner Hype Cycle and how current experts assess blockchain to be around its peek of expectation. Further research is required to elucidate the various aspects of the technology, and more thorough analyses of the business processes are needed to investigate the possibilities of blockchain in more detail.

The increase in IoT applications requires new ways to solve security concerns. In the attempt to alleviate trust issues between parties, Brenzikofer (2017) highlights the properties of blockchain. Recorded data from IoT sensors are suggested to be timestamped and stored on a blockchain for verification purposes. The immutable nature of the data, along with public access would reduce the trust gap between different actors. It would be easy to discover any attempts at manipulating the information stored on the blockchain. Further research is required to gain a better insight into the practicalities of such a system. It also shines a light on the vast architectural and design opportunities

available for businesses, eager to investigate how to best reap the benefits of blockchain.

The promising opportunities of blockchain encourages business to adopt the technology. In spite of the benefits, Sadu (2018) highlights that blockchain introduces a new set of challenges that must be comprehended and addressed before effective use. New threats including the 51% attack, endpoint vulnerabilities and legal complications must be addressed accordingly. The elevated security requirements and the system decentralization increases the complexity of developing and maintaining blockchain platforms (Bosu, Iqbal, Shahriyar, & Chakroborty, 2018). Halaburda (2018) discusses the integral components of blockchain and their individual benefits. She distinguishes between three concepts: (1) smart contracts, (2) encryption, and (3) distributed ledgers. These, she argues, must be considered independently in order to find the best uses in terms of cost and benefits. They do not depend on each other, but can be implemented independently (e.g., smart contracts may be used on a centralized system). This highlights the need of studying the technology in depth so to customize its design to the specific application area. Further inquiries, linked to how well the benefits of a distributed ledger counteracts the extra cost of slow verification and data duplication, are also put forward.

3. Methodology

Here follows an elaboration of the methodology used in this thesis. The framework is empirically constructed following a design science research (DSR) methodology over three cycles, in accordance with Wieringa (2009) and the guidelines provided by Hevner, March, Park, and Ram (2004). The architecture framework itself constitutes that artifact of the design science study. It is constructed following the international standard defined in ISO/IEC/IEEE (2011) and the additional recommendations by Hilliard (2014).

3.1 Design science research

This research follows the design science research methodology as described by Wieringa (2009). It is an iterative problem solving approach, with the ambition to successively build a complete artifact, as new knowledge is acquired. Figure 3.1 illustrates the the regulative cycle and its consecutive steps. This cycle is repeated three times throughout this work. The design revolves around an artifact that represents the problem at hand. The artifact can take on different forms including constructs, models, methods, and system instantiations (Hevner et al., 2004). According to Wieringa (2009), the artifact should be motivated by a business need. Given the blockchain technology anticipation presented in Chapter 1 and the lack of implementations standards, there is an apparent business value in creating a framework providing architects with blockchain development practices. The incremental construction of the artifact allows for continuous evaluation, ensuring that progress is made accordingly. Each cycle takes into account the evaluation provided by the end of the previous iteration. This ensures, Wieringa (2009) argues, controlled knowledge acquisition with the purpose of deepening the understanding about the artifact, while ensuring that progress is made according to stakeholder requirements.

3.1.1 Problem type

Design science research is a problem solving methodology and it is therefore important to define the problem under investigation. Wieringa (2009) makes a distinction between *practical* and *knowledge* problems. The classification of a problem is necessary so that an appropriate evaluation criteria can decide upon. A practical problem is concerned with how the world can be modified to satisfy certain stakeholder requirements. Efforts taken to solve a problem of this category require the evaluation against defined stakeholder needs to determine their appropriateness. In contrast, a knowledge problem is concerned with obtaining information about the world yet unknown to the stakeholders. The accuracy of the acquired answers must be evaluated against domain experts proficient in the field.

The ambition to create an architecture framework falls under the knowledge problem domain. It can be described as a subproblem of an architect's practical problem of creating a blockchain application.

It does not result in a blockchain solution by itself; rather it aspires to aid architects by providing them with knowledge about appropriate blockchain implementation practices. As stated in Wieringa (2009), solving such problems involve identifying relevant research and interviewing domain experts; two strategies pervading this thesis.

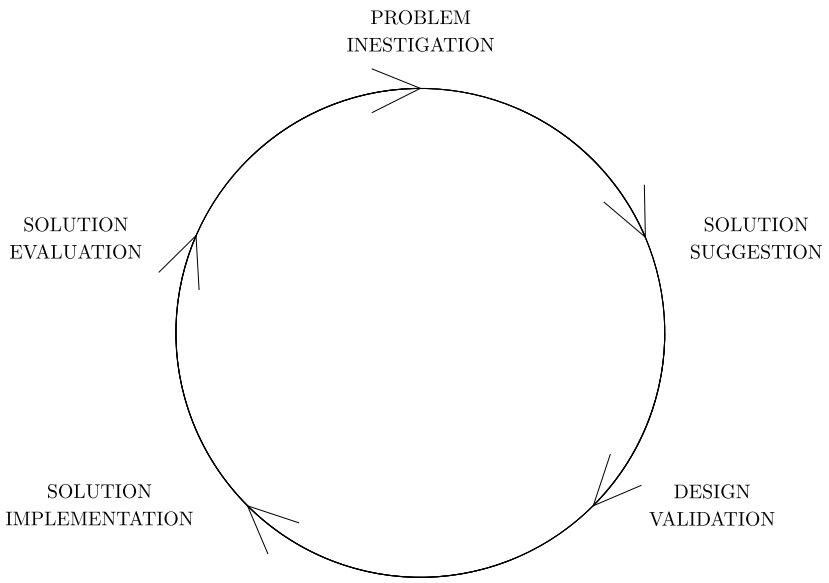

Figure 3.1: *The cyclic Design Science Research methodology.*

3.1.2 Artifact

In line with the recommendations by Hevner et al. (2004), a valuable artifact with the ambition to solve a business challenge has been defined. Each iteration will revolve around and incrementally construct an architecture framework in accordance with the standard defined by ISO/IEC/IEEE (2011). The framework's composition corresponds well with the design science research methodology's cycle objectives, addressing problems, solutions and architectural evaluations. More specifically, the framework requires the inclusion of the information depicted in Figure 3.2 which is further described below.

The **stakeholders** represent all the possible participants with an interest in the specific blockchain application. As illustrated in Figure 3.2, each stakeholder has a set of **concerns** or challenges that could either be unique for the particular group or shared among many different stakeholders. The framework makes it easy for stakeholders to identify and address their specific challenges.

The framework's **viewpoints** represent different system perspectives. A viewpoint should, according to ISO/IEC/IEEE (2011) *"include information on architecting techniques used to create, interpret or analyze a view governed by this viewpoint."* Each viewpoint frames a subset of the concerns held by the different stakeholders. A viewpoint's **correspondence rules** elaborate on its relations between other viewpoints.

The **model kinds** are tools that architects are recommended to use when constructing specific

system views from the viewpoints. According to the ISO/IEC/IEEE (2011) standard, *"an architecture view expresses the architecture of the system-of-interest in accordance with an architecture viewpoint."* In other words, a view is a concrete representation of the actual system, constructed from a specific viewpoint.

According to Hilliard (2014), a complete viewpoint must also include an **operations on views** section, providing guidance regarding model construction methods, interpretation methods, analysis methods, and/or implementation methods. Thus, this section supports architects to reason about the viewpoint's challenges, stakeholders, and models. For this framework specifically, this section is used to propose alternative architectural approaches, while elaborating on their respective benefits and drawbacks.

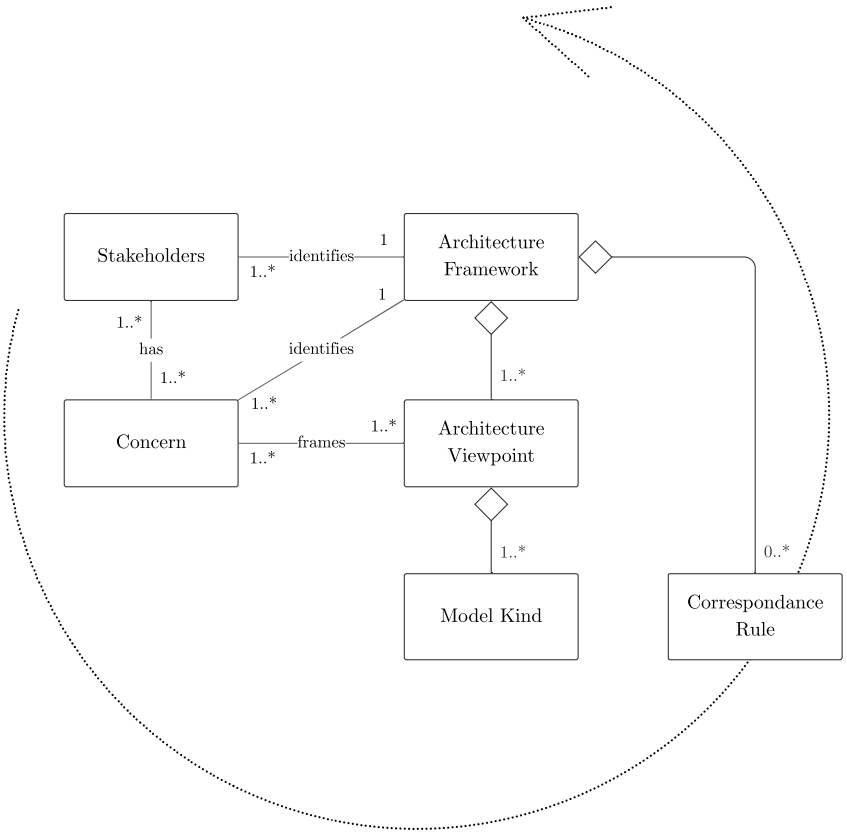

Figure 3.2: *The architecture framework artifact.*

3.1.3 Design science research components

As shown in Figure 3.1 each cycle constitute five major objectives: (1) problem investigation, (2) solution design, (3) design validation, (4) solution implementation, and (5) evaluation. These objectives are intertwined with the architecture framework and are addressed throughout the project but with different emphasis in each iteration. In Iteration 1, the investigation of blockchain challenges will be of greater importance to lay a foundation of knowledge and to discern where the most critical aspects reside. In Iteration 2, challenges will be accompanied by models, operations on views, and

correspondence rules in the architecture framework. Challenges will be mapped based on a dedicated literature review and a range of interviews. A set of viewpoints will be derived from the identified challenges and discussed in more detail. Lastly, in Iteration 3, the architecture framework will be refined and evaluated in discussion with case companies, interviewees, and at a conference identifying its potential to address challenges in industrial blockchain implementation projects.

Problem investigation

This phase is primarily concerned with acquiring information to facilitate understanding of the problem at hand. Knowledge is obtained without the wish to change (Wieringa, 2009). This investigation will be tightly linked to answering RQ1 in Section 1.4. As with most early technologies, research and practitioners are quick to produce content that highlights areas of improvement. The ambition with the problem investigation is to cover a wide range of research, touching upon most identified challenges, to later enable a justified prioritization.

Wieringa (2009) proposes four different rationales behind problem investigations:

- **Problem-driven investigation:** tries to unearth and diagnose the the problem; a prerequisite before the problem can be sufficiently addressed.
- **Goal-driven investigation:** does not necessarily have to address a certain problem but could be performed with the ambition to achieve a change or higher wish.
- **Solution-driven investigation:** analyzes a technology's potential to improve or some some particular problem.
- **Impact-driven investigation:** focus on evaluating the impact of previous activities.

During the first iteration, the problem investigation was primarily *problem-driven*. Literature investigations and online courses, together with case company discussions and five initial interviews comprised this phase. The ambition was to establish a sufficient knowledge base and to start identify some general problems with blockchain. As the work progressed, the focus shifted toward *solution-driven* investigations during the second iteration. By assessing the problems identified in the literature review and from nine additional interviews, appropriate guidelines were identified during the second iteration. Furthermore, by acquiring knowledge from literature and industry about the current use and future potential of blockchain, functionality and performance requirements were identified. The problem investigation during the third and last iteration was primarily *impact-driven*, considering how to improve the industrial value of the framework.

Solution suggestions

As depicted by Wieringa (2009), *solution suggestions* address means to reach stakeholder requirements. These means can take many forms: diagrams, sketches, blueprints, mathematical models, scale models, prototypes, etc. However, during this phase these solutions remain *suggestions*, given that neither implementation nor validation has yet been performed. It might appear, after implementation, that a suggested solutions work against some stakeholder requirements and thus need to be reconsidered. It is a "commitment act" between involved stakeholders, where different parties can share and explain

their their thoughts of the final product. These include stakeholders' ambitions to improve the artifact in ways they deem appropriate. In this study, the architecture framework will provide viewpoint and models to aid decision making regarding blockchain implementations and will thus be the means by which the stakeholders can achieve their goals. Literature, interviews and workshops have been used throughout the study to discuss solution suggestions.

Design validation

According to Wieringa (2009), the *design validation* phase serve the purpose of investigating whether the proposed solutions indeed brings the stakeholders closer to their goals. In this context, different blockchain approaches were analyzed with respect to their ability to address the identified challenges and help stakeholders to achieve their goals. While analyzing the architecture framework design, three aspects highlighted by Wieringa (2009) were considered:

- **Internal validity:** To what extent and efficiency can the challenges in the investigation be satisfied, would the proposed solution be implemented?
- **Trade-offs:** How does different design alternatives compare to each other?
- **Sensitive analysis:** Howe well would the suggested design satisfy its requirements in different contexts.

Regular discussion with case company representatives have been used to validate the finding in the light of their use case. The use of semi-structured interviews made it possible for interviewees to endorse or reject other's proposals and ideas. The large pool of literature investigated during the three iterations also serve as a benchmark against which solution suggestions were compared.

Solution implementation

Implementation are described by Wieringa (2009) as the process of carrying out the proposed and verified design. In the purpose of this study, this relate the construction of the actual viewpoints. The framework has been incrementally updated throughout each cycle. During iteration one, there was a heavy focus on challenges and stakeholders. The focus shifted toward viewpoint definitions during the second and third iteration.

Solution evaluation

Wieringa (2009) does not explicitly propose an *evaluation* section in the DSR. Hevner et al. (2004) emphasize, however, the importance of proper artifact evaluation; only by evaluating the artifact's *"utility, quality, and efficacy"* can future work be coordinated accordingly. The evaluation in each iteration was not postponed to the very end. As mentioned previously, regular case company discussions were scheduled to evaluate the findings in the light of the specific use case.

A workshop conducted during the second iteration was held together with four Biswise representatives. Challenges, potential solutions and an evaluation of the partial framework were discussed within one and a half hour. Support was also provided by the academic supervisor and a fellow researcher with experience in the actual construction of architecture viewpoints.

Table 3.1: *Prioritization of research articles.*

	Number of documents	
Inclusion/exclusion criterion	**Excluded**	**Included**
Search from 2016 to 2019	0	229
Limit results based on relevance	-129	100
Exclusion based on title and abstract	**-79**	**21**
Documents with unrelated titles		
Documents with inconceivable abstracts		
Documents with a too narrow focus		
Documents with poor linguistic standards		
Criteria based on full text	**-4**	**17**
Exclude documents that do not include clear challenges		
Exclude documents that do not include clear solution proposals		
Exclude documents with inconceivable purpose		
Exclude documents with ambiguous purposes		

During the third iteration, two one hour industry evaluation sessions assessed the framework. As the framework was presented, industry participants were encouraged to evaluate the findings and provide feedback on each viewpoint. Lastly, an interviewee survey (see Appendix B) was also sent out to study participants, and a SAPSA conference talk with subsequent feedback from impartial industry representatives further evaluated the work.

3.2 Literature review

Though not being a comprehensive systematic literature review, the process adopted in this study were inspired by the procedure described by Kitchenham and Charters (2007). After defining a search strategy and an inclusion criterion, the selection process was performed (see Table 3.1). The resulting subset of relevant articles was further evaluated based on the articles respective content. In this particular case, IEEE Xplore Digital Library and Engineering Village were used as the two database sources. A total of 229 primary studies were extracted based on the keywords: *blockchain*, *IoT*, and *challenges*. Based on each article's abstract and title, a subset of 21 were prioritized. Another four of these were excluded due to inadequate content. The challenges identified in the remaining 17 articles are presented in Table 4.2.

3.3 Interviews and workshop

Beside the literature review, qualitative data was obtained from interviews, workshops, and one survey. These collection techniques all fall under what Runeson and Höst (2008) call first degree methods: data is acquired in direct communication with the information sources. According to Runeson and Höst (2008), these methods can be easily customized so to increase the precision of the collected data, improving data quality. By investigating different blockchain technologies and interviewing various experts, the ambition was to reach a broad understanding of the technology

and how it could be accustomed to stakeholder's desired needs. Triangulation was used to improve the accuracy of the investigation by studying the problems and solutions from different perspectives. Triangulation is considered to be of great importance in qualitative research to counterbalance the often comprehensive but imprecise data (Runeson & Höst, 2008).

Table 3.2: *Study participants.*

Name	**Role**	**BW**	**I1**	**I2**	**I3**
Interviews					
Interviewee A	Blockchain developer		✓		✓
Interviewee B	Blockchain devotee		✓		
Interviewee C	Associate lawyer		✓		
Interviewee D	Junior consultant			✓	
Interviewee E	Junior consultant			✓	
Interviewee F	CEO			✓	
Interviewee G	Technology advisor			✓	✓
Interviewee H	Solution director			✓	
Interviewee I	CEO			✓	
Mail correspondence					
Interviewee N	Technology lawyer		✓		
Interviewee O	Researcher		✓		
Interviewee P	Blockchain developer			✓	
Interviewee Q	Researcher			✓	
Interviewee R	Team leader			✓	
Workshop					
Interviewee K	Consultant	✓			
Interviewee L	Senior partner	✓			
Interviewee M	Consultant	✓			✓
Supervisors					
Interviewee J	CEO	✓	✓	✓	✓
Interviewee S	Head of software		✓	✓	✓

Purposeful sampling was used to identify appropriate interviewees. In line with Palinkas et al. (2015), purposeful sampling is not a probabilistic sampling technique; rather, it aspires to locate well informed participants that can contribute with expertise knowledge within their fields. This strategy was appropriate given the limited time frame and availability of industry representatives. Moreover, in qualitative studies it is recommended to aim for a greater dispersion of interviewees. In accordance with Runeson and Höst (2008), the selection process thus emphasized variation and a diverse set participants were interviewed. Using the professional network LinkedIn, blockchain

proficient developers, project managers, technology advisers, legal counsellors and other devotees were identified and contacted. Some were further identified using a snowball strategy that involved asking interviewees whether they knew other relevant people (Palinkas et al., 2015). A complete list of study participants can be seen in Table 3.2.

To make sure that the interview template used did indeed comply with the investigation purpose and the three research questions, it was reviewed by case company representatives and the academic supervisor. This was done continuously, since the interview template was updated as the study proceeded and new directions were identified. The template (see Appendix A) follows a semi-structured approach and consists of both open-ended and more closed queries, taken inspiration from the funnel design described by Runeson and Höst (2008). The use of an interview guide helped ordering the responses and allowed for comparison between interviewee replies. Moreover, for the purpose of investigating new challenges and potential solutions, the combination of a semi-structured approach together with open-ended questions guaranteed sufficient flexibility. As emphasized by Doody and Noonan (2013) these types of interviews are appropriate when the researcher requires the possibility to investigate new areas of interests as they surface.

Following the recommendations provided by McLafferty (2004), the number of workshop participants were restricted to a smaller amount, so to increase interaction. To counteract the potential increase in moderating efforts, both researchers took part in the workshop; one with the responsibility of taking notes and the other managing the discussion, making sure it progressed on time. The same strategy was applied when performing the interviews. One researcher toke notes as the other conducted the interview. The strategy to transcribe each interview was considered initially. However, due to the limited time frame and the large number of planed interviews, it was decided that rigorous note taking would suffice.

A thorough review of the notes were performed instantly after each interview while the information could still be clearly recalled by the researchers. Whenever any ambiguity arose, the interviewee was contacted once again and asked to clarify the particular concept or statement. After getting familiar with the collected data, challenges and potential solution suggestions were abstracted from the responses. As suggested by Schutt (2009), the data was display using a spread sheet matrix that simplified the categorization of responses and identification of themes. The subsequent definition of challenging scenarios, viewpoints and guidelines were further evaluated during the workshop and the regular meetings held together with case company representatives.

4. Findings

This section aims to construct an empirically derived architecture framework in accordance with the ISO/IEC/IEEE (2011) standard. The standard describes a formal structure, by which new viewpoint definitions should abide. The required sections include questions and concerns, model kinds, operations on views, and correspondence rules. The intention of the framework is to formalize principles that architects are recommended to follow when implementing blockchain solutions. Current research is moderately disordered and the framework's ambition is to structure the identified architectural aspects in a coherent and accessible manner.

The section begins by identifying the main stakeholder categories of blockchain networks and their most prominent challenges. This exposition is followed by five viewpoints derived from these challenging scenarios. Lastly, an evaluation of the framework is presented assessing its relative quality from an industrial standpoint.

4.1 Research question 1

The sections of the architecture framework are aligned with each research question. The stakeholders and their concerns about blockchain implementations are extracted with an IoT perspective in mind. To revive the motive behind the first research question, it is restated here:

> *Based on stakeholders and their concerns, which challenging scenarios are important to consider when assessing the use of blockchain?*

4.1.1 Stakeholders

In this section, a generic set of stakeholder categories has been constructed in accordance with research question one from blockchain literature and interview responses. A network member can take on one or multiple roles with varying characteristics. Each category is defined in broad terms and can be further decomposed; nonetheless, they provide a foundation on which to derive concerns.

- **Blockchain platform developers:** The blockchain platforms currently on the market are open-source and maintained by devoted communities. Founders, investors and developers etc. are responsible for the existence of these platforms and are thereby essential stakeholders. There are large communities maintaining and updating different platforms such as Bitcoin, Ethereum and Hyperledger.
- **Network managers:** Depending on the implementation specifics, a blockchain application falls on the continuum from completely distributed to centrally governed. That said, some applications will require the establishment of a governance system. This is true for private

and permissioned blockchains, in which there exist one or more network administers with configuration privileges.

- **Blockchain application developers:** This category is responsible for designing and implementing the specific blockchain application on top of the chosen platform. It includes the system architect as well as the developers, charged with implementing the application logic and the development of smart contracts.
- **Node participants:** The decentralized nature of blockchain is realized by networks of dedicated computers, participating in the consensus protocol and appending the chain with new blocks of transactions. Nodes that keep a copy of the complete ledger and take part in the verification process are called *full* nodes. The *mining* nodes are full nodes entrusted with the additional task of appending new blocks to the ledger (Panarello et al., 2018). Other clients require blockchain access but do not have the computational resources to neither store nor process the chain's data. IoT devices fall into the this category. These *lightweight* nodes are able to send and request specific data from the blockchain through a base station or by communicating directly with a full node (Danzi et al., 2018).
- **Data producers:** The essence of blockchain is to manage data in a distributed and secure manner. There is no limit as to what data can be stored on a blockchain (Conoscenti et al., 2016) and the producers generating data can range from miniature sensors in user gadgets to large institutions recording financial transactions between corporations.
- **Data consumers:** The data consumers are stakeholder with an interest in the blockchain data. These include, for example, individuals or corporations trading monetary assets within the network, or actors using the verified data as input to their data algorithms.
- **Legal associates and other regulators:** Lawyers are currently involved in dialogues regarding the future development and use of blockchain. Recent initiatives such as the GDPR and regulations alike paint the regulatory environment in which blockchain applications will have to be revised. In an era of increased privacy awareness, legal aspects are gaining significance.

Each blockchain use case will involve a unique set of stakeholders, each of whom can take on one or several of the aforementioned roles. The stakeholders identified during the workshop with Biswise will serve as an illustrative example of a potential network composition. Table 4.1 provides a list of the stakeholders brought up during the discussion. Being the initiator of the battery trading platform, GreenStar Marine will serve as the *network manager*. They will also help develop the system, as well as analyze the battery data loaded onto their platform. This suggests that they will take on the additional role as *node participant*, *application developer* and *data producer* and *consumer*.

GreenStar Marine will neither be developing the application from scratch, nor by themselves. They will work in collaboration with technology partners providing them with a suitable blockchain platform and sufficient expertise. These technology partners fall in the categories of *platform* and *application developers*.

The main idea behind the platform is to collect and verify battery data to accurately assess its

Table 4.1: *Potential stakeholders and suggested role(s) based on the workshop discussion.*

Stakeholder	Platform developer	Network manager	Application developer	Node participant	Data producer	Data consumer	Regulator
GreenStar Marine		✓	✓	✓	✓	✓	
Technology partners	✓		✓				
OEMs				✓		✓	
Financial institutions				✓		✓	✓
Insurance companies				✓		✓	✓
Interest organizations				✓		✓	
Boat owners				✓	✓	✓	
Battery buyers				✓		✓	
Retailers				✓			

quality. In order to obtain the necessary data, GreenStar Marine will have to work with battery manufacturers and other OEMs, with the ambition to integrate sensors or establish other interfaces that will enable battery data to be communicated onto the blockchain. It is assumed that the manufacturers would at least request access to the data generated by their components for the purpose of product development, making them *data consumers*. The boat owners are *data producers*, generating data as they drive along. They are also likely to wish for information about their own use, making them *data consumers* as well. Prospecting sellers are interested in finding secondary batteries that would match their respective use cases. As *data consumers* they would acquire historic data and quality assessments of batteries to find a good match.

There are additional actors likely to take part in the network including financial institutions, insurance companies, interest organizations and retailers. They are all presumed to consume data. For example, by giving retailers access to battery status, they can approach boat owners with tailored offers when it is favourable to upgrade their battery packs. Financial institutions and insurance companies may also want to take part in the verification of new data.

It should be noted that other, more detailed, consortium constellations are possible. Nevertheless, it sheds some light on the rather complex task of gathering necessary participants and establishing a fully functional network. It is important that the roles are identified early on so that each category and their associated challenges can be addressed accordingly.

4.1.2 Concerns

The challenges are extracted from literature and interview discussions before grouped into wider scenarios. A list of 16 anticipated challenges of blockchain were identified in the literature review and based on 17 articles (see Table 4.2). The *No. of References* row indicates the total number of articles in which the specific problem was mentioned or discussed. As a complement to the literature review, challenges were investigated during interviews with practitioners including developers, project managers, technology advisers, legal counsellors and other devotees. In total, 14 interviews

were conducted showcasing where the most pressing challenges reside at present. For the sake of comprehension, these challenges have been grouped under the same categories as in Table 4.2 and can be further studied in Table 4.3. As in Table 4.2, the *No. of References* row indicates the total number of interviews in which the specific problem was explicitly mentioned or discussed.

Table 4.2: *Literature review.*

	Technology ignorance	Applicability issues	Large cost	Lack of proficient developers	Smart contract development	Scalability limitations	Demanding consensus protocols	Limited transaction throughput	Transaction latency	Bandwidth limitations	Storage restrictions	Large energy consumption	Constrained IoT devices	Intermittent connectivity	Security issues	Legal complications
Challenge No.	**C1**	**C2**	**C3**	**C4**	**C5**	**C6**	**C7**	**C8**	**C9**	**C10**	**C11**	**C12**	**C13**	**C14**	**C15**	**C16**
Ramachandran and Krishnamachari (2018)	✓						✓	✓	✓	✓			✓	✓	✓	
Christidis and Devetsikiotis (2016)		✓			✓		✓	✓	✓						✓	✓
Kumar and Mallick (2018)	✓			✓		✓	✓				✓					✓
Zheng et al. (2017)						✓	✓	✓	✓						✓	
He, Guan, Lv, and Yi (2018)						✓		✓	✓	✓	✓	✓			✓	
Gao, Hatcher, and Yu (2018)						✓	✓	✓	✓		✓				✓	
Ferrag et al. (2018)		✓													✓	
Chalaemwongwan and Kurutach (2018)				✓								✓				
Reyna et al. (2018)					✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓
Yu, Li, Tian, and Liu (2018)															✓	
Bosu et al. (2018)	✓			✓	✓										✓	
Fernández-Carmés and Fraga-Lamas (2018)		✓			✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓
Fabiano (2017)													✓		✓	✓
Song, Demir, Prevost, and Rad (2018)							✓	✓	✓							
Roy, Ashaduzzaman, Hassan, and Chowdhury (2018)						✓	✓	✓				✓	✓		✓	
Panarello et al. (2018)			✓	✓		✓		✓	✓		✓		✓			
Conoscenti et al. (2016)						✓	✓	✓			✓	✓	✓		✓	
No. of References	**3**	**3**	**1**	**4**	**4**	**9**	**10**	**11**	**9**	**4**	**7**	**6**	**7**	**1**	**13**	**5**

Blockchain is a sophisticated and multifaceted technology that advances at the forefront of computer development. As blockchain continues to allure skilled entrepreneurs, computer scientists, and software engineers, the ecosystem surrounding the technology will continue to grow. The increased participation will spur new application areas and identify novel technology adoptions. The check marks in Table 4.3 are clustered toward the left hand side of the table under technology ignorance and applicability issues. There is also an emphasis on the complexity of smart contract development. This suggests that scarce competence and unfitting use cases requirements are causing companies of today to remain hesitant and unable to see its potential. Although there exists an awareness of the technology, many practitioners highlight the inability to embrace and adapt blockchain.

Progress within the area of blockchain is nurtured in symbiosis between researchers and practitioners. As can be seen in Table 4.2 and Table 4.3, however, the two categories have slightly different focus. The research community encourages innovation and a continuous expansion of the current pool of knowledge. Two of the most heavily referenced problems in the set of articles relate to

Table 4.3: *Interview review.*

	Technology ignorance	Applicability issues	Lack of proficient developers	Smart contract development	Scalability limitations	Demanding consensus protocols	Constrained IoT devices	Security issues	Legal complications	User experience issues
Challenge No.	**C1**	**C2**	**C4**	**C5**	**C6**	**C7**	**C13**	**C15**	**C16**	**C17**
Interviewee A	✓	✓	✓	✓			✓	✓	✓	
Interviewee B	✓	✓								
Interviewee C		✓						✓	✓	✓
Interviewee D	✓	✓		✓				✓		
Interviewee E	✓	✓		✓				✓	✓	
Interviewee F		✓			✓	✓		✓	✓	
Interviewee G	✓		✓	✓			✓	✓		✓
Interviewee H	✓	✓			✓			✓		✓
Interviewee N									✓	
Interviewee O			✓							
Interviewee P		✓		✓						
Interviewee Q				✓		✓	✓			
Interviewee R	✓	✓		✓						
No. of References	**7**	**9**	**3**	**7**	**2**	**2**	**3**	**7**	**5**	**3**

(1) the inefficient consensus protocols and (2) the security aspects of the technology. This may not be so surprising considering the fact that each relate to one of blockchains most fundamental principles. The consensus protocol is the enabler of decentralized governance and cryptography pervades the blockchain infrastructure. Distributed systems and cryptography are mature areas of computer science and their strengths and weaknesses are therefore covered extensively in literature. Blockchain is fascinating computer scientists all around the world who are working on enhancing its technological composition. There is a heavy focus on improving scalability and throughput, while reducing transaction latency, storage restrictions and the constraint imposed by computationally restrained IoT devises. These are highly relevant areas when refining the underlying infrastructure of blockchain and its compliance with IoT. However, Table 4.2 indicates that literature show less concern for the actual business value of the technological developments.

This is what practitioners are currently struggling with. They are conducting a form of development themselves but with the purpose of adopting the technology to generate business value. Henceforth, consumers will put higher demands on product and service transparency. Current research, however, seems to make progress without sufficient regard to user experience and the commercial use of blockchain. There is an evident lack of focus on how to compile and derive user value from the technology. Practitioners have raised concern around the problem of processing and visualizing blockchain data in an efficient manner (as can be seen under the user experience issue in Table 4.3). There is also an apparent opinion among practitioners that the technological legislation is advancing and that attention to user privacy is needed.

4.1.3 Challenging scenarios

IoT practitioners struggle with technical, political, and ethical impediments. As previously discussed, blockchain technology has the potential to mitigate some of these hurdles and are perceived as an enabler of IoT advancement (Mittal et al., 2018). That said, there has been a shift away from cryptocurrencies toward a large range of conceivable application areas, ranging form large machinery and mobile devices to minute sensors (Wüst & Gervais, 2018). Particular use cases of IoT (e.g. logistics, assisted driving, and health care) prioritize different quality solution attributes and will therefore be differently suited for blockchain use. Also, the economics of such networks remain an intriguing challenge. Currently, resources are mainly dedicated to improving the underlying technology but future discussions will have to address the business aspects as well, such as how to establish network effects around these new systems.

Blockchain has the potential to reduce cost and improve efficiency through disintermediation and the automation of processes accomplished by smart contracts. Rodrigues et al. (2018) emphasizes, despite the evident benefits, that it is necessary to appropriately handle the technical implications and the changes inevitable to occur within current business processes. Here follows a set of critical scenarios (c.f. Table 4.4) that must be addressed by architects designing blockchain applications. Although these scenarios will be attended to differently depending on use case, they all relate to the first research question. These scenarios will aid the construction of the architecture viewpoints discussed later in this section.

The scenarios capture one or more of the challenges addressed during interviews and in literature (c.f. Table 4.2 and Table 4.3). By compressing the concerns into a limited set of challenging scenarios, the ambition is to make the work more manageable for architects. As can be seen in Table 4.2 and Table 4.3, each challenge is given a specific *Challenge No.* that is further referenced in Table 4.4, visualizing what challenges are addressed by the specific scenario. To systematize the output from interviews and literature, there was a need to define a number of general themes applicable to all blockchain applications. By analyzing the obtained data, five themes under which the scenarios are grouped were identified. Moreover, these five constitute the basis for each viewpoint addressed by the second research question in the subsequent section.

Scenarios relating to applicability

Novel technologies are often the subject of grand visions and conceptual ideas, rather than being treated with logic and sense – blockchain is no exception. Much research is still needed to discern in what situations and application areas the technology can provide an advantage over more mature and efficient technologies. The applicability viewpoint aims at addressing some of these challenges.

As an architect I want to properly evaluate the benefits of using blockchain technology over traditional distributed systems in the specific problem domain:

Considering the fact that the blockchain technology itself is so often the subject of misinterpretation and confusion, it is apparent that assessing its usability and implementing it successfully is challenging.

Table 4.4: *Challenging scenarios.*

Scenarios relating to applicability	
As an architect I want to properly evaluate the benefits of using blockchain technology over traditional distributed systems in the specific problem domain. (C2, C3)	As an architect I want to identify appropriate use cases and potential trade-offs to mitigate risks through blockchain illiteracy. (C1, C2, C3)
Scenarios relating to ecosystem competencies	
As an architect I want to actively identify the level of blockchain competency in the application ecosystem so that I can manage the technical risks. (C1, C4)	As an architect I want to actively manage the level of domain and legal knowledge to define and test smart contracts so that I can manage the technical risks. (C1, C4, C5)
Scenarios relating to quality attributes	
As an architect I want to improve transaction latency and throughput to enable large data processing. (C7, C8, C9)	As an architect I want to enable application scalability to be predisposed toward future expansions. (C6, C7, C11, C12)
As an architect I want to satisfy performance by making the application compatible with constrained IoT devices and their low bandwidth and intermittent connectivity. (C10, C13, C14)	As an architect I want to ensure sufficient security by selecting an appropriate blockchain platform. (C15)
Scenarios relating to legal abidance	
As an architect I want to define network responsibility to avoid legal disputes and financial penalties. (C15)	As an architect I want to ensure sufficient privacy to prevent any exposure of personal data. (C15, C16)
As an architect I want to enable sufficient data manipulation to comply with the data owner's rights to rectification and erasure. (C15, C16)	
Scenarios relating to end user experience	
As an architect I need to consider the interface between the user and the blockchain application to properly visualize and derive value from data verification and transparency. (C17)	

There are many parameters that must be considered and weighed against each other including performance, trust imbalances, and robustness (Fernández-Carmés & Fraga-Lamas, 2018). Companies often make an effort to leverage on blockchain without properly considering the particular use case applicability (Interviewee A, 2019). The initial question strategists must answers is therefore whether blockchain can or should be applied to a specific use case scenario or not.

As an architect I want to identify appropriate use cases and potential trade-offs to mitigate risks through blockchain illiteracy:

The blockchain technology has the potential to transform the way humans and corporations interact and collaborate. It can be described as building on top of two technological foundations, the first being asymmetric cryptography and the other distributed systems. These technological branches have been around for several decades. From that point of view, blockchain is not a revolution; rather, it is an evolutionary technology that disintermediates current systems. Rodrigues et al. (2018) phrase blockchain as an evolutionary technology because of its potentially extensive impact within areas outside the technical sphere, including but not limited to business practice, process design, and legal compliance. Still, blockchain represents a novel concoction and its components' technological depth allow practitioners to submerge themselves in the technology's complexity and varying features. There are constantly new developments and technological advances to improve this technological merger.

It has been marked by industry practitioners that blockchain is still a fashionable term, whereas the technological understanding and its potential impact remain unexplored. Blockchain remains at the initial stages of the adoption cycle, involving innovators and early adopters. Interviewee H (2019) agrees that comprehensive frameworks and guidelines must be established. They are necessary to aid developers and corporations in identifying relevant use cases and help them to decide on appropriate techniques. As emphasized by Interviewee R (2019), it is important to embrace the properties of immutability and the transparency from the very start; the concern being that developers might try to fight these properties just because they are unfamiliar. Interviewee H (2019) is currently working with identifying relevant use cases and informing their enterprise customers about the technology's prospects. Moreover, Interviewee H (2019) confirms that the technological aspects dominate the blockchain discussion. Although important in itself, this narrow focus prevents the technology from gaining a foothold among practitioners who is left without guidance on how to apply the technology. According to Interviewee F (2019), there is only a few actors present on the market that help companies to liaise with others and educate them in the benefits of blockchain. Focus will have to shift toward business relevance and user experience to include these actors beside the engineering community. Furthermore, as stated by Interviewee C (2019), the technological progress has preceded the advancements within legal practices, but that gap is being narrowed. This illustrates the important recommendation of Rodrigues et al. (2018) to consider the *lasting* impact blockchain will exert on the specific use case.

Scenarios relating to ecosystem competencies

As an architect I want to actively identify the level of blockchain competency in the application ecosystem so that I can manage the technical risks:

Blockchains are highly sophisticated platforms of technology and to launch an entirely new blockchain network is exceedingly difficult (Interviewee D, 2019). Rather, applications can be built as forks on top of already existing networks. This requires less sophisticated developer expertise which is, however, still difficult to attract. In particular, competence revolving smart contracts are tough to come by (Interviewee A, 2019). Interviewee I (2019) exclaims that novel and complex technologies commonly experience these hurdles, simply because most sought after personnel are efficiently kept by large corporations and unwilling to face associated risks involved in novel technologies. Kumar and Mallick (2018) note that, although blockchain proficient workforce is hard to come by, developers with an expertise of the integration of blockchain with IoT is even more scarce.

> *" It's absurdly complex [to create a blockchain network from scratch] and therefore prohibitively expensive. As if that wasn't enough, the technical talent is scarce; it's tough enough to find smart contract developers for an existing platform. Developing a blockchain requires substantial amounts of technical expertise and game-theoretic analysis. "*
>
> *– Interviewee A (2019)*

The application domain, or the environment in which the blockchain system will function, determines many of its requirements (Gunter, Gunter, Jackson, & Zave, 2000). Thus, an appropriate architecture cannot be constructed before the architect has gained a sufficient understanding of the domain. A blockchain network can be compared to an ecosystem of cooperating actors leveraging on shared network effects. The distributed nature of blockchain and the many involved actors urge architects to map capabilities within the consortium; who should take responsibility for the different technical aspects? The architecture of the system itself will determine what competencies are needed.

As an architect I want to actively manage the level of domain and legal knowledge to define and test smart contracts so that I can manage the technical risks:

A majority of interviewees are of the opinion that smart contract development must be done correctly from the very start. Although it is possible to freeze contracts and make modifications after deployment, the room for improvement is limited (Interviewee O, 2019). The activities of testing and debugging smart contracts are strained by this inflexibility. According to Interviewee P (2019), although modifications are not ruled out completely, it is vital to consider all scenarios in advance. This is endorsed by Interviewee Q (2019), emphasizing that the inclusion of smart contracts must be considered from the start since they determine the rules of how external entities such as IoT sensors interact and perform transactions.

Smart contract development differs from what is often considered regular programming. Modular and memory efficient programming practices are being extended to include computational cost efficiency. The fact that the smart contracts are executed by external network nodes adds a new dimension to the development practice. In public blockchains, this execution is also performed at a cost. To sufficiently cope with an increase in network size, the computational effort to execute smart contracts cannot grow linearly with the number of network nodes. Rather, it has to be performed in a constant time perspective (Interviewee A, 2019). Furthermore, bugs will be more expensive and probably not

so easily corrected depending on the architecture, because of the decentralization; the code is not longer hosted on a single server.

> *" The range of the scenarios that are being recorded and coded in smart contracts makes it challenging to ensure the network is behaving the way as it should among all the nodes. Another reason is, of course, the performance of the smart contracts and how heavy computations and calculations that are done before the data is added to the ledger. "*
>
> *– Interviewee P (2019)*

Also, the fact that smart contracts are executed autonomously without any supervision increases the need for other guarantees. The contract's logic needs to be carefully scrutinized to address all possible outcomes and implement verified fail-safe mechanisms. Once the contract is deployed onto the immutable chain, the opportunity to modify it is severely limited (Christidis & Devetsikiotis, 2016). Fernández-Carmés and Fraga-Lamas (2018) highlight another difficulty, namely that of translating legal terms into executable software representations.

SCENARIOS RELATING TO QUALITY ATTRIBUTES

As an architect I want to improve transaction latency and throughput to enable large data processing:

Since the emergence of blockchain technology in 2008, the relevance of transaction latency and throughput has increased. The upload throughput is proportional to the block frequency which indicates the average time needed to mine one block. This frequency is held at around 10 minutes per block in the Bitcoin network, increasing the complexity of the PoW puzzle as the computational power on the network improves. In combination with the block size being delimited to 1MB, the Bitcoin network can only handle 7 transactions per second (Vukolić, 2016). There are, however, more efficient blockchain platforms. The Ethereum network, for example, mines a new block approximately every 12 seconds (Hertig, 2019). Furthermore, private blockchains, using more efficient consensus protocols, are even more effective. Permissioned blockchains in general, however, suffer from the increase in communication overhead (Shafagh, Burkhalter, Hithnawi, & Duquennoy, 2017). Despite recent improvements, transaction throughput remains a large concern independent of technology platform. In comparison, major credit-card companies handles 2000 transactions per second and are capable of handling 10000 transactions at peak capacity (Vukolić, 2016).

As an architect I want to enable application scalability to be predisposed toward future expansions:

Vukolić (2016) argues that the consensus efficiency is the essential driver of a range of performance challenges, such as scalability opportunities, storage restrictions, and energy consumption.

Proof of Work (PoW) is the standard consensus algorithm used by the majority of the public blockchain platforms. Its robustness, however, causes it to be inherently slow and resource demanding, resulting in reduced performance.

> *" The consensus, mining, and validation systems; these concepts have a deep impact on the power consumed by IoT nodes and infrastructure, so they should be chosen in order to meet the actual power requirements. "*
>
> *– Interviewee Q (2019)*

One could adjust block size, but that would effect the throughput and security. The larger block size, the longer it will take to process and validate each block. Roy et al. (2018) argues that the Proof of Stake (PoS) protocol could improve the prospects of a successful Blockchain-IoT application. PoS requires only a fraction of PoW's computational power. The reduction in requirements also improves transaction speeds. As these words are written, Ethereum is finalizing its alternative PoS platform which is expected to be done by the end of next month (Kim, 2019). Byzantine Fault-Tolerant (BFT) is another popular consensus family (Salimitari & Chatterjee, 2018).

From an IoT perspective there is also a need for low transaction costs. The majority of the large and public blockchains are often subject to high fluctuations. The openness causes the transactional cost to rise and fall with supply and demand, which recent events have shown to be rather precarious. According to Interviewee A (2019), microtransactions could be prohibitively expensive during periods of high load.

As an architect I want to satisfy performance by making the application compatible with constrained IoT devices and their low bandwidth and intermittent connectivity:

IoT sensors, and other data producers are constrained in their ability to continuously communicate with a potential blockchain platform. They are also incapable of holding any reasonably large sized blockchain locally, nor mining its transactions (Danzi et al., 2018). Furthermore, IoT devices often suffer from bandwidth constraints and intermittent connectivity, meaning that they seldom communicate regularly (Ramachandran & Krishnamachari, 2018).

As an architect I want to ensure sufficient security by selecting an appropriate blockchain platform:

Blockchain technologies are generally categorized with respect to their openness and network permissions. A private blockchain, as opposed to public networks, requires consent before entrance by a network administer. A permissioned blockchain follows an endorsement policy to restrict blockchain actions among different network participants (Christidis & Devetsikiotis, 2016). A private network constitutes a more tightly knitted community where the identity of every cooperating peer is known only within the network. The increased transparency guarantees a higher degree of accountability. On the contrary, there is an increase communication overhead and the concern of centralization, endangering the computational democracy among distributed peers (Reyna et al., 2018). Also, private networks are generally smaller which reduces the strength in numbers. Interviewee A (2019) declares that network security increases exponentially with the number of participating nodes. As more nodes take part in the verification of new blocks, it will become harder to carry out any form of attack; the easiest to conceptualize being the 51% attack. This threat has gained special attention after of the rise of large mining pools on the public blockchains. Parker (2015) estimates that the Bitcoin blockchain accommodates 5000 full nodes, half of which actively mine new blocks. A majority of

this computational power is managed by mining pools (Mining, 2018). Nonetheless, the consensus algorithms are strengthened and more capable of withstanding these adverse attacks as the number of participants grow.

Scenarios relating to legal abidance

According to Herian (2018), blockchain has become an enabler of data verification, security assurance, and distributed governance. Disregarding the praise, it is important to realize that the technology is currently gaining a strong foothold within financial, legal, and corporate systems, not all of whom are currently considering the technology's larger social impact. As Harari (2018) concisely puts it, *"humans were always far better at inventing tools than using them wisely."* Although one may not agree with this statement in its entirely, it do call attention to the importance of applying the technology appropriately. There is an expressed lack of understanding among politicians and legislators about the profoundness of blockchain and other uprising technologies in general (Interviewee H, 2019). This inability prevents us from obtaining a clear understanding of the technology's wider social implications as well as the extent to which it can be controlled and regulated (Harari, 2018). Although politics and jurisdictions are inescapably destined to straggle along behind technological advancements, recent initiatives such as the GDPR and the new copyright directive suggest a more aggressive approach toward regulating technology use. Much work is still required, however, to prepare the political and legal institutions for the radical innovations emanating from blockchain research (Lyons, Courcelas, & Timsit, 2019). Additionally, there is an expressed concern that lawmakers, due to their technological ignorance, will impede technological advancement. Herian (2018) advocates blockchain regulation to be shaped in concert with its technological evolution to encourage further research while simultaneously guaranteeing sufficient safety and virtue. This may prove difficult due to the rapid progression of technological innovation but future initiatives will inevitably have to bridge the gap between software engineers and lawyers (Interviewee C, 2019).

As an architect I want to define network responsibility to avoid legal disputes and financial penalties:

The lack of accountability in public blockchains imposes a legal difficulty since no one is liable for the network activity. The GDPR requires the existence of a data controller who is responsible for the data processing (Lyons et al., 2019). The extent of this challenge depends on the blockchain technology chosen. From a legal perspective, the architectural principles of private and permissioned blockchains enable more convenient compliance, due to the fact that control is limited and exercised among known participants. In public and permissionless blockchains, this issue get more complicated and could disqualify many application attempts (Interviewee C, 2019).

As an architect I want to ensure sufficient privacy to prevent any exposure of personal data:

Despite introducing one or more liable network owners, there is still other aspects of the GDPR that need solving. Blockchain's promise to ensure security is its main driver of attention. Some interviewees indicate that security is indeed better compared to other data storage solutions. Nevertheless,

security threats are forever changing and will continue to remain present. New methods of solving cryptographic algorithms have followed in their wake, suggesting that present techniques are likely to be insufficient in the future; quoting Prof. Snape:

> *" [Security concerns] are many, varied, ever-changing, and eternal. Fighting them is like fighting a many-headed monster, which, each time a neck is severed, sprouts a head even fiercer and cleverer than before. You are fighting that which is unfixed, mutating, indestructible. "*
>
> *– Prof. Snape (Rowling, 2005)*

Related to security, data anonymization must be considered. Complete data anonymization is today not yet possible and participants remain pseudoanonymous, which does not adhere to GDPR when processing personal data. In public blockchains, it has been demonstrated that transaction activities can be effectively analyzed and correlated to specific users (Roulin, Dorri, Jurdak, & Kanhere, 2018). Developers will have to study the specific use case and identify appropriate security measures. Approaches include the combination of on-chain and off-chain storage, where the actual data is stored off-chain but hashed against a blockchain (Posadas & Dalmacio, 2018). However, Interviewee N (2019) call attention to the fact that even such solutions may not be allowed under current interpretations.

> *" Pseudonymization does not solve anything because you still have personal data processing. What's more, even usually used techniques such as storing hashes only on the blockchain is ridiculously not considered a proper anonymization . . . Again, the hope is in privacy technologies applied to blockchains as well as common sense legal interpretations. "*
>
> *– Interviewee N (2019)*

Lastly, several papers reports on the issue of constrained IoT devices. This fact constitutes a major security threat. Although the blockchain provides security assurances, one cannot simply ignore malicious IoT devices interacting with the blockchain. Malicious activities are most likely to take place at the wallet and browser level where the system is most vulnerable (Fabiano, 2017). The risks also lies in other aspects of the technology, such as the impossibility to access an account with lost login credentials.

As an architect I want to enable sufficient data manipulation to comply with the data owner's rights to rectification and erasure:

The rights to rectify and erase personal data are especially difficult in a blockchain scenario considering its immutability. Sufficient approaches are still being developed and discussed among lawyers and remain, according to Interviewee N (2019), the most *"relevant question"* that needs solving.

Scenarios relating to end user experience

As an architect I need to consider the interface between the user and the blockchain application to properly visualize and derive value from data verification and transparency:

As mentioned in the analysis of Table 4.2 and Table 4.3, the primary focus on the underlying blockchain technicalities have overshadowed the commercial use of the technology. End users often

experience difficulties using blockchain applications. A conventional bank app is, for example, often much more comprehensible than a blockchain wallet. Interviewee H (2019) is concerned with this inadequate interface between users and blockchain. The objective is to improve the end user experience without mitigating the verification and transparency benefits. This becomes even more prominent in a blockchain based IoT scenario processing vast amounts of data from a plethora of different IoT devices. To make sense of and accessing this data in a sufficient manner will be challenging. It will involve big data management on top of a complex integration of hardware and supporting software systems (Interviewee G, 2019). The new level of data management sophistication followed by GDPR, and the higher demands on product quality verification force system architects to investigate novel ways to fulfill customer requirements of user-friendly and comprehensible interfaces.

It is difficult to quantify the importance and internal prioritization of each scenario described above. One might argue, however, that their individual relevance may vary depending on the current implementation stage. Scenarios relating to applicability and technological competency are considered to be more relevant in the initial implementation phases, before being outstripped by more technical and design related scenarios. Also, different applications may prioritize different user stories depending on its purpose. For example, the relevance of GDPR is different when handling corporate data compared to when managing personal information. Another example would be blockchain-IoT solutions and their critical performance concerns relating to transaction throughput, scalability and security.

4.2 Research question 2

This section provides five viewpoints to consider when implementing blockchain applications. Each viewpoint is derived from the scenarios presented in Table 4.4. The purpose of these viewpoints are to dive deeper into the five scenarios presented in the previous section and to give guidance on how to address the challenges held by the relevant stakeholders. To revive the motive behind the second research question, it is restated here:

> *What viewpoints are necessary to consider in blockchain implementations and how can different architectural principles address their respective challenges?*

Each viewpoint is structured in accordance with the recommended template proposed in ISO/IEC/IEEE (2011), which is further refined by Hilliard (2014). Under *models*, recommendations on modeling techniques are presented. The *operations on views* sections provide directions and suggested practices that architects should consider while constructing blockchain systems and evaluating alternatives. Lastly, some of the viewpoints presented have similarities with already established viewpoints from other frameworks. These relationships are further discussed under the *correspondence rules* sections. The five viewpoints are summarized as follows:

1. Applicability viewpoint
2. Ecosystem viewpoint
3. Infrastructure viewpoint
4. Legal viewpoint

5. End user viewpoint

4.2.1 Applicability viewpoint

> ***The applicability viewpoint describes the application domain and its involved stakeholders, evaluating their properties and requirements in the light of what blockchain has to offer.***

The initial question an architect must address is whether or not blockchain is a potential solution candidate. A majority of the individuals interviewed for this thesis, regardless of position, brought up the applicability concern (c.f. Table 4.3). The technology is envisioned to have the necessary power to create a secure IoT infrastructure (Kumar & Mallick, 2018) and some visualize it as the future glue that will power the next generation of Internet (Lyons et al., 2019). However, that remains to be seen. Due to the current intensive promotion, practitioners risk making uninformed decisions on how to apply the technology.

Key questions and concerns

The following key questions have been identified from the scenarios described in Section 4.1.3 and are the most anticipated challenges identified by practitioners and in literature.

Question 1: How can system architects properly evaluate the benefits of using blockchain technology over traditional distributed systems?

Today's ubiquity of data has left the single node of a blockchain insufficient in terms of scalability and performance. Distributed database technology has enabled higher resilience and improved flexibility when managing large volumes of data by leveraging on the collective storage and computational power of multiple database nodes. A centralized distributed database management system is used to interact with the data as if it were stored on a single node. This requires that creation, modification, and deletion of data are mirrored at every storage location so to maintain data consistency (Ray, 2009). Though blockchain is a specific distributed database technology of its own, there are more mature alternatives of distributed database storage that, if fulfilling the application requirements, are more suitable to use (Interviewee C, 2019). Oracle, Microsoft and IBM all offer support for distributed storage systems build on top of proven technologies and optimized software.

Question 2: What trade-offs must be considered when considering blockchain implementation?

Identifying blockchain as a potential candidate is the first step in the decision making process. The nature of the final conclusion is more complex. The choice of using a blockchain platform over another distributed technology will have several repercussions including systems governance, modularity, scalability and performance.

Stakeholders

The decision to adopt a blockchain solution befalls certain stakeholders in the list presented earlier in this chapter (c.f. Section 4.1.1). In particular, network managers and application developers play an integral role when assessing the use of blockchain.

- ***Network managers:*** A network manger is only needed when considering the adoption of a private and/or permissioned blockchain platform (i.e., a platform requiring a governance system). It is the network manager who defines the system and establish its boundaries. That said, they are responsible for making the final decision as to use blockchain or not. In the scenario of a public blockchain implementation, there will be no need for a network manager. Instead, regular managers issuing the application should be consider relevant stakeholders.
- ***Application developers:*** Developers on all levels will have to contribute with their expertise to properly evaluate the feasibility and potential configurations of a blockchain. They are likely to possess the necessary domain knowledge to make informed decisions, and are therefore essential to include in a technology selection discussion.

Model kinds and examples

Goal modelling can be a useful tool when evaluating the technology's potential. It aims at identifying early requirements which are complemented and modified as new problems and stakeholders are discovered. As architects learn more about the application domain in which the system will operate, it is easier to evaluate and compare different technological approaches. The first step in the construction of a goal model is the identification of stakeholders and their respective goals. Goals are a form of high level requirements that capture an actor's wishes, ideas, and opinions (Huysegoms et al., 2013). The often opposing goals and incompatible requirements expressed by different stakeholders call for proper consideration of technological alternatives and the accurate identification of trade-offs (Mylopoulos, 2006). The aim is to identify potential solutions and acquire the necessary knowledge required to make an informed decision. After selecting an approach, the succeeding actions should be clear. There are currently numerous goal modeling languages in use; two of the more popular ones are KAOS (Keep All Objectives Satisfied) and i*. Ignoring the syntactic differences, they can all be used to identify stakeholder goals, from which functional and non-functional requirements of the system can be derived.

Practitioners can break down high level objectives into smaller, more conceivable goals. By mapping stakeholder requirements, an architect could potentially answer the four questions of blockchain relevance listed under operations on views. If the identified goals include or allude to the assurance of trusted data sharing and update between mistrusted actors, data provenance, and the removal of intermediaries, a potential blockchain use case is identified. That said, only by identifying the application goals can the four questions be answered, and goal modeling effectively aids the architect in this endeavour.

An illustrative example related to the case study is provided in Figure 4.1. The high level objective

of *battery quality evaluation* is broken down into three subgoals including the development of an *evaluation algorithm*, the *removal of intermediaries*, and the availability and constant *retrieval of battery data*. By replacing intermediaries with public verifiability, the business process is enhanced in terms of processing times and cost. Also, a prerequisite for accurate battery evaluation is that available data include historical records (data provenance). This example gives an idea of how goal modelling can be used to identify and answer the questions in the subsequent list and Figure 4.3.

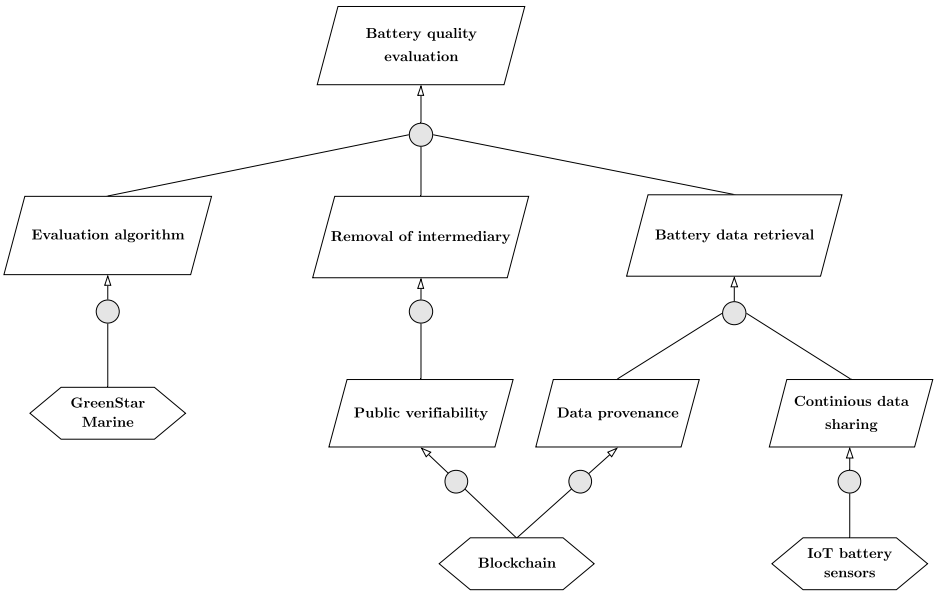

Figure 4.1: *A sample goal model using KAOS constructs.*

Another goal-model for technological selection using GRL (Goal-oriented Requirement Language) can be studied in Alwidian, Amyot, and Babin (2017). In this scenario, the technological alternatives are quantitatively assessed. The set of potential technologies are evaluated and parameterized according to their resource demand and ability to perform the tasks necessary to fulfill the derived system requirements held by the stakeholders.

In the Biswise workshop, the ***business model canvas*** (Osterwalder & Pigneur, 2009) was used to relieve inhibition among participants and to explore GreenStar Marine's proposed blockchain implementation. The resulting canvas of the use case's inner workings can be used to reason about blockchain's applicability. Figure 4.2 depicts canvas derived from the workshop.

The *key partners* and *customer relations* section indicates the network composition and creates awareness of who will participate in the platform creation. This aids architects answering the questions in Figure 4.3, determining the choice of blockchain platform. Similarly, *key Activities* indicate if the application requires functionality corresponding to the four prerequisites enumerated under the subsequent operations on views section. For example, the *verification of battery data* relate to the second point of verifying data between multiple entities. The *value proposition* section, in this case including transparency, trust and accurate resale value creates awareness of the application's purposes and could all be facilitated by a blockchain. Although being a rather high level management tool, the workshop showed that the business model canvas is a good way to map the system's stakeholders,

requirements, and their interactions. By careful analysis of the business model itself, one would be able to better analyze how well blockchain could be applicable to a particular use case.

Figure 4.2: *Business model canvas.*

Operations on views

Here follows two guidelines on how stakeholders should reason when addressing each of the two challenges presented in this viewpoint. These recommendations will have to be analyzed in the light of the business model canvas and constructed goal models.

Guideline 1: The properties of blockchain have to be evaluated against the application objectives and the environment in which it will be applied.

A proper evaluation of a potential blockchain should include a thorough investigation of various aspects. To start with, the following four prerequisites must coincide for a blockchain to even be relevant:

1. Is there a need for storing shared data between mistrusted entities?
2. Are multiple entities updating and verifying the state of this data?
3. Is subsequent activity governed by the sequence of previous transactions?
4. Is there no need of (or simply does not exist) a trusted intermediary?

Interviewee G (2019), a technology advisor, has noticed how CIOs and CTOs often start by asking themselves *how* rather than *why* they should deploy a blockchain solution. Asking why a blockchain implementation is feasible – in accordance with the list above – will avoid future complications and

can be facilitated by the models described above. For example, van Lamsweerde (2001) illustrates how new goals can be derived and visualized using KAOS by questioning the reasons behind already existing goals.

A key aspect to consider is whether there is *mistrust* between consortium participants, or as Interviewee O (2019) puts it:

> *" The key need for blockchain arises when nodes or data must cross organizational trust boundaries. "*
>
> *– Interviewee O (2019)*

These organizational trust boundaries do not have to be between separate corporations but can also be between departments within a single firm. According to Interviewee A (2019), a blockchain should be considered only if there is a need to eliminate central points of failure or authority. If so, blockchain can bring unquestionable benefits. By moving the responsibility of trusted mediators onto the blockchain itself, processes can be automated and accelerated while reducing costs (Rodrigues et al., 2018). Centralized solutions trade decentralization for speed, whereas blockchain solutions provide decentralized control and data validity at the expense of efficiency. Different guidelines have been proposed to help asses the applicability of blockchain in different scenarios. Wüst and Gervais (2018) has proposed an acknowledged set of decision rules for assessing whether blockchain is relevant and what type of platform configuration would be most appropriate. Another, well-cited list of conditions are proposed by Greenspan (2015). Though approaching the problem somewhat differently, both are based on the presence of mistrust and the inefficiencies of a trusted intermediary.

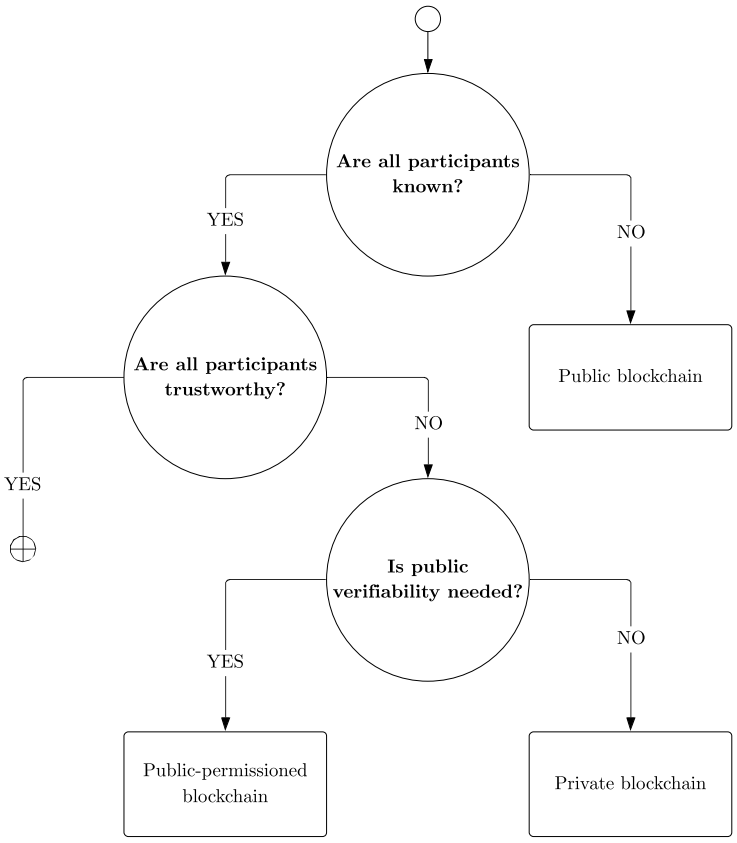

Figure 4.3: *A flowchart for blockchain platform selection (extracted from Wüst and Gervais (2018)).*

If the four aforementioned conditions are satisfied, the next part involve the choice of an appropriate blockchain platform. Wüst and Gervais (2018) makes a distinction between public and private blockchains and their differences regarding permissions. The choice is partly determined by the extent to which all participants are known within the network and whether public verification is required. Figure 4.3 shows a flowchart, that can be used as an initial guide for selecting an appropriate blockchain technology. The first question to be asked is whether all participants are known to each other? If the network constitutes of peers unknown to each other, a public blockchain may be appropriate. If the network is a small-scale, tightly linked supply chain with entrusted manufacturers, a blockchain may prove redundant. However, if all participants are known, but not all of them are trustworthy enough, one must ask whether public verifiability is needed? If the network is established to increase transparency within a restricted scope, a private blockchain would be appropriate. Otherwise, adopting a public-permissioned blockchain would be the recommended approach. The illustration in Figure 4.3 is an initial guideline; the specific use case will have to be studied in more detail in terms of performance and usability before a conclusion can be made.

Furthermore, Rodrigues et al. (2018) provides a somewhat ambiguous outlook when discussing private blockchains, suggesting that a traditional database would be preferable before arguing why a private blockchain may improve transparency and auditability. The decision can once again be related back to the level of mistrust. Complete trust may never be achieved, but would the increased transparency of a blockchain outbalance the reduction in efficiency? The private blockchain discussion touches upon one of the most pronounced disagreements within the blockchain community: that of ideology. Some seem to be of the opinion that private blockchains oppose the fundamental ideology of the technology and that many such application would benefit from applying more mature distributed database technologies instead. Considering the definition originally proposed by Nakamoto (2008), such a view would make sense. Naturally, there exists an opposition that argues for why the technology has to be modified so to fulfill the requirements governing enterprise relations.

> *" Private blockchains might seem to oppose the idea of decentralized technology at first, but my views on this statement would be vice-versa. It allows faster transactions, better scalability, consensus efficiency, and compliance support, which are some of the factors that are musts for enterprises. So in my point of view, it is a big plus, even though it shifts from the original blockchain idea. By permissioned networks, you use blockchain for what it is good for and behind a curtain which protects it more than the public network. "*
>
> *– Interviewee P (2019)*

Though ideology is important, it is of higher relevance to choose a platform, whether private or public, that has the best prospects of fulfilling the purpose of one's application.

Guideline 2: Properly evaluate the trade-off between performance/privacy and decentralization/resilience.

After having raised the questions about blockchain relevance and applicability, the use case requirements will need to be evaluated against a set of system attributes. It is inevitable that some attributes

benefit from the implementation of blockchain whilst some are impaired. Rodrigues et al. (2018) acknowledge a set of four categories:

Performance: Traditional databases perform significantly better than most blockchain solutions. That is because of the additional work enforced by the consensus protocol including transaction accumulation, block pervasion, and chain validation.

Privacy: Although blockchains are secure, the privacy concern is frequently highlighted, as can be seen in Table 4.3. Despite the promise of anonymization, statistical analyses of transactions over public blockchains have been proven successful in identifying patterns of user activities (Fabiano, 2017). Furthermore, in permissioned and private blockchains, the identity of each participant is known within the network. In use cases requiring transaction transparency and auditability, blockchain may prove appropriate. Ensuring privacy, however, is not what the technology was designed to do. Despite rigorous cryptography abilities that increases security, privacy is more easily obtained with a centralized database.

Distribution of control: The distribution of control relate to disintermediation and the previous choice between a public and private blockchain. In private and permissioned networks there has to be at least one network administrator that accepts new participants onto the network and distributes appropriate permission rights. Nevertheless, the network activity and the verification of new blocks are still decentralized, unburden any third-party mediator.

System resilience: The nature of blockchain and the fact that the shared ledger is replicated among all network nodes ensures high system resilience. This makes it harder to tamper with compared to a centralized database.

Conclusively, if application performance and privacy are prioritized system requirements a blockchain solution may not be appropriate. On the contrary, if distribution of control and system resilience are properties valued over performance and privacy, blockchain should be included onto the list of technological candidates. Evaluating this trade-off together with a careful analysis of use case requirements following the checklist and decision tree above, architects will be able to make an informed decision of blockchain relevance.

Correspondence rules

This viewpoint corresponds to the business architecture viewpoint in the TOGAF framework (see Desfray and Raymond (2014)) and the context viewpoint constructed by Rozanski and Woods (2005). These viewpoints are used to define the organizational context and the system's coexistence with its environment, similar to this viewpoint. This viewpoint distinguish itself, however, from the others by focusing on blockchain's *appropriateness* within the application domain.

4.2.2 Ecosystem viewpoint

> ***The ecosystem viewpoint addresses how an architect's decisions and the development of smart contracts are affected by the blockchain ecosystem and current levels of relevant competence.***

The underlying technology is currently the centre of attention among blockchain practitioners. The release of Nakamoto's (2008) whitepaper, describing the inner workings of Bitcoin, triggered a vast interest in its fundamental operations. In resent years, an elevated technological fascination is evident in the large number of new blockchain platforms, tokens, and technological solutions. The blockchain community is still dominated by academia and computer scientists trying to improve the underlying technology. Entailed by the distributed nature of blockchain itself, there is a need to spread the word and instill the concept of blockchain in more industry domains. This is imperative for the success of the technology. A blockchain network can be compared to an *ecosystem* of multiple players, working together for the purpose of more secure and efficient communication that will improve individual processes by leveraging on shared network effects. This perspective derives the need of an established network. The benefits of blockchain are contingent upon a network of actors; actors who first will have to agree on the technological use. The establishment of such a network will require much effort, time and cost (Davies & Likens, 2019). The establishment of a business network is more of a managerial concern. However, this is in turn dependent on sufficient blockchain expertise and that the competencies and responsibilities are properly identified and mapped within the ecosystem. The architecture is a prerequisite for the ecosystem. Simultaneously, the ecosystem defines the structure of the architecture. Thereby, the ecosystem and the architecture are closely linked together and it is essential that both of them are defined correctly from the beginning. A blockchain system enhances the network interactions, whether it facilitates financial transactions, supply chain traceability or IoT security. It becomes inseparable from the application domain in which it functions and can be classified as an E-program, as defined by Lehman (1980). The application domain will inevitably affect the architecture and are thus important to consider.

Key questions and concerns

Concern 1: There is a current lack of blockchain expertise, enforcing architects to specify and map appropriate competencies within the ecosystem.

Proficiency revolving blockchain technology and smart contract development are currently scarce. The development of a functional blockchain application requires technical expertise in game theory, computational efficiency and distributed execution (Interviewee A, 2019). Both literature and interviews have showcased a distinct thirst for blockchain experts as well as the request for new software tools that could facilitate blockchain development (Interviewee D, 2019). Given that the architecture to some extent determine the required competencies, it befalls the architect to specify and map these capabilities within the ecosystem, advising general management on what it needed (Wieringa, van Eck, Steghuis, & Proper, 2009). This challenge is aggravated by the lack of current

standards on what skills to acquire as a developer.

Concern 2: There is a need to assist developers in implementing and verifying smart contract execution so to satisfy the high quality requirements.

One of the more challenging areas of blockchain development is the construction of smart contracts. Smart contracts have enabled blockchain platforms to perform more sophisticated tasks other than just recording transactions from source to receiver. These are autonomous procedures that are executed instantly and strictly by its specifications. Distributed execution and the innumerable input combinations necessitate developers taking new approaches to smart contract development and software validation. According to Interviewee P (2019), the smart contracts must be able to handle all possible scenarios which makes them very difficult to test and debug. Orcutt (2019) acknowledges vulnerabilities in current platforms and provides examples of destructive attacks that has been successfully performed on blockchain systems. A total amount of 3.6M Ether was illegally obtained during the infamous DAO hack as a result of a loophole identified in one of the smart contracts (Güçlütürk, 2018). One can imagine the consequences if this was not just monetary founds, but personal data.

Orcutt (2019) alerts blockchain developers about the plenitude of new bugs that have erupted in the wake of smart contracts. This is especially true for public blockchains where the contracts are publicly available for hackers to study, along with the limited opportunities for correction. Smart contract defects can to some extent be solved by deploying a remedial contract that restores security through its interactions with the old contract. Depending on the implementation specifics, smart contracts can also be frozen after deployment (Ramachandran & Krishnamachari, 2018). However these are only responsive solutions and do not address the problem preventively. Nikolic, Kolluri, Sergey, Saxena, and Hobor (2018) analyzed the magnitude of smart contract defects present on the Ethereum blockchain. Using a customized software, they were able to categorize and identify different types of smart contract flaws using iterative invocation. 34200 of 970898 contracts were identified as susceptible to attacks. Interviewee A (2019) further elaborates on how system performance is heavily dependent on the contracts' execution efficiency.

The intricate interaction between network peers is complex and smart contracts are required to consider the large range of possible scenarios. To ensure that the network functions properly, it is necessary to understand and translate ecosystem composition into smart contract logic. Only by knowing the ecosystem structure appropriate smart contracts can be defined to facilitate its interactions.

Stakeholders

The blockchain network is describe as an ecosystem of participants, which suggests that all stakeholders would be relevant to list here. However, this viewpoint is mostly concerned with identifying and mapping the blockchain expertise within the consortium. Therefore the three most relevant stakeholders would be *platform developers*, *network managers* and *application developers*.

Model kinds and examples

An architect could potentially find inspiration in the suggestion proposed by Potts (2014). She argues that ***actor-network theory (ANT) modelling*** could help architects to better understand the context in which the system will eventually function. An ANT model is a mapping of an ecosystem's actors, systems, and their relationships. It visualizes what kind of organizations, people, and technologies exist and how they interact. Potts (2013) argues that conceptualizing the network on a piece of paper or a whiteboard, connecting actors with lines (similar to a mind map diagram) will do initially. Potts (2014) further elaborates on how this technique could be successful in identifying the distribution of *competencies* and *expertise*. By knowing where different competencies reside in the network, architects can locate required knowledge and data. Wolfensberger, Drayer, and Volker (2014) further explains how ANT mapping can help network participants nurture the network activity by creating awareness of its composition. With an increased awareness of how network peers interact with the blockchain, the architecture can be better designed in terms of fulfilling system requirements, policies, and user interfaces (Potts, 2014).

Operations on views

Here follows a set of guidelines that will aid blockchain architects in their initial efforts to establish a blockchain network, and how to make the smart contract development process more accessible.

Guideline 1.1: Take part in and learn from established blockchain communities.

Considering the technological novelty, there is a need for architects to join forces and work together toward standards and principles that can guide future progress. There are alliances currently making progress in this field (e.g., the Ethereum Enterprise Alliance). Interviewee H (2019) explains how his team has regular contact with technological ambassadors and external advisers providing them with blockchain expertise by discussing current advancements and successful use cases.

Guideline 1.2: Start with a small scale application, encouraging experimentation and development of competencies.

Interviewee D (2019) recommends blockchain architects to start with a small scale application, encouraging experimentation and development of competencies, before aiming at something more exhaustive. Though it is heavily related to deployment, it is ill-advised to deploy a complete distributed solution from the very start. Rather, when developing a blockchain system, Interviewee I (2019) advice architects to decentralize the blockchain application progressively so to maintain full control during the initial development phase. This further simplifies development, testing, and modification.

Guideline 2: Sophisticated verification and testing tools should be used by the architect to aid developers in creating high quality smart contracts.

New verification and testing tools are required to speed up and improve the development of smart contracts (Bosu et al., 2018). Interviewee A (2019) emphasizes the importance of formal verification and proper code audits. Formal verification is assumed to play an important role in the endorsement

of smart contract code. It is difficult to define proper state machines because of the complex logic of smart contract execution. The contract is not executed row by row; rather, the state depends on all nodes that runs the code including external nodes that are not part of the transaction itself. A popular verification language is Dafny, which offers first order logic constructs that can be used to reason about and prove the correctness of programs (Leino, 2017). The language allows developers to construct pre- and post conditions, loop invariants and lemmas that help the verifier to complete the proofs ensuring program correctness. Obsidian, proposed by Coblenz (2017), is an initiative to create a more secure blockchain programming language, which compiler is aspired to enable the translation of smart contracts into Dafny verifiable code (Coblenz & Kanal, 2017). Sánchez (2018) has expended upon this idea and developed a programming framework that enable actors to prove the correctness of a smart contract through formal verification *before* executing and *without* requiring code access.

Correspondence rules

This viewpoint corresponds to the context viewpoint defined by Rozanski and Woods (2005). It addresses the often overlooked relationship between the system and the application domain itself. A blockchain system enhances the interaction between a network of peers, whether it facilitates financial transactions, supply chain traceability or IoT security. Thus, it is inseparable from the application domain in which it functions and thereby falls under the E-program classification, as defined by Lehman (1980). The application domain will inevitably affect the architectural requirements imposed on the system itself and are therefore essential to consider. Nevertheless, this viewpoint distinguishes itself from the above by focusing on how blockchain expertise can be *mapped* and *managed* within the consortium. The identification of application constraints and network standards to ensure system coherence require a clear definition of the consortium's disperse requirements, limitations and capabilities. Lastly, this viewpoint is also important to consider in parallel with the applicability viewpoint presented previously. A properly defined ecosystem is a prerequisite for any blockchain application.

4.2.3 Infrastructure viewpoint

> ***The infrastructure viewpoint describes the performance related aspects of different blockchain platforms and provides examples of novel techniques that can be used to improve scalability and IoT integration.***

Blockchain poses one of the most ambitious attempts to facilitate alternatives to centralized data storage. No blockchain implementation has yet fully challenged conventional storage solutions however. The essence to this backlash lies in blockchain's many performance issues that prevent the technology to facilitate proper scalability. Cryptocurrencies have gained popular attention; however, they are not at a sufficient level to compete with current FIAT currencies. The increased interest in blockchain technologies has resulted in disperse solutions that address different challenges. It is important to investigate different solutions to find an appropriate combination for a specific use case. This section will elaborate on some of the most apparent performance challenges of blockchain. It is designed to

give a thorough understanding on why these challenges arise, and to provide architects with examples on how these can be alleviated.

Key questions and concerns

The choice of blockchain platform will have distinct repercussions on all performance aspects. It is therefore paramount to make sufficient evaluations of different alternatives before settling for one particular blockchain and therein evaluate the relevant performance aspects.

Concern 1: As the blockchain network grows, security is improved but transaction throughput abated.

There is a clear linkage between network size and reduced performance. Increased network size imply a more sophisticated distribution of stored information, granting higher security (Interviewee A, 2019). Public blockchains host large networks of participating nodes, enabling high transaction security, but at a cost. Both Ethereum's and Bitcoin's completely public and distributed networks require the use of a robust PoW consensus algorithm which, due to its heavy computational requirements, will inevitably be too slow for scalable applications (Reyna et al., 2018). At present, no application that requires close to immediate responses will find public blockchians easy to work with. Naturally, IoT applications that rely on frequent uploads will not work sufficiently. Private and permissioned blockchains can use more efficient consensus protocols. They also permit the network manager to define and reduce the number of full nodes responsible for validating each block. Nevertheless, Zheng et al. (2017) points out the distribution of blocks over the network will still reduce transaction throughput and increase latency.

Concern 2: Higher demands on computational power has resulted in network centralization and reduced scalability.

The pressure on blockchain's scalability will improve as accumulated IoT data increases (Dennis, Owenson, & Aziz, 2016). It is already becoming an issue to have conventional users commit their computational resources to such a cause, leading to a reduction in decentralization. As the chain will continue to grow indefinitely, the data storage capacity for full and mining nodes will eventually be reached. Also, the larger the chain, the longer time it will take for new users to synchronize the ledger (Reyna et al., 2018). In the age of IoT, petabytes of storage will need to be managed ("Data is giving rise to a new economy", 2017). In its current state of storage restrictions, blockchains do not remotely support applications at this scale. Potential solutions to these issues are surfacing, but until they do so effectively, one must consider the possible data load of particular applications and evaluate against current storage restrictions.

A private chain enable the use of more efficient consensus protocols since there is already some trust between network participants. However, note that the private blockchains are (as discussed under applicability) more centralized by its nature. The fact that participation require a consent from a network administer makes it inevitably harder to grow the network (Interviewee A, 2019). Furthermore, permissioned blockchains in general, however, suffer from the increase in communication

overhead (Shafagh et al., 2017).

Concern 3: The integration between the blockchain and IoT devices needs proper customization to alleviate the latter's constrained capacity.

In several use cases, blockchain does not have the potential to live up to the performance requirements. This is often true in relation to IoT devices of various kinds. Some devices, require constant connection with large transaction volumes and others intermittent connectivity and low bandwidth. It is important to note that a blockchain does not exist in isolation; rather, it will be integrated and synchronized with other technologies. Ramachandran and Krishnamachari (2018) illustrates a layered IoT approach in which end-devices equipped with sensors and actuators communicate with the server backend through an edge-device. The authors emphasize that the blockchain integration must be guided by each layer's resource constraints. Presently, only the backend server layer has the computational capacity to maintain a blockchain.

Question 4: How well does blockchain live up to its promises of security?

As more devices are connected to the Internet, the challenge to improve and implement identity control mechanisms is growing. Today, identity fraud and data breeches happens regularly to a large extent ("How to think about data in 2019", 2018). Blockchain has the potential to bring something new to this problem but it is required rather quickly as the rapid IoT adoption very much indicate. It should be stressed that the fundamental concepts of blockchain with its consensus and cryptography address security in its very nature. It is important to highlight that blockchain will not ensure security across all stakeholders interacting with a system. This is manifested by the many hacks on cryptocurrency exchanges that has taken place recently (Orcutt, 2019). In a setting with IoT enabled devices acting as data producers, implementing blockchain alone will not suffice.

Though there are some concerns regarding infrastructure security and the threat from large mining pools, Interviewee A (2019) argues that attacks are most likely to occur at the system boundaries where the system is most exposed. Interviewee C (2019) and Interviewee H (2019) are of the same opinion, raising the concern of hacked edge devices. IoT devices make up a large portion of these units and are located at the edge of the blockchain network, writing and potentially reading its data. In terms of security and privacy, Interviewee A (2019) highlights that the IoT sphere is far behind the blockchain community. The blockchain application would render quite futile if the data stored or referenced to by the chain is of low quality. Interviewee G (2019) argues for how blockchain, IoT, big data and artificial intelligence will have to work in unison so to address security threats on all levels. Lastly, as discussed by Interviewee A (2019) and Interviewee C (2019), account accessibility with asymmetric key encryption poses another, almost paradoxical, threat. Cases on public chains have protruded where millions of dollars have been made inaccessible due to the demise of the "keeper of keys". If no form of backup exists, once the password of a key-holder is lost there is no way to access that data.

Stakeholders

The viewpoint is mainly concerned with performance related attributes and platform security. Is the responsibility of *network managers* to ensure that the *platform* and *application developers* develop a blockchain application that fulfills the required quality attributes.

Model kinds and examples

Deciding upon a proper blockchain platform is a balancing act. The extensive nature of blockchain networks, involving multiple participants, requires the use of quality assessment tools to find a good balance between different stakeholder requirements. The quality attributes of a blockchain system can be measured and assessed using different ***software quality models***. These are objective means of reassuring that the system is built according to stakeholder requirements (Samadhiya, Wang, & Chen, 2010). The models help architects to break down high level quality attributes to subattributes that can be measured and quantified.

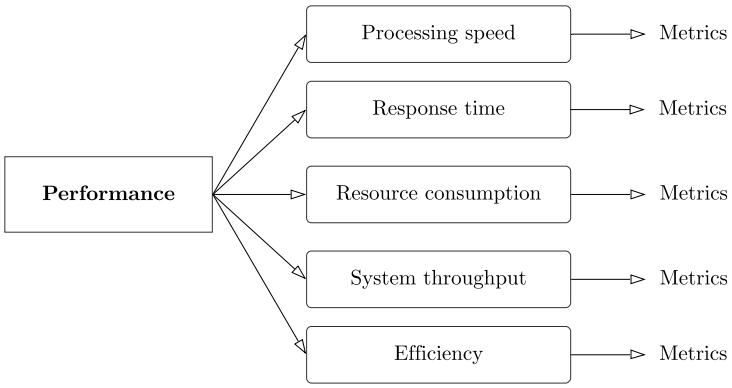

Figure 4.4: *A system quality breakdown.*

There exists different quality models that can be used in combination or separately. Samadhiya et al. (2010) studies the quality models of McCall, Boehm and Dromey, as well as the ISO 9126 and FURPS quality models. Of all these, only the FURPS quality model explicitly takes performance into account, supporting architects in prioritizing and measuring attributes such as response time, resource consumption and system throughput. On the contrary, however, it captures the least number of quality characteristics compared to the other alternatives. Both McCall's and Boehm's quality models organizes their quality attributes in high level perspectives that are broken down into more specific characteristics. Each addresses the general utility of the system and the end user requirements. Thus, these quality models share correspondence with the models presented in the end user viewpoint under Section 4.2.5. Conclusively, there's no one-size-fits-all approach to quality modeling. It is advisable to select the model or combination of models that best captures the system requirements.

An architect can map the domain of a software system by using ***structural models*** (Crispen & Stuckey, 1994). These models illustrate the application design and does so by representing the trade-offs between maintainability and performance, quality and efficiency. The structural models

help architects to break down the application into structural elements addressing separate problems (Batman, Howard, & Schelker, 1992). They can be used for exploring different designs and evaluate their separate performance in terms of certain quality attributes. The domain components, attributes, and inter-relationships are consistently enforced by the structural model, improving comprehension. It constitutes a set of limitations on the application's solution space (e.g. by elaborating on how entities relate, how data is exchanged, and how to organize the structure). All in all, the structural model should determine: (1) what system entities will be part of the design, (2) what constitute an entity, and (3) what entities communicate with each other (Crispen & Stuckey, 1994). According to Batman et al. (1992), structural models are usually compared and evaluated against a few dominant non-functional requirements. Structural modelling can be combined with software quality models to comprehensibly reason about the trade-offs between maintainability and performance, quality and efficiency. By analysing how data flows between system components and the way in which stakeholders interact, structural modelling can be useful in determining an appropriate blockchain configuration.

Blockchain networks are involved and difficult to analyze properly. To obtain a more comprehensive understanding and appreciation of how a decision will affect the performance and scalability of a system, it could be advisable to simulate the network. Rahman, Pakštas, and Wang (2009) describes and categorizes a set of popular ***modelling and simulation*** tools that can be used for network analysis. Initial efforts have been made to create blockchain specific simulators (e.g., see Gervais et al. (2016)). SimBlock is a recent initiative to create a more comprehensive blockchain simulator; it is currently under development and will be publicly available later this year (Aoki, Otsuki, Kaneko, Banno, & Shudo, 2019). The simulated values have shown to be consistent with those of real blockchains. The first version supports configuration on the network, node, and block level. These include the network bandwidth, the number of nodes and their behavior, the block size and the generation interval etc. It can also be configured to act in accordance with different consensus protocols. There also exists *testnets* that allow architects to run modelled application prototypes on simulated blockchain networks. There exists both public and private testnets for the Ethereum network, which allow system engineers to test system properties and smart contract execution (Angrish, Craver, Hasan, & Starly, 2018). Ganache CLI is an emulation tool that enables the creation of a simulated Ethereum network on the hosted computer alone. It allows developers to set up different accounts and perform test calls to the blockchain (Mohanty, 2018).

Operations on views

Guideline 1: Look into hardware cryptography to improve performance and privacy.

Cryptographic software can be replaced by designed hardware that can perform the task more efficiently (IBM Knowledge Center, 2019). Reyna et al. (2018) argue for how advancements in this area would be able to improve privacy and performance in applications involving constrained IoT devices. The most secure cryptographic processes of today are computationally heavy and can not be feasibly integrated into restricted sensors and smaller devices. This cryptographic processing is instead placed onto gateways and other networks that have more sufficient computational power. Though further research is still required, initial initiatives have showed promising performance improvements

in terms of consensus efficiency (Vukolić, 2016).

Guideline 2: Apply blockchain storage optimization techniques to challenge memory restrictions and scalability.

Computer scientist are currently trying to address the storage restrictions of blockchain. First of all, implementers should refrain form storing too much data onto the blockchain in the first place. According to Interviewee A (2019) data should be stored off-chain while letting the blockchain manage the data verification process alone. Interviewee A (2019) exemplifies this with distributed storage technologies such as IPFS (Interplanetary File System) and Swarm that are able to support an off-chain solution built on hashes. The data would remain equally verifiable and the chain would only have to store a 32-byte hash independently of the data size. This idea is further supported by Interviewee H (2019) and Interviewee D (2019) who believe in a hybrid solution that plays to blockchain's strengths, while letting traditional databases remain responsible for the actual data storage.

Bruce (2017) proposes a novel cryptocurrency scheme that involve deletion of old transactions to free up storage space. The technique prunes away old transactions while an account tree is updated to hold the balance of all addressees. The blockchain remains immutable and transparent by leaving the block headers intact. A *mini-blockchain* is a condensed versions of the fully fledged blockchain. By storing only the current state of every user, the mini-blockchain's size is unaffected by the number of transactions made in the network. Only when a new user joins the blockchain will the size increase. Fernández-Carmés and Fraga-Lamas (2018) elaborates on this idea and argue for its optimal use case within IoT applications, since the devices can interact directly with the chain without having to store too much information locally.

Guideline 3: Reallocate computationally heavy task from IoT devices to more capable network nodes.

Reyna et al. (2018) elaborates on how fog and edge computing will play an important role in unburden constrained IoT devices and the blockchain itself from computationally heavy work. As discussed in the Biswise workshop, edge computing could be used to address the well known concept of garbage in, garbage out. Information from IoT devices could be analyzed and pruned by edge computers so that the data reaching the application core are of high quality and free from noise, ready to be analyzed effectively.

Indeed, according to Interviewee A (2019), nodes with the sole purpose of *sending* transactions can be lightweight nodes. This is how most wallets are currently implemented and require a limited amount of computational and storage resources. It would therefore be an appropriate approach to establish an IoT device connection with the blockchain. This line of thought is expanded on by VerSum, a scheme invented by van den Hooff, Kaashoek, and Zeldovich (2014). It allows lightweight clients to forward heavy data computations to more sophisticated nodes. Thus, weak nodes would bypass computations beyond their capacity which would further increase scalability opportunities.

Guideline 4: Create a modular blockchain application that allow for flexible platform configurations.

Song et al. (2018) provide examples of layered blockchain solutions in which different blockchains are used on different network levels, customized to the capacity of each layer. The separate blockchains can be regularly synchronized by different verification processes. In line with Vukolić (2016), the consensus protocol has the largest impact on the platform's performance attributes. Gaetani et al. (2017) suggest a layered blockchain solution that combines a consensus efficient platform achieving high performance with a PoW platform guaranteeing data security. This provides an example of how an efficient blockchain, using a less sophisticated consensus algorithm, can be tailored for edge devices and securely backed up by a more robust blockchain located at the cloud level.

Consensus algorithms are indeed at the heart of many performance issues. A more straight forward approach that enables easy platform configuration is to select a modular blockchain technology. The most acknowledged platform of this type is Hyperledger. It is described as a private, modular blockchain platform that offers a high level of customization. This includes, of course, the selection of an appropriate consensus protocol (Androulaki et al., 2018).

Correspondence rules

This viewpoint corresponds to the legal viewpoint that in more detail elaborates on security and privacy concerns. The choice of blockchain platform is a pervasive one that will affect aspects in terms of all viewpoints. The viewpoint can be seen as a blockchain specific customization of the development viewpoints proposed by Kruchten (1995) and Rozanski and Woods (2005). These viewpoints are highly relevant from this infrastructure perspective as well, including their models: *component* and *package diagrams*. However, this viewpoint advocates the need for more sophisticates models to tackle the distributed nature of blockchain applications.

4.2.4 Legal viewpoint

> ***The legal viewpoint describes the system requirements from a legal point of view and how personal data should be processed in accordance with GDPR and regulations alike.***

The blockchain technology is designed to infuse trust among peers. Numerous authors emphasizes the growing concern about privacy and the management of personal information (see Table 4.2). This problem gets even more pronounced in an IoT network. The creation of data is increasing at full tilt. Gartner (2017) forecasts a total of 20 billion smart things by 2020, excluding computers and mobile phones. As a consequence of this ubiquitous connectivity, the annual generation of data is predicted to reach 180 zettabytes in 2025 ("Data is giving rise to a new economy", 2017). A large extent of this extensive pool of data is directly or indirectly related to individual data subjects. Tech giants, such as Google, Facebook, Netflix and Amazon has made it their business to collect and analyze their customer's personality traits and appeals. However, recent data breaches directed toward multiple

notable corporations have increased the privacy concern among individuals, making data protection one of the most prominent subjects of debate among engineers, politicians, and legislators ("How to think about data in 2019", 2018). Fernández-Carmés and Fraga-Lamas (2018); Ferrag et al. (2018); Panarello et al. (2018); Reyna et al. (2018) argue convincingly for why blockchain has the potential to tackle security and privacy issues considering its immutability, transparency and auditability. The technology shows promising potential and the research has covered a lot of ground in recent years and continues to make convincing progress.

According to Harari (2011), the reason behind human superiority is our ability to effectively collaborate within larger groups. This has been made possible by our ability to imagine and share a fabricated truth. Government, courts and corporations can flourish and function properly simply because there exists a common belief in these institutions. The disintermediation of trusted third-parties forces practitioners to put their trust in the abstract network instead. Interviewee C (2019) supports this line of thought, arguing that the wider adoption of blockchain is contingent upon the participants' trust in the technology. Moreover, Interviewee C (2019) asserts that regulation will play an important role in establishing and facilitating this trust. The jurisdictional regulations concerning blockchain are, however, still scarce and legislation that do exist have raised concerns. According to Interviewee C (2019), the current jurisdiction has been extended so to manage traditional distributed databases but falls short on a few substantial aspects introduced with blockchain.

The succeeding viewpoint is not an exhaustive recitation of the many legal factors that will inevitably surface during a blockchain implementation; each application will have its own set of attributes that will require legal consideration. Rather, the viewpoint aims to portray the current regulatory landscape with a focus on GDPR and the general features that blockchain applications, especially those managing private data, will have to take into account.

Key questions and concerns

The legal challenges identified in articles and interviews all fall under GDPR compliance, assuming the application in question is managing personal data. This regulation came into force May 25th, 2018 and is an initiative to protect individual citizens' private data. It aims to increase individual protection and ensure the ethical use of such data. Before diving into the various legal aspects and how they affect blockchain technology, it is important to note that the regulation does not try to limit the exploitation of the blockchain technology. Lyons et al. (2019) stress that the blockchain itself is not the subject of this regulation; rather, it is the way in which it is used and implemented that matters. Still, some doubt has been cast upon the compatibility of blockchain and GDPR. Posadas and Dalmacio (2018) discusses the inadequacies of current IoT sensors, a majority of which do not fulfill the required security and privacy measures. In a blockchain-IoT application, compromised data will have a highly destructive impact on the network's trustworthiness. Also, the distribution of private data among multiple nodes abate the sense of privacy, since the data subjects do not know where their personal data is being kept. Although arguing for the many opportunities blockchain can bring the IoT ecosystem, Posadas and Dalmacio (2018) state that the increased complexity required to fortify complex IoT data increases the risk of personal data being reviled, due to the large number

of distributed access points.

Taking the long view, there is cause for optimism about compliance but the current beliefs are scattered. The present uncertainty results in a general reluctance toward putting personal data on the blockchain (Interviewee N, 2019).

Concern 1: Blockchain network responsibility needs to be defined and allocated among participants.

Blockchain applications will inevitably have to answer for how legal aspects are fulfilled. Each blockchain project is required to carefully investigate how legal demands placed upon its data use will affect the architectural design. With the broad definition of private data in mind, the chances of having to conform to the principles of GDPR are significant. System architects will have to make sure that personal data is stored, processed and managed accordingly (Lyons et al., 2019).

The GDPR was not formulated with blockchain in mind. Several technical aspects of the blockchain technology contradict the presumptions endorsing the regulation; the most significant being the decentralization and the dispensability of recognized data controllers. According to Posadas and Dalmacio (2018), the regulation assumes the presence of an official administrator in charge of governing data access and processing. This challenge is especially pronounced in the use of a public-permissionless blockchain platform that leaves out any form of centralized configuration, remaining completely distributed. Public blockchains have gained much interest in later years and remain largely unconstrained (except for the rules enforced by the technological infrastructure). Note the double-edged sword of such platforms: the lack of imposed censorship has mounted to an incitement of illicit behaviours (Reyna et al., 2018; Yin, Langenheldt, Harlev, Mukkamala, & Vatrapu, 2019).

The Council of European Union (2018a) defines the data controller as *"the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data"*. It is the data controller who is responsible for making sure that all involved parties abide by the regulation and is therefore crucial to identify as soon as possible. Violations can result in sanctions of 20 million euro or 4% of a company's global annual revenue (Lyons et al., 2019). There is a current debate on who should be held accountable for the data processing. Since a completely decentralized system cannot be held accountable itself, different suggestions are currently being discussed. These include the software developers, the network nodes, or even the individual data producers themselves (Interviewee N, 2019).

Concern 2: Innovative designs and techniques are necessary to make personal data completely anonymous.

The GDPR only pertain to personal data that is related to a natural individual. However, this data domain is extended beyond what many would assume to be private information, including data that indirectly or combined with other sources may lead to the procurement of personal data and the identification of an individual (Posadas & Dalmacio, 2018). If the data can be made completely anonymous, however, it will no longer fall under the rules laid out by the regulation. Nevertheless, the encryption and security measures built into today's blockchain systems render the

data pseudoanonymous which is not good enough according to present benchmarks (Interviewee N, 2019).

Two key components of the blockchain technology is the individual's public and private keys as well as the extensive use of hashing. Each transaction is hashed and stored in blocks that are chained together by the hash of the previous block. Despite the sophisticated logic behind these techniques, neither asymmetric key encryption nor hashing can guarantee complete anonymity (*Opinion 05/2014 on Anonymisation Techniques*, 2014). For example, despite the presence of one-way cryptographic hash functions, knowledge of the input format will allow the perpetrator to try all possible input values and comparing each result to the real hash.

It has also been proven that personal data can be linked or inferred by pattern analysis. Roulin et al. (2018) demonstrate how machine learning algorithms can be applied on historic transactions to classify device type and identify user activity. Smart IoT networks classify appliances so to make sure the network functions properly. The combined success of classifying the device and the identification of its owner offers malicious actors an edge. Fabiano (2017) takes an example in which the frequency of data sent from an IoT compatible fridge could be used to determine when the occupants of the house were vacant. Roulin et al. (2018) highlight that transactions alone are sufficient input to such hostile schemes, which do not require information of the actual data, possibly stored off-chain.

Concern 3: Rectification and manipulation of data must be possible despite the indelible nature of blockchain transactions.

Architects must ensure that data processing is properly exercised. Lyons et al. (2019) present the six pillars of data security that pervades the regulation and define the proper ways to process personal data: personal data should be (1) legitimately collected and (2) processed according to the intended purpose; (3) the data should be kept accurate and (4) selected solely on the basis of its relevance; (5) it should be stored securely (6) within the necessary time frame, after which it must be removed. The application must also support the individual's right to rectify, remove, access and monitor the use of their personal data. In a blockchain context, the two most apparent activities are that of rectification and deletion of data. Once data is saved onto the blockchain, that information will remain invariable. Engineers and regulators are currently pondering the matter of data removal and what measures can be considered sufficient for abiding by the regulation (Lyons et al., 2019).

GDPR is also concerned with the automated processing of personal data and therefore the execution of smart contracts. The data subject has the right to be notified whenever their personal data is processed automatically and should be allowed to intervene, even after the execution of a smart contract (*Solutions for a responsible use of the blockchain in the context of personal data*, 2018). This would erode some of the trust endowed by the otherwise unbending contracts.

Concern 4: The incorporation of legal aspects into smart contracts is complex and inadequately covered in regulatory law.

Interviewee Q (2019) regards the translation of physical contracts into code, that can be carried out without intervention, as the biggest challenge in terms of smart contract development. Despite

being a domain of wordy precision and accuracy, jurisdiction often require proper judgments. How can such soft skills be implemented into software? In legal disputes one is often faced with equivocal aspects (e.g., the level of scienter) that require a level of sophistication not yet reached by artificial intelligence and machine learning (Interviewee C, 2019). Though artificial intelligence has showed promising prospects of becoming an integral part of the judicial system, potentially delivering final verdicts, this is an implausible reality within the near future (Rizer & Watney, 2018).

Stakeholders

All stakeholders fall under the legal umbrella. The current judicial environment and the ongoing debate makes it difficult to study the impact the jurisdiction will have on each stakeholder group. The most prominent discussion is that of network liability. In public blockchains, who should be held accountable for the data processing over the network? The suggestions are varied and are aimed toward different network participants. This indetermination makes the legal effects diffuse.

Model kinds and examples

As can be derived from the discussion above, the legal viewpoint is mainly concerned with the processing of personal data. It is therefore advisable to make use of models that could simplify and visualize the flow of information. The decentralization of nodes causes the information to ramify throughout the network, making the data flow nonlinear. A ***data flow diagram*** helps conceptualizing how data is exchanged between system components and involved actors (Li & Chen, 2009). The model can be applied on a notional level and do not require a completely specified system. It would be possible to visualize how data flows within a company and where it crosses organizational boundaries. In a blockchain context, it would be relevant to study how the information flows between user applications, APIs and the ledger, as well as between different layers of the architecture. To further illuminate to whom, where, when and how data is shared in the application, additional interaction diagram such as ***sequence diagrams*** and ***collaboration diagrams*** could be applied. They enable closer investigation into the shared data and the actors involved, respectively (vad der Aalst & van Hee, 2002).

Being a distributed database technology, standard approaches applied on traditional databases can arguably be used to model blockchain systems. Except for being interested in how data is moving within and outside the application, the data format is of significant relevance. One of the most popular and widely used data modeling techniques in industry is the ***Unified Modeling Language Diagram*** (UML). Though mainly perceived as a tool for notating object-oriented software, it can be used for designing databases (Garcia-Molina, Ullman, & Widom, 2009). The UML concepts of classes, associations, sub-classes, aggregations and compositions are well suited for describing the format of the data stored on the blockchain.

According to Interviewee S (2019), data will be generated on four levels: A *battery cell* constitutes the smallest component of a battery. This is the most detailed data source and require close collaboration with the battery manufacturers. A *battery* is composed of several battery cells. From these units,

information like voltage and internal temperature can be obtained. The *on-board diagnostic adapter* generates generic battery data such as charging status and other diagnostic information. Then there will be various *background data* that can be used for estimation purposes and includes details about in what vehicles and environments the battery has been used previously. Data from these separate sources must be organized and stored appropriately before queried by the evaluation algorithm. Structured in this manner, it is easy to reason about and identify sources of information with personal associations.

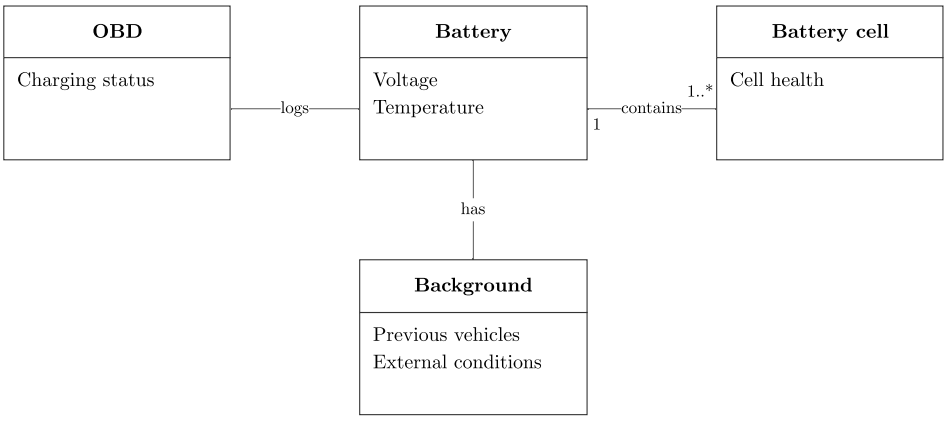

Figure 4.5: *A UML domain diagram for a battery data.*

Operations on views

Here follows a set of guidelines and recommendations that will aid involved stakeholders reasoning about legal challenges they face in blockchain implementations.

Guideline 1: Network accountability should be identified early on in the blockchain project and given to the actors uploading personal data.

Lyons et al. (2019) argues that the identification and distribution of network liabilities has to be determined depending on the use case, and is especially difficult in public-permissionless blockchains. These are often open source platforms being evolved by committed developers, contributing to a technology they believe in. It is therefore inadvisable to hold software developers accountable. The technology is not the subject of the regulation itself and along the same logic, the developers should not be held accountable for how their invention is used by others. With the sole aspiration of mining tokens and securing the open network, nodes should also not be held liable for the network activity. Although they exercise some influence by accepting protocol updates, the difficulty to coordinate their actions is problematic. Lyons et al. (2019), conclude it would be reasonable to make the actors *submitting* personal data onto the blockchain responsible for the processing. The distribution of liabilities will have to be customized for each use case and the decision is likely to be nonlinear. Each network participant including developers, miners, lightweight clients and individuals themselves will have a responsibility for upholding an honest and lawful network Interviewee C (2019).

Interviewee A (2019) initially suggested that the inconveniences of GDPR can be removed simply by establishing a public blockchain that cannot be held accountable. The problem is that the law *requires* the existence of a liable data controller. Despite the initial straggle, the judicial system will catch up and engineers will have to take that into consideration from the very start – GDPR being a featured example. Interviewee C (2019) strongly recommends taking contact with a legal counsellor at the beginning of a blockchain project.

Guideline 2.1: Adopt a privacy by design approach to infuse privacy into the system.

As pointed out by Fabiano (2017) and formalized in The Council of European Union (2018b), it is necessary to take a privacy by design approach when constructing a blockchain solution. It essentially advocates the use of state-of-the art technology to guarantee an individuals safety and privacy. Fabiano (2017) makes a clear distinction between security and privacy, suggesting that applying strong security measures is not enough to guarantee privacy. In many instances, personal identification is necessary to guarantee high security. Vast quantities of private data can be obtained by cross-referencing databases and directories containing personal information. However, such techniques can be used to track down malicious activities and prevent them from happening. This turns the two principles of privacy and security against each other. Is the integrity of the individual more important than the protection value of the larger society? Fabiano (2017) argues that one must not rule out the other. However, architects and system designers must start with the integrity of the individual. This will force privacy to be infused into the design, on top of which various security measures can be added.

Guideline 2.2: Refrain from storing personal data on the blockchain.

The discouragement of storing private data onto the blockchain is broadly acknowledged. The immutable nature of the chain contravene many of the rights assigned to data subjects by the GDPR legislation. Rather, it is advisable to store sensitive data off-chain, which would also allow for more efficient data management (Reyna et al., 2018). The blockchain should only be able to verify the existence of the external data without having to manage the storage itself. Besides alleviating the constrained storage capabilities, such an approach is better aligned with the current legal regulation. Several off-chain approaches have been proposed; the one most often referred to being that the blockchain should only store hashed values of the encrypted data (Ramachandran & Krishnamachari, 2018).

By only maintaining a reference to the external data stowed away using database technologies more suited for storing large quantities, the performance of the chain is improved in terms of latency and privacy (c.f. Section 4.2.3) (Panarello et al., 2018; Ramachandran & Krishnamachari, 2018). Though such an approach has proved promising in the light of GDPR, other solutions may prove even more compatible. Interviewee D (2019), describes an innovative approach applied in an initiative to make historical medical data more easily attainable, regardless of the current location of the patient. Instead of hashed values, the *location* of the external data is stored onto the blockchain. This solution enables even more flexibility in terms of rectification and erasure of data because of the loose coupling between the data stored on and off the chain. Any modification will leave the location address intact, whereas if hashed values were used, this would render the blockchain data invalid. Note that this

technique is inadequate in verifying the data. It is important to note that each use case will require a separate solution investigation and that the aforementioned proposals are considered in the light of the specific application.

Guideline 2.3: Anonymize data to the furthest extent to improve the integrity of participants.

Independent of blockchain technology, if data can be completely anonymized, the application is no longer subject to the GDPR regulation. The *Opinion 05/2014 on Anonymisation Techniques* (2014) paper elaborates on various anonymization techniques and sets the bar extraordinarily high in terms of what is accepted. There is neither a silver bullet that can guarantee complete anonymization, nor is it expected by legislators. The discussion on what anonymization techniques may be sufficient must be done on a case-by-case basis. Since blockchain's inception there has been a stream of innovative solutions to increase data integrity.

In public blockchains, the exposure of ones public key has shown to be a privacy risk. It is therefore important to make sure that adequate measures is taken to hide such trails. *Key regeneration* can be used to assign a different public key for each new transaction, making it harder to discover associated transactions and relevant patterns (Christidis & Devetsikiotis, 2016; Roulin et al., 2018). *Key aggregation* is a more sophisticated method that generates a public key from a set of individual keys. This discharges the participants from revealing their separate public keys. The composite key is sufficient for validating transactions. Several schemes offers support for key aggregation (see e.g., Maxwell, Poelstra, Seurin, and Wuille (2019)). A similar technique is *third-party indirection service* which moves the responsibility onto an intermediary who gathers multiple transactions before adding them onto the blockchain under its own public key, functioning solely as a cloak (Lyons et al., 2019). *Ring signatures* can be used to hide the true signer by letting a set of predefined group members endorse the transaction with their own signatures (Lyons et al., 2019).

Another popular technique is *mixing* which obfuscates the trail of a transaction by merging its input and output with other unrelated transactions. There are different techniques as to how this can be implemented but they all try making it more difficult to find transaction patterns (Zheng et al., 2017). Nevertheless, mixed transactions can still be vulnerable to statistical disclosure attacks and are often dependent on an intermediary performing the mixing (Fernández-Carmés & Fraga-Lamas, 2018).

Lastly, to bring down the possibility off device classification, Roulin et al. (2018) suggest that a random time delay is added to each transaction and that fragmented data-packages are combined into a single transaction. This will make it more difficult for an attacker to analyze transaction logs to the properties of different IoT appliances.

The aforementioned techniques aim to hide the identity of the transaction initiator. Here follows a set of acknowledged techniques that has proven effective in terms of hiding the data itself. A fascinating area of research is *zero-knowledge proofs*. They enable a data subject to prove his identity or the possession of other information without reveling it directly (Fernández-Carmés & Fraga-Lamas, 2018). Interviewee N (2019) evaluates the long-term prospects and are convinced of their high potential to

establish a safe environment for data processing. Moreover, Interviewee N (2019) takes the example of zero-knowledge Succinct Non-interactive Arguments of Knowledge (zk-SNARKs) which is used, among others, by Zerocash (Sasson et al., 2016). This cryptography allow participants to make transactions without publicly having to disclose the sender, receiver or the amount. Likewise, this cryptographic technique is aligned with the principles of the privacy-by-design methodology.

Homomorphic encryption is another promising method that allow external peers to process and validate encrypted data without accessing it directly (Moore, O'Neill, O'Sullivan, Doröz, & Sunar, 2016). As noted in Fernández-Carmés and Fraga-Lamas (2018) both zero-knowledge proofs and homomorphic encryption is computationally demanding. The use of customized hardware will make future adoptions more feasible (c.f. Section 4.2.3) (Moore et al., 2016).

Guideline 3: Make use of private and permissioned blockchains when managing private data.

When considering legal factors, some are prone to discuss the ideological features of the technology. Interviewee C (2019) describes a blockchain project as having three essential pillars that must be balanced for a successful implementation: (1) technology, (2) ideology, and (3) jurisdiction. As discussed under blockchain relevance, the blockchain community is divided around the technology's ideology. Interviewee C (2019) is of the opinion that a true blockchain application should be altogether distributed, leveraging on complete democracy of computing power. From that perspective, the technology exists but the ideology require certain modifications to be covered by today's laws; complete decentralization is impractical with current legislation. Interviewee C (2019) concludes that private blockchains, such as Hyperledger, have the largest potential for business applications in the near future. Hyperledger also offers the benefit of restricting the execution of smart contracts to a subset of permissioned nodes (Androulaki et al., 2018). One will still have to consider the individual's right to intervene in automated data processing; however, personal data is kept more private by limiting the data disclosure.

In general, there is a common conviction that private and permissioned blockchains are preferable when managing private data. However, this does not rule out the use of public platforms. Interviewee H (2019) comments on the advancements enabling blockchains to communicate between each other. This would allow private blockchains to manage sensitive information, while other interactions are managed by large-scale, public blockchains. The GDPR is not the end of public-permissionless blockchains. Once again, it is about the way the technology is used that matters. Due to the unsettled debates and the tightened jurisdiction it would be ill-advised, however, to use anything else but a private-permissioned blockchain for managing personal data. The ability to configure responsibilities and rules makes it easier to oversee the network activity. Stricter writing and reading permissions will increase privacy. Removing the threat exerted by large mining pools, an increased centralization would on the contrary impair data integrity (Reyna et al., 2018). Private blockchain do not require as computationally heavy consensus algorithms either, leaving room for stronger cryptographic algorithms. One way to establish some governance control in public blockchains is the introduction of a governance token (Interviewee I, 2019).

Guideline 4: Make use of dual integration to anchor the legal terms of traditional agreements into smart contracts.

Efforts to reconcile smart and legal contracts are being proposed. Dual integration is a hybrid solution to enforce lawful use of smart contracts and entail the attachment of the legal contract to the smart counterpart. The contracts are entwined by a duplex reference, linking them together. Christidis and Devetsikiotis (2016) describes how this integration is done in three steps:

1. Deploy the smart contract onto the network and save its address in the real contract.
2. Hash the real contract before storing it safely somewhere outside the blockchain.
3. Write a transaction containing the hash in its metadata and send it to the smart contract.

Correspondence rules

This viewpoint is strongly correlated to other viewpoints including security and privacy. It also relate to the information viewpoint defined by (Rozanski & Woods, 2005), which aims to describe what information is stored, and how it is processed and passed between system components. The models listed under this viewpoint include: *information life cycle models*, *information ownership models*, *information quality analysis*, and *metadata models*. Though focus is on privacy and anonymity, this viewpoint is extended to include the measures imposed by law and how the system can cope with the regulatory landscape under development.

4.2.5 End user viewpoint

> ***The end user viewpoint describes the interface between the user and the blockchain application, highlighting the user's perceived experience.***

This viewpoint is concerned with the interface between the end users and the blockchain technology itself. The end user will most likely never communicate directly with the chain itself, but through an application layer (e.g., through distributed apps and wallets). Nevertheless, the blockchain will influence, and to some extent determine, the way users interact with and view blockchain based applications.

Key questions and concerns

Blockchain's impact on user experience is a topic requiring further research. The following two questions must be considered by architects and will have to undermine the decision making process.

Concern 1: Blockchain applications need to visualize its data neatly, leveraging transparency while hiding technical aspects, making it more accessible for the everyday user.

The immutable property of blockchain together with the increased awareness of data privacy are likely to discourage individuals from connecting to and storing personal data on a blockchain (Interviewee N, 2019). In addition to this integrity problem, there is a current scarcity of research around human-blockchain interactions. Interviewee H (2019) points out that current efforts are allocated unevenly

and that the *"visual aspects of blockchain"* are often neglected. How should the data be processed and visualized? Pesale (2017) paints a rather gloomy picture while speculating on blockchain's impact on user experience. The new regulations and blockchain's distributed nature will require the end users to become more involved in the processing of their own data, an undertaking that many are not used to. People having a limited knowledge of blockchain may hesitate to accept a general permission to share their personal data, reducing the value of the chain. According to GDPR, individuals will either have to grant a general permission allowing the blockchain owner to forward their data, or they will face numerous requests of more specific data sharing. The added responsibility may result in people making decisions to avoid having to be included onto a blockchain.

Furthermore, the concept of a transaction being recorded and verified over a blockchain is not as intuitive as an ordinary message between two peers. This complexity may have a negative impact on user comprehension, especially for those completely new to blockchain. The technical language and the abstract concepts surrounding blockchain distantiate the technology from the everyday user (Interviewee H, 2019). The user experience should also try to leverage the principle concept of blockchain: verification and transparency. Designers will have to think carefully about how this value is converted into proper designs (Colina, 2017). Two other factors that may have an undermining effect on user experience is the aforementioned problems of transaction latency and the fact that each user on public platforms is hidden behind an illegible address.

Concern 2: Blockchain verified products and services are likely to become a necessity in the foreseeable future.

In contrast to the dispirited concern described above, there are practitioners and enthusiasts who believe that communities soon will consider blockchain as a necessity. Though few customers will understand the underlying technology, the increased awareness of data integrity and validity may make "blockchain stamped" products a requirement for a successful business. Blockchain brings extra value by ensuring customers of the product value. The ability to provide transparency and verified knowledge about products have a significant brand appeal in times of increased awareness (Interviewee G, 2019). Such initiatives have already been implemented in versatile industries such as wine production (see *Wine Blockchain* (2019)) and diamond forging (see *Do You Know Your Diamond?* (2019)). The use case pervading this thesis is another excellent example of how boat owners will be more prone to buy batteries which qualities have been verified by a blockchain. One of the largest concern for the end user, as mentioned during the Biswise workshop, is anticipated to be the additional cost of a blockchain subscription.

Stakeholders

- ***Application developers:*** will have to make sure that the end user experience is sufficiently satisfied. Even if the blockchain is part of the underlying technology, its repercussions will pervade to the outer edges of the system and affect the way users interact with the system. UX designers will have to work in close proximity to the blockchain engineers so to learn and coordinate their designs properly.

- ***Data consumers/producers:*** are concerned about how their data is managed and how the interaction with the application will be affected by the use of blockchain.

Model kinds and examples

A ***use case diagram*** illustrates in which ways users interact with the system and can be used as an initial tool for analyzing end user experience. Such a diagram could be a part of a ***user's manual*** which has shown to be an effective way to identify system user requirements, encouraging the requirement analyst to take an end user's perspective rather than focusing on the implementation specifics (Berry et al., 2004). A use case diagram or a user's manual solely addresses the interfaces through which the user will interact, leaving the developer free to decide upon an appropriate implementation strategy. A comprehensive user's manual has the benefit of providing the developers with a set of scenarios, from which test cases can be derived. The user's manual is not appropriate in all scenarios and will have to be complemented by other requirement models that studies other technical requirements in more depth (e.g., security, performance, and scalability). Nevertheless, it is a good tool for analyzing the system from the perspective of the user. Figure 4.6 illustrates a use case diagram for GreenStar Marine's battery trading platform.

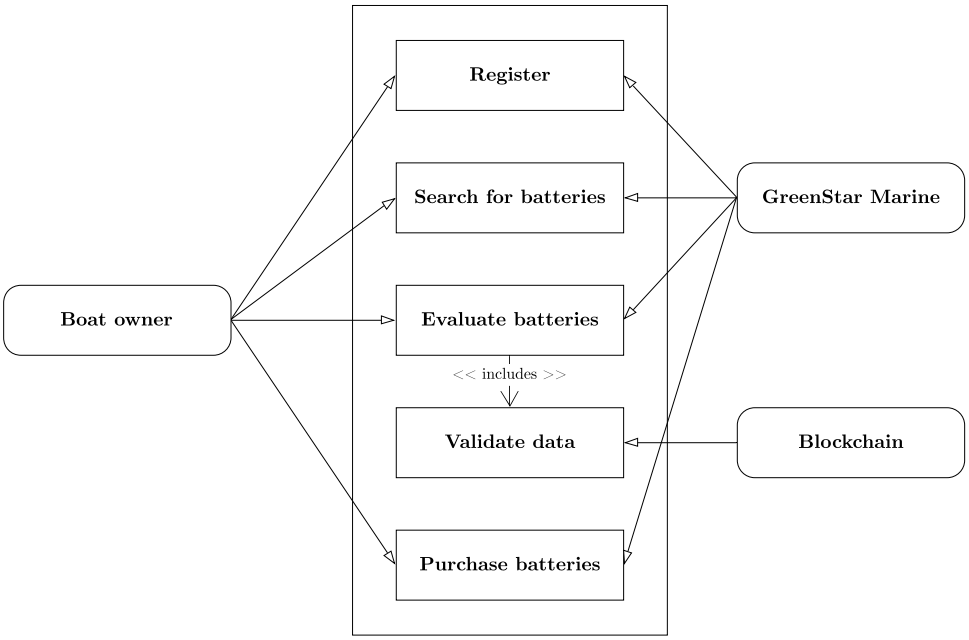

Figure 4.6: *A use case diagram for GreenStar Marine's battery trading platform.*

A ***system context diagram*** is a representation of the system and shows how it interact with its environment. As noted by Kossiakoff, Sweet, Seymour, and Biemer (2011b), *"the interactions of the system with its environment form the main substance of system requirements"*. It is therefore necessary to study the ways in which external entities interact with the system so that the blockchain can process these requests satisfactorily. To further study how the system process a user's request, a ***sequence diagram*** could be used.

Operations on views

Guideline 1: Be transparent and continue to experiment.

As an architect, it is important to remain transparent about how one approaches system design and in what ways this will affect the end user. A transparent development process in which the end user are involved and co-create the system could be valuable (Magnusson & Nilsson, 2014).

Guideline 2: Provide the end-user with a technology agnostic experience.

The scientific jargon and the complex features of blockchain should be hidden from a user's point of view so to avoid unnecessary confusion (Colina, 2017). To encourage this, Interviewee O (2019) recommends that the application should be developed and designed to be as agnostic as possible to the underlying blockchain protocol.

Correspondence rules

This view relate to the logical view of the 4+1 architectural view model by Kruchten (1995). Its main objective is to study the system's serviceability from an end user perspective and how its services are to be fulfilled accordingly. Models included in this viewpoint include *class diagrams*, *communication diagrams*, and *sequence diagrams* (Cooper-Bland, 2019). What distinguishes this viewpoint from the latter is the emphasize on data visualization. This viewpoint also relates to its legal counterpart since recently enhanced regulations will make the application usage and sharing of data more complex from the perspective of the end user.

4.3 Research question 3

A thorough evaluation of the derived framework is presented below with the ambition to assess its ability to aid blockchain implementations. Each viewpoint and its individual sections is scrutinized separately. The evaluation is divided into three parts: (1) *workshop evaluation*, (2) *interviewee evaluation*, and (3) *conference feedback*. To revive the motive behind the third research question, it is restated here:

> *By evaluating the framework derived from the preceding two questions, to what extent does the framework aid practitioners when reasoning about blockchain?*

A workshop, distributed over two sessions, was scheduled together with industry peers from GreenStar Marine and Biswise. During these sessions, quality was prioritized over quantity. As the framework was presented, participants were asked to ponder upon the recommendations and provide feedback along the way.

A survey (c.f. Appendix B) was sent out to all interviewees involved in the study, asking them to evaluate the concerns, models and guidelines provided in the respective viewpoints. It was intended to include a larger set of participants but did not reach a sufficient response rate. Only two interviewees completed the survey with a partial viewpoint focus. Nevertheless, the results provided valuable feedback and are therefore included in the report.

The conference evaluation followed a presentation held at the 2019 SAPSA conference in Göteborg. It was intended to involve people unrelated to the thesis, given that participants in the workshop and the survey evaluation where somewhat emotionally connected. The evaluation was based on a framework presentation before some 50 industry professionals from various industries spanning from automotive to health care. After the presentation, feedback was collected from a subset of the attendees.

4.3.1 Workshop evaluation

The framework was well received by the use case participants. According to Interviewee S (2019):

> *" The framework is of great value to companies considering implementing blockchain in their businesses ... It helps [architects] formalizing and structuring their often disordered ideas and thoughts. "*
>
> *– Interviewee S (2019)*

Interviewee J (2019), further elaborates on its significant depth that would make the viewpoints interesting, not only for companies novel to blockchain, but also for projects that have passed the initial phase of determining the technology's relevance. The technical guidelines are expected to be referenced in GreenStar Marine's efforts to customize their application in more detail.

Concerning the *challenging scenarios* listed in Table 4.4. According to Interviewee S (2019) and Interviewee J (2019), they are all important to consider and, more importantly, represent a comprehensive list of the challenges architects can be expected to face during blockchain implementations. Interviewee J (2019) points out that there are of course other challenges beyond the once presented and that the generic layout of the framework lessens the viewpoints' precision. Nonetheless, each use case can be expected to have its own set of specific challenges, and to cover them all in one framework is neither possible nor feasible. From that point of view, Interviewee S (2019) is of the opinion that the challenging scenarios are sufficient in covering the essential aspects of blockchain architectures that require immediate attention.

For companies novel to the technology, a large portion of the relevance surrounding blockchain implementation resides in the *applicability viewpoint* (c.f. Section 4.2.1). Interviewee S (2019) adds that considering why blockchain should be used is of course relevant, but that implementing the technology should not be a goal in itself. A company's intellectual property ought not be based on whether blockchain is implemented or not; rather, it should be considered as a potential tool that could help in achieving one or more of the value propositions stated in the business model canvas. This is important to keep in mind while considering the applicability viewpoint. Interviewee J (2019), along with the other Biswise workshop participants, appreciated the use of the business model canvas and the value derived from it. Along with the business model canvas, Interviewee S (2019) highlights the benefits of using goal models. These models have the ability to explicitly arrange and illustrate application goals. Furthermore, Interviewee S (2019) stresses their strength in unambiguously communicating goals to the relevant stakeholders.

The ordering of viewpoints was brought up after the presentation of the *ecosystem viewpoint*.

Interviewee S (2019) argues that defining the ecosystem surrounding the blockchain application should be done *before* constructing specific goal models for the applicability viewpoint. The composition of network participants will affect the way the system is implemented. This confirms the importance of considering the ecosystem viewpoint in parallel with the applicability viewpoint. A properly defined ecosystem is a prerequisite for any blockchain application. For this, the Actor-Network Theory (ANT) is according to Interviewee J (2019) expected to be an appreciated model, considering its capacity to map stakeholders and their relations. Interviewee S (2019) emphasized its value during the early stages of a blockchain platform construction, when the network of partners, suppliers, and customers are undefined. This problem was most prevalent during their initial planing as well. Interviewee S (2019) recommends, however, that the ANT model is complemented with other models such as a value stream model. This model maps transactions, shared directly and indirectly, among network participants. As recognized by Interviewee S (2019), by identifying closed loops of value streams that generate value to the business, network managers can make sure that every outgoing value stream is mapped to an incoming value stream. By structuring the system environment and its information flows, architects will be able to identify in what instances blockchain may be relevant and how it should be customized.

Moreover, Interviewee S (2019) and Interviewee J (2019) are positive toward competence development through small scale application development. This encourages a "fail fast, break things" approach and advocates developers to establish a complete data flow, generating testable values as early as possible in the development process. This flow can later be complemented and improved with additional functionality. In other words, setting up the fundamental blockchain functionality is crucial in the beginning. Any primitive modules can later be replaced by more sophisticated and accurate implementations.

The Infrastructure viewpoint (c.f. Section 4.2.3) was well perceived by both Interviewee S (2019) and Interviewee J (2019). The quality models should be considered as an evaluation tool to identify relevant quality attributes but must be complemented with other modeling techniques. The suggested use of network simulation to evaluate performance parameters were perceived as an intriguing strategy. Interviewee S (2019) did not, however, evaluate it to be a feasible option. Instead, an agile approach is recommended in which an early version is developed and later improved. The effort required to create a simulation model could be better spent on developing the real system itself. The viewpoint's technical guidelines regarding blockchain platform selection and IoT integration are expected to be a valuable support in their future endeavours.

Interviewee J (2019) and Interviewee S (2019) appreciate the inclusion of the *legal viewpoint* (c.f. Section 4.2.4). This was emphasized during previous meetings as well, and especially with respect to their selection of an appropriate blockchain platform.

> *" The legal aspects are of great importance and need consideration from the very start. "*
>
> *– Interviewee S (2019)*

The legal aspects of blockchain are not trivial and the guidelines provided in the legal viewpoint

are judged by both Interviewee S (2019) and Interviewee J (2019) to be of great value to GreenStar Marine. This viewpoint is especially important since the blockchain's legal complications are not yet sufficiently exploited. The viewpoint highlights aspects that engineers alone cannot tackle and, in line with the recommendations, Interviewee S (2019) has been in contact with a legal advisor early on.

The guidelines in the *end user* viewpoint addresses relevant concerns, but the guidelines provided are considered by Interviewee J (2019) to be relatively vague. The first one regarding experimentation and transparent development could be perceived as a platitude that does not bring sufficient value. The second one, that of technology agnosticism, is most relevant but could benefit from being anchored more strongly to blockchain. As for now, the recommendation is mainly generic.

Finally, Interviewee J (2019) calls attention to the academic format of framework that may hamper commercial use. The formal structure and the careful argumentation is considered to be of high academic quality. For industry to reap the benefits of the presented findings, however, the framework will have to be condensed. At present, it may prove difficult in terms of time and effort to digest the material and the theoretical reasoning. A concise version of the report would significantly increase the framework's industrial usability.

4.3.2 Survey evaluation

Six questions were asked in the light of each viewpoint. These can be studied in more detail in Appendix B. As mentioned in the introduction of this section, only two interviewees completed the survey with a partial viewpoint focus, covering only the *ecosystem* and *infrastructure* viewpoints. No quantitative or statistical analysis could therefore be performed on the data. The evaluation include feedback from Interviewee A (2019) and Interviewee G (2019). The received information were nonetheless valuable and are included. Future ambitions should aspire to complement these findings with more quantitative analyzes.

Interviewee A (2019) considers the *infrastructure viewpoint* to be very comprehensible, assigning it the highest score. Although giving the viewpoint structure a score of five, Interviewee A (2019) argues that it could be made more diverse by examining additional public blockchain models and decentralized storage technologies.

Interviewee A (2019) gives the challenge section a score of four. Essentially, blockchain is a protocol for proof and verification, and should not be considered as a storage protocol. Where the actual data is stored is according to Interviewee A (2019) not pertinent and continues by arguing that an off-chain approach is commonly used today. Furthermore, the loss of a access key does indeed prevents access to an account, but numerous countermeasures have been developed to mitigate this issue, for example: hardware wallets, multi-signature keys, and ring signatures. The infrastructure guidelines further recommend the use of lightweight nodes, which according to Interviewee A (2019) could be more emphasized in the framework. The rapid progress in blockchain suggests that the viewpoint is likely to change as performance related aspects are improved and new technological ideas invented.

The listed stakeholders and the operations on views receives a score of six and five, respectively. Interviewee A (2019) concludes the evaluation by advocating the addition of some extra models to complement the proposed quality models.

Interviewee G (2019) evaluates the *ecosystem viewpoint* and evaluated it to provide great value for blockchain implementers. Interviewee G (2019) assigns the highest score in response to all questions. However, no elaborating comments on the scoring is provided, which would have been valuable when evaluating the response. Nonetheless, the overall impression is interpreted as very positive.

4.3.3 The SAPSA conference evaluation

SAPSA is a non-profit organization with the objective to facilitate the exchange of knowledge, contacts, and experience around SAP products and services (*Vad är SAPSA*, 2019). An opportunity to share the concept of the battery management use case and to present the main pillars of the architecture framework was given together with Interviewee J (2019) and Niki Scaglione (director of SAP Cloud Platform in North EMEA). After the presentation, a handful industry representatives shared their feedback. Here follows a compiled summary of their thoughts.

They all appreciated the applicability concern and emphasized the difficulty of creating and establishing a blockchain network. They highlighted that support from larger actors such as IBM and SAP could help in linking participants together, while managing the technological underpinnings. From an industrial point of view, generalizing the technological concepts leads to a better comprehension and framework usability. Two representatives from Volvo and Perstorp Specialty Chemicals AB respectively also emphasized the educational value of the findings. Each called attention to the current lack of accessible material. One expressed threat was that of falling for the blockchain hype, owing to poor understanding of the concept itself. The framework's objective attempt to analyze blockchain relevance was therefore appreciated. Conclusively, from the procured feedback, the framework can be assumed to provide value to industry practitioners.

5. Artifact

The following chapter will abbreviate the architectural framework that has been derived from the study's three research questions. The aim is to condense the framework into a more easily digested format. Beyond well-known viewpoints traditionally used in architecture frameworks, five new viewpoints can be derived from the findings. The *applicability* viewpoint addresses the pronounced difficulty of finding an appropriate use case and using blockchain properly. The *ecosystem* viewpoint addresses the current scarcity of blockchain expertise and the means by which developers can be supported in their smart contract development. The *infrastructure* viewpoint highlights how the system's topology and platform parameters affect its performance, scalability and transaction throughput. The *legal* viewpoint helps architects taking informed decisions in the light of current legal developments with an accent on network liability and private data management. Lastly, the *end user* viewpoint attend to the identified prioritization of technological aspects over customer value.

5.1 Architecture framework

The fundamental premise of the architectural framework are the challenges identified in industry and literature, enumerated in Table 4.2 and Table 4.3 respectively. Five viewpoints have been derived from these challenges:

Table 5.1: *The applicability viewpoint.*

Definition:	Describes the application domain and its involved stakeholders, evaluating their properties and requirements in the light of what blockchain has to offer.
Concerns:	How can system architects properly evaluate the benefits of using blockchain technology over traditional distributed systems? What trade-offs must be considered when considering blockchain implementation?
Models:	Goal modeling and business model canvas.
Guidelines:	The properties of blockchain have to be evaluated against the application objectives and the environment in which it will be applied: ◦ *Is there a need for storing shared data between mistrusted entities?* ◦ *Are multiple entities updating and verifying the state of this data?* ◦ *Is subsequent activity governed by previous transactions?* ◦ *Is there no need of (or simply does not exist) a trusted intermediary?* If the application abide by these four prerequisites, then one of these three cases are advised to proceed with: Case 1: *Are there unknown participants?* → *Public blockchain* Case 2: *Is public verifiability needed?* → *Public-permissioned* Case 3: *Else* → *Private blockchain* Properly evaluate the trade-off between performance/privacy and decentralization/resilience: If decentralization and system resilience are valued over performance and privacy, blockchain should be considered.
Stakeholders:	Network managers and application developers.

Table 5.2: *The ecosystem viewpoint.*

Definition:	Describes the blockchain application's ecosystem and addresses the current scarcity of blockchain expertise and the means by which developers can be supported in their smart contract development.
Concerns:	There is a need to educate software developers in blockchain technologies. There is a need to assist developers in implementing and verifying smart contract execution so to satisfy the high quality requirements.
Models:	ANT modeling and value stream mapping.
Guidelines:	Take part in an learn from established blockchain communities. Start with a small scale application, encouraging experimentation and development of competencies. Sophisticated verification and testing tools should be used to aid developers in creating high quality smart contracts.
Stakeholders:	Platform developers, network managers and application developers.

Table 5.3: *The infrastructure viewpoint.*

Definition:	Describes the performance related aspects of different blockchain platforms and provides examples of novel techniques that can be used to improve scalability and IoT integration.
Concerns:	As the blockchain network grows, security is improved but transaction throughput abated. Higher demands on computational power has resulted in network centralization and reduced scalability. The integration between the blockchain and IoT devices needs proper customization to alleviate the latter's constrained capacity. How well does blockchain live up to its promises of security?
Models:	Software quality models and simulation models.
Guidelines:	Look into hardware cryptography to improve performance. Apply blockchain storage optimization techniques to challenge memory restrictions and scalability (e.g. off-chain storage, blockchain pruning and mini-blockchains). Reallocate computationally heavy task from IoT devices to more capable network nodes using fog and edge computing. Add lightweight client support with schemes like VerSum. Create a modular blockchain application that allow for flexible platform configurations.
Stakeholders:	Application developers, network managers, data producers and data consumers.

Table 5.4: *The legal viewpoint.*

Definition:	Describes the system requirements from a legal point of view and how personal data should be processed in accordance with GDPR and regulations alike.
Concerns:	Blockchain network responsibility needs to be defined and allocated among participants. Innovative designs and techniques are necessary to make personal data completely anonymous. Rectification and manipulation of data must be possible despite the indelible nature of blockchain transactions. The incorporation of legal aspects into smart contracts is complex and inadequately covered in regulatory law.
Models:	Data flow diagram, sequence diagram, collaboration diagram and UML diagram.
Guidelines:	Network accountability should be identified early on in the blockchain project and given to the actors uploading personal data. Adopt a privacy by design approach to infuse privacy into the system. Refrain from storing personal data on the blockchain. Anonymize data to the furthest extent to improve the integrity of participants using techniques such as key regeneration, key aggregation, third-party indirection service, zero-knowledge proofs and homomorphic encryption. Make use of private and permissioned blockchains when managing private data. Make use of dual integration to anchor the legal terms of traditional agreements into smart contracts: *1. Deploy the smart contract onto the network and save its address in the the real contract.* *2. Hash the real contract before storing it safely somewhere outside the blockchain.* *3. Write a transaction containing the hash in its metadata and send it to the smart contract.*
Stakeholders:	All stakeholders fall under the legal umbrella.

Table 5.5: *The end-user viewpoint.*

Definition:	Describes the interface between the user and the blockchain application, highlighting the user's perceived experience.
Concerns:	Blockchain applications need to visualize its data neatly, leveraging transparency while hiding technical aspects, making it more accessible for the everyday user. Blockchain verified products and services are likely to become a necessity in the foreseeable future.
Models:	Use case diagram, user's manual and system context diagram.
Guidelines:	Be transparent and continue to experiment. Provide the end-user with a technology agnostic experience.
Stakeholders:	Application developers, data producers and data consumers.

5.2 Industry relevance

It is difficult to quantify the framework's actual value. Nevertheless, it is reasonable to ask for both academic and practical merit. Industry peers, interviewees, and conference attendees have given an evaluation of the framework's potential ability to help architects in blockchain projects. The framework has been partially assessed in the light of the specific use case and qualitatively evaluated by practitioners. Further research is needed when implementing the framework on a broader perspective.

Even though the framework is not yet widely tested, it is evaluated to be comprehensive both from a general and specific perspective. Despite the wide range of blockchain challenges, the framework is assessed to capture the initial concerns architects will have to consider, without loosing focus on details. As the technological and legal environments are being developed, each viewpoint will have to be modified to fulfill future requirements. As for now, however, the framework satisfies the present needs. As for usability, a lot of energy is put into arguing for the framework's composition at the expense of accessibility. To encourage commercial use, the framework will have to be simplified and condensed.

6. Discussion

Interviewee J (2019) draws special attention to the often convoluted architectural documentation used within companies. An excessive amount of architectural views, models and guidelines causes inflexibility and impedes architectural modifications. In the moment of selecting an appropriate framework and defining relevant models and views, architects may run the risk of putting too much importance to details, feeling the need for constructing an excessive amount of documentation. Consequently, it is important to widen the perspective and be able to prioritize. An architectural description needs to consider all stakeholder requirements on different abstraction levels but without unnecessary bureaucracy. The proposed framework requires further refinement, e.g. through the addition of extra viewpoints from other frameworks. Nevertheless, the ambition to establish initial implementation practices is satisfied. It is anticipated to be of industrial value, based on evaluations provided by the case company and other industry representatives. The framework represents a comprehensive compilation of current blockchain research and industry investigations. It provides the case company with a toolkit that can be used through their blockchain implementation, from customizing their use case down to the selection of appropriate security measures.

The contribution of scholars and architectural pioneers have resulted in a plenitude of architectural frameworks and individual viewpoints available for system architects. Generally, the frameworks differ in scope and can be thought of falling along the continuum between *enterprise*, *system* and *software level* architectures. *Enterprise architecture frameworks* provide businesses with structured practices supporting the management of their competences and help them leverage on their data through appropriate application of technology. Two well known enterprise architectures are The Open Group Architecture Framework (Desfray & Raymond, 2014) and the Federal Enterprise Architecture (Johnson, 2016). Narrowing the scope down to the systems themselves, *system architecture frameworks* are used to define and manage the overall structure of a single system or a conglomerate of different subsystems working together. A few of the more well known frameworks in this category include the Department of Defense Architecture Framework and the British Ministry of Defence Architecture Framework (Kossiakoff, Sweet, Seymour, & Biemer, 2011a). The *software architecture frameworks* are usually applied to a single system and aid in structuring the definition of subcomponents. The most widely known framework of this type is the the 4+1 View Model of Architecture defined by Kruchten (1995).

The framework presented here falls close to this last category. It touches upon the blockchain system's security and privacy attributes along with other non-functional requirements including performance, IoT integration, and usability. However, it combines these aspects with a larger system perspective.

6.1 Implications to research

As technologies become more complex still, the task of developing, maintaining and properly using them are becoming increasingly difficult. As Vanderburg (2018) puts its, *"our intimate use of technology makes us more vulnerable to its misuse."* From a blockchain perspective, this is especially pronounced as the responsibility of upholding trust is shifted from intermediaries onto the network itself. Only a narrow range of applications are likely to derive value from blockchain inclusion. Currently, there is an insufficient emphasis on the importance of technology selections in the frameworks mentioned above, a void addressed by the inclusion of the applicability viewpoint (c.f. Section 4.2.1). Proper guidance and supportive frameworks help reduce the number of futile attempts, while increasing the quality of the successful ones.

The ecosystem viewpoint (Section 4.2.2) moves away from technical requirements toward higher level, consortium related concerns. It addresses the current scarcity of blockchain expertise and the means by which developers can be supported in their smart contract development. The viewpoint requires further refinement but touches upon a crucial aspect of blockchain implementations. It is expected to become even more prominent as connectivity and network complexity increases. For example, as stated by Harvey and Stanton (2014), system-of-systems involve cross-border interactions between actors and subsystems on different organizational and technical levels. The variability among these disparate actors and systems increases the complexity and introduces disputes of system intent. As such networks become every more interconnected, architects will need to maintain a sufficient understanding of how data is transferred within the consortium and who process what information. Only by knowing the ecosystem structure can an appropriate blockchain architecture be defined to facilitate its interactions.

The legal viewpoint (c.f. Section 4.2.4) correlates to any security and/or privacy aspect that an architect may define. The development view in the 4+1 framework, meanwhile, is predominantly concerned with the source code structure and the modularity of its subcomponents. It does,z however, appraise the software quality in terms of its reuse, portability and security (Kruchten, 1995). All the same, the fundamental security and privacy aspects of blockchain promotes the establishment of a separate viewpoint. Furthermore, software has become an integral part of human endeavour, pervading almost every part of life. Harari (2018) provides an important insight into technology development, arguing that *"the revolutions in biotech and infotech are made by engineers, entrepreneurs and scientists who are hardly aware of the political implications of their decisions, and who certainly don't represent anyone."* As technology becomes ever more entrenched in society, it continues to embrace new domains originally outside the technical sphere. The merge of technology and law is currently a topic receiving close review. The 4+1 framework does not sufficiently capture the system's legal context. On the contrary, the legal viewpoint helps architects taking informed decisions in the light of current legal developments. It is a step toward bridging the gap between engineers and legal practitioners, which is predicted to be necessary for a sustainable technological development.

The infrastructure viewpoint (c.f. Section 4.2.3) and the end user viewpoint (c.f. Section 4.2.5) have

similarities with, the process, physical and logical view of Kruchten's (1995) model. In his logical view example, Kruchten (1995) uses the Booch method to define and model the view. Decomposing the Booch approach into its five consecutive macro processes, one could argue that the end user viewpoint relates heavily to the first one, namely that of requirement identification (Booch et al., 2007). The interview and literature findings did not, however, extensively cover user experience but is still identified as a main concern of blockchain. Furthermore, the studied use case was still predominantly concerned with defining the basic blockchain infrastructure. Thus, the end user viewpoint requires further refinement. Still, it raises the identified disregard for user value of a blockchain implementation and provides the architect with a few recommended practices. The infrastructure viewpoint can be seen as a merge of Kruchten's (1995) process and physical view. It highlights how the system's topology and parameters of different blockchain platforms affect its performance, scalability and transaction throughput. As emphasized by Interviewee S (2019), blockchain is a tool used within a larger system to achieve certain system characteristics of a higher objective. Except for developing smart contracts, any blockchain implementation should be primarily concerned with selecting the most appropriate platform, which is emphasized in the viewpoint.

6.2 Validity concerns

As stated in Runeson and Höst (2008), there are four types of validity that affect the trustworthiness of the results: *construct*, *internal*, and *external* validity, along with the study's *reliability*. The study is examined below and evaluated in the light of these four categories.

Construct validity: To make sure that the operational measures used did comply with the investigation purpose and the three research questions, the interview template was reviewed by case company representatives and the academic supervisor. They were asked to read through the questions and identify any missing, ambiguous, or excessive content; acting as a pretest of the template. Case company representatives served as key informants and could, in the light of their blockchain expertise, provide valuable recommendations. This was done continuously, since the interview template was updated as the study proceeded and new directions were identified. The varying level of expertise among interviewees may have had an unfavourable effect on the construct validity. In an attempt to address this concern, the interview questions were sent out to all interviewees beforehand to give them time to prepare and raise any uncertainties. The findings obtained were thereafter discussed and reviewed during weekly evaluation sessions together with case company representatives.

Internal validity: Three types of triangulation were used to increase the study's internal validity. To improve the accuracy of the problem and solution investigation, the challenges were studied from different perspectives. *Data triangulation* was achieved by interviewing a broad range of industry participants including developers, project managers, technology advisers and legal counsellors. The fact that two students have taken part in this study imply the use of *observer triangulation*. Lastly, *Methodological triangulation* was applied by using interviews and workshops as two different data sources. Bearing in mind that each method is largely qualitative may, however, argue against the proper use of this type.

External validity: Parts of the result were derived from discussions revolving the use case. Thus, the result is in some measure biased toward the specific industry example. The limited number of interviewees and the use of a specific use case hampers the generalizability of the framework. Nevertheless, the disparate interviewees may increase the study's external validity. The study also describes in detail how conclusions are drawn from the interviews. This allows for recoverability and comparison of other findings that fellow researchers may find in different contexts.

Reliability: The scope of the interviews is clearly expressed and further depicted in Table 3.2, indicating the participants position and involvement. Data gathered from interviewees was assessed during the sessions with case company peers and further compared with information found in resent literature. The evaluation of the framework was heavily influenced by the feedback provided by the case company. This may reduce the reliability of the results. Requesting evaluative opinions from interviewees and conference participants helped ameliorate this concern.

7. Conclusion and Outlook

The architecture framework of this study has been constructed to facilitate blockchain implementation. It is anticipated to be of special value in the initial stages of architectural, facilitating informed decision making. The framework is structured around five viewpoints, each addressing a set of identified challenges.

This study has shown that the spectrum of applications that benefit from a blockchain implementation is reasonably narrow. The applicability viewpoint (c.f. Section 4.2.1) addresses the pronounced difficulty of finding an appropriate use case for the technology and using it wisely; there is a strong relationship between low performance and high system security. The infrastructure viewpoint (c.f. Section 4.2.3) highlights how the system's topology and platform configurations affect its performance, scalability and transaction throughput. Many are of the opinion that legal restrictions are blockchain's worst threat. The legal viewpoint (c.f. Section 4.2.4) helps architects taking informed decisions in the light of current legal developments with an accent on network liability and private data management. Though yet largely unattended by researchers and professionals, the user's perceived experience and expectations are likely to be affected by a blockchain integration. The end user viewpoint (c.f. Section 4.2.5) attend to the identified prioritization of technological aspects over customer value.

Conclusively, as the technology evolves, so will its challenges and corresponding solutions. The framework is not absolute and is likely to be refined and complemented with viewpoints from other frameworks. New models and guidelines will have to be appended to its structure so to address future advancements.

References

Alwidian, S., Amyot, D., & Babin, G. (2017, May). Evaluating the Potential of Technology in Justice Systems Using Goal Modeling. In *E-Technologies: Embracing the Internet of Things* (p. 185-202). Springer International Publishing. Retrieved from https://link.springer.com/chapter/10.1007%2F978-3-319-59041-7_11

Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., Caro, A. D., ... Yellick, J. (2018). Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. *CoRR*. Retrieved from http://arxiv.org/abs/1801.10228

Angrish, A., Craver, B., Hasan, M., & Starly, B. (2018). A Case Study for Blockchain in Manufacturing: "FabRec": A Prototype for Peer-to-Peer Network of Manufacturing Nodes. *Procedia Manufacturing*, *26*, 1180-1192. Retrieved from http://www.sciencedirect.com/science/article/pii/S2351978918308308

Aoki, Y., Otsuki, K., Kaneko, T., Banno, R., & Shudo, K. (2019). SimBlock: A Blockchain Network Simulator. *CoRR*. Retrieved from http://arxiv.org/abs/1901.09777

Baliga, A. (2017, April). *Understanding blockchain consensus models* (White Paper). Persistent Systems. Retrieved from https://www.persistent.com/whitepaper-understanding-blockchain-consensus-models

Batman, J., Howard, L., & Schelker, B. (1992). *An Introduction to Structural Models* (Report). Pittsburgh, Pennsylvania: Software Engineering Institute. Retrieved from https://apps.dtic.mil/dtic/tr/fulltext/u2/a268151.pdf

Berry, D. M., Daudjee, K., Dong, J., Fainchtein, I., Nelson, M. A., Nelson, T., & Ou, L. (2004, February). User's manual as a requirements specification: case studies. *Requirements Engineering*, *9*(1), 67-82. Retrieved from https://doi.org/10.1007/s00766-003-0181-1

Booch, G., Maksimchuk, R. A., Engle, M. W., Young, B. J., Conallen, J., & Houston, K. A. (2007). *Object-Oriented Analysis and Design with Applications* (3rd ed.). Addison-Wesley.

Bosu, A., Iqbal, A., Shahriyar, R., & Chakroborty, P. (2018, November). Understanding the Motivations, Challenges and Needs of Blockchain Software Developers: A Survey. *CoRR*. Retrieved from http://arxiv.org/abs/1811.04169

Brenzikofer, A. (2017). *Decentralized Trusted Timestamping* (White Paper). Zurich: Supercomputing Systems AG.

Bruce, J. (2017). *The Mini-Blockchain Scheme* [Technical Report]. Retrieved from https://cryptonite.info/files/mbc-scheme-rev3.pdf

Cadwalladr, C., & Graham-Harrison, E. (2018, March). Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach. *The Guardian*. Retrieved from https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election

Chalaemwongwan, N., & Kurutach, W. (2018, January). State of the art and challenges facing consensus protocols on blockchain. In *2018 International Conference on Information Networking* (p. 957-962). Retrieved from https://ieeexplore.ieee.org/document/8343266

Christidis, K., & Devetsikiotis, M. (2016, May). Blockchains and Smart Contracts for the Internet of Things. *IEEE Access*, *4*, 2292-2303. Retrieved from https://ieeexplore.ieee.org/document/7467408

Coblenz, M. (2017). Obsidian: A Safer Blockchain Programming Language. In *Proceedings of the 39th International Conference on Software Engineering Companion* (p. 97-99). Piscataway, New Jersey: IEEE Press. Retrieved from https://doi.org/10.1109/ICSE-C.2017.150

Coblenz, M., & Kanal, E. (2017). *Obsidian: A Safer Blockchain Programming Language* (Research Review). Pittsburgh, Pennsylvania: Carnegie Mellon University. Retrieved from https://resources.sei.cmu.edu/asset_files/Poster/2017_020_001_506446.pdf

Cole, A., & Gorman, D. (2019). *IBM Blockchain Foundation for Developers.* Retrieved from https://www.coursera.org/learn/ibm-blockchain-essentials-for-developers (Online course offered by IBM)

Colina, V. (2017, November). The user experience of blockchain. *Inside Design*. Retrieved from https://www.invisionapp.com/inside-design/user-experience-blockchain/

Conoscenti, M., Vetrò, A., & Martin, J. C. D. (2016, November). Blockchain for the Internet of Things: A systematic literature review. In *2016 IEEE/ACS 13th International Conference of Computer Systems and Applications* (p. 1-6). Retrieved from https://ieeexplore.ieee.org/document/7945805

Cooper-Bland, C. (2019). *Views and Viewpoints.* Retrieved from https://iasaglobal.org/itabok/capability-descriptions/views-and-viewpoints/

Crispen, R. G., & Stuckey, L. D. (1994). Structural Model: Architecture for Software Designers. In *Proceedings of the Conference on TRI-Ada* (p. 272-281). New York, New York: ACM. Retrieved from http://doi.acm.org/10.1145/197694.197729

Danzi, P., Kalør, A. E., & Popovski, C. S. P. (2018). Delay and Communication Tradeoffs for Blockchain Systems with Lightweight IoT Clients. *CoRR*. Retrieved from http://arxiv.org/abs/1807.07422

Data is giving rise to a new economy. (2017, May). *The Economist*. Retrieved from https://www.economist.com/briefing/2017/05/06/data-is-giving-rise-to-a-new-economy

Davies, S., & Likens, S. (2019). *What it takes to build your blockchain.* Retrieved from https://www.pwc.com/gx/en/issues/blockchain/blockchain-in-business/build-an-ecosystem.html

Dennis, R., Owenson, G., & Aziz, B. (2016, October). A Temporal Blockchain: A Formal Analysis. In *2016 International Conference on Collaboration Technologies and Systems (CTS)* (p. 430-437). Retrieved from https://ieeexplore.ieee.org/document/7871020

Desfray, P., & Raymond, G. (2014). TOGAF: General Presentation. In *Modeling Enterprise Architecture with TOGAF* (p. 1-24). Morgan Kaufmann. Retrieved from http://www.sciencedirect.com/science/article/pii/B978012419984200001X

Dieterich, V., Ivanovic, M., Meier, T., Zäpfel, S., Utz, M., & Sandner, P. (2017, November). Application of Blockchain Technology in the Manufacturing Industry. *FSBC Working Paper*. Retrieved from http://explore-ip.com/2017_Blockchain-Technology-in-Manufacturing.pdf

Doody, O., & Noonan, M. (2013). Preparing and conducting interviews to collect data. *Nurse Researcher*, *20*(5), 28-32. Retrieved from https://search-proquest-com.proxy.lib.chalmers

.se/docview/1443469489?accountid=10041

Do You Know Your Diamond? (2019). Retrieved from https://diamonds.everledger.io/

Fabiano, N. (2017, June). Internet of Things and Blockchain: Legal Issues and Privacy. The Challenge for a Privacy Standard. In *2017 IEEE International Conference on Internet of Things and IEEE Green Computing and Communications and IEEE Cyber, Physical and Social Computing and IEEE Smart Data* (p. 727-734). Retrieved from https://ieeexplore.ieee.org/document/8276831

Fernández-Carmés, T. M., & Fraga-Lamas, P. (2018). A Review on the Use of Blockchain for the Internet of Things. *IEEE Access*, *6*, 32979-33001. Retrieved from https://ieeexplore.ieee.org/document/8370027

Ferrag, M. A., Derdour, M., Mukherjee, M., Derhab, A., Maglaras, L. A., & Janicke, H. (2018, June). Blockchain Technologies for the Internet of Things: Research Issues and Challenges. *CoRR*. Retrieved from http://arxiv.org/abs/1806.09099

Gaetani, E., Aniello, L., Baldoni, R., Lombardi, F., Margheri, A., & Sassone, V. (2017, January). Blockchain-based database to ensure data integrity in cloud computing environments. In *Italian Conference on Cybersecurity*. Retrieved from https://eprints.soton.ac.uk/411996/

Gao, W., Hatcher, W. G., & Yu, W. (2018, July). A Survey of Blockchain: Techniques, Applications, and Challenges. In *2018 27th International Conference on Computer Communication and Networks* (p. 1-11). Retrieved from https://ieeexplore.ieee.org/document/8487348

Garcia-Molina, H., Ullman, J. D., & Widom, J. (2009). *Database Systems: The Complete Book* (2nd ed.). Prentice Hall.

Gartner. (2017). *Leading the IoT* (Tech. Rep.). Stamford, The United States: Author. Retrieved from https://www.gartner.com/imagesrv/books/iot/iotEbook_digital.pdf

Gaži, P., Kiayias, A., & Russell, A. (2018, June). Stake-Bleeding Attacks on Proof-of-Stake Blockchains. In *2018 Crypto Valley Conference on Blockchain Technology* (p. 85-92). Retrieved from https://ieeexplore.ieee.org/document/8525396

Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the Security and Performance of Proof of Work Blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (p. 3-16). Retrieved from http://doi.acm.org.proxy.lib.chalmers.se/10.1145/2976749.2978341

Greenspan, G. (2015). *Avoiding the pointless blockchain project.* Retrieved from https://www.multichain.com/blog/2015/11/avoiding-pointless-blockchain-project/

Gunter, C., Gunter, E., Jackson, M., & Zave, P. (2000). A reference model for requirements and specifications. *IEEE Software*, *17*(3), 37-43. Retrieved from https://ieeexplore.ieee.org/document/896248

Güçlütürk, O. G. (2018, August). The DAO Hack Explained: Unfortunate Take-off of Smart Contracts. *Medium*. Retrieved from https://medium.com/@ogucluturk/the-dao-hack-explained-unfortunate-take-off-of-smart-contracts-2bd8c8db3562

Halaburda, H. (2018, July). Economic and Business Dimensions Blockchain Revolution without the Blockchain? *Communications of the ACM*, *61*(7), 27-29. Retrieved from http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com.proxy.lib.chalmers.se/login.aspx?direct=true&db=buh&AN=130478716&site=ehost-live&scope=site

Harari, Y. N. (2011). *Sapiens: A Brief History of Humankind.* Harper.

Harari, Y. N. (2018). *21 Lessons for the 21st Century.* Spiegel & Grau.

Harvey, C., & Stanton, N. A. (2014). Safety in System-of-Systems: Ten key challenges. *Safety Science*, *70*, 358-366. Retrieved from http://www.sciencedirect.com/science/article/pii/S0925753514001684

He, Q., Guan, N., Lv, M., & Yi, W. (2018, June). On the Consensus Mechanisms of Blockchain/DLT for Internet of Things. In *2018 IEEE 13th International Symposium on Industrial Embedded Systems* (p. 1-10). Retrieved from https://ieeexplore.ieee.org/document/8442076

Herian, R. (2018). *Regulating Blockchain: Critical Perspectives in Law and Technology.* Routledge.

Hertig, A. (2019). How Ethereum Mining Works. *Coindesk*. Retrieved from https://www.coindesk.com/information/ethereum-mining-works

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design Science in Information Systems Research. *MIS Quarterly*, *28*(1), 75-105. Retrieved from http://www.jstor.org/stable/25148625

Hilliard, R. (2014). *Architecture viewpoint template for iso/iec/ieee 42010.* Retrieved from https://creativecommons.org/licenses/by/3.0/

How to think about data in 2019. (2018, December). *The Economist*. Retrieved from https://www.economist.com/leaders/2018/12/22/how-to-think-about-data-in-2019

Huysegoms, T., Snoeck, M., Dedene, G., Goderis, A., & Stumpe, F. (2013). Visualizing Variability Management in Requirements Engineering through Formal Concept Analysis. *Procedia Technology*, *9*, 189-199. Retrieved from http://www.sciencedirect.com/science/article/pii/S2212017313001758

Hyperledger Fabric. (2019). Retrieved from https://hyperledger-fabric.readthedocs.io/en/release-1.4/whatis.html#hyperledger-fabric

IBM Knowledge Center. (2019). *Hardware cryptography.* Retrieved from https://www.ibm.com/support/knowledgecenter/en/SSB23S_1.1.0.15/gtps7/hdcrypt.html

ISO/IEC/IEEE. (2011, December). *Systems and software engineering: Architecture description* (International Standard). Geneva, Switzerland: Author. Retrieved from http://cabibbo.dia.uniroma3.it/asw/altrui/iso-iec-ieee-42010-2011.pdf (ISO/IEC/IEEE 42010:2011(E))

Johnson, L. (2016). Security Component Fundamentals for Assessment. In *Security Controls Evaluation, Testing, and Assessment Handbook* (p. 531-627). Syngress. Retrieved from http://www.sciencedirect.com/science/article/pii/B9780128023242000117

Juels, A. (2006, February). RFID security and privacy: a research survey. *IEEE Journal on Selected Areas in Communications*, *24*(2), 381-394. Retrieved from https://ieeexplore.ieee.org/document/1589116

Kim, C. (2019). Code For Ethereum's Proof-of-Stake Blockchain to Be Finalized Next Month. *Coindesk*. Retrieved from https://www.coindesk.com/code-for-ethereums-proof-of-stake-blockchain-to-be-finalized-next-month

Kitchenham, B., & Charters, S. (2007, July). *Guidelines for performing Systematic Literature Reviews in Software Engineering* (EBSE Technical Report No. 2.3). The United Kingdom:

Keele University and Durham University. Retrieved from https://www.elsevier.com/__data/promis_misc/525444systematicreviewsguide.pdf

Kossiakoff, A., Sweet, W. N., Seymour, S. J., & Biemer, S. M. (2011a). Architecture Frameworks. In *Systems Engineering Principles and Practice* (2nd ed.). John Wiley & Sons. Retrieved from https://app.knovel.com/hotlink/khtml/id:kt011AW466/systems-engineering-principles/architecture-frameworks

Kossiakoff, A., Sweet, W. N., Seymour, S. J., & Biemer, S. M. (2011b). System Boundaries: The Context Diagram. In *Systems Engineering Principles and Practice* (2nd ed., p. 52-55). John Wiley & Sons. Retrieved from https://app.knovel.com/hotlink/khtml/id:kt011AVX12/systems-engineering-principles/system-boundaries-context

Kruchten, P. B. (1995). The 4+1 view model of architecture. *IEEE Software*, *12*(6). Retrieved from http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=edsgic&AN=edsgcl.18092619&site=eds-live&scope=site

Kumar, N. M., & Mallick, P. K. (2018). Blockchain technology for security issues and challenges in IoT. *Procedia Computer Science*, *132*, 1815-1823. Retrieved from http://www.sciencedirect.com/science/article/pii/S187705091830872X (International Conference on Computational Intelligence and Data Science)

Lehman, M. M. (1980, September). Programs, life cycles, and laws of software evolution. *Proceedings of the IEEE*, *68*(9), 1060-1076. Retrieved from https://ieeexplore.ieee.org/document/1456074

Leino, K. R. M. (2017, November). Accessible Software Verification with Dafny. *IEEE Software*, *34*(6), 94-97. Retrieved from https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8106874

Li, Q., & Chen, Y.-L. (2009). Data Flow Diagram. In *Modeling and Analysis of Enterprise and Information Systems: From Requirements to Realization* (p. 85-97). Springer. Retrieved from https://doi.org/10.1007/978-3-540-89556-5_4

Lisk. (2019). *Proof of stake.* Retrieved from https://lisk.io/academy/blockchain-basics/how-does-blockchain-work/proof-of-stake

Lyons, T., Courcelas, L., & Timsit, K. (2019, March). *Blockchain and the GDPR* (Thematic Report). Brussels, Belgium: EU Blockchain Observatory & Forum. Retrieved from https://www.eublockchainforum.eu/sites/default/files/reports/20181016_report_gdpr.pdf

Magnusson, J., & Nilsson, A. (2014). *Enterprise System Platforms: Transforming the Agenda* (1st ed.). Studentlitteratur AB.

Markets and Markets. (2019, March). *Blockchain IoT Market by Offering, Application, End User, and Geography: Global Forecast to 2024.* Retrieved from https://www.researchandmarkets.com/reports/4760815/blockchain-iot-market-by-offering-hardware#pos-0

Maxwell, G., Poelstra, A., Seurin, Y., & Wuille, P. (2019). Simple Schnorr multi-signatures with applications to Bitcoin. *Designs, Codes, and Cryptography*. Retrieved from http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=edselc&AN=edselc.2-52.0-85060972932&site=eds-live&scope=site

McLafferty, I. (2004). Focus group interviews as a data collecting strategy. *Journal of Advanced Nursing*, *48*(2), 187-194. Retrieved from https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2648.2004.03186.x

Meng, W., Tischhauser, E. W., Wang, Q., Wang, Y., & Han, J. (2018). When Intrusion Detection Meets Blockchain Technology: A Review. *IEEE Access*, *6*, 10179-10188. Retrieved from https://ieeexplore.ieee.org/abstract/document/8274922

Mining, B. (2018). Retrieved from https://www.bitcoinmining.com/bitcoin-mining-pools

Mittal, S., Tam, T. W., & Ko, C. (2018, June). *Internet of Things: The Pillar of Artificial Intelligence* (Sector Briefing). Asian Insights Office: DBS Group Research. Retrieved from https://www.dbs.com/aics/templatedata/article/generic/data/en/GR/062018/180625_insights_internet_of_things_the_pillar_of_artificial_intelligence.xml#

Mohanty, D. (2018). *Ethereum for Architects and Developers : With Case Studies and Code Samples in Solidity*. Apress. Retrieved from https://link.springer.com/book/10.1007%2F978-1-4842-4075-5

Moore, C., O'Neill, M., O'Sullivan, E., Doröz, Y., & Sunar, B. (2016, June). Practical homomorphic encryption: A survey. In *2014 IEEE International Symposium on Circuits and Systems* (p. 2792-2795). Retrieved from https://ieeexplore.ieee.org/document/6865753

Mylopoulos, J. (2006). Goal-Oriented Requirements Engineering, Part II. In *14th IEEE International Requirements Engineering Conference.* Retrieved from https://ieeexplore.ieee.org/document/1704043

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system.* Retrieved from http://bitcoin.org/bitcoin.pdf

Nikolic, I., Kolluri, A., Sergey, I., Saxena, P., & Hobor, A. (2018). Finding The Greedy, Prodigal, and Suicidal Contracts at Scale. *CoRR*. Retrieved from http://arxiv.org/abs/1802.06038

Opinion 05/2014 on Anonymisation Techniques (Tech. Rep.). (2014, April). Brussels, Belgium: Article 29 Working Party. Retrieved from https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2014/wp216_en.pdf

Orcutt, M. (2019, February). Once hailed as unhackable, blockchains are now getting hacked. *MIT Technology Review*. Retrieved from https://www.technologyreview.com/s/612974/once-hailed-as-unhackable-blockchains-are-now-getting-hacked/

Osterwalder, A., & Pigneur, Y. (2009). *Business Model Generation.* Self-published.

Palinkas, L. A., Horwitz, S. M., Green, C. A., Wisdom, J. P., Duan, N., & Hoagwood, K. (2015, September). Purposeful Sampling for Qualitative Data Collection and Analysis in Mixed Method Implementation Research. *Administration and Policy in Mental Health and Mental Health Services Research*, *42*(5), 533-544. Retrieved from https://doi.org/10.1007/s10488-013-0528-y

Panarello, A., Tapas, N., Merlino, G., Longo, F., & Puliafito, A. (2018, August). Blockchain and IoT Integration: A Systematic Survey. *Sensors*, *18*(8). Retrieved from http://www.mdpi.com/1424-8220/18/8/2575

Parker, L. (2015, May). *Number of Bitcoin Miners Far Higher Than Popular Estimates* (Technical Analysis). Zurich: Brave New Coin. Retrieved from https://bravenewcoin.com/insights/number-of-bitcoin-miners-far-higher-than-popular-estimates

Pesale, E. (2017, October). Blockchain technology may not be the best solution for GDPR compliance. *IDG Contributing Network*. Retrieved from https://www.csoonline.com/article/3232386/blockchain-technology-may-not-be-the-best-solution-for-gdpr-compliance.html

Posadas, J., & Dalmacio, V. (2018). The Internet of Things: The GDPR and the Blockchain May Be Incompatible. *Journal of Internet Law*, *21*(11). Retrieved from http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com.proxy.lib.chalmers.se/login.aspx?direct=true&db=buh&AN=131103583&site=ehost-live&scope=site

Potts, L. (2013). *Mapping Experiences with Actor Network Theory.* Retrieved from https://www.slideshare.net/LizaPotts/mapping-experiences-with-actor-network-theory

Potts, L. (2014). *Social Media in Disaster Response: How Experience Architects Can Build for Participation.* Routledge.

Rahman, M. A., Pakštas, A., & Wang, F. Z. (2009). Network modelling and simulation tools. *Simulation Modelling Practice and Theory*, *17*(6), 1011-1031. Retrieved from http://www.sciencedirect.com/science/article/pii/S1569190X09000197

Ramachandran, G. S., & Krishnamachari, B. (2018, May). Blockchain for the IoT: Opportunities and Challenges. *CoRR*. Retrieved from http://arxiv.org/abs/1805.02818

Ray, C. (2009). *Distributed Database Systems.* Pearson.

Reyna, A., Martín, C., Chen, J., Soler, E., & Díaz, M. (2018). On blockchain and its integration with IoT. Challenges and opportunities. *Future Generation Computer Systems*, *88*, 173-190. Retrieved from http://www.sciencedirect.com/science/article/pii/S0167739X17329205

Richardson, M. (2017, November). Ethereum vs Hyperledger. *Blockchain Training Alliance*. Retrieved from https://blockchaintrainingalliance.com/blogs/news/ethereum-vs-hyperledger

Rizer, A., & Watney, C. (2018). Artificial Intelligence Can Make Our Jail System More Efficient, Equitable, and Just. *Texas Review of Law & Politics*, *23*(1), 181-227. Retrieved from http://proxy.lib.chalmers.se/login?url=https://search-proquest-com.proxy.lib.chalmers.se/docview/2193091511?accountid=10041

Rodrigues, B., Bocek, T., & Stiller, B. (2018). The Use of Blockchains: Application-Driven Analysis of Applicability. In *Blockchain Technology: Platforms, Tools and Use Cases* (Vol. 111, p. 163-198). Elsevier. Retrieved from http://www.sciencedirect.com/science/article/pii/S006524581830024X

Roulin, C., Dorri, A., Jurdak, R., & Kanhere, S. (2018). On the Activity Privacy of Blockchain for IoT. *CoRR*. Retrieved from http://arxiv.org/abs/1812.08970

Rowling, J. K. (2005). *Harry Potter and the Half-Blood Prince.* Bloomsbury.

Roy, S., Ashaduzzaman, M., Hassan, M., & Chowdhury, A. R. (2018, Dec). Blockchain for IoT Security and Management: Current Prospects, Challenges and Future Directions. In *2018 5th International Conference on Networking, Systems and Security* (p. 1-9). Retrieved from https://ieeexplore.ieee.org/document/8631365

Rozanski, N., & Woods, E. (2005). *Software systems architecture: working with stakeholders using viewpoints and perspectives.* Addison-Wesley.

Runeson, P., & Höst, M. (2008, December). Guidelines for conducting and reporting case study research in software engineering. *Empirical Software Engineering*, *14*(2), 131. Retrieved from https://doi.org/10.1007/s10664-008-9102-8

Sadu, I. (2018, December). Auditing Blockchain. *Internal Auditor*, *75*(6), 17-19. Retrieved from http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=buh&AN=133555219&site=ehost-live&scope=site

Sagirlar, G., Carminati, B., Ferrari, E., Sheehan, J. D., & Ragnoli, E. (2018). Hybrid-IoT: Hybrid Blockchain Architecture for Internet of Things - PoW Sub-blockchains. *CoRR*. Retrieved from http://arxiv.org/abs/1804.03903

Salimitari, M., & Chatterjee, M. (2018). An Overview of Blockchain and Consensus Protocols for IoT Networks. *CoRR*. Retrieved from http://arxiv.org/abs/1809.05613

Samadhiya, D., Wang, S.-H., & Chen, D. (2010, October). Quality models: Role and value in software engineering. In *2010 2nd International Conference on Software Technology and Engineering* (Vol. 1, p. 320-324). Retrieved from https://ieeexplore.ieee.org/abstract/document/5608852

Sasson, E. B., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2016, May). Zerocash: Decentralized Anonymous Payments from Bitcoin. In *2014 IEEE Symposium on Security and Privacy* (p. 459-474). Retrieved from https://ieeexplore.ieee.org/abstract/document/6956581

Schutt, R. K. (2009). *Investigating the Social World: The Process and Practice of Research*. SAGE Publications Inc.

Shafagh, H., Burkhalter, L., Hithnawi, A., & Duquennoy, S. (2017). Towards Blockchain-based Auditable Storage and Sharing of IoT Data. *CoRR*. Retrieved from http://arxiv.org/abs/1705.08230

Solutions for a responsible use of the blockchain in the context of personal data (Tech. Rep.). (2018, November). Paris, France: CNIL. Retrieved from https://www.cnil.fr/sites/default/files/atoms/files/blockchain.pdf

Song, J. C., Demir, M. A., Prevost, J. J., & Rad, P. (2018, June). Blockchain Design for Trusted Decentralized IoT Networks. In *2018 13th Annual Conference on System of Systems Engineering* (p. 169-174). Retrieved from https://ieeexplore.ieee.org/document/8428720

Sánchez, D. C. (2018, August). Raziel: Private and Verifiable Smart Contracts on Blockchains. *CoRR*. Retrieved from http://arxiv.org/abs/1807.09484

Tapscott, D., Tapscott, A., & Kirkland, R. (2016, May). How blockchains could change the world. *McKinsey & Company*. Retrieved from https://www.mckinsey.com/industries/high-tech/our-insights/how-blockchains-could-change-the-world

The Council of European Union. (2018a). *Council regulation (EU) no 2016/679*. Retrieved from http://www.privacy-regulation.eu/en/article-4-definitions-GDPR.htm (Article 4(7))

The Council of European Union. (2018b). *Council regulation (EU) no 2016/679*. Retrieved from http://www.privacy-regulation.eu/en/article-25-data-protection-by-design-and-by-default-GDPR.htm (Article 25(1))

vad der Aalst, W., & van Hee, K. (2002). Appendix B: Workflow Modeling Using UML. In *Workflow Management - Models, Methods, and Systems*. MIT Press. Retrieved from https://app.knovel.com/hotlink/khtml/id:kt009E0B14/workflow-management-models/workflow-modeling-using

Vad är SAPSA. (2019). Retrieved from https://www.sapsa.se/om-sapsa/#sugen

van den Hooff, J., Kaashoek, M. F., & Zeldovich, N. (2014). VerSum: Verifiable Computations over Large Public Logs. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security* (p. 1304-1316). Retrieved from http://doi.acm.org.proxy.lib

.chalmers.se/10.1145/2660267.2660327

Vanderburg, E. (2018). Retrieved from https://www.digitalscouting.de/2018/07/19/our-intimate-use-of-technology-makes-us-more-vulnerable-to-its-misuse-eric-vanderburg-cybersecurity-expert/

van Lamsweerde, A. (2001, August). Goal-oriented requirements engineering: a guided tour. In *Proceedings Fifth IEEE International Symposium on Requirements Engineering* (p. 249-262). Retrieved from https://ieeexplore.ieee.org/document/948567

Violino, B. (2003, March). Benetton to tag 15 million items. *RFID Journal*. Retrieved from https://www.rfidjournal.com/articles/view?344

Vukolić, M. (2016, May). The Quest for Scalable Blockchain Fabric: Proof-of-Work vs. BFT Replication. In *Open Problems in Network Security* (p. 112-125). Springer International Publishing. Retrieved from https://link.springer.com/chapter/10.1007/978-3-319-39028-4_9

Wieringa, R. (2009). Design Science as Nested Problem Solving. In *Proceedings of the 4th International Conference on Design Science Research in Information Systems and Technology* (p. 1-12).

Wieringa, R., van Eck, P., Steghuis, C., & Proper, E. (2009). *Competences of IT Architects*. Nederlands Architectuur Forum. Retrieved from https://pdfs.semanticscholar.org/f79e/cdda868cd90d5cbb535862de80daff9dfc34.pdf

Wine Blockchain. (2019). Retrieved from https://www.ezlab.it/case-studies/wine-blockchain/

Wolfensberger, M. V. C., Drayer, L., & Volker, J. J. M. (2014). *Pursuit of Excellence in a Networked Society*. Waxmann.

Wüst, K., & Gervais, A. (2018). Do you Need a Blockchain? In *2018 Crypto Valley Conference on Blockchain Technology* (p. 45-54). Retrieved from https://ieeexplore.ieee.org/document/8525392

Yin, H. H. S., Langenheldt, K., Harlev, M., Mukkamala, R. R., & Vatrapu, R. (2019). Regulating Cryptocurrencies: A Supervised Machine Learning Approach to De-Anonymizing the Bitcoin Blockchain. *Journal of Management Information Systems*, *36*(1), 37-73. Retrieved from http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com.proxy.lib.chalmers.se/login.aspx?direct=true&db=buh&AN=135672064&site=ehost-live&scope=site

Yu, Y., Li, Y., Tian, J., & Liu, J. (2018, December). Blockchain-Based Solutions to Security and Privacy Issues in the Internet of Things. *IEEE Wireless Communications*, *25*(6), 12-18. Retrieved from https://ieeexplore.ieee.org/document/8600751

Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017, June). An Overview of Blockchain Technology: Architecture, Consensus, and Future Trends. In *2017 IEEE International Congress on Big Data* (p. 557-564). Retrieved from https://ieeexplore.ieee.org/document/8029379

Zheng, Z., Xie, S., Dai, H.-N., Chen, X., & Wang, H. (2018). Blockchain challenges and opportunities: A survey. *International Journal of Web and Grid Services*, *14*(4), 352-375. Retrieved from http://proxy.lib.chalmers.se/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=edsgao&AN=edsgcl.558536945&site=eds-live&scope=site

List of Figures

3.1 The cyclic Design Science Research methodology 14
3.2 The architecture framework artifact 15

4.1 A sample goal model using KAOS constructs 37
4.2 Business model canvas 38
4.3 A flowchart for blockchain platform selection (extracted from Wüst and Gervais (2018)) 39
4.4 A system quality breakdown 48
4.5 A UML domain diagram for a battery data 56
4.6 A use case diagram for GreenStar Marine's battery trading platform 62

List of Tables

3.1 Prioritization of research articles 18
3.2 Study participants 19

4.1 Potential stakeholders and suggested role(s) based on the workshop discussion 23
4.2 Literature review 24
4.3 Interview review 25
4.4 Challenging scenarios 27

5.1 The applicability viewpoint 69
5.2 The ecosystem viewpoint 70
5.3 The infrastructure viewpoint 71
5.4 The legal viewpoint 72
5.5 The end-user viewpoint 73

A. Interview Template

As described in the Section 3.3, a semi-structured interview format was employed. This permitted investigation into more broad questions while allowing the interviewee to dwell deeper into the areas of his or her expertise. The following interview template was used and sent out in advance. A larger set of questions were sent to the case company which feedback, advising us to exclude the too detailed and use case specific questions.

1. *How does one decide whether or not to use blockchain technology? Considering the current issues and limitations, what business aspects must be present for blockchain to be relevant?*
2. *Given the vast amount of technological solutions and architectural approaches, how do one navigate through the maze of options?*
3. *What questions must be considered when reasoning strategically about blockchain based implementations?*
    - *Which information is currently missing but would be valuable?*
    - *When developing blockchain solutions, which design decisions are critical early on?*
    - *How important is it to get the blockchain implementation and the development of smart contracts right form the very start, considering the immutable nature of the blockchain and the required security?*
4. *Which quality attributes would you emphasize in a blockchain application and how would you prioritize them?*
5. *What are the benefits/downsides of creating a blockchain network from scratch compared to building upon an already existing network?*
6. *Can blockchain be configured to comply with GDPR and regulations alike?*
7. *What are the prospects of the blockchain nodes to efficiently manage the explosion of data volume anticipated to occur? Although lightweight clients, mini-blockchains, and off-chain storage are proposed solutions, full- and mining nodes are required to maintain the entire blockchain.*
8. *What are the main hurdles that will need to be solved before blockchain-IoT solutions become widespread?*
9. *What security and privacy aspects must be considered when thinking about blockchain in IoT applications?*
10. *Is it feasible to establish blockchain-IoT networks without a cryptocurrency providing computational incentives? Assuming constrained IoT sensors are implemented as light weight nodes, how will the network otherwise motivate full nodes to carry out work for their light weight counterparts?*

B. Survey Template

The survey contained six independent questions covering the sections in each viewpoint. The closed questions were answered using Likert scales from 1 to 6, larger scores indicate higher contentment. Participants had the opportunity to add additional comments under each question, would they have concrete ideas on how the framework could be improved. The survey questions are enumerated below:

1. How comprehensible is the viewpoint? Is it easy to follow and understand its intention?
2. How well is the viewpoint structured?
3. How well does the challenges discussed in the viewpoint resemble your observations?
4. How well does stakeholders in the viewpoint resemble your observation?
5. How valuable are the operations on views in the viewpoint?
6. How useful are the models in the viewpoint?