A Blueprint for a Trustworthy Health Data Platform Encompassing IoT and Blockchain Technologies

Dennis Przytarski[1], Christoph Stach[1], Clémentine Gritti[2], and Bernhard Mitschang[1]

[1] Universität Stuttgart, IPVS, Universitätsstraße 38, 70569 Stuttgart, Germany
firstname.lastname@ipvs.uni-stuttgart.de
[2] University of Canterbury, CSSE, Christchurch 8041, New Zealand
clementine.gritti@canterbury.ac.nz

Abstract

eHealth provides great relief for patients and physicians. This means, patients autonomously monitor their condition via IoT medical devices and make these data available to physicians for analyses. This requires a data platform that takes care of data acquisition, management, and provisioning. As health data are highly sensitive, there are major concerns regarding data security with respect to *confidentiality*, *integrity*, and *authenticity*. To this end, we present a blueprint for constructing a trustworthy health data platform called *SEAL*. It provides a *lightweight attribute-based authentication mechanism* for IoT devices to validate all involved data sources, there is a *fine-grained data provisioning system* to enable data provision according to actual requirements, and a *verification procedure* ensures that data cannot be manipulated.

1 Introduction

The *Quantified Self* movement is significantly facilitated by the increasing popularity of the *Internet of Things* (*IoT*). For this purpose, people use *IoT devices*—i. e., everyday objects equipped with sensors and the ability to share data—to gather and analyze data about their lives. However, this is not just a gimmick as these data have a considerable value. For people with chronic diseases, it is a great relief if they can carry out the daily routine monitoring of their health data on their own and thus minimize the number of visits to the physician. This not only increases the self-reliance of patients, but also relieves the physicians considerably, as they can focus on emergencies [11].

To enable this, a *health data platform* is required that manages the captured data and provides it to the appropriate physicians. As a health data platform contains highly sensitive information, data engineers have to consider special requirements regarding data security when constructing such a platform. It has to be ensured that only authorized parties have access to the health data and the patient's privacy is guaranteed. Only then, patients may entrust their data to such a platform [13]. In addition, physicians must be able to rely on the integrity of the provided data, i. e., the data have been captured correctly using valid measuring instruments and have not been tampered with since.

A *trustworthy* health data platform must therefore consider the following aspects: *a)* Physicians (i. e., *data consumers*) must rely that the data comes from trustworthy sources. For instance, it must be ensured that they have not been manipulated and that they have the required accuracy. *b)* Patients (i. e., *data producers*) must trust that only authorized consumers get access to their data and that as little sensitive information as necessary is disclosed. For instance, the quantity or quality of available data can be restricted. *c)* Both, producers and consumers must trust that neither the platform nor third parties can manipulate the data. For instance, this can be ensured by means of a verification mechanism.

As existing platforms do not meet these requirements, we introduce a blueprint for constructing a trustworthy health data platform called *SEAL*, which is intended to assist data engineers in the design of such platforms. To this end, we make these contributions:

A) We apply a **lightweight attribute-based authentication mechanism** for IoT devices. Sources sign their data with verifiable attributes characterizing the measurement parameters (e. g., a device's firmware version or the duration of a measurement). The signature ensures data integrity during transmission. This covers aspect a).

B) We introduce a **fine-grained data provisioning system**. It not only regulates who has access to which data, but also introduces privacy filters. These filters can conceal certain features in the data without impairing the quality of the remaining data. This covers aspect b).

C) We outline a **verification procedure** to ensure the integrity of the data stored in SEAL. Hash values of the data are stored in a public blockchain so that everybody can check whether the health data have been tampered with. All data managed by our platform are *sealed* (in terms of encrypted) to ensure confidentiality and integrity. This covers aspect c).

Although the focus of this paper is on health-related use cases, SEAL enables trustworthy data acquisition, management, and provisioning for use cases in any IoT domain.

The remainder of this paper is as follows: In Section 2, we introduce an application scenario and derive requirements towards a health data platform. Section 3 discusses related work. We introduce SEAL in Section 4 and subsequently assess it in Section 5, whether it meets the requirements. Finally, Section 6 concludes this paper and gives an outlook on future work.

2 Application Scenario

Patients suffering from chronic diseases such as diabetes have to monitor certain health data periodically, e. g., their blood sugar level. The IoT enables *gamification* by incorporating the monitoring into a game, so it is less inconvenient. As IoT devices are interconnected, an IoT medical device can provide the health data not only to the game, but also to the physician in charge. These health data can additionally be annotated with further data, such as location data for each measurement. These metadata are not only interesting for physicians, but also for other stakeholders, such as city planners who want to create healthier cities [11]. Such applications benefit from a health data platform that handles data management and distribution.

In our application scenario (Fig. 1), a diabetic patient is equipped with a *continuous glucose monitoring sensor* (*CGM*) that regularly measures the blood sugar level. Each measurement is transmitted to a health data platform, where it is stored, linked to further data from the patient (e. g., location data), and made available to third parties. For instance, physicians get access to the measurements of their patients, so that they can monitor their health and adapt their treatments accordingly. Also, other stakeholders such as insurance agents and researchers can benefit from these data. Yet, patients express great concern regarding *confidentiality*. Therefore, third parties must only receive information about patients that is sufficient for their respective

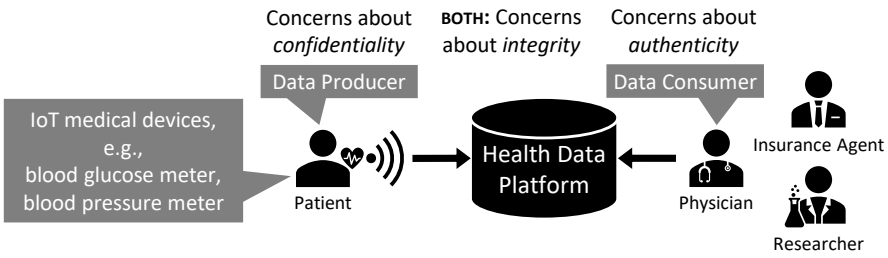

Figure 1: Schematic Representation of our Application Scenario.

use cases, e. g., insurance agents have access to aggregated information about the CGM status while researchers have access to aggregated and anonymized data only. In return, these data consumers expect that data *authenticity* is ensured, i. e., they originate from valid sources such as approved medical devices. In addition, data producers and data consumers have great interest that the *integrity* of the health data is ensured., i. e., no one can tamper with the data unnoticed.

This application scenario follows the approach of Steimle et al. [21], in which patients supplement their medical record via self-measurement.

Potential attack vectors against such a health data platform include, for instance: *i)* a user uses a manipulated fitness tracker in order to get better health insurance rates; *ii)* a hacker alters electronic health records of a hospital in order to extort money; *iii)* a physician subsequently adds treatments in order to receive more money from the health insurance; and *iv)* a health insurance company wants to access health data that violates the patient's privacy.

Despite all the benefits provided by such an application, several security and privacy concerns arise on the part of data producers (i. e., patients) and data consumers (e. g., physicians, insurance agents, or researchers), as our scenario illustrates. In this respect, data engineers have to take the following requirements into consideration in each of their areas of activity, namely acquisition, management, and provision of health data in order to build trust in their work:

R1 Authenticity. All data must originate from valid sources. This applies to both, the medical device itself (e. g., it has to be medically approved and its firmware has to be up to date) as well as the actual measurement procedure (e. g., the duration of the measurement). This ensures that no falsified data are uploaded to the health data platform.

R2 Confidentiality. By default, only data producers have access to their data. To share data with third parties, a permission management is required. Yet, it must be ensured that these authorized parties only have access to the information that is necessary for their specific use cases. This includes the application of privacy filters to reduce the information derivable from the data. Any data access by unauthorized parties must be averted.

R3 Integrity. Manipulations to measurements before transmission or subsequently in the health data platform must be detectable. This includes both, tampering with measured values as well as deleting data from the platform.

3 Related Work

There are several (partly discontinued) health data platforms such as the *Microsoft HealthVault*, enabling any application to feed data into them and providing these data to third parties [16]. Yet, patients (i. e., data providers) demand comprehensive data security measures when entrusting their sensitive health data to such a platform [13]. In recent years, blockchain-based approaches have turned out to be particularly suitable for this purpose [10]. In these approaches, however,

integrity of the shared data is the main focus. This is largely ensured by the immutability and tamper-resistance of the blockchain itself. More advanced approaches such as the ones by Azbeg et al. [1] and Dwivedi et al. [4] also include data acquisition and data provisioning in their security considerations. Yet, they do not provide a full end-to-end security solution, as required to build trust in such a platform. Azbeg et al., for instance, do not enable a fine-grained filtering of data for data provisioning in order to realize the privacy requirements of patients more effectively. Dwivedi et al. use digital signatures to ensure the integrity of the data. Yet, as data sources independently generate the keys for these signatures, it cannot be ensured that the sources themselves have not already been compromised and therefore provide corrupted data.

As none of these approaches provides a full end-to-end security solution ranging from gathering authentic IoT measurements to providing these data in a privacy-friendly manner, we discuss in the following some approaches that can be used to improve the authenticity, confidentiality, and integrity of such a platform.

Authenticity. In dynamic environments, such as the IoT, attribute-based authentication methods are well suited [22]. Since these approaches are heavyweight, while IoT devices are limited in resources, Karati et al. [9] introduce lightweight authentication certificates that also comprise user attributes. Yet, as users specify by themselves which attributes are included, the reliability of authentication is questionable. Plus, the usage of too much personal data for authentication poses a privacy risk. Idalino et al. [7] address the problem that data records have to be updatable (e. g., measurements have to be added to a medical record). To this end, they introduce a modular (i. e., expandable) signature. Yet, as this enables any *authorized* entity to modify the data, data authenticity is not ensured. Taylor [23] propose a method to detect inauthentic data (e. g., fake news). However, the data can only be validated in retrospect with some risk of possible misjudgment.

Confidentiality. Yang et al. [25] introduce a secure data storage for confidential data, which allows distributed access to these data. Yet, access is granted on an all-or-nothing basis. Also, role-based authentication is too coarse-grained [24]. *DISPEL* [18] allows users to use privacy filters to distort their data in order to reduce the amount of disclosed information. Since the control over these filters is solely in the hands of the users, the data can be rendered useless if they are distorted too heavily. This is prevented by using differential privacy techniques [8]. Yet, these techniques can only be used for statistical queries and do not enable analyses of the data of an individual user.

Integrity. Blockchain approaches are well suited to ensure integrity when sharing data with third parties [10]. Yet, full integrity is only ensured for public blockchains, which entail high transaction fees. Liang et al. [12] therefore store only a fingerprint of the data in the actual blockchain. Yet, integrity is only considered after the data have been stored, i. e., it is not ensured that the received data were genuine in the first place. Homomorphic encryption allows performing analytics on encrypted data without decrypting them first [14] and Chen et al. [2] introduce a searchable encryption scheme to make such fingerprints queryable. Yet, homomorphic encryption reveals some information about the data, which is not tolerable for health data.

Combining authenticity, confidentiality, and integrity into a holistic approach for a trustworthy health data platform is the novelty in our approach that differentiates from related work.

4 The Trustworthy Health Data Platform SEAL

Since none of these approaches fully meet the requirements towards a trustworthy health data platform regarding confidentiality, integrity, and authenticity, we come up with a novel blueprint for constructing a trustworthy architecture called SEAL.

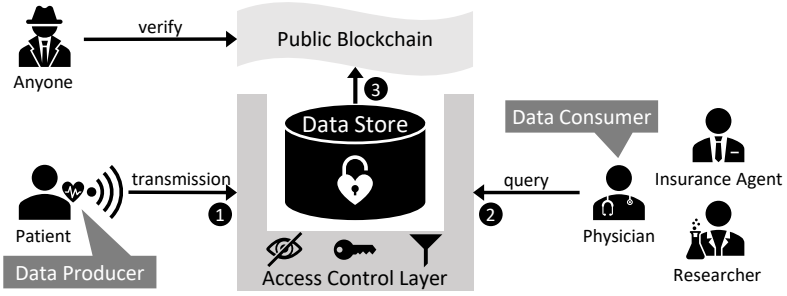

Figure 2: Data Acquisition, Management, Provision, and Verification via SEAL.

SEAL consists of three main components (Fig. 2): an *access control layer*, a *secure data store*, and a *blockchain*. The access control layer monitors and regulates all incoming and outgoing data flows. When a data producer (i. e., a patient) wants to upload new data to SEAL (e. g., new measurements from his or her CGM), s/he has to authenticate towards this access control layer first. For this purpose, we use a *lightweight attribute-based authentication mechanism* ❶ (Section 4.1). It ensures that the measurement data is valid, i. e., the measuring device is approved, the applied software is suitable, and the measurement was carried out correctly. If the authentication was successful, the data are stored in SEAL.

To this end, a secure data store is used. In this store, all data are encrypted and only the access control layer has the key, i. e., we can rely on symmetric cryptography as no costly and vulnerable key exchanges are required. This way, SEAL prevents the loss of confidentiality even if an attacker is able to bypass the access control layer and gets direct access to the data store. For more information on the design of such a secure data store, please refer to literature [19].

Data consumers (e. g., physicians, insurance agents, or researchers) get access to the data via the access control layer. An attribute-based access control approach is also used to this end. Via a *fine-grained data provisioning system* ❷ (Section 4.2), data consumers receive only those data that are required for their specific purpose (with respect to both, quantity and quality) to ensure minimal information disclosure. To this end, various privacy filters are applied.

Although it is ensured that the data cannot be manipulated by third parties, SEAL itself could manipulate the data unnoticeable. For example, an inside attacker like a SEAL administrator may tamper with the data either via SEAL APIs or SEAL administrator tools. Therefore, SEAL additionally provides a *verification procedure* ❸ for all stored data (Section 4.3). To this end, hash values for all data from the secure data store are stored on a public blockchain—also known as the anchoring technique. With these hashes, anyone is able to check the authenticity of the (encrypted) data. This way, SEAL ensures confidentiality, integrity, as well as authenticity and thus represents a trustworthy health data platform.

4.1 Lightweight Attribute-based Authentication Mechanism

For the authentication of data sources, we apply a mechanism based on Gritti et al. [5]. This approach is shown in Fig. 3. In an initial step, a trusted authority creates a private key for each source. These keys characterize all identifying attributes of the source. In SEAL, these attributes comprise not only features of the medical device (e. g., its firmware version *Attr(fmv)*) but also features of the application used to perform a metering (e. g., its serial number *Attr(snr)*) as well as features of the metering itself (e. g., how long it took *Attr(len)*). These keys are only valid if the source complies with these identifying attributes.

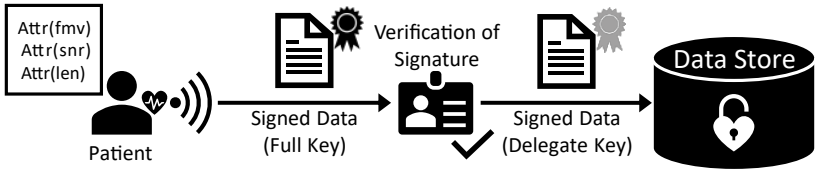

Figure 3: Authentication Process for Data Sources Applied in SEAL.

Data sources sign all data they send to SEAL with their respective private key. The access control layer receives the data and verifies the signature against an authentication policy. For this purpose, it has a public key that corresponds to the private key, i. e., it describes the expected attributes of the source. That way, SEAL ensures, e. g., that the device is approved for medical use, that the user is not using invalid or outdated firmware or software, and that the measurement was performed correctly and for an adequate length of time. If SEAL successfully verifies the signature, i. e., the data authenticity is ensured, the payload is stored encrypted in the data store. In addition, the integrity of the transmission is guaranteed since the data cannot be manipulated due to the signature.

Since the attributes used for the signature could reveal information about the source (and thus the data producer), it must not be revealed to the data consumer. In order to enable data consumers still to verify that the source has been authenticated, an approach called *delegated authentication* is applied [6]. To this end, the access control layer modifies the signature in terms of filtering out all compromising attributes. This reduced signature can then be verified by the data consumers against a reduced authentication policy.

4.2 Fine-grained Data Provisioning System

To enable needs-based data provision, it is necessary to identify what information can be derived from which sources and which data are required to fulfill a purpose. These correlations and requirements are specified in an *EPICUREAN* model [20]. One simple use case could be that an insurance agent has to audit whether a diabetic patient follows a diet. Via the EPICUREAN model, SEAL can determine which data sources the insurance agent needs to have access to, e. g., the data of the CGM. However, to ensure that the patient's privacy is not unnecessarily violated, the data quality can be reduced to such an extent that only the blood sugar level progression is visible, but not accurate individual values (e. g., by adding noise). This way, SEAL determines which restrictions regarding data quantity and quality are acceptable for which use case.

Using these specifications, data producers can define fine-grained access permissions. Such permission rules consist of four components: *(1)* recipient (i. e., date consumers), *(2)* context (i. e., purposes), *(3)* information (i. e., data sources), and *(4)* privacy filters. With the first three components data producers specify who gets access to which data and for what purpose. By adding privacy filters, the data quality can be reduced, or certain data records can be filtered out prior to provision. Yet, the privacy filters are not selected by the data producer, but are automatically derived from the EPICUREAN model, i. e., SEAL ensures that these filters comply with the requirements of the data consumer.

Figure 4 shows an example of how such a filter operates. It is based on the *SNIL* algorithm [3]. This filter is intended for generating noise in time series data—as patients suffering from chronic diseases have to monitor their health continuously, this is a relevant type of data type in our context. The filter decomposes the data series—blood glucose data in the example—into

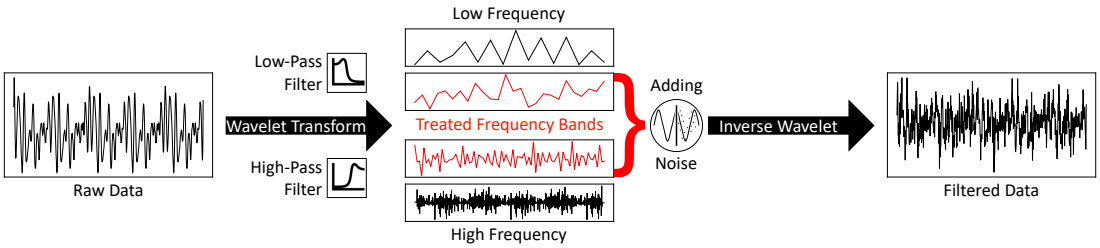

Figure 4: Application of a SNIL-based Privacy Filter to Blood Glucose Data.

individual frequency bands using discrete wavelet transform. Noise is added to the mid-range frequency bands (high frequencies are immune against noise, while noisy low frequencies affect the result too much). The frequency bands are then recomposed using inverse wavelet transform. As a result, it is still possible to monitor the progression, but no individual blood sugar levels. This technique is robust against noise filters, i. e., the raw data cannot be restored.

There are further privacy filters tailored to other types of data such as trajectories, i. e., filters to conceal location data. This allows SEAL to select an appropriate filter for any use case to minimize the disclosed information and thus maximize confidentiality. At the same time, integrity is ensured for data consumers, as it is guaranteed that the provided data have the quantity and quality as needed for their purposes.

The whole data request process is shown in Fig. 5. When a data consumer sends a query to SEAL, the access control layer identifies the consumer via his or her attributes (recipient) and derives the purpose (context) as well as the relevant data sources (information) from the query. SEAL then performs a policy lookup, whether the data producer specified a permission rule for such a query and whether a privacy filter has to be applied. Based on these permissions, SEAL identifies all required data records within its data store and assembles the query results. If necessary, privacy filters are applied by the access control layer before the results are provided to the data consumer.

4.3 Verification Procedure

As falsification of sensitive data (i. e., health data) might have grave consequences, the goal is to eliminate all unnoticed modifications to the health data stored in SEAL. This includes *technical errors* such as bit flips altering stored measurement values that might provoke an unnecessary medical review of a patient's health data; *human errors* such as unintentionally deleting health data; and also, *inside* and *outside attackers* altering health data to harm a patient.

The blockchain technology offers characteristics such as immutability and tamper-resistance, which ensures data integrity of all data stored on a blockchain. These are important characteristics that eliminate the aforementioned attacks on the data's integrity. In particular, a public blockchain has the advantage that the other participants in the blockchain network are practically always unknown to oneself. Unless one has the majority of the mining power in the blockchain network, it is practically impossible to convince other (unknown) participants to alter the blockchain's history for one's own benefit.

Therefore, a public blockchain ensures data integrity as all data in it is practically immutable and tamper-resistant. However, public blockchains may cause high transaction fees due to several factors such as the current network congestion (e. g., many unconfirmed transactions compete for the limited storage space of the next block). In SEAL, we minimize the data stored on the

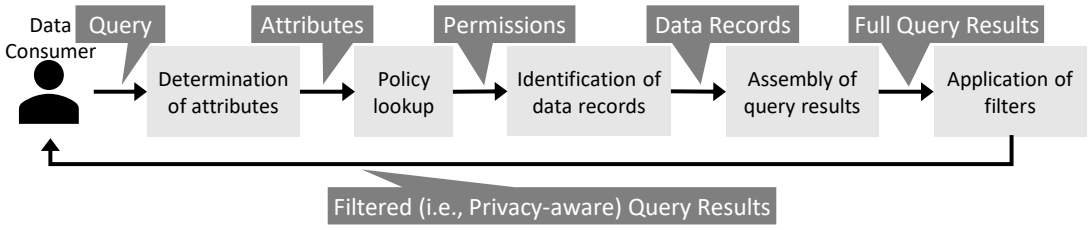

Figure 5: Query Process and Application of Privacy Filters.

public blockchain by using the anchoring technique. For each stored data record in the data store of the health platform, a hash value is computed and stored on the public blockchain—the actual data are stored encrypted in a secure data store such as *CURATOR* [19].

The integrity of the stored data can be verified by anyone via these hashes, i. e., subsequent manipulations can be detected. To this end, SEAL also provides an interface to directly validate stored data. The encrypted data is used to compute hashes and compare them with the ones stored on the blockchain. When these hashes match, data integrity is guaranteed. As the hashes do not contain any confidential information, this approach does not compromise confidentiality.

5 Assessment

In the following, we assess whether SEAL is indeed a *trustworthy* health data platform, i. e., whether it fulfills the three key requirements outlined in Section 2.

Regarding Requirement R1 (**Authenticity**), data consumers are concerned that the data do not originate from valid sources, so that falsified or incorrect data are uploaded to SEAL. This can happen, when there is a problem with the medical device itself (e. g., use of an unapproved medical device or outdated firmware) or with the actual metering (e. g., the duration of a measurement is shorter than specified by the manufacturer). In this respect, we use a *lightweight attribute-based authentication mechanism* to authenticate data sources towards SEAL's access control layer. They sign all data with their identifying attributes such as the firmware version of the medical device, the serial number of the application used, and the duration of the metering itself. As SEAL rejects any unsigned data or data with an inadequate signature (e. g., an outdated firmware or the duration of a measurement is outside of the specified range), it is ensured that only authentic data from valid sources are stored in SEAL.

Regarding Requirement R2 (**Confidentiality**), data producers are concerned that illegitimate third parties get access to sensitive data. To this end, every data access in SEAL is handled exclusively by the access control layer. Here it is checked whether the requesting party is authorized for access while unauthorized parties are directly rejected. Even if an unauthorized party is able to bypass the access control layer and directly access the data store, this does not constitute a breach of confidentiality since all data are encrypted. But also, in case of an authorized access, confidentiality is maintained. For this, there is a *fine-grained data provisioning system* in SEAL. The applied permission rules define who (recipient) may access which data sources (information), for what purpose (context). In addition, privacy filters are used, which are tailored to the respective use case. These filters reduce the data quantity and quality so that a need-based data provision is possible without compromising the data producer's confidentiality.

Regarding Requirement R3 (**Integrity**), both, data producers and data consumers are concerned that data could get manipulated either during transmission or at rest in SEAL. Since

all data are digitally signed by the data sources, any data manipulation during transmission is detected by SEAL when the signature is verified. As the data are then immediately encrypted by SEAL, third parties cannot manipulate them at rest. Yet, as SEAL has access to the key, it must be ensured that the platform itself cannot manipulate the data. To this end, SEAL provides a *verification procedure*. For each stored data record in the data store, a hash value is computed and stored on a public blockchain (e. g., Ethereum[1]). This allows anyone to verify the data's integrity by computing the hash values of the data in question and compare them with the ones on the public blockchain. We use this anchoring technique, because it is effective in preventing potentially harmful data sabotage on data at rest. However, each anchoring point to the public blockchain entails transaction fees. In order to reduce these transaction fees, multiple data records may be clustered (e. g., the measurement values of a day). Now only the cluster's hash value is stored on the public blockchain. On the downside, however, just one manipulated data record in a cluster corrupts the whole clustered data set. Therefore, an acceptable balance must be struck between the cluster volume and the transaction fees to be paid.

So, SEAL dispels all concerns regarding authenticity, confidentiality, and integrity, whereby it is trustworthy for both, data producers and data consumers.

6 Conclusion and Future Work

In this paper, we present a blueprint for constructing a trustworthy health data platform called *SEAL*. A *lightweight attribute-based authentication mechanism* for IoT devices ensures that only authentic data are stored in SEAL. A *fine-grained data provisioning system* ensures the confidentiality of a patient's health data, as it regulates who has access to which data and applies privacy filters, if necessary. A *verification procedure* ensures the integrity of the data stored in SEAL, i. e., measurement values cannot be tampered with. Thus way, SEAL dispels the main concerns towards a health data platform regarding *authenticity*, *confidentiality*, and *integrity*.

Although prototypes for the individual components of SEAL were developed as part of our previous work (e.g., for the implementation of a lightweight attribute-based authentication mechanism see [6], for the implementation of a fine-grained data provisioning system see [17], and for the implementation of secure and verifiable data stores see [19]), which proves the feasibility of SEAL, a full implementation of the blueprint is yet to be realized. Therefore, it is planned to integrate these available components and to come up with a fully functional prototype of SEAL as part of future work. A particular focus lies on the investigation of the merger of the data store and a private blockchain system with a refined data model enabling more efficient query processing capabilities [15]. Hereby, even fewer anchoring points to a public blockchain are required resulting in fewer transaction costs while still enabling comprehensive data verifications required by SEAL.

Acknowledgements. This work is funded by the DFG project DiStOPT (252975529).

References

[1] Kebira Azbeg et al. Blockchain and IoT for Security and Privacy: A Platform for Diabetes Self-management. In *Cloudtech '18*, 2018.

[2] Lanxiang Chen et al. Blockchain based searchable encryption for electronic health record sharing. *Future Gener Comput Syst*, 95:420–429, 2019.

[1] see https://ethereum.org

[3] Mi-Jung Choi et al. Publishing Sensitive Time-Series Data under Preservation of Privacy and Distance Orders. *IJICIC*, 8(5(B)):3619–3638, 2012.

[4] Ashutosh Dwivedi et al. A Decentralized Privacy-Preserving Healthcare Blockchain for IoT. *Sensors*, 19:326:1–326:17, 2019.

[5] Clémentine Gritti et al. CHARIOT: Cloud-Assisted Access Control for the Internet of Things. In *PST '18*, 2018.

[6] Clémentine Gritti et al. Privacy-Preserving Delegable Authentication in the Internet of Things. In *SAC '19*, 2019.

[7] Thaís Bardini Idalino et al. Modification Tolerant Signature Schemes: Location and Correction. In *INDOCRYPT '19*, 2019.

[8] Noah Johnson et al. Towards Practical Differential Privacy for SQL Queries. *Proc VLDB Endow*, 11(5):526–539, 2018.

[9] Arijit Karati et al. Provably Secure and Generalized Signcryption With Public Verifiability for Secure Data Transmission Between Resource-Constrained IoT Devices. *IEEE IoT J*, 6(6):10431–10440, 2019.

[10] Seyednima Khezr et al. Blockchain Technology in Healthcare: A Comprehensive Review and Directions for Future Research. *Appl Sci*, 9(9):1736:1–1736:28, 2019.

[11] Martin Knöll et al. *Spontaneous Interventions for Health: How Digital Games May Supplement Urban Design Projects*, pages 245–259. Springer, 2014.

[12] Xueping Liang et al. Integrating blockchain for data sharing and collaboration in mobile healthcare applications. In *PIMRC '17*, 2017.

[13] Emily C. O'Brien et al. Patient perspectives on the linkage of health data for research: Insights from an online patient community questionnaire. *Int J Med Inform*, 127:9–17, 2019.

[14] Timothy Oladunni and Sharad Sharma. Homomorphic Encryption and Data Security in the Cloud. In *SEDE '19*, 2019.

[15] Dennis Przytarski. Using Triples as the Data Model for Blockchain Systems. In *Block-SW/CKG@ISWC '19*, 2019.

[16] Ton Spil and Richard Klein. Personal Health Records Success: Why Google Health Failed and What Does that Mean for Microsoft HealthVault? In *HICSS '14*, 2014.

[17] Christoph Stach. VAULT: A Privacy Approach towards High-Utility Time Series Data. In *SECURWARE '19*, 2019.

[18] Christoph Stach et al. Bringing Privacy Control back to Citizens: DISPEL – A Distributed Privacy Management Platform for the Internet of Things. In *SAC '20*, 2020.

[19] Christoph Stach and Bernhard Mitschang. Curator — A Secure Shared Object Store: Design, Implementation, and Evaluation of a Manageable, Secure, and Performant Data Exchange Mechanism for Smart Devices. In *SAC '18*, 2018.

[20] Christoph Stach and Frank Steimle. Recommender-Based Privacy Requirements Elicitation - EPICUREAN: An Approach to Simplify Privacy Settings in IoT Applications with Respect to the GDPR. In *SAC '19*, 2019.

[21] Frank Steimle et al. Extended provisioning, security and analysis techniques for the ECHO health data management system. *Computing*, 99:183–201, 2017.

[22] Tobias Straub and Ulf Schreier. Distributed Access Control for the Internet of Things. In *SummerSOC '18*, 2018.

[23] Lamar S. Taylor. Trusted Validity: Combating Fake News with Distributed Ledger Technology. In *SEDE '18*, 2018.

[24] Arundhati Wahane and Ying Jin. A Graph Database Approach for XACML Role-Based Access Control Implementation. In *SEDE '18*, 2018.

[25] Yang Yang et al. Privacy-preserving smart IoT-based healthcare big data storage and self-adaptive access control system. *Information Sciences*, 479:567–592, 2019.