Decentralised Provenance for Healthcare Data

Andrea Margheri[a], Massimiliano Masi[b,*], Abdallah Miladi[b], Vladimiro Sassone[a], Jason Rosenzweig[c]

[a] *University of Southampton, University Road, Southampton, SO17 1BJ, UK*
[b] *Tiani "Spirit" Gmbh, DC Tower 1, Donau-City-Straße 7, 1220 Vienna, AT*
[c] *Cisco Healthcare, USA*

Abstract

Objective. The creation and exchange of patients' Electronic Healthcare Records have developed significantly in the last decade. Patients' records are however distributed in data silos across multiple healthcare facilities, posing technical and clinical challenges that may endanger patients' safety. Current healthcare sharing systems ensure interoperability of patients' records across facilities, but they have limits in presenting doctors with the clinical context of the data in the records. We design and implement a platform for managing provenance tracking of Electronic Healthcare Records based on blockchain technology, compliant with the latest healthcare standards and following the patient-informed consent preferences.

Methods. The platform leverages two pillars: the use of international standards such as Integrating the Healthcare Enterprise (IHE), Health Level Seven International (HL7) and Fast Healthcare Interoperability Resources (FHIR) to achieve interoperability, and the use of a provenance creation process that by-design, avoids personal data storage within the blockchain. The platform consists of: (1) a smart contract implemented within the Hyperledger Fabric blockchain that manages provenance according to the W3C PROV for medical

*Corresponding author
Email addresses: a.margheri@soton.ac.uk (Andrea Margheri), massimiliano.masi@tiani-spirit.com (Massimiliano Masi), abdallah.miladi@tiani-spirit.com (Abdallah Miladi), vsassone@soton.ac.uk (Vladimiro Sassone), jason.rosenzweig@gmail.com (Jason Rosenzweig)

document in standardised formats (e.g., a CDA document, a FHIR resource, a DICOM study, etc.); (2) a Java Proxy that intercepts all the document submissions and retrievals for which provenance shall be evaluated; (3) a service used to retrieve the PROV document.

Results. We integrated our decentralised platform with the SpiritEHR engine, an enterprise-grade healthcare system, and we stored and retrieved the available documents in the Mandel's sample CDA repository[1], which contained no protected health information. Using a cloud-based blockchain solution, we observed that the overhead added to the typical processing time of reading and writing medical data is in the order of milliseconds. Moreover, the integration of the Proxy at the level of exchanged messages in EHR systems allows transparent usage of provenance data in multiple health computing domains such as decision making, data reconciliation, and patient consent auditing.

Conclusions. By using international healthcare standards and a cloud-based blockchain deployment, we delivered a solution that can manage provenance of patients' records via transparent integration within the routine operations on healthcare data.

Keywords: healthcare, data provenance, blockchain, interoperability, FHIR

1. Introduction

Electronic healthcare systems are deployed worldwide and are changing the way medical treatments are prescribed and administered. One critical element of such *eHealth* systems is the management of patients' medical data, so-called *Electronic Healthcare Record* (EHR), across geographically distributed, usually non-interoperable, healthcare centres. Another is allowing patients to provide consent for doctors to access their medical records regardless the healthcare facility used. The availability of EHRs is indeed paramount to ensure patients safety and continuous health treatment.

[1] See `https://github.com/jmandel/sample_ccdas`

A fundamental shortcoming of current data sharing solutions is the lack of clinical context attached to patients' records: such context would permit, e.g., validating clinical plausibility or assessing the data sources [1]. Tracking the *provenance* of healthcare data stored across distributed EHRs would support such data-informed medical decision-making and clinical research [2, 3, 4]. Also, tracking provenance for healthcare data access empowers patients to have full control on the secondary use of personal data, i.e., creating awareness of where their data goes (e.g., public health enquiries, clinical trials).

In this paper, we apply blockchain technology to provide a *practical and ready-to-be-deployed solution for the decentralised management of the provenance of healthcare data*. We base our solution on a principled integration of a private permissioned blockchain and the W3C standard PROV [5] for provenance management. PROV allows practical creation of highly expressive provenance annotations, e.g., based on casual and temporal relationships of the activities on the monitored data. Permissioned blockchain frameworks can offer a fully decentralised platform that can cope with the distribution of EHRs, while balancing privacy, interoperability and performance needs.

Structure of the Paper. We first introduce background concepts and related works (Section 2), then we describe the architecture of our provenance system (Section 3) and its eHealth applications (Section 4). We comment on main challenges addressed (Section 5) and then conclude the paper (Section 6).

2. Background Concepts and Related Work

We introduce the context of eHealth systems, data provenance and the W3C PROV standard, and blockchain technology. We report related work as well.

eHealth Systems. The cornerstone of eHealth systems is the management of EHRs. Records are shared using internationally adopted standards such as IHE and HL7. We will use the IHE XDS messaging model [6] (shown in Figure 1) to submit, query, and retrieve medical documents across organisations. These messages are defined by IHE profiles [6, Table 3.20.4.1.1.1-1] (e.g. CREATE

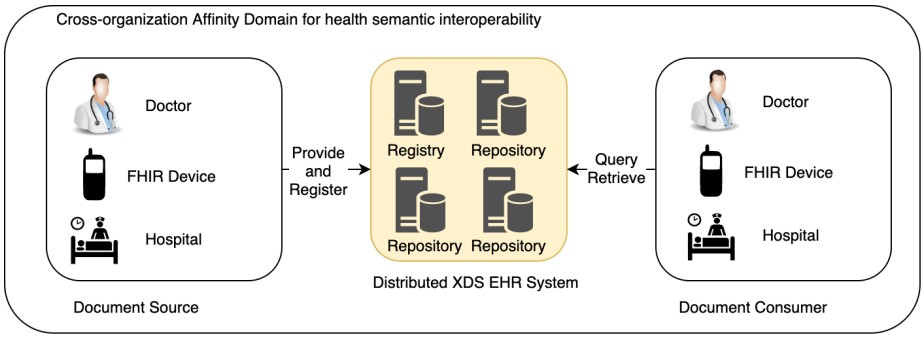

Figure 1: **Electronic Healthcare Record (EHR) systems**: exchange health documents of various format (e.g., CDA, FHIR) according to standardised XDS messages providing enforcement of patient-informed consent and health semantic interoperability based on shared attributed part of cross-organisation affinity domains.

and UPDATE) and can contain different types of health documents: 1) HL7 Clinical Document Architecture (CDA) [7]; 2) medical reports saved as PDF; 3) other standardised formats like DICOM [8] medical imaging information or HL7 FHIR [9] medical information formatted as JSON resources. In addition, XDS builds on the concept of *affinity domain* which provides the building blocks for the semantic interoperability of exchanged health data.

To implement and validate our system, we use an enterprise XDS system: the *Spirit Electronic Health Record* (SpiritEHR), an eHealth product provided by Tiani Spirit[2] for the secure management and sharing of health data.

Data Provenance. Provenance is used to support the assessment of the quality of data, by, for example, identifying source of errors or attribution of sources. For healthcare data, the US Office of the National Coordinator for Health Information Technology (ONC)[3] defines

> *Provenance as attributes about the origin of health information at the time it is first created and tracks the uses and permutations of the health information over its lifecycle*

[2] See `http://www.tiani-spirit.com`

[3] See `http://wiki.siframework.org/Data+Provenance+Glossary`

Practically, tracking the provenance of some data corresponds to build a graph of semantically connected concepts that describe the entities, activities and agents that were involved with such data.

Current eHealth systems manage provenance via convoluted methods based on audit trails [6] or digital signature (e.g. the IHE DSG[4]). These methods deliver security mechanisms that provide guarantees on the sources of data; however, they have significant shortcomings as they rely on Trusted Third Party and are prone to semantic interoperability issues due to logging across different organisations. To overcome these deficiencies, we selected the W3C PROV [5] provenance standard. W3C PROV provides design and implementation means for sharing semantically interoperable provenance attributes. Also, significant healthcare bodies, such as IHE and HL7, support PROV[5].

Blockchain. It offers a distributed, fault-tolerant data storage and computing platform whereby the design ensures decentralised data control. The crypto and distributed consensus mechanisms used to regulate changes to the stored data ensure immutability of data, of code (so-called *smart contract*) and that distributed data replicas are consistent.

Blockchain is being explored in the eHealth domain for many purposes, including data sharing, identity and access management [10, 11, 12, 13]. Although current blockchain systems need improvements in terms of security, the benefits for healthcare applications are evident, especially in terms of interoperability [14]. Our provenance system applies blockchain to overcome the need of Trusted Third Parties for verification and assessment of provenance attributes.

Specific to provenance, FHIRChain [13], MeDShare [15] and SmartProvenance [11] propose provenance tracking functions; with SmartProvenance that uses the Open Provenance Model [16] approach, the previous version of PROV. Our system differs significantly with all reported as we build on international eHealth standard to achieve by-design interoperability of our solution. Also, dif-

[4]IHE DSG - `https://wiki.ihe.net/index.php/Document_Digital_Signature`

[5]`https://wiki.ihe.net/index.php?title=Query_for_Existing_Data_for_Mobile`

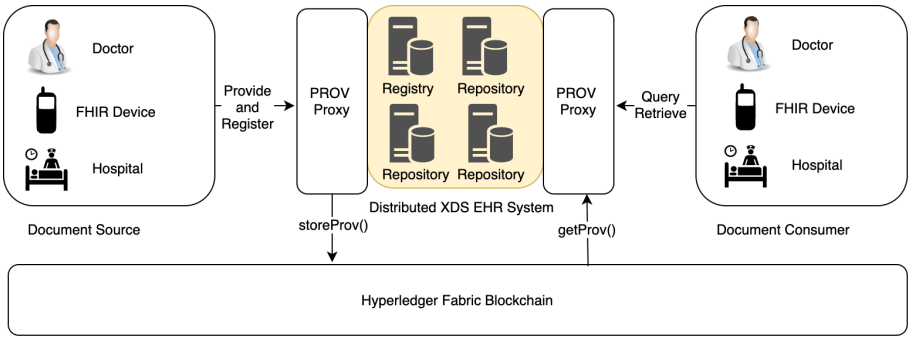

Figure 2: **Blockchain health provenance**: the Distributed XDS EHR System integrated with our PROV proxy that delivers decentralised provenance via the blockchain. Provenance is stored in the blockchain when documents are provided by the XDS Document Source and retrieved upon XDS Document Consumer queries.

ferently from us, MeDShare and SmartProvenance use a public blockchain which can pose significant issues for data compliance (see, e.g., this EU report [17]).

3. Healthcare Data Provenance System

Figure 2 illustrates the proposed provenance blockchain system architecture. To ensure *modularity* and *interoperability* with current EHR systems, we designed a *PROV proxy* that can transparently intercept XDS messages by enriching them with provenance data managed with a blockchain smart contract.

We developed the proxy in Java and integrated with the SpiritEHR system and with FHIR resources (e.g. patient devices) via the FHIR framework HAPI[6]. We use Hyperledger Fabric[7] to develop the blockchain smart contract.

Due to the high sensitivity of health data and the vast patient-consent requirements needed to handle EHR securely and lawfully, we opt *not to store health documents on blockchain*, but only privacy-preserving digital evidence of the documents. Given a document *doc* in a standardised medical format (e.g.

[6]HAPI FHIR - `https://hapifhir.io`. The HL7 API (HAPI) has been started in 2001 by `https://www.uhn.ca` and it is the most complete open source implementation available for the HL7 specifications, including FHIR.

[7]See `https://www.hyperledger.org/projects/fabric`

CDA or DICOM), our system creates and manages *provenance pairs* of the form

$$\langle h_{doc}, Prov_{doc} \rangle$$

where h_{doc} is a unique, tamper-proof signature of the document, and $Prov_{doc}$ is the provenance annotation in the PROV format for the document. We obtain the h_{doc} as the hash of the digital signature of the canonicalised version of the document[8].

Notably, a CDA document *doc* can group multiple CDA documents; in order to define provenance for these documents, CDAs are divided into (a list of) sections doc_{sec} via the DS4P method [18] according to patients' privacy consent and legal requirements. For each section doc_{sec}, the canonicalisation, signing and creation of the PROV document are carried out individually by obtaining a list of tuples $\langle h_{sec}, Prov_{sec} \rangle$.

In the following, we present the PROV template for health documents (Section 3.1), and interactions of the Proxy with the smart contract (Section 3.2).

3.1. Provenance for Health Documents

We describe here the constituent parts of the PROV documents $Prov_{doc}$.

Entity. The entity models the health document of interest. An entity contains the *document type* (e.g. CDA, PDF, DICOM), the reference to the *organisation* from where the document can be retrieved, and the *hash* of the digital signature.

Activity. The activity corresponds to the operation executed within the EHR system, i.e. one among the standardised IHE XDS actions.

Agent. The agent models the subject bearing the operation on the health document. Agents are identified from the EHR system based on the user identity attributes. An agent contains the identity attributes used by the federated

[8] Canonicalisation prevents that document formatting (e.g. spurious spaces, tabulations, or namespaces) could alter the hash value of the signature. For instance, two XML files are considered equivalent even when carriage returns are different. Therefore, canonicalisation ensure that lexical formatting and special characters does not impact the signing process.

identity providers within the EHR system; this includes, e.g., the role in the organisation such as doctor or pharmacist.

Relationships. The relationships map the semantic dependencies that led a subject to operate on a document. In particular, we define: 1) `wasGeneratedBy` which states the *time* of when the entity was created by an activity, and the *purpose* reporting the medical grounds for which the transaction happened; 2) `wasAssociatedWith` and `wasAttributedTo` which state how, resp., the activity and the entity relates to the agent; 3) in case of CDA sections, `used` which states that the activity utilised the CDA master to operate on an internal section, and `wasDerivedFrom` which states that such section was obtained from the CDA master.

3.2. Managing Health Document Provenance with Blockchain

The proxy intercepts health documents exchanged within XDS messages and, by computing the hash h_{doc}, it invokes the smart contract to either store or retrieve the provenance annotation.

Canonicalised Signing. For an intercepted document *doc*, the proxy performs the procedures of canonicalisation and digitally signing via the dedicated *aDES technique—viz., XaDES [19], PaDES [20], and CaDES [21] for, respectively, PDF, CDA and any other type. The result h_{doc} is the hash of (the signature of) canonicalised document that is used by the smart contract to uniquely identify and proof the integrity of *doc*.

Provenance Creation and Storage. When EHR systems exchange messages that modify a patient's document *doc* (i.e. create or update actions), we store a new provenance tuple as per the steps outlined in Algorithm 1.

After computing the hash h_{doc} (Line 3), the proxy retrieves the information on the agent *ag* (inherited from the EHR system) executing the XDS action *act* (e.g. CREATE) on the entity *doc* . The proxy retrieves location l from the EHR system (Line 4). The function *storeProv()* of the smart contract *sc* is then

Algorithm 1: Provenance creation and storage.

```
1  Function createProvDoc(Agent ag, Action act, HCDocument doc):
       Data: An agent ag; an EHR action act; an health document doc
       Result: Storage in the smart contract sc of ⟨h_doc, Prov_doc⟩ and, for each section sec
               of CDA doc, of ⟨h_sec, Prov_sec⟩
2      begin
           /* Canonicalisation and digital signing                          */
3          h_doc = getCanonicalisedHash(doc);
           /* Retrieve location of doc from the EHR                         */
4          l_doc = getLocation(doc);
           /* Create and store Prov_doc                                     */
5          sc::storeProv(ag, act, l_doc, h_doc);
6          if doc instanceof DataTypes.CDA then
               /* Sectioning of internal CDA documents                      */
7              List<CDASection> secList = cdaSectioning(doc);
8              foreach sec : secList do
9                  h_sec = getCanonicalisedHash(sec);
10                 l_sec = getLocation(sec);
                   /* Storage of PROV sections                              */
11                 sc::storeProv(ag, act, l_sec, h_sec, h_doc);
12             end
13         end
14     end
```

invoked to create the PROV document $Prov_{doc}$ and store it with h_{doc} within the blockchain (Line 5).

When the document *doc* is a CDA, the DS4P sectioning is applied (Lines 6-7). For each section *sec*, the corresponding hash h_{sec} and PROV document $Prov_{sec}$ is created and stored by linking it back to the master document *doc* via its hash h_{doc} (Lines 9-11). At the end of the execution of the smart contract, we persist the provenance of an health document *doc* as JSON objects using the key-value store of Hyperledger Fabric blockchain.

Provenance Retrieval. Once a health document *doc* is retrieved from an EHR system, the proxy intercepts the message and, by calculating the corresponding hash h_{doc}, invokes the function *getProv()* of the smart contract to obtain the relevant PROV documents. When the retrieved document is a CDA section, the smart contract also returns the master PROV document. The proxy will

then perform the necessary actions required by the EHR system (e.g. decorating the XDS message with PROV information or creating the corresponding FHIR resource) to guarantee transparent message exchanges.

4. Results

The transparent integration of provenance management in EHR systems empower multiple patient-centric services (Section 4.1), without introducing significant time overhead to typical EHR operations (Section 4.2).

4.1. Using Provenance for Patient-Centric Services

Data reconciliation. The distribution of healthcare systems led patients' EHRs to be scattered across multiple locations, causing eHealth systems to use sophisticated data discovery routines to fully retrieve patients' data. This challenge is amplified by mobile eHealth and mobility of patients, e.g. in cases of clinical encounter spanning over several healthcare facilities.

Our provenance system can trace complex health documents, such as CDAs, across different facilities. The use of provenance relationships connecting all the constituent parts of health documents, such as CDA sections, allows our system to support the identification (and, if authorised, the consequent retrieval) of the entirety of patients' EHRs. Furthermore, the integration with FHIR-based resources allows us to handle complex patients mobility situations where people use mobile devices to create and authorise document sharing of several medical encounters in different facilities[9].

Decision Making. Provenance documents can enable better-informed decision-making processes by providing contextual information on, for example, where and when specialist treatments were carried out [2]. Similarly, provenance documents can improve traceability and reproducibility of medical studies [22].

[9] See, e.g., the mobility use case in Section 45.4.2.1.1 of `https://www.ihe.net/uploadedFiles/Documents/ITI/IHE_ITI_Suppl_mXDE.pdf`

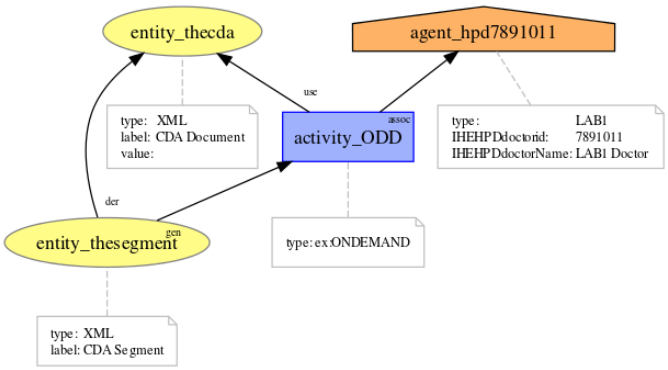

Figure 3: **Patient's data provenance graph**: Example of PROV graph presenting how a CDA document is used. Relationships `assoc`, `der`, and `use` correspond to the relationships `wasAssociatedWith`, `wasDerivedFrom`, and `used` respectively, as introduced in Section 3.1.

Patient Consent Auditing. The access, sharing and secondary use of patients' data in EHR systems critically rely on authentication and authorisation procedures. Our provenance system enables auditing procedures to validate enforcement of patient-informed consent requirements [23].

Importantly, the use of PROV allows the creation of provenance graphs to present to patients the audit on how EHR systems use their data. For instance, Figure 3 illustrates the provenance graph describing an entity ("`thecda`") that has been derived by another CDA section ("`thesegment`") with an On-Demand XDS activity ("`ODD`") triggered by the doctor with identifier 7891011 of the Healthcare Provider Directory (HPD) of the laboratory "`LAB1`".

4.2. Performance evaluation

This evaluation aims to measure the time required by our system to manage provenance of real-world CDA health documents. The three core operations of our systems, described in Section 3.2, are individually evaluated.

We base the tests on Hyperledger Fabric 1.0.1 deployed in the Microsoft Azure Cloud. The blockchain network is formed by four nodes (three nodes representing two healthcare organisations, the remaining providing the blockchain system functionality) deployed in four virtual machines Intel(R) Xeon(R) CPU

E5-2673 v3 @ 2.40GHz. We run tests as Java applications and use Apache Santuario[10] as the canonicalisation library. We ran the client application on a MacBook pro 2.9 GHz Intel Core i9. All the code is available on GitHub[11].

We executed the tests over a dataset of 750 Continuity of care CDA (CCDA) document samples; the dataset is publicly available online[12] and, to authors' knowledge, the biggest among those not containing protected health information. Included documents mimic real-world cases and were provided by organisations as the NIST, HL7, and by independent vendors.

Figure 4 shows the results of the performance evaluation. Specifically, Figures 4a and 4b report the *Canonicalised Signing* and *Provenance Creation and Storage* times, respectively; in both cases, the average is around 2msec (a thread asynchronously waits for transaction execution receipts). The spikes in Figure 4b are due to the Java garbage collector used by the Fabric Java SDK; this was confirmed by analysing the Java system time using JVisualVM. Figure 4c shows the sequential *Provenance Retrieval* of the provenance of every CCDA; the retrieval operation is around 200msec in average.

The evaluation shows that in a distributed concurrent setting our blockchain system will introduce approximately 4msec overhead to XDS messages creating new EHR data, and approximately 200msec to XDS messages retrieving EHR data. An enterprise XDS system deployed on cloud, e.g. the SpiritEHR, takes up to 3sec to retrieve a health document[13], with significant increments for multi-parameter document searches [24]. Therefore, we can conclude that our provenance management does not impact the overall performance of eHealth systems.

[10] See: `http://santuario.apache.org`. Apache Santuario is an open source implementation of XML Canonicalisation and Signature for Java, interoperable with many other libraries and tools.

[11] `http://github.com/mascanc/ProvenanceContract`.

[12] `https://github.com/jmandel/sample_ccdas`

[13] The times reported do not include other security services such as integrity check of the document and access control that may add additional milliseconds.

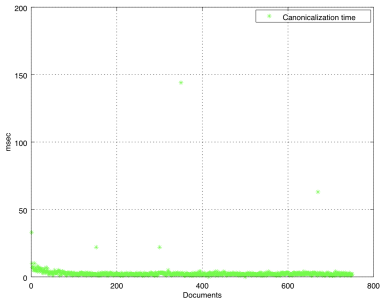

(a) Canonicalised Signing (avg. 2msec)

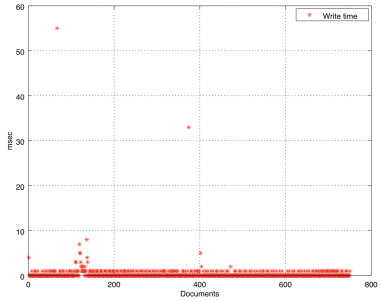

(b) Provenance Creation and Storage (avg. 2msec)

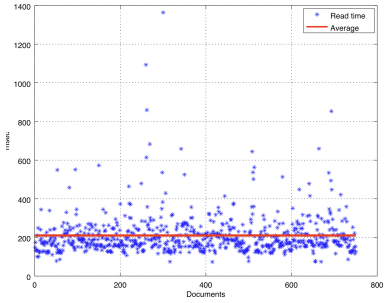

(c) Provenance Retrieval (avg. 200msec)

Figure 4: **Performance evaluation results**: execution time of the core system functions over the 750 CCDA samples.

5. Discussion

This section comments on some of the main challenges emerging from the adoption of our blockchain system in the eHealth domain: 1) interoperability with standards and legacy, 2) the role of data privacy.

Interoperability. The critical requirement for EHR systems is interoperability: limiting the exchange of health data across organisations may endanger the safety of patients. Multi-year efforts of international consortia such as IHE and HL7 have led to ubiquitous eHealth services accessible across different facilities. Therefore, we base our system on the internationally widely used IHE/HL7

standards[14] to achieve interoperability.

Besides the IHE standards, several healthcare legacy protocols exist and are currently in operation daily, such as the *secure mail exchange*[15]. Users can adapt our system to this standard by using a mail gateway so to be able to intercept and interpret mail disposition notifications.

Also, our system natively supports FHIR via the HAPI framework so can map FHIR resources to PROV as per available guidelines[16].

Privacy. The confidentiality of health documents is of utmost importance, and, fundamentally, the use of blockchain must not pose any impediments to protect patients' privacy.

Storage of EHR within the blockchain produces the first privacy threat. The use of encryption and semi-identifiers for storage of medical data within blockchain have been investigated [12], and it appears that these approaches may not ensure adequate levels of confidentiality. This inadequacy is due to, e.g., brute force attacks against encrypted data or linkage attacks against semi-identifiers [25]. By-design our approach avoids writing any personal data directly into the blockchain while ensuring its availability and reconciliation.

Our approach builds on the strong security controls of the underlying EHR systems to protect patients' privacy from system misuse. First of all, only subjects already in possession of a health document (hence authorised by the EHR system) can query and retrieve the corresponding provenance from the blockchain by computing the canonicalised hash of the document. Also, the provenance documents stored on the blockchain are designed only to contain *opaque identifiers*[17] (e.g. URI or URL according to the health document type).

[14] For example, the Austrian national healthcare exchange, ELGA, `http://www.elga.gv.at`, and the European Patient Summaries and Electronic Prescription exchange, `https://ec.europa.eu/cefdigital/wiki/display/EHOPERATIONS/`. More generally, IHE profiles have been mentioned by the EU Commission Decision 2015/1302 as eligible for being used in eHealth EU public procurements.

[15] See, e.g., the mail-based health data exchange by the Direct Project `http://directproject.org`

[16] `https://www.hl7.org/fhir/provenance-mappings.html`

[17] That is, identifiers that do not expose information about the referred data.

These identifiers will be available only to authenticated users (typically doctors) and will only give authorised user access to the linked EHRs according to the enforced patient-informed consent.

Finally, it is worth noting that patients can opt-out from the creation of provenance documents by submitting the corresponding dissent to the underlying EHR systems. The system will enforce the policy during standard XDS handling. Since the provenance documents do not hold any patient identifiable information, their cancelation should not be needed, although the blockchain legal framework is still in development [17].

6. Conclusions

In this paper, we proposed a blockchain-based system for managing provenance of health documents that can seamlessly integrate with existing EHR deployments. We base the provenance tracking process on the web open standard PROV. It integrates a digital signature process to ensure by design privacy-aware provenance management of healthcare documents on the blockchain. The provenance tracking can manage any medical document in a standardised format, including complex ones such as HL7 CDA; the fact that the system architecture builds on IHE/HL7 standards ensures interoperability with current eHealth systems and technical sustainability [26]. The integration with FHIR-based resources allows operators to use the system for multiple purposes, including data reconciliation across different organisations (e.g. to support the creation of longitudinal health records) and patient consent auditing. The experimental evaluation with enterprise-level blockchain on the cloud has shown that our system introduces up to 6% overhead to typical EHR access operations and 0.1% in the creation of new medical documents.

In future work, we aim to exploit the PROV documents to enable reproducibility of clinical research [14]. A pilot project jointly created with Microsoft [27] has already started. The objective of the pilot is to evaluate the metrics of usability, performance, and user acceptance of the proposed solution.

Summary Points

What was already known.

- Existence of data provenance standards and how to retrieve provenance records: W3C PROV[18], IHE mXDE[19], and HL7 Provenance resource[20]
- Legacy methods to reconstruct the history of a data access: audit trails (ATNA profile[21]), and Non Repudiation services[22]

What this study added.

- Legacy methods require an enormous work of human coordination and data mining to reconstruct the data provenance by accessing centralised databases (such as Audit Record Repositories or Evidence Storage Systems of Hospital Information Systems). We proposed and implemented decentralised means that collect and make provenance data available. Such a system enables both patients, doctors, and prosecutors to have immediate access to provenance records mediated by the underlying existing Hospital Information System access control systems.
- The usage of the system adds a minor overhead to data manipulation transactions (in the order of milliseconds) and seamlessly integrates with existing deployments.

[1] M. G. Kahn, T. J. Callahan, J. Barnard, A. E. Bauck, J. Brown, B. N. Davidson, H. Estiri, C. Goerg, E. Holve, S. G. Johnson, A harmonized data quality assessment terminology and framework for the secondary use of electronic health record data, eGEMs 4 (1).

[18] https://www.w3.org/TR/prov-overview/

[19] https://wiki.ihe.net/index.php/Mobile_Cross-Enterprise_Document_Data_Element_Extraction

[20] https://www.hl7.org/fhir/provenance.html

[21] https://wiki.ihe.net/index.php/Audit_Trail_and_Node_Authentication

[22] https://ec.europa.eu/cefdigital/wiki/display/EHOPERATIONS/10.01.+Handle+the+non-repudiation+mechanism

[2] V. Curcin, E. Fairweather, R. Danger, D. Corrigan, Templates as a method for implementing data provenance in decision support systems, Journal of Biomedical Informatics 65 (2017) 1 – 21.

[3] V. Curcin, Embedding data provenance into the learning health system to facilitate reproducible research, Learning Health Systems 1 (2) (2017) e10019.

[4] A. Hasselgren, K. Kralevska, D. Gligoroski, S. A. Pedersen, A. Faxvaag, Blockchain in healthcare and health sciences—a scoping review, International Journal of Medical Informatics 134 (2020) 104040.

[5] P. Missier, K. Belhajjame, J. Cheney, The W3C prov family of specifications for modelling provenance metadata, in: Proceedings of the 16th International Conference on Extending Database Technology, EDBT '13, Association for Computing Machinery, New York, NY, USA, 2013, p. 773–776.

[6] IHE, IHE technical framework, Webpage (2017).
URL `http://www.ihe.net/uploadedFiles/Documents/ITI/IHE_ITI_TF_Vol2b.pdf`

[7] HL7, CDA: Clinical Document Architcture, Webpage (2009).
URL `http://hl7.org`

[8] National Electrical Manufacturers Association, DICOM: Digital Imaging and COmmunications in Medicine, Webpage (2019).
URL `https://www.dicomstandard.org/`

[9] HL7, FHIR: Fast Healthcare Interoperability Resources, Webpage (2018).
URL `http://hl7.org/fhir`

[10] X. Liang, J. Zhao, S. Shetty, J. Liu, D. Li, Integrating blockchain for data sharing and collaboration in mobile healthcare applications, in: Personal, Indoor, and Mobile Radio Communications (PIMRC), IEEE, 2017, pp. 1–5.

[11] A. Ramachandran, M. Kantarcioglu, Smartprovenance: A distributed, blockchain based dataprovenance system, in: Data and Application Security and Privacy, CODASPY '18, ACM, New York, NY, USA, 2018, pp. 35–42.

[12] CACM Staff, Access controls and healthcare records: who owns the data?, Commun. ACM 62 (7) (2019) 41–46.

[13] P. Zhang, J. White, D. C. Schmidt, G. Lenz, S. T. Rosenbloom, Fhirchain: Applying blockchain to securely and scalably share clinical data, Computational and Structural Biotechnology Journal 16 (2018) 267 – 278.

[14] T. McGhin, K.-K. R. Choo, C. Z. Liu, D. He, Blockchain in healthcare applications: Research challenges and opportunities, Journal of Network and Computer Applications 135 (2019) 62 – 75.

[15] Q. Xia, E. B. Sifah, K. O. Asamoah, J. Gao, X. Du, M. Guizani, MeDShare: Trust-Less Medical Data Sharing Among Cloud Service Providers via Blockchain, IEEE Access 5 (2017) 14757–14767.

[16] L. Moreau, B. Clifford, J. Freire, J. Futrelle, Y. Gil, P. Groth, N. Kwasnikowska, S. Miles, P. Missier, J. Myers, B. Plale, Y. Simmhan, E. Stephan, J. V. den Bussche, The open provenance model core specification (v1.1), Future Generation Computer Systems 27 (6) (2011) 743–756.
URL https://eprints.soton.ac.uk/271449/

[17] EU Parliament, Blockchain and the general data protection regulation: Can distributed ledgers be squared with european data (2019).
URL https://www.europarl.europa.eu/RegData/etudes/STUD/2019/634445/EPRS_STU(2019)634445_EN.pdf

[18] HL7, Implementation guide: Data segmentation for privacy, Webpage (2014).
URL https://wiki.hl7.org/HL7_DS4P_Document_Library

[19] W3C, XAdES, XML advanced electronic signatures (2003).
URL https://www.w3.org/TR/XAdES/

[20] ETSI, Pdf advanced electronic signature profiles (2009).
URL http://www.etsi.org/deliver/etsi_ts/102700_102799/10277801/01.01.01_60/ts_10277801v010101p.pdf

[21] ETSI, Cms advanced electronic signatures (2013).
URL https://www.etsi.org/deliver/etsi_ts/101700_101799/101733/02.02.01_60/ts_101733v020201p.pdf

[22] S. S. Sahoo, J. Valdez, M. Kim, M. Rueschman, S. Redline, Provcare: Characterizing scientific reproducibility of biomedical research studies using semantic provenance metadata, International Journal of Medical Informatics 121 (2019) 10 – 18.

[23] H. Petritsch, Break-Glass - Handling Exceptional Situations in Access Control, Springer, 2014.

[24] G. Duftschmid, C. Rinner, M. Kohler, G. Huebner-Bloder, S. Saboor, E. Ammenwerth, The EHR-ARCHE project: satisfying clinical information needs in a Shared Electronic Health Record system based on IHE XDS and Archetypes, International journal of medical informatics 82.

[25] C. Esposito, A. De Santis, G. Tortora, H. Chang, K. R. Choo, Blockchain: A panacea for healthcare cloud-based data security and privacy?, IEEE Cloud Computing 5 (1) (2018) 31–37.

[26] H. Aranha, M. Masi, T. Pavleska, G. P. Sellitto, Securing mobile e-health environments by design: A holistic architectural approach, in: 2019 International Conference on Wireless and Mobile Computing, Networking and Communications, WiMob 2019, Barcelona, Spain, October 21-23, 2019, IEEE, 2019, pp. 1–6.

[27] D. Houlding, Improve collaborative care and clinical data sharing with blockchain, Webpage.

URL https://azure.microsoft.com/it-it/blog/improve-collaborative-care-and-clinical-data-sharing-with-blockchain/