HealthNet™: A Blockchain Enabled Population Health Platform for Enterprise Health Systems

By Consensus Networks

1. Introduction

Despite rapid advancements in information technologies, half of the 30 billion health data transactions that occur in the United States every year continue to be performed via the fax machine to avoid Electronic Medical Record interoperability issues[1]. Moreover, providers often have little guarantee that the record they receive is complete and not lacking pertinent information about a patient such as medication allergies. The lack of interoperable IT infrastructure is estimated to cost the U.S healthcare system $18.6 Billion and 500,000 preventable deaths per year.[2] Moreover, in 2016, researchers at John Hopkins found that preventable medical errors are now the third leading cause of death in the U.S, representing nearly 10% of all deaths yearly.[3] While there are many contributing factors in fatal medical errors, the lack of interoperable IT systems and resultant incomplete medical information plays a large role. The inability to effectively share data has also delayed the U.S Healthcare System's transition to a preventative, value-based care model. With poor access and incentive structures in place to promote data sharing, effectively coordinating care can be near impossible, or seen as too costly or resource-intensive in the short term despite the long-term patient benefits.

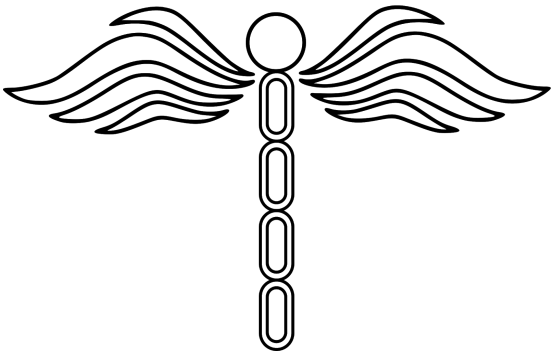

Figure 1: HealthNet Logo

In an effort to improve population health management and overcome the difficulties of a fragmented healthcare system, Consensus Networks has created HealthNet: a blockchain enabled universal patient index (UPI) that allows providers

[1]www.getreferralmd.com/2016/08/30-healthcare-statistics-keep-hospital-executives-night/
[2]www.medium.com/mit-media-lab-digital-currency-initiative/medrec-electronic-medical-records-on-the-blockchain-c2d7e1bc7d09
[3]www.washingtonpost.com/news/to-your-health/wp/2016/05/03/researchers-medical-errors-now-third-leading-cause-of-death-in-united-states

to securely track, access, and manage patient records through a collection of APIs designed to integrate with existing Electronic Medical Record (EMR) systems. Authenticated providers can transfer patient medical information using HL7 FHIR and an encrypted audit trail of metadata for updates to a patient's record is created on the Factom blockchain to verify they have a patient's most complete record. The following technical milestones have been completed to date: successfully normalizing data transacting on the network to FHIR, novel cryptographic implementations for protecting records and users private keys, creation of a blockchain network with public or private functionality using the Factom blockchain to track encrypted metadata for updates to patient records, and commenced R&D on encrypted hardware components for secure role access control.

2. Technical Overview

2.1 FHIR Standards

The Health Level Seven International (HL7) healthcare standards institute recently announced its fourth version of the Fast Healthcare Interoperability Resources (FHIR) standards. First released in 2014, FHIR is the healthcare industry's response to interoperability challenges by providing an increasingly agreed upon framework of data formats for interacting with patient records.[4] FHIR uses existing logical and theoretical models to make it easy to implement while still providing well defined formats for exchanging data between healthcare applications. All information exchanged using FHIR is defined as a resource with the following characteristics: a common definition and representation, a common set of metadata, and a human readable portion. Therefore, either a single resource, or collection thereof, can satisfy most use cases. FHIR not only makes transferring electronic medical records between legacy systems in hospitals easier, but also across different devices such as computers, tablets, or mobile.[5]

While FHIR promotes some level of interoperability between healthcare IT systems, it is not without shortcomings. For example, data often must still be mapped to or from the format of other, non-FHIR based IT systems, which can be problematic. Effectively transferring healthcare data requires achieving three different levels of interoperability: foundational, structural, and semantic interoperability.[6] Foundational interoperability refers to the successful exchange of data between two different platforms without a need for interpretation. Structural interoperability builds upon foundational interoperability and encompasses data transfer with defined formats for the exchanged information. Lastly, semantic interoperability refers to the exchange of data in a way that maintains a human readable format across the devices used in the transfer. Achieving

[4] https://www.hl7.org/fhir/overview.html
[5] https://healthitanalytics.com/news/4-basics-to-know-about-the-role-of-fhir-in-interoperability
[6] https://www.dre.vanderbilt.edu/~schmidt/PDF/blockchain-bookchapter-2018.pdf

all three levels of interoperability is vital to ensuring that data quality and safety are maintained in a HIPAA compliant manner and for physicians to derive clinical insight from a patient's healthcare information. Foundational and structural interoperability are a prerequisite for semantic interoperability, however, semantic interoperability also requires clinical domain knowledge. HealthNet can currently achieve functional and structural interoperability, and semantic interoperability can be reached with any FHIR native source a clinical domain expert would use. We intend to work alongside bioinformaticians within research teams at health to extend the abilities of HealthNet to attain semantic interoperability regardless of data format.

2.2 Data Storage:

While Blockchain is renowned for the security and immutability it provides the data it stores, it is an inappropriate storage solution for entire medical records. First, hashing or storing genuine medical records to a blockchain poses tremendous security risks to both patients and healthcare providers alike. While the record itself would likely remain uncorrupted if stored on a blockchain, unauthorized network users would have an increased ability to access sensitive information and violate patient privacy. Consequently, providers may be unable to maintain HIPAA or GDPR compliance.. Moreover, medical records are simply too large to be practically stored on a blockchain. Electronic health records hold tremendous amounts of data about an individual and can range anywhere from 1 MB - 5 GB in size depending on the file.[7] Blockchains are intended to record transactional data by design, making the storage large files unfeasible on many networks. For example, some Blockchains, like Bitcoin, have a capped block size of 1 MB.[8] Other ledgers like Ethereum theoretically have no block size limitations, but the computational power required to hash such a large volume of data to the blockchain would be impractical at scale.

A more appropriate solution is storing the actual electronic medical records within databases housed in HIPAA or GDPR compliant servers or in a secure cloud-based data warehouses. This is done already in many hospitals or Health Information Exchanges (HIEs) Doing so mitigates security and privacy concerns associated with HIPAA & GDPR and is a more practical solution for hospitals. Either of these systems are designed to encrypt and store large volumes of patient data in a manner compliant with regulatory bodies. Moreover, these are more cost efficient and pragmatic solutions for storing medical records. Any protective benefits that a decentralized, blockchain enabled solution offers for medical record storage is marginal compared to using distributed databases and improving security. Similar read and write permissions and levels of decentralization can be achieved at a far lower compute cost with more efficiency taking these measures as opposed to trying to 'blockchain' medical records. To put things in perspective,

[7] https://philip.greenspun.com/blog/2013/01/11/electronic-medical-records-dont-save-money/
[8] https://cointelegraph.com/news/satoshis-best-kept-secret-why-is-there-a-1-mb-limit-to-bitcoin-block-size

Ethereum is one of the most widely used and developed on blockchains, has been around for 5 years, and just recently reached a terabyte of data in size. As a result, the network has been experiencing tremendous bloat and the core developer team is now redesigning the entire network in hopes that it will be able to allow the network to scale.[9] By comparison, the volume of health data is growing at 48% per year and is expected to reach over 2000 exabytes by the end of next year.[10] That is over 2,000,000 times as much data as the entire Ethereum blockchain, much of which comes from the increasing size of electronic medical records. Hence, all medical records will be stored off chain in secure servers within health systems or Consensus Networks' secured SOC 2 data centers.

2.3 Factom as a Blockchain Selection

Multiple candidates were identified as potential blockchain network layer: Hyperledger Fabric and a private Ethereum or Factom Network implementation. These protocols provide the speed, flexibility, and need-to-know access to data demanded by HealthNet private network layer. API accessibility and functionality, compute requirements, and more were tested to determine the best candidate for HealthNet A comparison of the different private network characteristics for each protocol may be observed below in **Table 1**.

Table 1: Comparison of Private Network Characteristics

Characteristics	Hyperledger Fabric	Ethereum	Factom
1. Public or Private	No	Yes	Yes
2. Smart Contract Functionality	Yes	Yes	Yes
3. On Channel Encryption	Yes	Yes	Yes
4. MultiChannel Enabled	Yes	No	Not Required
5. Flexible Payload	Yes	Yes	Yes
6. Extended Transport	Partial	No	No
7. Interoperability with existing data structures	Yes	Some	Yes
8. Availability of SDKs in common languages	Yes	No	Yes

[9] https://hackernoon.com/the-ethereum-blockchain-size-has-exceeded-1tb-and-yes-its-an-issue-2b650b5f4f62

[10] https://cmte.ieee.org/futuredirections/2018/05/18/the-future-of-health-care-is-tied-to-big-data/

As discussed in the preceding paragraphs, HealthNet is capable of exhibiting both private and public functionality. To minimize interoperability concerns, a blockchain that was operable both publicly and privately was desired, which Ethereum and Factom were capable of supporting. Smart contract functionality was also desired to help manage the patient-provider-billing-insurerer loop. Although smart contracts were not a design criteria of Factom, they are possible and all three of the analyzed protocols met this criteria. Ethereum is not capable of maintaining multiple private blockchains simultaneously, however, multichannel enablement is supported by Hyperledger Fabric and not required for Factom due to its 'multi-chain' functionality. Factom allows for multiple chains of data to be created within the ledger itself, removing the need to maintain separate blockchains for various parties and simplifying implementation.[11] Health systems already have legacy information technology stacks that are critical to their daily operations and are reluctant to perform major overhauls. Hence, the ability to incorporate blockchain without requiring hospital IT teams to learn a new programming language or modify their stack was crucial. Factom and Hyperledger both have clients and software development kits (SDKs) in a range of different commonly used programming languages, making them far easier to integrate.

Figure 2: Factom Logo

Taking all of the above into consideration, Factom was selected for both the private and public blockchain layers for HealthNet. Factom is designed to serve as a public audit trail for contracts and other assorted datum that is also tethered to both the Bitcoin and Ethereum blockchains. Factom allows users to create their own chains for a project which they can then populate with relevant information and all transactions are then submitted to the network for validation and resolved to the public Factom Blockchain. Moreover, after blocks are added to the Factom blockchain, proofs for the network state are then written to the Bitcoin and Ethereum blockchains with appropriate timestamps. Hence, Factom essentially has an equal or greater level immutability to the Bitcoin or Ethereum blockchains while minimizing the cost of writing the transaction.[12] Factom charges $.0001 per

[11] https://coincentral.com/factom-beginner-guide/
[12] https://www.factom.com/assets/docs/Factom_Whitepaper_v1.2.pdf

write, compared with approximately $3.00 per kb of data written using Ethereum.[13] Low cost write operations grant network participants flexibility with what type of data is written to the blockchain and provides them with a predictable pricing model. Factom's ability to leverage Ethereum and Bitcoin's immutability was also a critical factor in selecting it as a public network layers. To date, Bitcoin is generally considered the most secure distributed digital ledger and has suffered minimally from internet attackers.[14] As a result, it is likely the safest place to hold a hashes for transactions corresponding to the current state of the HealthNet Network.

2.4 HealthNet Architecture:

HealthNet is a blockchain network with private and public functionality that is built on top of the Factom protocol. A private, permissioned blockchain network layer is required for HealthNet to securely and effectively manage patient data for two reasons: speed and the need-to-know basis on which parties can access information. Only permissioned users can gain access to this type of network, meaning that only parties that need access to a patient's data (i.e physicians, insurers, patients, etc.) will be able to join. This implies that there is an inherent 'trust' between all members on the HealthNet network, even if all of the providers do not know one another. The pseudo-trust benefits that a private blockchain network layer adds in conjunction with having fewer parties on the network also increases speed. Seeing as only trusted parties are allowed to participate on the HealthNet network, nodes can arrive at consensus much faster and transactions can be recorded rapidly in a manner using little computational power.

Private, permissioned networks also have a lesser degree of immutability allowing trusted members on the network to elect to change previous entries on the ledger as long as the remainder of the nodes come to consensus. While this creates room for an internal malicious actor to manipulate data on the ledger, it allows other trusted participants to have the flexibility to change entries if needed. For example, the network could agree to a proposal to remove entries for a deceased patient from the network to free up storage space on the nodes hosting the ledger. Permitting this type of flexibility is critical for maintaining an accurate ledger when accounting for the complexities of healthcare data. While the flexibility and speed offered by a private network are essential for the transactional aspects of HealthNet, inclusion of a public, proof-of-work network layer can also benefit HealthNet for security purposes.

Public, proof-of-work blockchain networks are unparalleled in terms of security, accuracy, and immutability when compared to other types of blockchain networks. The extensive nodal networks competing to validate blocks on these blockchains makes it nearly impossible for a single participant to manipulate the ledger due to compute power limitations, making them ideal

[13] https://factomize.com/
[14] https://theoutline.com/post/1618/how-hackable-is-bitcoin?zd=1&zi=hiwmllcg

for maintaining an immutable, time-stamped audit trail of data. As such, a public Factom network layer was also incorporated into the network design to serve as an anchor point for the private network layer(s) on HealthNet so that the overall network state is maintained. Data privacy is maintained as only a cryptographic hash generated from information regarding the state of the private network layer(s) is recorded on the public network layer. These hash serve as a timestamped anchor for the authenticity of the ledger up to that point and as a way for the private network(s) to query for data across the entire network. Hence, if a malicious actor were to manipulate information on the private network, the network participants could use the anchor point to restore the state of the private network to the last known time point that represented an accurate version of the ledger. Any changes made by a malicious actor could be negated and the network would be restored accordingly. Additionally, if a provider needed to request information for a patient not in their system, they could use this as a way to query the entire network and see if any provider on the network as of the most recent network state recording had the record to send. A visualization of a potential network architecture may be viewed below in **Figure 3**.

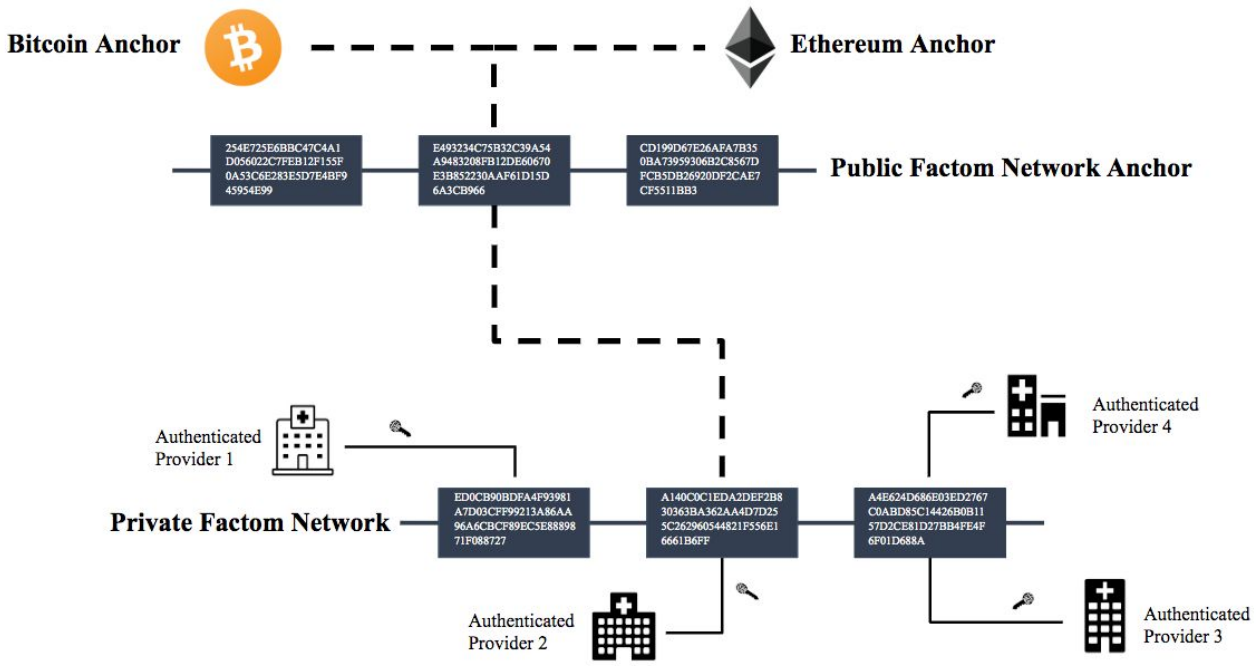

Figure 3: Potential HealthNet Network Architecture

2.5 Data Flow and Network Operations

For security and privacy reasons, patient data is not stored on the blockchain as it is accessible by a multitude of participants. Instead, patient records are stored in encrypted form in secure, HIPAA/GDPR compliant servers. The servers will not have access to the key material necessary to decrypt the records, but rather be used as a secure storage medium to facilitate record exchange as authorized by each patient. Any given patient record may be decrypted by only the

health provider who created the record, the patient to whom the record corresponds, and any other health providers (if any) to whom the patient explicitly grants permission to have access to the data. In addition to records being encrypted, they are also authenticated by means of cryptographic signatures to ensure that all updates come from a trusted source.

Blockchain, in this design, functions as a sort of immutable bulletin board that keeps track of all updates a patient's medical record. When a health provider retrieves encrypted records from HealthNet servers, it verifies that all updates to the patient records that it receives were also reflected on the blockchain. Each value recorded on a blockchain reveals nothing about the patient, the health provider, or any other information that could compromise patient privacy. Only authenticated providers who were granted access to data of a specific patient will be able to determine whether a value stored on the blockchain corresponds to the patient in question. Even in such cases, the provider can only use the blockchain for a confirmation that the patient data she received matches the cryptographic digest stored on the blockchain.

A significant component of the network design relates to key management to ensure that sensitive medical data is only accessible to authorized parties, similar to the way health data is accessible to health providers today based on patient consent. Each patient and health provider will have unique identifiers in the system and any credentials and permissions they have will be associated with these IDs. For any given patient, there will be multiple encryption keys; one for each unique set of permissions that grant access to his or her data. For example, suppose a patient has records accessible only by her primary physician, records accessible only by a specialist that she sees, and records accessible by both. With this proposed schema, there would be three different types of permissions and, consequently, three corresponding keys, with the physician having access to the first and third keys and the specialist having access to the second and third keys.

When a new patient's record is created or modified, the update is cryptographically signed by the corresponding health provider (and the patient herself when meaningful). The update is subsequently encrypted with the appropriate key and stored in the appropriate servers in encrypted form. A separate token is sent to the blockchain which permits verification of the update in the future and also allows one to verify completeness of the retrieved records. The token includes the patient and provider IDs and the digest (computed using a cryptographically strong hash function) of the update protected with the patient-specific and time-dependent key. The token reveals no information of any kind to anyone without proper key material. Also, because blockchain tokens are transmitted without revealing information about the originating health provider, it is not possible to learn any partial information about patients to whom the updates correspond.

The most non-trivial component of the design stems from the fact that permission can change over time when a patient grants (or revokes from) a health provider access to some of his or her records that another provider produced. With permission changes, there might be a need to re-key (and consequently re-encrypt records) to enforce the property that no party has access to more data than what the patient specified. This presents a conflict with the immutable nature of blockchains, which we resolve by separating key material used for record encryption and token protection on blockchain. The blockchain uses date-dependent keys for each patient which do not change over time and are computed using two hash chains to permit access to ranges of dates.

An in depth explanation of the way patient data is stored is included below. To increase system reliability, the key material associated with each patient does not have to be reliably maintained by the patient himself. Instead, it is backed up to the server properly encrypted with the patient's master key. The patient's master key is stored at their end on a secure hardware token, with additional provisions in place to permit its recovery and replacement in the case of its loss by the patient. Also, each key used for record encryption is recorded with structured metadata, *desc*, which specifies the set of providers to whom the data is accessible and other data characteristics such as the range of dates to which it corresponds. The use of this information permits for automatic processing of permission changes when, e.g., a patient gives consent to share certain types of his/her records with a health provider.

2.5.1 Patient Creation:

When a new patient is added to the system, the patient will be assigned a unique patient ID. The patient creates a public-private key pair for signing (pk_S; sk_S) and another symmetric encryption (i.e., AES) key sk_E for key management purposes. The binding between patient ID and his/her public key pk_S is reliably stored by the HealthNet servers, while sk_S and sk_E are private patient's keys stored at their end (on a smart card or other hardware security module).

We also create a separate symmetric encryption key k_{info} for protecting information about the patient, which will be accessible to all providers who work with the patient. This key will be used to guard the patient's personal information such as address, phone, etc., possibly including the list of known allergies. This key is stored together with its description on HealthNet servers encrypted under sk_E, i.e., the servers store $Enc_{skE}(k_{info}; desc_{info})$. The description $desc_{info}$ associated with the key needs to specify that this is a special-purpose key with no originating provider. The set of providers to whom this key is accessible is initially empty and the range is set to be unlimited.

Lastly, the patient also creates keys K_1 and K_2 for protecting information stored on the blockchain. These keys are private to the patient and can also be stored encrypted (under sk_E) at the HealthNet servers. The patient's (encrypted) key material is marked with the patient's ID and

cannot be retrieved or updated by anyone other than the patient during normal system functioning (i.e., non-emergency situations).

2.5.2 Patient Public-Key Query:

Upon providing a patient ID, a provider is able to obtain the patient's public key from HealthNet servers.

2.5.3 Provider Creation:

When a new provider is being established or added to the system, it is assigned a unique provider ID. The provider creates a public-private key pair for signing (pk_{P} ; sk_{P}). The binding between the provider ID and the public key pk_{P} is reliably stored by the HealthNet servers while the private key is only known to the provider. If the provider desires key backup capabilities the way they are implemented for patients, this could also be added to the system.

2.5.4 Provider Public-Key Query:

By specifying a provider ID, another provider is able to obtain the public key associated with the specified provider ID from HealthNet servers.

2.5.5 First Patient Visit:

If a patient visits a new provider for the first time such that there is no pre-established key material for the records generated by the provider for this specific patient, the patient and the provider generate a new symmetric (i.e., AES) encryption key k and the corresponding description *desc*. In particular, the description sets the current provider to be the originating provider with no other providers to whom the data is accessible for a typical health provider. The start date is the current date and the end date is open. The patient stores this new key together with its description in encrypted form to HealthNet servers. That is, the client stores $Enc_{\mathrm{skE}}(k;$ *desc*) with the servers. The patient also generates from K_1 and K_2 and shares with the provider $K_{d1}^{(1)}$ and $K_{d2}^{(2)}$, where *d1* is the current date and *d2* is a future date within a fixed time interval from *d1*. Lastly, the patient shares k_{info} with the provider and updates its corresponding description $desc_{\mathrm{info}}$ to include the current provider on the list of providers with whom the key is shared.

2.5.6 Updating Patient Data:

The provider needs to retrieve the key k it uses for encrypting the patient's data. Note that the provider might have more than a single key for that patient (each key is associated with a different set of permissions) and needs to identify the proper key k to use for that type of data. Once the provider has data (i.e., a FHIR record) with which it would like to update the patient's records, it will sign the data. The signed data should include the patient ID, the provider ID, and the date when the update is issued. Depending on the nature of data update, the provider might

also collect the patient's signature on the data that the patient can certify (e.g., the fact that the patient had an appointment of a particular type with that provider on a specific day). These data together with the corresponding signatures form an update denoted by C.

The provider encrypts C with k, sends $Enc_k(C)$ to HealthNet servers, and instructs them to associate this update with the ID of the patient and with key hash $h(k)$. In addition, the provider forms record D for the update C as specified below and places D on the blockchain.

When patient updates are placed on the blockchain, the format of the data is: $D = (d_1, d_2, d_3) = (r, PRF_{Kdate}(r) \oplus metadata, h(C))$ where r is a randomly chosen value of sufficient length (i.e., $r \in \{0,1\}^{\kappa}$ where κ is the security parameter), $PRF : \{0,1\}^{\kappa} \times \{0,1\}^{\kappa} \rightarrow \{0,1\}^{\kappa}$ is a pseudorandom function (instantiated with AES), *metadata* contains the information that a provider searching for a patient would need to know to identify C such as the patient ID, issuing provider ID, and the date, and h is a cryptographic hash function with the one-way property (meaning that it is not possible to learn information about its input from the output). Note that it is not feasible to retrieve any information about the patient or the data from the fields of D without access to the key K_{date}. Figure 1 below, which is a screenshot of D written to the Blockchain, shows an example of data written. Note the hashed values of d_1, d_2, and d_3 as well.

Figure 1: Demonstration of Healthcare data written to the blockchain

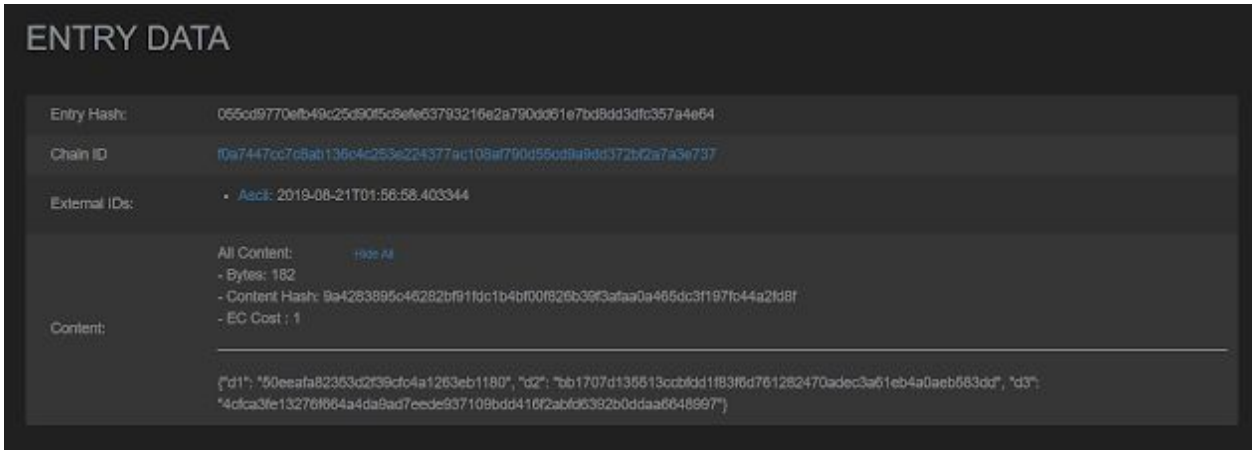

2.5.7 Retrieving Patient Data:

When a provider would like to query records associated with the patient, the provider needs to retrieve the key k associated with the records. If the records were generated by a different provider, the patient will provide the necessary key material during granting access to the records. Alternatively, the provider who generated the records can query them as well using the appropriate key. As before, there might be multiple keys k that a provider maintains for a patient

and would need to choose the correct key for the task (e.g., based on key descriptions *desc*). The provider would also need to have access to $K_{d1}^{(1)}$ and $K_{d2}^{(2)}$ associated with the patient, where the interval $[d_1; d_2]$ must cover the range of dates that the provider is retrieving.The provider will now query the (encrypted) records associated with ID of the patient stored under key hash *h(k)*. Upon retrieving the records, the provider decrypts them to obtain $C_1;C_2$; ... and verifies signatures in each C_i (querying public keys associated with the signing providers if necessary).

The provider also needs to retrieve blockchain data and verify that the blockchain recorded information about the same set of C_i that it retrieved from the servers (i.e., each C_i was recorded in the blockchain and the blockchain doesn't store an update for which no corresponding C_i can be found). For that purpose, the provider derives K_{date} for each blockchain update issued on date date (within the interval for which the provider is retrieving records) and determines whether the record matches a specific patient's update C_i issued on that date.

2.5.8 Resilience and Speed Testing:

Secure VPN interconnects between hardware and cloud components achieve low latency round trips of less than 10ms and can contain payloads over 10,000tps peak. Hardware components pass failover, disaster recovery, and penetration testing and are designed to maintain this functionality for at least two years of production operations. Further testing, with added nodes, to verify resistance against an inside threat or denial of service attack will be performed in Phase II and is described in greater detail in the research plan

2.6 Possible Uses

2.6.1 Chronic Disease Management

The initial intended application use case for HealthNet is chronic disease management. As many as 133 million people in the United States have at least one chronic disease. Chronic diseases like diabetes and heart disease are the leading causes of death and disability in the U.S and primary sources of its $3.3 trillion healthcare spend.[15] These patient's care networks could use HealthNet to securely share relevant information with all parties involved and more effectively coordinate care. Researchers could then be included to assess what factors are contributing to these disease states at the population level so providers can act accordingly.

Effective care coordination is critical to the success of a chronic disease patient's outcome. However, HealthNet does far more for the care coordination process than simply enable access to information. Seeing as the blockchain component of HealthNet serves as a single source of truth for updates to chronic disease patients' records, care coordination teams will have access to a transparent, agreed upon record of what tests have and have not been run on a patient. Hence,

[15] www.cdc.gov/chronicdisease/resources/infographic/chronic-diseases.htm

health systems will be able to prevent redundant testing, saving them substantial amounts of capital, lower their overall operating costs, and hasten the patient treatment cycle. There also exists a unique opportunity for population health researchers to collaborate across their region to look for regions of high density areas of chronic disease patients so that they may assess the environmental, socioeconomic, and geographic factors surrounding these 'hot spots' and make interventions. Doing so will turn chronic disease management from a reactionary treatment approach that costs providers millions, to a more preventative, interventionary model that will improve population health overall and save providers and payers substantially.

2.6.2 Prescription Drug Management

Another possible application for HealthNet is prescription drug tracking. Medication nonadherence is responsible for almost 125,000 deaths, 10% of all hospitalizations, and $300 billion in unnecessary expenditures.[16] Providers could work together with pharmacists and use HealthNet as an audit trail to track if patients are being adherent and coordinate interventions more effectively in the event of nonadherence. Additionally, the U.S is in the midst of an opioid epidemic that has claimed over 2 million lives and cost the U.S over one trillion dollars since 2001.[17] Similarly, providers and pharmacists could use HealthNet to track patients who are prescribed opioids and coordinate interventions if they believe the patient is abusing the drugs.

2.6.3 Insurance Claims Tracking

HealthNet could also be used by health insurers to track claims. While only a fraction of the billions of health insurance claims in the U.S are fraudulent, they are difficult to identify and cost the healthcare system between $68 - $200 billion every year.[18] Insurers could use HealthNet as an encrypted audit trail to track claims, verify their legitimacy, and take action against fraudulent sources much faster, saving billions of dollars.

3. Previous Approaches & Improvements

Healthcare often ranks highly among industries that are poised to be disrupted by blockchain. As such, numerous approaches have been made to tackle the problems associated with data management in healthcare, especially as it relates to medical records. Few, however, have seen a high degree of adoption from a commercial standpoint. Some projects that have tried to tackle these sorts of problems are: MedChain, MedRec, MedicalChain, & FHIRChain. While addressing a similar problem to these networks, HealthNet builds upon their approaches to provide a more secure, efficient, and practical solution for adoption in health systems.

---

[16] www.nacds.org/news/the-cost-of-medication-non-adherence/
[17] www.hhs.gov/opioids/about-the-epidemic/index.html
[18] www.economist.com/united-states/2014/05/31/the-272-billion-swindle

3.1.1 MedRec

MedRec was launched in August of 2016 through a joint effort by MIT Media Lab and Beth Israel Deaconess Medical Center in Massachusetts. In the original whitepaper, MedRec was proposed as a "Novel, decentralized record management system to handle EHRs, using blockchain". The goal of this project was to make it easier for patients to access their medical information across different healthcare providers and leverage blockchain to create an immutable log of where the most recent record for the patient is located. In their original design, MedRec utilized a proof-of-work, Bitcoin network that rewarded medical stakeholders to participate as miners. In exchange for mining, medical researchers were rewarded with access to the aggregated, anonymized data on the network. In the newest version of their offering, MedRec is transitioning its network to a proof-of-authority Ethereum based one, but their general focus remains on improving cross-hospital EHR sharing.

MedRec is currently the oldest player in the space and has expanded into several hospital system in Massachusetts. However, it is difficult to discern if the project is still active as little development effort has been made in recent months. MedRec was pivotal in starting to formalize a way of thinking about blockchain in healthcare, yet it had several shortcomings that HealthNet seeks to improve upon: security, not directly handling interoperability, and compute requirements.

While placing medical records in the hands of patient makes it easier to move records between hospital systems, it also opens up network vulnerabilities. Consequently, the lack of focus on maintaining secure endpoints in the network leaves an opportunity for hackers to undermine the network and compromise patient data. Additionally, MedRec does not try to address the interoperability problem in health systems directly, instead electing for health systems to figure out the record exchange. This presents an obstacle for health systems looking to adopt their network, as it does not address the core problem associated with interoperability: a universal data standard. Medrec's design choice to not handle data mapping between health systems still requires them to deploy extensive resources and capital to handle it accordingly. Additionally, the compute requirements for mining Bitcoin are reasonably intense, making HealthNet a fairly steep expense for already budget tight health systems. Even if a transition to Ethereum is completed, the compute requirements for this network are extensive as well and can have a high degree of variability depending on the bloat of the network. The lack of predictable pricing to mine on either of these networks in addition to the requirement for Health systems to hold cryptocurrencies presents another adoption barrier to already risk averse hospitals.

3.1.2 MedicalChain

MedicalChain is a small, Swiss based startup founded in 2016 that is also exploring ways blockchain can be used to improve patients' access to their health records, with the goal of shifting the current electronic medical record management system to a user centric one. MedicalChain provides two main services: a secure storage environment where patients can store and have sovereignty over their data and a marketplace where third parties can sell services to patients in exchange for access to their health information. Their network and ecosystem is powered through their associated cryptocurrency: MedTokens. Essentially, a patient can grant a healthcare professional access to their data and any changes would be tracked on MedicalChain's ledger. The marketplace they are creating then allows different parties to build applications that improve a user's healthcare experience on the network by leveraging their data. Therefore, patients can see the benefits of having power over their data and be compensated financially for its use. One example is a telemedicine application on their platform where patients can grant a doctor access to their medical record and then receive an online consultation directly through some type of smart device like a tablet in addition to MedTokens. MedicalChain also announced a working agreement with Mayo Clinic to explore potential uses and benefits from blockchain in healthcare in the future.[19] However, no subsequent updates have been published since the June 2018 announcement regarding this collaboration.

MedicalChain's platform has done much to shape the thinking of the future of the patient/provider relationship and how it can evolve into this sort of data economy. They also have done much to advance the field as it relates to the creation of a secure platform/network for health applications to be built on that leverages blockchain. However, similar to MedRec, MedicalChain does not address the underlying systemic interoperability problems within health systems. It relies on patient's requesting their records from a provider, storing it on their phone or other device, and then permissioning access to their information from there. MedicalChain made a critical jump in realizing the need for secure storage and offering it as a service to their customers, but this only accessible after completing the previously mentioned steps. Hence, security vulnerabilities are still accessible, and there is no way of ensuring that a provider's system can access the patient's record since no interoperability mechanisms are built in. Additionally, healthcare through MedicalChain's platform is contingent on the adoption and use of its native cryptocurrency. While such an approach is feasible on an individual level, it is highly uncertain at a systemic level within health systems.

---

19 https://medium.com/medicalchain/medicalchain-announces-joint-working-agreement-with-mayo-clinic-9cfb474dcf0f

3.1.2 MedChain

MedChain is an early stage startup from Colorado that is similar to MedicalChain. They brand themselves as a globally compliant healthcare blockchain with a secure dApp ecosystem for health record storage. Similar to MedicalChain, their ecosystem is powered by their own cryptocurrency, which is called MedChain Tokens. These too are an Ethereum derived ERC20 token just like the MedTokens used by MedicalChain. participants on the network mine MedChain tokens and can then use them for storage and validation. MedChain's goal is to improve the overall operability of conventional electronic medical record systems and promote better communication between healthcare professionals, patients, and anyone else who needs access to the records. They plan on storing health information across a network of HIPAA compliant servers to improve security and state that the records cannot be hacked unless a user on their network approves it.[20] MedChain works by encrypting and fragmenting electronic medical records across their network where they are stored securely in federated servers. Users on the network can then permission access to their data through smart contracts. MedChain's platform is accessible to users through an App on iOS and Android platforms.

MedChain took the notion of secure storage a step further than MedicalChain by utilizing secure, federated servers within health systems as a storage medium for health records, similar to MedRec. MedChain also does not have as heavy compute requirements as Bitcoin mining does, yet, being an ERC20 token, it is a derivative of the Ethereum network and can still experience some volatility in compute requirements. While using existing healthcare infrastructure to manage record storage, MedChain does require that patients and providers hold cryptocurrency. Similar to MedRec and MedicalChain, this does nothing to address the underlying interoperability problems between hospitals. Consequently, regardless of if records are stored in secure, HIPAA compliant servers housed within hospitals, there is little that can be done to send them to another health system unless significant resources and capital are deployed.

3.1.3 FHIRChain

FHIRChain is a research project out of Vanderbilt University that is investigating secure and scalable clinical data exchange using blockchain. FHIRChain utilizes a minimalistic blockchain architecture, where all data shared on the network is normalized to FHIR standards to make exchange between providers easier and more interoperable in addition to making the document highly readable but the only thing recorded to the blockchain is a hashpointer to the datasource where a record is stored. All data sources are housed within a provider's HIPAA compliant server(s) and are securely connected to the FHIRChain where they can grant access to the particular piece of information requested. The network is built using a private Ethereum network, and has a corresponding telemedicine dApp built on top of it through the use of an Ethereum

[20] https://medchain.global/doc/Medchain%20Whitepaper%20v1.1.pdf

smart contract. FHIRChain was designed to meet the following ONC technical requirements for secure and scalable sharing of clinical data: Verifying identity and authenticating all participants, storing and exchanging data securely, permission to access data sources context, consistent data formats, and maintaining modularity.

Relative to the other networks mentioned above, FHIRChain is the most holistic approach to addressing population health management and EMR interoperability. Blockchain is only a minor component of FHIRChain, simply serving as a system of pointers to where data is and then permissioning access to the data within the FHIR datasource. Though data normalization to FHIR, FHIRChain is able to promote interoperability and enhance readability so data can be securely shared and allow clinical decision makers coordinate care more effectively. However, there are a few limitations to FHIRChain in clinical practice that present difficulty in adoption. Having been built on top of Ethereum, transaction size is a practical concern and limitation. The larger the exchange or greater the volume of transactions on the network, the greater the transaction cost to request or write data on FHIRChain. This makes it limited at scale in terms of network size and ability to operate across multiple health systems. Additionally, in order to run the network, health systems are forced to set up Ethereum miners. Running these miners is computationally intensive at scale and requires health systems to hold cryptocurrency, which may be perceived negatively to risk averse hospitals.

3.2 Improvements Upon Previous Research

HealthNet builds upon the approaches and innovations of the aforementioned protocols and provides a new population health data exchange network that, among other features, provides a greater level of security, flexibility, and addresses all facets of the interoperability problem. HealthNet's architecture is most analogous to FHIRChain, where blockchain functions as only an encrypted, thin data layer that allows authenticated providers to decrypt information about the location of a patient's record and request the information. A FHIR formatted version is then presented to the provider so that they may make better informed clinical decisions when providing treatment. HealthNet also does not require health systems to hold volatile cryptocurrencies and provides predictable compute and cost requirements. This is because HealthNet is built on top of the Factom blockchain. Factom utilizes a two token model in which there is the token associated with the network, known as the factoid, and then a separate token used for writing data to the blockchain called an entry credit. The factoid's price varies with the cryptocurrency market, but entry credits are price stable, allowing for a fixed compute cost. Entry credits cost $0.001 per write to the Factom blockchain. Therefore, health systems on the HealthNet Network only have to hold entry credits which presents a lower amount of risk and volatility. Other improvements and advantages of HealthNet relative to blockchain enabled solutions and other general purpose interoperability solutions are outlined below.

One of the most distinct advantages of HealthNet compared to competing solutions is its unparalleled level of security. Other blockchain-enabled record sharing solutions require patients to store medical records on their phone to circumvent HIPAA regulations which exposes vulnerabilities to both the patient held record, as mobile phones are notoriously unsecure, and the healthcare networks to which they are connected. HealthNet also takes far greater security measures surrounding who can access a patient's record than the current industry standard: secure fax. When a record is faxed between hospital systems there is an assumed trust that the individual receiving the record is authorized to access it, however, the sending party has almost no way of proving this is the case. HealthNet uses novel key management technologies and encryption schemes to mitigate this threat vector by requiring a provider must authenticate herself with a private key stored on a secure hardware device. Only then can she sync to the network, request access to the portion of the record needed, and decrypt its contents. Thus, malicious actors never have an opportunity to manipulate the data.

HealthNet also allows providers to exchange complete data much faster. Customer interviews revealed that faxing records between two health systems can take up to 60 days due to the credentialing process both sides undergo to authenticate themselves. While other blockchain solutions allow patients to permission who may access their record, most leave it to the hospitals to figure out how to facilitate the transfer which often resolves back to faxing. Moreover, there is no way of proving that the record being sent is the most up to date information on a patient as opposed to only a fragment. HealthNet addresses all of these concerns by automating the credentialing and exchange processes through its key management system, direct integration with existing infrastructure, and tracking updates to patient records on the Factom blockchain using encrypted metadata. Then, when a provider authenticates herself, the blockchain will be traversed and all instances of updates to that patient's record will be decrypted so she will know which record is the most up to date, request access accordingly, and receive a FHIR formatted version of the record in minutes.

HealthNet also provides greater flexibility for health systems than competing solutions. Every health system is different and so are their interoperability needs, but providers at all of them have habitual work flows that create massive operational costs if interrupted. HealthNet's collection of APIs are designed to integrate into existing infrastructure within health systems so providers are not required to learn a new interface. All hospitals need to do to map their data to the FHIR standards or use Factom as an audit trail to track record updates is add a few extra lines of code into the backend of their EMR system.

4. Directions for Future Research

4.1 Expand Hardware Security Capabilities of the HealthNet System

The use of blockchain enables enhanced assurance of recorded data. However, endpoints and nodes remain possible attack vectors from entities intent on disrupting the network. A poorly secured node could result in system downtime or worse, compromisation of network identifier keys, which could allow the attacker to maliciously modify or steal data. In order to build enhanced security into nodes, we intend on exploring the integration of secure elements and other hardware security devices into node servers to evaluate their efficacy, similar to a PKI token, for healthcare provider authentication.

Key Features:

1. Increased digital security of devices. Secure element hardware contains more sophisticated key generation and storage methods than a CPU. A NIST certified random number generator guarantees the adequacy of private and authentication keys. Additionally, keys and other secure data is stored off RAM and CPU in the secure element, which is digitally inaccessible even if the node is compromised.
2. Increased fidelity of data. A secured identifier system utilizing keys generated on the secure element will ensure a transparent record of health data transactions, data updates, and procedures.
3. Workflow will consist of integrating code libraries to HealthNet software and integrating the secure element with server architecture.

4.2 Smart Permissioning, consent and revocation, emergency access to data

4.2.1: Consent Granting and Revoking.

The initial implementation of HealthNet is a basic form of permission granting based on the range of dates needed and with no constraints on how long data is available. Future work will expand to support the full range of functionality including the ability to grant access to specific types of records, and specifying constraints for the record use. The system will be able to facilitate enforcement of some constraints such as granting access for a limited period of time, where access to a record can be automatically revoked from the provider once the legitimate period for the record use has elapsed. Additionally, this functionality could allow granular sharing of data; for example, tracking a prescription refills could be recorded directly to the patient's health record.

4.2.2: Supporting Emergency Access to Patient Data.

Quick access to a patient's health records in the case of emergencies is an important goal of this project. Our intent is to use secret sharing to transfer each patient's key material across multiple servers or entities and grant emergency room (ER) facilities with the ability to retrieve a patient's records in extreme circumstances when the patient is unable to grant access to his or her records via the conventional mechanism. Design of this methodology will determine (i) what records are to be shared with ER in case of emergencies, (ii) how to efficiently enforce the temporary nature of this operation, and (iii) what mechanisms should be in place to prevent abuse of this capability (as a way to obtain unrestricted access to the records of a desired individual).

4.2.3: Customizing the System for Different Types of Health Providers.

The initial HealthNet Proof of Concept assumes a single type of health provider, while in a complete system it will be desirable to customize the system to support different types of health providers according to their role in the health ecosystem. To that extent, we will create slightly different interfaces and capabilities to medical offices, hospitals, and testing facilities, as they often interact with patients in different ways. For example, test results produced by testing facilities are typically sent to the medical office who ordered the tests and thus are automatically shared with the originating medical office or doctor. This has implications for access to this type of medical data, which in the current design is associated with the provider who prescribed the tests. However, because testing facilities are a certain type of health provider, they would need to be represented in the system design.

An additional goal is to incorporate pharmacies into the system design as it has the potential to improve information sharing and facilitate collaboration between pharmacies and providers who prescribe medications. For example, if medical offices are granted access to information about when patients fill prescriptions that they ordered, they may be in a position to make better or more informed decisions for treating their patients. Pharmacies, however, will need to be represented differently from other types of health providers in the system because of their distinct role.

4.2.4: Differentiating Data Accessibility to Patient.

This methodology would allow practitioners control over the access to certain data record types if needed.

4.3 Build Population Healthcare Data Toolkit

Intelligent sharing of healthcare data is an essential step to make the healthcare system smarter and improve the quality of service provided. While it is intuitive that merged healthcare data across multiple sources with quicker access will have multiple benefits, there has not been much research to quantify and evaluate the improvement in the quality of healthcare for patients and

the benefits for medical research and data analytical tools that eventually benefits patients and medical provider (e.g., more accurate diagnosis tools). Leveraging the in-depth development of cooperative R&D projects between HealthNet and the Statistics and Data Science expertise at the University of Notre Dame, we will collect data to answer the following research questions that have a multitude of practical implications. Specifically, we will address the following questions with appropriate analytical methods and tools.

1. Estimation and quantification of the shortening in time from making an access request for patient data to receiving the data or being granted access with the adoption of blockchain for data sharing, which will be achieved using survival data/time to event analysis. Similarly, the shortening in time for medical provider to make critical decisions about patients (diagnosis, treatment plan, or others) will also analyzed in a similar manner
2. Estimation and quantification of the increase in the proportion of making the right decision on diagnosis and treatment with more patient data. This will be quantified using categorical data analysis techniques.
3. Quantification in cost saving for medical providers after the adoption of blockchain through better insurance claim coordination with treatment rendered. This will be summarized for each medical provide, broken down by cost types. The cost data can also be analyzed using longitudinal data analysis to see the rate of decreasing over time and when it will reach plateau and equilibrium.
4. Artificial intelligence (AI) potentially can provide improved disease diagnosis with sufficient data inputs. Blockchain technology would allow people to share their medical data with researchers easily and securely and retain control over it as a way to protect and tune their tolerance level for sharing their private and sensitive medical information. To quantify the benefit, we will compare the prediction accuracy of trained models for predicting patient outcome or diagnosing diseases with vs. without shared patient data to show the potentials of the former in improving prediction accuracy.

4.4 Further Security Analysis

We will conduct a formal analysis of security guarantees of the system design under different types of attacks. Note that the current design targets high security guarantees where a malicious participant (be it a malicious insider at a medical provider or a blockchain participant) is unable to learn any information about patient data beyond what they can legitimately possess. Additionally, reliability of the system, including failover testing of hardware components is complete. The system, however, is designed to maintain security even in the face of physical break-ins into healthcare network servers which maintain (encrypted) patient records. To that extent, our goal is to formally evaluate information available to an adversary who breaks into one or more servers and obtains a snapshot of the system (one-time snapshot) and to an adversary who breaks into one or more servers and is able to continuously monitor the servers. Because

medical records are always stored in an encrypted form using key material that the servers do not possess, it is not possible to recover medical records themselves, but certain metadata might be available to such an adversary. Our design makes it difficult to link (encrypted record) updates to patients, but a formal treatment of the adversarial view is still desirable as this type of a system will be considered a desired target of attacks.

In addition, we will evaluate the system with respect to its susceptibility to denial of service by different types of malicious players (including patients, health providers, and external parties). We will devise protection mechanisms that limit the effectiveness of denial of service attacks by each type of participant including automatic responses and collecting general use data to analyze for misbehaving or compromised participants.

4.5 Evaluate and Integrate Smart Contract Functionality

Smart contract functionality has potential benefits for medical information sharing. In most cases, a 'stateless' function can be used for medical transactions. For example, a patient could share their record with a provider or researcher in a one time transaction for a medical appointment or study. By utilizing a 'stateful' function (smart contract), increased control over the sharing event could be achieved. A patient would be able to grant time limited access to their record or trigger payment to a provider upon completion of a procedure. Initially, we will evaluate a single system-wide contract that encodes the logic for all possible ways to grant or revoke access to one's medical records which all patients will use. Giving or withdrawing consent to share medical records with a provider will consist of specifying inputs to that smart contract and recording this fact. The inputs will include information about the parties involved (patient ID, provider ID, etc.) as well as information about the permission changes themselves (i.e., the terms in which permissions are granted, the affected dates, whether this is a grant or revoke operation, etc.).

The inputs will need to be kept private from the external parties (i.e., parties not affected by the permission changes), but we can record the fact that a permission changing smart contract has been triggered, i.e., private inputs to it have been provided, in a similar way to the way we record updates to medical records on the blockchain. If private smart contracts with the input hiding property are supported, this functionality would enable any party to verify contract execution on private inputs. In this scenario, proof of procedure could be shared with an insurance company as well, ensuring payment to the provider for the procedure while minimizing release of private patient data.[21]

[21] https://www.trustkernel.com/uploads/pubs/shadoweth.pdf