Design and Implementation Considerations for Blockchain for Health Records

Prashant Khambekar

Abstract

This chapter provides the motivation for using blockchain-based systems for health records as the currently prevalent electronic health records are inadequate for providing complete care and a smooth experience. The organization of such systems is discussed for developers intending to proceed on this path. Real examples from multiple such systems indicate that the challenges faced by developers in creating and deploying successful systems are not trivial. The issues are described here in depth and the corresponding solutions are discussed.

1 Introduction

Health records are being maintained in electronic systems for many years now. So, why is blockchain for health records desirable?

Consider a simple patient situation. Patient Paolo is sent by his family dentist to get one of his molars examined by an endodontist for a possible root canal. Paolo has to provide all the dental history to the endodontist and inform about the tolerance to local and general anesthesia as well as allergies to certain classes of medicines. The next year Paolo is referred to a periodontist for an unrelated gum problem and Paolo again has to provide the dental history, the anesthesia information, and the allergy information. This problem arises because every practitioner keeps their patient information in their own silo. Due to lapses in memory, and possibly not realizing the importance of all relevant history, Paolo only provides some of the information. Paolo thus faces the unnecessary risk of side effects.

P. Khambekar (✉)
Harbinger Systems, Philadelphia, PA, USA
e-mail: prashant.khambekar@harbingergroup.com

C. Lu, M. Tanniru (eds.), *Blockchain in Healthcare*, Future of Business and Finance,
https://doi.org/10.1007/978-3-031-45339-7_7

Consider a slightly more complicated situation. Patricia complains about discomfort in her feet. She undergoes surgery in a hospital and is provided with a cast. When she is discharged from the hospital, she has to continue the pain medication. At the same time, she needs to continue to be seen by her family physician for her chronic, low-impact problems. If the family physician does not get the complete picture of her medication and whether it has affected her overall health, there can be unnecessary complications. Soon, Patricia must start physiotherapy for the foot condition and the physiotherapist needs to see her X-ray images. So, Patricia has to remember to carry to the physiotherapist the image CD provided by the hospital to her. For the proper continuity of care, accurate health data records need to be shared across doctors.

Health records are related to one person, but they are locked into siloes. And for proper care by a collection of doctors, the complete sharing of information is desired. In some regions of the world, doctors do not trust the accuracy of the data residing with other physicians and health systems. A blockchain can provide the confidence that the data is accurate and is neither missing nor faked. The trust enabled by the blockchain allows physicians to share data (Peral et al., 2020; Azaria et al., 2016).

HL7 is a mechanism for sharing health record data from one party to another. DICOM is a standard for health image data which enables the sharing of health images. These form the basis of the transfer of health data across systems. However, they do not ensure continuity of care and complete care because not everyone shares the information. Doctors and hospitals maintain siloes of information and updates may not be shared leading to un-synchronized siloes. There is no guarantee that data is not faked. And even though health data is related to one person, there can be multiple identification keys—sometimes within a single system—leading to a mess. In contrast, the trust enabled by the blockchain encourages the sharing of health data.

A few companies have come forth with blockchain-based storage of health records. (Confidentiality agreements restrict disclosing the names of customer companies and fairness prevents the mention of other companies, but they can be found by searching on the Internet.) Patients can view their records. Doctors can view the records to which patients have given them permission to read. Patients can either add new records or permit the import of health records from EHRs and similar systems. Data can be shared with other doctors for continuity of care or expert opinions or second opinions. When such data is accumulated and available, medical research groups, pharmaceutical companies, and clinical research groups are interested in studying the data for understanding diseases and for developing new therapies. (Some use cases are given in Sect. 3 below.) This data can be either the raw data of individual records or it could be aggregated, anonymized data. Typically, patients need to permit part or the whole of their data to be provided to such entities. The users of such a system are patients, doctors, and research organizations along with a few administrators for overseeing the smooth sharing of data.

This chapter covers the organization of blockchain-based health record systems and the issues faced by developers and entrepreneurs. Section 2 covers the basic structure of the system. Section 3 starts off with simple use cases and how to map that to the blockchain system, and then moves to more complicated use cases. The rest of the sections address issues in order of simple to complex. Section 4 talks about how the systems are organized based on the peculiarities of the blockchain concept. Section 5 addresses which blockchain to use out of the available blockchains or whether one can create a blockchain from scratch. Section 6 talks about the overall user experience and thinking of the whole system as it lives and is used by users. Section 7 covers legal and governance issues.

2 Basic Organization of System

The basic organization of the system is simple. As in all software applications, there is the user interface with some front-end logic, there is a back-end business layer, and there is a database.

The main database of interest here is the blockchain. The patient health data is stored on the blockchain. A blockchain is not only a storage mechanism but there is code inherently associated with it for the proper handling of data and transactions. Whereas the blockchain acts as a record of the writes of health data, the logging of reads of the health data by various users and organizations is also important and is recorded on the blockchain. The blockchain is replicated across computers; each such computer is called a node. The blockchain is accessed by the server-side business layer and presented to the users via the user interface. There is a separate database for basic information such as users, their system passwords and access privileges.

In general, there is back-end interface for the input and output of data. There is an intake of data from external sources. In some cases, there is direct user entry of data via the user interface but in most cases, especially for historical data, the data is pulled in from external systems with permission. Bulk data desired by doctors or organizations such as research groups would be provided to them based on the permissions granted (Fig. 1).

3 Getting Started with Blockchain Development

3.1 Some Simple Use Cases

To understand how blockchain development is carried out, a set of simple scenarios can serve as a good start.

One blockchain was created for the referral of patients to specialists for the continuity of care. The patient health data was placed on the blockchain by the primary care doctor, and the specialist could then access the data and add to the health data as treatment progressed. A very similar blockchain was created for second opinion by expert doctors. The data would be segmented so that different experts could view the data of one patient and provide their opinion on the best treatment. Both these blockchain-based systems were created for customers in the USA.

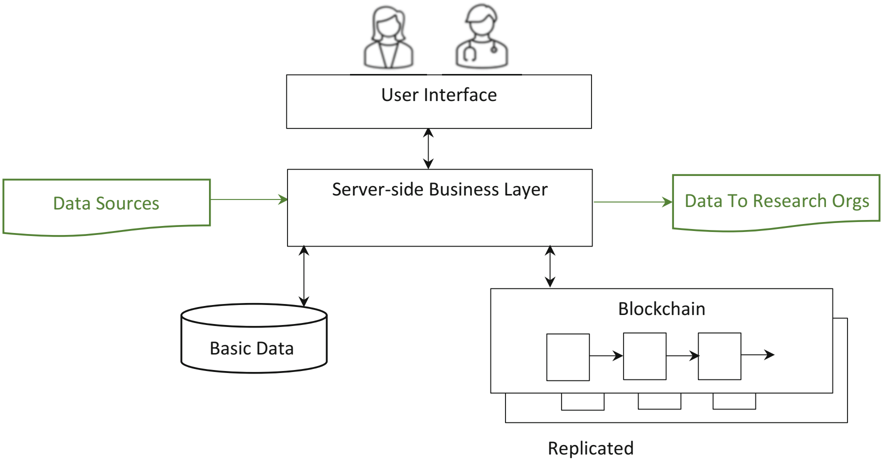

Fig. 1 Simple organization of Blockchain-based system

From these examples, the development process consisting of analysis of the problem, the solution design, and the implementation of the blockchain system can be studied.

3.1.1 Analysis: What Is to Be in Blockchain and What Is Outside the Blockchain

To proceed with the system, the main question to be asked is to what aspect is the blockchain applicable? In these two situations, the patient health history and current problems are of interest, so they are placed on the blockchain. The payment for services or insurance is not of critical importance, so that continues to be with the doctors' current systems. There is no need for data to be fed to research organizations, so they do not come into the picture.

3.1.2 Design: Define Roles and the Care Coordination Workflow

Now, the user roles and workflow are examined. The user roles are the Primary Doctor and the Expert Doctor. There can be sub-roles such as a doctor's assistant for doing some of the tasks. There is System Administration for the creation and exit of Primary Doctors and Experts. The main workflow is the addition of the health data—the historical and the current by the Primary Doctor, and the reading of that data and the addition to that data by the Expert (Ali, 2021). From an application viewpoint, the actions are View Health Data, Add Health Data, Assign Expert to Health Data (a set of related health records) plus some administrative actions such as View Primary Doctors, View Experts, Add Primary Doctors, View Usage Statistics, etc.

3.1.3 Implementation: Define the Detailed Parts—Blockchain and Business Logic

The next big activity is to map the software application world to some of the peculiarities of the blockchain world. This is a matter of terminology. The following are the main parts.

There is code for the blockchain to do its work. It is called smart contract, chaincode or similar. Essentially, it enables writes to the blockchain, reads from the blockchain, doing conditional coding (if-then) and looping. The basic unit of work is a transaction. So, the writing of one health data record is a transaction. As the recording of permissions is important, that is a transaction. The recording of who read what is important, so that is a transaction.

The identities that write and read are called as addresses. So, writes are done on behalf of a user, that is, one address. A read is done on behalf of the same user or another user, so those addresses come into play. Reads are done with the public part and writes are done with the private part of the public–private key pair that constitutes an address. The server-side back-end code needs to provide this to the blockchain as part of the transaction request.

All transactions can be read with appropriate permissions. In a general blockchain like Bitcoin, anyone can read the transactions between two entities, which is two addresses. That cannot work with private data such as health data. So, for reading a health data record, firstly the permission has to be read from the blockchain. It is possible that permission was given and then revoked. The latest permission setting is the one that comes into play.

The transactions and blockchain reads are submitted to the nodes of the blockchain. Depending on how nodes are laid out, the user entities may or may not have their own node. (See section "Copies of the Blockchain" below.) If a user entity has their own node, the server-side back-end code can submit the request to it. Otherwise, the server-side back-end logic needs to determine to which node a request should be submitted. This can be determined based on geographical nearness, the current load, or similar such considerations.

Some blockchain implementations utilize a fee for carrying out a transaction. In broad applications such as health records, such micro-fees for each transaction probably confuse the big picture of providing proper healthcare. But as the blockchain demands the fee, the server-side back-end code needs to provide it. This needs to be resolved at a higher level within the system—either by periodically topping up the fees for each user, or, accumulating and settling them over a period such as a month.

It can be seen from the above that there is blockchain code which does the job of storing, maintaining, and retrieving the blockchain data properly. There is server-side code that feeds the blockchain as needed and accepts the data that the blockchain provides. The server-side code connects to the user interface and to the external data sources and data sinks.

Separate from the blockchain data there is a need for a small database to keep track of the users who are the Primary Doctors and the Experts. The retrieval and update of current and historical users is done using this database. If users need passwords in addition to public–private key pairs, then passwords are stored in this database.

For the segmentation of data for second opinion by Experts, the server-side business layer is slightly different. It needs to handle the splitting of data and possible determination of the appropriateness of the expert to whom the data is permitted (for example, renal data to a renal specialist verses an orthodontist). The blockchain code is unaware of such business layer considerations.

3.1.4 Summary of Development

The process for handling of the health records between the Primary Doctors and the Experts is, thus, the identification of what will be placed on the blockchain, the user roles, the workflows between the user roles, and the separation of what the server-side code does with the blockchain and what it does with the basic database. The blockchain code for managing the health application needs to be written to ensure the privacy of the health records, which is different from that of other blockchain applications. The server-side code interaction with the blockchain needs to be in accordance with the terminology and requirements of the blockchain (addresses, public key, private key, which node, etc.).

3.2 More Complicated Use Cases

One blockchain application was created for the utilization of patients' health data for clinical trials. Clinical Research Organizations reach out to patients for their data and offer payment for the data. With the data being on the blockchain, patients can provide access to the organizations to the parts of the historical data or current data as per mutual agreement. For the ease of use by patients, this included a mobile phone front-end for popular phone operating systems. This was for a US-based customer.

In addition, two different blockchain applications were created for the aggregation of health data for research into diseases. The research organizations pay the patients for access to their health data in a manner like clinical trials as in the above paragraph. One of the applications was for a US-based company whereas the other was for a Europe-based organization.

3.2.1 The Development Process and Sub-Parts

The development process is the same as indicated in the previous section.

What part needs to be on the blockchain? The health data, for it to be trusted by the entities involved. The payment information, for the purposes of completeness, traceability, and trust. Going "upstream" from the payments because payments depend on the permissions given to the various organizations for the access to the pieces of the health data, the permissions should be on the blockchain.

The number of entities is much larger. To keep it simple, the user roles are Patient, Doctor, and Buyer, where Buyer is the organization requesting the data for either clinical trials or disease research.

3.2.2 The Workflows

The workflows are sightly complicated. In addition to the write of the health record, the read permissions and the payment need to be handled.

The following gives a glimpse into the permission workflow. The permission can be provided by the patient to one Buyer on a per-record basis, or it can be to one Buyer on a blanket basis, or it can be to a slate of Buyers on a blanket basis. If permission is not already provided, then a Buyer may request permission from a Patient; again, this can be on a per-record basis or a blanket basis. The time aspect needs to be considered; requests for permission could be ignored, reminded in a timely manner, or, could time out. All permissions need to reside on the blockchain (Waghmare, 2020).

For payments, the following is a glimpse into the workflow. When a request is made to read a health record, the permission is checked. If permission is provided, the health record is read from the blockchain, the read event is recorded, the payment is computed and recorded on the blockchain, and notifications are sent to the Patient and the Buyer.

The blockchain code is written with all the above requirements. That development is similar to the one already described in the previous section. The complexity introduces issues to be considered; these are given in the sections below.

3.2.3 Business Logic and User Interface

As is to be expected, the server-side code is significantly more complicated. It includes all the aspects of client–server system creation. This covers the handling of incoming bulk data and outgoing bulk data too. One aspect that could be unique is interfacing with crypto exchanges if cryptocurrency is utilized (but also see section “The cryptocurrency aspect” below). The blockchain connection likely throw up issues; these are given in the sections below.

The user interface needs to sensitize the users about possible transaction times (see section “Speed of the Blockchain” below). The user profile also needs to be considered. Whereas mobile phone applications are common now, if the user interface is mobile phone, then users may need some training and guidance on how and why things are different when a blockchain is involved.

3.3 Stepping into Detailed Considerations

The different aspects to be considered for actual implementation of such a system and taking it into sustained and successful production are given below. The aspects are organized from a low-level to a high-level.

Getting all the health record data is a complicated topic in itself. As it is not exclusively related to the blockchain technology, it is not covered here.

4 Organization of the Blockchain

4.1 What Is on the Blockchain and What Is Kept Separate?

Whereas some health records are small in size, some files such as radiology images can be huge. They are typically 5–16 MB but can be as big as 50 MB (Ohal, 2021; Seibert, n.d.). Blockchain blocks are typically limited in size; the block size is based on considerations such as the rate of arriving transactions, and the computation and coordination needed to create the blocks. Typical block sizes are 1 to 8 KB. Because of this, voluminous files need to be stored separately. They are stored in an encrypted manner in global, replicated databases such as IPFS or Cosmos DB. The hash of the data in the file, the date of creation, the file identifier, and other key characteristics are stored in the blockchain (Fig. 2).

4.2 Small Health Records Can Be Large for the Blockchain

Some health records may simply be larger than one block or a few blocks. A single, non-image medical record can be 4 to 60 KB in size. Typical block sizes are 1 to 8 KB. One technique to manage this is to increase the block size if this can be accommodated along with the rate of arriving transactions, and the coordination needed to create blocks.

That may not suffice for most health records. Health records need to be split so that they span multiple blocks. Typical health records have multiple sections such as identifying demographics, medications, problems, etc. One way to split would be by sections, as given in the figure below (Fig. 3).

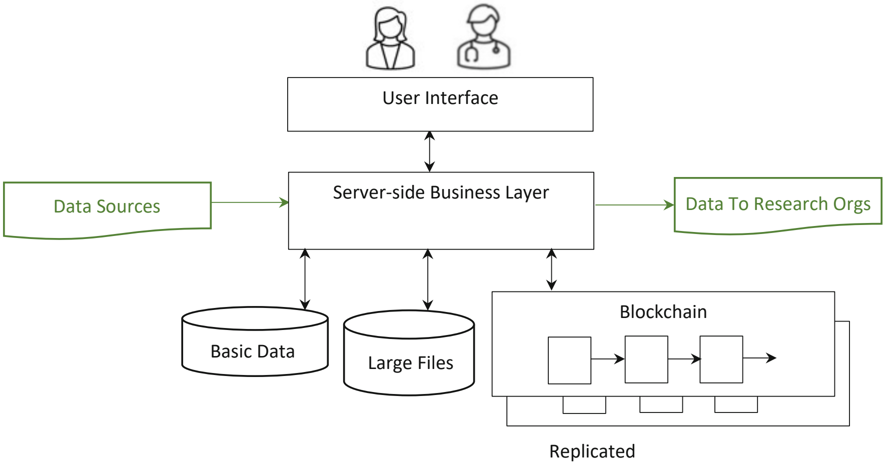

Fig. 2 Blockchain system with large files on global database

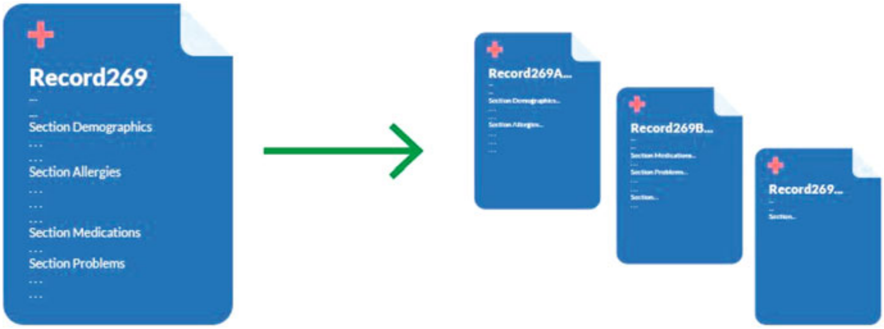

Fig. 3 Splitting a health record

Another way of splitting would be to simply chop the record into pieces that each are just under the block size and add some identifying information to each part (Sharma, 2017).

The parts need to be re-assembled into the complete health records during a read. The server-side back-end logic needs to do the splitting and re-assembly. Also, see “redacting” in section “Compliance as per country and state” below.

4.3 How Many Copies of the Blockchain?

The thinking among the entrepreneurs and visionaries often is that as they add partner organizations such as research groups or non-profit organizations to be part of the ecosystem, each of those can maintain a copy of the blockchain. Given the prevalence of cloud systems in today’s world there is no need for such one-to-one equivalence between partner organizations and the copies of the blockchain.

However, when the system is spun up, there may not be any partner organizations or there may be just one or two. For the authenticity of the blockchain, the blockchain should have at least five copies.

Moreover, if the number of users and the number of records are small in the initial ramp up phase, then for the sake of the trust and confidence in the blockchain, the number of copies should be larger. Nine or more copies would generate sufficient confidence in the user and partner community.

These copies are best distributed as widely as possible geographically. Whereas countries have restrictions that data should not cross the boundaries (see section “Compliance as per country and state” below), the distribution should be as spaced out as possible within a country.

4.4 How About Blockchain on Mobile Phone?

Given that people use their mobile phones a lot there are attempts to put copies of a blockchain on phones. This may work for some types of records and data. However, a blockchain does not simply have one individual's data. Each block of a blockchain has data from several users. All that data is tightly bound with the hash of that block. Given that health records will be for many users, that each user's health data spans many decades and also given that some records will be large, keeping health data blockchain copies on a phone is impractical.

To get a feel for the volume, a rough computation can be carried out. A single, non-image medical record can be 4 to 60 KB in size. A patient generates close to 80 MB of health data per year including imaging data. Keeping the image data off the blockchain with only the key information on the blockchain perhaps reduces each patient's yearly data on the blockchain to about 0.5 MB. Whereas a patient's lifetime is fairly long, let us assume that medical records are for a period of 20 years. So, that makes a single patient's data to be 10 MB. For 100,000 patients, which is a fairly small subset of any population (by location, by community, by disease, etc.), this amounts to 1 TB of data.

Thus, for other than a few, niche, health record blockchain applications, keeping blockchain copies on a mobile phone is infeasible.

4.5 Proper Handling of Bulk Data

The ability of the blockchain to handle bulk data input needs to be tested beforehand, and if needed, strategies for handling bulk input need to be planned.

The situation may arise when users upload their historical health data records. If the records are split and submitted to different nodes in a concurrent manner, then the different nodes may decide to prioritize separate sets of records, that is, blockchain transactions. Normally, the consensus mechanism will ensure fairness and handle the transactions in a timely manner, that is, process first-in-first-out in an overall sense (though not in an exact sense for every transaction request). However, if the load is too even, the blockchain may create ommer blocks (uncle blocks) and then spend significant time trying to resolve those, leading to incoming transactions piling up.

Another situation is when sustained bulk load is applied as input, then the blockchain performance falls off and can take a few minutes to recover.

If this situation is possible and is identified during performance testing, then strategies of load balancing, introducing gaps in the load or higher-level prioritization may be needed in the business layer to ensure that all health records are handled in a timely manner.

Another possible solution that may work in some situations is to use elevated CPU or GPU processing power. This can be a temporary solution for a few days if a backlog of health records is to be processed. This has significant cost implication and is not a long-term solution (Kulkarni, 2019a). This will work only if there is control over all the nodes of the blockchain but not if the nodes are distributed across parties (see section "How many copies of the blockchain").

5 Selection of the Blockchain Technology

The basic concepts of the underlying blockchain technology are fairly clear—gather transactions, have the nodes create consensus regarding which transactions are to be placed into a block, create the block including its cryptographic hash, share this block with the other nodes so that every node then connects the block to the previous block and thus maintains the chain. Should you then create your own blockchain technology or select one of the available ones?

5.1 Speed of the Blockchain and Record Retrieval

The time that a blockchain takes to generate one block can be an important consideration. Because multiple transactions are assembled and put together, and different nodes of the blockchain could assemble different sets of transactions, they need to come to consensus about which transactions to assemble. This implies that any single transaction takes a much longer time to be completed/committed compared to the currently prevalent centralized transaction processing systems. Typical blockchains take 1 second to 20 seconds to assemble a block. Some transactions will "miss" the current block and will have to be taken up for a subsequent block. This needs to be factored when uploading bulk data as well as when an individual user is submitting data (Patel, 2020). And also, when a user does meta-level actions such as giving permission to an organization for accessing the data—because such transactions are also on the blockchain.

Retrieval of data needs to go backwards through the blockchain assembling records as they are found. Older records take longer to be retrieved (Joshi, 2020).

As the blockchain gets longer, its response time slows down. This is for writes as well as reads.

All these aspects must be considered as part of the system characteristics as well as from a user experience viewpoint.

5.2 Widely Used Blockchain Technology or a Professional-Seeming One?

The intrinsic blockchain technology needs to be very carefully engineered. It is easy to have holes and gaps in the technology which may manifest—suddenly and severely—in security and performance. Malicious actors try to exploit defects and new aspects are discovered every year. As a result, creating your own blockchain technology is a highly challenging activity. You must rely on existing blockchain technology created by other organizations. Ethereum (ethereum.org) and

Hyperledger (Hyperledger.org) are among the well regarded multipurpose blockchain technologies.

The technology needs to be open source so that multiple people around the world have already examined the code, pointed out flaws, and the code has been updated to repair those flaws. Some blockchain technologies keep most of the code open while withholding some code. Trusting such technology could turn out to be highly problematic and is best avoided.

Both Ethereum and Hyperledger have been used for the example customer systems described in the above sections. Ethereum has the ability to plug in different consensus algorithms. Hyperledger has two strong implementations in Fabric and Sawtooth. Sawtooth is more flexible than Fabric (Anwar, 2021). Sawtooth supports Byzantine Fault Tolerance, four consensus algorithms including Proof of Elapsed Time (POET), and supports the coding languages Rust, JavaScript, Go, and Python.

5.3 The Cryptocurrency Aspect

Blockchains can use tokens for transactions, that is, accept payments in the form of tokens for record updates and for record retrievals. It is tempting to float a new cryptocurrency which will gain value as the blockchain adoption increases over the years. However, any currency must deal with multiple aspects of economics such as speculation, inflation, convertibility with other currencies, attacks on the currency, and so on. If the main aim of the blockchain is to handle the health records for the benefit of patients and the medical community, then it is best not to embark on the cryptocurrency path and simply use ordinary, fiat currency.

6 Thinking of the Complete System

6.1 Secure Access to the System

Whereas the data on the blockchain is secured by the nature of the blockchain, all access points to the system should be secure too. This includes the databases, the portal, and all interfaces.

Malicious agents will try to exploit the system for financial gain as health data is quite valuable. Rather than grabbing the data from the blockchain directly, they would simulate the actions of legitimate users. They would attack the basic database of logins or the user interface and exploit vulnerabilities there. By acting like legitimate users, they may siphon off the data.

Other reasons for malicious attacks include revenge, sabotage, vandalism, corporate espionage, or quite simply the challenge involved. A very basic attack could be a Denial-of-Service attack by which unrelated requests flood the system leaving genuine users unable to access services (Cloudflare, n.d.; Fortinet, n.d.). To prevent all such attacks, not only should the application security should be verified by penetration testing but also the system should be secured with complete website

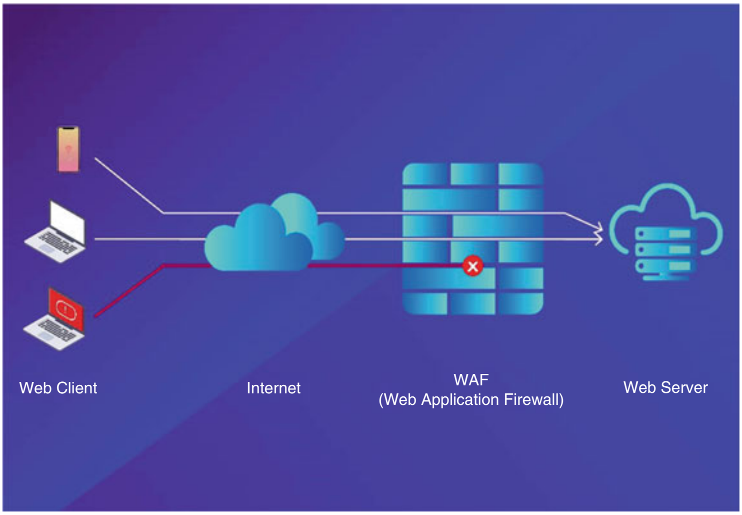

Fig. 4 Website security software

security software (Shekhawat, n.d.). This is shown in the figure below, where the security software handles legitimate access and bars suspicious access (Fig. 4).

6.2 The User Experience Matters

The main reasons for a blockchain-based system of health records are the trust and sharing of records. Whereas that is a necessary basis for the system, the whole system should be designed properly for the users and partner groups. Performance and reliability are important. The user interface should be pleasing and easy to use. There should be notifications and warnings as appropriate. There should be proper online help as well as error messages. The handling of bulk data imports and exports should be proper with adequate messages and with retries as needed.

For all the customer systems mentioned in the above sections, the user interface was tried to be kept as close to the current world-class application systems as possible. This was irrespective of a web front-end (Kulkarni, 2019b) or a mobile phone front-end (Sharma, 2018).

If users get turned off by any part of the system or its performance, they are likely to stop using the system. After that recovering from such losses and getting other users on to the system is an uphill battle.

The user experience matters a lot.

6.3 User Sensitization and Training

Some aspects related to the intrinsic blockchain need to be introduced, explained, and reinforced to the users.

Because a blockchain transaction can take time, the users need to have asynchronous submission of records and requests. The users submit first and can then do other, unrelated actions within the system. When the health record is committed by the blockchain in its entirety (possibly over multiple blocks, as in section "Small health records can be large for the blockchain" above) or a health record is read, then a notification can alert users so that they can choose to take the next step. This delay is not intuitive and users may not understand the underlying cause, so it needs to be explained when they start off (Waghmare, 2018).

Users need to be told clearly and often that their private–public key pair is unique to them, and it must not be lost. Whereas there are some ways to seed keys, most systems do not do this at the outset.

If cryptocurrency is involved, then users may need to be told about the risks in cryptocurrency (see section "The cryptocurrency aspect" above).

In general, users are not blockchain enthusiasts; they are doctors, patients, researchers, and so on. Hence, detailed aspects such as the tracking of transaction queues, the completion of transactions, the transaction fees, a node going down and so on by viewing the details of the blockchain activity need to be handled carefully. Either this should be kept hidden from all users (by allowing actions such as View Health Data, Add Health Data, Assign Expert to Health Data as in section "Getting Started with Blockchain Development" above) or only self-professed, advanced users may be allowed to view these aspects.

6.4 Responding Quickly to User Issues

When a system is in production then issues like down time and software quality come to the fore.

Keeping the system working properly as per Service Level Agreements promised to the users is extremely important. Down time for upgrades must be carefully coordinated across nodes and must be kept minimal. Blockchains are built such that one or more nodes may go down; when nodes come up, they copy the blocks already created and catch up. However, excessive down time for any node can become problematic.

All systems have defects. Even if testing is fairly robust, users may use the system in ways that are unexpected and unplanned. When issues crop up, they need to be handled quickly. A system may have good intentions and may draw in users in a fantastic manner, but users should not be lost due to bad service or unsatisfactory user experience.

7 Compliance and Governance

7.1 Compliance as per Country and State

Compliance with the laws of the countries in which a system operates seems to be too obvious to even be stated.

If your system spans multiple countries, you may want a separate instance of the blockchain for the users and organizations in each distinct country. That would need to be coupled with country-specific configuration of the system, possibly along with the business layer using country-specific rule engines.

Whereas raw data may not be transferable across country borders, aggregate data and insights can possibly be shared and sold. Thus, there would need to be a distinction between the type of data access that is provided, and the users need to be properly told and continuously guided on the distinction.

Even within a country, different states may have different laws. For example, the USA has HIPAA (Health Insurance Portability and Accountability Act) as the main law regulating health data. However, the states of California, Colorado, Connecticut, Utah, and Virginia have comprehensive laws related to consumer data privacy (NCSL, 2022). Some states may emphasize SOC (System and Organization Controls), especially SOC2, related to trusted services.

GDPR (General Data Protection Regulation) and related laws indicate that a user may choose to have some of their data deleted. Whereas it is not possible to simply delete old data out of the blockchain (as that would affect the hash of that block and all the subsequent blocks of data), a "redact" mask would have to be applied to that data. Thus, if a patient had a health record in 2020 that was requested to be "deleted" in 2022, then every read access by anyone after that request would need to retrieve the redact information first and mask out that specific health record of 2020.

7.2 Audit Access to the Blockchain

For building trust in the blockchain, audit access must be built into the blockchain. The auditors would be agencies other than the user community and the partner community. They could be governmental agencies or professional third-party audit companies.

The actual data that is added and shared must not be shown to the auditors. The data seen by the auditors would be the high-level data stored on the blockchain such as the date-time of data reads, and, whether the data reads correspond to the permissions granted to the reading user/organization. The "if and only if" constraint on the data sharing should be viewable by the auditors (Fig. 5).

The figure shows the blockchain on the left. The blockchain has health records as well as permissions given to other users. An audit, as shown on the right, examines some or all the permissions.

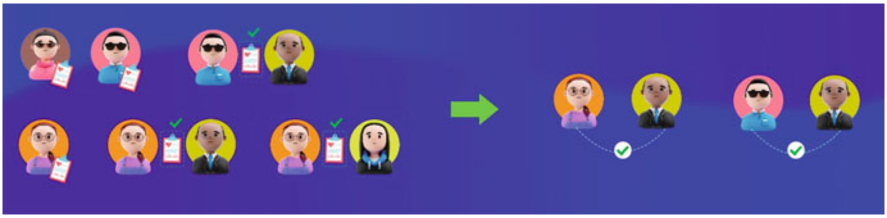

Fig. 5 Auditability of the Blockchain

This auditability is a fundamental characteristic of a system for something as crucial and regulated as health records. It needs to be built-in at the start. It cannot be retrofitted on the system.

7.3 Governance

For further building trust in the system a governance structure is important. Whereas the governance structure may not be present when the system is originally conceived and created, it is important for guiding the changes that need to be made to accommodate evolving types of users or the evolving needs of different stakeholders.

The governance could be through a body of directors, with the inclusion of some independent directors. This is the business norm for for-profit companies as well as non-profit organizations. This can be utilized for the blockchain-based system too.

In accordance with this, one of the customers has a governance structure wherein independent non-profit patient associations are a part of the body of directors, along with people from reputed, international pharmaceutical companies.

For blockchain governance, instead of a body of directors, there could possibly be a decentralized autonomous organization, or DAO (Weston & Beginner’s Guide To Decentralized Autonomous Organization Or DAO, 2021). The basic concept is that the voting power is decentralized among all or a large number of stakeholders. At this time, there are different types of DAOs and the goodness of the DAO concept is not clearly established across businesses.

8 Future Directions

As blockchain-based systems for patient records get created and released for usage by patients, doctors, researchers, and the public further issues of efficiency and usability will get discovered. They will require immediate engineering solutions in order to satisfy the users. They will also force the development of engineering solutions that will prove robust for the next few years. Concepts in blockchain such as consensus algorithms and efficient storage will continue to evolve. This indicates that those involved in designing and implementing blockchain-based system need to keep abreast of the latest developments.

9 Conclusions

The main motivation for creating a blockchain-based system is to utilize the advantages of blockchain technology. That is straightforward. For a successful implementation of the system, the aspects mentioned here need to be kept in mind, addressed, and planned for. Whereas the focus here is on how the characteristics of health records impact the usage and the success of the system, some of the aspects have general applicability and would need to be considered for blockchain-based system that aim to provide solutions in other domains and verticals.

References

Peral, J., Gallego, E., Gil, D., Tanniru, M., & Khambekar, P. (2020). Using visualization to build transparency in a healthcare blockchain application. *Sustainability, 12*(17), 6768. https://www.mdpi.com/2071-1050/12/17/6768

Azaria, A., Ekblaw, A., Vieira, T., & Lippman, A. (2016). MedRec: Using blockchain for medical data access and permission management. In *Proceedings–2016 2nd international conference on open and big data, OBD*. https://doi.org/10.1109/OBD.2016.11

Ali, T. (2021). *Harbinger systems private correspondence on organization and data interface*.

Waghmare, S. (2020). *Harbinger systems private correspondence on permission workflow*.

Ohal, P. (2021). *Harbinger systems private correspondence on sizes of radiology images*.

Seibert, J. A. (n.d.). *Archiving: Fundamentals of storage technology, chapter 2: Medical image data characteristics*. SIIM, University of California, Davis. https://siim.org/page/archiving_chapter2

Sharma, M. (2017). *Harbinger systems private correspondence on data record splitting*.

Kulkarni, R. (2019a). *Harbinger systems private correspondence on CPU-GPU cost*.

Patel, N. (2020). *Harbinger systems private correspondence on performance*.

Joshi, S. (2020). *Harbinger systems private correspondence on performance*.

Anwar, H. (2021). *Hyperledger Sawtooth Vs. Fabric: How are they different?*. https://101blockchains.com/hyperledger-sawtooth-vs-fabric/

Cloudflare. (n.d.). *What is a DDoS attack?* https://www.cloudflare.com/learning/ddos/what-is-a-ddos-attack/

Fortinet. (n.d.). *What is the difference between DoS attacks and DDoS Attacks?* https://www.fortinet.com/resources/cyberglossary/dos-vs-ddos

Shekhawat, V. S.. (n.d.). *What is a DDoS attack and how to mitigate it*. https://www.loginradius.com/blog/engineering/how-to-mitigate-ddos-attack/

Kulkarni, R. (2019b). *Harbinger systems private correspondence on user experience*.

Sharma, M. (2018). *Harbinger systems private correspondence on user experience*.

Waghmare, S. (2018). *Harbinger systems private correspondence on user sensitization and training about the blockchain experience*.

NCSL. (2022). *State Laws related to digital privacy (originally, state Laws related to internet privacy)*. https://www.ncsl.org/technology-and-communication/state-laws-related-to-digital-privacy

Weston, G. (2021). *Beginner's Guide To Decentralized Autonomous Organization Or DAO*. https://101blockchains.com/decentralized-autonomous-organization-dao/

Prashant Khambekar has worked extensively in commercial software, after obtaining a Computer Science doctorate in distributed computing. The work encompasses engineering, process, and business and touches wide-ranging domains such as supply chain, healthcare, and human resources. Prashant Khambekar has spearheaded innovation in Internet of Things, Data Science, and blockchain. The work in healthcare software spans several years and includes claims processing, Internet of Things for healthcare, patientcare software, pharmacy software, clinical trials, telehealth, etc. The interest in applying blockchain in healthcare goes beyond patient records to billing, claims and systems for nurses.