Access Controls and Health Care Records

Who Owns the Data?

case study

**A DISCUSSION
WITH
DAVID EVANS,
RICHARD
MCDONALD,
AND TERRY
COATTA**

Y*ou needn't be an expert with years of health care data—management experience to conclude the field is a hot mess. One visit to a hospital, clinic, or pharmacy can convince you of that. Burdened by legacy and fragmented into silos so alien from one another that they can scarcely communicate, health care recordkeeping has for decades frustrated any and all efforts to unify it.*

The underlying reason couldn't be more obvious: Each clinic, hospital, practice, and pharmacy operates its own isolated record-management system. The platforms and techniques vary from organization to organization, with almost no provisions having been made to share any of the information.

But what if these records were handled in more of a patient-centric manner, using systems and networks that allow data to be readily shared by all the physicians, clinics, hospitals, and pharmacies a person might choose to share them with or have occasion to visit? And, more radically, what if it was the patients—rather than the providers—who were considered to actually own the data?

It was with thoughts like these that a Toronto-based startup called HealthChain set out to create a platform for managing a patient's medication profile on the basis

of relationships established between patients and their various providers.

That vision became the challenge that HealthChain CTO David Evans took on, drawing on 25 years of work in the financial industry on portfolio management and quantitative research systems. In the years that led up to his transition to health care, he found himself increasingly intrigued with the possibilities of applying emerging digital identities and blockchain technologies to the creation of more efficient government services. Now he has an opportunity to put some of those ideas to the test.

To provide some insight into how HealthChain is addressing the medication profile–management challenge, Evans is joined in discussion here by Richard McDonald, a recently retired IBM Distinguished Engineer, and Terry Coatta, the CTO of Marine Learning Systems, a Vancouver-based startup working to develop a learning platform.

RICHARD McDONALD As people who've had occasion to visit doctors' offices—or even hospitals—from time to time, we all know just how important recordkeeping is to those operations. Historically, that has taken the form of paper records kept in overstuffed filing cabinets. But increasingly, it now seems also to include electronic records as the medical profession is making a belated push to fully enter the 21st century. As that process continues to move forward, questions come up as to who actually owns those records, who looks after them, and who needs to have access to them. What do you see as some of the key issues with respect to the custodial responsibilities these various medical organizations now need to address?

DAVID EVANS As you say, at this point there are doctors who use computer systems and others who don't. The ones who keep computer-based records generally use some localized instrumentation called an EMR (Electronic Medical Record) system [also known as an EHR (Electronic Health Record) system, depending on country and usage]. In terms of who owns the records kept on these things, any doctor who enters details about a patient into an EMR system is accountable for the accuracy of that information. With that said, when it comes to *ownership* of the full medical record for that patient, there's growing sentiment that this information properly belongs to the patient.

The truth on the ground right now is that any patient's records are actually scattered among any number of siloed databases since, as you go from one clinic to another or from one pharmacy to another, each is going to create a fresh record for you. When it comes to determining who owns those records, I'm inclined to distinguish the information that each practitioner is responsible for maintaining on his or her own system from the sum total of all the medical information collected about some particular patient over the course of that patient's lifetime. While many believe this is information that belongs to the patient, the current reality is that the patient doesn't actually even enjoy *access* to that data and has little to no control over how it's used.

TERRY COATTA Just to be clear about those silos, are you talking about centralized databases or just a lot of small individual databases maintained in different doctors' offices?

DE It's essentially a mix of the two. There are lots of independent doctors' offices out there, and to the degree

they're using computer systems at all, it may amount to little more than a database running on a desktop or laptop.

By the same token, there are health care teams that span multiple locations. One of the largest in Ontario has 37 different clinic locations that all share a single EMR implementation hosted in the cloud. But I believe even that implementation is organized such that doctors are able to see information for only those patients who visit their particular clinic location.

RM What does this look like from the perspective of an admin working at one of these clinics?

DE Even where working EMR applications are in place, there's a lot of aging technology to contend with. Also, early on, we were cautioned not to walk into a doctor's office just expecting to be allowed to change the workflow. There's a reason these technologies haven't necessarily advanced all that much since the '80s. The doctors are used to them. They reflexively know they need to hit the tab three times on this one particular page, owing to some archaic design choice. And yes, it may be stupid, but it's a design they've long since grown accustomed to.

TC Is there an integration story among these different sites? Obviously, if you've got all these data silos floating around in the world, you'd think there would be some standards related to information interchange. Or is it basically just the Wild West out there?

DE The most promising global standard for this is FHIR, which stands for fast health care interoperability resources. It's a standard that has gone through a number of iterations over the years, and it's focused on interoperability.

The only kind of consent that actually seems to exist is implied consent.

It isn't without its flaws, but FHIR is definitely the best resource for interoperability between different systems right now. Still, there's room for lots of data duplication, and it's essentially intended only for one-time transfers of data. That is, there isn't an FHIR network in use now where all these different databases are kept in sync in any way. Or at least there sure isn't anything along those lines that I've ever seen.

RM Am I right to assume that patients have very little say over what happens with their information?

DE That's absolutely correct. If you look at health-privacy standards, they all emphasize patient consent. But for all practical purposes, the only kind of consent that actually seems to exist is *implied* consent, since, as things currently stand, there's no practical mechanism patients can use either to provide or withdraw consent as to how their data is to be used. Which is to say the patient is almost completely out of the picture.

The other side of this is that information also isn't shared among providers in any sort of way patients might reasonably expect. If you always just go to the same clinic, they already know you and have your records readily at hand. But if, for some reason, you find you need to go to some other clinic or end up in an emergency room, chances are you're a blank slate for anyone who treats you there. They're not going to know what allergies you have. They're not going to know what prior conditions you have. They're not going to know what medications you're on. So, that means they're going to have to scramble around to scrounge up all the information they can from either you or a family member.

TC That sounds like a complete mess, so what part of that problem are you now trying to address?

DE To borrow a term from the blockchain world, we're working to deliver a *shared ledger* to the medical community. Our goal is to provide a view of a patient's records that not only doctors and pharmacists are able to share, but that can also be available to the patient. This is something that's actually possible today.

One of our key objectives is to keep track of all this information from the patient's perspective: What pharmacies do they use? What clinics do they visit? Which doctors treat them at those clinics? And how is it that each of these participants in the patient's Circle of Care—as we call it—is authorized to access the patient's records? Moreover, can we provide transparency and some control for patients in terms of how their data is being used and accessed?

TC That sounds like an admirable goal, but can you actually make this happen? It sounds like you might be trying to move the immovable object here.

DE Actually, I wouldn't say that we're working to move or replace anything. In fact, by regulation, we're precluded from replacing the existing EMR systems. We're definitely not aiming to capture all the data an EMR needs to retain since each custodian organization—that is, each health care provider—remains legally responsible for maintaining its own records.

However, as they continue to update the medication profile in some amount of detail for each patient, what we're hoping to do is to *integrate* with all those EMRs so

they can keep shared state about these patients' records up to date while also leveraging that information in such a way that everyone within a patient's Circle of Care can readily review it.

We're definitely not looking to change the world here, but only to make it possible for doctors, pharmacies, and patients to share a common view of a patient's prescription history across all the different providers the patient has used over time. That's a big enough challenge in its own right and it's an important goal, but it's also a lot more practical than attempting to replace all the EMR systems out there.

RM It would seem that one of the keys to this problem has to do with keeping track of who the patients are and having some way to identify them across all these different parties—while also managing access control, obviously. Can you elaborate a little on this, especially the issues around digital identity?

DE The system requires unique identifiers, of course. So far, that means we're able to recognize patients by way of the health care card numbers issued to them. By the same token, we recognize doctors by their license numbers and pharmacies by their accreditation numbers. But we also need to be able to accept some of the other identifiers accepted out in the world today, so we're working to come up with a more robust registration process. What this really comes down to, though, is taking whatever steps are necessary to guard against having multiple profiles on the system for the same patient, doctor, or pharmacist.

Built on a foundation of proven, familiar open-source software—Hyperledger Fabric and Hyperledger Composer—HealthChain presents no obvious impediments to universal deployment.

Access to a network of patients and providers formed using HealthChain, however, is limited only to credentialed participants, who in turn are granted access only to certain information assets on which they're allowed to perform a specific set of functions.

This means that, ultimately, access-control lists may prove to be the key to resolving the longstanding health care data-management stalemate, since they're not only the means by which access to objects can be bound to each participant type, but also the means for defining the operations permitted on any given object.

RM Now that you've told us what you set out to accomplish with your application, tell us what you actually did.

DE Well, first off, we're using Hyperledger Fabric and Hyperledger Composer, both of which come out of a blockchain-related project backed by the Linux Foundation. Fabric basically provides a bunch of tools around a blockchain component, which effectively gives us an indelible record of all the idempotent transaction requests that come in. We think of these transactions as "smart contracts," but in any event they're transactions that update state. Which is to say Hyperledger Fabric gives us a way to maintain a global state database.

The underlying database technology is CouchDB, which gives us an object-store for handling JSON (JavaScript Object Notation) objects, with MapReduce being used to

Block-chains are generally really good at handing out a bunch of keys. Unfortunately, they're not all that good at managing those keys.

apply indexes and perform fast queries against all these JSON object structures. Another important technology for us that came along as part of Fabric is Kafka, which provides for high-performance message streaming. That provides for event notifications between peers as well as for the delivery of notifications to other processes.

Hyperledger Composer, meanwhile, is a framework built on top of Fabric to help you define what the Hyperledger world thinks of as a business network, where all the participants within a network, as well as all the various assets the network is expected to manage, are identified, along with the control rules that govern access between all the different participants and their respective assets. In our case, we leverage this business-model language to define rules for a business network expressly designed for the management of prescription records.

RM What are some of the novel things your architecture does to accomplish this?

DE Hyperledger Fabric is what's known as a permissioned—or private—blockchain, which requires everyone to have an identity recognized by that blockchain. Once you're issued an identity, you're also issued PKI (public key infrastructure) certificates, which you can then use to identify yourself. I should add that blockchains are generally really good at handing out a bunch of keys. Unfortunately, they're not all that good at managing those keys. So, if you're building a system that utilizes blockchain—or PKI in general—you're going to need to come up with your own way to manage those keys.

We've built a microservice architecture, which resides on top of this whole Hyperledger stack, that is focused on

matching authenticated users with their respective keys or certificates. Once that's accomplished, the idempotent transactions can then be executed and digitally signed using the key belonging to that individual.

Another dimension is that one person can have multiple credentials on the same network. This, too, has to be managed. Obviously, everyone can become a patient at some point, so all of us qualify for patient credentials. Only some people qualify as doctors, however, so they need to present appropriate credentials that the system recognizes. There also will be clinical administrative staff that need to use the system. Even though they don't have the credentials to write prescriptions, they might be able to enter draft prescriptions for a doctor to approve later. What this means is that one user could potentially be associated with multiple keys—each of which defines a different set of things they're allowed to do within the business network.

This is where access controls come in. Each participant type associated with a person on the network determines who is allowed to perform particular functions. If you happen to be someone who has multiple keys on the platform, each key represents a different set of things you're credentialed to do. For example, if you're a patient and you query the system through HealthChain's Portage APIs to "show me all patients," the only record you'll get back will be your own. A doctor, on the other hand, might say, "Show me all patients named Bob." But that doesn't mean the doctor should be shown the records for all the patients named Bob across the entire system. It just means the doctor should be shown the records of those patients

named Bob who have listed that particular doctor as part of their Circle of Care.

TC But if I have multiple keys, would you manage all those for me within the context of a single identity?

DE This is a bit of a slippery concept. In addition to defining participant types, we define the different application types that can integrate with the HealthChain network. This, for example, allows us to ensure that if you're connecting from an EMR application at a clinic, you'll be allowed to connect only as a doctor or a clinic administrator. It just wouldn't make sense for you to connect using your patient credentials.

As another example, it's common for doctors to connect from multiple clinic applications, so the certificate associated with an application instance ensures that we know which clinic the doctor is connecting from, since the authorization token issued to a doctor at the time of authentication is unique to that application instance. Regardless of which clinic application Dr. Jones might happen to be connecting from, it will be Dr. Jones's credentials that are used to sign transactions on the network.

RM But if Dr. Jones has a relationship with some particular pharmacy or clinic, and yet also happens to work at some particular hospital, how are you going to be able to sort that out? It sounds like there might be an organization tag associated with each particular transaction.

DE Exactly. We keep records on which clinic or hospital locations each prescriber is authorized to act on behalf of. And each application certificate is bound to a specific clinic location. In this way, we're able to guard against doctors spoofing the locations they're communicating from.

Although policy arguments have long been made for letting patients themselves determine who should be given access to their health care records, it turns out the most compelling reason may prove to be a very simple technical one. The traditional premise is that this information belongs to the providers, so determining who can receive updates requires either wading through data for a large number of patients or relying upon some number of data joins. By instead looking at this conundrum from the patient's perspective, it quickly resolves into simply determining which providers the patient has had dealings with previously and just leaving it at that—a far more elegant approach.

It just so happens the patient-centric approach could also lead to all sorts of new possibilities as more metadata accumulates around the defined patient-provider relationships. At minimum, this could make it possible to tune access controls at a much more granular level than is currently possible.

RM What was it about the Hyperledger technology that initially drew you to the platform?

DE First off, the CouchDB database technology underlying the platform is something that's proved to be quite robust and very capable of scaling. Data is managed in JSON format, which FHIR and other data standards generally support.

It also was reassuring to find that Kafka was being used to handle the messaging. Over the years I worked in the finance space, Kafka proved to be a solid contender in high-speed trading systems. Which is all to say I became

pretty enamored with Hyperledger Fabric's underlying technology stack.

And, of course, the Hyperledger Fabric stack contains a blockchain data structure that captures an indelible record of transaction requests. People associate a lot of different things with blockchain, but, at heart, it's just a data structure for transactions that's replicated, with the benefits being security and redundancy. So, here you've got Fabric, a technology focused on making sure a state database stays in sync with all the transactions it has processed, while also ensuring each peer gets a replicated copy of state. If, for some reason, someone should manage to get access to one of those peers directly and try to modify the data, that node will be dropped from the network. It's this tamper-evident property of blockchain that gives the underlying data yet another dimension of security. These security characteristics—combined with the robustness of the underlying stack—sold us on the Hyperledger Fabric technology.

RM A little earlier you cited an example that suggested the ability to adjust access permissions for a user according to the particular requirements of a situation. That tells me there's either a fair amount of work yet to be done programmatically to reflect policies capable of covering those situations or that you have some provisions that allow clinic administrators to deal with these sorts of outlier situations on a case-by-case basis. In general, how have you dealt with implementation issues like these as they've come up?

DE One of the challenges we faced right off had to do with figuring out the right model for managing the relationships

One big advantage of taking a patient-centric approach is that the patient profile provides all the information needed to enforce access controls.

between patients and their respective prescribers and medical providers. We initially took a traditional relational data-modeling way of thinking about things from the provider's perspective. But what we discovered was that, by taking the provider's view, every pharmacist, clinic location, and doctor would end up having a reference to any patient they had ever treated or served. As a result, we would have had to deal with huge arrays of patients associated with each provider.

From there, we moved along to a different relational concept where we started to think in terms of joining tables or joining objects that would keep references to each of the parties—effectively creating central relationship objects. The challenge we ran into there, though, had to do with the constraints of the Hyperledger Composer access controls. Basically, any given access control defines a participant, an asset, the type of access that participant ought to have to the asset, and a constraint. The problem with this was that by introducing a third mapping object at the time the access controls for these assets are evaluated, we would have been unable to access all the data required by the mapping object to properly evaluate whether or not that access ought to be granted.

This led us to take another stab at the problem by looking just at who really owns these relationships. And that's when things started to fall into place more naturally. We now keep those relationships as part of the patient profile. One big advantage of taking a patient-centric approach is that the patient profile provides all the information needed to enforce access controls. Then,

if somebody without a relationship to the patient wants to take a look at their records, we can just ask, “Does the patient have a relationship with this person who’s trying to get access to the records?”

RM This is quite a departure from the norm of developing something that looks at things primarily from the perspective of the clinic or the pharmacy.

DE Exactly, and that’s what makes it really exciting. It’s pretty clear this could lead to all sorts of new possibilities since we’re only just starting to attach more metadata to these relationships. For example, think about a scenario where a patient shows up in the emergency room of a hospital the patient has never visited before. This means the doctors there can’t call up any records for this patient. For exactly this sort of time-sensitive situation, it would be good to include some sort of provision that allows for break-the-glass, time-limited access to the patient’s records—maybe only for the following 24 hours or so.

Another possibility we’re exploring would amount to more of a lockbox-based security mechanism. Basically, for each patient, we already have a record of all the organizations or individuals that the patient has some sort of connection to, so we could just combine that knowledge with the cryptographic keys associated with each of those relationships. Then, using those keys, we could go one step further so that the identifier presented to a particular provider looking up a patient’s health care number would actually be unique to that particular patient relationship. This, in effect, would amount to a mechanism for anonymizing data access right at the source.

The point here is that, as we look at attaching more

metadata to these relationships, we can begin to really fine-tune these access controls at a very granular level.

RM What are some of the other challenges you've addressed along the way?

DE We were frustrated for a while because we were having a hard time getting performance out of our CouchDB implementation. We ended up introducing an Elasticsearch capability that leveraged events from transactions to trigger updates that would allow for a faster search capability. That was a bit of a hack, though. So, I'm happy to say we've since figured out how to optimize and apply indexes to CouchDB more effectively, meaning we've managed to reduce the need to use that Elasticsearch enhancement. Still, along the way, we came to a much better understanding of CouchDB and the possibilities that indexing offers there.

TC Now that you know much more about Hyperledger Fabric and what might be done to optimize CouchDB performance, what advice would you offer other developers looking to follow a similar path?

DE DevOps, DevOps, DevOps. And that's because extreme automation is really what's called for here. We've been very fortunate to have a team that's quite strong on DevOps. We were able to migrate our testnet environment from an IBM infrastructure to AWS (Amazon Web Services), while managing to preserve all our data and keys—along with the access they imply. Now that we're on the AWS infrastructure, we're going to be able to scale up even more. Similarly, as other people start to move in this direction, they're going to find almost all the infrastructure involves Docker containers. We also ensure

that any given microservice will accept only TLS (Transport Layer Security)-secured communications from other trusted services.

And because we're now operating in a world like this, it's not enough just to be a developer. You also need to be focused on automating your development environment and understanding your release process since things are just evolving too quickly *not* to be on top of that. There are a lot of headaches to go through before you learn how to manage keys effectively. And you need to be thinking in terms of seeding your environment for testing purposes right from the start so you'll later have multiple configuration options you can readily use to perform various scalability and performance tests.

As for all those developers who expect to find themselves in some sort of blockchain environment at some point, I'd say it's important to remember that these environments need to be kept alive forever. You start with a genesis block and then go on from there, with keys being distributed to each and every user and organization that has ever been a part of the chain. There has to be some reasonable means for managing that... *forever* essentially. So, as I say: DevOps, DevOps, DevOps is key.

Copyright © 2019 held by owner/author. Publication rights licensed to ACM.