Do Blockchains Have Anything to Offer Identity?

22 February 2018

Introduction

We wrote this paper because Identity and access management (IAM) has become central to our online interactions. Like a lot of infrastructure, when well designed and implemented, IAM is largely invisible to users. Because of this, many who are not actively engaged on this topic are not fully aware of its ubiquity and impact on our daily lives.

Blockchain is a broad class of relatively new data security methods, with certain properties of potential value in IAM. A great deal of excitement has come with blockchain. Many IAM start-up companies have launched identity registration solutions "on the blockchain", while others are developing new blockchain-inspired infrastructure for distributing attributes, which is a key element of IAM. Faced with a growing amount of associated hype and scepticism, we seek here to provide a balanced perspective, and to clarify the ways in which blockchain technologies may or may not serve the needs of IAM. Perhaps most importantly, we hope to provide a useful lens through which to evaluate current and new blockchain-based IAM solutions as they come along.

In thinking about how these exciting technologies can help with IAM, the starting point should be to appreciate what the first blockchains were designed to do, and then build carefully on that, with precise systems thinking. Thus, the paper should help those devising new IAM solutions, and those who may be acquiring solutions and needing to evaluate new blockchain-based approaches.

This paper is intended for information technology staff and management, and any others working on or with IAM technologies, and who are curious about blockchain and how it may impact IAM. First, we give background on blockchain, then on IAM, and then to recommendations to tie them together.

Authors
Steve Olshansky
Internet Society

Steve Wilson
Lockstep Consulting
Steve Wilson is an independent researcher, innovator, adviser and analyst in digital identity and privacy.
See http://lockstep.com.au

Contributor
Robin Wilton
Internet Society

Summary

IAM has become essential infrastructure for our interactions online. It is evolving rapidly, the stakes are high, and enterprises face an increasingly complex and puzzling digital identity landscape. There is growing concern that many digital businesses know too much about us, and control over identity should somehow be reclaimed by end users. So IAM is a hot topic, with new architectures, business models, and philosophies all in play.

Blockchain technology[1] is gaining attention, positive and negative, and no small amount of hype. Proponents are advocating its use for a wide variety of use cases, including IAM.

Blockchain came along at a critical time, with a suite of security related promises, many of them apparently applicable to identity, if only loosely. Numerous IAM companies have sprung up "on the blockchain," and broad claims have been made that this new family of solutions will disrupt traditional IAM. While use cases and requirements vary, simply adding blockchain technologies to an IAM system without careful consideration of all factors will not necessarily bring real improvement and may in fact have the opposite effect. Like any technology, blockchains are simply tools, to be carefully considered in the context of your particular environment. This paper presents a range of analysis and design issues to be considered when applying blockchain technologies to IAM, in consultation with relevant technical staff.

This paper seeks to impartially explore the fit between blockchain technology and IAM for personal identity use cases. We discuss the evolution of blockchain and its applicability to IAM, single out relevant aspects of the problem space, and identify key issues to be addressed when considering blockchain. Blockchain technologies are evolving rapidly, and the fit with IAM must be expected to change, but at this stage there is certainly no need to rush in, and no reason to fear missing out. We show that the original blockchains are generally not a good fit for identity management, and we try to sketch out, at this early stage, where new blockchain developments could suit the needs of IAM.

We presume familiarity with both IAM and blockchain. For anyone interested in learning more about either of these topics, there are some good references listed at the end.

1 At the time of writing, the terminology in this new field is still in flux. For the purposes of this paper, we use "blockchains" broadly to embrace the public distributed blockchains created for Bitcoin and other cryptocurrencies, as well as more recent developments, sometimes known as Distributed, or Shared, Ledger Technologies.

Blockchain, Blockchains, and "Distributed Ledgers"

Why is the blockchain (or, more broadly, "blockchain technologies") being promoted as a useful solution to current problems in online IAM? While there is a great deal of hype associated with it, and concurrently people proposing its use in a variety of settings, there are issues particular to identity that some view as solvable with blockchain algorithms and related spin-offs.

The Characteristics and Security Properties of Evolving Blockchain-Related Technologies

First generation public blockchains e.g. Bitcoin, Ethereum	**Advanced /special purpose blockchain-technologies** e.g. Corda, Fabric, Plenum, Hashgraph
Highly distributed, highly decentralised networks	Tend to be more concentrated; smaller numbers of participants or nodes
Permissionless	Access controls for write and/or read
Public peer-to-peer network of nodes	Can be private physical or virtual network
Immutable by sheer weight of numbers in the network	With a smaller pool of participants, tamper resistance can require traditional security at the nodes
Free and Open Source	Can be proprietary software, and/or commercial network

The original blockchain was developed to solve a particular problem: the "double spend" of cryptocurrency without a central administrator. The intrinsic problem with purely virtual currency is that nothing inherently prevents money from being duplicated; some sort of oversight is necessary to prevent double spend. While various electronic cash schemes have existed for decades (one of the first being *Digicash* founded in 1989), they always entailed a central authority to monitor double spends. Blockchain enabled the first peer-to-peer cryptocurrency – namely, Bitcoin – to operate with no intermediaries and no digital "reserve bank." A massive public network oversees every Bitcoin movement, maintaining an ever-growing append-only ledger – the blockchain – of every transaction ever made. Network nodes each run open source blockchain software and are rewarded for their participation by random allocations of Bitcoin.[2]

The original Bitcoin blockchain is highly distributed,[3] with no single point of control (or failure) and cannot be altered once written to. Decentralized, practically immutable, and cryptographically verifiable, this type of blockchain appears to lend itself to countless applications beyond payments, including IAM, to reduce fraud, remove bottlenecks, and trace the provenance of complex multi-party data. These properties are important to IAM and so there's been a rush of research & development around blockchain for IAM, among a

2 The details need not concern us here, but in brief, certain "full" nodes of the Bitcoin network are rewarded by a bounty payment which goes to the first node that completes a brute force "proof of work" task, the difficulty of which makes the ledger immune to tampering and counterfeiting. The odds of winning the bounty rise with the amount of computing power put into the task; running a full node is referred to as "mining."

3 Note carefully that there are several aspects to decentralization in blockchain technologies, and that the degree of decentralization envisioned by its original designers has not been even. The issue is too broad to cover at any length in this paper, but in brief, blockchain architectures can decentralize the *storage and availability* of the ledger, and/or the *process for reaching consensus* about the state of the ledger. The latter, known as "mining" in public blockchains, was hoped to remain highly distributed and thus resistant to corruption, but as things turned out, Bitcoin mining became so lucrative that it has brought massive concentration of this activity, and consequential compromises to Bitcoin's security, at least in theory. Another aspect of decentralisation is *governance*, which has yet to be properly organized in the public blockchains of Bitcoin and Ethereum. Decision-making for important software maintenance has proven difficult or even intractable with Bitcoin, and on occasion has fallen to just one person with Ethereum.

great many other things. The past four or five years have seen frenetic evolution. The original Bitcoin system and its closely related derivatives represent one class of *public* blockchains. More advanced descendent algorithms, developed for more complex use cases than cryptocurrency, deliver different combinations of properties.

Authentication and Authorization

As noted above, IAM is evolving rapidly, the stakes are high, and enterprises face an increasingly complex and puzzling digital identity landscape. When it is working well, most of IAM's mechanics remain neatly hidden from users, who are generally more interested in convenient and reliable access to services than being "identified" as such. Mobile technologies with powerful built-in cryptography and biometrics have become popular authenticators. At the same time, there is widespread frustration and growing concern that many digital businesses know too much about us, and that control over our information and our identity should somehow be reclaimed by end users.

When considering the disruptive potential of technologies like blockchain, it's all the more important to be clear about the problem we're trying to solve. If blockchain is thought to have the potential to improve the quality and availability of information about the parties we're trying to deal with, then let's first review what authentication and authorization are fundamentally about.

The central question in IAM may be framed like this: In a particular context, what do you need to know about a counterparty in order to be able to deal with him or her (i.e. accept a transaction or digital artefact from them)? In most business settings, it's less important to know *who* someone is than *what* they are. That is, for instance, what is their professional qualification? Or their membership of an organization, relationship with a service provider, country of origin, entitlements to government service, standing as a trade customer, or age, as applicable? These are the sorts of data (aka attributes) that are used in fine-grained (or attribute-based) access control decisions.

These sorts of questions should be asked at design time, when undertaking a risk assessment of the intended transaction, and analysing authentication and authorization requirements. Different ways can be explored for transaction systems to get the necessary identity attributes at the right time, for example when users register for services, or when they transact. This introduces a further set of design decisions: up-front identification when registering need not be as rigorous, for example, if other risk mitigations (such as real-time risk scoring to detect fraud) are available. When designing identification systems, we must decide what quality of information is needed, where that information will be obtained, and how it will be validated.

Terminology: The main actors in IAM
Identity Management analysis and design revolves around a set of actors or roles, as follows:

Subject (aka user) is the person or entity being identified or named in a transaction and typically being provided services according to their identity or attributes. Subjects are typically customers, employees, account holders and so on.

Relying Party (RP, aka Service Provider - SP) is an entity transacting with a Subject, providing services, and usually depending on a third party to confirm the identity or attributes of the Subject. Typical RPs are retailers, employers, financial institutions, government agencies, utilities and the like.

Identity Provider (IdP) is a party which vouches for a Subject's identity according to some agreed identification protocol. In classical IAM architectures, the IdP is classically a third party (like a government agency, educational institution, employer, or a commercial identification service), but in conventional business, many RPs such as banks take responsibility for identification and therefore act as their own IdPs.

Attribute Authority/Provider is often the IdP but may also be a separate entity outside the direct control of the IdP. Similarly, there could be external "attribute brokers" which procure and aggregate attributes from any number of sources. It is ultimately up to the RP to make a decision as to how much trust to place in attribute sources.

Some IAM settings involve separate Attribute Providers and additional intermediaries. In general, Relying Parties tend to bear most of the risk in a transaction, and usually have the final say in whether an identification protocol is fit for purpose or not. The contractual obligations of IdPs to RPs, warranties, and liability arrangements are perennial topics of discussion in IAM. Hence in conservative or higher risk applications like government, healthcare and financial services, RPs often act as their own IdPs.

Claims / Attributes / Assertions

In IAM, the things we need to know about counter-parties are variously referred to as claims, attributes, or assertions. Many risk management formalities, like banking's customer identification rules, focus on common subsets of attributes, such as "legal name" from an official government document, date of birth, and residential address.

By framing IAM to concentrate on specific attributes in different contexts, we can reduce the accumulation of extraneous information (which increasingly represents a liability for many enterprises, in light of the epidemic of data breaches). The "Need to Know Principle"[4] applies, which is good for privacy, through the minimization of data collection and disclosure. And the relevance to authentication of blockchain technologies should become clearer. Depending upon the risks involved in the intended application, IAM designers must decide how much confidence is needed in attributes presented (asserted) by users. Is self-asserted data sufficient, or is external validation by a trusted third party necessary? The World Wide Web Consortium (W3C) Verifiable Claims Working Group[5] is doing some promising work to enable external verification of claims, regardless of where they are stored.

4 https://security.berkeley.edu/need-know-access-control-guideline
5 https://www.w3.org/2017/vc/charter.html

Provenance

One of the hot topics in IAM now is *provenance*. If we concentrate on the precise attributes that are used to authenticate the people we deal with, then how do we know the attributes are reliable? That is, what do we need to know about the attributes? The most obvious authentication metadata of interest includes the attribute issuer: Which school issued the university degree? Is someone's age obtained from a public registry of births, a department of motor vehicles (DMV), or a social network? It can also be important to know the age of an attribute, its expiry date, where and how it's been stored and protected from tampering, and how the attribute is bound to the subject.

For some personal attributes, there is an obvious authoritative source. Professional qualifications for example are usually issued by professional bodies, and credit card numbers are created by issuing banks. But other claims, like employer or home address, might well come from a number of sources. The concept of an "Attribute Economy" is emerging in some personal data discussions[6] and we can expect to see a contestable market of attribute providers to emerge. But in all cases, an attribute is only as reliable as our certainty about where it came from and our confidence in the source.

Blockchain Evolution

Fundamentally, the problems inherent to peer-to-peer cryptocurrency like Bitcoin are different from those of identity and access management; these differences must be understood before attempting to match blockchain technologies to IAM. Attempts to address these problems for different use cases have driven significant changes in the way blockchains operate and perform.

Blockchains arose as a novel way of overseeing certain transactions between people who don't need to know or trust each other, and who choose not to rely on any central administrator. In a more general sense, blockchain technologies can be used to establish consensus about the state of a shared data set, composed of multiple real-time contributions, without central oversight. With Bitcoin, consensus is specifically about the order in which attempts are made to move cryptocurrency, in order to detect and prevent detect attempted double spends. The ability to obtain and record consensus about the order of data stored in a record can be useful in contexts beyond cryptocurrencies, and this has been one of the key drivers for other proposed uses — including for IAM.

In some use cases, the participants in a complex transaction are either competitors (like trading banks, or pharmaceutical companies) or come from far flung sectors with no common supervision (such as the diverse shippers and suppliers involved in international trade). Blockchain technologies promise to streamline the way transactional data sets like trade manifests, supply chain records, and complex financial deals are assembled in real time and settled.

Peer-to-peer cryptocurrency is a highly specialised application, with unusual design assumptions and constraints. As more complex use cases for blockchains emerged, the features and design options of the architectures shifted markedly, as explained in the following sections.

6 https://www.linkedin.com/pulse/youre-become-part-attribute-economy-nathan-kinch

Public or Private

The Bitcoin blockchain is a public and famously "immutable" data structure. To support the monitoring of double spends, the blockchain preserves all Bitcoin transactions for all time, with no restrictions on who can read the history. But when enterprise applications for blockchain technology were first contemplated, the primary consideration was confidentiality, and thus, some of the earliest blockchain spinoffs were *private* blockchains of various forms, with access controls over who can read and/or write to the record. There are many subtleties to do with blockchain permissions, which we explore further below.

"Trustless" or Managed

The first blockchains were avowedly "trustless." The philosophy of peer-to-peer cryptocurrency rejects central reserve banks, government oversight, and indeed all administration. The singular achievement of the original Bitcoin blockchain was to allow total strangers to reliably move real value without knowing anything about each other, and without relying on any third party. Conventional wisdom holds that any security system rests on the triad of People, Process and Technology. Bitcoin transactions are secured by technology alone, and that's what "trustless" means in this context.[7]

The original blockchain also had no need for off-chain cryptographic key management. Most crypto systems require certainty about which keys go with which users (and which key metadata, like key lifetime and revocation status). And they need *key lifecycle management*, to renew, revoke and replace users' keys as necessary. But Bitcoin needs none of that. Bitcoin account holders self-register (infamously sidestepping financial regulators' customer identification rules) and they accept full responsibility for safeguarding their wallets and private keys. It's everyone for themselves; if you lose your Bitcoin wallet key, there are no people or processes to guarantee your deposit or help you recover.

Use cases beyond cryptocurrency become much more complex. For one thing, they usually require *permissions*, because it's not normally acceptable for corporate records to be public, nor do businesses usually outsource their operations and software maintenance to anonymous volunteers. Permissions for reading and writing to a managed blockchain requires the sort of management which Bitcoin did away with for its purposes. When management must be folded back in, different consensus algorithms from Bitcoin's "proof of work" may be more efficient, and the system can be concentrated in just a few nodes instead of distributed across thousands. Bitcoin's requirement for a fully distributed ledger, with no centralized authoritative source, simply doesn't work for many enterprise or corporate use-cases.

Another driver to for hybrid, administered blockchains comes from the reality that higher risk IAM use cases often necessitate trusted third parties to validate user identities and/or attributes. The original blockchain does not support external validation of claims, but rather is purpose-built simply to provide a verifiable and practically immutable distributed ledger.[8]

---

7 Of course, any software involves processes and people at the *design level*. Users at large must trust that the software developers know what they are doing, are on the lookout for – and will promptly fix – the inevitable bugs or errors as they come to light and are committed to the proper and efficient operation of the system as a whole. Or, if we don't really trust a software developer, we might rely on the assurances of an independent "trusted" auditor. So, trust is inescapable at some level, and with that caveat in mind, "trustless" is an apt descriptor for the essentially automatic *operation* of public blockchains and the lack of relationship required between parties transacting these cryptocurrencies.

8 See discussion of "Claims / Attributes / Assertions" in Section 4 above.

The primary objective of the original public blockchains, to reach consensus about a ledger in the absence of any administrator, can become moot if, after all, a third party is given a central role in the system. Some of the recent IAM-focused blockchain developments have followed detailed examination of available blockchains and a finding that they don't meet the needs of identity management.

Decentralised or Concentrated

The massive distributed networks of the archetypal cryptocurrency blockchains provide great resilience and redundancy. One of the design assumptions underpinning the original blockchain is that a simple majority of the network nodes will always remain independent; hence one of the vulnerabilities of Bitcoin is known as the "Fifty One Percent Attack"[9], in case of which the record can be interfered with, specifically by covertly distorting consensus.[10] As mentioned, the more enterprise-focused blockchains, like R3 Corda[11] and Hyperledger Fabric[12], are not massively distributed but *concentrated*, and operated privately. Rather than assuming a majority of nodes will remain uncorrupted, security in private blockchains requires more conventional approaches. IBM for example implements its blockchain-as-a-service as a pool of virtual nodes, which can physically run on just one mainframe computer, with protective measures including hardware security modules, containerization, and highly vetted (trusted) operators[13].

"Immutability"

This is one of the storied properties of the original blockchains. It is true that the effort required to subvert a public blockchain network and then counterfeit and re-install the entire history of blocks, is utterly infeasible.[14] This extreme tamper resistance is a means to the end of resisting double spend, for according to the Bitcoin philosophy, the community needs to be able to reference every transaction ever made, without the benefit of a central transaction log. For the Bitcoin community, the cost of maintaining a decentralised ledger is accepted as the price paid for a currency system free from central banks and regulators (with the "proof of work" consensus algorithm, the cost translates directly into gigantic computation loads and power consumption). In other applications however, the huge expense incurred by *public* blockchains may be disproportionate, and traditional tamper resistance measures may be adequate. It should also be noted that the immutability of public blockchains is qualified by the power retained by their software maintenance teams to create branches (or "forks") through software updates, which can render old records inactive from the date of the fork.

9 https://en.bitcoin.it/wiki/Weaknesses#Attacker_has_a_lot_of_computing_power

10 Remember that with the public blockchains, "consensus" is reached about the order of entries and the overall veracity of the ledger; a Fifty One Percent attack can in principle see the ledger tampered with, but there are other ways to create fraudulent transactions, such taking control of an account holder's wallet and private key(s).

11 http://www.r3cev.com/blog/2016/4/4/introducing-r3-corda-a-distributed-ledger-designed-for-financial-services

12 https://hyperledger.org/projects/fabric
https://www.ibm.com/developerworks/cloud/library/cl-top-technical-advantages-of-hyperledger-fabric-for-blockchain-networks/index.html

13 http://www-03.ibm.com/press/us/en/pressrelease/51840.wss

14 However, tampering directly with a ledger is not the only means of attack. It should be noted that, in the words of security leader Bruce Schneier, only amateurs seek to attack technology; professional criminals attack people. No blockchain is immune to attacks on, for example, individual users' keys; the immutability of data written to a blockchain does not prevent fraudulent transactions being formed off-chain and injected into the ledger. https://www.schneier.com/crypto-gram/archives/2000/1015.html

Blockchain Technologies and Authentication

At present we see a contest of ideas around the methods of delivery of attributes and mechanisms to prove their provenance. For a decade or more, the classic federation frameworks[15] anticipated that *Attribute Authorities* would operate alongside *Identity Providers* (IdPs) and serve up attribute information in real time. An alternative approach is to equip end users with Personal Data Stores or Attribute Wallets, based either in the cloud or on mobile devices, and arrange for pertinent details to be transmitted more or less directly to service providers (SPs) on demand.

And now blockchains provide another type of platform for distributing attributes. One advantage of many blockchains, especially the public instances, is discoverability. Their distributed nature and transparent, open source software, installed across the world, means that finding records is straightforward and requires no central directory or addressing scheme.[16]

It is worth noting at this point that the transparency of public blockchains creates tensions with privacy principles. The original Bitcoin architecture exposes every single blockchain entry to the world, so that oversight of the currency system could be crowdsourced; the entire transaction history has to be available for all to see. When these types of blockchain are re-purposed for applications like IAM, additional privacy controls become necessary, such as separately encrypting transaction payloads before they are slotted into or referenced from blockchain entries or wrapping extra access control layers around the native blockchain algorithm to restrict who can read from (or write to) the ledger. Other privacy challenges are presented, such as the uniqueness of account holders' key pairs. Every transaction undertaken with an account is indelibly recorded for all time; a Bitcoin key or "address" might be un-named and, as such, anonymous, but it forms a permanent index (or "correlation handle") for one's blockchain history and is thus a significant privacy risk. One can find all records using a particular key and correlate them to a particular profile.

One special topic in IAM has resonated especially with blockchain concepts – *Self Sovereign Identity* (SSI). SSI proponents reject the tight control usually exercised by governments and big business over citizen and customer identities and call for greater self-determination in the way individuals represent and reveal themselves online, and decentralization of identity issuance. However, for some higher risk use cases, a trusted external authority to validate claims or assertions is needed. The goals and tenets of SSI have been a long time coming, and mostly predate blockchain, but have gelled around the emergence of practical distributed ledgers inspired by the early public blockchains. The quality of "self-sovereignty" evokes literal or metaphorical ownership of identities by the people concerned, and a regaining of control.[17] The decentralisation and availability of blockchains are thought by many to be a good fit for SSI, and a great deal of R&D is currently underway; see for example the *Sovrin Foundation*[18] and its new consensus

15 See for example the foundational "Identity Metasystem" from Microsoft https://www.identityblog.com/?p=355, the US National Strategy for Trusted Identities in Cyberspace NSTIC https://obamawhitehouse.archives.gov/blog/2010/06/25/national-strategy-trusted-identities-cyberspace, and the GOV.UK Verify program, https://www.gov.uk/government/publications/introducing-govuk-verify/introducing-govuk-verify.

16 It should be noted that finding records on a public blockchain is one thing but interpreting them is another. With cryptocurrencies, transactions and their meaning are simple, but when more complex information is stored in blockchain entries, the semantics needs to be worked out on behalf of all users. The rule setting tends to require a central policy authority, which was something the first blockchain designs sought expressly to banish. Semantic interoperability of blockchain records across complex use cases such as IAM needs careful consideration.

17 The Self Sovereign Identity movement has much in common with Vendor Relationship Management (VRM); see https://cyber.harvard.edu/projectvrm/Main_Page

18 https://sovrin.org

algorithm *Plenum* which focuses on the reliability of certain attributes, especially relating to one's online graph. Important R&D is also being done under the *Distributed Identity Foundation*[19] to manage IAM metadata like revocation status, and by the *Rebooting Web of Trust* project to distribute public keys without reliance on central authorities, thus reducing single points of failure.[20]

Important Elements of Blockchains

If you are evaluating candidate IAM solutions involving blockchains or undertaking your own original R&D in identity management, you should explore the following issues. For authentication and authorization purposes, the following properties of blockchain technologies are especially important:

Availability and resilience: The public blockchains are massively distributed and almost universally accessible (they have to be highly available in order to constantly support their cryptocurrencies). Synchronisation and replication are taken care of automatically, with blockchains maintaining an agreed state of the record at all nodes.

Discoverability (of attributes): one of blockchain's more under-stated properties is discoverability. No complicated addressing or directory schema is needed to reach most blockchains; all participants' software knows where the ledger is. This accessibility may be useful for globally scalable IAM systems; users' attributes will be available round the clock in a uniform virtual location. On the other hand, **the semantic meaning of attributes on a blockchain need to be understood and agreed to *off-chain (i.e., separately)***. While a distributed blockchain is technically available to all users, it seems possible that separate communities in the broad system might have their own unique interpretations (or classification codes) of what their attributes means; so semantic *interoperability* does not necessarily follow from discoverability.

Guidance on Blockchain Technologies for Authentication

The purpose of the cryptocurrency blockchains was to have a leaderless network maintain consensus about the state of a ledger, so that reliable transactions can occur without any central administrator. Decentralization is expensive; in a sense, it is a 'high energy state' which requires constant effort to support. The shared overhead of the consensus algorithm is the price paid by Bitcoin users in foregoing management. Most IAM applications are intrinsically different from the decentralized ideal of cryptocurrency.

Consider off-chain processes. An identity management system usually bridges several domains, in order to bring information about users to the systems they are interacting with. Decisions need to be made about which identity data is pertinent, who vouches for it, and how it is kept current. These design and operational processes often involve third parties or authorities of some sort, which may not work with the decentralisation of a blockchain. Remember that the public blockchains utilize enormously compute-intensive networks, principally as a result of the assumption that no third parties or administrators are involved. If off-chain authorities are in fact required in an IAM system, then the architects must

19 http://identity.foundation

20 http://www.weboftrust.info/

accept at least a degree of central administration, and the philosophy of a distributed blockchain might not be so important or advantageous.

What do you need to reach consensus about? Consensus has been an important albeit niche topic in computer science for many years, and numerous algorithms predate the now-famous "Proof of Work" utilized by the first blockchains[21]. Database designers and gaming developers have long grappled with the question of ruling on what happened first. Remember that with cryptocurrency and most transaction-oriented blockchains, consensus concerns the *order of events*, in order to resolve double spend without an arbitrator. On the other hand, in identity management, there may not be the equivalent of a "double spend" problem. Fundamentally, identity is not as transactional as currency. So, in IAM use cases, take time to review what decentralised consensus would be about, and how much of a factor it would be.

Consider too the different security and recording requirements through the identity lifecycle. When a user is enrolled into an IAM system, certain facts about them need to be (usually) verified and recorded to be accessible later. When the user needs to access a system, some of those facts may need to be presented to counterparties and checked by them in real time. Certain transactions require that the user's identity or other attributes be bound to digital artefacts; audit logs may need to be maintained to various standards and referred to later with reasonable integrity. All these types of activities place different demands on IAM records, be they traditional directories or newer blockchains. One special consideration is a significant time delay entailed by some consensus algorithms. The Bitcoin blockchain famously takes *on average* of 10 minutes to refresh the ledger, which can constrain the response time for certain identity lifecycle operations.

Consider Key Lifecycle Management – meaning, in general, making sure the right keys are in the right hands and they stay there – is core to most identity management, but irrelevant to the first blockchain platforms. With Bitcoin, no one is supposed to care who they are transacting with, so the system doesn't need to make any assurances about private key custody, nor the association of public keys with certain individuals. For this reason alone, marrying IAM with public blockchains can create wasted effort: it is not obvious that decentralizing consensus is beneficial when some authority is needed for key management. The newer IAM-specific blockchain technologies should be alert to this need for fundamentally stronger key lifecycle controls than required by the first blockchains.

Consider private key safety. Related to key management is the specific problem of end user's private keys. The Bitcoin blockchain system infamously takes no interest in how its end users look after their private keys (i.e. their cryptocurrency wallets). Once it became apparent that private keys could be lost or stolen by hackers, a vigorous market of wallet solutions emerged, including cloud-based key stores, mobile phone storage, backup services, and personal hardware security modules. In the more advanced blockchains, hardware key management is also becoming topical.

Blockchain maintenance. One of the most conspicuous points of difference between the first blockchains and their descendants is management of the core software. Public blockchains tend to be maintained by open source volunteers while some of the newer platforms are closed or proprietary (or at least they can start out that way before being

21 https://www.khanacademy.org/economics-finance-domain/core-finance/money-and-banking/bitcoin/v/bitcoin-proof-of-work.

open-sourced). While the goal of open source tends to dominate, a practical concern for some enterprise IAM implementers is the dependability of software maintenance. When bugs or urgent design improvements arise, businesses can want certainty as to when fixes will be deployed. With the Bitcoin blockchain, certain design problems have taken years to be resolved; with Ethereum, one major bug led to a unilateral decision by the founder to "fork" that blockchain, leading to multiple incompatible records and variations of the currency.[22]

Conclusion

Blockchain technologies are collectively a work in progress. Despite early excitement about their general security properties, on closer inspection we find that the original public blockchains are generally not a good fit for Identity and Access Management. The objective of cryptocurrency – to exchange electronic cash without intermediaries and without trust – is fundamentally different from that of enterprise IAM, which typically requires much more rigorous key lifecycle management and access controls than public blockchains offer. On the other hand, several new blockchain technology developments show promise for improving particular aspects of IAM, such as the provenance of identity attributes and keys. We recommend that any ongoing examination of blockchain technologies for identity begin with a clear problem statement, and an appreciation of the nuances in blockchain security.

22 https://www.coindesk.com/ethereum-executes-blockchain-hard-fork-return-dao-investor-funds.

References & Further Reading

Want to really understand how bitcoin works? Here's a gentle primer"
https://arstechnica.com/tech-policy/2017/12/how-bitcoin-works/

"Immutable agreement for the Internet of value," Sigrid Seibold & George Samman, KPMG, 2016 https://assets.kpmg.com/content/dam/kpmg/pdf/2016/06/kpmg-blockchain-consensus-mechanism.pdf

"Blockchain plain and simple," Steve Wilson, Constellation Research, January 30, 2017
https://www.constellationr.com/blog-news/blockchain-plain-and-simple

"Blockchain Security for Digital Identity," Adam Migus, September 2016
https://medium.com/@amigus/blockchain-security-for-digital-identity-e10c8750cf9c

Decentralized Identity Foundation - DIF
http://identity.foundation

"Corda: An Introduction," Richard Gendal Brown, James Carlyle, Ian Grigg & Mike Hearn, R3, 2016 https://docs.corda.net/_static/corda-introductory-whitepaper.pdf

"Overview of Swirlds Hashgraph," Leemon Baird, Swirlds, 2016
http://www.swirlds.com/wp-content/uploads/2016/06/2016-05-31-Overview-of-Swirlds-Hashgraph-1.pdf

"Still don't understand blockchain? Let's untangle the wires"
https://www.weforum.org/agenda/2017/11/blockchain-bitcoin-ethereum-tech-explained/

"Do You Need a Blockchain?," Morgan E. Peck, IEEE Spectrum, September 2017
https://spectrum.ieee.org/computing/networks/do-you-need-a-blockchain

"Introduction to Identity Management"
https://meetings.internet2.edu/media/medialibrary/2014/11/06/20141028-dors-intro-to-idm.pdf

Online Identity: Who, Me?
https://www.internetsociety.org/resources/doc/2016/online-identity-who-me/

Internet Society Identity Resources
https://www.internetsociety.org/issues/identity/