Journal of Information Technology Management

ISSN #1042-1319

A Publication of the Association of Management

BLOCKCHAIN BASICS AND SUITABILITY: A PRIMER FOR PROGRAM MANAGERS

DAVID C. CHALLENER
JOHNS HOPKINS APPLIED PHYSICS LABORATORY
david.challener@jhuapl.edu

MARIA E. VACHINO
EASY DYNAMICS, INC.
mvachino@easydynamics.com

JAMES P. HOWARD, II
JOHNS HOPKINS APPLIED PHYSICS LABORATORY
james.howard@jhu.edu

CHRISTINA K. PIKAS
JOHNS HOPKINS APPLIED PHYSICS LABORATORY
christina.pikas@jhu.edu

ANIL JOHN
U.S. DEPARTMENT OF HOMELAND SECURITY
anil.john@hq.dhs.gov

ABSTRACT

Blockchain is a peer-to-peer transaction system popularized by cryptocurrency applications. Today, many new platforms and use cases for blockchain technology have been proposed. In this article, we provide a primer for program managers in both government and industry for deciding when a blockchain might be an appropriate technical solution. This is important for these program managers because, despite the hype, blockchain is not a panacea and in fact can create exceptional risk, particularly when misapplied. In some cases, as we describe, risks associated with blockchain are mitigable and the technology may be worth the investment. Regardless, program managers need to understand the technology to make informed risk-based decisions about its use.

Keywords: blockchain, program management, enterprise architecture

INTRODUCTION

Given all the media hype, it is not surprising that expectations are riding high for blockchain technology. Since 2014, more than 100,000 stories have appeared in various media, most of them favorable (see Figure 1), and many characterizing blockchain as a game changer, a revolutionary technological innovation with the potential to disrupt and transform many industries. However, as is often the case, the level of excitement surrounding an innovation does not always match its eventual utility, nor does it necessarily translate into quick adoption or the realization of all promises made.

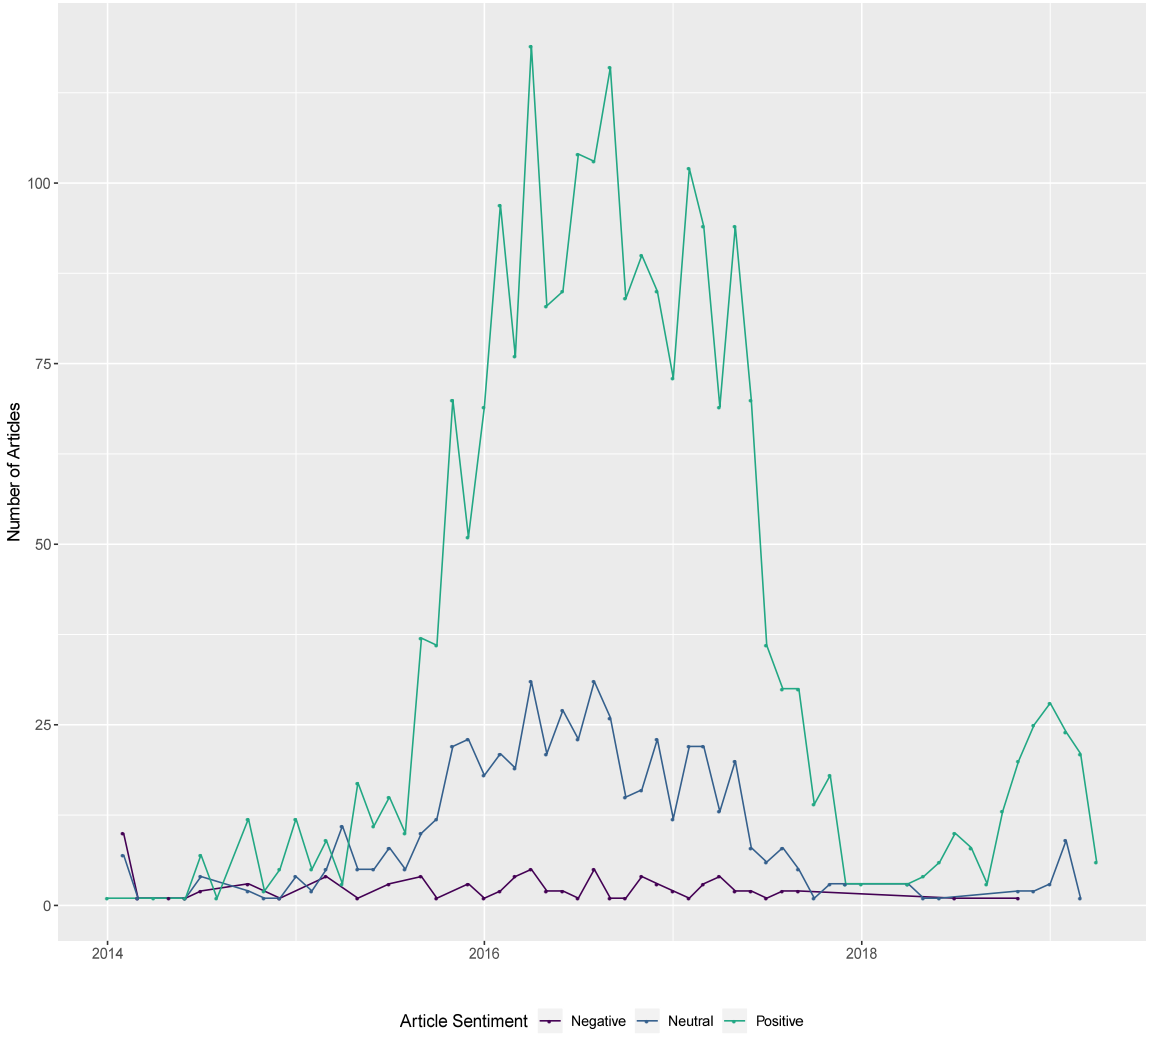

Figure 1: Sentiment Scores of News Stories on Blockchain in the Past 4 Years

This is particularly true of blockchain technology, which was popularized by the proposal to use it for decentralized cryptocurrency in "A Peer-to-Peer Electronic Cash System" [10]. The money made by some early investors in Bitcoin [6] and other cryptocurrencies has led to levels of enthusiasm not typically seen for innovations in computer science [9]. However, that excitement has also led to an unusually high proliferation of misinformation about the capabilities and limitations of the technology. Although many of the promises made by

blockchain enthusiasts will never come to fruition, some may, and the unique characteristics of blockchain may likely also lead to innovations that have not yet been anticipated.

This paper provides the information needed to understand what those capabilities and limitations are so users can determine whether blockchain may be suitable for their use case and organization. This information is based on lessons learned during a 4-year investment in blockchain technologies by the Cybersecurity Division (CSD) of the Department of Homeland Security (DHS) Science and Technology (S&T) Directorate.

This paper first provides a high-level overview of what a blockchain is and how it fits in the data storage ecosystem. We present an outline of what should go into a blockchain proof of concept (POC) and what considerations should be monitored. We also present a flowchart to facilitate choosing an appropriate use case for blockchain.

What is a blockchain?

A blockchain is a distributed database:

- That is not owned, controlled, or managed by a central authority;
- Where the copies are held by multiple independent parties, or nodes, and are kept synchronized;
- Where new records are validated and added using consensus [1];
- Where each record, or block, is strongly linked to the prior one, forming a strong chain of records; thus
- Ensuring that records cannot be altered or removed.

This design can solve some problems previous technology could not, but it comes with several limitations that must be evaluated before blockchain adoption is considered. It is important to note that a blockchain is not a single "chain" of "blocks," as its name implies. A single copy of a blockchain is nothing more than a very limited database. Successful blockchain applications require many independent parties to each have a copy of the same blockchain or book of records, and to all work together to keep those copies validated and synchronized.

Essentially, blockchain is a new type of peer-to-peer shared data storage technology where the shared data can be a record of a transaction, a digitally signed document, a hash of sensor or other data, or even a small computer program (a "Smart Contract" or "distributed app").

Blockchains and the Government

Although blockchain is the enabling technology behind Bitcoin and other cryptocurrencies, its potential applications are far more diverse in U.S. Government (USG) and enterprise settings.

From a government perspective, there are a number of non-cryptocurrency applications of blockchain technology that are promising. These include greater supply chain visibility and efficiency; the potential for enhanced transparency and auditing of public service operations; increased confidence in the data from cameras, sensors, Internet of Things (IoT), and Network of Things (NoT) devices; the mitigation of forgery and counterfeiting of official licenses and certificates; and the facilitation of international trade and customs processes.

Conversely, there are a number of use cases where blockchain technology is being applied that would be better solved using existing technologies such as conventional databases, perhaps connected to service-oriented architecture (SOA) components such as web services, micro-services, and Enterprise Service Buses (ESBs). In some cases, an even simpler technology may be the best solution. When compared to traditional databases, blockchains are far less efficient in terms of computing power, time, storage, and network traffic. That is one of the reasons that blockchain may be the most efficient choice only in cases where multiple records signed by different authorities are used in transactions. In such cases, the use of blockchain could result in a reduction in paperwork and multi-party processing time.

Of greater concern are use cases where the inappropriate application of blockchain technology could lead to data privacy breaches or a reduction in security. To help program managers decide whether blockchain technology is appropriate for their use case, DHS S&T created a flowchart that outlines the core questions to ask when considering a blockchain-based solution (Figure 2).

Example Anti-Use Cases – When is a blockchain not appropriate?

- Storing sensitive information
  - Permanent availability of blocks enables adversaries to decrypt blocks should the encryption used to store data on a blockchain ever become vulnerable.
    - Data stored on a blockchain cannot be re-encrypted for any reason, including a stolen or compromised encryption key.

        - Encryption methods available today will not necessarily withstand attacks over the long term.
    - Potentially sensitive information should NOT be stored directly on a blockchain, even if encrypted; this includes medical records, personnel records, and personally identifiable information (PII).
- One entity controls all access to the database
    - A centrally managed database is far more efficient than a blockchain.
- Record immutability is not desirable
    - Because blocks cannot be altered or deleted, incorrect or bad records live forever. These errors can lead to confusion, embarrassment, or even liability.
- Systems where the historical record is not important or a change record is unnecessary
- Public databases that require high-transaction volumes or real-time updates
    - Blockchains update at random, although fairly regular, intervals. For example, Bitcoin updates occur approximately once every 10 minutes, and not all transactions made during that window are guaranteed to write in the next update. The fastest block- chains can update thousands of records a second. If that is not sufficient, a blockchain is not appropriate.

Once you have identified a use case where the inclusion of blockchain technology may be appropriate, the next step is to consider planning for a POC study.

POC demonstrations are required to determine how blockchain can be integrated into existing infrastructures and applications, and to understand and plan for the complexity of implementing a peer-to-peer data sharing technology. Without first completing a POC, it is unlikely that a full implementation will be successful because implementing a blockchain requires significant multi-party coordination and agreement, in addition to the typical challenges associated with integrating new technology into existing environments.

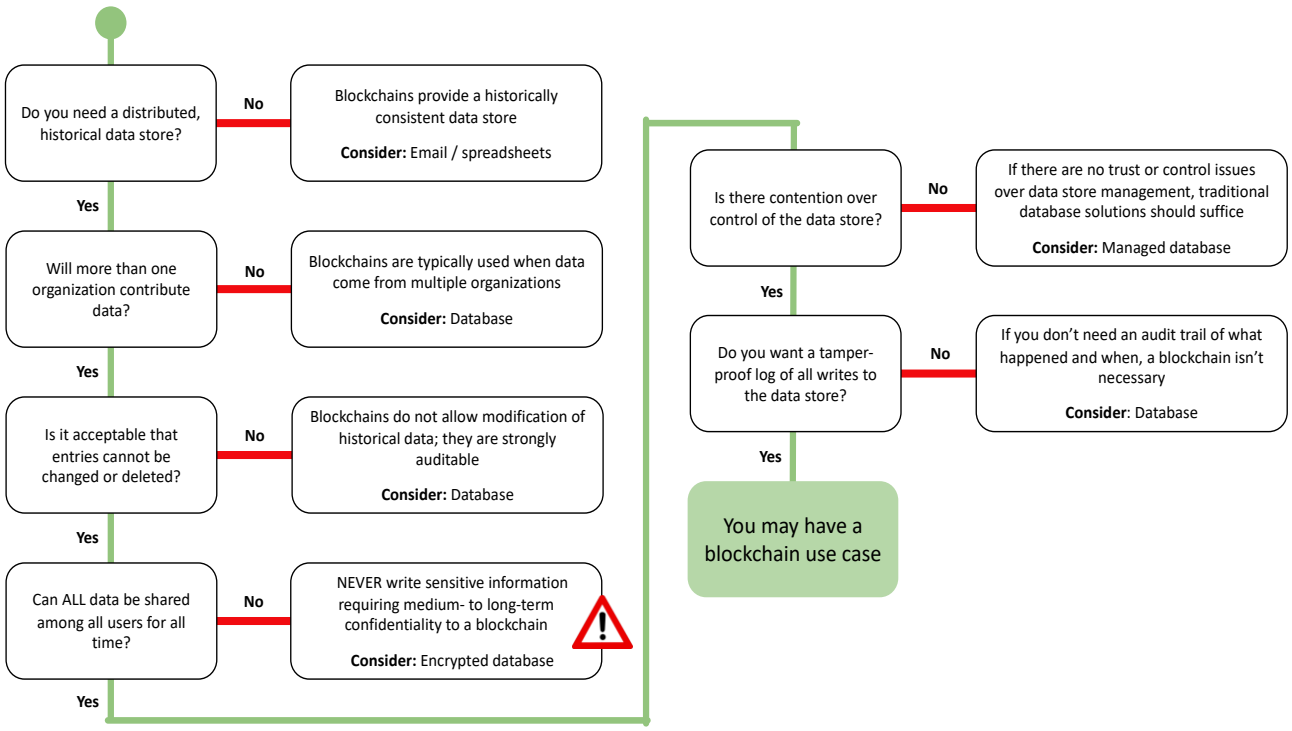

Figure 2: Do You Need a Blockchain?[1]

[1] Note: An earlier version of this figure was published in NISTIR 8202 [12]

The idea that blockchains remove or reduce the need for trust has become popular; however, the reality is that the lack of central management requires that a number of agreements be collaboratively reached with all organizations who will use the blockchain. The effort required to reach multi- party consensus on the required policy, governance, data payload definitions, decentralized key management frameworks, and data privacy requirements for a blockchain solution can be substantial. Therefore, although the technology is still maturing and not yet ready for critical operational deployments, organizations considering a blockchain solution should start planning for a POC now.

TECHNOLOGY OVERVIEW

At a high level, a blockchain is an immutable distributed book of records, or ledger, with certain unique properties. A ledger is an ordered record or log of events or entries (see Table 1). New items are appended to the bottom or end, and change history is maintained indefinitely. In comparison, a database typically shows only the current information after additions, modifications, and deletions, and deletions are usually permanent. Some blockchain implementations also allow small programs to be embedded in the ledger. These are called Smart Contracts and may provide some additional capabilities. For more detail, see Appendix B.

Table 1: A simple ledger

Date	Description	Debit	Credit
4/5/2016	Credit card payment	1023.77	
4/5/2016	Quick Haircuts	15.00	
4/6/2016	Jewelry to Go	723.05	
4/7/2016	Burger Shack	7.23	
4/15/2016	Mortgage payment	1200.62	
4/15/2016	Other loan	345.23	
4/17/2016	Payday		3675.43

General Blockchain Properties

The fundamental properties of all blockchain implementations are described as follows:

- A blockchain ledger is “append only” and immutable
  - Records cannot be changed or deleted.
  - Previous records exist as part of the permanent history.
- All blocks are ordered.
- Everyone (public blockchain), or all members (private blockchain), can read all records.
- All participants may have a copy of the entire blockchain.

Blockchain implementations may be either permissioned or permissionless, depending on whether participants must be authorized prior to adding a record to the chain. A permissioned blockchain is essentially a private blockchain that allows only accepted parties to add or read data from the chain, whereas a permissionless blockchain allows anyone to add or read data. Permissioned and private distributed ledger technologies may be more suitable for leveraging existing business relationships and regulatory frameworks, which form the majority of United States government use cases.

Consensus Methods

Because of the distributed nature of data access with blockchains, the technology requires a way to determine what is accepted for addition and what may be rejected. These processes are collectively called consensus methods, and different types of consensus methods yield different properties of the blockchain [4].

Proof of Work Consensus

Many public blockchains, such as Bitcoin, require a "miner" to complete a proof of work to add a new block. Blockchains that use proof of work to achieve consensus:

- Take a fixed average amount of time to add a single entry to the blockchain.
- Experience an increase in security as the number of independent nodes increases, and a decrease in security if the number of independent nodes decreases.
- Allow blocks of records to be added by whichever node first solves a complex puzzle, which is an energy- intensive process.

For these and other reasons, proof of work is unlikely to be used in a government blockchain implementation. A possible exception would be an implementation of a private blockchain that periodically records a hashed copy of its information into a larger public blockchain, such as Bitcoin.

Voting Consensus

The limitations of proof of work consensus have led to the development of additional consensus mechanisms based on voting among the blockchain nodes. Consensus methods that use voting:

- Can add blocks much faster, and
- Are far less energy-intensive [11].

Proof of Stake is the most well-known class of voting consensus algorithm.

Voting consensus algorithms need additional study to understand their vulnerabilities and limitations.

Proof of Authority

Other, more experimental, consensus mechanisms have been proposed for private blockchains, including proof of authority (PoA) [8]. At this time, it is unclear what advantage a PoA blockchain would provide over a traditional permissioned database. Anyone considering a PoA blockchain should therefore first perform an analysis of alternatives study that includes traditional database technologies before planning for a POC.

Data Validity

The immutability of a blockchain does not suggest the data recorded on the blockchain are true any more so than information found on the Internet is true. Immutability simply means that data cannot be changed or deleted once it is on the blockchain.

Therefore, users must take more care when writing data to the blockchain than they would with any other data storage technology. Organizations must ensure that parties with blockchain write privileges will validate their own data and follow all privacy regulations. If a Social Security number, for example, makes its way onto a blockchain, there is no method to remove it.

Another limitation of blockchain is that one cannot retroactively apply changes required by new privacy regulations. A database, on the other hand, does permit data to be updated or removed should new regulations be implemented.

Smart Contracts—Distributed Applications

Static records are not the only type of data that users can enter into a blockchain. Some blockchain implementations also allow small programs to be written to a block. Such programs are called smart contracts; however, the term must be distinguished from a legal contract. It is a logical, not a legal, contract, and represents a relatively immature technology. Because smart contracts are logical contracts, they might be more appropriately termed distributed applications.

For example, a legal contract might state "Sue Smith will get $1000 from bank account B when she turns 18 to help pay for college." When Sue turns 18, she presents this document to a court and obtains a court order to transfer the money. A logical or smart contract might state "Transfer $1000 from account B to Sue Smith's account when proof is received that Sue Smith is at least 18 years old." When proof is posted that Sue Smith is 18, the next time the blockchain is queried, Sue's bank account will have grown by $1000. In this case, calculating the state of the smart contract replaces going to court. Although this method can speed processes enormously, it is accompanied by a lack of flexibility. For example, what happens if Sue Smith graduates high school early and needs the money to attend college when she is 17? Smart contracts have to be extremely detailed to account for every potential circumstance, whereas with a legal contract such unanticipated occurrences can be adjudicated by a court or mediator.

Furthermore, there is the problem of interfacing the digital world with the physical world. What happens if the account number changes? How does the blockchain know that Sue Smith is 18? Someone or something must have entered that data into a block of the blockchain. Is that entity trustworthy? Is their information accurate?

In reality, smart contracts are (inflexible) small computer programs, not legal agreements. And because of the indelible nature of blockchains, such programs cannot be updated or removed if they contain errors. Any smart contract should therefore be carefully analyzed using formal methods if placing that program onto a blockchain is better than alternative solutions.

Smart contracts will execute on all nodes of the blockchain network and must achieve identical conclusions across all nodes to be valid. Accordingly, smart contracts are constrained in both scope and application and have limited usefulness.

QUESTIONS TO ASK BEFORE STARTING A BLOCKCHAIN PROOF OF CONCEPT

Before starting a blockchain POC, the first step is to clearly identify the problem to be solved, and then determine whether the use of a blockchain may be an appropriate part of a solution (see Figure 2).

The next step is to thoroughly document the use case to identify the organizations and individuals who will need to read from, and write to, the blockchain. Because blockchain deployments typically involve the cooperation of multiple organizations, it is critical to identify all stakeholders early in the process. At a minimum, the stakeholders must agree on which blockchain to use, how to distribute and manage the digital signature keys used for writing to the blockchain, which data will be stored on the blockchain, and which data will be stored off-chain. Cooperative governance is essential for successful

blockchain implementations; therefore, it should be planned for early in the process.

Determining which data will be stored on the blockchain, and which should be stored off-chain, requires careful consideration. No sensitive information should be placed on the blockchain (see Encrypted Blocks). The exception is when the data are sensitive for only a short period. When sensitive data are stored elsewhere, references to these data can be stored on the blockchain proper. At this point, the user needs to consult security and blockchain experts to help determine the blockchain design and algorithms appropriate for the solution.

Selecting a Blockchain Implementation

There has been an explosion in the variety of blockchain implementations since Bitcoin's arrival, with different types supporting particular data structures or implementations. For example, some implementations support smart contracts; others do not. Some can add thousands of blocks in one second; others average a new block every 10 minutes. Public blockchains permit anyone to add blocks; private blockchains restrict access to groups of users or organizations. Selecting the right blockchain is critical to the successful implementation of a blockchain-based solution.

It is difficult (and sometimes impossible) to move data from one blockchain to another, making it imperative to get it right the first time. Thus, it is highly recommended that users conduct a POC before attempting a full operational implementation. Blockchain standards are still emerging; if a blockchain choice is not compliant with an eventual standard, the implementation may need to be abandoned and replaced by a standards-compliant implementation. Expert advice may be necessary when making these choices.

What needs to be done prior to implementing a blockchain solution?

Data Payload Definition

Agreement must be reached regarding the data payload that each blockchain transaction record will contain.

For each blockchain entry, the following questions must be asked: What are the required data elements? What are the optional elements? What are the acceptable values? Which data elements should be salted and hashed? It is important to remember that sensitive data should never be stored directly on a blockchain, even if it is encrypted. There are, however, use cases for placing salted hashes of sensitive data on a blockchain.

It is also important to determine which data elements will be stored on the blockchain itself and which will be stored off-chain, then referenced on the blockchain. Storing off-chain data requires a traditional database in addition to the blockchain, as well as adequate Application Programming Interfaces (APIs) and protocols for locating and accessing the data in the database. Currently, there is no standard way of doing this [3]. For data referenced by a blockchain, careful change management for the location of the externally stored data must be provided because the records on the blockchain cannot be updated.

Policy/Governance

A consortium-based blockchain requires multiparty cooperation and consensus because it is a shared piece of infrastructure. Policies must be agreed upon in advance regarding the governance of the blockchain, including the mechanism for deciding when and how to update the blockchain software. Processes and policies for running a blockchain node, and for adding data to the blockchain, must also be agreed upon.

Policy: Key Management

Key management is critical to any blockchain implementation and can become complex if participating organizations do not share the same Public Key Infrastructure (PKI). In a decentralized environment, agreements must be reached regarding acceptable key management practices, digital signature algorithms, and revocation policies.

To address these issues, DHS S&T is funding the development of a decentralized key management system framework based on National Institute of Standards and Technology (NIST) standards [1].

ADDITIONAL CONSIDERATIONS

There are several additional concepts relating to blockchain technology that program managers should understand before recommending or selecting a blockchain-based solution.

Implications of a Quantum Computer

When a large general-purpose quantum computer (LGPQ) becomes available, one with thousands of qubits, all current blockchain implementations will be impacted. All widely used asymmetric signing and key passing algorithms (e.g., RSA and ECC) will be broken. Symmetric algorithms used to encrypt data will be weakened (so key sizes must double), and hash algorithms may also need to double in size to maintain their current level of security.

All blockchains depend on hash algorithms. Changing to a larger hash algorithm is comparatively

easy, but it is best to choose a quantum-resistant hash algorithm (at least 384 bits long) initially. However, many consensus algorithms depend on digital signatures to determine whether a block can be added to the chain. These will need to be abandoned or migrated to a quantum-resistant digital signature mechanism or protocol. There have been a number proposed, but none have been approved by NIST [3].

Encrypted Blocks

As noted in "Example Anti-Use Cases," sensitive information that requires long-term protection should not be stored on a blockchain. The most pressing reason for this is that cryptography and cryptanalysis are always being improved. For example, the Data Encryption Standard (DES) was state of the art when first released in 1977. It used a 56-bit key and was widely considered secure enough for commercial as well as Government use. By 1998, however, demonstrations indicated that DES keys could be recovered in 56 hours. One year later, this time was more than halved to 22.25 hours. By 2017, chosen-plaintext attacks could recover DES keys in 25 seconds. Essentially, for a person born in 1977, any PII records stored in DES are now completely insecure.

DES was able to be defeated by increases in traditional computing power. The advent of general-purpose quantum computers will have an even more dramatic effect on cryptography. It is expected that these computers will be able to break the public-key algorithms currently approved by NIST, including all of the public-private key pairs used by blockchains. Large-scale quantum computers will also be able to defeat the smaller keys currently used for AES symmetric- key cryptography. Unfortunately, advances in computing power are not the only threat to encryption on a blockchain. Even cryptographic algorithms invulnerable to increased computing power, including quantum-resistant algorithms, are still vulnerable if not implemented correctly. Many cryptographic applications are weakened by poor or incorrect implementations, a reliance on flawed random number generators, or backdoors introduced by developers, all of which can result in exploits even without advances in computing power.

Improvements in computing power and flaws in cryptographic systems are not the only reasons why sensitive data requiring medium to long-term encryption should not be stored using blockchain technology. There is the additional risk in blockchain if the key is stolen or leaked. In a blockchain system, the encrypted data are available in multiple independent and complete copies, including copies held by potential adversaries. If the key is stolen, the encrypted data become broadly available. Using a modern database system, full access to data is not common, and stealing the key is only part of the process of gaining access.

Information encrypted and stored in a database can be decrypted, then re-encrypted with a new key whenever the need arises, regardless of whether that need is driven by a lost or stolen key, a flawed cryptographic implementation, or an advance in computing power. Unfortunately, this is not also true for information stored on a blockchain, which cannot be rekeyed for any reason. Therefore, the data are at permanent risk of key theft (see Figure 3). The write-only nature of blockchains is a strength for some use cases and an unacceptable weakness for others.

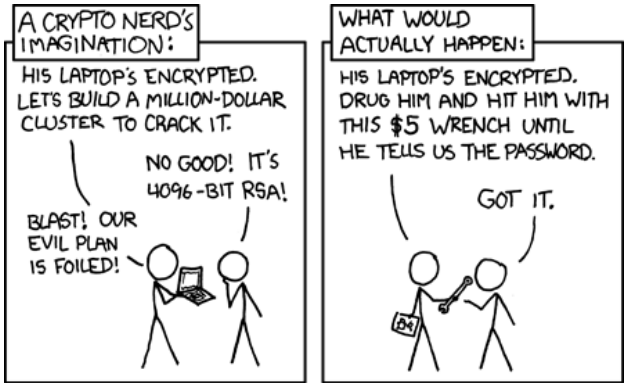

Figure 3: Security Risks (via xkcd.com)

CONCLUSION

Blockchain is a nascent technology and, as such, is progressing rapidly. It can be very useful when different organizations with different authorities wish to share a data store without sharing administrative authority. Blockchain can also be useful when the historical accuracy of entries must be maintained. This increases the value of conducting POCs, and it increases the risk of early adoption. Blockchains are not complete applications that can be acquired off-the-shelf and installed. Instead, a blockchain is a peer-to-peer network infrastructure component and foundational technology. It must be integrated into existing infrastructures, and client applications must be created to read from, and write to, the blockchain. Because this requires the distribution and management of keys, the individuals responsible for key management should be involved early in the POC process.

Furthermore, competition in the market place is fierce. Several large companies are developing blockchain technologies, and it is still unclear which technologies

will become widely adopted and which will be abandoned. The lack of current standards in the blockchain space is also a risk. In both cases, current solutions may need to be replaced in the future. These are indeed some of the challenges of adopting a new and emerging technology; however, the peer-to-peer nature of blockchain technologies amplifies these challenges.

Although blockchain technology will likely create a number of revolutionary methods for recording and sharing data, it is not appropriate for every situation. Decision makers seeking a blockchain solution are well advised to first investigate the characteristics of blockchains (strengths and weaknesses), understand where they are best used, and evaluate which type of blockchain would best meet their needs before making a final selection. Even with a good fit, it is best to start with a POC and carefully consider what will be required if a change (especially to the basic blockchain architecture) is necessary in the future.

REFERENCES

[1] Barker, E., Smid, M., Branstad, D., Chokhani, S. "A Framework for Designing Cryptographic Key Management Systems," National Institute of Standards and Technology, Gaithersburg, Maryland, SP 800-130, August, 2013.

[2] Cachin, C. "Blockchains and Consensus Protocols: Snake Oil Warning," 13th European Dependable Computing Conference (EDCC), Geneva, 2017, pp. 1-2.

[3] Challener, D. "Quantum Resistant Active Code Signing Using Blockchains," Technical Report, Johns Hopkins Applied Physics Laboratory, (March 6, 2018).

[4] David, B., Gaži, P., Kiayias, A., and Russell, A. "Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain," in Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2018, pp. 66-98.

[5] Fairley, P. "Blockchain world-Feeding the blockchain beast if bitcoin ever does go mainstream, the electricity needed to sustain it will be enormous," IEEE Spectrum, Volume 54, Number 10, 2017, pp.36-59.

[6] Holub, M. and Johnson, J. "The impact of the Bitcoin bubble of 2017 on Bitcoin's p2p market," Finance Research Letters, Volume 29, 2019, pp. 357-362.

[7] Kishigami, J., Fujimura, S., Watanabe, H., Nakadaira, A. and Akutsu, A., "The blockchain-based digital content distribution system." Fifth International Conference on Big Data and Cloud Computing, IEEE, 2015.

[8] Li, X., Jiang, P., Chen, T., Luo, X., and Wen, Q. "A survey on the security of blockchain systems," Future Generation Computer Systems, in press, 2017.

[9] Li, Z.Z., Tao, R., Su, C.W., and Lobonț, O.R., "Does Bitcoin bubble burst?" Quality & Quantity, Volume 53, Number 1, 2019, pp. 91-105.

[10] Nakamoto, S. "Bitcoin: A peer-to-peer electronic cash system," https://bitcoin.org/bitcoin.pdf, 2008.

[11] Vukolić, M. "The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication," International workshop on open problems in network security, Springer, 2015, pp. 112–125.

[12] Yaga, D., Mell, P., Roby, N., and Scarfone, K. "Blockchain technology overview," National Institute of Standards and Technology, Gaithersburg, Maryland, Interagency/Internal Report 8202, October. 2018.

ACKNOWLEDGEMENTS

Information in this paper is based on research funded by DHS S&T. Any opinions contained herein are those of the authors and do not necessarily reflect those of DHS S&T. For more information, please contact Anil John at anil.john@hq.dhs.gov.

AUTHOR BIOGRAPHIES

David C. Challener is a Principal Staff member of the Johns Hopkins University Applied Physics Laboratory. He has worked in the security field for 20 years. His main fields of interest involve Trusted Computing in its varied forms. His PhD in Applied Mathematics was from the University of Illinois in 1984.

Maria E. Vachino is currently a Director at Easy Dynamics where she is focused on IT Modernization for federal agencies. She began researching blockchain technologies in 2015 while she was the Technical Lead for the DHS S&T Cybersecurity Division's Identity Management RDT Program. Maria has a BS in Computer Science from UMBC and an MS in Cybersecurity from the Johns Hopkins University Whiting School of Engineering.

James P. Howard, II is a data scientist at the Johns Hopkins Applied Physics Laboratory. His work focuses on the applications of mathematics and statistics to critical challenges. He has a PhD in public policy from the University of Maryland Baltimore County.

Christina K. Pikas is a librarian and information scientist at the Johns Hopkins University Applied Physics Laboratory. She has a BS in Physics, an MLS, and a PhD from the University of Maryland. Before becoming a librarian, she served as a surface warfare officer in the United States Navy.

Anil John is the Technical Director of the DHS S&T Silicon Valley Innovation Program (SVIP), which works with innovation communities across the nation and around the world to adapt, develop and harness cutting-edge technologies and capabilities that are commercially sustainable while simultaneously meeting government needs. In this role, he identifies and conducts due diligence on technologies, companies, products and capabilities that could be adapted to meet Homeland Security operational needs and informs and educates the global innovation ecosystem including startups, accelerators, incubators, venture capital community and many others regarding the diverse Homeland Security challenges and opportunities available to them through the Program.

APPENDIX A: DEPARTMENT OF HOMELAND SECURITY (DHS) SCIENCE & TECHNOLOGY BLOCKCHAIN INITIATIVES: A SHORT HISTORY

Since December 2015, the Cybersecurity Division of the DHS S&T Directorate has invested in research and development (R&D) to understand the relevance of blockchain technology to the Homeland Security Enterprise (HSE). The initial goal of the research was to understand the security and privacy implications of blockchain to support capabilities that increase security and productivity and decrease cost and security risk. DHS S&T's current R&D investments are focused on using customer-driven POCs to identify integration points and the return on investment for blockchain deployments, and to help drive the development of the globally interoperable specifications required for blockchain to become a successful operational technology.

The current DHS S&T Identity Management blockchain portfolio includes:

- Immutability of IoT data and enterprise integration practices
- Identity and anti-spoofing of non-person entities
- Privacy enhancing population scale attribute delivery and enterprise integration
- Decentralized Key Management System Frameworks
- Verifiable Credentials Data Model
- Decentralized Identifiers
- Reducing friction in international trade through the use of blockchain technology

APPENDIX B: BASIC TECHNICAL DETAILS ON HOW A BLOCKCHAIN WORKS

The idea behind blockchain is relatively simple. There is a class of functions known as one-way hashes or crypto- graphic hashes. These functions produce an output for a given input that acts as a digital fingerprint of the input. There are many different cryptographic hashes in common use, but a cryptographic hash should have two key properties:

- One-way—It is computationally infeasible to find any input that maps to any pre-specified output.
- Collision resistant—It is computationally infeasible to find any two distinct inputs that map to the same output.

In practice, this means for a given pre-specified output, the only way to find an input that hashes to that output is to test all possible input values. The difficulty in accomplishing this is the key to blockchain security.

Creating a hash chain

A given block, called B_n, contains several data points, including ledger entries, notes, and other information. The block also contains the hash of the previous block, called B_{n-1}. This ties the two blocks together. A hash for the current block is also included. This hash ensures the data in the current and preceding blocks remain unchanged. Similarly, the hash of block B_{n-1} ensures the data of B_{n-1} and B_{n-2} are unchanged. This process works all the way back to the first block, known as the genesis block. This linked list of blocks is shown in Figure B-1.

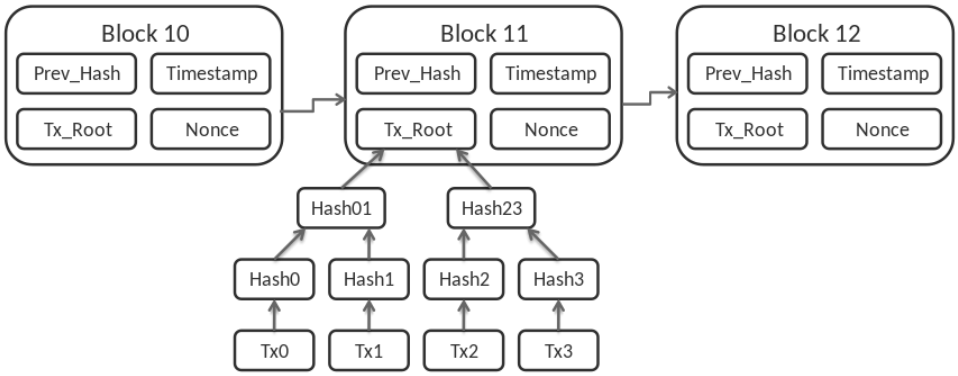

Figure B-1: Example Blockchain Transaction Blocks (via Matthäus Wander and Wikimedia Commons)

It is not enough to merely have a hash chain—one must have a way of determining if the hash chain is valid. This is done by having a mechanism for determining if an added block to a hash chain is valid. Mechanisms for doing this are commonly called "consensus mechanisms." There are two dominant types of consensus mechanisms. The first is called proof-of-work and is used in Bitcoin and most public blockchains. The second is a voting consensus, for which there are many varieties. Voting consensus mechanisms are used in all private blockchains and some public ones. Consensus mechanisms are designed to increase the profitability for those who mine or vote to protect the system instead of breaking it.

Types of Consensus

Proof of Work

In a proof-of-work system, the block contains an additional field called a nonce. A nonce is a value that is essentially random and has no meaning beyond the context of being a value [7].

In a proof-of-work system, dedicated computers called miners search for a nonce value, that when added to the block, leads to the hash of a block starting with some number of zeros. (In the case of Bitcoin, many leading zeros are necessary to make the process of finding that random nonce value so difficult that it takes, on average, 10 minutes for one of the many miners to find it.) Computers around the world try random nonce values, and if the nonce value, combined with the block's data, leads to a hash with the requisite number of leading zeros, the computer that found it writes that block to the chain, and a new block is begun. In the case of a tie, users wait until a block is added to one of the tied blocks. If two (or more) chains of blocks are proposed to be added, then the longest chain wins.

Trust in each block is then based on trust that the distributed set of miners:

- Have not colluded in sufficient numbers to prevent blocks from being mined, and
- Have not colluded in sufficient numbers to "re-write" the blockchain by having sufficient computing power to create longer chains that start earlier in the blockchain than a block they wish to replace. Because "the longest chain wins," creation of such a chain would replace earlier blocks.

This process is computationally intensive, requiring the use of specialized equipment, especially to participate in the Bitcoin network. Computationally intensive also implies energy intensive, and by the end of 2017, global Bitcoin mining required the same amount of energy used by the country of Denmark and was still growing [5].

Democratic voting consensus

A voting consensus blockchain has a number of holders of the right to vote, and they vote to decide whether or not a new block should be added to the chain. When enough votes are achieved, the block is added to the chain. If more than 50 percent of the voters are needed, then there should not be any competing chains. If there are competing chains, the length of the chain is used as a tiebreaker.

Trust in the blockchain is based on there not being collusion among a sufficient number of voters to cause the blockchain to be untrustworthy.

Proof of stake voting consensus

This type of consensus is based on the idea that those with the most stake in the blockchain being trusted

should have the greatest input on whether a block should be added. Stake can be measured in many ways (e.g., number of coins associated with the block, age, and number of transactions taken using the blockchain). Then consensus on when a block is added can be achieved in a number of ways; for example, voting weight can be apportioned by amount of stake, or people can randomly be asked to vote; however, those with a larger stake are asked to vote more often. There are many sub-varieties of proof of stake, as well, including leased proof of stake, delegated proof of stake, and proof of importance.

Regardless of which method is used to a write a block, the process begins again after that block is written. At this point, the network collects and adds all transactions submitted since the last write to the new block. Readers looking for advanced technical details should refer online to NISTIR 8202 [12].