Enhancing IoT Security and Privacy with Distributed Ledgers - A Position Paper

Paul Fremantle[1], Benjamin Aziz[1] and Tom Kirkham[2]
[1]*School of Computing, University of Portsmouth, Portsmouth, U.K.*
[2]*Science and Technology Funding Council, Harwell, U.K.*

Keywords: Internet of Things, Blockchain, Distributed Ledgers, Security, Privacy, Blockchains.

Abstract: The Internet of Things has a number of well-publicised security flaws, resulting in numerous recent attacks. In this paper we lay out a framework for looking at how distributed ledgers and Blockchain technology can be used to enhance the security, privacy and manageability of IoT devices and networks. A significant concern is the inability to process blockchains on small devices. We propose an architecture for IoT security and privacy based on blockchains that addresses this and other issues. We look at related work and propose areas of further research.

1 INTRODUCTION

The Internet of Things (IoT) consists of the set of Internet-connected devices that collect sensor data and allow control of actuators to measure and affect the physical world.

Concerns over the security and privacy of the Internet of Things reached a new high when the Mirai botnet caused massive disruption to the Internet in September 2016. In Fremantle and Scott (2015), a number of issues are identified for security of IoT. The direct cause of the Mirai attack was the use of a dictionary password attack on devices connected to the Internet that offered direct access. Other security concerns about IoT devices include significant challenges in updating devices; lack of clear and effective registration processes; lack of well defined identity models; use of IoT devices as attack vectors into more secure systems; and attacks on physical systems such as Stuxnet (Langner, 2011).

Distributed Ledger technology, which is based on the *Blockchain* (Nakamoto, 2008) construct, provides a new class of distributed technologies. These systems are famous for being the basis of cryptocurrencies, where there is a fully decentralised system with no central bank. Such systems are characterised by a distributed, secure, ledger that provides an immutable, assured record of a set of transactions. While some blockchains such as Bitcoin have restricted transaction models, others, such as *Ethereum*, support more complex transaction models. Effectively, each transaction recorded in the blockchain can support arbitrary logic which is coded in a scripting language, which is Turing complete. This ability to create new transactions through the use of scripting languages is commonly known as *Smart Contracts*. Together these three constructs — *Distributed Ledger*, *Cryptocurrency* and *Smart Contracts* — create an environment where a set of parties can share in the governance of the system, even when the identities and reputation of the parties is not known. Each party participates in guaranteeing the good behaviour of the others and of the set of transactions. The benefit of this is that it provides a fair and unbiased system whereby users can trust in the integrity of the system independent of any single parties' overall behaviour. We propose that this shared governance model has significant benefits for the security and privacy of the Internet of Things. In this work, we argue that the specific characteristics of the blockchain provide unique approaches to solving problems in the IoT.

The contributions of this paper are: a model for reasoning about how blockchains can improve privacy and security in IoT; a set of approaches for improving security and privacy of IoT with blockchains, derived from the model; and a proposed architecture for creating distributed trust in a blockchain on low-power devices.

The rest of the paper is laid out as follows. In Section 2 we propose reasons why distributed ledger technologies can help with IoT privacy and security. In Section 3 we outline a model for reasoning about

Fremantle, P., Aziz, B. and Kirkham, T.
Enhancing IoT Security and Privacy with Distributed Ledgers - A Position Paper.
DOI: 10.5220/0006353903440349
In *Proceedings of the 2nd International Conference on Internet of Things, Big Data and Security (IoTBDS 2017)*, pages 344-349
ISBN: 978-989-758-245-5
Copyright © 2017 by SCITEPRESS – Science and Technology Publications, Lda. All rights reserved

security and privacy for IoT, and use that model to identify areas where distributed ledgers can assist. In Section 4, we propose an implementation plan that addresses these issues. In particular we aim to solve the problem of how to provide trust in a blockchain from a device that lacks the power to participate in the blockchain. In Section 5 we compare this approach to related work, and outline a plan of further work in this area.

2 MOTIVATION

Because of the widespread deployment of IoT together with the use in many areas where personal information can be collected or deduced — including smart homes, health monitoring, smart cars, and fitness tracking — there are ongoing major concerns about privacy. These include: lack of informed consent for publishing data; fingerprinting of hardware; and de-anonymisation, amongst others.

The Web and Internet have been driven by heterogeneity of clients and services: any browser will work with any website; any FTP client with any FTP server. This choice enables several security and privacy benefits. Users can migrate away from insecure or privacy-leaking systems and services. It also encourages both client and service providers to produce security updates and improvements to the core protocols and to the websites and services provided. While there are problems, this has led to incremental improvements and choices. For example, users concerned about email integrity and spamming can adopt anti-spam measures such as the Domain Key Identified Mail standard (Crocker et al., 2011). Users concerned with confidentiality can encrypt email. Users concerned about insecure messaging systems can choose more secure instant messaging systems such as *Signal*[1] and *Threema*[2]. In contrast, the Internet of Things is fundamentally controlled by the manufacturers of devices, and does not offer choice for users to migrate to more secure services. IoT devices are hard-coded with firmware that typically connects to a single service. In many cases the only true privacy control a user has is to completely disable a device. Even when systems use standard protocols, they may not document the usage.

As discussed above, blockchains create a distributed ledger that creates a shared governance. Blockchains rely on proof that the parties are behaving in a consistent and correct manner, since the parties are not known and are assumed to be inherently untrustworthy. Bitcoin, the most famous blockchain-based system, relies on a concept called *Proof of Work* to ensure that parties are behaving properly. Other blockchain systems rely on differing proofs, including *proof of stake*, *proof of storage*, or combinations of different proof types. The proof is used to guarantee the immutability of an *block* and each block guarantees the previous blocks (typically using a construct known as a *Merkle Tree*) — creating a chain of blocks known as a blockchain. Because the proof requires some form of value (whether it is ownership of a cryptocurrency, expended work, storage of data, etc), the system can provide rewards for those who participate. This is the basis of cryptocurrencies. For example, Bitcoin issues new coins to the participating system that first correctly produces the work that verifies the latest block.

Therefore, the distributed ledger offers the promise of creating an environment for IoT networks where there can be trust, anonymity, and effective contracts between parties without any single vendor being in charge, and without requiring any party to trusted above any other.

However, there remains a major concern. The processing, memory and code requirements of blockchains makes them inappropriate to run on IoT devices and gateways. For example, the current Bitcoin database is around 80Gb and it takes at least 512MB of RAM and a 1Ghz processor to participate [3]. The *Simple Payment Verification* (SPV) model in Bitcoin supports lighterweight clients that can connect to a random sample of servers, but this is susceptible to Sybil attacks. Recent work in Frey et al. (2016) has made it feasible to participate in the Bitcoin network with a smartphone, but to have an effective approach for IoT the system must support common cheap IoT devices. For example, the ESP8266 device is a common device target that offers 1Mb of program memory and 80Kb of variable memory. Even with the improvements from Frey et al, this is still insufficient even to validate the Bitcoin system. As a result, we consider this a significant issue that needs addressing before the vision of using blockchains as the basis of IoT security and privacy can be achieved. To solve this problem, we propose an avenue of research that can provide trust between IoT devices and blockchains without requiring the device to actively participate in the blockchain.

[1] https://whispersystems.org/

[2] https://threema.ch/en

[3] https://bitcoin.org/en/bitcoin-core/features/requirements

3 MODEL

Spiekermann and Cranor (2009) offer a model for looking at user privacy. In their model, they identify three spheres: the *User Sphere*, the *Joint Sphere* and the *Recipient Sphere*. The User Sphere is completely in the control of the user (e.g. a laptop). The Joint Sphere refers to areas that may seem to be in the user's control, but may have some significant control by a third-party. For example, a cloud email account may seem like the user can delete emails, but the cloud provider may in fact back these up and keep a copy. Finally, once data has been transferred to a third-party, it is assumed to be in the Recipient Sphere, where the only controls are legal and contractual.

In the model, a device that offers the user full control is firmly in the *User Sphere*. However, we would argue that many current devices are actually in the *Joint Sphere*. This is where the device appears to be in the control of the user but in fact is in the control of a third-party. To give an example, the Google Nest device offers users the opportunity to apply smart heating controls to their house. While a number of user-centred controls give the user the impression that it is in the User Sphere, there are two key reasons to counter this: firstly, the data logged by the device is extensive and cannot be controlled by the user; secondly, the device auto-updates itself based on commands from Google rather than based on user input.

Using this model, we can propose clear approaches that strengthen each of the privacy and security controls available in each sphere. Figure 1 provides an overview of this model.

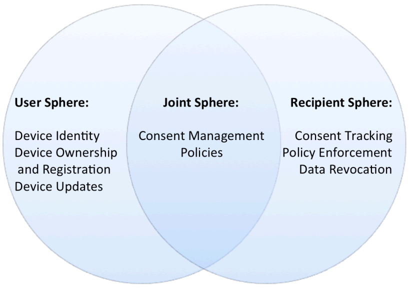

Figure 1: IoT Blockchain applied to Privacy Model.

3.1 User Sphere

Moving privacy and security controls back to the users inherently strengthens the User Sphere and provides greater choice, thereby allowing more secure approaches to flourish. We have identified a number of controls which can be managed by a distributed ledger.

Devices need to have secure identities (Fremantle et al., 2014), and currently these are either not provided, or provided by the device manufacturer. Systems such as NameID [4] provide a model whereby secure Web identities can be provided with no central trusted system. Instead a distributed ledger (in this case Namecoin) provides a neutral shared system. This approach can be extended to IoT to issue each device with a secure identity based on a distributed ledger.

A second, related issue, is the ownership of devices. The Mirai botnet spread because dictionary attacks allowed attackers to take ownership of devices. Some systems offer models of taking ownership securely (e.g. Bluetooth, NFC). In (Fremantle and Aziz, 2016), a QR code is used in conjunction with a Web-based system. We propose that the cryptocurrency model of blockchains provides a safer and clearer model of ownership. Distributed ledgers implicitly support ownership: a Bitcoin transaction is the transfer of some bitcoins from own owner to another. Similarly, a blockchain for IoT could support the transfer of ownership of a device from the manufacturer to the user, and from one user to another. The security of this is therefore based on the security of the blockchain.

A third issue within the User Sphere is updating the device firmware. A number of attacks have originated in lack of updates. One issue is that device manufacturers are incentivised to create new products but not to update old products. In (Tindall, 2015), a model is proposed whereby devices can pay for updates in a cryptocurrency. In addition, blockchains could be used as the basis for building trust in updates and also in validating updates.

3.2 Joint Sphere

Recall that the Joint Sphere is the parts of the system where the user has some form of control over their data and systems, but the provider also shares control. For example, a health-monitoring device may upload data to an Internet-based system and then may offer users controls on how they share data. A major change in legislation around this is the European Union's General Data Protection Regulation (GDPR) which requires much stronger consent controls. Many systems offer forms of user consent for sharing data with third parties, but these lack significant requirements. For example, many users are not aware of how to revoke consent. Similarly, there is no clear place a user can identify all the consents they have approved

[4] http://nameid.org

across different devices. Consent is not just about privacy. IoT devices often include actuators that can act based on *Commands*, and the security of a device includes ensuring that only authorised systems can issue commands to devices.

We propose that consent be treated as a contract in a distributed ledger system. This has several benefits. Firstly, this creates a clear ledger of all issued consents. Secondly, this creates a clear contractual obligation on the recipient of the consent to behave according to the contract, which could (for example) specify penalties for breaking the contract. Finally, this model creates a clear opportunity for revocation of consent, giving a single place to do so.

A related area is that of policies. In this meaning a policy is a computer-readable expression of rights and obligations. For example, a consent approval may refer to a policy: the user might approve sharing of data to a website based on the fact that the website promises not to share the data to any other body. Languages such as XACML (Godik et al., 2002) allow complex access control policies to be encoded in XML or JSON. We propose that a distributed ledger could be used to store and validate policies in an immutable model that could be then used by IoT networks and consent models. In addition, the policies could themselves refer to transactions in a distributed ledger. We will describe this in the Recipient Sphere.

3.3 Recipient Sphere

The Recipient Sphere is the area where the user's data is now out of their control. Ultimately, the user must rely on legislation or legal contracts in order to maintain control of this data. However, we propose that distributed ledgers can offer a significant benefit in this space. For example, suppose that a user shares data with a third-party using a smart contract to enable consent. The contract refers to a policy that says the third-party can only share that data with others (*on-bound sharing*) if they inform the user of that on-bound sharing. This policy could point to a new smart contract that must be executed at the point of on-bound sharing. By executing this second contract, the on-bound sharing is logged immutably into a distributed ledger. Of course, it is hard to police this recipient sphere: it is possible that the third-party website will share data without executing the contract. There are two defences against this. The first is that they are explicitly breaking a contract and penalties can occur. In fact, we can envision that these penalties could be implemented directly in the ledger by smart contracts. If the user can prove data was shared illicitly, the payment would be automatic.

Spotting these illicit data shares can possibly be done using a concept of a *Trap Street*. This is the habit that map-makers have of including incorrect data to see if others copy it. Similarly, IoT devices could deliberately share incorrect data to specific parties to see if it leaks out without a ledger entry to record it.

A similar capability that could be implemented using distributed ledgers and smart contracts is that of *Data revocation*. This would be a situation where a smart contract is used to trigger the secure deletion of specific data.

In summary, we have seen a number of ways where distributed ledgers can enhance privacy and security of IoT. We now look at a proposed implementation.

4 PROPOSED IMPLEMENTATION

The biggest concern we have regarding the implementability is the inability of small devices to participate in blockchains. Many blockchains provide lighter-weight models of validation such as the Bitcoin SPV[5] and the Ethereum Light Client Protocol[6]. However, even these may require more processing than an IoT device can provide, and this requirement may also increase in the future with the growth in the blockchain ledger.

There already exists a concept in the blockchain system of an *Oracle*. An oracle is a system that reports on the world to the blockchain in a reliable fashion. For example, a smart contract may require payment when a certain condition is met, and the oracle is used to report to the blockchain when that condition exists. We propose that the IoT and Blockchain require the exact opposite - a trusted intermediary that reports on the state of the blockchain on behalf of the IoT device. Such an entity, which we call a *Pythia*, could interact with the blockchain on behalf of IoT devices and do so in a trusted fashion. Therefore, it would act both as an oracle to the device, as well as an oracle to the blockchain.

In order to allow the Pythia to be trusted, we propose creating an Open Source codebase that would be run using *Intel Secure Guard eXtensions* (SGX) (Costan and Devadas, 2016). SGX provides a secure enclave within modern Intel processors that provides two key benefits that are required by the Pythia model. Firstly, the code is sealed and pro-

[5] https://en.bitcoin.it/wiki/Thin_Client_Security

[6] https://github.com/ethereum/wiki/wiki/Light-client-protocol

tected from attack from other code running in the same processor. Secondly, the code can be validated using *Remote Attestation*, which allows code external to the processor to verify the specific codebase running within the enclave. This verification still requires some processing power as it is based on Diffie-Hellman Key Exchange (DHKE). Many devices do perform aspects of Public Key Cryptography, but we have not yet assessed the ability of typical devices to perform DHKE, and that is one of our ongoing goals. The remote attestation is then used to validate that the Pythia is indeed running a specific version of the open source codebase. In turn, this allows the device to trust that the Pythia will act as intended with the blockchain. This model is shown in Figure 2.

The Pythia model is that it offers a set of APIs to IoT devices, and that these APIs offer capabilities around device identity, ownership, data sharing consent, etc. A number of these APIs are already in place in systems such as OAuthing (Fremantle and Aziz, 2016), which utilises the OAuth2 APIs to provision identities, consent scopes and ownership tokens to devices. We believe that this model can be extended to interface with blockchain smart contracts. One challenge is that code running within an SGX enclave only runs efficiently when it fits into 128Mb, which may prove to be a challenge when building the Pythia.

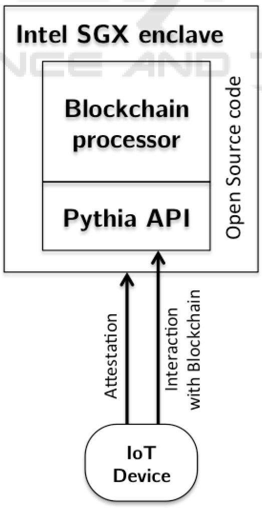

Figure 2: Proposed Pythia Model.

Given this infrastructure, we propose building a set of smart contracts that implement the flows identified in Section 3. These are:

- Secure Device Identity
- Device Ownership
- Device Updates
- Data sharing and command consent
- Consent logging
- Policy sharing and validation
- Data revocation

We intend to use the Ethereum blockchain as the basis for this work.

5 RELATED WORK AND CONCLUSIONS

In Christidis and Devetsikiotis (2016) there is considerable discussion of using blockchains together with IoT, but this does not focus specifically on the privacy and security challenges of IoT. In Milutinovic et al. (2016), an innovative blockchain is proposed based on running within an SGX enclave. In Fremantle and Aziz (2016) there is a framework for device identity, registration and consent that is independent of vendors. However, this still relies on trust in the central party, which can be obviated with the use of blockchains and smart contracts. Town Crier (Zhang et al., 2016) proposes an approach where SGX extensions are used to feed data into a blockchain in a trusted manner, acting as an oracle. This does not propose the opposite flow where the SGX enclave provides verified data from the blockchain as in the proposed Pythia.

5.1 Conclusions and Further Work

At the moment, this vision of a blockchain-based IoT is just a preliminary proposal. Clearly, implementing such a blockchain based on existing distributed ledgers that enable Smart Contracts such as HyperLedger or Ethereum is possible. Implementing an SGX-based blockchain client to provide trusted data from the blockchain is also clearly possible, although given the limitations on SGX memory, this may need significant tuning to be effective and efficient. At this point, we have not evaluated the ability of a device to perform remote attestation of an SGX enclave, nor the complexities of the key distribution requirements to make this possible. We see this as the first step in the ongoing research plan, together with evaluating the possibilities of running an Ethereum node in an SGX enclave.

In this paper we have proposed using a distributed ledger to provide a shared governance model for IoT devices, networks and cloud systems. Using Spiekermann and Cranor's Three Layer Privacy model, we have outlined an approach for evaluating the use of blockchains in IoT. We have identified a number of key areas where blockchains could be used to improve privacy and security. We have proposed an system that we call a Pythia whereby IoT devices can trust the transactions of a blockchain without requiring the memory and processor capabilities to actively participate in the blockchain.

REFERENCES

Christidis, K. and Devetsikiotis, M. (2016). Blockchains and smart contracts for the internet of things. *IEEE Access*, 4:2292–2303.

Costan, V. and Devadas, S. (2016). Intel sgx explained. *IACR Cryptology ePrint Archive*, 2016:86.

Crocker, D., Hansen, T., and Kucherawy, M. (2011). Domainkeys identified mail (dkim) signatures. Technical report.

Fremantle, P. and Aziz, B. (2016). Oauthing: privacy-enhancing federation for the internet of things.

Fremantle, P., Aziz, B., Scott, P., and Kopecky, J. (2014). Federated Identity and Access Management for the Internet of Things. In *3rd International Workshop on the Secure IoT*.

Fremantle, P. and Scott, P. (2015). A security survey of middleware for the internet of things. *PeerJ PrePrints*, 3:e1241v1.

Frey, D., Makkes, M. X., Roman, P.-L., Taïani, F., and Voulgaris, S. (2016). Bringing secure bitcoin transactions to your smartphone. In *Proceedings of the 15th International Workshop on Adaptive and Reflective Middleware*, page 3. ACM.

Godik, S., Anderson, A., Parducci, B., Humenn, P., and Vajjhala, S. (2002). Oasis extensible access control 2 markup language (xacml) 3. Technical report, Tech. rep., OASIS.

Langner, R. (2011). Stuxnet: Dissecting a cyberwarfare weapon. *IEEE Security & Privacy*, 9(3):49–51.

Milutinovic, M., He, W., Wu, H., and Kanwal, M. (2016). Proof of luck: an efficient blockchain consensus protocol. In *Proceedings of the 1st Workshop on System Software for Trusted Execution*, page 2. ACM.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

Spiekermann, S. and Cranor, L. F. (2009). Engineering privacy. *IEEE Transactions on software engineering*, 35(1):67–82.

Tindall, K. (2015). How bitcoin might fix the broken internet of things medium. https://freo.me/2jNZRBm. (Accessed on 01/20/2017).

Zhang, F., Cecchetti, E., Croman, K., Juels, A., and Shi, E. (2016). Town crier: An authenticated data feed for smart contracts. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 270–282. ACM.

SCITEPRESS
SCIENCE AND TECHNOLOGY PUBLICATIONS