



UniquID

A Quest to Reconcile Identity Access Management and the IoT

Giaretta, Alberto; Pepe, Stefano; Dragoni, Nicola

Published in:

Proceedings of 51st International Conference on Methods and Tools

Link to article, DOI:

[10.1007/978-3-030-29852-4_20](https://doi.org/10.1007/978-3-030-29852-4_20)

Publication date:

2019

Document Version

Early version, also known as pre-print

[Link back to DTU Orbit](#)

Citation (APA):

Giaretta, A., Pepe, S., & Dragoni, N. (2019). UniquID: A Quest to Reconcile Identity Access Management and the IoT. In M. Mazzara, B. Meyer, J-M. Bruel, & A. Petrenko (Eds.), *Proceedings of 51st International Conference on Methods and Tools* (pp. 237-251). Springer. https://doi.org/10.1007/978-3-030-29852-4_20

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

UniquID: A Quest to Reconcile Identity Access Management and the Internet of Things

Alberto Giaretta¹, Stefano Pepe², and Nicola Dragoni^{1,3}

¹ Centre for Applied Autonomous Sensor Systems (AASS), Örebro University, Örebro, Sweden. alberto.giaretta@oru.se

² UniquID Inc., San Francisco, USA. pepe@uniquid.com

³ DTU Compute, Technical University of Denmark, Kongens Lyngby, Denmark
ndra@dtu.dk

Abstract. The Internet of Things (IoT) has caused a revolutionary paradigm shift in computer networking. After decades of human-centered routines, where devices were merely tools that enabled human beings to authenticate themselves and perform activities, we are now dealing with a device-centered paradigm: the devices themselves are actors, not just tools for people. Conventional identity access management (IAM) frameworks were not designed to handle the challenges of IoT. Trying to use traditional IAM systems to reconcile heterogeneous devices and complex federations of online services (e.g., IoT sensors and cloud computing solutions) adds a cumbersome architectural layer that can become hard to maintain and act as a single point of failure. In this paper, we propose UniquID, a blockchain-based solution that overcomes the need for centralized IAM architectures while providing scalability and robustness. We also present the experimental results of a proof-of-concept UniquID enrolment network, and we discuss two different use-cases that show the considerable value of a blockchain-based IAM.

Keywords: IAM · Identity management systems · Blockchain · Internet of Things · IoT · Machine-to-machine · M2M

1 Introduction

Information Technology (IT) has radically changed throughout history. In just a few decades, computers evolved from bulky, standalone machines that filled rooms to small and powerful devices capable of gathering information from other devices over the Internet. This explosion of capabilities has led to a concomitant expansion of complexity in a variety of areas, including identity access management (IAM).

Conventional IAM systems are essential for traditional local networks and businesses, but not well-suited for large networks of complex, highly distributed devices, such as the combination of Internet of Things (IoT) and machine to machine (M2M) communication. The main reason is that IAM systems were designed for human beings, not devices. Until a decade ago, a typical scenario

involved a known number of terminals and a comparable number of human users with well-defined roles. Accounts were issued for each person and access rights stored on a central server, making access control (AC) relatively easy to manage. Today, IT is much more complex than it once was.

The Internet of Things is causing the old user-centered paradigm to shift toward a device-centered one. Previously, accounts were tied to human beings and devices were just a means to the end of accomplishing all of the tasks those humans performed. Today, devices are increasingly becoming the actors themselves, with the tasks narrowed down from general purpose functions to very specific operations. With this in mind, traditional IAM solutions are expensive, make maintaining cross-domain consistency challenging, and represent a critical single point of failure in organizations due to their centralized nature.

In the 1990s, PKI was considered one of the key technologies for dealing with security issues and enabling trust among parties [5]. But adoption of the technology never took off as envisioned, due to critical issues such as privacy and liability concerns, management complexity, and high costs. PKI technology is compartmentalized by design, as shown in Figure 1, and therefore irremediably ill-suited for devices that need cross-domain interactions or technical solutions that rely upon the federation of independent services. PKI lacks the flexibility needed for IoT and M2M, whereas certificate-based authentication is well-suited for these uses.




Fig. 1: PKIs inherently introduce compartmentalization, which increases the complexity of identity reconciliation over cross-domain federated architectures. In this example, the identity of a single Internet-enabled car has to be separately managed by several different PKIs.

This paper describes UniquID⁴, a solution based on an infrastructure that takes advantage of certificates and blockchain technology to overcome the difficulties in reconciling IoT credentials and cross-domain IAM. The goals of this paper include the following:

- Describing a cheaper and simpler alternative to traditional IAM systems.
- Illustrating the implementation of cross-domain identities for IoT devices to circumvent account reconciliation.

⁴ <http://www.uniquid.com>

- Showing how the proposed design removes single points of failure from the trust structure.
- Demonstrating direct peer-to-peer (P2P) authentication and authorization among IoT devices.
- Showing how an IoT device, empowered to locally confirm a smart transaction, deals with partitioning issues as defined by the CAP Theorem.

1.1 Paper Outline

The paper is organized as follows. Section 2 briefly presents the main players in the blockchain-based IAM landscape, how UniquID addresses some of the distributed IoT challenges, and how it is linked to the CAP Theorem. Section 3 describes the general concepts behind UniquID, and Section 4 shows the experimental evaluation of an enrolment proof-of-concept. Section 5 describes two example case studies, and Section 6 lays out the paper conclusions and discusses future work.

2 Related Work

In the past, blockchain has already been advocated for more secure IAM systems. For example, Kshetri [11] suggests that blockchain can help in strengthen the IoT in different ways, such as preventing DDoS and IP spoofing attacks. Furthermore, Gartner estimates that by 2020 the IoT will require up to 1000 times the 2016 network capacity [13]. This entails that the centralized IAM paradigm might not scale enough to tolerate such requirements, and this is where the decentralized IAM can provide both more security and scalability.

In Roman et al. [17], the authors lay out the main challenges of distributed IoT. They argue that identity and authentication are primary concerns due to the inherent dynamism introduced by device mobility, unstable connections, and related problems. Throughout this paper, we show how UniquID provides direct identification and authentication, which in turn enables efficient M2M resource negotiation. Another issue raised is security, which UniquID ensures through asymmetric encryption, adding symmetric cipher-based encryption for larger data streams. Furthermore, depending on the chosen backbone, privacy can be partially or even totally lost. In cases where privacy is needed, it can be enforced through non-interactive zero-knowledge proofs such as zk-SNARKs [18] and Bulletproofs [4], but this involves considerable overhead that may not be appropriate for resource-constrained devices.

Several parallel projects have been proposed to address the problem of identity management over blockchain, yet most of them do not address the access management part. For example, IBM Hyperledger Indy aims to provide an SDK solution to manage identities over distributed ledgers [7]. Even though UniquID similarly provides an SDK, our proposal addresses both the identity and access management parts of the equation, whereas Indy addresses only identity.

Scholars proposed some solutions as well. Le and Mutka propose CapChain [12], a blockchain-based access control framework that enables IAM on public blockchains, ensuring at the same time privacy; in particular, the authors built a proof-of-concept over Monero source code and ran processing time local benchmarks. Ouaddah et al. [16] propose a blockchain-based access control named FairAccess, and implemented a proof-of-concept which uses Bitcoin OP_RETURN as a storage field. Finally, Novo [15] proposes a blockchain-based access management that utilizes a management hub (e.g., an edge node) as a midpoint between IoT devices and the blockchain.

Even though all the aforementioned papers prove that blockchain-based IAM is possible, they lack a thorough discussion on how the underlying blockchain choice can deeply affect operational costs, data immutability, immediacy, and efficiency. Moreover, none of such works envisage a way to verify a policy without Internet connection, failing to take full advantage of the M2M paradigm. In particular, no one highlights the strong relationship between Brewer’s CAP Theorem and IAM solutions, nor motivations and implications of choosing consistency over availability (or vice-versa) in their proposals.

2.1 Consistency, Availability, and Partition Tolerance (CAP)

According to Brewer’s CAP Theorem [3], one cannot ensure all three of the following in a system at any given time: *consistency (C)*, *availability (A)* and *partition tolerance (P)*. CAP is often misunderstood, with people thinking that a distributed system is always unable to assure all three requirements. In reality, the choice is only between consistency and availability when a partition or a failure occurs; under normal circumstances, all three can be assured simultaneously. Another misunderstanding is about consistent-available (CA) systems, which are simply not possible in a distributed scenario. According to CAP, a system could be designed to be CA, but it would require a network that ensures no packet is ever dropped at any moment in time. For a fixed partition tolerance requirement, the only real choice is between consistency and availability.

Consistency, which is a property related to read operation, can be either strong or eventual. In a consistent-partition-tolerant (CP) scenario typical of an RDBMS, the system ensures that every commit to the database is propagated and kept consistent throughout all database replicas, so that every read operation returns the most recently updated result. In an available-partition-tolerant (AP) scenario typical of NoSQL, the read operation does not ensure that the user receives the most up-to-date result. But even though AP sounds problematic, consistency is eventually achieved, and this approach is common in many non-critical applications due to its strong support for partition tolerance and availability.

It is important to keep in mind that partitions and latency are strongly related, to the point that we can define a partition as a function of latency: the developer can decide the latency threshold beyond which partitioning occurs. By indefinitely retrying communications, one is essentially choosing a CP solution over an AP one, whereas replying right away to a user request means choosing

AP over CP, as the data might be stale (inconsistent with the current state). The idea of selecting either CP or AP is a false dichotomy, however, as tuning the time threshold makes it possible to switch from a CP solution to an AP one after a chosen amount of time.

UniquID aims to provide similar flexibility depending on the application domain. This can be easily done by deciding for how long an IoT device will try to download an up-to-date smart contract from the blockchain (i.e., CP behavior) before conceding the resources to the client based on the locally stored smart contract (i.e., AP behavior). In this particular case, the IoT device is a UniquID node that can locally read a smart contract even though it does not store the whole blockchain. The ability to locally read a smart contract means that any network failure happening between the node and the rest of the UniquID network is a partition. Therefore, the CAP theorem applies and a thoughtful decision must be made between consistency and availability.

3 UniquID Overall Architecture

Figure 2 shows how different networks interact with each other in a traditional IAM infrastructure. An IoT device in Network A must go through the IAM platform to authenticate on a different IoT network, such as Network B. In the event of network unavailability, a device cannot authenticate, showing that this solution does not provide good availability.

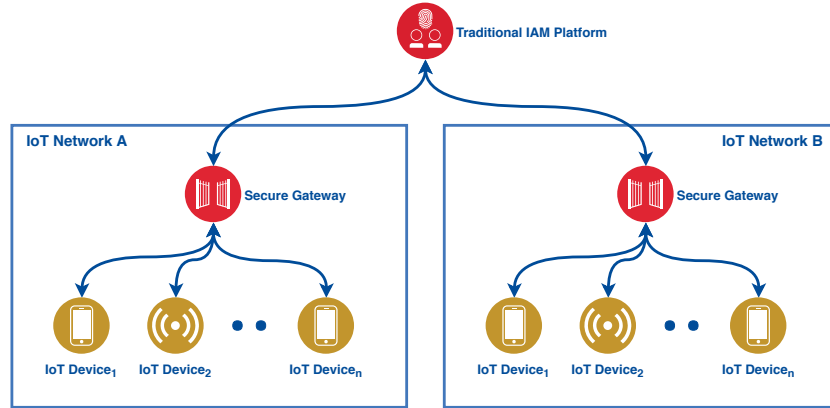


Fig. 2: Traditional IAM systems impose a hierarchical structure over managed devices. This introduces potential single points of failure in the architecture and strongly impedes direct M2M communications.

Furthermore, as shown in Figure 3, this kind of architecture involves a considerable number of message exchanges. This overhead decreases the overall responsiveness of the system and means that it cannot ensure availability in the

event of network partitioning. Indeed, we can classify a traditional IAM architecture as CP-compliant without any possibility of choosing an AP solution. This system was acceptable in the old days of fixed and stable architectures, but IoT networks require greater flexibility and interoperability.

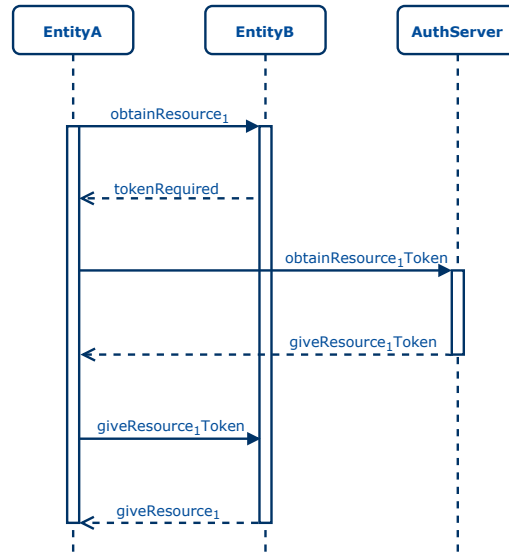


Fig. 3: With traditional IAM systems, resource negotiation between a requestor (EntityA) and a provider (EntityB) involves an additional exchange between EntityA and an authentication server (AuthServer).

A primary goal of UniquID is to replace traditional IAM structures with a less expensive architecture that is more flexible and easier to manage. As shown in Figure 4, the system uses an infrastructure where devices can directly authenticate each other without a trusted third party IAM platform, as envisioned by the PGP Web of Trust [6,8].

Figure 5 shows the UniquID workflow, illustrating two main differences from the IAM setup depicted in Figure 3. First of all, fewer messages are required than in a traditional approach, which entails that UniquID provides better performances than a centralized IAM, assuming that AuthServer and Blockchain have comparable response times. Second, and more important, the entities use smart contracts stored in a tamper-proof public blockchain instead of tokens. This not only ensures higher security but also enables Entity B to store the contract so it can still authenticate Entity A in the event of a network failure.

This might pose some security issues. Suppose a malicious person M steals Entity A, the laptop of an important CEO. The network administrator immediately revokes all Entity A permissions by issuing a blockchain transaction. If Entity B is instructed to allow authentication without double-checking the




Fig. 4: UniquID removes the authentication server, which has two important implications. First, the potential single point of failure disappears from the network. Second, scalability greatly improves, as M2M communication is possible and a central authority is no longer essential for negotiation.

currentness of the permission, M might try to gain as much time as possible to cause damage by initiating a network fault that isolates Entity B from the blockchain. This would enable M to authenticate despite the staleness of the Entity A permissions stored within Entity B.

At the same time, the ability to verify a smart contract offline might be invaluable for some other applications. Take as an example a public transportation service. Assume that user U pays for a one-hour ride ticket that allows him to take any city bus line he likes. Assume that the city buses pull this ticket from the blockchain, and temporarily store it locally. After U has done all his errands, he goes to the bus stop and tries to get on the bus, but a network fault occurs. If the buses have been configured to authenticate him without double-checking the current status of the ticket on the blockchain, U will be able to board the bus. If, on the other hand, the buses need to perform a double-check, U will be left waiting until the network fault is resolved, which might take a long time. Again, a malicious user M might isolate the bus he is boarding and extend his ticket for some time. However, this is a small economic risk that the company would be willing to take given the more serious consequences that a network fault would have on all of its transportation infrastructure.

Each scenario has different requirements. One of the strengths of UniquID is that the end users can decide whether a CP system or an AP system is best suited to their needs.

3.1 Imprinting Ceremony

With the goal of excluding any kind of PKI, we envision an approach similar to the Web of Trust paradigm [6,8] for the node initialization phase. In UniquID, the imprinting ceremony resembles a PGP key-signing party. To be enrolled in UniquID, every node must be initialized, which happens in different phases. To minimize the chances of man-in-the-middle (MitM) attacks, the following steps should take place as soon as possible in the production chain.




Fig. 5: UniquID communication is more straightforward than in traditional IAM. EntityA can negotiate for resources directly with EntityB, which can interrogate the blockchain to confirm that the request is valid. But this interrogation is not mandatory, since EntityB can store hash trees and locally confirm certificate validity thanks to the Merkle tree data structure.

First of all, each node generates its own public key, which is directly (i.e., not through the Internet) exchanged with other devices in the same local network. After this phase is complete, the identities are stored in a blockchain through the imprinting process. An *imprinter* is a node designated to collect all the new identities, and forward them to the blockchain of choice.

In details, the imprinter is appointed to perform some critical tasks. First, the imprinter generates a special smart contract, the imprinting contract (IC), which links the device to the public key of its administrator. Once the generation is done, the imprinter collects the ICs and announces them to the blockchain, in order to ensure data immutability. From this moment on, these enrolled devices are able to interact with the UniquID infrastructure. If the device is transferred to another entity (e.g., sold to a customer), the administrator (e.g., the manufacturer) signs a contract and transfers the administrative rights by replacing its public key with that of the new owner. Ideally, due to the critical role of the imprinter, the manufacturer deploys a number of different imprinting nodes, to both speed up the imprinting process, and to avoid single point of failures in the architecture.

In Figure 6, we show the main tasks performed during the imprinting phase. After that a node has created its identity (i.e., the public key), it announces it to the imprinter which is appointed to create the IC and to communicate it to the blockchain. At that point, the IC will go under the underlying blockchain processes, until it will be included in a new block and become almost immutable.




Fig. 6: Sequence diagram that depicts a standard imprinting ceremony.

4 Experimental Evaluation

In this section, we present the experimental results of a UniquID enrolment instance, executed on a cloud service. Even though we measured the performance of the identity generation phase, as well as the imprinting phase, our goal is not to assess the performance of such instance. Indeed, the system scalability is tightly tied to the utilized resources, and the small setup we used for our experiments cannot represent a fully operating UniquID network. The aim of these experiments is to show that our solution is feasible, and that it successfully stores immutable identities on a public blockchain. On a larger scale, this would be enough to make PKIs, passwords, and certificates unnecessary.

To perform our experiments, we created 7 parallel clients, designated to generate 1000 virtual IoT identities and communicate such identities to 1 imprinting node. Every identity is announced as soon as it is created, and the receiving imprinting node is appointed both to create the IC transactions and forward them to the Litecoin Testnet blockchain. Imprinter and clients ran over AWS T2.Micro instances, burstable performance instances equipped with 1 Gb of RAM. As a communication protocol between the imprinter and the clients, we used MQTT (Message Queue Telemetry Transport), an ISO standard (ISO/IEC PRF 20922) publish-subscribe-based messaging protocol [10], designed for lightweight communications.

As aforementioned, we chose Litecoin as a storing public blockchain, which ensures one mined block every 2.5 minutes. Considering that a Litecoin block is 1 Mb and that a UniquID transaction is 400 byte, this design choice entails a theoretical upper bound of 2500 enrolled devices per 2.5 minutes, or 1000 devices per minute. More in general, we can define the theoretical upper bound of enrolments per minute with the following equation:

$$\frac{\textit{Block Size (bytes)}}{400 \textit{ (bytes)} \cdot \textit{Average Mining Time (m)}}, \quad (1)$$

where 400 bytes are the size of a UniquID transaction, and the other parameters depend on the underlying blockchain.

4.1 Identity Generation

In the first evaluation phase, we measured the time required to generate an identity on a virtual client. As shown in Figure 7a, on average it took 6.61 ± 0.03 milliseconds to generate an identity.

As stated before, in our experiments we delegated the generation task to the virtual clients. In an ideal scenario, the IoT devices would create and announce their identity themselves, therefore we performed an additional experiment, in order to evaluate the generation time on a low power device. The device we used was equipped with an *arm926ejste* CPU, 256 MB RAM, and a *mlinux 3.3.6* OS. On this IoT device, the generation process took 627 milliseconds, 44.7% of CPU and 0.4% of RAM. The footprint left on the device storage was 6.70 MB.

While the generation time of 627 milliseconds is substantially higher than the one obtained with a virtual client (6.61 ± 0.03 milliseconds, as showed in Figure 7a), generating an identity well under 1 second is still an excellent result. Moreover, the resources used by the whole process are reasonably low and the resulting footprint is small enough to fit into any modern IoT device.


4.2 Imprinting

As a second part of our experiment, we analysed and measured the necessary time to imprint new identities on a public blockchain. As aforementioned, our experimental setup is composed by 6 clients that created 1000 identities, and announced them to a single imprinter node. Again, the imprinter has the role to create an IC transaction for each identity and submit all the transactions to the blockchain. In Figure 7b, we show that the average time to imprint one identity is 9.06 ± 0.35 minutes.

Our experiments clearly show that the imprinting took a considerably longer time than the generation of the identities. The former took minutes, whereas the latter took milliseconds.

4.3 Summing Up: Enrolment

Our experiments show that, in total, it took around 4.5 hours to automatically enrol (i.e., create, announce, and imprint) 1000 identities, without any human intervention. The reason why this practical result is considerably under the theoretical upper bound of 1000 enrolments per minute, is that our virtual clients on Amazon AWS could open no more than 5 concurrent sockets, per each. This resulted in a bottleneck, where already generated identities were announced with



(a) Time is in *milliseconds*, error bar shows the Standard Error (SE). (b) Time is in *minutes*, error bar shows the Standard Error (SE).

Fig. 7: Average times for generating and imprinting identities, respectively.

considerable delays. Figure 8 shows the magnitude of this bottleneck: on average, it took 115.86 ± 2.19 minutes to announce an identity to the imprinter, which is a significant amount of time compared to the time for identity generation (6.61 ± 0.03 milliseconds) and imprinting (9.06 ± 0.35 minutes).




Fig. 8: Average time to announce one identity. Announce time is in minutes, error bar shows the Standard Error (SE).

Internal Experiment. In order to further investigate this communication bottleneck, we ran some internal tests with exactly the same network setup. These tests showed that we can easily run up to 50 concurrent threads per MQTT client, overcoming the communication bottleneck and increasing our capacity to 400 enrolments per minute. Still, at the present stage UniquID does not enable us to get near the theoretical upper bound, but this is purely due to software limitations that are currently addressed.

To sum up, our experiments successfully show that it is possible to automatically enrol IoT devices in a blockchain-based IAM. In practice, this means that enrolled devices no longer need to rely upon PKIs, centralized CA authorities, certificates, nor passwords. Moreover, our experiments lead to immutable iden-

tities, publicly verifiable by anyone on the Litecoin Testnet ⁵. Once showed that our approach works in practice, the next step will be to optimise the performance to get as close as possible to the theoretical upper bound.

5 Case Studies

In this section, we present a couple of case studies to highlight UniquID strengths and variety of applications.

5.1 Smart Vehicle

Let us assume a connected vehicle with 3G cellular capability and a mobile app designed to unlock the doors or start the engine. In a typical scenario, the smartphone app connects to a RESTful cloud service responsible for authenticating the user and forwarding the command to the vehicle, leveraging an encrypted session pushed through an available 3G connection. This architecture presents three main issues:

- In the absence of 3G signal coverage, the cloud service cannot perform any remote command and control of the vehicle, leaving the user locked out [2].
- In large-scale deployments, the cloud service becomes the main bottleneck of the system, introducing latency and potential downtime during peak hours.
- Every vehicle-side application is exposed to Internet connectivity and thus must be maintained against zero-day vulnerabilities and malware and ransomware attacks [9].

UniquID allows the two endpoints (i.e., the smartphone and vehicle) to independently synchronize with the ledger, leveraging the hashcash PoW cost-function [1] to verify the integrity of a new block. Moreover, transactions stored as Merkle tree leaves allow efficient verification of the veracity of a new authorization [14]. Once this information is locally stored on both ends, handshake and authorization are performed through low-energy and proximity protocols (e.g., BLE) without relying on Internet connectivity or remote command-and-control services, as previously shown in Figure 5. Figure 9 shows the resulting architecture.

5.2 Industrial Sensor Network

Consider an industrial Internet application where a large network of remotely installed, battery-operated sensors collect time series data in a low-connectivity environment. These sensors do not have cellular signal coverage, and satellite uplink is not a feasible option. The data from these sensors is periodically "harvested" by human operators through a portable device, such as a laptop or rugged tablet, using an ad-hoc Wi-Fi connection to the sensors. This architecture presents two main challenges:

⁵ For instance, the first enrolment is verifiable at: <https://testnet.litecore.io/tx/feac5a1dc645c701ea17ceb1657541d7094fbb43f18749cd9bd8a54014bd0197>




Fig. 9: UniquID enables M2M offline verification of smart contracts, which makes possible to start a loaned smart vehicle even in case of network failure.

- IAM for the sensors is achieved through Wi-Fi passwords, which require complex maintenance and periodic rotation (e.g., in case of operator change).
- Multi-tenancy scenarios require an additional layer of authentication at the application level, which needs to manage passwords stored on the sensor itself.

UniquID approach pushes the storage of these credentials onto the distributed ledger, enabling a remotely controlled, asynchronous IAM solution. A typical implementation leverages a central orchestration interface that sends microtransactions containing the ACL between one sensor and its authorized users on the distributed ledger. As a result, no passwords are needed: combining the ledger pseudonymous identity from the wallets (installed on the operator’s device and the sensor), and blockchain-stored transactions that contain the ACLs, a sensor can recognize the operator’s device and provide the data that the user is authorized to collect.

Furthermore, the operator acts as the dispatcher of the most recent blocks on the ledger. PoW is sufficient to verify the integrity of any received block, reducing the risk of forged blocks (which would require network consensus) and removing the need for a trusted session between peers. In this way, even if always disconnected from the Internet, sensors can be kept up to date with the latest authorizations stored on the ledger. Figure 10 shows the resulting architecture.

6 Conclusion

In this paper, we first briefly covered why current IAM systems and PKIs are too cumbersome for the challenges posed by the IoT paradigm. We also showed how the CAP theorem strongly applies to blockchain-based IAMs, and how developers can independently decide how their applications react to network partitions.

We illustrated the overall UniquID architecture and how its SDK aims to simplify both the trust architecture as a whole and resource negotiation flows.

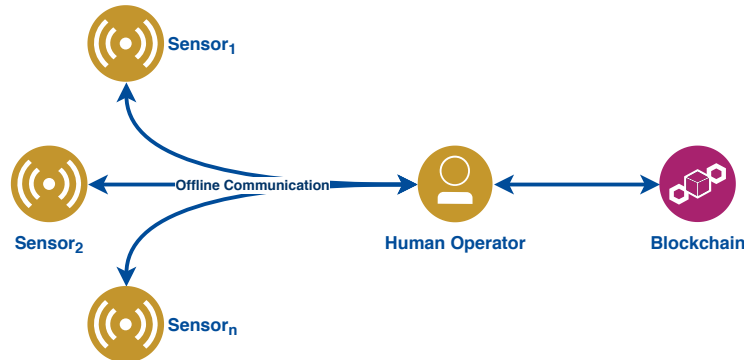


Fig. 10: UniquID enables a human operator to authenticate against non-connected sensors. Furthermore, the user dispatches to the sensors the latest blocks of the shared ledger.

We also discussed in detail how the backbone architecture of a blockchain-based IAM can be designed, as well as the benefits and drawbacks of each solution.

We provided an overview of the essential features that devices need to implement to participate in UniquID, as well as how the initial imprinting ceremony of such devices happens. We performed a proof-of-concept experiment which led to generating and imprinting 1000 identities on the Litecoin Testnet, virtually immutable and publicly verifiable by anyone. Last, but not least, we provided two relevant case studies that showed how UniquID could help unlock the potential of IoT. In future works we will conduct extensive experiments to assess the robustness and responsiveness of UniquID under a number of adverse conditions, such as man-in-the-middle and denial-of-service attacks.

Having shown the feasibility of our approach through practical experiments, future work will focus on optimising the performance of the system to get as close as possible to the theoretical upper bound of 1000 enrolled devices per minute.

Acknowledgements

We would like to thank Charles Kozierok for his help to proofread the manuscript.

References

1. Back, A., et al.: Hashcash-a denial of service counter-measure (2002), <http://www.hashcash.org/papers/hashcash.pdf>
2. Boyle, D.: Tesla driver gets stranded in the desert after leaving his keys behind (2017), <http://www.dailymail.co.uk/news/article-4128220/Tesla-driver-stranded-desert-forgot-keys.html>
3. Brewer, E.: Cap twelve years later: How the "rules" have changed. *Computer* **45**(2), 23–29 (Feb 2012). <https://doi.org/10.1109/MC.2012.37>

4. Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wulle, P., Maxwell, G.: Bulletproofs: Short proofs for confidential transactions and more. 2018 IEEE Symposium on Security and Privacy (SP) **00**, 319–338. <https://doi.org/doi.ieeecomputersociety.org/10.1109/SP.2018.00020>
5. Carayannis, E.G., Turner, E.: Innovation diffusion and technology acceptance: The case of pki technology. *Technovation* **26**(7), 847–855 (2006)
6. Caronni, G.: Walking the web of trust. In: *Enabling Technologies: Infrastructure for Collaborative Enterprises, 2000.(WET ICE 2000)*. Proceedings. IEEE 9th International Workshops on. pp. 153–158. IEEE (2000)
7. Dhillon, V., Metcalf, D., Hooper, M.: The hyperledger project. In: *Blockchain Enabled Applications*, pp. 139–149. Springer (2017)
8. Grandison, T., Sloman, M.: A Survey of Trust in Internet Applications. *IEEE Communications Surveys Tutorials* **3**(4), 2–16 (2000). <https://doi.org/10.1109/COMST.2000.5340804>
9. Greenberg, A.: Hackers remotely kill a jeep on the highway with me in it (2015), <https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/>
10. ISO/IEC: Iso/iec 20922:2016 - information technology – message queuing telemetry transport (mqtt) v3.1.1 (2016), <https://www.iso.org/standard/69466.html>
11. Kshetri, N.: Can blockchain strengthen the internet of things? *IT Professional* **19**(4), 68–72 (2017). <https://doi.org/10.1109/MITP.2017.3051335>, [doi.ieeecomputersociety.org/10.1109/MITP.2017.3051335](https://doi.org/10.1109/MITP.2017.3051335)
12. Le, T., Mutka, M.W.: Capchain: A privacy preserving access control framework based on blockchain for pervasive environments. In: *2018 IEEE International Conference on Smart Computing (SMARTCOMP)*. pp. 57–64 (June 2018). <https://doi.org/10.1109/SMARTCOMP.2018.00074>
13. van der Meulen, R.: Gartner says 6.4 billion connected "things" will be in use in 2016, up 30 percent from 2015 (2015), www.gartner.com/newsroom/id/3165317
14. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008), <https://bitcoin.org/bitcoin.pdf>
15. Novo, O.: Blockchain meets iot: An architecture for scalable access management in iot. *IEEE Internet of Things Journal* **5**(2), 1184–1195 (April 2018). <https://doi.org/10.1109/JIOT.2018.2812239>
16. Ouaddah, A., Abou Elkalam, A., Ait Ouahman, A.: Fairaccess: a new blockchain-based access control framework for the internet of things. *Security and Communication Networks* **9**(18), 5943–5964. <https://doi.org/10.1002/sec.1748>, <https://onlinelibrary.wiley.com/doi/abs/10.1002/sec.1748>
17. Roman, R., Zhou, J., Lopez, J.: On the features and challenges of security and privacy in distributed internet of things. *Computer Networks* **57**(10), 2266 – 2279 (2013). <https://doi.org/https://doi.org/10.1016/j.comnet.2012.12.018>, <http://www.sciencedirect.com/science/article/pii/S1389128613000054>, towards a Science of Cyber Security Security and Identity Architecture for the Future Internet
18. Sasson, E.B., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., Virza, M.: Zerocash: Decentralized anonymous payments from bitcoin. In: *2014 IEEE Symposium on Security and Privacy*. pp. 459–474 (May 2014). <https://doi.org/10.1109/SP.2014.36>