

# SOFIE

## Secure Open Federation for Internet Everywhere

Arto Karila, Yki Kortensniemi,  
Dmitry Lagutin, Pekka Nikander  
and Santeri Paavolainen  
Aalto University, Helsinki, Finland

Nikos Fotiou, George C. Polyzos,  
and Vasilios A. Siris  
Athens University of Economics and Business,  
Athens, Greece

Theodore Zahariadis  
Synelixis, Athens, Greece

**Abstract**—According to some studies, there are well over 300 different Internet of Things (IoT) platforms and several dozens of (so-called) standards. In practice, this has led to a situation where most of the deployed IoT systems are closed and largely incapable of communicating with other IoT systems. Furthermore, it is likely that at least a few different basic IoT communication protocols will co-exist, including the Constrained Application Protocol (CoAP), Message Queue Telemetry Transport (MQTT), and HTTP.

In this paper, we introduce SOFIE, Secure Open Federation for Internet Everywhere. SOFIE is a proposal for applying distributed ledger technology (DLT) to securely and openly federate IoT platforms. SOFIE’s approach is based on the idea of using interconnected distributed ledgers as a cornerstone to build decentralised business platforms that support the interconnection of diverse IoT systems. Among other things, the ledgers are used to provide openly accessible metadata about the individual IoT platforms, to define business and other rules on how to connect to the platforms, and to securely record audit trails that can be used to resolve disputes.

The SOFIE concept will be prototyped and studied in an EU Horizon 2020 funded project, commencing in early 2018 and running for three years.

The SOFIE concept will be prototyped and studied in an EU Horizon 2020 funded project, commencing in early 2018 and running for three years.

**Index Terms**—Internet of Things, Federation, Blockchains, Distributed Ledger Technology, Inter-Ledger

### I. INTRODUCTION

Fragmentation and lack of security are among the biggest problems of IoT systems. Most IoT platforms are vertically oriented closed systems, dedicated to specific application areas. There are today more than 360 different IoT platforms [1]. While many of these platforms are built on standardised interfaces, interoperability between most of them is still only a dream. IoT related standardisation suffers from similar fragmentation, with tens of competing, and sometimes self-assigned, standardisation organisations and well over a hundred different standards. As usual, proper end-to-end security

and privacy remain areas with the least amount of interoperability.

The main goal of the SOFIE approach is to enable diversified applications to utilise heterogeneous IoT platforms across technological, organisational and administrative borders in an open and secure manner, making reuse of existing infrastructure and data easy. Secure open federation is the key concept of SOFIE’s approach, aiming to enable creation of open business platforms, building upon existing IoT platforms and distributed ledgers, without needing to negotiate with any gatekeeper or an overlay interconnection entity.

The SOFIE approach is based on the following four principles:

*Federation:* The baseline of the SOFIE approach is federation, not integration. This means that each IoT platform, siloed or not, remains internally intact. (See Section III for further details.)

*Openness:* From the business point of view, anyone can join an open business platform, as there are no gatekeepers, in either a technology or business sense. At the technical level, virtually any IoT platform can join the federation, provided that it has some open interfaces, often even without support from the technology vendor. In a business sense, it is possible through federation to provide sufficient rules and conditions for allowing anyone to join, often without human intervention.

*Security:* We exercise security by design. Our goal is to build modules with the necessary security and privacy features that offer protection against cyber-attacks through the unforgeable Distributed Ledger Technologies (DLTs), thereby establishing transparency and accountability, and giving users better control of their data. (See Section IV)

*Data sovereignty:* In SOFIE, data is shared in a controlled way, within the bounds of security and privacy policies defined by the owners of the data. While strongly coupled with security, this aspect has characteristics and requirements which go well beyond what is addressed by traditional information security. (Section IV-A)

From the technology point of view, the main idea behind SOFIE is to define and implement a technology-agnostic approach to federate existing IoT systems, utilising blockchains and other distributed ledger technologies (DLTs). System interoperability will be largely based on the W3C WoT standards [2] and, where feasible, FIWARE IoT technology [3],

extended with any necessary new security mechanisms. A major effort in SOFIE is to define and implement cross-standard security mechanisms that enable open federation by recording security related information in administratively decentralised blockchains and other DLT systems.

The remainder of this paper is organised as follows. First, in Section II we briefly cover baseline work on the top of which SOFIE plans to contribute. Section III explains the main SOFIE approach, federation. In Section IV we discuss the so-far identified main security and privacy issues. Section V outlines the contributions we expect from SOFIE. Finally, Section VI summarises the paper.

## II. RELATED WORK

The technology baseline for SOFIE consists of existing IoT standardisation and consolidation efforts, decentralised security, blockchain based IoT actuation, and simultaneous use of multiple blockchain technologies. In the next few sections we very briefly present the state of the art in these fields.

### A. Existing IoT platforms

Currently, there are some 360 different IoT platforms in various application domains [1]. Many of these platforms are proprietary and either fully or partially closed-source. There is no unified way of addressing security across different platforms. For most deployments, security is dealt separately, with TLS or DTLS “tunnels”, which does not provide proper end-to-end security.

Below we discuss a few representative IoT systems, and specifically UNIFY-IoT, OMA oneM2M, OCF IoTivity and related AllJoyn, and OMA LWM2M, all of which aim to provide connectivity across platforms.

The UNIFY-IoT project [4] has defined an eight layer IoT stack, starting with the physical and network layers, up to the data abstraction, service, application and collaboration layers, roughly corresponding to the OSI 7-layer model. The stack includes several potential federation or integration points: At the lowest network layer there are CoAP and/or HTTP ReSTful APIs. The semantic, processing, and storage layers provide more APIs, such as the FIWARE NGSI 9/10 or W3C WoT servient APIs.

The oneM2M standards [5] focus on providing semantic interoperability between different IoT platforms, introducing Common Services Entities (CSE) between applications (data processing) and networks (communications capabilities). FIWARE includes an open source oneM2M IoT agent. Interoperability between FIWARE and oneM2M has been demonstrated several times and is available as open source.

In addition to W3C WoT and oneM2M, IoTivity [6] by the Open Connectivity Foundation (OCF) is an open source software framework that aims at facilitating device-to-device connectivity; the AllJoyn framework [7] by AllSeen alliance is another open source approach that allows device and application discovery independent of the transport technology. After the late 2016 when the AllSeen Alliance and OCF merged,

further development of the AllJoyn framework ceased, apparently in favour of IoTivity.

OMA LightweightM2M (LWM2M) [8] defines an object and interaction model that can be used for managing things that support the CoAP protocol.

### B. Decentralised security

To the best of our knowledge, so far there are no genuinely decentralised federated security solutions. Some decentralised solutions, like Simple Public Key Infrastructure (SPKI), were developed in the late 1990s, but were never deployed. However, with the success of Bitcoin and Ethereum, decentralised solutions have now become socially accepted (and even hyped). In this new environment, SOFIE aims to achieve real decentralisation and will implement security solutions that do not require any “root of trust”, by leveraging blockchain technology. Beyond being just a theoretical advantage, this is a key property that will increase trust and acceptance, allowing fast adoption and organic growth of SOFIE-based business platforms.

### C. IoT actuation

Regarding IoT actuation, there are already a number of proposals to combine IoT and blockchains, such as IBM Watson IoT with blockchain, Tileplay, Catenis, Chronicled, slock.it, TransactiveGrid, Gridgularity, SolarCoin, Farmshare, Provenance, Chain of Things, Iota, and Flowchain. All these solutions are vertical and bound to a particular blockchain technology. SOFIE’s inter-ledger transaction layer aims to allow the development of a multitude of applications across domains, technology and business silos, while also allowing the use and interoperation of multiple blockchain technologies. SOFIE also plans to develop DLT based mechanisms that can correlate, in a non-repudiable manner, the physical world sensing of the results of actuation with the corresponding actuation.

### D. Inter-ledger approaches

Regarding using multiple ledgers simultaneously, the Inter-ledger protocol (ILP) [9] and the W3C Inter-ledger Payments Community Group [10] focus on payments across DLT systems. A very interesting related proposal is Polkadot [11] by Gavin Wood, the author of the Ethereum yellow paper. Other proposals in the area include Überledger [12], Cøsmos Network [13], and Blocknet [14]. SOFIE’s goal is to combine the best aspects of the above approaches to produce a solution for the IoT that goes beyond supporting simple payments. Also, SOFIE’s target is to be compatible with at least ILP and potentially also Polkadot.

## III. FEDERATION

A major challenge in the evolving IoT world is the fragmentation of vertically oriented, closed systems, architectures and applications, and moving on towards open systems and platforms. Interoperability between most of the above is for the most part still only a dream and sometimes even shunned

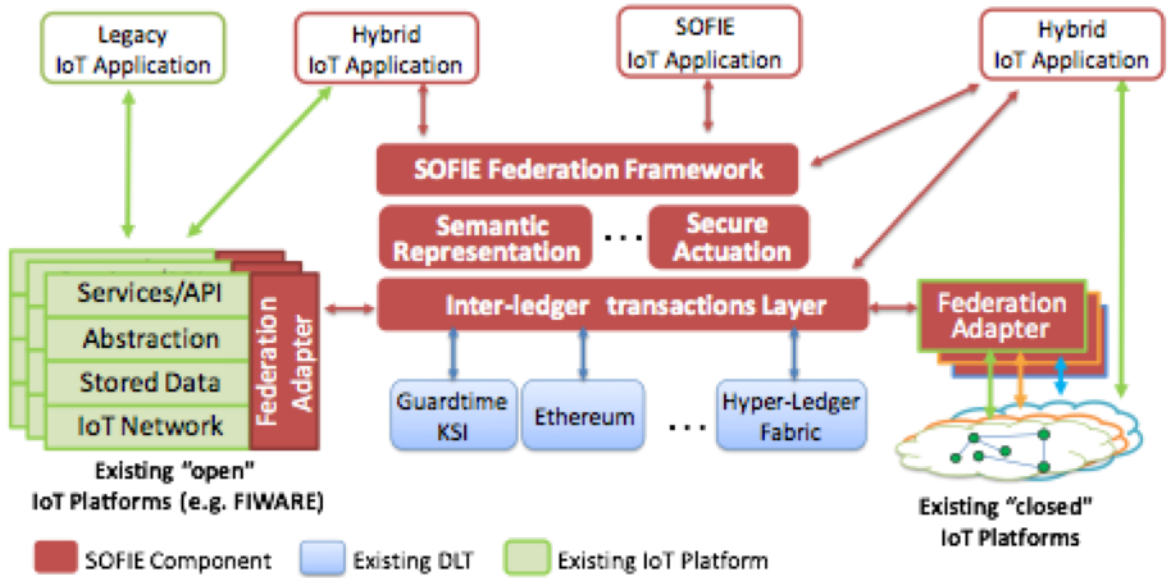


Figure 1: Overview of the SOFIE federation architecture, its components and relationships to existing systems

by vendors. The SOFIE strategy is to achieve interoperability through technology-agnostic, secure, open federation, allowing existing systems to interact in a secure and privacy-preserving manner.

The SOFIE federation approach allows for on-the-fly adaptation and interpretation of data and control messages while keeping the data in their respective IoT platforms. Being based on federation and DLTs, we expect SOFIE to scale to support billions of devices.

Our approach is to initially create a federation architecture, then define a federation framework based on the federation architecture, and finally, based on the federation architecture and framework, create open business platforms for three pilots in three different sectors: food chain, gaming, and energy. The federation framework is expected to consist of mostly existing components, though some new components will also be developed, for example, for data privacy and sovereignty. To facilitate adoption, we will define well-documented procedures to use the components to create new business platforms.

For federating security, we propose a novel open approach, based on blockchains and other DLTs. Furthermore, instead of using a single blockchain, SOFIE’s federation uses an inter-ledger approach, where multiple independent and diverse distributed ledgers may be used in parallel, while also allowing the replacement of one ledger technology with another, if needed.

SOFIE enables the combination of several IoT platforms and ledgers into a federated IoT platform supporting the reuse of existing IoT infrastructure and sharing of data by various applications and businesses. Figure 1 illustrates the overall architectural approach.

#### A. Inter-ledger transactions

The completely decentralised core of the architecture is formed by the inter-ledger transaction layer, which allows

transactions to be recorded into multiple blockchains or other ledgers in a logically atomic manner. We will build upon existing cutting-edge work, including the W3C-associated Inter-ledger Protocol (ILP), adapting existing solutions to the IoT domain, and developing them further. The logically atomic transactions will be implemented as multi-stage smart contracts whose resolution is determined by the requirement that the transactions are correctly recorded in all the participating ledgers, without however requiring that all the ledgers support smart contracts.

The inter-ledger transaction layer will be used for achieving three main targets:

- Describe the data (“things”) in the existing IoT platforms, thereby enabling financially tied IoT actuation between organisations. For data description, our tentative plan is to use and augment W3C WoT Thing Descriptions and/or use representations such as RFC6690. Whether the descriptions include the full representation of things or just their hashes and (short) URLs involves different tradeoffs, which will be investigated, and is related to where “thing events” need to be handled.
- Enable secure and traceable IoT actuation. The idea is to negotiate and use smart (micro-) contracts that may span multiple ledgers to record intention or desire to actuate, to trigger actuation, to permanently record both actuation instances and the corresponding sensor values, and to trigger any financial transactions, thereby supporting smart behaviour.
- Enable interoperability between diverse existing IoT platforms. In many cases there is a need to augment the existing IoT platforms with a federation adapter. While the adapter may in some cases be somewhat complex, we believe that if the IoT platform supports existing open standards, the adapter will mostly deal with security

and privacy, with little or no actual application-level functionality.

SOFIE's federation approach is designed to be technology-agnostic, allowing systems with different APIs and data formats to interoperate, while satisfying the applicable security policies. Some of the existing IoT platforms already support interoperability across different protocols and standards. Examples of this include FIWARE through its IoT adapters, such as the LWM2M and oneM2M adapters, and W3C WoT, where the IoT servient concept supports both proprietary APIs and various protocol adapters.

#### IV. SECURITY AND PRIVACY

Most existing solutions already provide decent end-to-end security within a single system and system-specific authentication. Therefore, SOFIE focuses on innovation in the areas of data sovereignty, privacy, federated key management, authentication, and authorisation.

The security of current and emerging IoT platforms is mainly based on Transport Layer Security (TLS), Datagram TLS (DTLS), JSON Web Tokens (JWT), and JSON Web Encryption (JWE). We will re-use the existing security solutions and add federated inter-platform key and certificate management. We will also closely follow the developments of the IETF's Constrained RESTful Environments (core) working group in this domain. Instead of using traditional certificate revocation techniques, SOFIE's approach is to record certificate revocation information into the blockchains. Blockchains can also be utilised for implementing secure firmware upgrades of IoT devices, offering device-level protection against malware. Storing important data to blockchains, including the combination of physical world sensing of actuation results and the corresponding actuation requests, also provides traceability and non-repudiation.

Authentication and authorisation are two important security functions, whose efficient implementation in the context of the IoT is very challenging (RFC7744). SOFIE will utilise technologies for authentication and authorisation that already facilitate federation, such as OpenID Connect and OAuth. Indeed, adaptations of these two technologies are considered by the Authentication and Authorisation for Constrained Environment IETF WG. Moreover, using blockchain-based immutable audit trail makes the detection and post-mortem analysis of cyber-attacks more reliable than in current systems.

##### A. User-oriented privacy

User control over their data is important in the IoT, both for security and privacy. Currently the vendors of various IoT systems typically store user data in proprietary Clouds and have complete control over the data, which is not acceptable. Such approaches also result in a tussle over the control of data between multiple stakeholders (e.g. the utility company, the owner of an office building, the tenant of the office space, the individual working in the office), which we will need to carefully consider during the project.

IoT data can often be personal, the use and processing of which has been tightly controlled within the EU based on the EU Data Protection Directive (Directive 95/46/EC). The new General Data Protection Regulation (GDPR) that becomes enforceable in May 2018 puts strict requirements on the security and privacy of personal data and complements them with significant sanctions for parties failing to meet the requirements. Thus, ensuring compliance with the GDPR is a major design requirement for the SOFIE architecture.

One tool to achieve the above data privacy requirements is MyData, an approach to allow individuals to better control how their personal data is used. MyData provides GDPR compliant tools and open source implementations for use cases where the individual consents to sharing data from one service to another while making sure that both legal restrictions and secure data sharing requirements are met. Furthermore, we are already involved in ongoing work to augment the MyData environment with blockchains, such as the Sovrin Foundation's identity blockchain. We will also consider using some results of the Hub of All Things (H.A.T.) project, which allows storing personal data in a user-controlled manner.

In order to support data sovereignty and privacy, SOFIE adopts a three-level approach to the storage of data. First, there is a private data store managed entirely by the stakeholder. A private blockchain (such as Guardtime's KSI) forms the second level, containing data that is shared between collaborating stakeholders (for example: producer, reseller, and supermarket, in the food chain use case). Finally, some data will be stored in a public blockchain, such as Ethereum or Bitcoin. Such an approach allows fine grained control of the data supporting policies that can range from total openness (e.g. to bring transparency to certain public services) to very tight access control (e.g. to protect trade secrets or the privacy of people). In all cases, integrity and non-repudiation of the data is important.

##### B. Privacy-related challenges

Immutability of the data stored in blockchains can give rise to privacy challenges. Indeed, the increasing pool of data can be mined for insights and dedicated techniques, such as correlation attacks, can reveal even obfuscated information. Therefore, careful analysis is needed to determine what information should and should not be stored in the blockchain and what methods should be used to protect this information. In most situations only hashes of the actual data (e.g., the hash of the root of a Merkle tree containing the actual transactions) will be stored in private or public blockchains. Extremely sensitive data will only be stored privately. Furthermore, SOFIE's inter-ledger layer will consider privacy preserving distributed ledger implementations (such as z-cash) and it will utilise covert-channel techniques designed for communication networks to hide blockchain access patterns.

#### V. EXPECTED CONTRIBUTIONS

We expect that SOFIE will introduce novel contributions in the following areas:

- Industry-wide open business platforms for the IoT.
- Secure and open federation utilising multiple blockchains and other ledger technologies.
- Extension of inter-ledger transactions from payments to the IoT world, adding support for device actuation.

The expected main technical contributions of the planned research can be summarised as follows:

The SOFIE federation architecture provides concepts, methods and tools that allow simultaneously secure and open federation of both open and closed IoT platforms, which include existing and evolving sensing, actuating, energy harvesting, networking and interface IoT technologies.

In the highly distributed, heterogeneous and dynamic IoT environment, the SOFIE security architecture provides end-to-end security (data confidentiality and integrity), identification, authentication and authorisation, cyber-attack detection and resilience, while supporting users' privacy and control over their data. This will be achieved by utilising both currently available solutions and new approaches, such as the MyData concept, and blockchain properties combined with traditional security mechanisms.

A particular focus will be on secure actuation because of the potential real-world risks and the added complexity from needing external certification of actuation outcomes. This will be the main domain of our cyber-attack prevention and containment work.

SOFIE will combine decentralised user-controlled identities, "self-certified" authorisation, smart contracts, and MyData-based (or similar) privacy, together with the open business models made possible with Bitcoin and Ethereum, thereby allowing federated systems to be fully open and secure at the same time.

From the business platform point of view, utilising the SOFIE secure open federation approach will enable new business models and increased application innovation. We will study both the technical and other scalability hindrances of open blockchains, and apply the results to allow the inter-ledger transaction layer to support open business platforms in a scalable way. The starting point for this work is the so-called equilateral governance model that defines how the existing open blockchains are governed.

Another area where we expect to provide innovative contributions is in implementing IoT actuation through inter-ledger transactions. We will address it in cross-disciplinary manner, considering related technical, contractual, and legal issues. For resilience, the actuation itself and the sensors that measure the real-world effects of actuation must be separated, so that the measurement results can be used for contract arbitration without potential conflicts of interest. We expect secure IoT actuation to take place through smart contracts, entered through the inter-ledger transaction layer. This makes it possible to provide simultaneously secure and open APIs where users can pay for or otherwise contractually request and verify local and remote IoT actuation in an accountable and non-repudiable manner.

#### A. An open approach

For us, a major goal is sustained availability of the federation approach and framework. We intend to release the entire SOFIE federation framework as open source under the Apache V2.0 license, hence encouraging interested project partners and any external parties, including individual hobbyists and startup companies, to pick up, contribute, and build upon our work. In this way we seek to establish SOFIE as one of the leading frameworks for innovative IoT solutions.

To maximise our stakeholders' interest and uptake, we will adopt the following practices:

- We will commence open development from the project's start, using well-established code management platforms that ensure powerful collaboration, code review and code management. The most likely candidate for that is GitHub.
- We will set up a SOFIE community, linked with the FIWARE and other IoT communities, to support the platform and proposed approach.
- We will develop documentation and guidelines for prospective contributors to accelerate their familiarisation with SOFIE's vision and platform. A SOFIE wiki will be set up early, containing precise guidelines for a) contributing modules/elements to the SOFIE federation framework and b) creating applications with the associated SDKs.

## VI. SUMMARY

SOFIE, "Secure Open Federation for Internet Everywhere", is a technical and business approach for building open business platforms, whose target is to federate existing heterogeneous IoT systems and platforms. SOFIE is based on the idea of using lightweight IoT system adapters to "open up" IoT functions by describing their data model, protocol accessibility, and business accessibility in a DLT-based or DLT-like open directory, with the goal of allowing at least trivial access with "smart contracts" running in some DLT. In addition to using DLTs for opening existing IoT platforms, DLTs will also be used to record usage contracts, for pay-as-you-go actuation scenarios and for providing a non-repudiable audit trail. A large part of the forthcoming work will focus on balancing openness, security, and privacy aspects.

The SOFIE project started in January 2018, and will run for three years. Our goal is to perform the work and establish tight liaison within open source and other communities so that SOFIE's solutions will be sustainable even after the project's end. To achieve this, SOFIE will conduct the whole framework development in a fully open source manner from the start of the project.

## ACKNOWLEDGMENT

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 779984.

## REFERENCES

- [1] IoT Analytics. *IoT Platforms Market Report 2015-2021*. 2015. URL: <https://iot-analytics.com/product/iot-platforms-market-report-2015-2021-3/> (visited on 2017-09-22).
- [2] W3C. *Web of Things Working Group*. 2017. URL: <https://www.w3.org/WoT/WG/> (visited on 2017-08-01).
- [3] Juanjo Hierro. *FIWARE, the Standard That the IoT Needs*. 2016-09-26. URL: <https://iot.telefonica.com/blog/2016/09/en-fiware-standard-iot> (visited on 2017-08-01).
- [4] UNIFY-IoT. *Creativity, Innovation and Interaction for the Internet of Things*. 2017. URL: <http://www.unify-iot.eu/> (visited on 2017-08-01).
- [5] oneM2M. *Standards for M2M and the Internet of Things*. 2017. URL: <http://www.onem2m.org/> (visited on 2017-08-01).
- [6] IoTivity. *A Linux Foundation and Open Connectivity Foundation Open Source Project*. 2017. URL: <https://www.iotivity.org/> (visited on 2017-08-01).
- [7] AllSeen Alliance. *AllJoyn Framework*. 2016. URL: <https://allseenalliance.org/framework>.
- [8] Open Mobile Alliance. *Lightweight M2M (LwM2M)*. 2017. URL: <http://openmobilealliance.org/iot/lightweight-m2m-lwm2m> (visited on 2017-08-01).
- [9] Interledger. *Interledger, a Protocol for Connecting Payment Networks*. 2017. URL: <https://interledger.org/> (visited on 2017-08-01).
- [10] W3C. *Interledger Payments Community Group*. 2018. URL: <https://www.w3.org/community/interledger/> (visited on 2018-02-12).
- [11] Gavin Wood. "PolkaDot: Vision for a Heterogeneous Multi-Chain Framework". In: (2017-07-24). URL: <https://polkadot.io/>.
- [12] Überledger. *Überledger, Disintermediating Inter-Blockchain Transactions*. 2017.
- [13] Interchain Foundation. *Cosmos Network - Internet of Blockchains*. 2018. URL: <https://cosmos.network/> (visited on 2018-02-12).
- [14] Blocknet. *Blocknet - Decentralized Application Platform*. 2018. URL: <http://blocknet.co/> (visited on 2018-02-12).