https://doi.org/10.3011/ESARDA.IJNSNP.2021.6

Three properties of Distributed Ledger Technology systems applied in the nuclear sector adding value to safeguards - immutability, timestamping and auditability

Marco Sachy, Roberto Spigolon and Stefan Nonneman

European Commission, Joint Research Centre (JRC), Ispra, Italy
Directorate G - Nuclear Safety and Security
Nuclear Security Unit
E-mail: marco.sachy@ec.europa.eu, roberto.spigolon@ec.europa.eu, stefan.nonneman@ec.europa.eu

Abstract

We present our preliminary findings from the exploratory research project Shared Ledger Technologies for Nuclear Safeguards (SLT4SFG) conducted at the Joint Research Centre of the European Commission. As long as nuclear fuel (fresh or spent) is stored in facilities of a state, that state should always comply with the international non-proliferation treaty. Safeguards' verification processes will be affected also by the digital transformation. The exploratory research project SLT4SFG, has the objective to provide an evidence-based analysis on the benefits and challenges of Distributed Ledger Technology (DLT) systems for the nuclear sector, with special emphasis on safeguards. After presenting background knowledge on DLT systems, we discuss three key properties adding value to nuclear safeguards processes: (1) practical immutability can improve security of sensors identity management; (2) data anchoring through decentralised timestamping can provide proofs of existence and non-alteration of relevant data; and (3) auditability can increase efficiency in data sharing and become a source of forensic evidence to help determine legal liability. We inferred these three properties of DLT systems by endorsing a deductive methodology for the definition of nuclear sector use cases on containment and surveillance and radiation protection. Our aim is to frame a Proof-of-Concept strategy of software implementations intended to be ultimately exploited in nuclear safeguards. We conclude pointing to future research on performance tests simulated on the JRC Experimental Infrastructure for Internet Contingencies. We will aim at offering metrics to quantitatively measure the performance and derive added value of DLT systems properties applied to nuclear safeguards.

Keywords: Distributed Ledger Technology; blockchain; safeguards; immutability; timestamping; auditability.

1. Introduction

In the nuclear industry, the regulatory requirement to avoid the diversion of nuclear material from its intended uses shall be met for many years to come [1]. Indeed, nuclear safeguards will be required in both the near and distant future to ensure that nuclear material will be used only within regulatory constraints. In order to lower operational costs and increase efficiency in the management of nuclear safeguards business processes such as – but not limited to – containment and surveillance and nuclear material accountancy, the digital transformation in the nuclear industry will arguably have a significant impact also on safeguards.

Among many technologies such as robotics, the Internet of Things or still Artificial Intelligence, the European Commission considers also Distributed Ledger Technology (henceforth, DLT) systems, such as blockchains, as innovations with a high transformative potential. In this context, the European Union is promoting initiatives such as the European Blockchain Partnership [2], the EU Blockchain Observatory and Forum [3] together with open consultations for the European Blockchain Services Infrastructure [4] and the European Blockchain Pre-Commercial Procurement [5].

As a contribution to a strategy on "continuity of knowledge and data for very long periods", the European Commission's Joint Research Centre is providing insights and foresight on DLT systems in various domains. While we leave political and legal considerations to future research, the scope and purpose of this paper is exclusively to present our preliminary evidence-based technical research outputs from the exploratory research project Shared Ledger Technologies for Nuclear Safeguards (SLT4SFG) [6]. The objective of the SLT4SFG explorative research project is to provide evidence-based answers on whether and to what extent DLT systems and their properties can improve and add value to nuclear safeguards business processes.

Roughly put, DLT systems implement peer-to-peer networks deployed to validate digital assets' transaction history on an append-only and tamper-evident log, replicated to all participating nodes. DLT systems leverage applied cryptography and distributed computing to achieve consensus on global system state among either completely or partially distrusting parties.

Sachy, M., Spigolon, R., & Nonneman, S. (2021). Three properties of Distributed Ledger Technology systems applied in the nuclear sector adding value to safeguards – immutability, timestamping and auditability. *ESARDA Bulletin - The International Journal of Nuclear Safeguards and Non-proliferation*, 62, 60-79. https://doi.org/10.3011/ESARDA.IJNSNP.2021.6

While they periodically experience 'hype' phases, DLT systems form a new family of technologies that has not been yet thoroughly explored and validated. Our preliminary evidence-based findings suggest that individual properties of DLT systems can already add value to nuclear safeguards business processes by becoming a features' layer firstly added to, and in the future possibly replacing, legacy systems. In other words, as we will argue more in detail in the sections below, DLTs systems' properties can add a synergistic layer of functionality to legacy systems, viz. a series of complementary modules that can be initially identified by the three properties discussed in this paper. In the future, especially in the case where DLT systems will be validated and standardised, various properties that today we propose to deploy autonomously in concert with legacy systems could entirely replace them.

Accordingly, below we discuss three key DLT systems' properties, i.e. practical immutability, decentralised timestamping and structural auditability. In our view, they offer benefits and add value to nuclear safeguards data management techniques. These properties emerged as the most relevant ones for testing use cases in different nuclear sector's domains: nuclear safeguards - with a focus on containment and surveillance - and radiation protection. For the use cases on radiation protection, we indicate how the same principles could be relevant for nuclear safeguards.

Firstly, we introduce practical immutability as a property related to public distributed ledgers through a use case in the context of containment and surveillance that does not involve any sensitive data. We propose to leverage practical immutability of data stored on a public blockchain for Public Key Infrastructure management to prevent that a malicious insider forge a valid digital identity for a sensor used in safeguards containment and surveillance without leaving any trace. Alongside increased cybersecurity, adding a blockchain-based module to Public Key Infrastructure management can contribute to deploy new systems for remote monitoring (e.g. with the use of digital seals). This could desirably impact inspections planning activities with potential cost savings related to the deployment of human resources.

Secondly, we introduce decentralised timestamping as another property inherent to public distributed ledgers that can be used as a data anchoring service for datasets, proving both the existence of data at a certain moment in time and the absence of alterations. We identified the benefits of this property in the radiation protection context, as an integrity layer to prove in the future that historical data on absorbed doses of personnel will not have been modified. The same property can be leveraged by nuclear installations operators in the context of nuclear material regular mailbox declarations, especially in the case where such declarations are not sent to the inspectorate but remain inside the facility. Moreover, decentralised timestamping offers a supplementary layer of integrity for databases and datasets backups.

Thirdly, we introduce structural auditability as a property of distributed ledgers in general. We studied the benefits of this property on the digitalisation of the radiation passbook used by workers (and inspectors as well) when travelling between different nuclear sites. By virtue of its internal structure as a chain of transactions, a digital radiation passbook implemented with a permissioned DLT system structurally provides an auditable trail of the history of records related to workers exposed to ionizing radiation. Considering also the non-repudiation property (i.e. who committed a transaction cannot repudiate it), DLT systems can thus become a source of forensic evidence that can be used to help determine legal liability in case of disputes.

Moreover, the deployment of smart contracts (i.e. computer programs encoding a business logic whose output is stored on each node of a distributed ledger) automates and enforces the execution of workflows and lowers the rate of clerical errors. Because information exchange would take place on a commonly shared system, rather than through different centralised databases still processed with a significant degree of human intervention as for current practices, it follows that a DLT-based system would be even more easily auditable from a backend perspective. The same kind of approach could be applied also in the domain of Nuclear Material Accountancy and Control as a way to integrate and coordinate the execution of the workflows related to nuclear material accountancy, ease both data sharing procedures and the auditability of the whole process.

Although not all of them are strictly related to nuclear safeguards, the use cases presented below have been selected by endorsing a broad deductive methodology. From the DLT systems' properties, we inferred use cases through a top-down scientific experimental approach to test them in the nuclear sector. For each use case, we then endorsed the best fit-for-purpose software development methodology to implement a Proof-of-Concept strategy. In this way, we could better assess whether and to what extent our deductions were corroborated by evidence to prove the correctness, or lack thereof, about our ideas on the applicability of DLT systems properties to the nuclear sector with special focus on added value for nuclear safeguards.

The remainder of this paper is structured as follows. Section 2 briefly presents related work and elicits background knowledge on DLT systems by providing an analysis of their general benefits and challenges. Section 3 concisely elicits our methodological choices. Section 4 analyses three key DLT systems' properties, their benefits and added value to nuclear safeguards. We conclude the paper in Section 5 pointing to a potential way forward for future research on performance testing with the emulation of use case implementations on the JRC Experimental Platform for Internet Contingencies (EPIC). This will enable us to reproduce the performance of DLT systems behaviour in real world network conditions, under a fully controllable experimentation environment. In this way, we will aim to establish more granular quantitative metrics to measure both benefits and added value of DLT systems' properties applied to nuclear safeguards.

2. Related work and background knowledge on DLT systems

2.1 A brief introduction to DLT systems

The European Commission is not the only organization currently exploring the applicability of DLT systems in the nuclear safeguards domain. To our knowledge, there are other active actors in the field, especially in research institutes in the United States of America, also with international collaborations, for instance experimenting on nuclear safeguards data management [7], transit matching [8] and on aspects of DLT systems deployment for UF6 cylinder tracking and process monitoring [9].

First, to our knowledge the Stimson Centre is currently the main actor in the United States landscape that is exploring DLTs added value applications in this domain [10] through its "Blockchain in practice" program. Together with the University of New South Wales (UNSW) and the Finnish Radiation and Nuclear Safety Authority (STUK), they launched the SLAFKA prototype: a permissioned blockchain system that enables nuclear facilities to record nuclear material assets on a distributed ledger. It is implemented using Hyperledger Fabric [11], an open source permissioned DLT project maintained by the Linux Foundation. SLAFKA has been implemented to test DLT and how such technology performed in handling safeguards transactions: instead of having a primary role of the regulators to settle the transactions, nuclear facilities would be able to transact assets whilst being supervised by regulators. SLAFKA's prototype outcomes are:

1. The introduction of a distributed networking approach to safeguards reporting.
2. A way to reduce reconciliation time among State and operators.
3. A single source of truth for the management of safeguards information.

Secondly, the Pacific Northwest National Laboratory (PNNL) simulated a transit matching system based on DLT, experimenting with both Hyperledger Fabric and Ethereum [8]. Their main goal was to understand whether a DLT system could bring benefits in comparison to the current IAEA approach. The outcomes on this prototype are:

1. a DLT system could improve the efficiency of the process through real-time match attempts of all transactions posted to the ledger.
2. Using "graded scores" applied to match attempts could represent a useful source of information for increasing the effectiveness of safeguards inspections.
3. Since the DLT system is a tamper-evident record of transactions, transit matching operations performed on such a system could lead to an increased confidence on IAEA safeguards conclusions through transparent reconciliation of transit matching reports.

Because they are converting a legacy system into a DLTs-based one, they also highlight that among these three findings, only the last one is really dependent on the technology used (i.e.: a distributed ledger), while the first two could be achieved also with "traditional" technologies.

Finally, Sandia National Laboratories built their DLT-based prototypes on the field. According to information shared during the 2020 Institute of Nuclear Material Management (INMM) Annual Meeting [9], they built a prototype based on a private version of Ethereum where they stored together Inventory Change Reports data and sensors data such as gamma-ray events and video cameras recordings. The idea was to enable workflows, e.g. retrieving of sensors data to validate an inventory change.

In terms of background knowledge on DLT systems, the past decade witnessed significant advancements in integrated and applied cryptography for the innovation of distributed computing with the introduction of public blockchains such as Bitcoin [12] and more in general DLT systems [13] as blockchains are a subset of this larger class. Fully aware that the cryptologic history of these systems dates back decades, here we limit ourselves to an overview on this family of technologies by referring to the conceptual genealogy of the notion of 'blockchain', both in applied cryptography and in the nuclear industry.

Because a genealogy of a concept researches its original meaning to then provide current definitions, we will begin by presenting an etymology of this term. In applied cryptography, the term 'block-chain' can be traced back to block cyphers modes of operation algorithms [14]. In particular, the algorithm for Cipher Block Chaining (CBC) mode is defined such that "the plaintext is XORed with the previous ciphertext block before it is encrypted" [15]. Similarly, in the nuclear sector, the idea of using cryptographic techniques to ensure that data acquired to verify treaty compliance be trustworthy is also not new [16]. In effect, during the 20th century, cryptographers operating in the nuclear sector developed techniques to solve problems of mutual distrust, whereby "data as well as the redundant identifying information would be block-chain encrypted" [16].

What has been novel in the more recent past, perhaps relies in the fact that cipher block-chaining evolved and coalesced with other advancements in cryptology for digital cash applications such as e-cash [17] [18] and hashcash as a Proof-of-Work system [19] [20] into increasingly popular public blockchain protocols. After an initial focus on the seminal application, i.e. cryptocurrency, it became increasingly clear that the underlying blockchain technology had farther reaching implications and the potential to bring innovation to entire industries.

From a genealogical point of view, the reference implementation for DLT systems, i.e. the Bitcoin blockchain is a data structure representing the ledger of transactions that everyone participating to the network can store to acknowledge a common transaction history:

“a blockchain data structure is an ordered, back-linked list of blocks of transactions. Each block within the blockchain is identified by a hash, generated using the SHA256 cryptographic hash algorithm on the header of the block” [21].

This technical arrangement enables network participants to share a common transaction history among a group of distrusting peers or nodes. In the Bitcoin reference implementation [10], nodes compete to generate the next block and acquire a reward by consuming resources, i.e. electricity, to run the consensus algorithm, in the case of Bitcoin based on a Proof-of-Work mechanism. This mechanism is put in place to avoid that block producers named miners assign to themselves extra coins and engage in double spending. Participants can freely join the network by downloading the client without the need for human identity verification. Consequently, access to the Bitcoin blockchain is public and permission-less.

Following the deployment of the Bitcoin network in 2009, second generation blockchains such as Ethereum [22] added the possibility to execute smart contracts [23] on top of the blockchain layer. Smart contracts are computer programs that enable to perform transactional semantics more powerful than mere monetary exchange either directly on a distributed ledger by requesting all nodes to execute complex business logics or more simply to record their outputs on a distributed ledger. Smart contracts are designed to impersonate the role typically attributed to trusted third parties. As an example, we can consider the typical escrow use case, where some currency funds are managed by a smart contract and sent to the recipient only after specific conditions are met. More in general, smart contracts enable to program business logics on a distributed ledger, ensuring that their execution is not manipulated.

It then followed a plethora of implementations, some detaching from strict Proof-of-Work consensus blockchains and proposing alternative data structures, for instance Directed Acyclic Graphs, e.g. IOTA [24], Hashgraph [25] and Keyless Signature Infrastructure [26]. “Blockchain” has been then reclassified as a special case of DLTs systems as there are many other possible ways to achieve distributed consensus on the transaction history tracking in principle any type of digital data and asset, without relying on a central authority as a single source of truth.

As opposed to public permission-less distributed ledgers such as the Bitcoin or Ethereum blockchains designed for highly distrusting environments, a private, semi-private or permissioned distributed ledger which can be either a blockchain or another type of data structure, leverages on already existing trust and collaboration among stakeholders and processing units belonging to a shared operational environment. This type of DLT systems is shared by the members of either a single company or an industrial consortium. Stakeholders can agree on the type of consensus mechanism to order transactions, the governance of the infrastructure to run nodes participating to consensus rounds and define read/write access permissions for different types of participants tasked to maintain the distributed ledger.

Concluding our genealogical exercise on background knowledge, we adopt this overarching working definition of DLT system:

“A system of electronic records that (i) enables a network of independent participants to establish a consensus around (ii) the authoritative ordering of cryptographically-validated (‘signed’) transactions. These records are made (iii) persistent by replicating the data across multiple nodes, and (iv) tamper-evident by linking them by cryptographic hashes. (v) The shared result of the reconciliation/ consensus process - the ‘ledger’ - serves as the authoritative version for these records.” [13]

We propose this definition, because it is general enough to include public blockchains while leaving open the opportunity to explore other data structures such as permissioned distributed ledgers.

2.2 Key beneficial properties emerging from DLT systems

The general DLT systems properties that are considered beneficial in the literature referred to in the table below are summarised in Table 1.

2.3 Challenges emerging from DLT systems

By contrast to DLT systems’ benefits and desirable features, we identified a set of challenges that currently curb their widespread adoption. Indeed, research efforts are underway in governments, industry and academia to overcome the limits of such a relatively immature technology. As for the benefits elicited above, also with regard to challenges, alternative classifications are possible as a matter

Property	Description
DECENTRALIZATION	The reason why DLT systems, and in particular public blockchains, were designed in the first place is to eliminate the need of either a trusted third party or an intermediary responsible for validating and settling transactions [12]. A decentralized consensus mechanism, instead of an independent and centralized validator, enables all participants to the network to agree on the validity of transactions.
AUDITABILITY	A practically immutable audit trail of all identities comprising a record of related operations and any changes is kept [27]. This quality attribute can be achieved by DLT systems on a per case basis. Either all (e.g. in public blockchains) or some types of participants (e.g. in permissioned contexts) can maintain, store and access the full history of transactions or parts of it, depending on the opacity needs of the system.
AUTOMATION	Smart contracts enable programmable business logics. They are triggered when certain conditions are met. Their data output or a message digest referring to such data can either be stored off-chain in a local database or be written on a distributed ledger either before or after consensus on transaction ordering, execution and validation is reached by the distributed network. Some DLT systems require the installation and execution of the smart contract on each node of the network, while others pre-process the smart contract locally and broadcast to all nodes only the output of the smart contract.
STRONG IDENTITY MANAGEMENT	By design, DLT systems implement strong identity management mechanisms. This ensures that only who is allowed to perform a certain transaction/operation can actually do so. The identity of a user is verified through strong asymmetric cryptography algorithms.
TIMESTAMPING	DLT systems embed timestamping to coordinate nodes in an asynchronous way. This property can be applied to prove that information existed when the timestamp had been created.
RESILIENCE	Because DLT are distributed systems made by a reasonable high number of nodes, they are highly resilient systems: each node stores all the history of transactions solving the single points of trust and failure problems affecting, by definition, any centralised system. This applies also to physically segregated systems as there can always be the possibility for insider threats. For instance, by outsourcing parts of a digital identity management infrastructure to an independent network such as a public blockchain, DLT systems structurally reduce the attack surface available to malicious insiders.
IMMUTABILITY	DLT systems offer by design practically immutable records of their transactions history primarily by employing both cryptographic hashing functions and consensus mechanisms. Indeed, depending on several characteristics of a distributed ledger such as the number of nodes, the type of consensus mechanism and cryptographic hashing functions used, it is considered practically unfeasible also for well-funded adversaries to dispose of the necessary computational resources to hijack a majority of the nodes, corrupt data by changing the whole chain of blocks and rewrite transactions history. Accordingly, DLT systems resistance to quantum computing attacks is currently being researched and developed [28].
NON-REPUDIATION	Non-repudiation is a security service that provides unforgeable evidence that a particular action has occurred [29]. The service provides cryptographic evidence in electronic transactions so that, in case of disputes, it can be used as a confirmation of an action [27].

Table 1: DLT systems main features and quality attributes

Challenge	Public Blockchains	Permissioned DLT
Scalability	Public blockchains cannot process a high number of transactions per second, if compared to centralized systems. The throughput is limited to a few transactions per second. [27]. At the contrary, they have generally no problems in increasing the number of concurrent nodes connected to the network.	Permissioned DLT systems offer better throughput performance in terms of transactions per second. Nevertheless, they suffer of technical constraints in viably increasing the number of nodes above a certain threshold. [30]
Performance	Alongside throughput efficiency, latency in both data transmission and append-only operations can be a limit in some use cases, particularly in the domains of financial services and the Internet of Things.	Permissioned DLT systems usually perform better than Public Blockchain in terms of throughput efficiency and latency. However, they suffer from performance degradation when the number of nodes increases above a certain threshold. [30]
Adaptability	Especially in public blockchains, once the system is operational, it is hard to change configuration parameters such as consensus mechanisms.	This may not be the case for permissioned DLT systems designed to offer pluggable consensus capabilities.
Privacy	Public blockchains are inherently transparent and offer limited privacy usually achieved with pseudonymity. There are however newer public blockchains (i.e. Monero [31]) that are able to keep transactions details private using Zero Knowledge Proof [32] and related cryptographic methods	Some permissioned distributed ledgers are designed to embed a privacy layer to deploy use cases where information must be accessed on need-to-know basis. Nevertheless, there are claims that adding privacy layers on DLT is not straightforward and can add quite a significant amount of complexity [33]
Interoperability	It is still a challenge to have different public blockchain systems to seamlessly communicate with each other efficiently.	It is still a challenge to have different DLT systems to seamlessly communicate with each other efficiently.
Energy efficiency	Public blockchains based on Proof-of-Work are costly to run and burn significant amounts of electricity to secure the network against cyberattacks such as Distributed Denial of Service [27]. More recent public blockchains implementations, however, use different types of consensus mechanisms (e.g. Proof-of-Stake) that do not require high electricity consumption to operate.	Not having a Proof-of-Work consensus mechanism, permissioned DLTs do not have the same amount of energy expenditure as the Public blockchains based on that specific consensus mechanism.
Easiness of use	Interacting with a public blockchain is not usually straightforward for the average user. If we consider also that mistakes are, as well, immutable, we can see how this is still an area that needs improvements.	On a permissioned DLT system, the user interface can be built in a way that it is easier to interact. It should be noted however that also in this case mistakes will remain written on the ledger. Moreover, on permissioned DLT system it should be considered also the complexity of maintaining the network by IT personnel [33]
Transaction cost	Transacting on public blockchains can be expensive as the fees depend on the price of the underlying cryptocurrency or token, which is usually prone to high price volatility and the fee structure of each network, which depends on the incentives conditional on the consensus mechanism (e.g. requiring high transaction speeds results in increased transaction fees).	Many permissioned DLT systems do not rely on an underlying cryptocurrency as an incentive mechanism and, therefore, the cost per transaction is entirely dependent on the use case. However, on permissioned DLT systems, infrastructure setup cost is a fundamental parameter to consider for deployment and maintenance.
Limited storage space	As a consequence of the high number of data replicas (e.g. on a public blockchain network the data is replicated on every full node) and the fact that historical data is not deleted, the amount of data stored on a distributed ledger should be kept to a minimum. This is especially true for public blockchains, because the paid fees are proportional to data volumes.	Also in permissioned DLT systems the stored data on the ledger should be kept to a minimum as the stored data is replicated on different locations and historical data is not deleted, meaning that the storage need will grow over time.

Table 2: DLT systems main challenges

of course. However, we selected the challenges listed below as we maintain that they are mostly relevant in the nuclear sector.

Considering that these challenges have different applicability on public blockchains and permissioned DLT systems, we summarised them in Table 2 with a different description for the two families.

After briefly introducing a series of benefits and challenges of DLT systems, in the next section we will elicit our methodological choices.

3. Methodology – a Proof-of-Concept strategy

Starting from DLT systems' properties, we performed deductive inferences intended to define use cases in the nuclear sector that, in our view, more adequately adapted to such properties. We endorsed this type of methodological approach, because our primary goal was to evaluate in practice the tangible benefits and challenges that would arise by applying DLTs systems' properties to the nuclear sector in view of adding value to nuclear safeguards. In turn, we adopted a Proof-of-Concept [34] strategy for software design to implement use cases' requirements. In fact, both enterprises and institutions commonly use this strategy to test new products and technologies while avoiding putting too much financial effort at stake. Because the technology is new and in continuous evolution, hands-on experience is essential. In this view, the ultimate goal of a Proof-of-Concept strategy is to provide evidence-based recommendations on whether what we are attempting to achieve is actually feasible and could add tangible value to nuclear safeguards business processes.

In both the use cases on containment and surveillance and on the decentralized timestamping of radiation protection's data, we endorsed the MoSCoW method [35]: a categorization method for software design that uses the following modal adverbs to better clarify the requirements' priority: Must, Should, Could and Would. Moreover, considering the nature of the software artefacts to be developed, we are performing design and implementation using the Software Prototyping methodology [36] Short and fast iterations enabled us to quickly test the functionalities that we intended to deploy, and to have a working prototype that can be easily shown to domain experts and other stakeholders involved.

For the digitalization of the radiation passbook, we endorsed LEAN-UX [37], an Agile method for software requirements analysis. Standing for Lean User Experience, LEAN-UX enabled us to define hypothesis statements for each stakeholder type, creating use case scenarios in a bottom-up process and in a language accessible to both business and technical domain experts. In particular, a Radiation Protection Expert colleague at the Joint Research Center and his team members could directly inform use case requirements without leaving much space for our misinterpretation.

After requirements analysis and co-creation, software design and development together with testing and evaluation phases of our approach have been dealt with the endorsement of the software design and implementation methodology named Behaviour Driven Development (BDD) [38] [39]. With BDD we designed and developed sample smart contracts for the digitalization of the radiation passbook, according to the three phases arguments of this methodology: Given (validates the input), When (processes the contents); and Then (prints out the results).

In the next section, we will present three key DLT systems' properties that while they are partly derived from non-strictly safeguards domains of the nuclear sector, they nevertheless add, in our humble view, value to nuclear safeguards.

4. Key DLT systems properties adding value to nuclear safeguards

4.1 Practical immutability for sensors identity management

In nuclear safeguards, several types of equipment and devices such as sensors, cameras and seals are used to ensure that proper containment and surveillance is applied throughout the fuel cycle. Their deployment is crucial to help inspectorates in formulating safeguards conclusions, because their data can provide relevant information on what has happened inside a facility.

Especially in modern times, most of these devices are electronic in nature, can send remotely their data to an inspectorate, in case connectivity is present and the operator agrees to it, and are equipped with hardware security modules capable of performing asymmetric cryptographic operations on the data produced. Asymmetric cryptography [40] is based on the concept of private and public keys: these two keys are mathematically related to each other so that a message encrypted with one key can be decrypted with the other one and vice versa. The private key is secret and stored on a secure memory inside the device, while the public key, as the name suggests, can be shared with anyone. In our case, asymmetric cryptography is used to digitally sign data so that who receives them or retrieves them from the device's memory can verify both data provenance (i.e. the identity of who created the data) and integrity (i.e. the absence of data modification after the signature).

To achieve this outcome, who retrieves the data must know the public key associated to the private key used to sign them. This can be done in two ways: a simple approach prescribes that before installing the device, the inspectorate

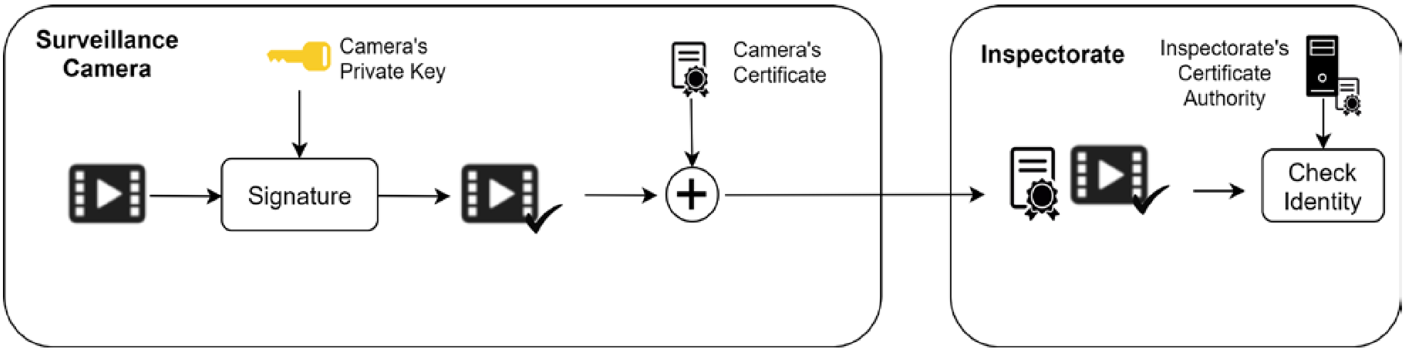

Figure 1: The original video authentication workflow

record on a master list the association between the device's identity and its public key. A more robust approach, currently used, prescribes instead that the public key be inserted into a digital certificate, created and signed by a Certificate Authority, in our case owned and controlled by the inspectorate. This second approach is generally known as a Public Key Infrastructure (henceforth, PKI) [41].

Both these approaches are vulnerable to an insider threat scenario. In fact, if someone is able to modify an entry on the master list, or to issue a valid while non-authorized certificate, s/he can do it so that even a tampered or fake device will have its own entry on the master list or its own valid certificate. If that is the case, data sent from the tampered or forged device will be recognized as authentic, leading to the possibility for a well-funded malicious entity to provide counterfeit data to the inspectorate with the goal of hiding material's diversion or other illicit operations. Moreover, the same person or group of people that has access to the systems used to record relations between a public key and a device's identity (i.e. the master list or the Certificate Authority) could also delete traces of their operations (e.g. log entries). All this could increase the chances that their actions will go indefinitely undetected.

This kind of issue is pervasive to every traditional PKI, having to rely on a Certificate Authority that is trusted by default [42] and consequently it is present in our case as well. In our exploratory research, we initially considered exclusively a specific type of safeguards device, i.e. surveillance cameras. Accordingly, we developed a Proof-of-Concept software implementation to address this specific issue by leveraging on the practical immutability property of data registered on a public blockchain.

More in particular, modern surveillance cameras deployed in nuclear safeguards produce digitally signed video streams. Key pairs are randomly initialized during the camera's setup and the public key is used to create a certificate signing request that, in turn, enables the inspectorate's internal Certificate Authority to generate a certificate. Two different algorithms for asymmetric cryptography are employed: the Digital Signature Algorithm (DSA) is used to digitally sign every single frame, while the Rivest-Shamir-Adleman (RSA) algorithm is used to digitally sign an entire daily video stream. The respective private keys used for the signatures are stored on a secure memory inside the camera, automatically zeroed in case a tampering attempt is detected, whereas the corresponding certificates (containing the public keys also stored within the camera) are embedded in the video stream file itself.

When the inspectorate receives a new video stream, it is therefore able to verify the correctness of each signature and the validity of the certificates (Figure 1). As mentioned above, while very robust, this process is vulnerable to attackers or malicious insiders whom, by having access to the Certificate Authority used to generate the certificates, are enabled to forge new, albeit unauthorized, certificates to be assigned to tampered cameras.

We therefore propose a scenario whereby every certificate (or public key, in case we are referring to the simpler master list approach) is hashed and registered on a public blockchain [43]. In applied cryptography, a hash results from a cryptographic operation that from an input of any size calculates an output of a fixed (usually smaller) size. For each input value there is a single output, but the opposite is not true: from an output, there are multiple (infinite in fact) input values that would lead to it. The hashing function is therefore mathematically irreversible. Thanks to this property of one-way cryptographic hashing functions, a hash therefore is not a sensitive information. Hence, the hash of a certificate can be recorded on a publicly accessible system without the danger to reveal confidential information. Moreover, even a single bit changed on the input causes a huge difference on the output. Such difference, unless the used hashing algorithm has vulnerabilities, is not predictable, so that there is no way other than pure brute forcing to try to obtain a specific output [44].

In other words, if these hashes are recorded somewhere where they cannot be modified (e.g. a blockchain), we can use them to check whether a certificate that we have received is exactly the same one that was initially generated during a surveillance camera's initialisation, by calculating again the hash and comparing it with the registered one.

By registering the hash of issued certificates on a public blockchain, we ensure that such hashes would benefit from the practical immutability that data stored on a public blockchain take on. This in turn ensures that there will always be a trace left of every certificate issuance operation, enabling us to keep this log of information monitored to detect suspicious certificate issuance operations, and leading eventually to flagging the corresponding camera (and its video streams) in order to raise attention.

Indeed, and conversely, if malicious actors could gain control of the inspectorate's internal Certificate Authority to issue a new digital certificate for a tampered camera, they would have to register such operation on the blockchain otherwise the new malicious certificate would not be recognized as valid. By doing so, they leave an undeletable trace of their malicious operation that can be detected.

Also in the simpler master list case, attackers, assuming that they are capable of penetrating and gaining control of the centralized system, could either tamper with or substitute a camera and modify the relative entry on the master list to match the substitution and make it looks like as if the substitution never happened. Storing such information on a blockchain would instead require rewriting part of the transactions history to match a locally stored public key with the one linked to the hash stored on-chain.

For this use case, we are focusing specifically on public blockchains and not on permissioned DLT systems. The rationale behind this decision is based on two considerations: firstly, having underlined that hashes are not sensitive information, from which it is practically unfeasible with current technology to retrieve the original information, we can claim that we do not need confidentiality in this use case. Secondly, the immutability property that data acquires on a blockchain can be considered fully achieved especially on public blockchains. We can argue this because either to alter or to delete stored information, attackers should maliciously operate a majority of the public blockchain's nodes to attempt rewriting transactions history with modified data. This operation requires increasing resources as time passes and blocks accumulate. Notwithstanding how well-funded attackers might be, it is considered practically unfeasible to successfully rewrite the transactions history on public blockchains with a considerably high number of nodes without being detected by the rest of honest nodes. As an example, at the time of writing, on the Bitcoin network nobody has ever been able to perform a successful attack of this sort.

This form of assurance on immutability is generally stronger than the one provided by permissioned distributed ledgers. In the latter case, there is a lower number of nodes and, consequently, data manipulation can occur, in theory, if the participants of the DLT system's network jointly agree to do so [45]. Obviously, such a scenario is hard to image in Nuclear Safeguards, but it could nevertheless happen in principle (e.g. a coalition of malicious states corrupting inspectorates' system administrators). Some permissioned DLT infrastructures partly address this by periodically publishing on a public repository (e.g. on social media, newspapers or websites) a hash that represents an anchor to their private data [26]. In this way, even in the case where all the participants to the permissioned DLT jointly decided to alter data stored in the distributed ledger, there would still be an unambiguously identifiable mismatch with the published hashes. While this solution is certainly valuable, especially in the case of confidential data, where it would be not possible to use a public blockchain, in our specific case it would represent too much of a burden considering that we do not need confidentiality. More importantly, by choosing to adopt a public blockchain, we do not have costs and other organisational issues related to the infrastructure setup and maintenance, because those public blockchain networks that we analysed are already existing and there are no barriers of entry.

To showcase this approach, and initially focusing on the simpler public keys master list case, we created three implementations using three different public blockchains, i.e. Bitcoin [12], Ethereum [22] and Algorand [46] for managing surveillance cameras' digital identities. Selection criteria considered three facets: (1) overall security of the blockchain network, (2) flexibility in terms of possibility to add features other than merely storing information (e.g. the possibility to authorize only specific accounts to store either certificate or public key's hashes), and (3) cost associated to transactions fees. We therefore identified the following public blockchains according to these rationales: the Bitcoin network has not been disrupted from its inception to the time of writing. This makes it the oldest and most secure public blockchain available for experimentation. Secondly, Ethereum offers one of the most widely used platforms for smart contracts deployment. Thirdly, Algorand has very low transaction fees and it enables to provide signature delegation capabilities for smart contracts. This last property is relevant to be tested for the use case at hand, because it could enable inspectorates from regulatory agencies such as IAEA and EURATOM to share responsibilities when installing common infrastructure.

A generalized architecture of the proofs-of-concept implemented with these three blockchain networks is depicted in Figure 2

The architecture spans across three loci: the inspectorate premises, the nuclear facility and the connection channels

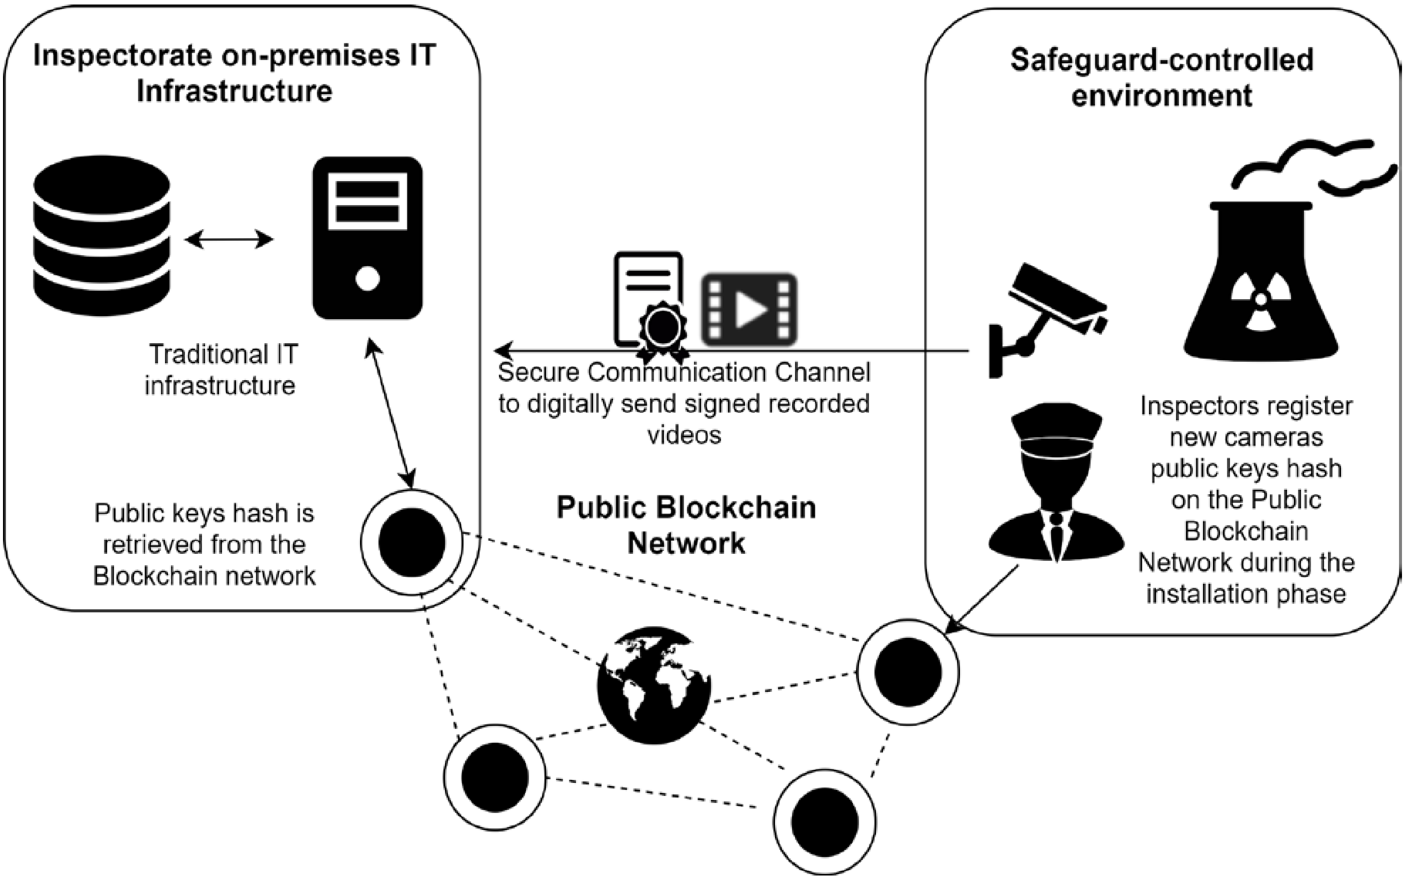
Inspectorate on-premises IT Infrastructure
Traditional IT infrastructure
Public keys hash is retrieved from the Blockchain network
Secure Communication Channel to digitally send signed recorded videos
Public Blockchain Network
Safeguard-controlled environment
Inspectors register new cameras public keys hash on the Public Blockchain Network during the installation phase

Figure 2: General high-level architecture of the process monitoring use case (source [43])

between the two. The novelty, with respect to the traditional master list approach, lies on the utilization of a public blockchain network to store hashes of a camera's public keys.

First, in the Bitcoin proof-of-concept implementation, we proceeded accounting for the protocol constraint that a maximum of 83 bytes of arbitrary data can be recorded in a Bitcoin transaction. While they are not sufficient for storing DSA and RSA public keys, they are enough for storing the public keys' hashes. The transaction identifier is a pointer to the Bitcoin blockchain where it is possible to locate the camera identifier and the hash of the corresponding public key. The drawback is that this adds the small overhead to set up a local database for managing transaction pointers, because searching the Bitcoin blockchain to retrieve public keys' hashes without a pointer would greatly reduce performance. It is nevertheless important to note that such transaction pointers' database is not a single point of failure, because it can be rebuilt in case of data loss for analysing the blockchain.

Secondly, on Ethereum, it is possible to implement the same model by using smart contracts. Contrary to Bitcoin, there is no need to set up a local database to store transaction pointers, as the smart contract is natively equipped with storage capabilities that enables quick information retrieval. However, compared to a traditional software, the implementation phase needs to put more emphasis on the testing part, as fixing bugs on smart contracts is more complicated. Indeed, once deployed a smart contract cannot be modified but must be deactivated and redeployed. In all cases, by virtue of the remarkable volatility of native cryptocurrencies on both Bitcoin and Ethereum networks, transaction costs can change swiftly. Volatility and transaction costs are thus important elements to take into account when defining added value and business viability of this solution at scale.

Finally, as a response to high volatility and transaction costs to curb the effect of highly volatile cryptocurrencies, Algorand can be an appropriate candidate as it offers a very low transaction cost of 0,0002 USD per transaction with storage capacity of up to 1 Kb of data per transaction. However, Algorand's smart contract semantics is less powerful than Ethereum's, but thanks to it, it also reduces the possibility that bugs are introduced in the system. The main drawback with Algorand is that storage capacity must grow quicker over time, if compared to Bitcoin and Ethereum.

In summary, the simple proof-of-concept implementations that we developed on the Bitcoin, Ethereum and Algorand public blockchain networks enabled us to confirm that the idea of using the immutability property of public blockchains to store hashes related to digital identities of safeguards sensors is correct and implementable. Such result shows that DLT not only can have a role in Nuclear Material Accountancy, as shown by the referenced research on DLT and Safeguards, but also in Containment and Surveillance.

At this stage we are not proposing a specific public blockchain network for further implementations as that choice would be dependent on specific application requirements that are usually formulated on a more advanced implementation stage (i.e. a pilot project). In case of further developments, it is indeed of the utmost importance to clearly identify and detail, together with the relevant stakeholders, what are the requirements and constraints that such a system should consider (e.g. total number of sensors, average number of maintenance operations in a sensor lifetime that require the reissuing of a certificate, advanced features like signature delegation, etc.). The final decision on the specific blockchain network to use shall be derived considering such detailed requirements and constraints.

Potential further applications in the containment and surveillance domain can be summarised as follows:

- Not only surveillance cameras, but all digital sensors embedding strong security features (e.g. fiber-optic seals, laser scanners and in principle all sensors used in safeguards) can benefit from the added security of their identities by applying what has been described in this section.
- Analogue sensors could also benefit from it as long as it is possible to derive a “unique” signature of the sensor via signal analysis. If that is possible, what has been described in this section can be applied by registering a hash of the signature on a public blockchain.
- Ad-hoc scripts/libraries could be developed for custom sensors to offer the identity registration on the blockchain as an additional feature.

4.2 Data anchoring through decentralised timestamping

A second property of DLT systems applied to nuclear safeguards is decentralised timestamping for data anchoring on public blockchains.

Data anchoring [47] refers to the process of taking every piece of meaningful data, calculate its hash and publish it on an immutable timestamped repository (e.g. a public blockchain). By doing so, without disclosing any sensitive information thanks to the nature of the hash (as explained in section 4.1), we have a secure way to check if a specific datum was modified simply by recalculating its hash and comparing it with the registered one. It is important to note that other than the immutability of the hash, also its associated timestamp is relevant in this process as it gives a precise indication on when that hash was registered and therefore from which point in time we can speculate on the integrity of the underlying data.

Applying this process to safeguards enables us to generate proofs of existence for nuclear safeguards data, files or events, i.e. that data existed at the time when the timestamp had been created and was not modified ever since. The added value of this process lies on the possibility to check with certainty if information has been modified. In particular, this feature not only provides an additional internal security measure for an inspectorate (e.g. as a way to check integrity of backups, archives and whatsoever relevant data), but also an increased layer of transparency to prove to external parties that the data used to draw safeguards conclusions have not been tampered (e.g. in case of disputes).

While it could be argued that integrity of digital data can be proved also with digital signatures, it should be noted that the difference in the two approaches (i.e. data anchoring on public blockchain versus digital signature) lies on the timestamp. Timestamping of data per se is not a novelty and can be already implemented by a Time Stamping Authority as it happens within the digital signature context. Considering that the timestamping is coming from a third party, there is however the concrete possibility that a Time Stamping Authority could make a mistake or misbehave, thus providing an incorrect timestamp [48].

On a public blockchain, by contrast, each new block is timestamped when it is created. Such timestamp cannot deviate from real time, because it must be temporally situated strictly after the previous block, but not too far in the future. For instance, on the Bitcoin blockchain a new block is discarded, even where formally valid, if its timestamp points to a time situated more than two hours after the latest block [21]. As a downside however, we need to consider that such timestamps cannot be considered extremely precise (e.g. on Bitcoin we should consider a timeframe of 2 hours uncertainty) and therefore they cannot be used for applications where timestamping precision is fundamental.

To explore the feasibility of the decentralised timestamping to ensure data integrity, in the domain of radiation protection (our first use case in this field) we explored the OpenTimestamps protocol to notarise dosimetry data stored on legacy systems at the Joint Research Center, i.e. the Unified Dosimetry System (henceforth, UDS). The OpenTimestamps protocol was firstly proposed by Peter Tood in 2012 [49]. It is an attempt at standardising and solving the scalability and cost issues of timestamping on a public blockchain. Merely notarising every single data element on the

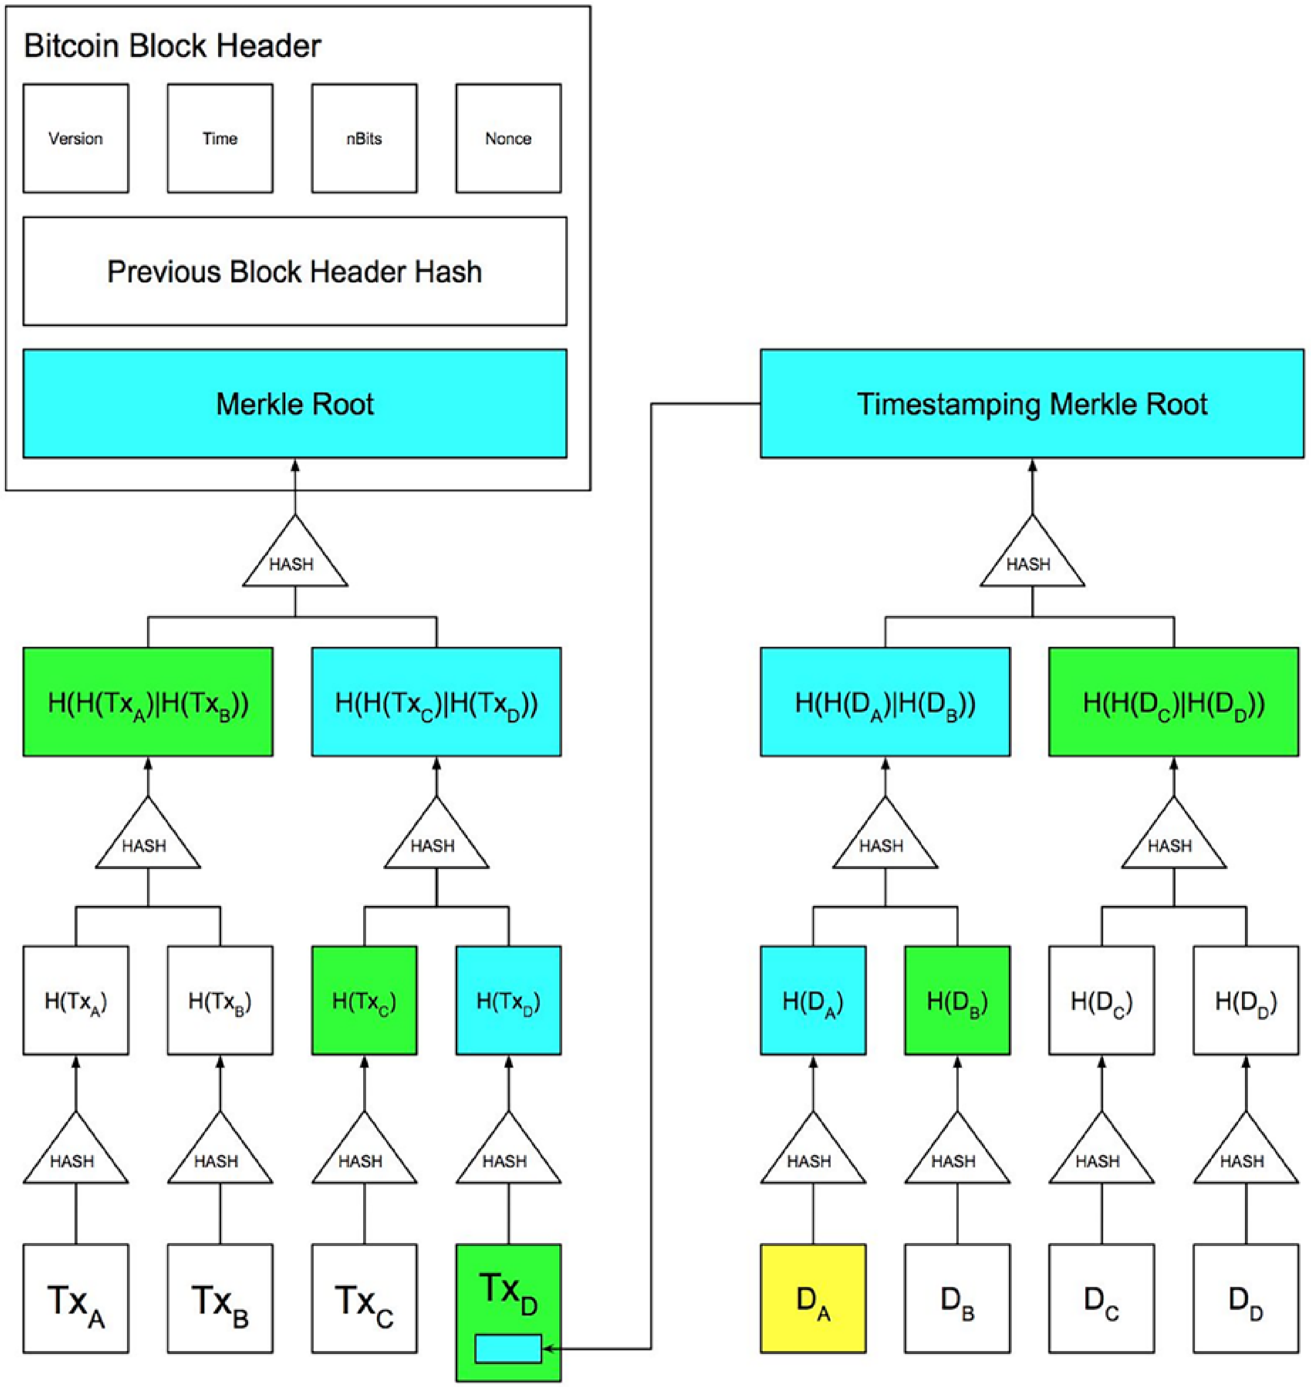

Figure 2: OpenTimestamps protocol: highlighting the path from the Document A (DA) up until the Bitcoin block. In green are highlighted the information contained on the registration proof, which enable to verify mathematically, together with the original data (in yellow) that the information has indeed been timestamped on the Bitcoin blockchain (source: [52])

blockchain, in fact, would lead to the creation of a high number of transactions that are both expensive (taken as a whole) and may clog the blockchain itself by increasing its size too much and by raising the average fee for a transaction to be committed. [50].

The idea proposed by the OpenTimestamps protocol is simple and efficient: instead of timestamping and registering on the blockchain individual hashes, the protocol aims at creating Merkle Trees of data hashes, registering only the root element of the tree on the blockchain as depicted in Figure 3. In cryptography, a Merkle Tree [51] is a data structure used for data verification. It is a binary tree where each leaf node (i.e. all nodes that do not have any child) contains the hash of some data block, and each non-leaf node contains the hash of its child nodes, up to the root element of the tree (i.e. the only element that has no parents).

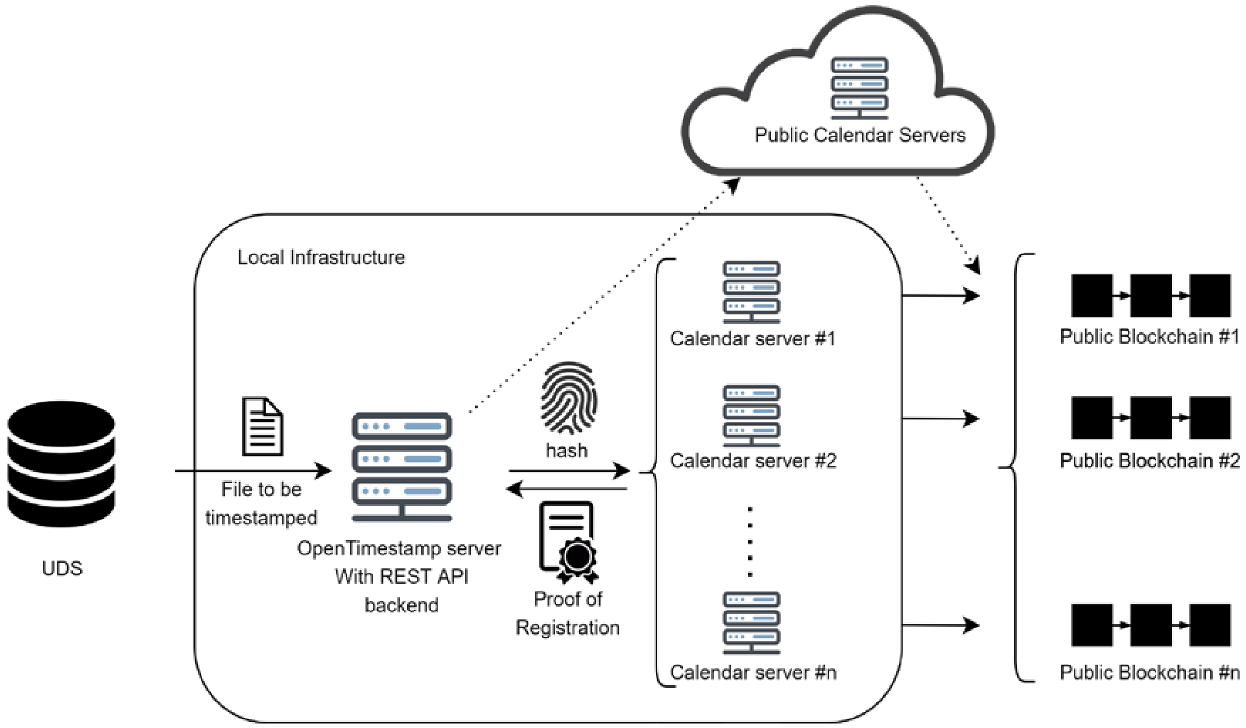

Figure 4: High-level view of how an OpenTimestamps architecture connected to Unified Dosimetry System

The properties of both hashes and Merkle trees enable anyone to verify mathematically that the tree's root element stored on the blockchain has been derived taking into consideration also the hash of a specific datum: to do so it is fundamental to have the list of operations applied starting from the original data to the root element.

To be useful this timestamping process requires that two preconditions be met:

- The original timestamped data must be available over time as is, with absolutely no modifications.
- The proof of registration must be stored together with the original data as the validity of the timestamp is confirmed only by having both of them.

If both requirements are satisfied, who requested the timestamp can mathematically prove the existence of that particular data at the timestamp calendar date.

The required infrastructure is composed at minimum only by a machine to implement the client-side of the OpenTimestamp protocol.

Such client will:

1. Calculate the hash of the data to be timestamped. To be more precise, such hash will be modified by appending a nonce (i.e. a random number) and rehashing the result to avoid involuntary information exposure on the original data, but for simplicity we can consider it as being a simple hash of the data.
2. Send the hash to a calendar server, which is a server responsible of aggregating multiple timestamping requests by generating a Merkle Tree, and of ultimately registering the tree's root element on the blockchain.
3. Receive back from the calendar server the information related to the operations performed starting from the original request that leads to the Merkle Tree root inserted on the blockchain transaction and up until the block header of the blockchain block. This information will be used together with the list of operations performed to generate the data's hash (see 1st step) to generate the proof of registration: a file that lists all operations performed to be stored together with the original data.
4. Independently verify, starting from the original data and the registration proof, if indeed by repeating all the operations we can confirm that there is a block on the blockchain that proves such data existed when the block was created.

As briefly explained, to perform all these operations, the client needs to interface itself with two other servers: a calendar server for the registration operations and a blockchain node (i.e. a machine which has downloaded the complete blockchain and keeps it synced through the consensus algorithm) for the verification.

To be completely secure from any man-in-the-middle attack, it is not recommended to use third party services, but instead to deploy an owned blockchain node. In fact, if the verification phase is performed through an external blockchain monitor (i.e. a webservice provided by a third party that permits to query easily the blockchain without having to download it), nothing can ensure us that the verification itself is not tampered.

The calendar servers instead could still be public (i.e. offered by third parties) considering that (1) they do not receive the original data and (2) the received proofs are independently verified. A private calendar server can nevertheless be deployed in case it is desirable to be completely independent from third party's services (e.g. for business continuity reasons).

The OpenTimestamps protocol, born with the Bitcoin blockchain, nowadays works also with other public blockchains. It is therefore possible to publish the same proof of existence on different blockchains to make it even more tamper resistant. To do so, it is necessary to add additional calendar servers, each one devoted to a particular public blockchain. Obviously, also for the verification part, one node for each blockchain used needs to be locally present to guarantee a better security.

In Figure 4, it is depicted how a high-level architecture of the system would be:

This approach enables the sustainable use of a public blockchain as a reliable while decentralised timestamping authority.

Further potential applications of data anchoring could be:

- All relevant data, declarations, reports, which an operator or the inspectorate might need to be able to prove their existence at a given time and un-alteration ever since, are suitable for this approach.
- Secure software update (e.g. of containment and surveillance devices) by timestamping executables.

Among these possible further applications, we will initially focus on the application of this approach for data integrity in the context of nuclear material accountancy. Indeed, individual nuclear facilities may voluntarily transmit reports or other form of information to EURATOM and IAEA via mailbox. There are cases, however, where this additional information is not sent directly to the inspectorates but must remain on the facility's premises to be retrieved manually by the inspectors. In this case, this type of information on nuclear material accountancy would add value to nuclear safeguards as it could be timestamped following the same steps of the process described above for radiation protection data anchoring to ensure that they are not modified once inserted in the system.

4.3 Structural auditability of dosimetry data

While the previous two properties of DLT systems that can add value to nuclear safeguards emerged from the domain of public blockchains, we inferred a third added value from another property of DLT systems, i.e. structural auditability of data on a permissioned distributed ledger. Also known as consortium blockchains, permissioned DLT systems enable the creation and the broadcasting of transactions only by nodes that have permissions to write new blocks on the ledger. The requirement that only authorised parties can manage certain types of data after permission is granted is normal practice in the use case on radiation protection data management that we will analyse in this section.

As we argued above, if compared to public blockchains, permissioned DLT systems can count on a definition of immutability only in the limited sense that it is difficult for a participant to modify data on all the nodes. There could be, however in principle, a potential collusion risk if all participants jointly decided to modify transactions history recorded on the permissioned distributed ledger.

Contrary to public blockchains designed for environments dominated by high distrust among participants usually operating with pseudonyms, deployment of permissioned distributed ledgers can nevertheless be beneficial in all the cases where higher levels of trust than those characteristic of public blockchains already exist among participating stakeholders (e.g. between nuclear operators and inspectorates). Moreover, permissioned distributed ledgers are recommended when stakeholders also share the requirement to preserve confidentiality of data exclusively accessible by those with permission.

In the use case scenario analyzed in this section, i.e. the digitalization of the radiation passbook on a single shared infrastructure, there is indeed already present a significant level of trust among clearly identifiable stakeholders: an authority issuing a radiation passbook, Radiation Protection Experts at every site, Medical Physics Experts and radiation protection workers. In our scenario, RPEs are the only actors with permission to create new blocks and write digitally signed transactions on the distributed ledger. In this specific use case, confidential data are directly written and shared on the ledger, and this is the reason for the selection of a permissioned DLT system and not a public permission-less one. In other words, in the test for the use case on containment and surveillance, because no confidential information was shared, a public distributed ledger was selected. While in this use case on radiation protection, a permissioned distributed ledger enabling

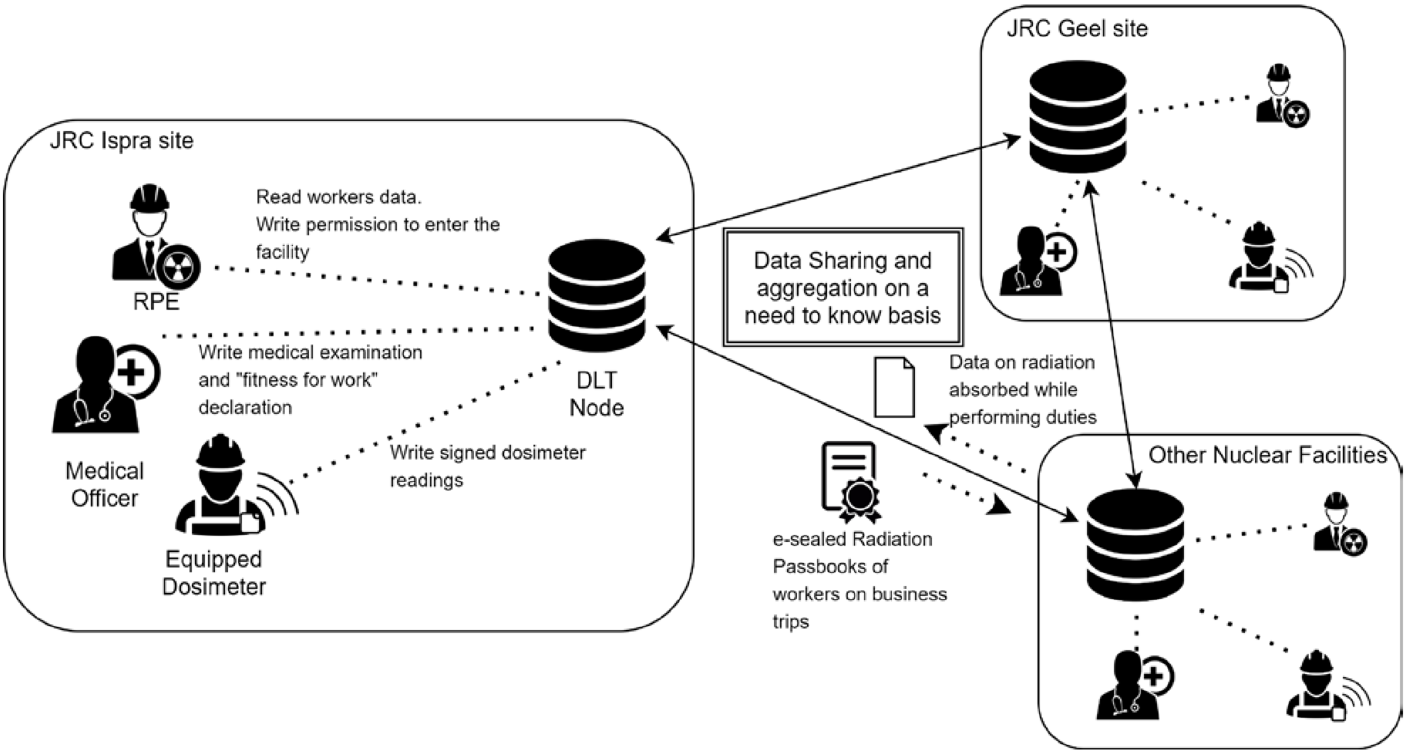
JRC Ispra site
Read workers data. Write permission to enter the facility
RPE
Write medical examination and "fitness for work" declaration
DLT Node
Medical Officer
Write signed dosimeter readings
Equipped Dosimeter
JRC Geel site
Data Sharing and aggregation on a need to know basis
Data on radiation absorbed while performing duties
e-sealed Radiation Passbooks of workers on business trips
Other Nuclear Facilities

Figure 5: High-level schema of the digitalisation of the radiation passbook

confidentiality (also providing immutability as for every DLT system) by design, was the correct technology selection outcome.

By virtue of their network and data structure, permissioned distributed ledgers embed auditability and non-repudiation properties, because all parties with access permission can write and read information on a need-to-know basis. Structural auditability is a property of DLT systems resulting from ordering and timestamping signed dosimetry data transactions that cannot be repudiated. The added value is threefold: (1) higher efficiency in information sharing; (2) lower number of clerical errors thanks to going paperless through smart contracts' automation; and (3) creation of forensic evidence that can help solve disputes on legal liability. As we will discuss below, these three added values emerged from a use case in the nuclear sector of radiation protection could also be applied to nuclear safeguards, specifically in the domain of nuclear material accountancy.

In particular, the radiation passbook is a paper booklet of the dimension of a conventional passport. It lists fields to record employees and inspectors' personal data, and a list of approved compilers, i.e. Radiation Protection Experts with their signatures. Moreover, the radiation passbook registers an employee's occupational exposure to radiations and any involvement in accidents. In turn, it records the 5-year dose limit, the dose assessment for the calendar year; the estimated doses in mSv in another employer's-controlled area(s); Whole Body Count; Radiotoxicologial monitoring; medical examinations; fitness-for-work in normal and in case of arduous conditions; and radiation protection training. A final section includes important addresses and telephone numbers, specifically Headquarters, Medical Service and Radiation Protection Experts.

The choice to explore DLT systems' applications for the radiation passbook was initiated by acknowledging the fact that nuclear safeguards workers travel to different locations, where they are potentially exposed to radiation. However, data related to absorbed doses during missions are rarely shared between their employer and the nuclear installation. Indeed, workers carry two separate dosimeters: exposure to ionising materials is measured using both their employer's and nuclear installations' dosimeters. Using a permissioned DLT system, it is possible to digitally record and share dosimetry data to track radiation exposure of workers as depicted in Figure 5:

As a hands-on proof-of-concept exercise, we implemented sample smart contracts using Zenroom [53], an output implementation funded by DG CNECT to research distributed applications for compliance with the EU General Data Protection Regulation [54]. Zenroom is a trusted execution environment with no external dependencies. It comprises a tiny library (1Mb) and requires low memory usage (600 Kb – 2 Mb). It runs smart contracts written with Zencode, a

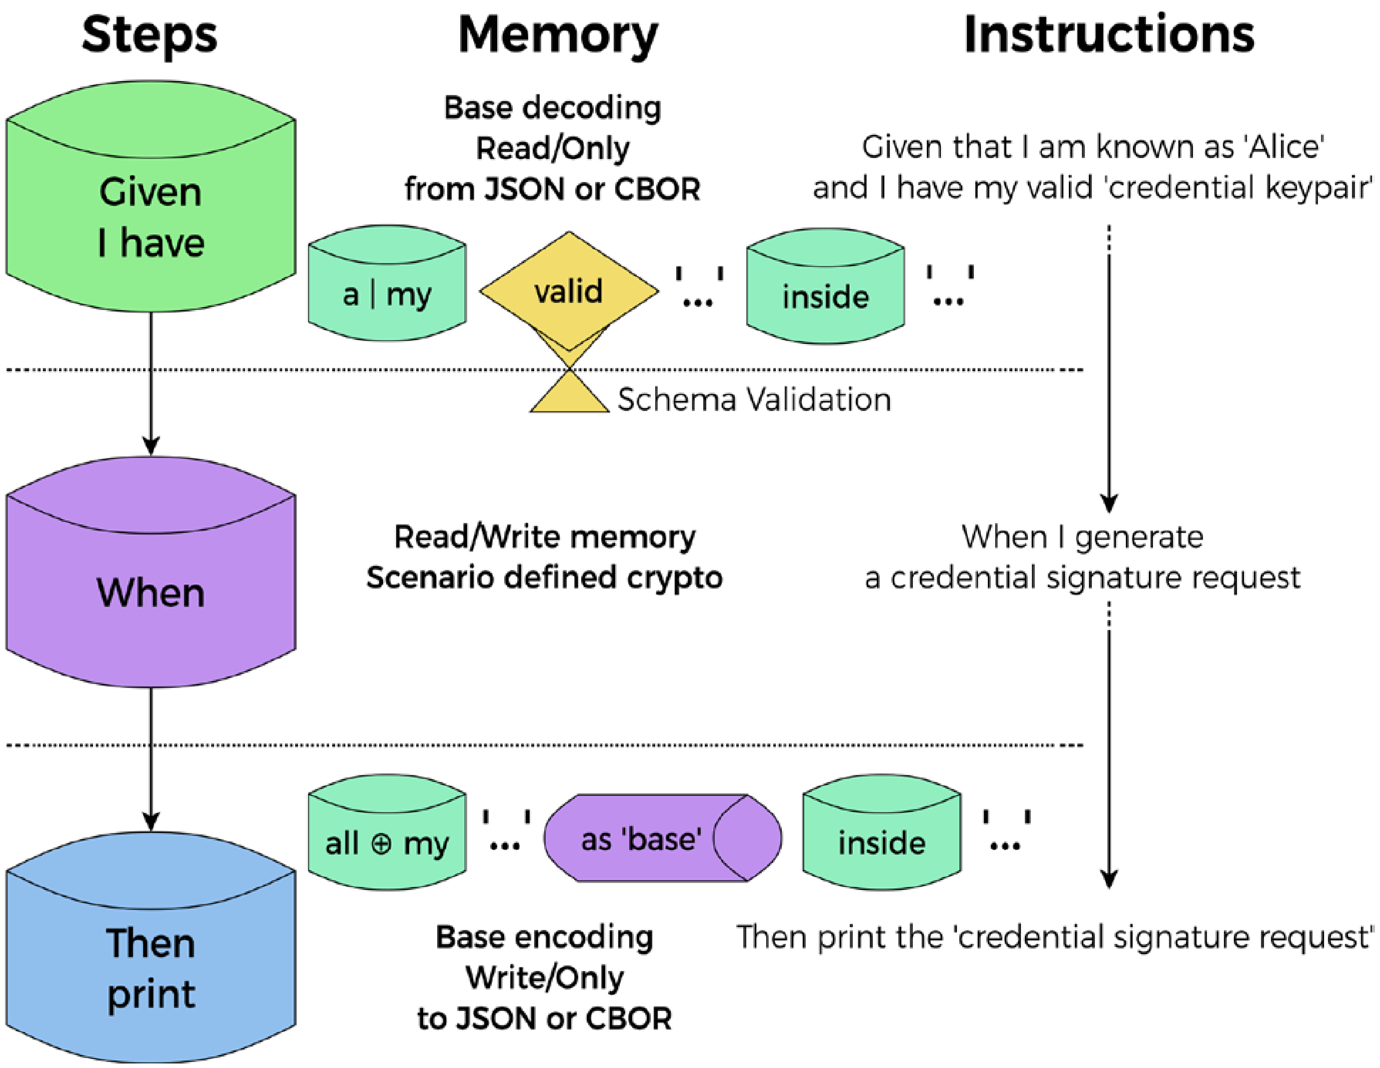

Figure 6: High-level schema of the Zenroom virtual machine's memory model (source [37]).

human readable, English-like Domain Specific Language [29]. Zencode language design followed the principles of language theoretic security [55] formally verifying that smart contracts are free of currently dominant classes of bugs and vulnerabilities.

This design approach implements a memory model based on Behaviour Driven Development that segregates from external calls input (Given), processing (When) and output (Then) phases of a smart contract's encoding and execution. The result is a highly efficient process virtual machine, which provides a powerful smart contracts semantics less prone to unexpected data change and control flow bugs [56], i.e. the classes of bugs that led, for instance, to the 2016 'DAO hack' on the Ethereum blockchain [38]. Figure 6 depicts the Zencode processing memory model:

In our implementation, deterministic smart contracts have been encoded to issue a digitalized version of the radiation passbook authenticated through electronic sealing by means of Elliptic Curve Digital Signature Algorithm (ECSDA) digital signatures. Moreover, we implemented smart contracts to process digital signatures in order to provide RPEs with the means to proof dosimetry data provenance, non-repudiation and ensure confidentiality. Signed dosimetry data can thus be stored on a permissioned distributed ledger.

We ran a test using Zenroom integrated as a transaction processor on the Hyperledger Sawtooth permissioned distributed ledger. Hyperledger Sawtooth offers the flexibility for pluggable consensus and smart contract virtual machines [57]. In particular, we implemented a version of Hyperledger Sawtooth v1.0.1 ordering transactions with Practical Byzantine Fault Tolerance [39] consensus, whereby no more than one third of the nodes (rounded down) can be unreachable or dishonest at any given time. In the future, we will aim to complement confidentiality and non-repudiation with a data aggregation module to update dosimetry data history on the radiation passbook each time a nuclear safeguards worker completes a mission.

As also stated by other research institutes in the nuclear sector [7] [8], we also believe that a system similar to the one described for this use case, based on a permissioned DLT system, could be implemented applied to add value to Nuclear Material Accountancy and Control business processes. For instance, using a permissioned DLT system to process nuclear material accounting records could enable near-real time reporting as data, e.g. an Inventory Change Report is shared to relevant stakeholders (strictly on a need-to-know basis) immediately after it is issued. Moreover, structural auditability would provide early data analysis with the goal of detecting red flags soon enough to promptly address potential criticalities. Secondly, smart contracts could automate data validation and reconciliation operations, i.e. transit matching. Finally, a permissioned DLT system would offer a single source of truth to all stakeholders involved, providing also in this case a source of forensic evidence to help address disputes in nuclear material accountancy.

However, we anticipate that DLT systems applications on nuclear material accountancy must be carefully studied and tested, because there are several constraints that could hinder their applicability. While a complete feasibility study on this specific topic is not within the scope of this paper, we report some examples of constraints, also analysed in related work [7], which should be carefully considered when designing permissioned DLT system applications in the context of nuclear material accountancy at an international level:

— International Regulations: since some regulations, for instance EURATOM Regulation 302/2005, details quite specifically how Nuclear Accountancy Reporting must be performed, new systems must accordingly be compatible with the standard procedure listed on this regulation.

— Data locality: some nation states forbid reliance on safeguards related data other than the official accountancy records (e.g. safeguards sensors data, mailbox declarations) from leaving either the country or the facility premises where they are produced. Consequently, a DLT system shall implement measures to provide data locality on some categories or data, or that such system should be used only to provide proof of existence and non-alteration and not to store the actual data.

— Data transmission format: some nation states forbid the usage of electronic communications media for sending accounting records to inspectorates. As stated in the previous point, this could mean that a DLT system should be used only as an additional integrity layer and not to store the actual data.

5. Conclusion and potential way forward

In this paper, we presented three key properties of DLT systems, whose application has been transversal to the nuclear sector in general with special emphasis on added value to nuclear safeguards: practical immutability, data anchoring through decentralised timestamping for public blockchains and structural auditability for permissioned DLT systems. Practical immutability adds value by increasing the level of cybersecurity, potentially impacting inspections plans thanks to a more efficient automation of safeguards business processes for containment and surveillance, specifically for surveillance cameras' Public Key Infrastructure management.

Moreover, decentralised timestamping and structural auditability, albeit tested in radiation protection, can add value to nuclear safeguards by lowering clerical errors through enhanced automation and consistency of business processes by virtue an additional layer of both data integrity and a real-time audit trail of transactions' history. These two key DLT systems' properties can also be seen as innovative sources of forensic evidence for legal dispute resolution not only in radiation protection but also in the domain nuclear material accountancy.

Future research aims at further testing these findings on the Experimental Infrastructure for Internet Contingencies (EPIC), maintained by JRC-Ispra E.3 Cyber and Digital Citizens' Security Unit. EPIC enables the re-creation of cyber-infrastructures for testing various configurations. It provides special physical equipment, such as a Programmable Logical Controller, enabling cyber-physical testing (max 356 nodes). These characteristics give significant advantages in terms of repeatability, scalability and controllability of experiments and tests.

We plan to run performance experiments on fiber-optic seals for containment and surveillance, and regular mailbox declarations for nuclear material accountancy. We will then plan to emulate the performance of a permissioned DLT system to process the digitalized radiation passbook by Radiation Protection Experts at both JRC sites and from nuclear installations. The goal of these tests is to build robust datasets for performance comparison between DLT-based systems and more traditional approaches. The objective is to generate quantitative metrics to more finely evaluate the benefits and added value of DLT systems applied to nuclear safeguards. The overarching ambition of these experimental exercises within the SLT4SFG exploratory research project is to define a general methodology to select key DLT properties for their applications to nuclear safeguards.

6. Acknowledgements

The authors would like to acknowledge all those who have reviewed previous drafts of this paper and contributed to its completion together with JRC-Ispra Units A.5 (Scientific Development), G.ii.7 (Nuclear Security), E.3 (Cyber and Digital Citizens' Security) and R.I.5 (Ispra site Safety & Security) for enabling, operationalising and supporting the exploratory research project Shared Ledger Technologies for Nuclear Safeguards.

7. References

[1] W. Janseens, "Nuclear Safeguards Challenges from a JRC Perspective," in International Cooperation for Enhancing Nuclear Safety, Security and Non proliferation - 60 Years of IAEA and EURATOM, Rome, 2018.

[2] European Commission, "European countries join Blockchain Partnership," [Online]. Available: https://ec.europa.eu/digital-single-market/en/news/european-countries-join-blockchain-partnership. [Accessed 14 May 2021].

[3] EU Blockchain Observatory and FOrum, "EU Blockchain Observatory and Forum," [Online]. Available: https://www.eublockchainforum.eu/. [Accessed 14 May 2021].

[4] European Commission, "EBSI - Experience the future with the European Blockchain Service Infrastructure (EBSI)," [Online]. Available: https://ec.europa.eu/cefdigital/wiki/display/CEFDIGITAL/ebsi. [Accessed 14 May 2021].

[5] European Commission, "European Blockchain Pre-Commercial Procurement," [Online]. Available: https://ec.europa.eu/digital-single-market/en/news/online-questionnaire-european-blockchain-pre-commercial-procurement-live. [Accessed 14 May 2021].

[6] S. Nonneman, G. Renda, L. Dechamp, M. Isabella, T. Jacobi and I. N. Fovino, "Distributed Ledger Technology in Nuclear non-proliferation Safeguards?," in Symposium on International Safeguards, Stresa, 2018.

[7] C. Vestergaard, G. Green, E. G. Obbard, E. Yu and G. D. Putra, "SLAFKA Demonstrating the Potential for Distributed Ledger Technology for Nuclear Safeguards Information Management," Stimson Center, 2020.

[8] S. Frazar, C. Joslyn, R. Goychayev and A. Randall, "Developing an Electronic Distributed Ledger For Transit Matching," in Proceedings of the INMM 61st Annual Meeting, Baltimore, MD, USA, 2020.

[9] N. Pattengale, "DLT Activities at Sandia National Laboratory (Presentation)," Sandia National Laboratory, 2020.

[10] L. Umayam and C. Vestergaard, "Complementing the Padlock: The prospect of blockchain for strengthening nuclear security," Stimson Center, 2020.

[11] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich and S. Muralidharan, "Hyperledger fabric: a distributed operating system for permissioned blockchains," in Euro Sys '18: Proceedings of the Thirtheenth EuroSys Conference, 2018.

[12] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2009. [Online].

[13] M. Rauchs, A. Glidden, B. Gordon, G. Pieters, M. Recanatini, F. Rostand, K. Vagneur and B. Zhang, "Distributed Ledger Technology Systems - a Conceptual Framework," Cambridge Centre for Alternative Finance, University of Cambridge, Judge Business School, 2018.

[14] N. Ferguson, B. Schneier and T. Kohno, Cryptography Engineering Design Principles and Practical Applications, Indianapolis: Wiley Publishing, Inc., 2010.

[15] B. Schneier, Applied Cryptography: Protocols, Algorithms and Source Code in C, John Wiley & Sons, Inc., 1995.

[16] G. J. Simmons, "How to Insure that Data Acquired to Verify Treaty Compliance are Trustworthy," in Proceedings of the IEEE, 1988.

[17] D. Chaum, "Blind signatures for untraceable payments," in Advances in Cryptology Proceedings 82, 1983.

[18] D. Chaum, A. Fiat and M. Naor, "Untraceable electronic cash," in Advances in Cryptology - CRYPTO '88 Proceedings, New York, 1988.

[19] A. Back, "Hashcash - A Denial of Service Counter-Measure," 2002.

[20] C. Dwork and M. Naor, "Pricing via Processing or Combatting Junk Mail," in Advances in Cryptology - Crypto '92, 1992.

[21] A. Antonopoulos, Mastering Bitcoin, O'Really, 2015.

[22] G. Wood, "Ethereum: A Secure Decentralised Generalised Transaction Ledger," 2014. [Online]. Available: gavwood.com/paper.pdf.

[23] L. Luu, D.-H. Chu, H. Olickel, P. Saxena and A. Hobor, "Making Smart Contracts Smarter," in Proceeding of

the 2016 ACM SIGSAC Conference on Computer and Communications Security, 2016.

[24] U. Sarfraz, M. Alam, S. Zeadally and A. Khan, “Privacy aware IOTA ledger: Decentralized mixing and unlinkable IOTA transactions,” Computer Networks 148, pp. 361-372, 2019.

[25] L. Baird, M. Harmon and M. Paul, “Hedera: A Public Hashgraph Network & Governing Council,” 2019. [Online]. Available: https://hedera.com/hh-whitepaper-v1.5-190219.pdf.

[26] A. Buldas, L. Risto and A. Truu, “Keyless signature infrastructure and PKI: hash-tree signatures in pre- and post-quantum world,” International Journal of Services Technology and Management, pp. 117-130, 2017.

[27] I. Bashir, Mastering Blockchain: A deep dive into distributed ledgers, consensus protocols, smart contracts, DApps, cryptocurrencies, Ethereum, and more, Birmingham: Packt Publishing, 2020.

[28] T. M. Fernandes-Camares and P. Fraga-Lamas, “Towards Post-Quantum Blockchain: A Review on Blockchain Cryptography Resistant to Quantum Computing Attacks,” IEEE Acess, vol. 8, pp. 21091-21116, 2020.

[29] E. A. Brewer, Towards robust distributed systems (Invited Talk), Portland, Oregon, 2000.

[30] Q. Nasir, I. A. Qasse, M. A. Talib and A. B. Nassif, “Performance Analysis of Hyperledger Fabric Platforms,” Security and Communication Networks, vol. 2018, 2018.

[31] The Monero Project, “About Monero,” [Online]. Available: https://www.getmonero.org/resources/about/. [Accessed 03 June 2021].

[32] S. Goldwasser, S. Micali and C. Rackoff, “The Knowledge Complexity of Interactive Proof Systems,” SIAM Journal on Computing, vol. 18, no. 1, pp. 186-208, 1989.

[33] L. Feagan, “Hyperledger Fabric Myths and Reality,” April 2020. [Online]. Available: https://objectcomputing.com/resources/publications/sett/april-2020-hyperledger-fabric-myths-and-reality. [Accessed 3 June 2021].

[34] M. K. Pratt, “proof of concept (POC),” [Online]. Available: https://searchcio.techtarget.com/definition/proof-of-concept-POC. [Accessed 14 May 2021].

[35] Wikipedia, “MoSCoW Method,” [Online]. Available: https://en.wikipedia.org/wiki/MoSCoW_method. [Accessed 14 May 2021].

[36] Wikipedia, “Software Prototyping,” [Online]. Available: https://en.wikipedia.org/wiki/Software_prototyping. [Accessed 14 May 2021].

[37] D. Roio and A. Dintino, “Smart Contracts for Data Commons integrated with GDPR compliant legal rules and tested in pilots,” European Commission DG CNECT, 2019.

[38] K. O’Hara, “Smart Contracts - Dumb Idea,” IEEE Internet Computing, vol. 21, no. 2, pp. 97-101, 2017.

[39] M. Castro and B. Liskov, “Practical Byzantine fault tolerance,” in OSDI ‘99: Proceedings of the third symposium on Operating systems design and implementation, 1999.

[40] M. E. Hellman, “An Overview of Public Key Cryptography,” IEEE Communications Magazine, 2002.

[41] C. Adams and S. Lloyd, Understanding public-key infrastructure: concepts, standards, and deployment considerations, 1999.

[42] J. A. Berkowsky and T. Havaineh, “Security Issues with Certificate Authorities,” in IEEE 8th Annual Ubiquitous Computing Electronics and Mobile Communication Conference (UEMCON), 2017.

[43] R. Spigolon, M. Sachy, S. Nonneman, R. Neisse, I. Maschio and I. Nai Fovino, “Using Public Blockchain for the management of Public Key Infrastructure to strengthen Nuclear Safeguards,” in Proceedings of the INMM 61st Annual Meeting, Baltimore, MD, USA, 2020.

[44] R. Sobti and G. Geetha, “Cryptographic hash functions: a review,” International Journal of Computer Science Issues (IJCSI), 2012.

[45] B. Putz and G. Pernul, “Trust Factors and Insider Threats in Permissioned Distributed Ledgers,” in Transactions on Large-Scale Data and Knowledge-Centered Systems XLII, vol. 11860, Springer, 2019, pp. 25-50.

[46] S. Micali and J. Chen, “Algorand,” 26 May 2017. [Online]. Available: https://algorandcom.cdn.prismic.io/algorandcom%2Fece77f38-75b3-44de-bc7f-805f0e53a8d9_theoretical.pdf.

[47] O. Konashevuch and M. Poblet, “Blockchain Anchoring of Public Registries: Options and Challenges,” in ICEGOV2019 Proceedings of the 12th International Conference on Theory and Practice of Electronic Governance, 2019.

[48] M. Vigil, D. Cabarcas, J. Buchmann and J. Huang, “Assessing trust in the long-term protection of

socuments," in 2013 IEEE Symposium on Computers and Communications (ISCC), 2013.

[49] P. Tood, "OpenTimestamps: Scalable, Trust-Minimized, Distributed Timestamping with Bitcoin," 15 September 2016. [Online]. Available: https://petertodd.org/2016/opentimestamps-announcement. [Accessed 6 November 2020].

[50] E. Strehle and F. Steinmetz, "Dominating OP Returns: The Impact of Omni and Veriblock on Bitcoin," Journal of Grid Computing, pp. 575-592, 2020.

[51] R. Merkle, Secrecy, authentication and public key systems/ A certified digital signature, Dept. of Electrical Engineering, Stanford University, 1979.

[52] R. Casatta, "Scalable and Accountable Timestamping (presentation)," [Online]. Available: https://bit.ly/321YcNy.

[53] Zenroom, "Zenroom," [Online]. Available: https://dev.zenroom.org/. [Accessed 14 May 2021].

[54] European Commission, "Decentralised Citizens Owned Data Ecosystem," [Online]. Available: https://cordis.europa.eu/project/id/732546. [Accessed 14 May 2021].

[55] F. Momot, S. Bratus, S. Hallberg and M. Patterson, "The Seven Turrets of Babel: A Taxonomy of LangSec Errors and How to Expunge Them," IEEE Cybersecurity Development, SecDev 2016, pp. 45-52, 3-4 November 2016.

[56] D. Roio and M. Sachy, "Initial Definition of Smart Rules and Taconomy," European Commission H2020 DECODE Project Deliverable, 2018.

[57] Bitwise IO, Inc , "Sawtooth PBFT," [Online]. Available: https://sawtooth.hyperledger.org/docs/pbft/releases/1.0.1/. [Accessed 14 May 2021].