

# GDPR-Compliant Use of Blockchain for Secure Usage Logs

Valentin Zieglmeier  
Technical University of Munich  
Munich, Germany  
valentin.zieglmeier@tum.de

Gabriel Loyola Daiqui  
Technical University of Munich  
Munich, Germany  
gabriel.loyola-daiqui@tum.de

## ABSTRACT

The unique properties of blockchain enable central requirements of distributed secure logging: Immutability, integrity, and availability. Especially when providing transparency about data usages, a blockchain-based secure log can be beneficial, as no trusted third party is required. Yet, with data governed by privacy legislation such as the GDPR or CCPA, the core advantage of immutability becomes a liability. After a rightful request, an individual's personal data need to be rectified or deleted, which is impossible in an immutable blockchain. To solve this issue, we exploit a legal property of pseudonymized data: They are only regarded personal data if they can be associated with an individual's identity. We make use of this fact by presenting P<sup>3</sup>, a pseudonym provisioning system for secure usage logs including a protocol for recording new usages. For each new block, a one-time transaction pseudonym is generated. The pseudonym generation algorithm guarantees unlinkability and enables proof of ownership. These properties enable GDPR-compliant use of blockchain, as data subjects can exercise their legal rights with regards to their personal data. The new-usage protocol ensures non-repudiation, and therefore accountability and liability. Most importantly, our approach does not require a trusted third party and is independent of the utilized blockchain software.

## CCS CONCEPTS

• Security and privacy → Pseudonymity, anonymity and untraceability; Privacy-preserving protocols; • Computer systems organization → Peer-to-peer architectures.

## KEYWORDS

Blockchain, GDPR, Privacy, Anonymity, Logging

### ACM Reference Format:

Valentin Zieglmeier and Gabriel Loyola Daiqui. 2021. GDPR-Compliant Use of Blockchain for Secure Usage Logs. In *Evaluation and Assessment in Software Engineering (EASE 2021)*, June 21–23, 2021, Trondheim, Norway. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3463274.3463349>

## 1 INTRODUCTION

Blockchains are used for a multitude of applications. They can offer advantages in distributed contexts with untrusted participants,

especially if immutability of data is required [31, 51]. One potential application that benefits from these properties is distributed secure logging. Multiple secure logs based on blockchain were developed in recent years [e.g., 19, 43, 47]. The technology has many advantages for these applications, as it is a lightweight way to guarantee immutability of stored entries (integrity) and functions even in unreliable distributed networks (availability).

Yet, the unique properties of blockchain mean that the confidentiality of stored data cannot be guaranteed by default. Even when storing minimal information, some form of identifier is required to denote ownership or association to entries. Without necessitating any information leak, the blockchain then allows any network participant to trace entries based on their identifier [40]. At best, users can try to hide their association to blockchain entries by keeping their identifier secret and creating new addresses. Even then, network participants can deduce information about users simply by analyzing publicly available data [5]. If the identifier is leaked or known to a third party, though, all respective entries can be retroactively associated with them. This has been identified as a problem of blockchain-based secure logs [19, 43], where confidentiality can be an important property [2]. More critically, recent privacy legislation such as the General Data Protection Regulation (GDPR) [20] of the European Union and the California Consumer Privacy Act (CCPA) [12] require those who store personal data to protect and, on request, even delete it.

One potential solution for this is to use permissioned (private) blockchains. Here, only parties with a legitimate interest in the stored data get access to the blockchain, managing it like any other database [4]. This solves the issue of confidentiality (to a point) and can enable GDPR-compliance, but it removes the key advantage of a permissionless blockchain by requiring trust in those who can write to it [4]. Therefore, in this case there is no advantage of using blockchain compared to any other distributed database [31].

Intuitively, it seems as if this means a fundamental conflict between the requirements of privacy legislation and permissionless blockchain technology. Therefore, the only solution would be not to utilize blockchain when storing data that may be governed by the GDPR or CCPA. We argue that this conflict can be solved differently, though. We want to combine the strengths of blockchain (such as immutability, decentralization, and fault tolerance [19]) with the requirements of privacy legislation (confidentiality, deletability).

*Use Case:* We are motivated by the idea of providing employees with more transparency regarding the usage of their data, e.g. by their employer. Typically when personal data are handled, their usage is covered by privacy policies or company agreements. These policies are hard to read and understand [35], calling into question whether individuals subjected to them truly understand their impact. While some usages of their data might be beneficial to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://www.acm.org).

*EASE 2021, June 21–23, 2021, Trondheim, Norway*

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9053-8/21/06...\$15.00

<https://doi.org/10.1145/3463274.3463349>

employees, giving these data up poses the risk of profiling and misuse. To give individuals more oversight and control in situations of such asymmetric knowledge, various authors have described the idea to track individual usages of data and make the usage log available to data owners. The concept has been discussed both in theory [see, e.g., 3, 11, 49] and in practice [see, e.g., 10, 43, 52, 53]. This of course requires *data owners* to trust that the usage log they are presented is complete and correct, necessitating a tamper-proof logging mechanism [43]. Ideally, a trusted third party such as the employer is not required, as they might be interested to modify the log and potentially remove incriminating evidence. To achieve this, Schaefer and Edman recently proposed utilizing a blockchain to serve as a secure usage log [43].

*Contribution:* Secure usage logs based on blockchain are fundamentally at odds with privacy legislation such as the GDPR. The metadata recorded in the blockchain are themselves personal data that need to be protected [13], an on request rectified or deleted [29]. To enable secure and private logging of usage data with blockchain, we therefore contribute the pseudonym provisioning system  $P^3$ . The system guarantees two properties: Proof of ownership, as required for deletion requests [27], and unlinkability, to provide users anonymity against adversaries [48]. Most importantly, it does not require any changes in the utilized blockchain software.

## 2 PSEUDONYMITY AND ANONYMITY

Two concepts are important to understand when discussing the applicability and implications of the GDPR: pseudonymity and anonymity.

Pseudonymized data are data where identifiers (such as names) have been replaced by pseudonyms, and the relationship between pseudonyms and identifiers is stored separately from the data themselves. As the availability of this link allows re-identification, these data are not anonymous and fall under the provisions of privacy legislation [29]. Anonymized data on the other hand are data that have been modified as such to make it impossible to re-identify an individual from them [23, 29]. Importantly: While pseudonymous data are regarded personal data, anonymous data are not [29]. This means that to fulfill a user’s legal right to erasure, we do not actually need to delete their personal data, as long as we can anonymize it by deleting their link to the pseudonym [27, 29].

## 3 $P^3$ : PRIVATE PSEUDONYM PROVISIONING

We present a pseudonym provisioning system to enable GDPR-compliant and secure data usage logging with blockchain. As we have described above, the same datum can be considered anonymous and pseudonymous, depending on the knowledge of the respective party. In the following, we will therefore deliberate the specific properties that the generated pseudonyms guarantee for each participating party.

First, we define our adversarial model and attacks that we consider. From our use case and the adversarial model, we derive requirements. Then, we describe the  $P^3$  system and discuss how it addresses these requirements.

## 3.1 Adversarial Model

A user  $u$  may be either in the set of *data owners*  $O = \{o_1, o_2, \dots, o_n\}$ , in the set of *data consumers*  $C = \{c_1, c_2, \dots, c_m\}$ , or both. The adversary  $\alpha$  can be any user and assume any role. Whenever a consumer  $c_i \in C$  accesses data of an owner  $o_j \in O$ , a usage  $u_{ij}(c_i \rightarrow o_j)$  is appended to the usage log  $U$  stored in the blockchain.

We assume that  $\alpha$  has limited computational capacity, and can therefore never assume control over the blockchain network. Yet, within their means, they aim for maximum damage and therefore do not “play fair”. As we utilize blockchain, we know that attacks on the integrity of the log are infeasible for  $\alpha$  [51]. Instead,  $\alpha$  is motivated to attack the confidentiality of the stored data by gaining access to data that are not meant to be accessible by them.

Specifically,  $\alpha$  tries to conduct the following attacks:

- (a) Derive from any entry  $u_{ij}(c_i \rightarrow o_j)$  with  $\alpha \notin \{c_i, o_j\}$  the identity of  $c_i$  or  $o_j$ .
- (b) Associate any two entries  $u_{ij}(c_i \rightarrow o_j), u_{ik}(c_i \rightarrow o_k)$  with each other, thereby leaking their association with a single *data consumer*.
- (c) Associate any two entries  $u_{ji}(c_j \rightarrow o_i), u_{ki}(c_k \rightarrow o_i)$  with each other, thereby leaking their association with a single *data owner*.
- (d) Leak the identity of  $c_j$  for a stored usage  $u_{jj}(c_j \rightarrow o_j)$  with  $\alpha = o_j$ , after  $c_j$  has legitimately exercised their right to erasure under the GDPR regarding  $u_{jj}$ .

## 3.2 Requirements

From our use case and adversarial model arise five main requirements: (1) The usage log needs to be tamper-proof, preventing e.g. *data consumers* from removing incriminating entries. (2) No trusted third party is required. (3) *Data owners* can query for arbitrary log entries concerning their data and view their content. Importantly, they can prove the association of the *data consumer* to the logged usage (non-repudiation). (4) *Data consumers* can query for arbitrary log entries concerning their usages and verify their content. (5) No third party can derive their identity or usage information from data stored in the blockchain.

In addition, the data stored in the usage log are governed by the GDPR for as long as they can be associated with the identities of users. From that follows an additional requirement: After a legitimate request, (6a) a user’s right to erasure [20, Art. 17] and (6b) a user’s right to rectification [20, Art. 16] can be fulfilled. We consider both issues a single requirement, as enabling the deletion of personal data indirectly enables rectification, by removing the incorrect entry and adding the rectified version. Furthermore, as we have discussed in Section 2, we can fulfill a user’s legal right to erasure by anonymizing their personal data. Therefore, we arrive at requirement (6): The ability to anonymize the personal data stored in the blockchain as such to make re-identification impossible.

## 3.3 The $P^3$ System


Our concept for the  $P^3$  system consists of four parts: The block structure, the new-usage protocol, the pseudonym generation algorithm, and the deployment architecture. First, the block structure describes how usage data are stored in the blockchain and how they are protected. Second, the new-usage protocol describes the

interaction between *data consumer* and *data owner* when a datum is accessed and the usage is logged. Third, the pseudonym generation algorithm enables the provisioning of private pseudonyms that guarantee unlinkability and proof of ownership. Finally, the deployment determines the required trust and computational resources.

**3.3.1 Block Structure.** First, we define the actual data stored in each block of the blockchain. As we have stressed above, our goal is to not require any changes to the underlying blockchain software, thereby allowing our solution to be used with existing blockchains.

As an example, we consider a new entry logging the usage  $u_{ij}(c_i \rightarrow o_j)$ . Based on requirements (3) and (4), we need to store a one-time pseudonym for both the *data consumer* and *data owner*, to allow each party to query for entries concerning them. We will discuss the properties these pseudonyms guarantee in Section 3.3.3. In addition, both parties need to be able to access the stored usage information, while preventing third parties from reading it, following (5).

As shown in Figure 1, each block therefore contains a payload with the *data consumer*'s pseudonym  $p(c_i)$ , the *data owner*'s pseudonym  $p(o_j)$ , and two copies of the usage data. Each copy is encrypted with an encryption function  $enc()$ , once with the *owner*'s and once with the *consumer*'s one-time public key. Importantly, this key is not shared with other parties and stored securely with the private key. The chosen encryption should be regularly updated, but we require asymmetric (public-key) encryption [14]. As of this point, we recommend RSA [42] with a key size of 3072 bits [24, 28]. Importantly though, each individual can choose their preferred key size, and therefore security level, themselves.



**Figure 1: The components of each block (adapted from [37]), for the exemplary usage  $u_{ij}(c_i \rightarrow o_j)$ . We replace the transaction history utilized in Bitcoin with a generic payload. The payload consists of the pseudonyms  $p(c_i)$  and  $p(o_j)$  of the *data consumer* and *data owner*, respectively. Furthermore, the logged usage is stored twice and encrypted once for the *consumer* and once for the *owner*.**

Besides allowing both parties to read it, storing the usage data twice is also important because the block is created by  $o_j$  (see Section 3.3.2).  $c_i$  then needs to be able to verify the validity of the block. In case  $o_j$  manipulates the stored entry,  $c_i$  can utilize their copy of the usage and the non-repudiation evidence (see Section 3.3.2) to defend themselves against the faked evidence.

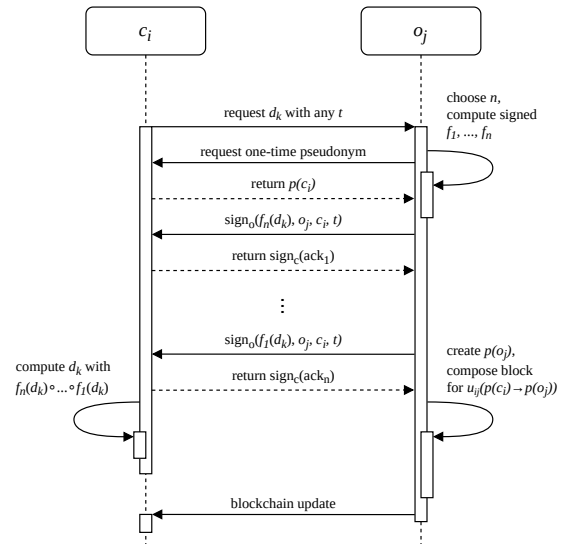
Data usages often follow a similar or even the same form. If a field is limited to few possible values, the deterministic nature of encryption algorithms will lead to the same ciphertext as long as the public key does not differ. This would allow a malicious user to group logs based on those fields, thereby reducing the confidentiality of the data. In order to mitigate this potential threat, the

used encryption method needs to support initialization vectors or optimal asymmetric encryption padding (OAEP). An example for such an algorithm would be RSA OAEP [18].

**3.3.2 New-Usage Protocol.** Many properties we aim for hinge on the specific protocol that is followed when a data usage occurs. Concretely, that means a *data consumer*  $c_i$  is accessing a datum  $d_k$  of a *data owner*  $o_j$  and this usage being logged to the blockchain. The protocol is designed to minimize the knowledge that each party has about the other party, while still ensuring the non-repudiability of their interaction.

Based on our use case, we consider the example of a department manager in the company accessing the completed tasks of one of their subordinates for a yearly report. We specifically choose this example due to the inherent power asymmetry that exists between these parties. Should the manager access data for improper purposes, they might be interested to use their power to remove any incriminating log entry.

The most important challenge is to guarantee non-repudiation of the occurred usage without a trusted third party. Therefore, we adapt the protocol designed by Markowitch and Roggeman [34] for our use case. In short: After the protocol is successfully completed, each party will possess proof of their interaction with the other party. Importantly,  $o_j$  can prove that  $c_i$  has received the datum  $d_k$ .



**Figure 2: The new-usage protocol, adapted from [34, 44].  $c_i$  requests access to the datum  $d_k$  from  $u_j$  with an arbitrary value  $t$  they choose. After receiving pseudonym  $p_i$ , the iterative non-repudiation protocol is performed, with a random number of steps  $n$  unknown to  $c_i$  [see 34, pp. 5–6]. Having received all messages,  $c_i$  can compute the datum  $d_k$  by composing the message payloads. Lastly, the block is created and distributed. Both parties store their respective non-repudiation evidence on their own machine.**

The protocol is described in Figure 2. For this scenario, we assume that  $o_j$  returns the requested datum  $d_k$  directly. Alternatively, they could also return the decryption key for a datum stored elsewhere.

The number of steps  $n$  is chosen at random by  $o_j$ . For less critical data, it can be reduced [34, p. 7], e.g. to lower energy consumption or improve scalability. Following Markowitch and Roggeman [34], as  $c_i$  cannot predict  $n$ , and if the chosen composition function takes long enough to compute, they will not be able to get any meaningful data when cheating [34, p. 5]. Only after having received all  $n$  messages, they can compose them to receive the resulting datum [34, 44]. Now, each party holds non-repudiation evidence of the interaction. For  $o_j$ , this is  $\text{sign}_c(\text{ack}_n)$ , for  $c_i$  it is the set of all messages  $\text{sign}_o(f_x(d_k), o_j, c_i, t)$  with  $x \in (1, n)$  [34, p. 5].

This protocol depends on the nodes being able to verify the authenticity of requests and, importantly, being protected against man-in-the-middle or eavesdropper attacks. Therefore, each request is signed by the sender. We do not aim to reinvent the wheel here, instead relying on the established HTTP over TLS standard [41]. This enables communication confidentiality and authenticity [26]. By utilizing the approach of a web of trust, as established in PGP [1], nodes are fully independent of any trusted third party to verify certificates. In that case, unknown certificates would be rejected and would need to be verified in-person. Alternatively, if sensible for the specific deployment, a certificate authority can be used to sign the individual certificates used by each node to sign and encrypt its requests. As these are often used in companies to enable the signing of internal emails or access to protected resources, no additional certification infrastructure is required in either case.

**3.3.3 Pseudonym Generation Algorithm.** The central step towards our goal is the ability to generate unique one-time pseudonyms that guarantee unlinkability and enable proof of ownership, without requiring a trusted third party. Unlinkability is required for the data we store to be able to qualify as anonymous data, as discussed in Section 2. Proof of ownership, on the other hand, enables the owners of the pseudonyms to exercise their rights of erasure and rectification as given by the GDPR.

Florian et al. [17] describe a pseudonym generation algorithm that serves as inspiration to our solution. Pseudonyms are guaranteed to be unlinkable to each other and to the real identity of the user. Furthermore, authenticity proofs enable proof of ownership, meaning that our requirements are met. Beyond those properties, their algorithm provides sybil-resistance, which is achieved by requiring additional computational steps for joining a network and creating new pseudonyms [17, p. 68–69]. In our case, the additional property of sybil-resistance is not required, as there is no inherent danger in a user creating multiple pseudonyms (see Section 3.3.2). We can therefore omit these additional complexities and simplify our algorithm accordingly. By that, we reduce its computational complexity and energy consumption to a minimum.

Therefore, we define our pseudonym generation algorithm as follows: As part of the new-usage protocol (see Section 3.3.2), the user has created a new RSA private-public key pair with a key size of 3072 bits. As discussed above, the chosen encryption method can be updated if a higher level of security is appropriate. Now, to generate the one-time pseudonym, the collision-resistant and cryptographic hash function BLAKE2 [9], specifically BLAKE2s [9, p. 121], is applied to create a cryptographic message digest. Concretely, the user hashes the public key of their key pair, with the resulting irreversible and cryptographically safe digest representing their


one-time pseudonym. BLAKE2s ensures a digest size of at most 32 bytes, which is important to minimize storage requirements.

The pseudonym generated with our algorithm then guarantees three important properties: First, the owner of the pseudonym, and only the owner, can prove the authenticity of the pseudonym (see below). Second, the unique properties of the hash function guarantee collision-resistance and therefore uniqueness [6]. Third, users only need to manage a single key pair for each block, significantly reducing the complexity of the operation and increasing its speed.

To enable proof of ownership, we make use of the asymmetric nature of the RSA key pair. When a user wants to prove their ownership of a pseudonym, they sign a message with their private key and make available the corresponding public key. The signed message proves that they are in possession of the private key. The public key is then hashed by the recipient with the BLAKE2 hash function. If the result is the correct pseudonym, the ownership is proven. As BLAKE2 is collision-resistant, it is infeasible for an attacker to guess a different string that would result in the same pseudonym. Making matters even more secure, they would in addition need to crack the utilized RSA algorithm, as a collision-inducing string alone would not suffice to allow them to also sign a message with a matching private key. We can therefore conclude that our algorithm is more than sufficiently secure.

**3.3.4 Deployment.** As we have hinted at above, the chosen deployment highly influences the privacy and security guarantees that can be given. The protocol can be flexibly adapted and supports both centralized and peer-to-peer architectures. We aim for our solution to not be dependent on any trusted third party. Therefore, our deployment architecture is fully decentralized.

The central component of the deployment is the  $P^3$  module that handles pseudonym provisioning and key management. Each node in the peer-to-peer network runs its own  $P^3$  instance as well as a private key store KS to store its RSA key pairs and pseudonyms (see Sections 3.3.2 and 3.3.3). The usage log blockchain is permissionless and therefore shared between all nodes participating in the network. Therefore, the architecture is fully decentralized, with each node communicating directly with other nodes (see Figure 3).



**Figure 3: Deployment in a fully decentralized peer-to-peer architecture. Each node runs its own  $P^3$  instance and private key store KS, while the blockchain copy is shared within the network [see also 37].**

The management of a large number of keys as required by our approach can itself become a privacy risk. Therefore, the chosen key management strategy of the key store component KS is relevant. Applying the Master & Sub Key Generation pattern [30] reduces key management complexity while maximizing privacy [30, p. 5]. To additionally protect the master key, the Key Shards pattern [30] can be utilized by users to protect against key loss [30, p. 7].

In theory, the nodes would not necessarily need to store their key pairs and used pseudonyms at all to reduce their attack surface. Yet, these are important to enable requirements (3), (4), and (6). To be able to query for entries concerning their usages, users need to know which pseudonyms belong to them. In theory, they could iterate all blocks in the blockchain and simply try to decrypt their content, but this quickly becomes infeasible. Additionally, to exercise their GDPR-awarded rights, users need to prove their ownership of a pseudonym, requiring them to be in possession of the respective private key for the specific block (see Section 3.3.3). Still, each user can choose on their own how to manage their keys. If they prefer the maximum level of security, they are free to delete new keys immediately after usage.

## 4 ANALYSIS

In the following, we analyze our concept based on three core aspects: Its compliance with the GDPR (our central requirement), its robustness against attacks, and its protocol security.

### 4.1 GDPR Compliance

We have discussed above (see Section 2) that data can be considered pseudonymous and anonymous. When a possibility for re-identification exists, they count as pseudonymous and therefore fall under the provisions of the GDPR. To comply with the GDPR, we need to enable data subjects to exercise their right to erasure as well as their right to rectification. We have also deliberated that by enabling the right to erasure, the right to rectification is also given, as we can simply delete the incorrect block and add a new, corrected version.

Each block in our approach gets its own transaction pseudonym. We have shown in Section 3.3.3 that these pseudonyms are unlinkable to the individual’s identity and to each other. For a recorded usage  $u_{ij}(p_x(c_i) \rightarrow p_y(o_j))$ , only  $c_i$  and  $o_j$  know the association of the other party’s single one-time pseudonym ( $p_x$  or  $p_y$ ) to their real-world identity. In fact,  $c_i$  and  $o_j$  have to prove their identity to each other in the first step of the protocol (see Section 3.3.2). Due to the guaranteed unlinkability, this link is only given for the single pseudonym created for that block, i.e.  $o_j$  only knows the link  $p_x(c_i) \leftrightarrow c_i$ . This case is covered by the existing GDPR provisions. The association is not stored in the blockchain, but only on the nodes of  $c_i$  and  $o_j$ . That means the request for deletion simply has to be forwarded to them. Considering an adversarial user, they might just not fulfill that request for deletion of course (see also Section 4.2). At first glance, this could imply that our protocol does not offer an advantage over just modifying the blockchain and asking all nodes to delete their old copy. Importantly, though, we do not deal with *unknown* nodes. When a deletion request is raised, the data subject *knows* the identity of the offending user, and can *prove* it (see Section 3.3.2). That means, the individual can then be

made responsible for deletion under the GDPR, and can be sued in case they do not follow through.

As soon as the association of the individual’s identity to the pseudonym has been deleted, the data stored in the blockchain are anonymized. Then, they do not qualify as personal data anymore (see also Section 2), satisfying the right to erasure.

### 4.2 Robustness Against Attacks

We have described an adversarial model in Section 3.1, with adversary  $\alpha$  trying to subvert the confidentiality of the stored logs. Specifically, they try to conduct four attacks. For each, we will discuss the robustness of our approach against the attack and potential implications.

Firstly,  $\alpha$  tries to (a) derive from any usage  $u_{ij}(c_i \rightarrow o_j)$  the identity of  $c_i$  or  $o_j$ . As the transaction pseudonyms for both  $c_i$  and  $o_j$  are created with the same algorithm, this attack depends on being able to reverse the employed pseudonym generation. The one-time cryptographic security of our utilized algorithm has been shown [see, e.g., 9, 32], guaranteeing it to be irreversible. Furthermore, as each transaction gets its own pseudonym, each would have to be cracked independently, meaning an exponential effort.

Secondly,  $\alpha$  tries to (b) associate any two usages  $u_{ij}(c_i \rightarrow o_j)$ ,  $u_{ik}(c_i \rightarrow o_k)$  with each other, revealing their association with a single *consumer* and (c) associate any two usages  $u_{ji}(c_j \rightarrow o_i)$ ,  $u_{ki}(c_k \rightarrow o_i)$  with each other, thereby leaking their association with a single *owner*. We can discuss both attacks together, as they hinge on the same security mechanism. The cryptographic security of our algorithm is guaranteed by the cryptographic security of the two underlying algorithms, which has been shown for both [see 32, 33]. Therefore, this again depends on the ability of  $\alpha$  to reverse the transaction pseudonym generation. As we have shown above, this can be considered infeasible.

Finally,  $\alpha$  tries to (d) leak the identity of  $c_j$  for a stored usage  $u_{jj}(p(c_j) \rightarrow p(o_j))$  with  $\alpha = o_j$ , after  $c_j$  has exercised their right to erasure. In the last section, we have detailed that the association of  $c_j$  to their pseudonym is known to  $o_j$  for blocks storing usages of data that  $o_j$  owns. We have no way to technically force  $o_j$  to delete this association when  $c_j$  exercises their right to erasure. Now, assuming  $o_j$  does not delete this data and wants to utilize their knowledge, e.g., by publishing the real identity of  $c_j$  and their one-time pseudonym  $p(c_j)$ . By itself, this proves nothing though, as there is no technical relationship between the pseudonym and the identity of  $c_j$  (see Section 3.3.3). To actually prove the association of  $c_j$  to  $u_{jj}$ ,  $o_j$  therefore needs to publish their non-repudiation evidence (see Section 3.3.2). As shown by [34], this evidence by design contains their own identity (through their signature) as well [34, pp. 5–6]. This means that  $o_j$  would automatically also leak their own identity, making them legally liable.

As this scenario is covered by the privacy legislation and can be prosecuted accordingly, we consider it a non-issue for most cases. Still, for the most secretive of environments, this might not be enough of a guarantee.

### 4.3 Protocol Security

For the highest-security deployments, our peer-to-peer architecture enables pseudonym generation and block creation without

necessitating a trusted third party, mitigating most attack vectors on the integrity and confidentiality of data. Two potential attack vectors on the confidentiality of the exchanged information remain: The messages sent on block creation, and the block update.

Firstly, when creating a new block for a usage  $u_{ij}(c_i \rightarrow o_j)$  following the protocol (see Section 3.3.2), communication between  $c_i$  and  $o_j$  has to occur. Even though HTTP over TLS is utilized, which prevents  $\alpha$  from listening in as an eavesdropper [26], they may still deduce that there is some usage association between the nodes. There are various options to address this. In our view, the easiest to implement would be to add an additional fake protocol. This would work much the same way, only that a special non-existent datum  $d_0$  is requested. Then, both parties understand that this is just a fake request, and no actual block is added to the blockchain. This fake protocol can be run by nodes in randomized intervals, choosing arbitrary other nodes to request  $d_0$  from. Thereby, we can hide real requests in the noise of these fake requests.

What remains then is the block creation. Even if fake protocols are run regularly,  $\alpha$  could simply watch for blockchain updates and derive from those which two nodes were responsible for the new block. The simplest mitigation of this is to add a random wait before the block is added. Then, plausible deniability is enabled, as there are a sufficient number of other potential users that might have been responsible. In fact, nodes might even wait for a certain number of block updates before publishing their update. Here, too, each node can decide itself the level of confidentiality it requires, and act accordingly. Furthermore, traditional blockchain algorithms already (indirectly) protect from  $\alpha$  understanding the originator of a blockchain update. As the architecture is designed to be peer-to-peer, the mere fact that a node sends a block update does not give  $\alpha$  any additional information about its creation. As nodes forward block updates to other nodes, the specific node that sends  $\alpha$  the update might simply have forwarded it [37, 51].

This shows that our protocol can be adapted to fulfill the highest requirements towards information security.

## 5 RELATED WORK

Various proposals for how to solve the conflict between the GDPR requirements and blockchain properties exist. In the following, we describe important works and discuss how they differ from our approach. The overview by Pagallo et al. [38] serves as a foundation. A recent systematic literature review [21] confirms its completeness.

### 5.1 Key Destruction

A trivial solution is to encrypt all data that are stored on the blockchain and delete the decryption key if the data are to be deleted, as discussed by Pagallo et al. [38].

While it is easy to implement, this approach is flawed. Encryption per se only guarantees pseudonymity of data [20, 27], therefore the data protection requirements still apply [27]. More problematically, though, if the full content of the block is encrypted, querying history becomes all but impossible, which is a requirement in secure logs for efficiently reading past entries. Only the affected parties themselves would be able to retrieve their entries, and only with high computational overhead by going through every block and trying to decrypt it.

Our proposal in contrast enables querying of entries based on the one-time pseudonyms, while enabling retroactive anonymization of data, which fulfills the requirements of the GDPR's right to erasure [29, 46].

### 5.2 Forgetting Blockchain

Farshid et al. propose to achieve a GDPR-compliant blockchain by automatically deleting blocks from the blockchain after a certain amount of time has passed [16].

As the described network no longer contains a genesis block, joining a network becomes a challenge. The authors propose to ask other nodes for the current block and just accept it if all the returned blocks are equal [16]. Since there is no way to verify that the received block reflects the true state of the network, joining it requires trust and does not satisfy the integrity constraint. Secondly, the nature of their approach prevents the existence of a chain history. Applications relying on the full history, specifically in the case of secure logs, would therefore not work with this algorithm. Furthermore, this proposal only achieves eventual GDPR compliance, since a block can only be deleted after the predefined time has passed. If a user requests deletion of their data, this request cannot be fulfilled immediately. For this reason, it is questionable if the presented idea is compatible with the GDPR. Most problematically though, the data are only actually deleted if all nodes behave honestly [16]. Any node can simply decide not to delete older blocks, meaning that no additional privacy guarantees can be given.

In contrast to the forgetting blockchain, our proposal does not require adaptation of the utilized blockchain software and is therefore easier to integrate into existing blockchains. Furthermore, we do not depend on the honesty of *arbitrary* and *unknown* nodes. In contrast, only one *known* party has *provable* access to additional identity information and can be held liable under the GDPR.

### 5.3 Redactable Blockchain

The reason that the immutability of data stored in a blockchain can be guaranteed is the utilized hash function: An ideal hashing algorithm guarantees hashes that are one-way, which means impossible to reverse, and collision-free. Then, blocks cannot be replaced without notice, as any change would result in a new hash, thereby invalidating the chain.

Redactable blockchains utilize so-called *chameleon* hash functions to determine the hash of a block. Such hash functions are collision-resistant as long as a secret known as *trapdoor* is not known. If one is in possession of said secret, they can efficiently compute colliding hashes [8, 15]. With the power to create hash collisions, any block can be replaced or even removed [8], making the blockchain effectively arbitrarily editable.

In order to function, such a redactable blockchain network needs a trusted third party that is in possession of the trapdoor and can decide which block to edit [8]. This constraint again requires trust, thereby calling into question the value of utilizing blockchain at all [38]. Furthermore, similar to the forgetting blockchain, every individual node needs to be trusted. Redactions are published as chain updates, allowing *arbitrary* nodes to make a copy of the removed or edited entry before updating their chain [38, 46]. This means that, effectively, no privacy guarantees can be given.

Our solution on the contrary does not require a trusted third party and functions even in the face of adversarial nodes.

## 5.4 Hashing Out

Hashing out refers to the practice of saving the hash of the data on the blockchain and the data itself off-chain [38]. This approach is one of the most commonly used ideas to guarantee GDPR compliance in blockchain solutions [see, e.g., 25, 43, 47]. This is because the on-chain hash does not contain any private or personal data and the off-chain data can be deleted or modified to comply with a data subject’s request.

There are two major downsides, though. Since the data are not stored on the blockchain, this solution is not truly decentralized and requires trust in the authority managing the data [22]. Furthermore, using this approach one can only be sure of the existence of entries, not of their content. That means it is effective only as long as the log is not tampered with, but does not guarantee accountability or non-repudiation.

We allow users to benefit from the maximum possible utility of the log data, specifically giving them access to the content of stored log entries that relate to them, while still guaranteeing the same level of confidentiality. Most importantly, we require no trusted third party.

## 6 DISCUSSION

Even with the best technical protections, individual users remain as an often-abused attack surface [7, 36, 39, 50]. For most data that we store, there is no danger of users unwillingly leaking information about other parties except for themselves, with one exception: *data owners* could be tricked or hacked to reveal the identities of *consumers* of their data. In our current architecture, it is impossible to prevent this case, yet we consider the attack surface to be acceptably small. To get access to a meaningful dataset about the usage pattern of a *data consumer*, an adversary would have to hack or phish each individual *data owner* of data accessed by said *data consumer*. We consider that infeasible.

Furthermore, we prioritize security. That in turn means that other properties, such as availability or scalability, are not optimized for. Regarding availability, each individual node manages its own data and has to be reached when accessing data. Should the node crash, be shut off, or otherwise disconnected from the network, the *data consumer* is prevented from continuing their work. In scenarios where the availability of the nodes is prioritized higher than their security and independence, we can imagine running user’s nodes, e.g., on virtual servers. While this adds an attack surface and removes control from the user, it can improve availability. Considering scalability on the other hand, the requirement to have a key pair for each individual access can quickly lead to a huge key store. This is one of the biggest weaknesses of our approach. Beyond pooling requests, the most sensible solution in less critical scenarios might be to reduce key size or reuse pseudonyms. As always, this is a trade-off depending on the specific requirements.

To expand on this point, there has been broader discussion on what constitutes “good enough” software security and how to make objective judgments about it [45]. Tøndel et al. suggest to not only consider the system from the perspective of the adversary (as we

have done), but to additionally factor in other perspectives such as those of users or operators [45, p. 364]. Following their proposal, it might therefore be sensible to conduct a broader analysis of the system that also covers these perspectives before deploying it. This could be important to ensure user acceptance and usability, as well as to address potential practical issues with deployment and operation, which might otherwise hinder adoption.

## 7 CONCLUSION

At first glance, the requirements of privacy legislation such as the GDPR seem to fundamentally contradict the technical properties of blockchain. This makes it seemingly impossible to store data governed by the GDPR in a blockchain. Yet, the unique advantages of utilizing blockchain technology can be beneficial in cases where a tamper-proof secure log is needed, such as when logging data usages. By design, such a secure usage log stores personal data governed by the GDPR. To solve this conflict, we have analyzed the requirements of the GDPR and developed a concept to address them. Our concept exploits a specific legal property of pseudonymized data: Depending on if they can be linked to the identity of a user, they count as personal data or as anonymous. We describe a pseudonym generation algorithm and protocol that guarantee unlinkability and anonymity, while enabling proof of ownership and authenticity. With those properties, the requirements of the GDPR can be fulfilled. Our block structure and deployment architecture then allow individuals to query for and read arbitrary usage log entries concerning their data, while protecting them from attacks on their confidentiality by adversaries. In our analysis, we deliberate how this allows us to fulfill the requirements of the GDPR by enabling confidentiality and the rights to erasure and rectification, while at the same time benefiting from the properties of blockchain, specifically guaranteeing the integrity of the logged data. Related works require either the use of a permissioned blockchain, necessitating a trusted third party, or modifying the utilized hashing algorithms or blockchain software to make the blockchain mutable. Both approaches entail effectively giving up the advantages of blockchain, thereby calling into question the use of blockchain in the first place. With the P<sup>3</sup> system, GDPR-compliant secure usage logs based on blockchain can be implemented. In addition, our approach does not require a trusted third party, can be used with existing blockchain software without requiring changes, and is secure and flexible enough to enable deployment even in highest-security settings.

## ACKNOWLEDGMENTS

This work was supported by the German Federal Ministry of Education and Research (BMBF) under grant no. 5091121.

## REFERENCES

- [1] Alfarez Abdul-Rahman. 1997. The PGP trust model. *EDI-Forum: The Journal of Electronic Commerce* 10 (1997), 27–31.
- [2] Rafael Accorsi. 2010. BBox: A distributed secure log architecture. In *Proceedings of the 2010 European Public Key Infrastructure Workshop*. Springer, 109–124.
- [3] Rakesh Agrawal, Jerry Kiernan, Ramakrishnan Srikant, and Yirong Xu. 2002. Hippocratic databases. In *Proceedings of the 28<sup>th</sup> International Conference on Very Large Databases*. Elsevier, 143–154.
- [4] Saifedean Ammous. 2016. Blockchain technology: What is it good for? (2016). Available at SSRN: <https://ssrn.com/abstract=2832751>.
- [5] Elli Androulaki, Ghassan O. Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. 2013. Evaluating user privacy in bitcoin. In *Proceedings of the 17<sup>th</sup>*

- International Conference on Financial Cryptography and Data Security*. Springer, 34–51.
- [6] Benny Applebaum, Naama Haramaty-Krasne, Yuval Ishai, Eyal Kushilevitz, and Vinod Vaikuntanathan. 2017. Low-complexity cryptographic hash functions. In *Proceedings of the 8<sup>th</sup> Innovations in Theoretical Computer Science Conference (Leibniz International Proceedings in Informatics)*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 7:1–7:31.
- [7] Iván Arce. 2003. The weakest link revisited. *IEEE Security & Privacy* 1, 2 (2003), 72–76.
- [8] Giuseppe Ateniese, Bernardo Magri, Daniele Venturi, and Ewerton Andrade. 2017. Redactable blockchain – or – Rewriting history in bitcoin and friends. In *Proceedings of the 2<sup>nd</sup> IEEE European Symposium on Security and Privacy*. IEEE, 111–126.
- [9] Jean-Philippe Aumasson, Samuel Neves, Zooko Wilcox-O’Hearn, and Christian Winnerlein. 2013. BLAKE2: simpler, smaller, fast as MD5. In *Proceedings of the 11<sup>th</sup> International Conference on Applied Cryptography and Network Security*. Springer, 119–135.
- [10] Eugene Bagdasaryan, Griffin Berstein, Jason Waterman, Eleanor Birrell, Nate Foster, Fred B. Schneider, and Deborah Estrin. 2019. Ancile: Enhancing privacy for ubiquitous computing with use-based privacy. In *Proceedings of the 18<sup>th</sup> ACM Workshop on Privacy in the Electronic Society*. ACM, 111–124.
- [11] David Brin. 1998. *The Transparent Society*. Basic Books, New York.
- [12] California Consumer Privacy Act. 2018. An act to add Title 1.815.5 (commencing with Section 1798.100) to Part 4 of Division 3 of the Civil Code, relating to privacy (California Consumer Privacy Act of 2018). *Assembly Bill 375* (2018), 1–24.
- [13] Yves-Alexandre de Montjoye, Erez Shmueli, Samuel S. Wang, and Alex Sandy Pentland. 2014. openPDS: Protecting the privacy of metadata through SafeAnswers. *PLOS ONE* 9, 7 (2014), 1–9.
- [14] Hans Delfs and Helmut Knebl. 2007. Public-key cryptography. In *Introduction to Cryptography*. Springer, Chapter 3, 33–80.
- [15] Nico Dötting, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. 2019. Trapdoor hash functions and their applications. In *Proceedings of the 39<sup>th</sup> Annual International Cryptology Conference*. Springer, 3–32.
- [16] Simon Farshid, Andreas Reitz, and Peter Roßbach. 2019. Design of a forgetting blockchain: A possible way to accomplish GDPR compatibility. In *Proceedings of the 52<sup>nd</sup> Hawaii International Conference on System Sciences*. 7087–7095.
- [17] Martin Florian, Johannes Walter, and Ingmar Baumgart. 2015. Sybil-resistant pseudonymization and pseudonym change without trusted third parties. In *Proceedings of the 14<sup>th</sup> ACM Workshop on Privacy in the Electronic Society*. ACM, 65–74.
- [18] Eiichiro Fujisaki, Tatsuaki Okamoto, David Pointcheval, and Jacques Stern. 2001. RSA-OAEP is secure under the RSA assumption. In *Proceedings of the 21<sup>st</sup> Annual International Cryptology Conference*, Vol. 17. Springer, 260–274.
- [19] Chunpeng Ge, Siwei Sun, and Pawel Szalachowski. 2019. Permissionless blockchains and secure logging. In *Proceedings of the 2019 IEEE International Conference on Blockchain and Cryptocurrency*. IEEE, 56–60.
- [20] General Data Protection Regulation. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union* 59 (2016), 1–88.
- [21] AKM Bahalul Haque, AKM Najmul Islam, Sami Hyrinsalmi, Bilal Naqvi, and Kari Smolander. 2021. GDPR compliant blockchains—a systematic literature review. *IEEE Access* 9 (2021), 50593–50606.
- [22] Luis-Daniel Ibáñez, Kieron O’Hara, and Elena Simperl. 2018. *On blockchains and the general data protection regulation*. Technical Report 422879. University of Southampton. 1–13 pages. <https://eprints.soton.ac.uk/422879/>
- [23] ISO 25237:2017 2017. *Health informatics – Pseudonymization*. Standard. International Organization for Standardization, Geneva, CH.
- [24] Mikko Kiviharju. 2017. On the fog of RSA key lengths: Verifying public key cryptography strength recommendations. In *Proceedings of the 2017 International Conference on Military Communications and Information Systems*. IEEE, 1–8.
- [25] Michael Kolain and Christian Wirth. 2018. Privacy by Blockchain Design: A blockchain-enabled GDPR-compliant approach for handling personal data. In *Proceedings of the 1<sup>st</sup> ERCIM Blockchain Workshop*. EUSSET.
- [26] Hugo Krawczyk, Kenneth G. Paterson, and Hoeteck Wee. 2013. On the security of the TLS protocol: a systematic analysis. In *Proceedings of the 33<sup>rd</sup> Annual Cryptology Conference*. Springer, 429–448.
- [27] Cedric Lauradoux, Konstantinos Limniotis, Marit Hansen, Meiko Jensen, and Petros Efthymiopoulos. 2021. *Data pseudonymisation: advanced techniques and use cases*. Technical Report. European Union Agency for Cybersecurity (ENISA). <https://www.enisa.europa.eu/publications/data-pseudonymisation-advanced-techniques-and-use-cases/>
- [28] Arjen K. Lenstra and Eric R. Verheul. 2001. Selecting cryptographic key sizes. *Journal of Cryptology* 14, 4 (2001), 255–293.
- [29] Konstantinos Limniotis and Marit Hansen. 2019. *Recommendations on shaping technology according to GDPR provisions – an overview on data pseudonymisation*. Technical Report. European Union Agency for Cybersecurity (ENISA). <https://www.enisa.europa.eu/publications/recommendations-on-shaping-technology-according-to-gdpr-provisions>
- [30] Yue Liu, Qinghua Lu, Hye-Young Paik, and Xiwei Xu. 2020. Design patterns for blockchain-based self-sovereign identity. In *Proceedings of the 2020 European Conference on Pattern Languages of Programs*. ACM, 1–14.
- [31] Sin K. Lo, Xiwei Xu, Yin K. Chiam, and Qinghua Lu. 2017. Evaluating suitability of applying blockchain. In *Proceedings of the 22<sup>nd</sup> International Conference on Engineering of Complex Computer Systems*. IEEE, 158–161.
- [32] Atul Luykx, Bart Mennink, and Samuel Neves. 2016. Security analysis of BLAKE2’s modes of operation. *IACR Transactions on Symmetric Cryptology* 2016, 1 (2016), 158–176.
- [33] Dindyal Mahto, Danish Ali Khan, and Dilip Kumar Yadav. 2016. Security analysis of elliptic curve cryptography and RSA. In *Proceedings of the 2016 World Congress on Engineering*, Vol. 1. 419–422.
- [34] Olivier Markowitch and Yves Roggeman. 1999. Probabilistic non-repudiation without trusted third party. In *Proceedings of the 2<sup>nd</sup> Conference on Security in Communication Networks*, Vol. 99. 25–36.
- [35] Aleecia M. McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society* 4, 3 (2008), 543–568.
- [36] Kevin D. Mitnick and William L. Simon. 2002. *The art of deception: Controlling the human element of security*. John Wiley & Sons, New York, NY.
- [37] Satoshi Nakamoto. 2008. Bitcoin: A peer-to-peer electronic cash system. (2008).
- [38] Ugo Pagallo, Eleonora Bassi, Marco Crepaldi, and Massimo Durante. 2018. Chronicle of a clash foretold: Blockchains and the GDPR’s right to erasure. In *Proceedings of the 31<sup>st</sup> Annual Conference on Legal Knowledge and Information Systems*. 81–90.
- [39] Marissa Randazzo, Michelle Keeney, Eileen Kowalski, Dawn Cappelli, and Andrew Moore. 2005. *Insider threat study: Illicit cyber activity in the banking and finance sector*. Technical Report CMU/SEI-2004-TR-021. Software Engineering Institute, Carnegie Mellon University.
- [40] Fergal Reid and Martin Harrigan. 2013. An analysis of anonymity in the bitcoin system. In *Security and Privacy in Social Networks*, Yaniv Altshuler, Yuval Elovici, Armin B. Cremers, Nadav Aharoni, and Alex Pentland (Eds.). Springer, New York, NY, 197–223.
- [41] Eric Rescorla. 2000. *HTTP Over TLS*. RFC 2818. RFC Editor. <https://www.rfc-editor.org/rfc/rfc2818.txt>
- [42] Ronald L. Rivest, Adi Shamir, and Leonard Adleman. 1978. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21, 2 (1978), 120–126.
- [43] Christian Schaefer and Christine Edman. 2019. Transparent logging with Hyperledger Fabric. In *Proceedings of the 2019 IEEE International Conference on Blockchain and Cryptocurrency*. IEEE, 65–69.
- [44] Adi Shamir. 1979. How to share a secret. *Commun. ACM* 22, 11 (1979), 612–613.
- [45] Inger Anne Tandell, Daniela Soares Cruzes, and Martin Gilje Jaatun. 2020. Achieving “good enough” software security: the role of objectivity. In *Proceedings of the 2020 International Conference on Evaluation and Assessment in Software Engineering*. ACM, 360–365.
- [46] David van de Giessen. 2019. *Blockchain and the GDPR’s right to erasure*. Technical Report. University of Twente.
- [47] Laurens Van Hove, Pieter-Jan Maenhaut, Tim Wauters, Bruno Volckaert, and Filip De Turck. 2019. Logging mechanism for cross-organizational collaborations using Hyperledger Fabric. In *Proceedings of the 2019 IEEE International Conference on Blockchain and Cryptocurrency*. IEEE, 352–359.
- [48] Stefan G. Weber. 2012. On transaction pseudonyms with implicit attributes. Cryptology ePrint Archive, Report 2012/568.
- [49] Daniel J. Weitzner, Harold Abelson, Tim Berners-Lee, Joan Feigenbaum, James Hendler, and Gerald Jay Sussman. 2008. Information accountability. *Commun. ACM* 51, 6 (2008), 82–87.
- [50] Ryan West, Christopher Mayhorn, Jefferson Hardee, and Jeremy Mendel. 2009. The weakest link: A psychological perspective on why users make poor security decisions. In *Social and Human Elements of Information Security: Emerging Trends and Countermeasures*, Manish Gupta and Raj Sharma (Eds.). IGI Global, 43–60.
- [51] Zibin Zheng, Shaoan Xie, Hongning Dai, Xiangping Chen, and Huaimin Wang. 2017. An overview of blockchain technology: architecture, consensus, and future trends. In *Proceedings of the 2017 International Congress on Big Data*. IEEE, 557–564.
- [52] Valentin Zieglmeier and Alexander Pretschner. 2021. Trustworthy transparency by design. arXiv:2103.10769 [cs.SE]
- [53] Guy Zyskind, Oz Nathan, et al. 2015. Decentralizing privacy: Using blockchain to protect personal data. In *Proceedings of the 2015 IEEE Security and Privacy Workshops*. IEEE, 180–184.