

Accountable Decryption made Formal and Practical

Rujia Li, *IEEE Member*, Yuanzhao Li, *IEEE Student Member*, Qin Wang, *IEEE Member*, Sisi Duan, *IEEE Senior Member*, Qi Wang, *IEEE Member*, and Mark Ryan



Abstract—With the increasing scale and complexity of online activities, accountability, as an after-the-fact mechanism, has become an effective complementary approach to ensure system security. Decades of research have delved into the connotation of accountability. They fail, however, to achieve *practical* accountability of decryption. This paper seeks to address this gap. We consider the scenario where a client (called encryptor, her) encrypts her data and then chooses a delegate (a.k.a. decryptor, him) that stores data for her. If the decryptor initiates an illegitimate decryption on the encrypted data, there is a non-negligible probability that this behavior will be detected, thereby holding the decryptor accountable for his decryption. We make three contributions. First, we review key definitions of accountability known so far. Based on extensive investigations, we formalize new definitions of accountability specifically targeting the decryption process, denoted as *accountable decryption*, and discuss the (im)possibilities when capturing this concept. We also define the security goals in correspondence. Secondly, we present a novel Trusted Execution Environment (TEE)-assisted solution aligning with definitions. Instead of fully trusting TEE, we take a further step, making TEE work in the “trust, but verify” model where we trust TEE and use its service, but empower users (i.e., decryptors) to detect the potentially compromised state of TEEs. Thirdly, we implement a full-fledged system and conduct a series of evaluations. The results demonstrate that our solution is efficient. Even in a scenario involving 300,000 log entries, the decryption process concludes in approximately 5.5ms, and malicious decryptors can be identified within 69ms.

Index Terms—Accountability, Decryption, Trusted Hardware

1 INTRODUCTION

Modern cryptographic systems mainly depend on preventive strategies such as passwords, access control mechanisms, and authentication protocols to ensure privacy and security. Prior to gaining access to sensitive data or performing any action that raises privacy or security concerns, individuals must demonstrate their authorization to do so. These preventive strategies can indeed offer a degree of protection against unauthorized users attempting to infiltrate the system. Nonetheless, as the scale and complexity of computer systems increase dramatically, it has become

increasingly evident that relying solely on a preventive approach is insufficient. In some break-glass scenarios [1], users have to bypass the preventive strategies to access sensitive information. Consequently, accountability becomes a crucial mechanism to supplement preventive strategies, as it can identify and punish malicious individuals who breach pre-established rules after accessing sensitive data [2].

In this work, we formalize the concept of *accountable decryption* [3][4], which is an innovative approach that allows users to share the ability to access a specific piece of information, simultaneously providing evidence of whether the information has been accessed or not. We consider a scenario with three roles, as illustrated in Figure 1.(a): *encryptor*, *decryptor*, and *judge*. An encryptor is a person who provides access for a decryptor to her encrypted data and defines policies to restrain how the decryption can be conducted. A decryptor decrypts the ciphertext and recovers the original data. Each decryption operation performed by the decryptor unavoidably produces evidence of the decryption which undergoes scrutiny by a judge. The judge verifies whether the conducted decryption complies with the predefined policies. If any policy violation is detected, the judge imposes appropriate punishments on the decryptor. In this way, the decryptor is accountable for his behavior.

Bringing accountability for decryption offers many benefits. For example, in electronic surveillance [5], the law-enforcement agency seeks access to users’ sensitive data from tech companies. Yet, confidential court-ordered inquiries are processed privately to ensure investigation integrity. This curtails oversight of privacy-sensitive government actions. Indeed, there are concerns about government power and privacy, as there is a risk of eavesdropping on legal citizens. Accountable decryption offers a potential solution in this case: a user (encryptor), law-enforcement agency (decryptor), and overseer (judge) work in tandem to strike a balance between investigation requirements and privacy requirements. On the one hand, accountable decryption allows law-enforcement agencies to access sensitive data in adherence to investigative orders. On the other hand, it facilitates the detection of potential abuses of the granted warrant, as overseers can verify whether law-enforcement agencies comply with the court order’s limits.

Beyond electronic surveillance, the practical accountability of decryption offers a range of general benefits for many applications (cf. Section 9.2). We summarize these benefits

Rujia Li, and Sisi Duan are with the Tsinghua University, Beijing, China; Email: {rujia, duansisi}@tsinghua.edu.cn.

Yuanzhao Li, and Qi Wang, are with the Southern University of Science and Technology, Shenzhen, China. Email: {12232416,wangqi}@mail.sustech.edu.cn.

Qin Wang is with the CSIRO Data61, Sydney, Australia. Email: qin-wangtech@gmail.com.

Mark Ryan is with the University of Birmingham, Birmingham, UK. Email: m.d.ryan@bham.ac.uk.

as follows.

- *Detection of unauthorized access.* Accountable decryption allows the encryptor to audit the decryption process [3].
- *Deterrents of illegal behaviours* [6]. The decryptor is responsible for its actions, as any violation of the access control to the data can be caught and punished [7].
- *Regulatory compliance.* One can also implement regulatory compliance with the pre-defined policies. The accountability of a decryptor’s behavior demonstrates his compliance with regulations, as a verifiable record of decryption events is provided.
- *Key leakage awareness.* If a decryptor discovers that her ciphertexts have been decrypted without her permission, she will be alerted of a potential leakage of decryption keys.

Reviewing accountability definitions in the literature.

The concept of accountability has been defined in different contexts. We reviewed accountability definitions in the literature, including those defined for identity authentication [8][9], non-tampering evidence [1][6], and policy violations [7][10]. Our observation is that while these definitions share similarities and follow the general principle of *after-the-fact behavior checking* and *misconduct blaming* [11][12], we still need a new notion for accountable decryption. In particular, prior definitions only provide a rather high-level abstraction of accountability. For instance, they fail to delineate what constitutes legitimate decryption activities. Furthermore, prior definitions often assume the existence of a trusted third party responsible for managing and tracking evidence of actions without addressing scenarios where this intermediary is compromised. Accordingly, we extend prior definitions and propose a new definition for accountable decryption.

Formalizing definitions and properties for accountable decryption. We propose new definitions of accountability that specifically target the decryption process. We emphasize two fundamental pillars for achieving ideal accountability of decryption (cf. Definition 11). First, our definitions ensure that dishonest decryptors can always be detected and punished (*non-repudiation*). Second, we ensure that honest decryptors are never wrongly accused and punished (*non-frameability*). This definition is derived from the synthesis of existing studies with various emphases. It encompasses faithful compliance at every stage, including policy setting, action performance, post-feedback, and punishments with appropriate punishments.




Figure 1: General review on our solution

The above definition, however, might seem sufficient for defining the accountability of decryption. In practice, the

decryption is conducted in the decryptor’s local client. A decryptor could decrypt the ciphertext silently and secretly without providing decryption evidence to the judge, and thus, the misbehavior cannot be captured [4]. To achieve the requirement that the decryptor’s actions are visible to the judge, we require another party called *trustee* (TS), which holds the decryption key. It receives requests from the decryptor and enforces the visibility of the requests to the judge. We show the solution in Figure 1.(b). Unfortunately, this reliance on specific roles may raise another concern for our definition, as the TS may become corrupted or untrustworthy. Thus, we propose two practical definitions for accountable decryption, considering a trusted TS (cf. Definition 13) and an untrusted TS (cf. Definition 16).

Building a practical accountable decryption scheme.

We construct an innovative scheme that achieves accountable decryption, called PORTEX. Our solution leverages trusted hardware as the trustee, specifically Trusted Execution Environments (TEEs) [13][14][15]. In particular, by integrating TEEs into the trustees’ infrastructure, the decryptor’s private key and the audit trail of data access are securely protected by trusted hardware. Every time decryption is required, the decryptor must initiate a request to the trusted hardware. In response, the trusted hardware not only supplies the necessary decryption key but also records evidence of the transaction. This process ensures that decryption activities are accountable, thereby fulfilling the criteria outlined in Definition 13.

We have noticed that TEEs are vulnerable to fault attacks [16], software attacks [17], and side-channel attacks [18], et al. While certain countermeasures (e.g., software patches) can mitigate some of these vulnerabilities, achieving a completely secure TEE remains uncertain. Recent research [19][20] has highlighted that a compromised TEE cannot maintain the foundational security properties expected of TEE-based systems. In light of these challenges, our approach deviates from the traditional reliance on the inherent trustworthiness of TEEs. Instead, we introduce two algorithms to cope with TEE failures. The first algorithm inspired by the key management solutions used in [21], [22], is designed such that TEE only stores a partial key while the user securely holds another portion. Even if attackers compromise the TEE, they cannot deduce a complete decryption key. The second algorithm adopts a *trust, but verify* model, where we trust TEE and use its service, but empower users (i.e., decryptors) to detect the potentially compromised state of TEEs. This minimizes losses upon detecting a compromised state. By these algorithms, our system is in accordance with Definition 16. To the best of our knowledge, we proposed the first *practical* solution that achieves accountability in TEE-based systems without assuming absolute trust in TEEs.

Providing full-fledged implementation and evaluations.

We provide a fully functional implementation¹² including encryption module, decryption module, tracing module, and conduct a series of evaluations for our implementation. Experimental results further demonstrate that our system is efficient. Even in an extreme scenario involving 300,000

¹Released at <https://github.com/portex-tee/portex>
²Public service at <http://ac-dec.com>

log entries, the decryption process concludes in approximately 5.5ms. Furthermore, a user can identify the malicious decryptor within 69ms, while the compromised TEE is detected in a mere 4.3ms. In addition, we benchmark our proposed system against state-of-the-art studies. The results indicate that our system offers a more practical solution, as it obviates the need for a centralized storage center for ciphertext backups, thereby reducing both the number of participants and the frequency of interactions.

2 PRELIMINARIES

Trusted hardware. We use the trusted execution environment in this work. TEE is a secure area within the main processor that operates as an isolated kernel. It ensures the confidentiality and integrity of sensitive data and computations. State-of-the-art TEE implementations include Intel Software Guard Extensions (SGX) [13], ARM TrustZone [14], RISC-V Keystone [15]. A TEE typically provides features including: *runtime isolation* and *local/remote attestation*. Without loss of generality, we use Intel SGX as the TEE instance to illustrate TEE’s key features. Runtime isolation guarantees that the code execution is isolated from untrusted memory regions. Attestation proves that the application is running within trusted hardware. By following several notions from [23], we define TEE as general trusted hardware HW.

Definition 1 (Trusted hardware, HW). HW for executing a probabilistic polynomial time (PPT) program Q consists of algorithms $\{\text{HW.Setup}, \text{HW.Load}, \text{HW.Run}, \text{HW.Run\&Quote}, \text{HW.VerifyQuote}\}$.

- $\text{HW.Setup}(\lambda)$: The algorithm inputs the security parameter λ , creates a secret key sk_{quote} for signing remote attestation quotes, and outputs public parameters pm_{shw} .
- $\text{HW.Load}(pm_{shw}, Q)$: The algorithm creates an enclave and loads the code Q into the created enclave. It outputs the handle hdl of an enclave.
- $\text{HW.Run}(hdl, in)$: The algorithm executes the program in the enclave with an input in .
- $\text{HW.Run\&Quote}(hdl, in)$: It takes as input hdl, in , and executes the program in enclaves. After obtaining the result, the enclave signs the output with sk_{quote} , and creates a quote $q = (md_{hdl}, H_Q, in, out, \sigma)$, where md_{hdl} is the enclave metadata, H_Q is a hash of Q , and σ is the signature of previous data.
- $\text{HW.VerifyQuote}(pm_{shw}, q)$: The algorithm verifies the quote. It verifies the signature σ and outputs *true* if the verification succeeds. Otherwise, it outputs *false*.

Integrity of TEE execution. TEE creates a secure, isolated area that can run the application and handle sensitive data. This area is a self-contained execution environment completely isolated from the rest of the system. The application code cannot be replaced or modified once it is successfully loaded in this area [24]. Thus, the execution of identical loaded code with congruent inputs shall yield consistent outputs. We formalized this property as follows.

Definition 2 (TEE Execution Integrity, Exelnty). Considering that an adversary \mathcal{A} and a challenger \mathcal{C} playing the following game.

$$\begin{array}{l} \mathbf{G}_{\text{Exelnty}}(\lambda) \\ \hline 1: \mathcal{C} \text{ Run } pm_{shw} \leftarrow \text{HW.Setup}(1^\lambda) \\ 2: \mathcal{C} \text{ and initialize } \mathcal{O} := \emptyset. \\ 3: \mathcal{A} \text{ runs } hdl \leftarrow \text{HW.Load}(pm_{shw}, Q) \\ 4: \mathcal{A} \text{ runs } q \leftarrow \text{HW.Run}(hdl, in) \text{ outputting } opt. \\ 5: \mathcal{A} \text{ adds } opt \text{ into } \mathcal{O} \\ 6: \text{repeat step 3-5} \end{array}$$

The advantage $\text{Adv}_{\mathcal{A}, \text{HW}}^{\text{G}_{\text{Exelnty}}}(\lambda)$ of \mathcal{A} winning the game is

$$\text{Adv}_{\mathcal{A}, \text{HW}}^{\text{G}_{\text{Exelnty}}}(\lambda) = \Pr[\text{HW.Run}(hdl, in) \notin \mathcal{O}].$$

HW is Exelnty secure if for all PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ such that $\text{Adv}_{\mathcal{A}, \text{HW}}^{\text{G}_{\text{Exelnty}}}(\lambda) < \text{negl}(\lambda)$.

Notes. In several TEE products (e.g., Intel SGX), to preserve signers’ anonymity, a quote is created using an anonymous group signature [13]. As anonymity is orthogonal to our formalization, we omit it in this work.

Cryptographic primitives. Our construction relies on standard cryptographic primitives, i.e., public key encryption scheme PKE, signature scheme S , and commitment scheme CM. We list them as follows.

Signature scheme. A signature scheme S consists of three PPT algorithms ($S.\text{Gen}, S.\text{Sign}, S.\text{Verify}$). $S.\text{Gen}$ takes as input λ and outputs a key pair (vk, sk) , written as $(vk, sk) \leftarrow S.\text{Gen}(1^\lambda)$. $S.\text{Sign}$ takes as input a private key sk and a message m and outputs a signature σ , written as $\sigma \leftarrow S.\text{Sign}_{sk}(m)$. The deterministic algorithm $S.\text{Verify}$ takes as input a verification key vk , a message m , and a signature σ . It outputs 1 if the signature is valid and 0 otherwise, written as $b \leftarrow S.\text{Verify}_{vk}(m, \sigma)$.

Definition 3. A signature scheme S is existential unforgeability under chosen message attack (EUF-CMA) secure if for all PPT adversaries \mathcal{A} , there exists a negligible function negl such that

$$\Pr[\text{Sig-forge}_{\mathcal{A}, S}^{\text{CMA}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where $\text{Sig-forge}_{\mathcal{A}, S}^{\text{CMA}}(\lambda)$ is an experiment defined as below:

- Run $S.\text{Gen}(1^\lambda)$ to obtain a key pair (vk, sk) .
- \mathcal{A} is given vk and access to a signature oracle $\mathcal{O}^{S.\text{Sign}_{sk}(\cdot)}$ and $\mathcal{O}^{S.\text{Verify}_{vk}(\cdot)}$. Then \mathcal{A} outputs a legal message m with a signature σ .
- Let \mathcal{Q} denote the set of messages queried by \mathcal{A} via the oracle $\mathcal{O}^{S.\text{Sign}_{sk}(\cdot)}$. The experiment’s output is 1 if $S.\text{Verify}_{vk}(m, \sigma) = 1$ and $m \notin \mathcal{Q}$, and 0 otherwise.

Commitment scheme. A commitment scheme CM consists of three PPT algorithms ($\text{CM.Gen}, \text{CM.Com}, \text{CM.Verify}$). CM.Gen takes as input λ and outputs common parameters g, h , written as $g, h \leftarrow \text{CM.Gen}(1^\lambda)$. CM.Com takes as input a random decommitment value r and a message m and outputs a commitment c , written as $c \leftarrow \text{CM.Com}(m, r)$. The deterministic algorithm CM.Verify takes as input a commitment c , a message m , and a decommitment value r . It outputs 1 if the commitment is valid and 0 otherwise, written as $b \leftarrow \text{CM.Verify}(c, m, r)$.

Definition 4. A commitment scheme CM is perfect hiding if for all PPT adversaries \mathcal{A} , satisfied

$$\Pr[\text{H-exp}_{\mathcal{A}, \text{CM}}(\lambda) = 1] = 1/2,$$

where $\text{H-exp}_{\mathcal{A}, \text{CM}}(\lambda)$ is defined as follows:

- The adversary \mathcal{A} is provided with the output of CM.Gen and selects two messages $m_0, m_1 \in \mathbb{Z}_p$.
- A bit $b \in \{0, 1\}$ and the value r is chosen uniformly at random, and the commitment c to m_b is computed as $c \leftarrow \text{CM.Com}(m_b, r)$.
- The adversary \mathcal{A} attempts to guess b based on the commitment c , yielding a guess b' .
- The experiment outputs 1 if $b' = b$, and 0 otherwise.

3 DEFINITIONS OF ACCOUNTABILITY

In this section, we review the development of the definitions. Our objectives are twofold: (i) to assess the applicability of the existing definitions to accountable decryption, and (ii) if they are not suitable, to identify commonalities in the current definitions that can serve as a foundation for formulating our new definition.

📌 Nissenbaum (1996) [25] first proposed the notion of accountability in computer science. It did not provide a clear definition of accountability but strengthened that accountability is crucial for building reliable computer science.

📌 Kailar (1996) [8] and Yumerefendi et al. (2005) [9] extended the definition of accountability by emphasizing its objectivity and establishing its strong connection to real entities. Their work emphasized the importance of binding accountability with a unique identifier, enabling the tracking and linking of actions to specific individuals or entities. This crucial exploration allows the detection of misbehavior.

Definition 5 (On linkage [8][9]). *An accountable system associates states and actions with identities and provides primitives for actors to validate the states and actions of their peers, such that cheating or misbehavior becomes detectable, provable, and undeniable by the perpetrator.*

📌 Subsequent endeavors have further enriched the concept of accountability by accentuating the significance of honest behaviors exhibited by participants. To achieve this, Bella et al. (2006) [1] and Haeberlen et al. (2007) [6] embrace accountability as the process of furnishing evidence to a principal, which can subsequently be presented to the judge. Buldas et al. (2000) [26] and Yumerefendi et al. (2007) [27] delved into its application and strengthened the importance of real-time detection of misconduct.

Definition 6 (On detection, rephrased). *Accountability has the following features: (a) reliable evidence [1][6]: It indicates the delivery to a principal of evidence that is later presented to the judge. The evidence can be validated and achieves fairness (i.e., that one protocol participant gets evidence if and only if the other one does) and non-repudiation; (b) enhanced misconduct detection [26][27]: A system should provide a means to directly detect and expose misbehavior by its participants, or, enable a principal to prove to the judge any detected fraud.*

📌 A parallel line of studies, comprising Lampson (2005) [12], and Feigenbaum et al. (2011) [11], further improved the accountability definition by emphasizing that accountability is actually a posteriori behavior. These studies emphasized the significance of consequences and deterrent

measures in shaping responsible behavior in various domains. Grant et al. (2005) [10] and Weitzner et al. (2008) [7] underscored the imperative of validating adherence to pre-defined policies.

Definition 7 (On punishment, rephrased). *Accountability in a computing system implies the following properties: (a) awareness of policy violation [7][10]: Some actors have the right to hold other actors to a set of standards, to judge whether they have fulfilled their responsibilities in light of these standards. (b) posterior penalty [10][11]: An entity is accountable with respect to some policy (or accountable for obeying the policy). Whenever the entity violates the policy, with some non-negligible probability, the entity will be punished.*

📌 Subsequent studies such as Haeberlen (2010) [28] have made strides in identifying the key factors that define the connotation of accountability. In addition to that, a few studies, such as Ishai et al. (2014) [29], have delved into establishing formal treatments of accountability.

Definition 8 (A summary by [28]). *A system is accountable if (a) faults can be reliably detected, (b) each fault can be undeniably linked to at least one faulty node, and (c) the faulty entities will be properly sanctioned.*

The above review shows that decades of research have thoroughly studied the concept of accountability with different focuses. However, they fail to fully capture the entities, components, and algorithms involved in accountable decryption. For example, none of the definitions explicitly define the actions of encryption operations or what constitutes legitimate decryption activities. This finding highlights the need for an accountable decryption definition.

General principle of accountability. We summarize the general principle of accountability in the above definitions. It holds the following features: (1) *linkage identities*: each action is linked to an entity that performed it; (2) *reliable evidence*: the system maintains a record of past actions such that entities cannot secretly omit, falsify, or tamper; (3) *policy compliance*: the evidence can be inspected for signs of faults; (4) *detection*: when a judge detects a fault, it can obtain an alert of the fault that can be verified independently by a third party; (5) *punishment*: a proper sanction can be applied to misconduct entities.

4 ACCOUNTABLE DECRYPTION

In this section, we formally define the accountability of decryption. Our definition fits in the scenario where users choose a delegate to store their encrypted data and are aware of any subsequent decryption actions of the encrypted data. The process involves multiple roles: *encryptor* (E), *decryptor* (D), and *judge* (J). E and D are entities that create ciphertexts and perform the decryption of ciphertexts. J is responsible for detecting the misbehavior and imposes penalties against D who conduct such misbehavior. Accountable decryption requires five protocols.

- **Setup.** $params \leftarrow \text{Setup}(1^\lambda)$; The system takes security parameters λ and outputs the public parameters $params$. Here, $params$ is the default input for the rest of the algorithms and is thus ignored for simplicity.

- **Encryption.** $(ct, \mathcal{P}) \leftarrow \text{Enc}(aux, m)$; An encryptor E executes this algorithm to generate a ciphertext ct by using a message m and decryptor's auxiliary data aux , and policies \mathcal{P} . Here, \mathcal{P} dictates what are legal actions. Notably, $params$ is a public parameter for all algorithms. In contrast, aux consists of auxiliary parameters for certain encryption tasks.
- **Decryption.** $(m, \pi) / \perp \xleftarrow{\tilde{e}} \text{Dec}(key, ct)$; A decryptor D executes this algorithm within a designated environment, denoted as \tilde{e} . It takes as input a private key key , a ciphertext ct , and the encapsulated environment variables of \tilde{e} . If the decryption is successful, it outputs a plaintext m and a piece of evidence π , which serves as substantiation for validating the decryption process. Otherwise, it outputs \perp , indicating an abort. Notably, \tilde{e} captures critical aspects of the event, including precise timing, the unfolding sequence, and the identities of participating entities. The symbol \Leftrightarrow denotes a bidirectional relationship. Specifically, $\pi \Leftrightarrow \tilde{e}$ signifies that π captures all occurrences of \tilde{e} , and simultaneously, every state of \tilde{e} is accurately represented and recorded in π .
- **Check violation.** $tag \leftarrow \text{CheckViolation}(\pi, ct, \mathcal{P})$; A judge J executes this algorithm to scrutinize the actions of the decryptor, ensuring compliance with the predefined policies. It takes as input the evidence of decryption π , the policy \mathcal{P} and outputs a result $tag \in \{true, false\}$. Here, $false$ indicates the decryptor's action is aligned with policies, and we define it as legitimate decryption.
- **Punishment.** $true/false \leftarrow \text{Punish}(tag, \mathcal{P})$; J imposes penalties against D in case of non-compliance. Here, $true$ indicates penalties are imposed successfully.

Our definition follows the general principle summarized in Section 3. It has the ability to *trace*, *identify*, and *punish* malicious decryptors. In particular, if D maliciously accesses the plaintext, J provides feedback on D's actions by checking evidence π , and then imposes penalties upon D for its misbehaved decryption.

4.1 Ideal Accountable Decryption

We formalize the ideal notion of accountable decryption in terms of two properties: *non-repudiation* which says that decryptors who did illegal decryption can be identified and penalized, and *non-frameability*, which says that decryptors who performed legal decryption should not be blamed or subjected to false accusations.

To formalize non-repudiation, we consider a game in which the adversary takes on the role of the decryptor and is allowed to interact with other honest roles. The adversary wins if it illegally decrypts a ciphertext (i.e., $\tilde{e} \notin \mathcal{P}$) without facing punishment (the *punishment* protocol outputs *false*). Here, \in indicates the decryption is authorized and complies with policy, while \notin signifies a deviation from the policy.

Definition 9 (Non-repudiation). Define $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{nrp}}}(\lambda) = \Pr[\text{G}_{\mathcal{A}}^{\text{nrp}}(\lambda)]$, where $\text{G}_{\mathcal{A}}^{\text{nrp}}(\lambda)$ is defined as follows:

$$\begin{array}{l} \text{G}_{\mathcal{A}}^{\text{nrp}}(\lambda) \\ \hline params \leftarrow \text{Setup}(1^\lambda) \\ (ct, \mathcal{P}) \leftarrow \text{Enc}(aux, m) \\ (m, \pi) \xleftarrow{\tilde{e}} \mathcal{A}^{\text{Dec}(key, ct)} \\ tag \leftarrow \text{CheckViolation}(\pi, ct, \mathcal{P}) \\ \text{return } (false = \text{Punish}(tag, \mathcal{P}) \wedge (\tilde{e} \notin \mathcal{P}) \wedge (\pi \Leftrightarrow \tilde{e})) \end{array}$$

A protocol satisfies non-repudiation, if for all probabilistic polynomial-time (PPT) adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ satisfying $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{nrp}}}(\lambda) < \text{negl}(\lambda)$.

Next, to formalize non-frameability, we consider an adversary aiming to frame an honest decryptor by generating evidence of its "misbehavior", to ensure the unjust punishment of the honest decryptor. We set up a game where the adversary can assume any role in the system, including that of the encryptor and the judge, and interact with an honest decryptor. The adversary wins by producing evidence of its "misbehavior" and inducing punishment upon the legal decryptor who decrypted a ciphertext in accordance with the policy i.e., $\tilde{e} \in \mathcal{P}$.

Definition 10 (Non-frameability). Define $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{nfrm}}}(\lambda) = \Pr[\text{G}_{\mathcal{A}}^{\text{nfrm}}(\lambda)]$, where $\text{G}_{\mathcal{A}}^{\text{nfrm}}(\lambda)$ is defined as follows:

$$\begin{array}{l} \text{G}_{\mathcal{A}}^{\text{nfrm}}(\lambda) \\ \hline params \leftarrow \text{Setup}(1^\lambda) \\ (ct, \mathcal{P}) \leftarrow \text{Enc}(aux, m) \\ (m, \pi) \xleftarrow{\tilde{e}} \mathcal{A}^{\text{Dec}(key, ct)} \\ tag \leftarrow \text{CheckViolation}(\pi, ct, \mathcal{P}) \\ \text{return } (true = \text{Punish}(tag, \mathcal{P}) \wedge (\tilde{e} \in \mathcal{P}) \wedge (\pi \Leftrightarrow \tilde{e})) \end{array}$$

A protocol satisfies non-frameability, if for all PPT adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ satisfying $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{nfrm}}}(\lambda) < \text{negl}(\lambda)$.

Definition 11 (Ideal Accountable Decryption, IAD). A system achieves IAD if it satisfies non-repudiation and non-frameability.

Definition 11 seems sufficient for defining the accountability of decryption. However, in practice, decryption occurs locally. If a decryptor withholds π or provides false π (i.e., π cannot faithfully reflect \tilde{e} , namely $\pi \not\Leftrightarrow \tilde{e}$) during the decryption process, misbehaviors will never be captured.

4.2 Practical Accountable Decryption

To address this challenge, we introduce a new role to manage key and π , and we call this role *trustee* (TS). Accordingly, we introduce one protocol (called, EvidenceKeyManage) with two interactive sub-protocols.

- **Management.** $(key', \pi') \xleftarrow{\tilde{e}_1} \text{Manage}(ct, \mathcal{P})$; A trustee TS executes this protocol, producing a key called key' and a piece of evidence π if $\tilde{e}_1 \in \mathcal{P}$.
- **Key generation.** $key \leftarrow \text{KeyGen}(key')$; A decryptor D derives its decryption key from key' .

In the scenario above, D must obtain the key from TS before executing decryption, preventing the issue of evidence withholding. However, this brings a new problem: TS may behave maliciously. Therefore, we propose two definitions

for accountable decryption, considering a trusted TS and an untrusted TS.

A trusted TS must fulfill a new requirement: TS should provide authentic keys and evidence. Formally, we define this requirement by *evd-key-soundness* game. In this game, we consider a game in which an adversary plays the role of TS and is allowed to interact with other honest roles. The adversary wins if it generates illegal evidence that cannot represent its decryption behavior ($\pi^* \not\Leftarrow \tilde{e}_*$) or generates an illegal key that cannot decrypt the ciphertext.

Definition 12 (Evd-key-soundness). Suppose $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{sound}}}(\lambda) = \Pr[\text{G}_{\mathcal{A}}^{\text{sound}}(\lambda)]$, where $\text{G}_{\mathcal{A}}^{\text{sound}}(\lambda)$ is defined as follows:

$$\begin{array}{l} \text{G}_{\mathcal{A}}^{\text{sound}}(\lambda) \\ \hline \text{params} \leftarrow \text{Setup}(1^\lambda) \\ (ct, \mathcal{P}) \leftarrow \text{Enc}(aux, m) \\ (key^*, \pi^*) \xleftarrow{\tilde{e}_*} \mathcal{A}^{\text{Manage}(ct, \mathcal{P})} \\ key \leftarrow \text{KeyGen}(key^*) \\ \text{return } (\pi^* \not\Leftarrow \tilde{e}_*) \vee (\perp = \text{Dec}(key, ct)) \end{array}$$

A protocol satisfies *evd-key-soundness*, if for all PPT adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ satisfying $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{sound}}}(\lambda) < \text{negl}(\lambda)$.

Definition 13 (Practical Accountable Decryption with trusted Trustee, PAD-tTS). A system achieves PAD-tTS if it satisfies *non-repudiation*, *non-frameability* and *evd-key-soundness*.

We now consider the worst-case scenario where TS is not fully trusted. (1) if TS is compromised, it attempts to decrypt the ciphertext using D's partial key, key^* ; (2) TS, when compromised, also tries to evade detection of its compromised status. To capture them, we define two additional properties: *compromise-security*, which says that even if TS is compromised, the compromised TS cannot learn the plaintext based on D's partial key key^* ; *compromise-awareness*, which say that it has a certain probability of detecting and learning the TS status of being compromised.

In the *compromise-security* game, we consider a game in which an adversary takes on the role of TS and is allowed to interact with other honest roles. The adversary wins if it can learn D's plaintext based on the user's partial key.

Definition 14 (Compromise-security). Define $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{csec}}}(\lambda) = \Pr[\text{G}_{\mathcal{A}}^{\text{csec}}(\lambda)]$, where $\text{G}_{\mathcal{A}}^{\text{csec}}(\lambda)$ is defined as follows:

$$\begin{array}{l} \text{G}_{\mathcal{A}}^{\text{csec}}(\lambda) \\ \hline \text{params} \leftarrow \text{Setup}(1^\lambda) \\ (ct, \mathcal{P}) \leftarrow \text{Enc}(aux, m) \\ (key^*, _) \xleftarrow{\tilde{e}_1} \mathcal{A}^{\text{Manage}(ct, \mathcal{P})} \\ key \leftarrow \mathcal{A}^{\text{KeyGen}(key^*)} \\ \text{return } m = \mathcal{A}^{\text{Dec}(key, ct)} \end{array}$$

A protocol satisfies *compromise-security*, if for all PPT adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ satisfying $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{csec}}}(\lambda) < \text{negl}(\lambda)$.

We further introduce an additional detection protocol to capture *compromise-awareness*, namely, $\text{true/false} \leftarrow \text{Trace}(key, D)$; The algorithm outputs *false* only when the provided key key is the same as the decryption key D .

$\text{false} \leftarrow \text{Trace}(key, D)$ means that the misbehavior \tilde{e} is not captured. We then consider a game in which an adversary takes on the role of TS and is allowed to interact with other honest roles. The adversary wins if it generates a π which cannot reflect \tilde{e} , namely $\pi \not\Leftarrow \tilde{e}$, and such misbehavior cannot be detected.

Definition 15 (Compromise-awareness). Suppose $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{aware}}}(\lambda) = \Pr[\text{G}_{\mathcal{A}}^{\text{aware}}(\lambda)]$, where $\text{G}_{\mathcal{A}}^{\text{aware}}(\lambda)$ is defined as follows:

$$\begin{array}{l} \text{G}_{\mathcal{A}}^{\text{aware}}(\lambda) \\ \hline \text{params} \leftarrow \text{Setup}(1^\lambda) \\ (ct, \mathcal{P}) \leftarrow \text{Enc}(aux, m) \\ (key^*, \pi^*) \xleftarrow{\tilde{e}_*} \mathcal{A}^{\text{Manage}(ct, \mathcal{P})} \\ D \leftarrow \text{KeyGen}(key^*) \\ key \leftarrow \mathcal{A}^{\text{KeyGen}(key^*)} \\ \text{return } (\pi^* \not\Leftarrow \tilde{e}_*) \wedge (\text{false} = \text{Trace}(key, D)) \end{array}$$

A protocol satisfies *compromise-awareness*, if for all PPT adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\lambda)$ satisfying $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{aware}}}(\lambda) < \text{negl}(\lambda)$.

Definition 16 (Practical Accountable Decryption with untrusted Trustee, PAD-uTS). A system achieves PAD-uTS if it satisfies *compromise-security* and *compromise-awareness*.

5 A SECURE CONSTRUCTION

Philosophy of our design. The main idea behind our construction is to establish a dependable party TS using TEEs [13][14][15]. Traditional methods of ensuring accountability in decryption rely heavily on a judge to verify the legality of each decryption. However, even if the judge is assumed to be fully honest, obtaining evidence is challenging if a decryptor refuses to comply with the protocol or provides falsified information. To address this issue, we introduce a new role called trustee operating within TEEs. It receives decryption requests from decryptors and ensures the visibility of these requests to the judge. By integrating TEEs, our construction improves the reliability of trustees, implementing tamper-evident measures during key distribution to enhance key security and traceability in the distribution process. We have observed that TEEs are susceptible to numerous attacks such as [16], [17], [18]. Achieving a completely secure TEE remains a challenge. In response, we propose a design that enables the detection of compromised trustees.

Strawman protocol. Our naive solution involves four steps. ① E, defines the decryption policy, encrypts a message, and sends the ciphertext to D. ② To decrypt the ciphertext, D submits a decryption request to TEE-based TS with a signature to prove the identity. ③ TEE retrieves the decryption key, generates a piece of evidence about the key request, stores it in a log, and then sends the decryption key to D. Then, D decrypts the ciphertext. ④ J checks the evidence and imposes penalties against malicious decryption. The above approach immediately achieves Definition 12 when we assume that TEE is fully trusted. Under this assumption, the codes are executed as loaded, ensuring the integrity of the execution; when a key request is made by the decryptor, both the evidence and the key are produced accurately.

Technical challenges. Definition 16 further assumes that the trustee can be compromised. Indeed, this is aligned with the fact that TEE products suffer from architectural vulnerabilities [30], side-channel attacks [31], and fault attacks [16]. Therefore, we need to provide a solution that satisfies Definition 16. However, a few technical challenges still exist in building a fully-fledged solution.

Challenge-1: How to protect decryption keys under compromised TEEs. In the above approach, if a TEE is compromised, the adversary can access private keys. This not only renders the accountable mechanism ineffective but also jeopardizes the confidentiality of ciphertexts. In fact, this issue is a major drawback of existing TEE-based accountability systems, such as those presented in [32], [4].

Inspired by the key management solutions used in [21], [22], we introduce a key-splitting mechanism wherein TEE only stores a partial key while the user holds another partial one. Even if an attacker accesses the key inside TEEs, it cannot obtain a full decryption key.

Challenge-2: How to detect the compromised TEE. Within the present design paradigms, TEE functions as an encapsulated entity (black box), precluding any interrogation of its internal states by unauthorized peripheral components. However, this attribute also engenders a notable challenge: the dearth of visibility for overseers seeking potential compromise or deviations from established protocols.

We develop a detection algorithm to find potential compromises by inspecting TEE’s outputs. Our algorithm is based on the concept of deception technology [33]: using tactics to deceive malicious TEE into attacking the wrong targets and thus producing potentially useful information. To create wrong targets, we introduce two new roles: encryption inspector P_e and decryption inspector P_d . They emulate conventional encryptors and decryptors, respectively. If TS outputs inconsistent results, P_e and P_d will learn potential compromise of TEE. Beyond that, by leveraging the commitment scheme, our newly proposed key generation algorithm guarantees that TS cannot generate a decryption key identical to the key requested by users. Thus, a pair of private keys, which are honestly generated under D’s request and maliciously generated by TS, can serve as evidence of a fraudulent TS.

Overview of our construction. In our design, TS comprises two essential elements: *private key generator* (PKG) and *log manager* (LM). These elements can operate on the same server but within distinct domains for specific functionalities. PKG is executed within TEEs and handles the generation of decryption keys. LM operates in untrusted environments and is responsible for maintaining and updating the evidence π related to the key extraction process.

Threat model. We assume that the decryptor’s public key can be linked to a unique identifier. This assumption is common in the literature and reasonable in practice. Instead of fully trusting the TEE like previous TEE-based solutions [3][4][34][32], we make TEE work in the “trust, but verify” model where a compromised state is detectable. Also, we require cryptographic primitives that we rely upon to be provably secure.

At a high level, our system studies as follows. Initially, E defines the condition (i.e., \mathcal{P}) of key extraction; TS runs

PKG in TEE while providing public interfaces. Then, ① E encrypts a message and sends the corresponding ciphertext ct and \mathcal{P} to D. ② for decrypting ct , D sends a key request message with a commitment for his random number to TS. ③ TS generates decryption keys and updates the evidence π of key extraction. Specifically, LM first stores D’s key request in a log and transforms it as evidence π , while PKG generates the partial decryption key based on the received commitment. ④ J traces log to find the misbehavior of decryption and imposes penalties against the decryptor. ⑤ the inspectors (i.e., E, D, J) identify and prove the guilty of dishonest/compromised TEE who suppressed the evidence/key and provided the forged evidence/key.

Our full design. Our solution generally follows the straw-man approach mentioned in Section 5. The concrete decryption scheme follows the design of identity-based encryption (IBE) [35][36]. We run PKG inside TEE, and thus, the master private key of IBE is securely protected. Every ciphertext is encrypted using the decryptor’s identity and a serial number. To perform decryptions, a decryptor D has to request the key from PKG. An essential aspect of this procedure is that the request is forced to be stored in a log, and PKG issues a decryption key only if the request is proven to be securely stored. The log is structured as an append-only Merkle tree, as used, for example, in certificate transparency [37] to improve security and efficiency. Our detection mechanism utilizes a *weak black-box tracing algorithm* to trace the decryption box generated from the compromised hardware. The technical details of our approach are described below.

- (1) *Key management.* E encrypts a message using the receiver’s identity (e.g. email addresses) and E-generated serial number. This brings two benefits. First, it eliminates the need for E to request a new key for every encryption operation and thus improves practicality. Second, due to the nature of IBE scheme, the dynamic serial number makes a fresh key. D never knows the serial number until he obtains the ciphertext, which forces him to request a new key each time for decryption and then leaves a piece of unique evidence. Meanwhile, we require TS to only maintain a *partial* private key. Through a local calculation process, D combines his partial key with that of TS to obtain the final private key. This design ensures that, even if an attacker fully gains access to the TEE-based PKG, it cannot deduce the complete decryption key (adhering to Definition 16.vi).
- (2) *Evidence management.* We use an append-only Merkle tree based log $\overline{\mathcal{LOG}}$ to store the decryption evidence π . Whenever a decryption occurs, trusted hardware updates the Merkle tree by extending the leaf from the rightmost branch. This unique structure enables the generation of two distinct types of proofs to ensure the integrity of the evidence: proof of presence ρ guarantees that the specific evidence is stored within a tree; the proof of extension ε proves that such a tree is an append-only extension of $\overline{\mathcal{LOG}}$. We give an example to explain how to verify the ρ and ε . As shown in Figure 2, it depicts the Merkle tree before and after the insertion of node N_7 . The ε is valid only if the root hash H is the hash of $\{H(6, 6), H(4, 5), H(0, 3)\}$, and the ρ is valid

only if the current root hash H' is the hash of $\{H(7, 7), H(6, 6), H(4, 5), H(0, 3)\}$.

- (3) *Decryption checking.* Based on the decryption activities, a judge J conducts log queries, thereby acquiring supplementary details pertinent to the decryption event. Subsequently, J compares these supplementary details with predefined decryption policies to validate the legitimacy of the decryption process.




Figure 2: Merkle tree update

Our design also considered the worst-case scenario that a compromised TS may manipulate evidence and suppress evidence (adhering to Definition 16). We accordingly propose a *detection algorithm* to detect such misbehaviors. Central to this algorithm lies the deployment of *inspectors*. In a bid to mimic a conventional encryptor and decryptor, P_e and P_d dynamically engage in standard interactions with TS (behaving like a challenger). Based on the observations meticulously gathered by P_e and P_d throughout these interactions, our algorithm possesses the capability to effectively discern whether TS has been compromised.

- (1) *TS provides forged evidence or forged key.* When P_d requests a private key, it is required to submit a signature on the identity and timestamp. Then, signatures are organized into a secure Merkle tree structure by extending the tree to the right, following an append-only policy. This empowers P_d to detect any forged evidence originating from TS. Meanwhile, if the issued partial key lacks the capability to yield a complete key, then P_d knows TS has been compromised, e.g., TS provides an invalid partial key.
- (2) *TS suppresses the evidence or key.* An inspector P_d pretends to be a regular D and asks for a private key from TS. If the request is granted, but TS fails to show the necessary evidence or the decryption key, P_d knows that TS has hidden them.

We now delve into the most tricky scenario: TS effectively conceals both the key and evidence. If TS refrains from any communication with P_d , its malicious actions may remain undetected. In our solution, the private key is generated based on P_d 's commitment to an undisclosed random number, unbeknownst to TS. In the event that a decryptor, P_d , subsequently uncovers a new key not linked to the commitment, and no traces are left behind, it becomes evident that this key must have been clandestinely generated by TS. Therefore, for a given identity, the existence of a pair of private keys stemming from different sources acts as evidence against a deceitful TS. Admittedly, this relies on probabilistic grounds since P_d cannot guarantee a 100%

probability of finding a malicious key. Nevertheless, it at least exposes P_d to the risk of being detected.

- (3) *TS, privately generates key without leaving any evidence.* When P_d requests a private key, it is required to use a commitment to conceal a random number. Under normal circumstances, the private key should be tailored to this specific commitment. If a new key that is not associated with the commitment is later discovered, P_d will learn TS must have conducted an invalid decryption.

Notably, existing TEE-based studies assume that the hardware is fully trusted. However, in practice, real-world scenarios diverge from this ideal, as even reliable hardware like TEEs exhibit vulnerabilities, as evidenced by studies such as [30][31]. Our work takes a step forward and considers accountable decryption under the bad cases of trusted hardware being compromised. Our construction is critical for building practical decryption with accountability.

6 DETAILED PROTOCOL

In this section, we present our detailed protocol (see Figure 3). We instantiate TEE using Intel SGX [13].

Setup the system. $\text{Portex.Setup}(\lambda)$: This is a preparation phase aiming to set up the system. The algorithm takes as input a security parameter λ , and outputs public parameters $pm.s$. In particular, it first calls HW.Setup to initiate the HW and generates $pm.s_{hw}$ (L.1). Then, it runs HW.Load to load PKG code (e.g., $\text{Portex.EvidenceKeyManage}$ as defined in Figure 3) into a generation enclave GE (L.2). Next, it runs (“init”, λ) to set up the enclave and then publish the public keys (mpk, vk_{GE}) (L.3). Meanwhile, the log manager LM calls $\text{MT.Init}(\lambda)$ to initialize the log and runs $\text{S.KGen}(\lambda)$ to generate a signature key pair (vk_{LM}, sk_{LM}) (L.4). For D , it generate a key pair $(pk_{cli}^{enc}, sk_{cli}^{enc})$ and signing key pair $(vk_{cli}^{sign}, sk_{cli}^{sign})$ for secure data transfer (L.5-L.6). Here, parameters $(pm.s_{hw}, mpk, vk_{LM/GE}, pk_{cli}^{enc}, vk_{cli}^{sign})$ are publicly accessible, and default inputs of the rest of algorithms.

① **The encryption phase.** $\text{Portex.Enc}(ID, m)$: In this phase, E encrypts a secret/message m and generates a ciphertext ct and corresponding policy \mathcal{P} . In particular, it first generates a random serial number SN . Then, it runs $\text{IBE.Enc}(mpk, ID|SN, m)$ to encrypt a message m using mpk and a combination of user identity ID and SN (L.2).

② **The decryption phase.** A decryptor D recovers the plaintext, before which he needs to request the corresponding private key. This is an interactive protocol between TS (including KM, LM), and D with two algorithms as follows.

$\text{Portex.EvidenceKeyManage}(ID, SN, \sigma_{cli})$: The algorithm takes as input ID , SN , and D 's signature σ_{cli} , and outputs a decryption key $pkey$, and the evidence π of key extraction. Inspired by [36], we separate the key generation into three sub-algorithms $\{\text{IBE.KGen}_{D1}, \text{IBE.KGen}_{PKG}, \text{IBE.KGen}_{D2}\}$. The detailed algorithms are as follows.

- (L.1-L.3): D first locally runs IBE.KGen_{D1} to select (t_0, θ) and to generate a commitment C . Then, it generates a signature σ_{cli} on ID and SN . After that, it sends (ID, SN, C) , and σ_{cli} to LM.
- (L.5-L.7): LM runs MT.Insert . If the verification of σ_{cli} fails, the algorithm aborts. Otherwise, it runs MT.Insert

<p>Portex.Setup(λ)</p> <hr/> <pre> 1: PKG runs $pms_{hw} \leftarrow \text{HW.Setup}(\lambda)$ 2: $hdl_{GE} \leftarrow \text{HW.Load}(pms_{hw}, Q_{GE})$ 3: $mpk, vk_{GE} \leftarrow \text{HW.Run\&Quote}(hdl_{GE}, ("init", \lambda))$ 1: $H \leftarrow 0$ 2: $mpk, msk \leftarrow \text{IBE.Setup}(\lambda)$ 3: $vk_{GE}, sk_{GE} \leftarrow \text{S.KGen}(\lambda)$ 4: return mpk, vk_{GE} 4: LM runs $vk_{LM}, sk_{LM} \leftarrow \text{MT.Init}(\lambda)$ 5: CLIENT runs $pk_{cli}^{enc}, sk_{cli}^{enc} \leftarrow \text{PKE.KGen}(\lambda)$ 6: $vk_{cli}^{sig}, sk_{cli}^{sig} \leftarrow \text{S.KGen}(\lambda)$ </pre> <hr/> <p>Portex.CheckViolation(π, ct, \mathcal{P})</p> <pre> 1: $(ID^*, SN^*, \tau^*) \leftarrow \mathcal{LOG}$ 2: <i>if</i>(($ID^*, SN^*, \tau^* = null$)); return false 3: $(N, H', H, \rho, \varepsilon) \xleftarrow{parse} \pi$ 4: $(ID, SN, \tau, \sigma_{cli}) \xleftarrow{parse} N$ 5: $\pi InLog = ((ID, SN, \tau) = (ID^*, SN^*, \tau^*))$ $\wedge (true = \text{MT.Verify}(\pi))$ $\wedge (true = \text{S.Verify}(vk_{cli}^{sig}, \sigma_{cli}, ID SN))$ 6: return $((ID, \tau) \notin \mathcal{P}) \wedge (true = \pi InLog)$ </pre> <hr/> <p>Portex.Trace(key, D)</p> <pre> 1: $s, s' \xleftarrow{\mathbb{S}} \mathbb{Z}_p^*$ 2: $C_1 = X^s, C_2 = (g^{ID} \cdot Z)^s, C_3 = e(g, h)^{s'}$ 3: $m' \leftarrow \text{IBE.Dec}(mpk, D, C)$ 4: return $m' = m$ </pre> <hr/> <p>Portex.Enc(ID, m)</p> <pre> 1: $SN \xleftarrow{\mathbb{S}} \mathbb{Z}_m^*$ 2: $ct \leftarrow \text{IBE.Enc}(mpk, ID SN, m)$ 3: return ct, SN </pre>	<p>Portex.EvidenceKeyManage(ID, SN)</p> <hr/> <pre> 1: CLIENT runs $C \leftarrow \text{IBE.KGen}_{D1}(mpk)$ 2: $\sigma_{cli} \leftarrow \text{S.Sign}(sk_{cli}^{sig}, ID SN)$ 3: CLIENT sends $(ID, SN, C, \sigma_{cli})$ to LM. 4: <i>if</i>($true \neq \text{S.Verify}(vk_{cli}^{sig}, \sigma_{cli}, ID SN)$), <i>abort</i> 5: LM runs $\pi \leftarrow \text{MT.Insert}(ID, SN, \tau, \sigma_{cli})$ 6: $ir = (\pi, C), \sigma_{ir} \leftarrow \text{S.Sign}(sk_{LM}, ir)$ 7: LM sends (ir, σ_{ir}) to PKG 8: <i>quote</i> $\leftarrow \text{HW.Run\&Quote}(hdl_{GE}, ("Manage", ir, \sigma_{ir}))$ 1: <i>if</i>($true \neq \text{S.Verify}(vk_{LM}, \sigma_{ir}, ir)$), <i>abort</i> 2: <i>if</i>($H \neq H_{old}$), <i>abort</i> 3: $(\pi, C) \xleftarrow{parse} ir$ 4: $(N, H_{new}, H_{old}, \rho, \varepsilon) \xleftarrow{parse} \pi$ 5: $(ID, SN, \tau, \sigma_{cli}) \xleftarrow{parse} N$ 6: <i>if</i>($false = \text{MT.Verify}(\pi)$), <i>abort</i> 7: $H \leftarrow H_{new}$ 8: $pkey \leftarrow \text{IBE.KGen}_{PKG}(msk, ID SN, C)$ 9: $ct_{pkey} \leftarrow \text{PKE.Enc}(pk_{cli}^{enc}, pkey)$ 10: $\sigma_{pkey} \leftarrow \text{S.Sign}(sk_{GE}, ct_{pkey})$ 11: return ct_{pkey}, σ_{pkey} 9: CLIENT receives <i>quote</i> 10: $(md_{hdl}, H_Q, in, out(ct_{pkey}, \sigma_{pkey}), \sigma) \xleftarrow{parse} quote$ 11: <i>if</i>($H_Q \neq H_{GE}$), <i>abort</i> 12: <i>if</i>($false = \text{HW.VerifyQuote}(pms_{hw}, quote)$), <i>abort</i> 13: <i>if</i>($false \leftarrow \text{S.Verify}(vk_{GE}, \sigma_{pkey}, ct_{pkey})$), <i>abort</i> 14: $pkey \leftarrow \text{PKE.Dec}(sk_{cli}^{enc}, ct_{pkey})$ 15: $key \leftarrow \text{IBE.KGen}_{D2}(mpk, ID SN, pkey)$ </pre> <hr/> <p>Portex.Dec(key, ct)</p> <pre> 1: $m \leftarrow \text{IBE.Dec}(mpk, key, ct)$ 2: return m </pre> <hr/> <p>Portex.Punish(tag, \mathcal{P})</p> <pre> 1: <i>if</i>(tag) impose punishment 2: return tag </pre>
--	---

Figure 3: PORTEX protocol and generation enclave

adds a tuple $N(ID, SN, \tau, \sigma_{cli})$ to an append-only log \mathcal{LOG} . To be specific, LM first records the timestamp τ of the D's request. Then, it calls MT.Insert and outputs evidence π , which contains the newly inserted tuple N , current tree root hash H_{new} , old tree root hash H_{old} , proof of presence ρ , proof of extension ε . Next, LM signs a tuple $ir(\pi, C)$ by calling $\text{S.Sign}(sk_{LM}, ir)$ and obtains a signature σ_{ir} . Finally, LM sends (ir, σ_{ir}) to PKG.

- (L.8.1-L.8.6): Once PKG receives (ir, σ_{ir}) , it calls the $\text{Portex.EvidenceKeyManage}$ inside GE. Specifically, it first verifies σ_{ir} using vk_{LM} . If the above steps fail, the algorithm aborts. Next, GE runs MT.Verify to check the validity of the proof of presence ρ and the proof of extension ε . (L.8.7-L.8.11): If the evidence π is valid, GE selects (r', t_1) from \mathbb{Z}_p^* and starts to calculate a partial decryption key $pkey$ by calling IBE.KGen_{PKG} , where $pkey$ consists of (d'_1, d'_2, d'_3) , and C is used to calculate d'_1 . Then, GE encrypts $pkey$ to ct_{pkey} with D's

public key pk_{cli}^{enc} . PKG sends *quote* of ct_{pkey} to D.

- (L.9-L.15): D finally runs HW.VerifyQuote to verify *quote*. If the verification fails, the algorithm aborts. Otherwise, it calculates the final decryption key by calling IBE.KGen_{D2} . The final decryption *key* is calculated from $pkey$, and consists of (d_1, d_2, d_3) .
- 🔑 $\text{Portex.Dec}(key, ct)$: Now, D can access the secret m by running IBE.Dec with mpk and *key* (L.1). The algorithm outputs m if it is encrypted with $ID|SN$. Otherwise, it outputs \perp .
- ③ **The checking phase.** $\text{Portex.CheckViolation}(\pi, ct, \mathcal{P})$: In this phase, J examines the actions of the decryptor based on the evidence π stored in \mathcal{LOG} . Precisely, it finds π in \mathcal{LOG} , and then checks whether π satisfies \mathcal{P} (L.1-L.5). If misbehavior is found, the protocol outputs *true*. Otherwise, it outputs *false*.
- ④ **The punishment phase.** $\text{Punish}(tag, \mathcal{P})$: The judge J imposes penalties against D who has conducted malicious de-

<p>IBE.Setup(λ, n)</p> <hr/> 1: select bilinear groups $(\mathbb{G}, \mathbb{G}_T, e, p)$, where the $p > 2^\lambda$ 2: $x \xleftarrow{\$} \mathbb{Z}_p^*$ 3: $msk \leftarrow x$ 4: $X \leftarrow g^x$ 5: $g, h, Y \xleftarrow{\$} \mathbb{G}$ 6: $Z \leftarrow (Z_0, Z_1, \dots, Z_n) \xleftarrow{\$} \mathbb{G}^{n+1}$ 7: $mpk \leftarrow (X, Y, h, Z)$ 8: return mpk, msk <hr/> <p>IBE.KGen_{D_1}(mpk)</p> <hr/> 1: $t_0, \theta \xleftarrow{\$} \mathbb{Z}_p^*$ 2: return $h^{t_0} \cdot X^\theta$ <hr/> <p>IBE.KGen_{PKG}(msk, ID, C)</p> <hr/> 1: $r', t_1 \xleftarrow{\$} \mathbb{Z}_p^*$; 2: $d'_1 \leftarrow (Y \cdot C \cdot h^{t_1})^{1/x} \cdot H_Z(ID)^{r'}$ 3: $d'_2 \leftarrow X^{r'}$ 4: $d'_3 \leftarrow t_1$ 5: $pkey \leftarrow (d'_1, d'_2, d'_3)$ 6: return $pkey$ <hr/> <p>IBE.KGen_{D_2}($mpk, ID, pkey$)</p> <hr/> 1: $r'' \xleftarrow{\$} \mathbb{Z}_p^*$ 2: $r \leftarrow r' + r''$ 3: $d_1 \leftarrow \frac{d'_1}{g^\theta} \cdot H_Z(ID)^{r''}$ 4: $d_2 \leftarrow d'_2 \cdot X^{r''}$ 5: $d_3 \leftarrow d'_3 + t_1$ 6: $key \leftarrow (d_1, d_2, d_3)$ 7: checks $e(d_1, X) = e(Y, g) \cdot e(h, g)^{d_3} \cdot e(H_Z(ID), d_2)$ 8: return key <hr/> <p>H(a, b)</p> <hr/> 1: $c \leftarrow a b$ 2: return $\Pi.\text{Hash}^s(c)$	<p>IBE.Enc(mpk, ID, m)</p> <hr/> 1: $s \xleftarrow{\$} \mathbb{Z}_p^*$ 2: $C_1 \leftarrow X^s$ 3: $C_2 \leftarrow H_Z(ID)^s$ 4: $C_3 \leftarrow e(g, h)^s$ 5: $C_4 \leftarrow m \cdot e(g, Y)^s$ 6: $ct \leftarrow (C_1, C_2, C_3, C_4)$ 7: return ct <hr/> <p>IBE.Dec(mpk, key, ct)</p> <hr/> 1: $(d_1, d_2, d_3) \xleftarrow{\text{parse}} key$ 2: $m \leftarrow C_4 \cdot \left(\frac{e(C_1, d_1)}{e(C_2, d_2) \cdot C_3^{d_3}} \right)^{-1}$ 3: return m <hr/> <p>MT.Init(λ)</p> <hr/> 1: $\mathcal{LOG} \leftarrow \{\}$ 2: $vk_{LM}, sk_{LM} \leftarrow \text{S.KGen}(\lambda)$ 3: return vk_{LM}, sk_{LM} <hr/> <p>MT.Insert(T)</p> <hr/> 1: $n \leftarrow \text{number of leaf nodes}$ 2: <i>insert</i> T to \mathcal{LOG} 3: $H_{old} \leftarrow H(0, n-1), H_{new} \leftarrow H(0, n)$ 4: $\rho \leftarrow \{\text{hash nodes covering } [0, n]\}$ 5: $\varepsilon \leftarrow \rho \cup H(n, n)$ 6: $\pi \leftarrow (N, H_{new}, H_{old}, \rho, \varepsilon)$ 7: return π <hr/> <p>MT.Verify(π)</p> <hr/> 1: $(N, H_{new}, H_{old}, \rho, \varepsilon) \xleftarrow{\text{parse}} \pi$ 2: $H_{poe}, H_{pop} \leftarrow \phi$ 3: foreach $H_i \in \varepsilon$ 4: $H_{poe} \leftarrow \Pi.\text{Hash}(H_i, H_{poe})$ 5: endforeach 6: foreach $H_j \in \rho$ 7: $H_{pop} \leftarrow \Pi.\text{Hash}(H_j, H_{pop})$ 8: endforeach 9: return $(\Pi.\text{Hash}(N) \in \varepsilon \text{ and } H_{poe} = H_{old} \text{ and } H_{pop} = H_{new})$
--	--

Figure 4: IBE algorithm and evidence management algorithm

crypton. These penalties include measures such as deposit forfeiture and credit downgrade, which are excluded here for brevity.

⚠ The detection for the compromised trustee. For identifying potential compromises of TS, we introduce two sub-algorithms: *deterministic detection algorithm* and *probabilistic detection algorithm*. The deterministic detection algorithm infers a compromised state through a challenge-response mechanism, while the probabilistic detection algorithm identifies potential compromises by comparing the final keys issued by D with those issued by TS.

Deterministic detection algorithm. We make the inspector P_d pretend to be a regular decryptor D (as a challenger) and ask for the private key and evidence from the trustee TS. If any inconsistency is detected, it suggests that TS has

been compromised. In particular, if P generates a partial key without evidence in \mathcal{LOG} , TS must have forged the evidence (L.4 in Algorithm 1); if P's evidence is aligned with \mathcal{LOG} , but the partial key lacks the capability to yield P's complete key, TS has attempted to manipulate P's key (L.7 in Algorithm 1); if P's partial key arises but no evidence is shown, decryption evidence suppression has happened (L.11 in Algorithm 1); if evidence appears without P's partial key issuance, it implies decryption key concealment by TS (L.14 in Algorithm 1).

Probabilistic detection algorithm. We now consider the trickiest scenario, where TS hides the newly generated key as well as evidence. We argue if TS leaks any information on its misbehavior, it runs the risk of being caught and prosecuted. In the decryption phase, as (t_0, θ) are chosen randomly

Algorithm 1 The detection algorithm

```

wrongdoing  $\in \{true, false\}$ 
1: -----deterministic detection-----
2: upon receiving  $\pi$  and  $pkey$  from TS
3:    $key = \text{IBE.KGen}_{D_2}(mpk, ID, pkey)$ 
4:   if  $\pi \notin \mathcal{LOG} \wedge \text{Dec}(key, ct) \neq \perp$  then
5:     wrongdoing = true ▷ TS must have forged  $\pi$ 
6:   upon receiving  $\pi$  and  $pkey$  from TS
7:     if  $\pi \in \mathcal{LOG} \wedge \text{IBE.KGen}_{D_2}(mpk, ID, pkey) = \perp$  then
8:       wrongdoing = true ▷ TS must have forged  $pkey$ 
9:   upon receiving  $pkey$  from TS
10:   $key = \text{IBE.KGen}_{D_2}(mpk, ID, pkey)$ 
11:  if  $\text{find\_evidence}(\mathcal{LOG}) = \emptyset \wedge \text{Dec}(key, ct) \neq \perp$  then
12:    wrongdoing = true ▷ TS must have suppressed  $\pi$ 
13:  upon receiving  $\pi$  from TS
14:    if  $\text{find\_key}(\pi) = \emptyset$  then
15:      wrongdoing = true ▷ TS must have suppressed  $pkey$ 
16:  -----probabilistic detection-----
17:  upon finding  $key'$ 
18:    if  $key \neq key' \wedge \text{Dec}(key', ct) \neq \perp$  then
19:      wrongdoing = true

```

by P that are unknown to PKG, PKG cannot generate a decryption key that is same as key . This laid the foundation for detecting the PKG's misbehaviors. For a valid decryption key, d_3 contains the parameter t_0 that only knows to P. If any P finds another valid key but does not equal key , then it demonstrates that D is maliciously generated by TS, and thus shows TS's misbehavior.

7 SECURITY ANALYSIS

This section provides a security analysis. We consider two scenarios: one where the TEE remains secure, and another where the TEE is compromised.

Security analysis under secure TEEs. Assuming that TEE and the used cryptographic primitives are secure, our construction achieves PAD-tTS defined in Definition 13. We prove the theorem by contradiction. If an adversary \mathcal{A} made our construction fail to achieve properties in PAD-tTS, we could use such abilities to break our assumptions.

Theorem 1. *If HW is Exelnty secure and S is existential unforgeability under chosen message attack (EUF-CMA) secure, our construction achieves PAD-tTS as defined in Definition 13.*

Proof. *We prove this theorem by demonstrating that the probabilities of the adversary winning the non-repudiation game, non-frameability game, and evd-key-soundness game are negligible. We discuss each of these in turn.*

Non-repudiation: To win the game $G_{\mathcal{A}}^{\text{nrrep}}(\lambda)$, an adversary \mathcal{A} must make $\text{Punish}(tag, \mathcal{P})$ return false, after it illegally decrypts a ciphertext, i.e., $(m, \pi) \xleftarrow{e} \text{Dec}(key, ct) \wedge (e \notin \mathcal{P})$. The algorithm $\text{Punish}(tag, \mathcal{P})$ returning false implies that $\text{CheckViolation}(\pi, ct, \mathcal{P})$ returns false. The algorithm CheckViolation returns false when π does not exist in \mathcal{LOG} or \mathcal{A} has produced new illegal evidence π' . We define these two events as E_1 and E_2 . Their corresponding probabilities denoted as $\Pr[E_1]$ and $\Pr[E_2]$. These two events are independent and mutually exclusive, and thus $\text{adv}_{\mathcal{A}}^{\text{nrrep}}(\lambda) = \Pr[E_1] + \Pr[E_2]$.

We now prove that the probabilities of E_1 and E_2 occurring are negligible, i.e., $\Pr[E_1]$ and $\Pr[E_2]$ are negligible.

E_1 : The happening of E_1 indicates that \mathcal{A} has successfully obtained the decryption key while LM does not provide valid decryption evidence. As shown in Figure 3, when LM cannot provide valid decryption evidence in the form of (ir, σ_{ir}) . the algorithms L.8.1 and L.8.6 are programmed to abort. This further leads to the failure of algorithm L.8.8.

To prove that the occurrence of E_1 is negligible, we transform an adversary \mathcal{A} against $G_{\mathcal{A}}^{\text{nrrep}}(\lambda)$ into adversaries against the Exelnty and EUF-CMA of \hat{S} . We describe a sequence of games following the methodology used in [38], [39].

Game \mathcal{D}_0 . This is an unmodified game. Trivially, the winning probability of this game equals the advantage of \mathcal{A} against the non-repudiation game, namely, $\text{adv}_{\mathcal{A}}^{\text{nrrep}}(\lambda)$. For simplification, we extract the key algorithms as follows.

$$\left[\begin{array}{l} \text{if } (true \neq S.\text{Verify}(vk_{LM}, \sigma_{ir}, ir)), \text{abort} \\ \dots \\ ct_{pkey} \leftarrow \text{PKE.Enc}(pk_{cli}^{enc}, pkey) \\ \sigma_{pkey} \leftarrow S.\text{Sign}(sk_{GE}, ct_{pkey}) \\ \text{return } ct_{pkey}, \sigma_{pkey} \end{array} \right]$$

Game \mathcal{D}_1 . We introduce a minor modification to the previously described game by removing the conditional statement. Now, the system will always abort regardless of the input provided.

$$\left[\begin{array}{l} \text{if } (\overline{true \neq S.\text{Verify}(vk_{LM}, \sigma_{ir}, ir)}), \text{abort} \\ \dots \\ ct_{pkey} \leftarrow \text{PKE.Enc}(pk_{cli}^{enc}, pkey) \\ \sigma_{pkey} \leftarrow S.\text{Sign}(sk_{GE}, ct_{pkey}) \\ \text{return } ct_{pkey}, \sigma_{pkey} \end{array} \right]$$

In **Game \mathcal{D}_1** , the code undergoes modification, which stands in contradiction to our initial assumption that the Trusted Execution Environment (TEE) maintains Exelnty security. Therefore, the integrity constraints of **Game \mathcal{D}_1** are directly tied to the adversary \mathcal{A} 's ability to breach the execution integrity of the TEE. Formally, this relationship is expressed as $\text{adv}_{\mathcal{A}}^{\mathcal{D}_1}(\lambda) \leq \text{Adv}_{\mathcal{A}, \text{HW}}^{\text{Exelnty}}(\lambda)$.

Game $\mathcal{D}_0 \approx \text{Game } \mathcal{D}_1$ This Transition is based on failure events. In **Game \mathcal{D}_1** , the system always aborts. These two games interact with the attacker in exactly the same way, except in the scenario where a valid signature σ_{ir} is provided. Thus, the winning condition for \mathcal{D}_0 is equal to the winning condition for \mathcal{D}_1 if and only if the σ_{ir} is valid. Suppose \mathcal{A} has successfully provided a forged signature of the log manager, i.e., σ_{ir} without the signing key. We can build another adversary \mathcal{B} to break the security of EUF-CMA of S . Thus, we have

$$|\Pr[\mathcal{D}_0] - \Pr[\mathcal{D}_1]| = \Pr[\text{Sig-forge}_{\mathcal{A}, S}^{\text{CMA}}(\lambda) = 1].$$

Putting the above games together, we get that

$$\Pr[E_1] = \Pr[\mathcal{D}_0] \leq \text{Adv}_{\mathcal{A}, \text{HW}}^{\text{Exelnty}}(\lambda) + \Pr[\text{Sig-forge}_{\mathcal{A}, S}^{\text{CMA}}(\lambda) = 1]$$

E_2 : The happening of E_2 indicates that \mathcal{A} has successfully obtained decryption key while \mathcal{A} has produced new illegal evidence π' . As shown in Figure 3, when a decryptor D cannot provide a valid signature in the form of σ_{cli} . The algorithm L.4 is programmed to abort. Suppose \mathcal{A} has successfully produced a new signature π' that can be accepted by LM, we can use this

ability to break the security of EUF-CMA of S. Thus, we have $\Pr[E_2] \leq \Pr[\text{Sig-forge}_{\mathcal{A},S}^{\text{CMA}}(\lambda) = 1]$.

Putting everything together, we have that

$$\begin{aligned} \text{adv}_{\mathcal{A}}^{\text{G}_{\text{A}}^{\text{nrp}}}(\lambda) &= \Pr[E_1] + \Pr[E_2] \\ &= \text{Adv}_{\mathcal{A},\text{HW}}^{\text{G}_{\text{A}}^{\text{Exelnty}}}(\lambda) + 2 \Pr[\text{Sig-forge}_{\mathcal{A},S}^{\text{CMA}}(\lambda) = 1] \\ &\leq \text{negl}(\lambda) \end{aligned}$$

Non-frameability: To win $\text{G}_{\mathcal{A}}^{\text{nrp}}(\lambda)$ game, an adversary must produce illegal evidence for an honest decryptor's "misbehavior" to ensure its punishment. According to the algorithm L.6, we know that `Portex.CheckViolation` returns true only when all three conditions defined in L.5 return true. We define three events for these conditions. In particular, we define the first event, where $(ID, SN, \tau) = (ID^*, SN^*, \tau^*)$ as event E'_1 , with its corresponding probability denoted as $\Pr[E'_1]$; the second event, where $(\text{true} = \text{MT.Verify}(\pi))$ as event E'_2 , with its corresponding probability denoted as $\Pr[E'_2]$; the third event, where $(\text{true} = \text{S.Verify}(vk_{cli}^{\text{sig}}, \sigma_{cli}, ID|SN))$ as event E'_3 , with its corresponding probability denoted as $\Pr[E'_3]$; The first and third events happen only \mathcal{A} generates a valid signature σ_{cli}^* without knowing sk_{cli}^{sig} , which can be used to break the security of EUF-CMA of S. The second event occurs only if the TEE codes are manipulated, which contradicts our assumption regarding the execution integrity of the TEE (Exelnty secure); Thus, we have $\text{adv}_{\mathcal{A}}^{\text{G}_{\text{A}}^{\text{nrp}}}(\lambda) \leq \text{Adv}_{\mathcal{A},\text{HW}}^{\text{G}_{\text{A}}^{\text{Exelnty}}}(\lambda) \leq \text{negl}(\lambda)$ and $\text{adv}_{\mathcal{A}}^{\text{G}_{\text{A}}^{\text{nrp}}}(\lambda) \leq \Pr[\text{Sig-forge}_{\mathcal{A},S}^{\text{CMA}}(\lambda) = 1] \leq \text{negl}(\lambda)$.

Evd-key-soundness: To win $\text{G}_{\mathcal{A}}^{\text{sound}}(\lambda)$ game, an adversary must produce a tuple (key^*, π^*) , that makes $(\pi^* \not\Leftarrow \tilde{e}_1)$ or $(\perp = \text{Dec}(key, ct))$. However, both events imply that an adversary has provided a valid signature σ_{cli} (as shown in EvidenceKeyManage algorithm L.4, without providing valid σ_{cli} , the protocol will abort), which breaks the security of EUF-CMA of S. Thus, we have $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{sound}}}(\lambda) \leq \Pr[\text{Sig-forge}_{\mathcal{A},S}^{\text{CMA}}(\lambda) = 1] \leq \text{negl}(\lambda)$. \square

Analysis under compromised TEE. We now consider cases of TS being compromised. We prove that even in this extreme case, our construction is still secure, and users can be aware of such compromises.

Theorem 2. If HW is Exelnty secure, the commitment scheme CM is perfect hiding, and S is EUF-CMA secure, our construction achieves PAD-uTS defined in Definition 16.

Proof. We prove this theorem by applying Lemma 1, Theorem 3, and Theorem 4. Specifically, we demonstrate that the probabilities of the adversary winning the compromise-security and compromise-awareness games are negligible. \square

Lemma 1. If CM is perfect hiding, then for any probabilistic polynomial-time (PPT) adversary \mathcal{A} , the adversary gains no information about t_0 .

Proof. In the decryption phase, the adversary \mathcal{A} receives the tuple (ID, SN, C) . Since the values ID and SN are independent of the committed value t_0 , they provide no information about t_0 . The commitment C is constructed using a perfect hiding commitment scheme [40]. Due to the perfect (unconditional) hiding property, the commitment C does not reveal any information about t_0 to the adversary. Thus, \mathcal{A} cannot gain any information about t_0 from the commitment. \square

Theorem 3. If CM is perfect hiding, HW is Exelnty secure, and S is EUF-CMA secure, our construction achieves compromise-security defined in Definition 14.

Proof. From the algorithm (L.2), we know that $m^* = C_4 \cdot (\frac{e(C_1, d_1)}{e(C_2, d_2) \cdot C_3^{d_3}})^{-1}$, where the provided key $key = (d_1 = (Y \cdot h^{t'+t_1})^{1/x} \cdot \text{Hz}(ID)^r, d_2 = X^r, d_3 = t' + t_1, msk = x)$. Here the provided key key is derived from decryptor's partial key key^* . Let the plaintext decrypted from ciphertext ct as m^* , the original message as m , and decrypt ct with the provided key key , then we have the following equation.

$$\begin{aligned} m^* &= C_4 \cdot (\frac{e(C_1, d_1)}{e(C_2, d_2) \cdot C_3^{d_3}})^{-1} = \frac{C_4 \cdot e(C_2, X^r) \cdot C_3^{t'+t'}}{e(C_1, \frac{d_1}{g^{\theta'}} \cdot \text{Hz}(ID)^{r'})} \\ &= m \cdot \frac{e(g, Y)^s \cdot e(\text{Hz}(ID)^s, X^r) \cdot e(g, h)^{s \cdot (t_1+t')}}{e(X^s, \frac{(Y R h^{t_1})^{1/x}}{g^{\theta'}} \cdot \text{Hz}(ID)^r)} \\ &= m \cdot \frac{e(g, Y)^s \cdot e(g, h)^{s \cdot (t_1+t')} \cdot e(X^r, \text{Hz}(ID)^s)}{e(g^{xs}, (\frac{Y R h^{t_1}}{g^{x\theta'}})^{\frac{1}{x}}) \cdot e(X^s, \text{Hz}(ID)^r)} \\ &= m \cdot \frac{e(g, Y)^s \cdot e(g, h^{t'} h^{t_1})^s}{e(g^s, \frac{Y h^{t_1} h^{t_0} \cdot X^{\theta}}{X^{\theta'}})} \\ &= m \cdot e(g^s, h^{(t'-t_0)} \cdot X^{(\theta'-\theta)}) \end{aligned}$$

The output satisfies $m^* = m$ when $t' = t_0$ and $\theta' = \theta$. From Lemma 1, for any plaintext $m \in \mathbb{G}_T$, there exists a value d_3 such that the ciphertext C_4 can be explained. The parameters r, t_0 , and t_1 are randomly selected from \mathbb{Z}_p^* , where the public parameters g and h are generators of a prime order group. \mathcal{A} outputs the value d_3 derived from the user-chosen t_0 with probability $1/p$. Thus, $\text{adv}_{\mathcal{A}}^{\text{G}_{\mathcal{A}}^{\text{csec}}}(\lambda) = \Pr[m^* = m] \leq 1/2^\lambda$. Therefore, the adversary \mathcal{A} wins the game $\text{G}_{\mathcal{A}}^{\text{csec}}(\lambda)$ with negligible probability. \square

Theorem 4. If HW is Exelnty secure and S is EUF-CMA secure, our construction achieves compromise-awareness defined in Definition 15.

Proof. We prove this theorem by demonstrating that the probability of the adversary winning the game $\text{G}_{\mathcal{A}}^{\text{aware}}(\lambda)$ is negligible. We consider two cases: (**Case-1**) TS does not provide any result at all, i.e., TS simultaneously suppresses the evidence and associated keys, and (**Case-2**) TS provides inconsistent or forged results.

Case-1: In this case, TS suppresses both the evidence and the associated keys. If TS fails to provide the private key or the corresponding evidence, or if it only provides the original evidence π^* , then the condition $(\pi^* \not\Leftarrow \tilde{e}_*)$ will evaluate to false. This allows the inspector \mathcal{P} to immediately detect that the TS is compromised. Hence, the adversary wins in this case with probability $\Pr[\mathcal{A} \text{ wins Case-1}] = 0$.

Case-2: We prove this case by demonstrating that the probability of false = Trace(key, D) is negligible. Given a decryption key $D = (D_1, D_2, D_3)$ generated by `Portex.EvidenceKeyManage`, where $D_1 = (Y \cdot h^{t_0+t_1})^{1/x} \cdot \text{Hz}(ID)^r$, $D_2 = X^r$, and $D_3 = t_0 + t_1$ for distinct $t_0, t_1 \xleftarrow{\$} \mathbb{Z}_p^*$ (as defined in `IBE.KGen`, Figure 4). From Lemma 1, t_0 is chosen uniformly at random from \mathbb{Z}_p^* , which implies that the probability of the provided key key satisfies false = Trace(key, D) is $\frac{1}{p}$, where $p > 2^\lambda$ (as defined in `IBE.Setup`, Figure 4) Therefore, the probability $\Pr[\mathcal{A} \text{ wins Case-2}] = \Pr[\text{false} = \text{Trace}(key, D)] \leq \frac{1}{p} \leq \frac{1}{2^\lambda}$.

Table 1: Theoretical complexity (L) and experimental time on IBE (R): (a) Line 2 of Portex.EvidenceKeyManage in Figure 3; (b) Line 9-14; (c) Line 4-7; (d) Line 1-8 of GE on input *Manage*; (e) It costs 3.0121 ms for running $KGen_{PKG}$ in CPU.

		Algorithm	Estimated Time (ms)	Experiment Time (ms)	Size	Environment	Symbol	Size (B)
Setup		HW.Load	-	342.9301	1.5MB	inside TEE		
		MT.Init	-	<0.001	<1KB	outside TEE	pk_{pke}	148
		PKE.KGen	-	1.0413	<1KB	outside TEE	sk_{pke}	180
		S.KGen	-	1.8399	<1KB	outside TEE	vk_{sign}	148
Enc		Enc.Setup	-	0.3119	-	outside TEE	sk_{sign}	180
		IBE.Enc	$2T_{bp} + 2T_{ep,g1} + 2T_{ep,gt} + T_{bm} \approx 1.8535$	2.0010	3.4KB	outside TEE		
EvidenceKeyManage	D	IBE.KGen $_{D1}$	$2T_{ep,g1} + T_{ep,gt} + T_{bm} \approx 1.2544$	1.3170	1.2KB	outside TEE		
		Manage.SendReq ^a	$T_{sign} \approx 1.6972$	7.1164	-	outside TEE		
		Manage.Verify ^b	$T_{verify} + T_{dec} \approx 7.5307$	17.9325	-	outside TEE	Symbol	Time (ms)
		IBE.KGen $_{D2}$	$4T_{bp} + 3T_{ep,g1} + T_{ep,gt} + 5T_{bm} \approx 2.9666$	3.1452	3.1KB	outside TEE	$T_{ep,g1}$	0.6028
	LM	Manage.LogGen ^c	$T_{sign} \approx 1.6972$	6.9579	-	outside TEE	$T_{ep,gt}$	0.0473
		Manage.LogVerify ^d	$T_{verify} \approx 6.0527$	10.6760	-	outside TEE	T_{bm}	0.0015
	PKG	IBE.KGen $_{PKG}$	$4T_{ep,g1} + 3T_{bm} + T_{hw} \approx 2.4252$	2.5676 ^e	3.0KB	inside TEE	T_{hw}	0.0096
		Manage.SendPkey	$T_{sign} + T_{enc} \approx 3.4893$	16.6631	-	outside TEE		
Dec		Dec.Setup	-	0.1128	-	outside TEE		
		IBE.Dec	$2T_{bp} + T_{ep,gt} + 3T_{bm} \approx 0.6035$	0.7079	3.9KB	outside TEE		
CheckViolation		CheckViolation	-	0.1599	-	outside TEE	Symbol	Time (ms)
Detection		Forged_evidence	-	1.0677	<1KB	outside TEE	T_{enc}	1.7921
		Forged_key	-	3.5497	<1KB	outside TEE	T_{dec}	1.4780
		Suppress_evidence	-	<0.001	<1KB	outside TEE	T_{sign}	1.6972
		Suppress_key	-	<0.001	<1KB	outside TEE	T_{verify}	6.0527
		LogInspect	$2 \cdot (4T_{bp} + T_{ep,gt} + 2T_{bm}) \approx 2.3076$	2.3586	3.7KB	outside TEE		

From the *Case-1* and *Case-2*, the adversary \mathcal{A} wins the game $G_{\mathcal{A}}^{\text{aware}}(\lambda)$ with negligible probability. \square

8 IMPLEMENTATION AND EVALUATION

We implement PORTEX in C++ and use Intel SGX SDK [24] as the TEE. The implementations of IBE algorithms, including key generation, encryption, decryption, and trace, are based on the Pairing-Based Cryptography (PBC) library [41]. Our Merkle Tree implementation is based on the merklecpp library [42]. The public key encryption and signature algorithms are based on the OpenSSL library [43]. The SGX Enclave does not support the original OpenSSL. Therefore, in the PKG program, we utilize SGXSSL as an alternative [44].

Performance&scalability. It involves two aspects: (a) evaluating decryption performance, misbehavior-checking performance, and detection performance for finding the compromised TEE; (b) examining how scalability is affected when the number of decryptors increases. We evaluate the performance of PORTEX on a machine with 3.5GHz Intel Xeon CPU (Ice Lake). The symmetric bilinear pairing $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is constructed on the curve $y^2 = x^3 + x$, and the security parameter λ is selected as 512. We report the average time of each algorithm by repeating the experiments 1,000 times.

Performance evaluation with static parameters. We first evaluate performance with static parameters. As an example, we consider the following scenario. We assume that there are 1000 decryptors, and each decryptor can launch

one decryption. Our evaluation shows that the full decryption process, including two sub-algorithms: *evidence-key management* and *decryption*, concludes within 10ms. In particular, evidence-key management takes approximately 1.31ms, 3.15ms, 2.56ms to finish IBE.KGen $_{D1}$, IBE.KGen $_{D2}$, and IBE.KGen $_{PKG}$, respectively; It takes about 1.69ms and 6.05ms to generate and verify the decryption evidence. Simultaneously, the decryption of the ciphertext exhibits remarkable efficiency, demanding less than 1ms: the Dec.Setup operation concludes within an incredibly short span of 0.1ms, while IBE.Dec finalizes in a mere 0.6ms.

Meanwhile, tracing malicious decryption is significantly efficient, as it only takes at most 0.002ms to trace the evidence of malicious decryption. The process of identifying potential forgery within the TEE unfolds in two stages. Specifically, detecting forged evidence and forged keys takes approximately 0.02 milliseconds and 0.01 milliseconds, respectively. Should the evidence or the key be suppressed, their compromise is swiftly ascertained. The probabilistic detection hinges on the methodology employed to distinguish these two elements. Remarkably, the comparison itself concludes in a mere 0.001 milliseconds upon the discovery by the decryptor D.

We then assess the execution time of each basic cryptographic primitive used in IBE, and report them in Table 1.right. In this table, we use the following symbols to represent the running time of each function. T_{bp} : The execution time of one bilinear pairing operation $e(P, Q)$, where $\{P, Q\} \in \mathbb{G}$; $T_{ep,g1}$: The execution time of one exponentiation operation P^x , where $P \in \mathbb{G}, x \in \mathbb{Z}_p$; $T_{ep,gt}$: Except for $P \in \mathbb{G}_T$, similar to $T_{ep,g1}$; T_{bm} : The execution

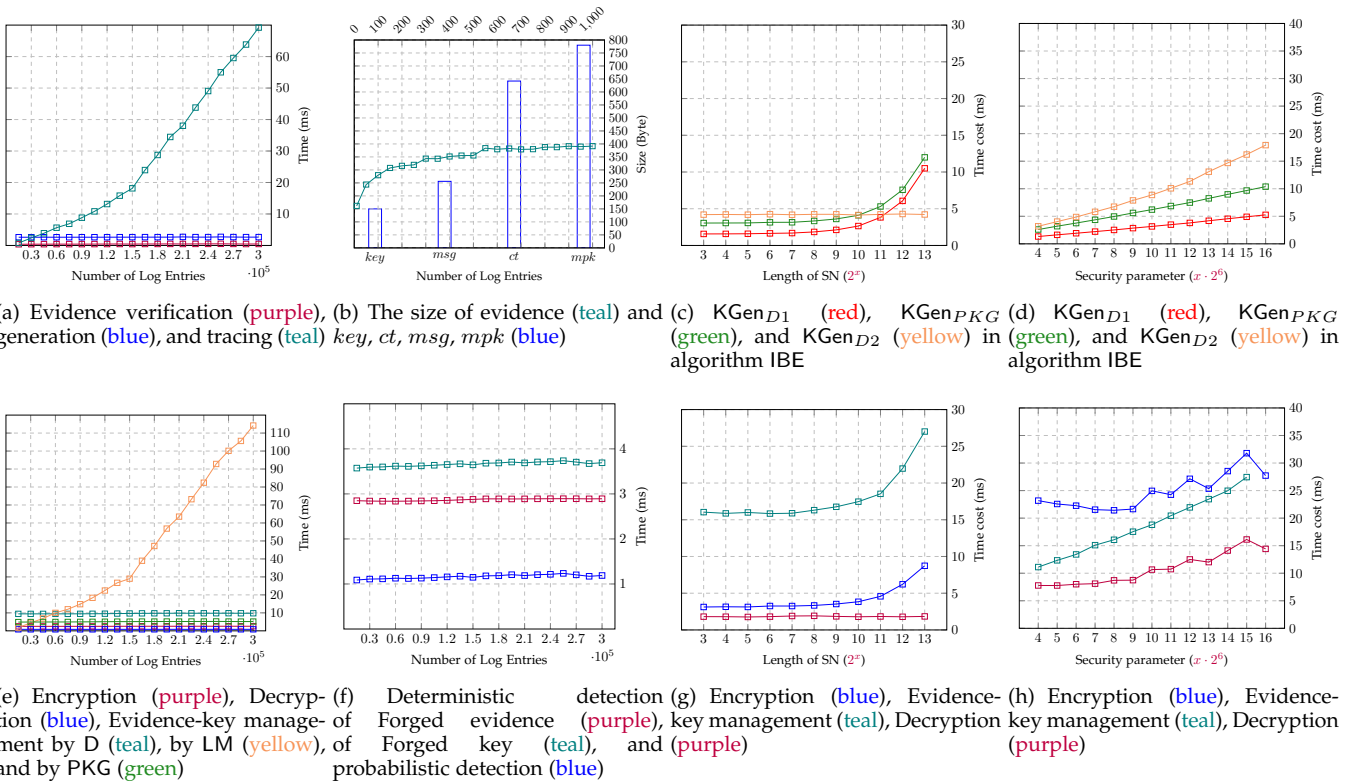


Figure 5: Experimental performance results

time of one scale multiplication operation $P \cdot Q$, where $\{P, Q\} \in \mathbb{G}$; T_{hvu} : The time of context switch to TEEs. These results closely align with our estimated times for each primitive, which offers robust evidence supporting the accuracy of our experimental findings.

Evaluation with dynamic parameters. To determine if deployment might be feasible in practice, we further evaluate how scalability is affected by dynamic parameters. We consider two scenarios: the performance overhead of decryption, evidence tracing, and compromise detection under (a) escalating sizes of security parameters and (b) an increasing number of decryption instances.

Intuitively, a larger security parameter enhances resilience against attacks. This enhancement comes at the cost of increased execution time for cryptographic operations. Empirical findings align with our expectations: *The time costs of encryption, evidence-key management, and decryption consistently increase with higher security parameter settings.* This phenomenon stems primarily from the augmented computational load of the underlying sub-algorithms. For example, the execution time of $IBE.KGen_{D_2}$ within the key-request algorithm expands from 4.2ms at a security parameter of 512 to 11.2ms at a heightened security parameter of 1024 (cf. Figure 5(c)). Nonetheless, the evidence tracing and compromise detection exhibit relatively steady performance, as they involve fewer time-intensive primitives.

We now consider the performance overhead arising from an escalating number of decryption operations (cf. Figure 5(a), 5(e) and 5(f)). To illustrate, we consider a specific scenario where the number of decryption operations grows from an initial 1,000 to a substantial 300,000. Our evalu-

ation reveals that encryption time remains invariant with increased decryption operations as it is independent with TS, relying solely on the decryptor’s ID and the encryptor’s sequence number provided. In contrast, the time required for decryption directly correlates with the historical count of decryption operations. In particular, the execution time of operations such as evidence generation and evidence verification (L.8.7-L.8.11 in Figure 3), both integral to TS, experiences an increment. Consequently, this rise in processing time within TS contributes to an overall elongation of the decryption process. Meanwhile, the execution time of the trace algorithm is directly correlated with the size of the log. As the log size increases, the time required to query a malicious decryption operation from the log also increases. For example, when the log contains 21,000 entries, the tracing algorithm takes 38.5ms to finish the query. Upon augmenting the log size of 300,000 entries, the execution time of the trace algorithm increases to 69.2ms. Notably, the time cost of compromise detection is stable. Also, to maintain the integrity of evidence during decryption, it is required to choose a new sequence number for each encryption. As depicted in Figure 5, when the sequence number length is increased, the running time for evidence-key management displays a clear increasing trend. Fortunately, the impact on the running time for encryption and decryption is relatively minimal.

Performance comparison with other solutions. Table 2 compares the performance of our proposed solution with that of existing works. As different works use different decryption algorithms, conducting a direct comparison based on actual runtime performance is unfair. Therefore, we examine the

Table 2: Comparison between ours and related studies by analyzing the computational complexity of functional operations, the operations including Inter-party Interactions (T_{inter}), Encryption (T_{enc}), Decryption (T_{dec}), Signing (T_{sign}), Verification (T_{verify}).

Solutions	Encryption	Decryption	Tracing
Ours \mathbb{P}_o	$T_{inter} + T_{enc} + T_{sign}$	$4T_{inter} + T_{enc} + 2T_{dec} + 3T_{sign} + 4T_{verify}$	$2T_{inter} + T_{sign} + 2T_{verify}$
Severinsen’s [32]	$2T_{inter} + T_{enc} + T_{sign}$	$6T_{inter} + T_{dec} + T_{sign} + 2T_{verify}$	$4T_{inter} + T_{sign} + T_{verify}$
Luo’s [45]	$2T_{inter} + T_{enc} + T_{sign}$	$7T_{inter} + T_{enc} + T_{dec} + T_{sign} + 2T_{verify}$	$5T_{inter} + T_{dec} + 2T_{sign} + 2T_{verify}$
Fialka [4]	$2T_{inter} + 2T_{enc} + 2T_{sign} + T_{verify}$	$N * T_{inter} + T_{dec} + T_{verify}$	$6T_{inter} + 2T_{dec} + T_{sign} + 2T_{verify}$

workflows of these protocols and quantify the frequency of critical operations within each protocol to analyze their computational costs. These critical operations contain inter-party interactions T_{inter} , encryption T_{enc} , decryption T_{dec} , signing T_{sign} , and signature verification T_{verify} . The inter-party interaction T_{inter} refers to the communication or exchange of information between different parties involved in a protocol, such as between the encryptor and potentially a third-party entity like a storage center or log manager.

Our solution has less interaction complexity compared to prior related works. For example, Severinsen’s protocol (\mathbb{P}_{sever}), Luo’s protocol (\mathbb{P}_{luo}), and Fialka (\mathbb{P}_{fialka}) mandate that the encryptor sends a copy of the ciphertext or auxiliary information to a storage center (termed as the Application Provider in \mathbb{P}_{sever} and \mathbb{P}_{luo} and Privacy-Preserving Smart Contracts (PPSC) in \mathbb{P}_{fialka} , cf. Figure 6) in addition to the intended recipient. In contrast, our protocol (\mathbb{P}_o) eliminates the need for encryptors to send an additional ciphertext backup, thereby obviating the requirement for a storage center and reducing interaction complexity. In \mathbb{P}_o , the encryption, decryption, and tracing operations require T_{inter} , $4T_{inter}$, and $2T_{inter}$, respectively.

Furthermore, we present the computational complexities associated with encryption and decryption in different protocols (see Table 2). The protocol \mathbb{P}_{sever} requires one encryption and one decryption ($T_{enc} + T_{dec}$), while the protocols \mathbb{P}_{luo} , \mathbb{P}_{fialka} , and \mathbb{P}_o each necessitate an additional encryption and decryption ($2T_{enc} + 2T_{dec}$). Although \mathbb{P}_{sever} involves fewer encryption and decryption operations, it does not consider the message transmission after encryption, rendering the protocol incomplete.

Space evaluation. We also consider the feasibility of storage requirements. First, we evaluated the storage usage across four types of data (namely, *key*, *ct*, *msg*, and *mpk*). Our evaluations were conducted using SHA256 hashing, which generates a 32-byte length digest. Results indicate that the storage load for *mpk* with 1,000 entries requires approximately 780 bytes of storage, which is the most significant usage we observed. Secondly, we evaluated the trend of evidence size as the number of log entries increased. We found that the trend remains relatively stable, indicating the evidence size will occupy a constant amount of storage.

9 DISCUSSION

We clarify the difference between the concepts of access control, traceability, and accountability, and provide potential applications to emphasize the necessity of our scheme.

9.1 Access Control, Traceability, and Accountability

Access control defines *who* can access data and resources and *when* this access is permitted, preventing unauthorized access before it occurs. Traceability records the actions related to data creation, modification, access, and transfer, coming into effect after resources have been accessed. Accountability ensures that individuals are held responsible for their actions, based on traceability, and emphasizes the consequences of inappropriate access (Table 3).

In the context of accountable decryption, the access control can be more fine-grained, e.g., one can even specify the legitimate decryption timeframe. For example, one can ask the decryptor to decrypt only between 3:00 PM and 5:00 PM on some day. Traceability underscores the recording of the specific decryption time to ensure a traceable trail. Meanwhile, accountability focuses on verifying whether decryption occurred within the specified time frame. If decryption takes place outside the permitted timeframe, the decryptor is accountable for its behavior and the owner of the data is aware of such behavior.

Table 3: Description and example

Concept	Description	Example
Access control	Governing who can access specific resources, defining user permissions, privileges, and restrictions.	A nurse may access patient records in his/her department but not in other departments.
Traceability	Providing a documented trail of how the data has been created, modified, accessed, or transferred.	A nurse accesses a patient’s record in his/her department. The system records the nurse’s name, time of access, and the reason for access (e.g., treatment, diagnosis).
Accountability	Checking the documented trail to make users accountable for any misuse.	A nurse has improperly accessed a patient’s records he/she is not allowed to access. The improper access to the data cannot be denied.

9.2 Potential Applications

Accountable ePHI. Electronic protected health information (ePHI) [46] refers to any health information that can identify a patient. Balancing ePHI privacy and accessibility is challenging. Sharing ePHI to all doctors increases the potential risk of information leakage. Conversely, sharing the needed health information with only authorized doctors and meanwhile making any improper access undeniable is crucial. While one can use conventional access control mechanisms to achieve the same goal, this can slow down treatment as additional access control procedures are required, which is not desirable during emergencies. Our system provides a solution to address the challenges. It allows for information sharing while auditing unauthorized access and potential

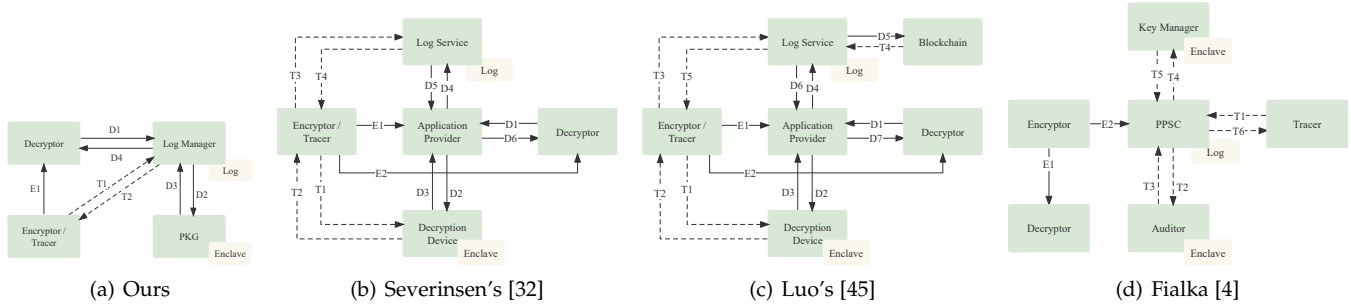


Figure 6: Comparing operational mechanism of Encryption-Decryption protocol ($\xrightarrow{E/D}$) and Tracing protocol (\xrightarrow{T})

breaches. We consider patients as encryptors (i.e., E) who possess sensitive data, such as congenital genetic disorders. E encrypts their data and stores it in the hospital’s cloud server, making it inaccessible to general practitioners (e.g., dental consultants). For instance, when a patient requires emergency treatment, dentists (i.e., D) can bypass rules and access ePHI. After that, a designated auditor J (e.g., regulatory authorities) accesses the audit trail to monitor whether patient data is being accessed for the right purpose.

Transparent KYC. Know Your Client (KYC) standards are employed in the investment industry to verify customer identities and assess their risk and financial profiles [47]. Customers are required to provide KYC information to their banks, enabling the latter to conduct anti-money laundering procedures. However, this practice raises concerns regarding the potential misuse of customer data. Our solution addresses these concerns by allowing the customers to upload encrypted data and become aware of who have accessed the data. This approach permits banks to access private information only when necessary.

Accountable warrant. In an ideal scenario, law enforcement agencies could access citizens’ data stored in the cloud center based on electronic warrants [48]. However, in practice, warrants are processed confidentially to maintain the integrity of investigations and abuse by malicious law enforcement agencies are not in place [49]. Our solution addresses these concerns by requiring the law enforcement (D) to maintain a documented trail when accessing an individual (E)’s encrypted data. This trail enables public verification of law agencies’ decryption activities and ensures the safe and legal use of warrants, thereby mitigating the risks of abuse.

10 RELATED WORK

Trusted hardware has been proposed as an effective tool for achieving accountability protocols. Pasture et al. [50] devised a secure messaging and logging library incorporating offline data access safeguarded by trusted hardware. Their approach ensures that access actions are irrefutable and revocations are verifiable. Ryan [3] proposed an accountable decryption protocol assisted by trusted hardware. Under the security of trusted hardware devices, the private key is stored inside TEEs, and the information of each decryption will be recorded in a transparent log. Severinsen [32] implemented Ryan’s protocol using Intel SGX [13], guaranteeing that encryptors can detect every decryption. In this scheme, decryptors can only decrypt the ciphertext after updating

the Merkle tree in logs. Liang et al. [51] proposed a decentralized accountable system for healthcare data using Intel SGX and blockchain. Luo et al. [45] applied the blockchain and SGX to an accountable decryption scheme. The decryptor is assisted by SGX, similar to Severinsen’s solution, but the log information is recorded on the blockchain. Fialka [4] employed a TEE-based confidential smart contract to make decryption accountable, where the transaction is used as evidence to trace a user’s decryption.

Different from prior studies that merely provide prototypes, our work formally defines the security properties of accountable decryption and proves the correctness of our concrete construction. Besides, this line of work often assumes that TEEs are fully trusted. If the hardware is compromised, it is not hard to see that the solutions no longer work. Worse still, overseers cannot learn whether the hardware has been compromised. To our knowledge, our scheme is the first work that considers compromised TEEs. This is achieved by introducing a detection algorithm as part of our solution.

11 CONCLUSION

In this paper, we focus on achieving the practical accountability of decryption operations. We introduce a novel set of definitions tailored for accountable decryption. Our definition aims to capture all possibilities and impossibilities in the formulation. We further construct a scheme aligning with the definitions by using trusted hardware. Our prototype and evaluations prove the practicability and efficiency.

REFERENCES

- [1] Giampaolo Bella, et al. Accountability protocols: Formalized and verified. *TISSEC*, 9(2):138–161, 2006.
- [2] Joan Feigenbaum, et al. Accountability and deterrence in online life. In *WEBSCI*, pages 1–7, 2011.
- [3] Mark D Ryan. Making decryption accountable. In *SPW*, pages 93–98. Springer, 2017.
- [4] Rujia Li, et al. An accountable decryption system based on privacy-preserving smart contracts. In *ISC*, pages 372–390, 2020.
- [5] Jonathan Frankle, et al. Practical accountability of secret processes. In *USENIX Security*, pages 657–674, 2018.
- [6] Andreas Haeberlen, et al. Peerreview: Practical accountability for distributed systems. *OSR*, 41(6):175–188, 2007.
- [7] Daniel J Weitzner, et al. Information accountability. *CACM*, 51(6):82–87, 2008.
- [8] Rajashekar Kailar. Accountability in electronic commerce protocols. *TSE*, 22(5):313–328, 1996.
- [9] Aydan R Yumerefendi, et al. The role of accountability in dependable distributed systems. In *Proceedings of HotDep*, volume 5, pages 3–3. Unseix Association, 2005.

- [10] Ruth W Grant, et al. Accountability and abuses of power in world politics. *APSR*, 99(1):29–43, 2005.
- [11] Joan Feigenbaum, et al. Towards a formal model of accountability. In *NSPW*, pages 45–56, 2011.
- [12] B Lampton. Notes for a presentation entitled “accountability and freedom.” <http://research.microsoft.com/en-us/un/people/blampton/sli des/AccountabilityAndFreedom.ppt>, 2005.
- [13] Victor Costan, et al. Intel SGX explained. *IACR Cryptology ePrint Archive*, 2016(86):1–118, 2016.
- [14] Sandro Pinto, et al. Demystifying ARM Trustzone: A comprehensive survey. *CSUR*, 51(6):1–36, 2019.
- [15] Dayeol Lee, et al. Keystone: An open framework for architecting trusted execution environments. In *EuroSys*, pages 1–16, 2020.
- [16] Zitai Chen, et al. VoltPillager: Hardware-based fault injection attacks against Intel SGX enclaves using the SVID voltage scaling interface. In *USENIX Security*, pages 699–716, 2021.
- [17] Jo Van Bulck, et al. A tale of two worlds: Assessing the vulnerability of enclave shielding runtimes. In *CCS*, pages 1741–1758, 2019.
- [18] Jo Van Bulck, et al. Foreshadow: Extracting the keys to the intel sgx kingdom with transient out-of-order execution. In *USENIX Security*, pages 991–1008, 2018.
- [19] Weili Wang, et al. Engraft: enclave-guarded raft on byzantine faulty nodes. In *CCS*, pages 2841–2855, 2022.
- [20] Fabian Fleischer, et al. Memory corruption attacks within android tees: a case study based on op-tee. In *ARES*, pages 1–9, 2020.
- [21] Craig Gentry. Certificate-based encryption and the certificate revocation problem. In *EUROCRYPT*, pages 272–293. Springer, 2003.
- [22] Sattam S Al-Riyami, et al. Certificateless public key cryptography. In *ASIACRYPT*, pages 452–473. Springer, 2003.
- [23] Ben Fisch, et al. Iron: Functional encryption using Intel SGX. In *CCS*, pages 765–782, 2017.
- [24] Intel SGX SDK. <https://software.intel.com/en-us/sgx-sdk>.
- [25] Helen Nissenbaum. Accountability in a computerized society. *Science and Engineering Ethics*, 2:25–42, 1996.
- [26] Ahto Buldas, Helger Lipmaa, and Berry Schoenmakers. Optimally efficient accountable time-stamping. In *PKCW*, pages 293–305. Springer, 2000.
- [27] Aydan R Yumerefendi, et al. Strong accountability for network storage. *TOS*, 3(3):11–es, 2007.
- [28] Andreas Haeberlen. A case for the accountable cloud. *OSR*, 44(2):52–57, 2010.
- [29] Yuval Ishai, et al. Secure multi-party computation with identifiable abort. In *CRYPTO*, pages 369–386. Springer, 2014.
- [30] Ferdinand Brasser, Urs Müller, Alexandra Dmitrienko, Kari Kostainen, Srdjan Capkun, and Ahmad-Reza Sadeghi. Software grand exposure: SGX cache attacks are practical. In *WOOT*, 2017.
- [31] Kogler Andreas et al. Minefield: A software-only protection for SGX enclaves against DVFS attacks. In *USENIX Security*, 2022.
- [32] Kristoffer Myrseth Severinsen. Secure programming with intel SGX and novel applications. Master’s thesis, 2017.
- [33] Fred Cohen. The use of deception techniques: Honey pots and decoys. *Handbook of Information Security*, 3(1):646–655, 2006.
- [34] Riccardo Paccagnella, et al. Custos: Practical tamper-evident auditing of operating systems using trusted execution. In *NDSS*, 2020.
- [35] Dan Boneh, et al. Identity-based encryption from the weil pairing. In *CRYPTO*, pages 213–229. Springer, 2001.
- [36] Benoît Libert, et al. Towards black-box accountable authority IBE with short ciphertexts and private keys. In *PKCW*, pages 235–255. Springer, 2009.
- [37] Ben Laurie. Certificate transparency. *Communications of the ACM*, 57(10):40–46, 2014.
- [38] Mihir Bellare, et al. The security of triple encryption and a framework for code-based game-playing proofs. In *EUROCRYPT*, pages 409–426. Springer, 2006.
- [39] Victor Shoup. Sequences of games: A tool for taming complexity in security proofs. 2006.
- [40] Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In *Annual International Cryptology Conference (CRYPTO)*, pages 129–140. Springer, 1991.
- [41] The pairing-based cryptography library. <https://crypto.stanford.edu/abc/>.
- [42] MerkleC++ library. <https://github.com/merklecpp/merklecpp>.
- [43] Openssl library. <https://www.openssl.org/>.
- [44] Intel SGX SSL. <https://github.com/intel/intel-sgx-ssl>.
- [45] Luo Yili, et al. Accountable data sharing scheme based on blockchain and SGX. In *CyberC*, pages 9–16. IEEE, 2019.
- [46] Health insurance portability and accountability act. <https://hipaa.yale.edu/security/break-glass-procedure-granting-emergency-access-critical-epi-systems>, 2004.
- [47] José Parra Moyano, et al. KYC optimization using distributed ledger technology. *Business & Information Systems Engineering*, 59:411–423, 2017.
- [48] Joshua A Kroll, et al. Secure protocols for accountable warrant execution. 2014.
- [49] Matthew Green, et al. Abuse resistant law enforcement access systems. In *EUROCRYPT*, pages 553–583, 2021.
- [50] Kotla Ramakrishna et al. Pasture: Secure offline data access using commodity trusted hardware. In *OSDI*, 2012.
- [51] Xueping Liang, et al. Towards decentralized accountability and self-sovereignty in healthcare systems. In *ICICS*, pages 387–398. Springer, 2017.