IoT Notary: Sensor Data Attestation in Smart Environment

Nisha Panwar, Shantanu Sharma, Guoxi Wang, Sharad Mehrotra, Nalini Venkatasubramanian,
Mamadou H. Diallo, and Ardalan Amiri Sani
University of California, Irvine, California, USA.

***Abstract*—Contemporary IoT environments, such as smart buildings, require end-users to trust data-capturing rules published by the systems. There are several reasons why such a trust is misplaced — IoT systems may violate the rules either deliberately or since they may fall victim to cyberattacks that hijack IoT devices transferring user data to a malicious third-party leading to the loss of individuals' privacy or service integrity. To address such concerns, we propose IoT Notary, a framework to ensure trust in IoT systems and applications. IoT Notary provides secure log sealing on live sensor data to produce a verifiable 'proof-of-integrity' of the data captured based on which a verifier can attest that sensor data capture adheres to the published data-capturing rules. IoT Notary is an integral part of TIPPERS, a smart space system that has been deployed at UCI to provide various real-time location-based services in the campus. IoT Notary imposes nominal overheads for verification, such that a user can verify his/her data of one day only in less than two seconds.**

I. Introduction

IoT devices (*e.g.*, camera, WiFi access-points, cell phones, bodyworn sensors, occupancy sensors, temperature sensors, and light sensors) capture user-related data for empowering existing systems with new capabilities. While fine-grained continuous monitoring offers numerous benefits, it also raises several privacy and security concerns (*e.g.*, smoking habits, gender, religious belief) [1]. To highlight the privacy concern, we first share our experience in building location-based services at UC Irvine using WiFi connectivity data.

Use-case: University WiFi data collection. In our on-going project, entitled TIPPERS [2], we have developed a variety of location-based services based on WiFi connectivity dataset. At UC Irvine, more than 2000 WiFi access-points and four WLAN controllers (managed by university IT department) provide campus-wide wireless network coverage. Whenever a device connects to the campus WiFi network (through an access-point), the access-point generates Simple Network Management Protocol (SNMP) trap for this association event. Each association event consists of access-point-id, say s_i, user device MAC address, say d_j, and the time of the association, say t_k. All SNMP traps having $\langle s_i, d_j, t_k\rangle$ are sent to access-point's controllers in realtime. Additionally, the access-point controller anonymizes device MAC addresses (to preserve the privacy of WiFi users in the campus).

TIPPERS collects WiFi connectivity data from one of the controllers that manage 490 access-points and receives $\langle s_i, d_j, t_k\rangle$ tuples for each connectivity event. However, before receiving any WiFi data, TIPPERS notifies all WiFi users about the data-capture rules by sending emails over a university mailing list and/or by pushing messages to user devices. Subsequently, based on WiFi connectivity data $\langle s_i, d_j, t_k\rangle$, TIPPERS provides various realtime applications. Some of these services, *e.g.*, computing occupancy levels of (regions in) buildings in the form of a live heatmap, require only anonymous data. Other services, *e.g.*, providing location information (within buildings) or contextualized messaging (to provide messages to a user when he/she is in the vicinity of the desired location), require user's original disambiguated data. To date, over one hundred users have registered into TIPPERS to utilize realtime services. A key requirement imposed by the university in sharing data with TIPPERS is that the system supports provable mechanisms to verify that individuals have been notified prior to their data (anonymized or not) being used for service provisioning. Also, an option for users to opt-out of sharing their WiFi connectivity data with TIPPERS must be supported. If users opt-out, the system must prove to the users that indeed their data was not shared with TIPPERS. TIPPERS use immutable log-sealing to help all users to verify that the captured data is consistent with pre-notified data-capture rules.

Our experience in working with various groups in the campus, *e.g.*, the committee of privacy and security, labor unions, is that (persistent) location data can be deemed quite sensitive by certain individuals with concerns about the spied upon by the administration or by others and being in ways detrimental to the individuals. As a result, mechanism for notification of data-capture rules, secure log-sealing, and verification components made a sub-framework, entitled IoT Notary, which has become an integral part of TIPPERS.

Data-capture concerns in IoT environments are similar to that in mobile computing, where mobile applications may have continuous access to resident sensors on mobile devices. In the mobile setting, data-capture rules and permissions [3] are used to control data access, *e.g.*, which apps have access to which data generated at the mobile device (e.g., PII, credit cards, location, contact list, accelerometer data, gyroscope data, camera, etc.), for what purpose and in what context. However, in IoT settings, the data-capture framework differs from that in the mobile settings, in two important ways:

(i) Unlike the mobile setting, where applications can seek user's permission at the time of installation, or when they request sensor data on the device, in IoT settings, there are no obvious mechanisms/interfaces to seek users' preferences about the data being captured by sensors embedded in the smart environment. Recent work [4] has begun to explore mechanisms using which environments can broadcast their data-capture rules to users and seek their explicit permissions.

(ii) Unlike the mobile setting, sensors/data-capture devices in IoT settings are outside the control of the user. While in mobile settings, a user can trust the device operating system not to violate the data-capture rules, in IoT settings, trust (in the environment controlling the sensors) may be misplaced. IoT systems may not be honest or may inadvertently capture sensor data, even if data-capture rules are not satisfied [5].

Our problem focuses on the above-mentioned second scenario and determines ways to provide trustworthy sensing in an untrusted IoT environment. Thus, the users can verify their data captured by IoT environment based upon previously notified data-capture rules. Particularly, the problem deals with three sub-problems, namely *secure notification* to the user about data-capture rules, *secure (sensor data) log-sealing* to retain *immutable* sensor data, as well as, data-capture rules, and *remote attestation* to allow a user to verify the sensor data against pre-notified data-capture rules, while being minimally involved in the attestation process.

Our contribution and outline of the paper. We provide:

• A user-centric framework to ensure trustworthy data collection in untrusted IoT spaces, entitled IoT NOTARY (§IV-B) that contains three entities (§III-A): (*i*) *infrastructure deployer* that installs p sensors ($s_1, s_2, \ldots, s_p$), (*ii*) a *service provider* (SP, *e.g.*, TIPPERS) that establishes a list of data-capture rules ($dc_1, dc_2, \ldots, dc_q$) that dictates the condition under which the sensor s_i can/cannot capture data and utilizes the sensor data to provide realtime services to the user, and (*iii*) *users* that use services provided by sensors, as well as, by SP, (if interested).

• Two models to inform the user about the data-capture rules (§V-A). In the first model, entitled *notice-only model* (NoM), SP informs the user about its data-capture rules without requiring an acknowledgment from the user. In the second model, entitled *notice-and-ACK model* (NaM), SP informs the user about the data-capture rules, but cannot capture data unless the user provides an explicit consent.

• A secure log-sealing mechanism (§V-B) implemented by secure hardware at SP that cryptographically retains logs, data-capture rules, sensors' state, and contextual information to generate a *proof-of-integrity* in an immutable fashion. It also prohibits a malicious SP to capture any data other than what qualifies notified rules to the user.

• A secure attestation mechanism (§V-C), mixed with SIGMA protocol [6], allowing a verifier (a user or a *non-mandatory* auditor) to securely attest the sealed logs as per the data-capture rules. Implementation results of IoT NOTARY on the university live WiFi data are provided in §VI.

II. Comparison with Existing Work

We classify the related work in the area of IoT attestation into the following three categories:

Attestation in the context of IoT. The existing remote attestation protocols verify the internal memory state of untrusted devices through a trusted remote verifier. For example, AID [7] attests the internal state of neighboring devices through key exchange and proofs-of-non-absence. In AID, the adversary can compromise communication channels, the internal state of the device, and the cryptographic keys. SEDA [8] attests low-end embedded devices in a swarm and provides the number of devices that pass attestation. However, SEDA attests neighboring peer devices only. Similarly, DARPA [9] and SANA [10] allow detection of physical attacks by using heartbeat messages and provide aggregate network attestation, with high computation and communication costs, which are quadratic in the network size. SMARM [11] protects against malware by scanning memory in a secret randomized order. However, it may require many iterations to eventually detect malware. In short, all such work only deals with attestation of sensor devices, and their methods cannot be used to verify sensor data against the data-capture rules.

In contrast, IoT NOTARY does not deal with the verification of internal state of sensor devices, since in our case, (WiFi access-point) sensors are assumed to deploy by a trusted entity (*e.g.*, university IT department). Of course, cyberattacks are possible on sensors to maliciously record data and that can also be detected by IoT NOTARY, while not verifying the sensors.

Attestation using secure hardware. [12] provided SGX-based attestation method for physical attacks on the sensor, *e.g.*, modifying memory and changing I/O signals. Fiware [13] provides secure key management through key vault running in SGX, thereby provides an alternate to PKI-based solutions. However, [12], [13] are unable to verify any sensor data. In [12], [13], if data-capture rules are mis-notified to the user, SGX cannot detect any inconsistency.

In contrast, IoT NOTARY does not deal with attacks on sensors, as well as, a specific key management protocol. However, IoT NOTARY can detect and discard the sensor data that does not comply with the notifications released earlier.

Integrity verification. [14], [15] proposed a privacy-preserving scheme based on zero-knowledge proofs to detect log-exclusion attacks. [14] provided solutions for accountability and auditing through hierarchical multi-party computation (MPC) and succinct zero-knowledge proof statements. [15] provided a privacy-preserving certificate transparency service, which signs a message four times, where an auditor can trace the certified logs. [16] proposed a Merkle tree-based history tree to prove the sequence of logs over time. [17] proposed a Bloom tree that stores proof of logs at an untrusted cloud. Further, access-pattern-hiding cryptographic techniques [18], [19] may be used to verify any stored log, since an adversary cannot skip the computation on some parts of the data, due to executing an identical computation on each sensor reading. Techniques, *e.g.*, function secret-sharing [18] or vSQL [19], may be used to verifying query results on cleartext. However, these techniques cannot detect log deletion. Also, all such techniques incur signification time. For example, vSQL took more than 4000 seconds to verify a SQL query. In addition, any end-to-end encryption model is not sufficient for verification.

In contrast, IoT NOTARY provides a complete security to sensor data and realtime data attestation approach. Unlike [15], IoT NOTARY requires only two signatures per file, where one is used to validate log completeness, and another is used to validate the log ordering with respect to adjacent logs.

III. Modeling IoT Data Attestation

A. Entities

Our model has the following entities, see Figure 1:

Infrastructure Deployer (IFD). IFD (which is the university IT department in our use-case; see §I) deploys and owns a network of p sensors devices (denoted by $s_1, s_2, \ldots, s_p$),

which capture information related to users in a space. The sensor devices could be: (*i*) dedicated sensing devices, *e.g.*, energy meters and occupancy detectors, or (*ii*) facility providing sensing devices, *e.g.*, WiFi access-points and RFID readers. Our focus is on facility providing sensing devices, especially WiFi access-points that also capture some user-related information in response to services. For example, WiFi access-points capture the associated user-device-ids (MAC addresses), time of association, some other parameters (such as signal strength, signal-to-noise ratio), and denoted by: $\langle d_i, s_j, t_k, param \rangle$, where d_i is i^{th} user-device-id, s_j is the j^{th} sensor device, t_k is k^{th} time, and $param$ is other parameters (this paper does not deal with $param$ field and focuses on only the first three fields). All sensor data is collected at a controller (server) owned by IFD. The controller may keep the sensor data in cleartext or in encrypted format; however, it only sends encrypted sensor data to the service provider.

Service Providers (SP). SP (which is TIPPERS in our use-case; see §I) utilizes the sensor data of a given space to provide different *services*, *e.g.*, monitoring a location and tracking a person. SP receives encrypted sensor data from the controller. Data-capture rules. SP establishes data-capture rules (denoted by a list DC having different rules $dc_1, dc_2, \ldots, dc_q$). Data-capture rules are conditions on device-ids, time, and space. Each data-capture rule has an associated *validity* that indicates the time during which a rule is valid. Data-capture rules could be to capture user data by default (unless the user has explicitly opted out). Alternatively, default rules may be to opt-out, unless, users opt-in explicitly. Consider a default rule that individuals on the 6^{th} floor of the building will be monitored from 9pm to 9am. Such a rule has an associated condition on the time and the id of the sensor used to generate the data. Now, consider a rule corresponding to a user with a device d_i opting-out of data capture based on the previously mentioned rule. Such an opt-out rule would have a condition on the user-id in addition to conditions on time and sensor-id. For sensor data for which a default data-capture rule is opt-in, the captured data is forwarded to SP, if there does not exist any associated opt-out rules, whose associated conditions are satisfied by the sensor data. Likewise, for sensor data where the default is opt-out, the data is forwarded to SP only, if there exists an explicit opt-in condition. We refer to the sensor data to have a *sensor state* ($s_i.state$ denotes the state of the sensor s_i) of 1 (or active), if the data can be forwarded to SP; otherwise, the sensor data is said to be in a sensor state 0 (or passive). In the rest of the paper, unless explicitly noted, we will assume that the default rule is opt-out for simplicity of discussion.

Whenever SP creates a new data-capture rule, SP must send a *notice message* to user devices about the current usage of sensor data (this phase is entitled *notification phase*). SP uses SGX [20],[1] which works as a trusted agent of IFD, for securely storing sensor data corresponding to data-capture rules. SGX keeps all valid data-capture rules in the secure memory and only allows to keep such data that qualifies pre-notified valid data-capture rules; otherwise, it discards other sensor data. Further, SGX creates immutable and verifiable logs of the sensor data (this phase is entitled *log-sealing phase*).

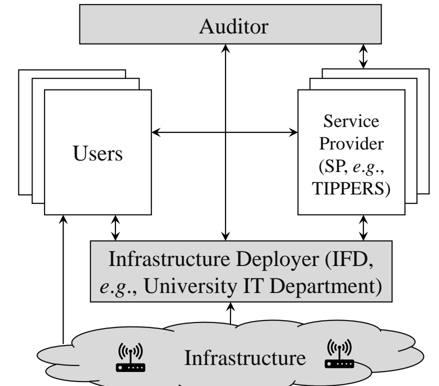

Figure 1: Entities in IoT Notary.

Note that the assumption of secure hardware at a machine (*e.g.*, SP that utilizes the data produced by sensors) is rational with the emerging system architectures; for instance, Intel machines are equipped with SGX [21]. However, existing SGX architecture suffers from side-channel attacks, *e.g.*, cache-line, branch shadow, page-fault attacks [22]–[24]. Due to these attacks, the adversary may learn some sensitive information inside SGX. However, we consider such attacks are outside the scope of this paper, and solutions (*e.g.*, T-SGX [25] or Sanctum [26]) can overcome these attacks in the future.

Users. Let $d_1, d_2, \ldots, d_m$ be m (user) devices carried by $u_1, u_2, \ldots, u_{m'}$ users, where $m' \leq m$. Using these devices, users enjoy services provided by SP. We define a term, entitled *user-associated data*. Let $\langle d_i, s_j, t_k \rangle$ be a sensor reading, and let d_i be the i^{th} device-id owned by a user u_i. We refer to $\langle d_i, s_j, t_k \rangle$ as user-associated data with the user u_i. Users worry about their privacy, since SP may capture user data without informing them, or in violation of their preference (*e.g.*, when the default was an opt-out rule or when a user opted-out from an opt-in default). Users may also require SP to prove service integrity by storing all the sensor data associated with the user (of course, in such a case user must have opted-in into the service), while minimally being involved in the attestation process and storing records at their sides (this phase is entitled *attestation phase*).

Auditor. An auditor is a *non-mandatory* trusted-third-party that can (periodically) verify entire sensor data against data-capture rules. Note that a user can only verify his/her data, not the entire sensor data or sensor data related to other users, since it may reveal the privacy of other users.

B. Threat Model

We assume that SP and users may behave like adversaries. The adversarial SP may *store* sensor data without informing data-capture rules to the user. The adversarial SP may *tamper* with the sensor data by inserting, deleting, modifying, and truncating sensor readings and secured-logs in the database. By tampering with the sensor data, SP may *simulate* the sealing function over the sensor data to produce secured-logs that are identical to real secured-logs. Thus, the adversary may hinder

[1] Intel Secure Guard eXtension (SGX) [20] provides a private process execution environment, called *enclave* at an untrusted system. An SGX-enabled processor partitions the trusted and untrusted environment, thereby a process executing outside of the enclave cannot observe the enclave's application/data.

the attestation process and make it impossible to detect any tampering with the sensor data by the verifier (that may be an auditor or a user). Further, as mentioned before that SP utilizes sensor data to provide services to the user. However, an adversarial SP may provide *false answers* in response to user queries. We assume that the adversarial SP cannot obtain the secret key of the enclave (by any means of side-channel attacks on SGX). Since we assumed that sensors are trusted and cannot be spoofed, we do not need to consider a case when sensors would collude with SP to fabricate the logs.

An adversarial user may *repudiate* the reception of notice messages about data-capture rules. Further, an adversarial user may *impersonate* a real user, and then, may retrieve the sensor data and secured-log during the verification phase. This way an adversarial user may reveal the privacy of the users by observing sensor data. Further, a user may infer the identity of other users associated with sensor data by potentially launching *frequency-count attacks* (*e.g.*, by determining which device-ids are prominent).

C. Security Properties

In the above-mentioned adversarial model, an adversary wishes to learn the (entire/partial) data about the user, without notifying or by mis-notifying about data-capture rules, such that the user/auditor cannot detect any inconsistency between data-capture rules and stored sensor data at SP. Hence, a secure attestation algorithm must make it detectable, if the adversary stores sensor data in violation of the data-capture rules notified to the user. To achieve a secure attestation algorithm, we need to satisfy the following properties:

Authentication. Authentication is required: (*i*) between SP and users, during notification phase; thus, the user can detect a rouge SP, as well as, SP can detect rouge users, and (*ii*) between SP and the verifier (auditor/user), before sending sensor data to the verifier to prevent any rouge verifier to obtain sensor data. Thus, authentication prevents threats such as impersonation and repudiation. Further, a periodic mutual authentication is required between IFD and SP, thereby discarding rouge sensor data by SP, as well as, preventing any rouge SP to obtain real sensor data.

Immutability and non-identical outputs. We need to maintain immutability of notice messages, sensor data, and the sealing function. Note that if the adversary can alter notice messages after transmission, it can do anything with the sensor data, in which case, sensor data may be completely stored or deleted without respecting notice messages. Further, if the adversary can alter the sealing function, the adversary can generate a proof-of-integrity, as desired, which makes the flawless attestation impossible. The output of the sealing function should not be identical for each sensor reading to prevent an adversary to forge the sealing function (and to prevent the execution of frequency-count attack by the user). Thus, immutability and non-identical outputs properties prevent threats, *e.g.*, inserting, deleting, modifying, and truncating the sensor data, as well as, simulating the sealing function.

Minimality, non-refutability and privacy-preserving verification. The verification method must find any misbehavior of SP, during storing sensor data inconsistent with pre-notified data-capture rules. However, if the verifiers wish to verify a subset of the sensor data, then they should not verify the entire sensor data. Thus, SP should send a minimal amount of sensor data to the verifier, enabling them to attest what they wish to attest. Further, the verification method: (*i*) cannot be refuted by SP, and (*ii*) should not reveal any additional information to the user about all the other users during verification process. These properties prevent SP to store only sensor data that is consistent with the data-capture rules notified to the user. Further, these properties preserve the privacy of other users during attestation and impose minimal work on the verifier.

D. The Attestation Protocol

This section presents the formal definition of the attestation protocol. Before that we discuss some assumptions, we made, as follows:

Assumptions. IoT NOTARY has the following assumptions: (*i*) The communication channels between SP and users, as well as, between SP and auditor are insecure. Thus, our solution incorporates an authenticated key exchange based on SIGMA protocol (which protects sender identity). When the verifier's identity is proved, the cryptographically sealed logs are sent to the verifier. (*ii*) Sensor devices are tamper-proof, and they cannot be replicated/spoofed (*i.e.*, two devices cannot have an identical id). In short, we assume a correct identification of sensors, before accepting any sensor-generated data, and it ensures that no rogue sensor device can generate the data on behalf of an authentic sensor. (*iii*) The sensor devices are assumed to be computationally-inefficient to locally generate a verifiable log for the continuous data stream as per the data-capture rules. (*iv*) By any side-channel attacks on SGX, one cannot tamper with SGX and retrieve the secret-key of SGX. (Otherwise, the adversary can simulate the sealing process.)

Definition: Attestation Protocol. Now, we define an attestation protocol that consists of the following components:

• $\mathit{Setup}()$: Given a security parameter 1^k, $\mathit{Setup}()$ produces a public key of the enclave (PK_E) and a corresponding private key (PR_E), used by the enclave to securely write sensor logs.

• $\mathit{Sealing}(PR_E, \langle d_i, s_j, s_j.state, t_k\rangle, dc_l)$: Given the key PR_E, $\mathit{Sealing}()$ (which executes inside the enclave) produces a verifiable *proof-of-(log)-integrity* ($\mathcal{PI}$) and *proof-of-integrity for user/service (query) verification* ($\mathcal{PU}$), based on the received sensor readings and the current data-capture-rule, dc_l.

• $\mathit{Verify}(PK_E, \langle *, s_j, s_j.state, t_k, \mathcal{PI}, dc_l\rangle, \mathit{Sealing}(PR_E, \langle d_i, s_j, s_j.state, t_k\rangle, dc_l))$: Given the public key PK_E, sensor data, proof, and data-capture rule, $\mathit{Verify}()$ is executed at the verifier, where $*$ denotes the presence/absence of a user-device-id, based on dc_l. $\mathit{Verify}()$ produces 1, iff $\mathcal{PI} = \mathit{Verify}(PK_E, \langle *, s_j, s_j.state, t_k, dc_l\rangle, \mathit{Sealing}(PR_E, \langle d_i, s_j, s_j.state, t_k\rangle, dc_l))$; otherwise, 0. Similarly, $\mathit{Verify}()$ can attest $\mathcal{PU}$.

Note that the functions $\mathit{Sealing}()$ and $\mathit{Verify}()$ are known to the user, auditor, and SP. However, the secret-key PR_E is only known to the enclave.

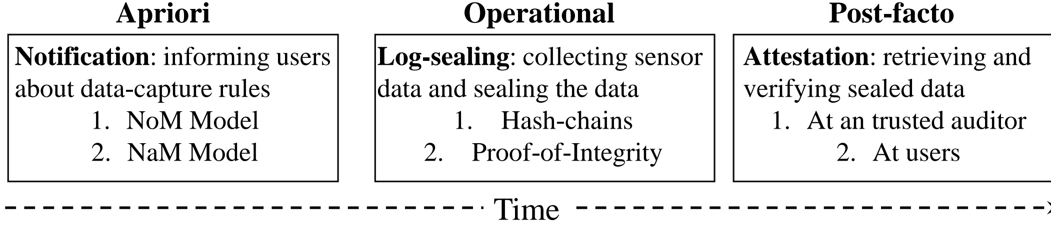

Figure 2: Phases in IoT Notary.

IV. IoT Notary: Challenges and Approach

This section provides challenges we faced during IoT Notary development and an overview of IoT Notary.

A. Challenges and Solutions

We classify the problem of data attestation in IoT into three categories: (*i*) secure notification, (*ii*) log-sealing, and (*iii*) log verification. Before going into details, we provide the challenges that we faced during the development and the way we overcome those challenges (Figure 2), as follows:

C1. Secured notification for data-capture rules. The declaration of data-capture rules requires a reliable and secure notification scheme, thereby users can verify the sender of notice messages. Trivially, this can be done through a unique key establishment between each user and SP. However, this incurs a major overhead at SP, as well as, SP can send different messages to different users.

Solution. To address the above challenge, we develop three solutions: One is based on secure notifier that delivers a cryptographically encrypted notice message, which is signed by the enclave, to all the users (see §V-A). The second solution uses time-based one-time passwords [27], [28] that remain valid for only a specific duration. The users can verify these temporal passwords to authenticate the sender, and hence, the validity of notifications sent (see Appendix A). The third solution uses an acknowledgment from the user and does not need any trusted notifier (see §V-A).

C2. Tamper-proof cryptographically secured-log sealing. The verification process depends on immutable sensor data that is stored at SP. A trivial way is to store the entire data using a strong encryption technique, mixed with access-pattern hiding mechanisms. While it will prevent tampering with the data (except deletion), SP cannot develop realtime services on this data. Thus, the first challenge is how to store sensor data at SP according to data-capture rules; hence, the verifier can attest the data. The second challenge arises due to the computational cost at the verifier and communication cost between the verifier and SP to send the verifiable sensor data; *e.g.*, if the verifier wishes to attest only one-day old sensor data over the sensor data of many years, then sending many years of sensor data to the verifier is impractical. Finally, the last challenge is how to securely store data when data-capture rules are set to be false (*i.e.*, not to store data). In this case, not storing any data would not provide a way for verification, since SP may store data and can inform the verifier that there is no data as per the existing data-capture rules.

Solution. To address the first challenge, we develop a cryptographic sealing method based on hash-chains and XOR-linked-lists to ensure immutable sensor logs, after the sealed logs leave the enclave. Thus, any log addition/deletion/update is detectable (see §V-B1). To address the second challenge, we execute sealing on small-sized chunks, which each maintains its hash-chain and XOR-links. The XOR-links ensure the log completeness, *i.e.*, a chunk before and after the desired chunk has not been altered (see §V-B1). To address the third challenge, we store the *device state* of the first sensor-reading for which the data-capture rule is set to false. Further, we discard all subsequent sensor-readings, unless finding a sensor-reading for which data-capture rule is to *store data* (§V-B2).

C3. Privacy-preserving log verification. In case of log-integrity verification, SP can provide the entire sensor data with cryptographically sealed log to the trusted auditor. But, the challenge arises, if a user asks to verify her user-associated data/query results. Here, SP cannot provide the entire sensor data to the user, since it will reveal other users' privacy.

Solution. To address this challenge, we develop a technique to non-deterministically encrypt the user-id before cryptographically-sealing the sensor data. However, only non-deterministic encryption is also not enough to verify the log completeness, we compute XOR operation on each sensor reading, and then, create XOR-linked-list (see §V-B3).

B. IoT Notary*: An Overview*

This section presents an overview of the three phases and dataflow among different entities and devices, see Figure 3.

Notification phase: SP to Users messages. This is the first phase that notifies users about data-capture rules for the IoT space using notice messages (in a verifiable manner for later stages). Such messages can be of two types: (*i*) notice messages, and (*ii*) notice-and-acknowledgment messages. SP establishes (the default) data-capture rules and informs trusted hardware (❶). Trusted hardware securely stores data-capture rules (❷, ❺) and informs the *trusted notifier* (❸) that transmits the message to all users (❹). Only notice messages need a trusted notifier to transmit the message (see §V-A).

Log-sealing phase: Sensor devices to SP messages. Each sensor sends data to the controller that sends encrypted data to SP (❻). As we assumed the sensors cannot be replicated or altered, the encrypted data flow between a sensor and the controller is secured. At SP, trusted hardware (Intel SGX) reads the encrypted data in the enclave (❼).

Working of the enclave. The enclave decrypts the data and checks against the pre-notified data-capture rules. Recall that the decrypted data is of the format: $\langle d_i, s_j, t_k \rangle$, where d_i is i^{th} user-device-id, s_j is the j^{th} sensor device, and t_k is k^{th} time. After checking each sensor reading, the enclave

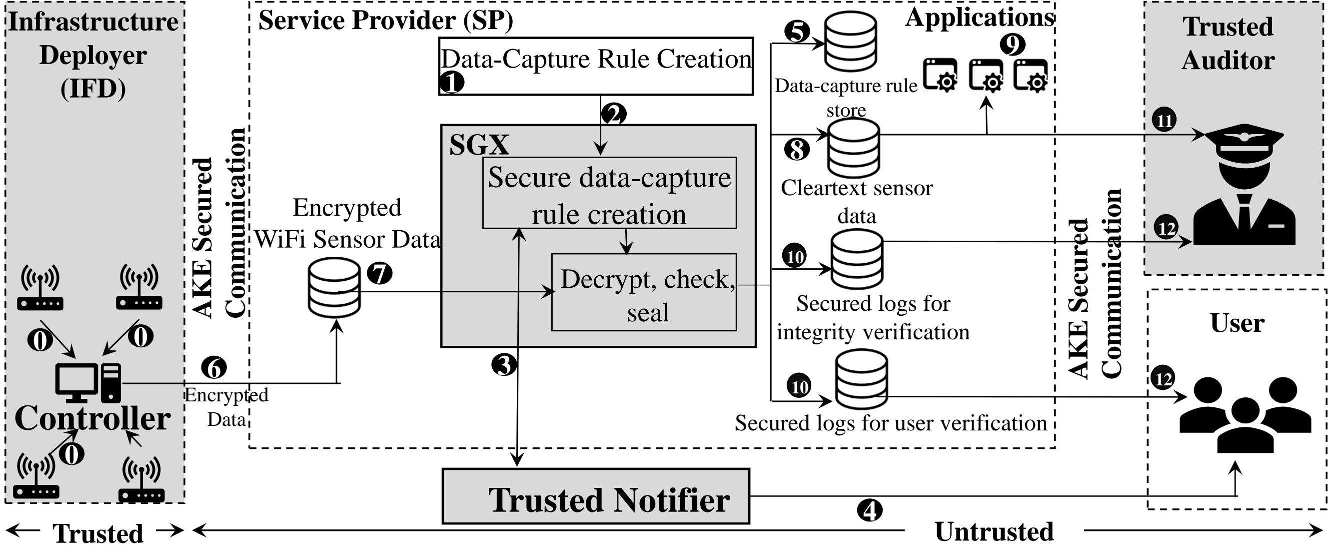

Figure 3: Dataflow and computation in the protocol. Trusted parts are shown in shaded box.

adds a new field, entitled *sensor (device) states*. The sensor state of a senor s_j is denoted by $s_j.state$, which can be `active` or `passive`, based on capturing user data. For example, $s_j.state$ = `active` or (1), if data captured by the sensor s_j satisfies the data-capture rules; otherwise, $s_j.state$ = `passive` or (0). For all the sensors whose $state = 0$, the enclave deletes the data. Then, the enclave cryptographically seals sensor data, regardless of the sensor state, and provides cleartext sensor data of the format: $\langle d_i, s_j, s_j.state = 1, t_k \rangle$ to SP (❽) that provides services using this data (❾). Note that the cryptographically sealed logs and cleartext sensor data are kept at untrusted storage of SP (❽, ❿).

Verification phase: SP to verifier messages. In our model, an auditor and a user can verify the sensor data. The auditor can verify the entire/partial sensor data against data-capture rules by asking SP to provide cleartext sensor data and cryptographically sealed logs (❽, ❿). The users can also verify their own data against pre-notified messages or can verify the results of the services provided by SP using only cryptographically sealed logs ⓬. Note that using an underlying authentication technique (as per our assumptions, given in §III-A), auditor/user and SP authenticate each other before transmitting data from SP to auditor/user.

V. Attestation Protocol

This section presents three phases of attestation protocol.

Preliminary Setup Phase. We assume a preliminary setup phase that distributes public keys (PK) and private keys (PR), as well as, registers devices into the system. We assume that the enclave has its own secret-key, SK_E, fused into the hardware. $\langle PK_{SP}, PR_{SP} \rangle$ denotes SP's keys, $\langle PK_E, PR_E \rangle$ denotes the enclave's keys, $\langle PK_N, PR_N \rangle$ denotes the trusted notifier's keys, $\langle PK_{di}, PR_{di} \rangle$ denotes keys of the i^{th} user device. Further, SGX and the controller share a secret-key. These keys can be revoked and renewed by the trusted authority only. Note that our key management is identical to the standard settings proposed for SGX (*e.g.*, [29]).

We assume a registration process during which a user identifies himself/herself to the underlying system. For instance, in a WiFi network, users are identified by their mobile devices, and the registration process consists of users providing the MAC addresses of their devices (and other personally identifiable information, *e.g.*, email and a public key). During registration, users also specify their preferred modality through which the system can communicate with the user (*e.g.*, email and/or pushing messages to the user device). Such communication is used during notification phase.

A. Notification Phase

The notification phase informs data-capture rules established by SP to the (registered) users by explicitly sending *notice messages*. We consider two models for notification, differing based on acknowledgment from users.

In the *notice-only model (NoM)*, SP informs users of data-capture rules, but users may not acknowledge receipt of the message. Such a model is used to implement policies, when data capture is mandatory, and the user cannot exercise control, over data capture. Since there is no acknowledgment, SP is only required to ensure that it sends a notice, but is not required to guarantee that the user received the notice. In contrast, a *notice-and-ACK model (NaM)* is intended for discretionary data-capture rules that require explicit permission from users prior to data capture. Such rules may be associated, for instance, with fine-grained location services that require users' location. A user can choose not to let SP track his location, but will likely not be able to avail some services.

Implementation of notification differs based on the model used. Interestingly, since NaM requires acknowledgment, the notification phase is easier as compared to NoM that uses a trusted notifier to deliver the message to users. Below we discuss the implementation of both models:

Notification implementation in NoM. NoM assumes that, by default, data-capture rules are set not to retain any user data, unless SP, first, informs SGX about a data-capture rule, (*i.e.*, SP cannot use the encrypted sensor data for building any

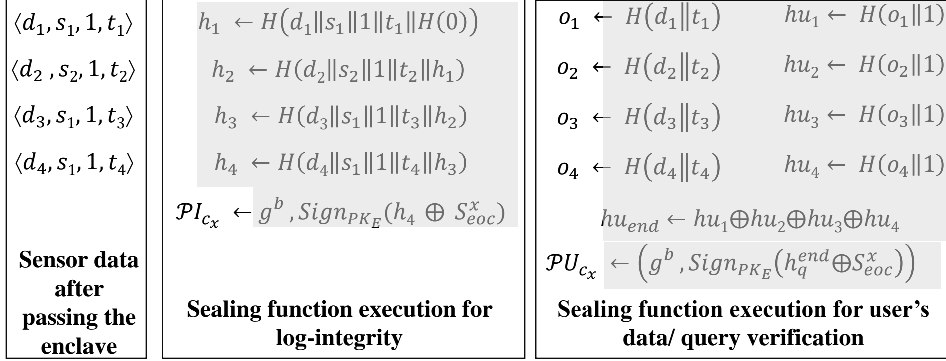

Figure 4: Cryptographically sealing procedure executed on a chunk, $\mathcal{C}_x$. Gray-shaded data is not stored on the disk. White-shaded data is stored on the disk and accessible by SP. Figure shows proof-of-integrity only for one chunk, $\mathcal{C}_x$; hence, some notations are abusively used.

application, see ❾ in Figure 3). When SP creates a new data-capture rule, SP must inform SGX. Then, the enclave encrypts the data-capture rule using the public key (*i.e.*, PK_N) of the notifier and informs the trusted notifier (via SP) about the encrypted data-capture rule by writing it outside of the enclave (in our user-case §I, the university IT department works as a trusted notifier). Data-capture rules are maintained by SP on stable storage, which is read by SGX into the enclave to check, if the sensor data should be forwarded to SP. SGX can retain a cache of rules in the enclave, if such rules are still valid (and hence used for enforcement).[2] Finally, the trusted notifier acknowledges SP about receiving the encrypted data-capture rule, and then, informs users of the encrypted data-capture rule via signed notice messages. On receiving the notice message, the users may decrypt it and obtain the data-capture rule.

To see the role of trusted hardware above, suppose that SP was responsible for informing users about data-capture rules directly. Data-capture rules are also required by SGX during log-sealing (PHASE 2). An adversarial SP may inform SGX, not to users, or may inform non-identical rules to users and to SGX. Hence, SP cannot inform the rule to users directly.

To see the role of trusted notifier above, suppose that SP can directly inform users about encrypted data-capture rules obtained from SGX. An adversarial SP may not deliver the data-capture rule to all or some of the users; thus, an encrypted data-capture rule is not helpful. Hence, a trusted notifier ensures that the notice message is sent to all the registered user. Note that the trusted notifier might be a trusted web site that lists all the data capture rules which users can access.

Implementation of notification in NaM. Unlike NoM, the notification phase of NaM does not require the trusted notifier. In NaM, by default, SP cannot utilize all those sensor readings having device-ids for which the users have not acknowledged. Likewise NoM, in NaM, SP informs data-capture rules to SGX that encrypts the rule and writes outside of the enclave. The encrypted rules are delivered by SP to users, unlike NoM. On receiving the message, a user may securely acknowledge the enclave about her consent. The enclave retains all those device-ids that acknowledge the notice message for log-sealing phase and considers those device-ids during the log-sealing phase to retain their data while discarding data of others.

[2] Since the enclave has a limited memory, the enclave cannot retain all the valid and non-valid data-capture rules after a certain size. Thus, the enclave writes all the non-valid data-capture rules on the disk after computing a secured hash digest over all the rules. Taking a hash over the rules is required to maintain integrity of all the rules, since any rule written outside of the enclave can be tampered by SP. Recall that altering a rule will make it impossible to verify partial/entire sensor data.

B. *Log Sealing Phase*

The second phase consists of cryptographically sealing the sensor data for future verification against pre-notified data-capture rules. The sensor data is sealed into secured logs using authenticated data structures, *e.g.*, hash-chains and XOR-linked lists (as shown in Figures 4, 5), by the sealing function, $Sealing(PR_E, \langle d_i, s_j, s_j.state, t_k\rangle)$, executed in the enclave at SP. Let us explain log-sealing in the context of WiFi connectivity data. The enclave reads the encrypted sensor data (❼ in Figure 3) and executes the three steps: (*i*) decrypts the data, (*ii*) checks the data against pre-notified valid data-capture rules, and (*iii*) cryptographically seals the data and store *appropriate secured logs*.

Below we explain our log sealing approach. To simplify the discussion, we first consider the case when all the sensor data satisfies some data-capture rule (*i.e.*, the state of all the sensor data is one), and hence, data is forwarded to and stored at SP §V-B1. Then, we adapt the protocol to deal with the case when some sensor data satisfies some data-capture rule (*i.e.*, the state of some sensor data is one, and hence, data is forwarded to and stored at SP), while remaining sensor data does not satisfy any rule (*i.e.*, the state of the remaining sensor data is zero, and hence, data is forwarded to SP) V-B2.

1) ***Sealing Entire Sensor Data****:* Informally, the sealing function executes a hash function on each sensor reading (or value), whose output is used to create a chain of hash digests. At the end of the sensor readings/values, the sealing function generates an authenticated proof-of-integrity by mixing a computationally-hard secure string. For example, consider four values: v_1, v_2, v_3, and v_4. The sealing function works as follows:

Value	Hash output
v_1	$h_1 \leftarrow H(v_1\|\|H(0))$
v_2	$h_2 \leftarrow H(v_2\|\|h_1)$
v_3	$h_3 \leftarrow H(v_3\|\|h_2)$
v_4	$h_4 \leftarrow H(v_4\|\|h_3)$
Proof-of-integrity	$\langle SS, sign_{PK_E}(SS \oplus h_4)\rangle$

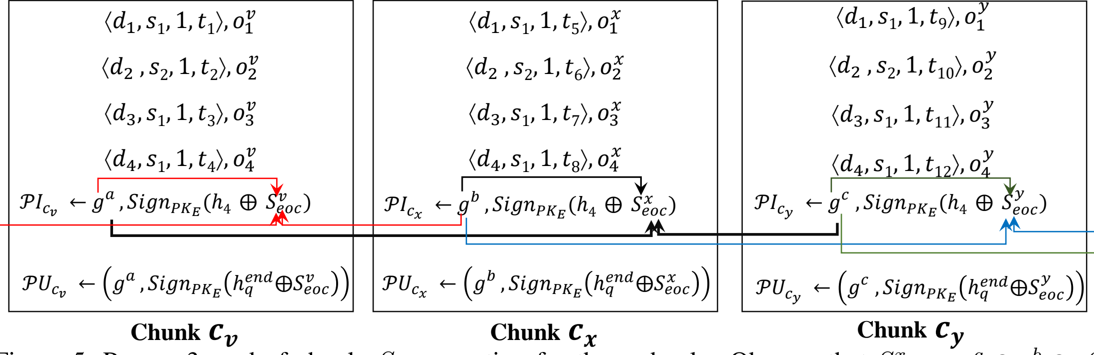

Figure 5: PHASE 3: end of chunk, S_{eoc}, creation for three chunks. Observe that $S^x_{eoc} = g^a \oplus g^b \oplus g^c$.

In this example, all the values are hashed while including the hash digest of the previous value, except the first value. Note that only the first value is hashed with the hash digest of zero. In the end, the proof-of-integrity is prepared by signing XORed-value of the hash digest of the last value and some secret-string, SS.

The sealing operation consists of the following three phases: (*i*) chunk creation, (*ii*) hash-chain creation, and (*iii*) proof-of-integrity creation; described below.

PHASE 1: Chunk creation. The first phase of the sealing operation finds an appropriate size of a chunk (to speed up the attestation process). Note that the incoming encrypted sensor data may be large, and it may create problems during verification, due to increased communication between SP and the verifier. Also, the verifier needs to verify the entire data, which have been collected over a large period of time (*e.g.*, months/years). Further, creating cryptographic sealing over the entire sensor data may also degrade the performance of $Sealing()$ function, due to the limited size of SGX enclave. Thus, we first determine an appropriate chunk size, for each of which the sealing function is executed.

The chunk size depends on time epochs, the enclave size, the computational overhead of executing sealing on the chunk, and the communication overhead for providing the chunk to the verifier. A small chunk size reduces the communication overhead and maintains the log minimality property, thereby during the log verification phase, a verifier retrieves only the desired log chunks, instead of retrieving the entire sensor data. Consequently, SP stores many chunks.

PHASE 2: Hash-chain creation. Consider a chunk, $\mathcal{C}_x$, that can store at most n sensor readings, each of them of the format: $\langle d_i, s_j, t_k \rangle$. The sealing function checks each sensor reading against data-capture rules and adds sensor state to each reading, as: $\langle d_i, s_j, s_j.state, t_k \rangle$. Since in this section we assumed that all sensor data will be stored, the sensor state of each sensor reading is set to 1. The sealing function starts with the first sensor reading of the chunk $\mathcal{C}_x$, as follows:

First sensor reading. For the first sensor reading of the chunk, the sealing function computes a hash function on value zero, *i.e*, $H(0)$. Then, the sealing function mixes $H(0)$ with the remaining values of the sensor reading, *i.e.*, sensor-id, device-id, sensor state, and time, at which it computes the hash function, denoted by $H(d_1||s_j||s_j.state||t_k||H(0))$ that results in a hash digest, denoted by h^x_1. After processing the complete first sensor reading of the chunk $\mathcal{C}_x$, the enclave writes cleartext first sensor reading of $\mathcal{C}_x$, *i.e.*, $\langle d_1, s_j, s_j.state, t_k \rangle$ on the disk, which can be accessed by SP.

Second sensor reading. Let $\langle d_2, s_j, s_j.state, t_{k+1} \rangle$ be the second sensor reading. For this, the sealing function works identically to the processing of the first sensor reading, *i.e.*, it computes a hash function on the second sensor values, while mixing it with the hash digest of the first sensor reading, *i.e.*, $H(d_2||s_j||s_j.state||t_{k+1}||h^x_1)$ that results in a hash digest, say h^x_2. Finally, the enclave writes the second sensor reading in cleartext on the disk.

Processing the remaining sensor readings. Likewise, the second sensor reading processing, the sealing function computes the hash function on all the remaining sensor readings of the chunk $\mathcal{C}_x$. After processing the last sensor reading of the chunk $\mathcal{C}_x$, the hash digest h^x_n is obtained.

PHASE 3: Proof-of-integrity creation. Since each sensor reading is written on disk, SP can alter any sensor reading, which makes it impossible to verify log integrity by an auditor. Thus, to show that all the sensor readings are kept according to the pre-notified data-capture rules without any alteration, the sealing function prepares an immutable proof-of-integrity for each chunk, as follows:

Let $\mathcal{C}_v$, $\mathcal{C}_x$, and $\mathcal{C}_y$ be three chunks (see Figure 5). For each chunk, the sealing function generates a secure random string, say g^a, g^b, g^c for $\mathcal{C}_v$, $\mathcal{C}_x$, and $\mathcal{C}_y$ chunks, respectively, by computing secret string exponents on g; where g is the generator of cyclic group G of prime order q, and a, b, and c are secret strings. Now, for producing the proof-of-integrity for the chunk $\mathcal{C}_x$, the sealing function: (*i*) executes XOR operation on g^a, g^b, g^c, whose output is denoted by S^x_{eoc}, where *eoc* denotes the end-of-chunk; (*ii*) signs h^x_n XORed with S^x_{eoc} with the private key of the enclave; and (*iii*) writes the proof-of-integrity for log verification of the chunk $\mathcal{C}_x$, as follows:

$$\mathcal{PI}_{\mathcal{C}_x} = (g^b, Sign_{PR_E}(h^x_n \oplus S^x_{eoc}))$$

Example. Please see Figure 4, where the middle box shows PHASE 2 execution on four sensor readings. Note that the hash digest of each reading is passed to the next sensor reading on which a hash function is computed with the sensor reading. After computing h_4, the proof-of-integrity, $\mathcal{PI}$, is created that included signed $h_4 \oplus S^x_{eoc}$ with a secure random string, g^b.

Note. g^* for the first chunk. The initialization of log sealing function requires an initial seed value, say g^*, due to the absence of 0^{th} chunk, which can produce its secret-string. Therefore, in order to initialize the secure binding for the first chunk, while there is no residual secret from the previous

chunk, the seed value is used as a substitute secret value.

2) ***Sealing Mixed State Sensor Data***: The protocol so far has assumed that all data has a sensor state of one. We next consider how it can be generalized to a situation when some sensor data may not satisfy the data capture rules (and hence, have a sensor state of zero).

Example. Assume following five sensor readings in a chunk. $\langle d_1, s_1, 1, t_1\rangle\langle d_2, s_2, 0, t_2\rangle\langle d_2, s_2, 0, t_3\rangle\langle d_3, s_2, 0, t_4\rangle\langle d_1, s_1, 1, t_5\rangle$ In this case, the enclave does not store the second, third, and fourth sensor readings on disk and delete them, while the first and fifth will be stored on the disk. But, the enclave will store only the sensor device and its state of the first reading for which state was zero and generate verifiable logs such that the verifier can verify that the three entries has been deleted. For example, the sealing function computes the hash-chain and proof-of-integrity as follows:

$$h_1 \leftarrow H(d_1||s_1||1||t_1||H(0))$$
$$h_2 \leftarrow H(s_2||0||t_2||h_1)$$
$$h_3 \leftarrow H(d_1||s_1||1||t_5||h_2)$$
$$\mathcal{PI} \leftarrow \langle secret\ string, sign_{PR_E}(secret\ string \oplus h_3)\rangle$$

Note that here we compute a hash function on the first, second and last sensor reading, while do not seal the third and fourth sensor readings, due to their passive (or zero) sensor states.[3]

Informally, the sealing function checks each sensor reading and deletes all those sensor readings for which sensor state is `passive`, either due to data-capture rules/context. In this case, it is not mandatory to seal each sensor readings, like §V-B1. Thus, the sealing function provides a *filter operation* that removes sensor readings whose device state is `passive`, while storing sensor readings whose device states are `active`. But, the sealing function cryptographically stores the data with minimal information of sensor readings whose device states are `passive`. Below we describe PHASE 2 for sealing mixed state sensor data. Note that PHASE 1 (chunk creation) and PHASE 3 (proof-of-integrity) creation is identical to the method described in §V-B1.

PHASE 2: Sealing operation. Formally, to create hash-chain for this case, the sealing function will do the following: For the first i^{th} sensor reading (for example, $\langle d_j, s_k, 0, t_i\rangle$) whose $state = 0$, the enclave executes two operations: (*i*) sealing function that computes a hash function, as: $H(s_j||s_j.state||t_k||h_{i-1})$, whose output is denoted by h_i, and (*ii*) filter operation that deletes the device-id d_j and stores $\langle s_k, 0, t_i\rangle$ on the disk. Now for all the successive sensor readings until encountering a sensor reading, say l, with $state = 1$, all the sensor readings are discarded (not stored on the disk), as well as, the hash function is not executed. However, the sealing function computes the hash function on the l^{th} sensor reading ($\langle d_x, s_k, 1, t_l\rangle$), as: $H(d_x||s_j||1||t_l||h_i)$ and stores l^{th} sensor reading on the disk.

3) ***Sealing Data for User Data/Service Verification***: While capturing *user-associated data*, users may wish to verify their user-associated data against notified messages. Note that the protocols presented so far will require entire cleartext data to be sent to the verifier to attest log integrity (it will be clear soon in §V-C). However, such cleartext data transmission is not possible in the case of user-associated data verification, since it may reveal other users' privacy. Thus, to allow verification of user-associated data (or service/query result[4] verification), we develop a new sealing method, consists of the three phases: (*i*) chunk creation, (*ii*) hash-generation, and (*iii*) proof-of-integrity creation. Chunk creation phase of this new sealing method is identical to the above-mentioned chunk creation phase 1; see §V-B1. Below, we only describe PHASE 2 and PHASE 3.

PHASE 2: Hash-generation. Consider a chunk, $\mathcal{C}_x$, that can have at most n sensor readings, each of them of the format: $\langle d_i, s_j, s_j.state, t_k\rangle$. The sealing function works on i^{th} sensor reading ($1 \leq i \leq n$) is explained next.

Our objective is to hide users' device-id and its frequency-count (*i.e.*, which device-id is prominent in the given chunk). Thus, on the i^{th} sensor reading, the sealing function mixes d_j with t_k, and then, computes a hash function over them, denoted by $H(d_j||t_k)$ that results in a digest value, say o_i. Note that hash on device-ids mixed with time results in two different digests for more than one occurrence of the same device-id. Note that o_i helps the user to know his presence/absence in the data during attestation, but it will not prove that tampering has not happened with the data. Then, the sealing function mixes o_i with the sensor state (to produce a proof of sensor state) of the i^{th} sensor reading, and on which it computes the hash function, denoted by $H(o_i||s_j.state)$ that results in a hash digest, denoted by hu^x_i. After processing the i^{th} sensor reading of the chunk $\mathcal{C}_x$, the enclave writes o_i on the disk, (with, of course, the i^{th} sensor reading value, if $s_j.state$ is not passive). After processing all the n sensor readings of the chunk $\mathcal{C}_x$, the sealing function computes XOR operation on all hash digests: $hu^x_1 \oplus hu^x_2 \oplus \ldots \oplus hu^x_n$, whose output is denoted by hu^x_{end}. (Reason of computing hu^x_{end} will be clear in §V-C)

PHASE 3: Proof-of-integrity creation for user. The sealing function prepares an immutable proof-of-integrity for users, denoted by $\mathcal{PU}$, for each chunk and writes on the disk. Likewise, proof-of-integrity for entire log verification, $\mathcal{PI}$ (§V-B1), for each chunk, the sealing function obtains S_{eoc}; refer to PHASE 3 in §V-B1. Now, for producing $\mathcal{PU}$ for the chunk $\mathcal{C}_x$, the sealing function: (*i*) signs hu^x_{end} XORed with S^x_{eoc} with the private key of the enclave, and (*ii*) writes the signed output with the secret-string of the chunk, g^b, as $\mathcal{PU}_{\mathcal{C}_x}$.

$$\mathcal{PU}_{\mathcal{C}_x} = (g^b, Sign_{PR_E}(hu^x_{end} \oplus S^x_{eoc}))$$

Example. Please see Figure 4, where the last box shows PHASE 2 execution on four sensor readings. Note that left-

[3]We can only compress $x > 1$ continuous sensor readings, say $\langle *_d, *_s, *_s.state = 0, *_t\rangle$ (where $*_d$, $*_s$, and $*_t$ denote any device-id, sensor-id, and time, respectively) to produce a proof that such x readings have been deleted. However, we cannot compress x sensor readings having $*_s.state = 1$, since it disallows to verify service integrity (*e.g.*, a user query, how many time a user has visited a location, cannot be verified, if x readings with $*_s.state = 1$ have been deleted).

[4]SP may develop applications based on sensor data, *e.g.*, location tracking or information about previously visited places, as mentioned in §I. The users, who access these services, may also wish to verify the query results. A user may ask queries, *e.g.*, past locations during time t_1 and t_2, or the current location of a friend (recall our objective in this paper is not to show how one can develop these services). However, SP may tamper with the data to show the wrong results.

hand-side hash digest is obtained for each device-id mixed with time to prevent frequency-count attacks, and then, this hash digest is mixed with the sensor state on which another hash function is executed. After computing hu_4, a hash function is executed on all hash digests hu_i ($1 \leq i \leq 4$) to produce hu_{end} that is mixed with the secret-string of the chunk to produce the proof-of-integrity for the user, $\mathcal{PU}$.

C. Attestation Phase

The attestation phase contains two sub-phases: (*i*) key establishment between the verifier and service provider to retrieve logs, and (*ii*) verification of the retrieved logs.

1) ***Key Establishment****:* The log attestation method embodies: (a) secure log retrieval by the auditor from an untrusted service provider, and, (b) secure proof reconstruction by the auditor on received logs. The secure log retrieval is crucial for proof validation and non-trivial in the presence of decentralized verifier model where anyone can fetch a remote request to attest the logs. Our log retrieval scheme is based on AKE where the auditor (i.e., verifier) and the service provider (i.e., prover) dynamically establish a session key (shown in Figure 6). In order to ensure that the for each attestation request (from verifier to prover) the sender must be authenticated prior to the session key establishment. This dynamic key establishment provides forward secrecy for all future sessions, such that, for any compromised session in future, all sessions in past remain secure. To achieve these properties we use SIGMA [6] protocol which is the cryptographic base for Internet Key Exchange (IKE) protocol. The family of SIGMA protocols are based on Diffie-Hellman key exchange. We only show a 3-round version of SIGMA as it provides sender-identity protection during the key establishment process. In our solution the prover is centralized service provider but the verifier can be anyone and therefore, we use this identity protection method to achieve verifier's identity privacy during the session.

Without the loss of generality, we first define the Computational Diffie Hellman (CDH) [30]. During the session the verifier and the prover must execute the computations within a cyclic group $G = \langle g \rangle$ that has generator g of prime order q. According to the CDH assumption the computation of a Discrete Logarithm DL function on public values (g, g^x, g^y) is hard within the cyclic group G. In particular, given the publicly known value g it is hard to distinguish g^{xy} from g^x and g^y without knowing the ephemeral secret x and y.

Figure 6 depicts the SIGMA based communication flow between a verifying user and the service provider. Initially, a verifying user choose a secret value x, computes g^x as a public exponent and send it to the prover. Similarly, prover choose a secret value y, compute g^y, and then receive the public exponent g^x from user. Next, the prover computes a joint secret value as $e = g^{xy}$ and also use it to derive a MAC key MAC_k. The prover compose a message structure as $[g^y, SP_{id}, MAC_k(SP_{id}|g^x|g^y)]$ where g^y is prover's public exponent for session key generation, SP_{id} is the identity of sender of this message ('or' the recipient of initial mes-

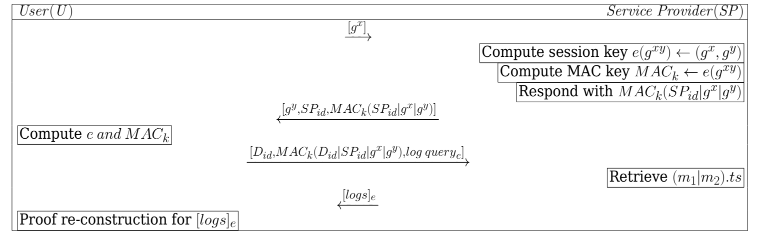

Figure 6: Secure log retrieval by the verifier.

sage from verifier), and, $MAC_k(SP_{id}|g^x|g^y)$ is the message authentication code generated on sender's identity and all public exponents used for session key derivation. Note that in order to generate this MAC a separate key MAC_k is derived from the session key e. Next, the verifier receive this message and retrieve the public exponent g^y to generate a local copy of session key e as well as the message authentication code generating key MAC_k derived from e. At this stage both parties have locally computed a secure session key e, however, there is still an identity disclosure required from the verifier to prove that the freshly generated key e binds to an authentic identity holder D_{id}. This message is consists of $[D_{id}, MAC_k(D_{id}|SP_{id}|g^x|g^y), log\ query_e]$ where D_{id} is the identity of key initiator (i.e., verifying user in this), $MAC_k(D_{id}|SP_{id}|g^x|g^y)$ is the message authentication code on initiator's identity D_{id}, responder's identity SP_{id}, and all public exponents exchanged so far, $log\ query_e$ is the log query protected with the freshly generated session key e. This step binds all public key exponents exchanged with the claimed identity holders together and marks the end of authenticated session key exchange process. The service provider then retrieve the corresponding logs $(m_1|m_2).ts$ as per the log query in last message and send the encrypted log response as $[logs]_e$.

2) ***Verification of Logs****:* This section presents procedures for log verification at the auditor and a user.

Verification process at the auditor. Recall that the auditor can verify any part of the sensor data. Suppose the auditor wishes to verify a chunk $\mathcal{C}_x$. Hence, entire sensor data (the data written in first box of Figure 4) of the chunk $\mathcal{C}_x$, secure random strings g^a, g^b, and g^c (see Figure 5), and proof-of-integrity ($\mathcal{PI}$) are provided to the auditor. The auditor performs the same operation as in PHASE 2 of §V-B3. Also, the auditor computes the secret end-of-chunk string $S^x_{eoc} = g^a \oplus g^b \oplus g^c$. At the end, the user matches the results of $h^x_n \oplus S^x_{eoc}$ against the decrypted value of received $\mathcal{PI}$, and if both the values are identical, then it shows that the entire chunk is unchanged.

Verification process at the user. If the user wishes to verify his data in a chunk, say $\mathcal{C}_x$, the user is provided all hash digested computed over device-id and time (o_i, see the last box in Figure 4), time, sensor state, secure random strings g^a, g^b, and g^c (see Figure 5), and the proof $\mathcal{PU}$ by SP. Since, the user knows her device-id, first, the user verifies her occurrences in the data by computing the hash function on her device-id mixed with provided time values and compares against received hash digests. This confirms the user's presence/absence in the data. However, to verify that no

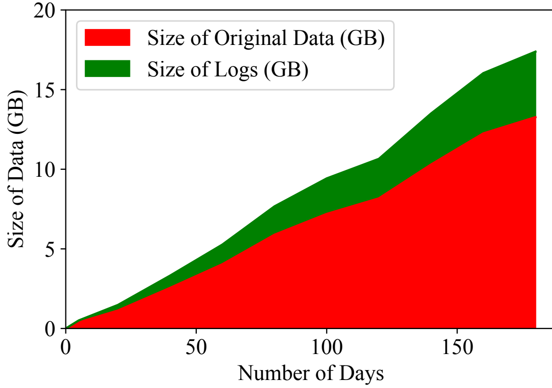

Figure 7: Exp 1: Storage overhead.

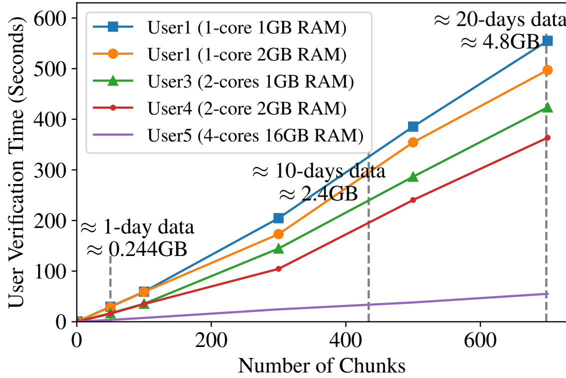

Figure 8: Exp 4: User verification time.

hash-digest is modified/deleted by SP, the user computes the hash function on the sensor state mixed with the received o_i ($1 \leq i \leq n$, where n in the number of sensor readings in $\mathcal{C}_x$) and computes $hu^x_{end} = h^x_1 \oplus h^x_2 \oplus \ldots \oplus h^x_n$. Finally, the user computes $hu^x_{end} \oplus S^x_{eoc}$ and compares against the decrypted value of $\mathcal{PU}$. Note that if any hash-digest o_i is deleted by SP, this comparison operation will fail.

VI. Experimental Evaluation

This section presents our experimental results on live WiFi data. We execute IoT Notary on a 4-core 16GB RAM machine equipped with SGX at Microsoft Azure cloud.

Setup. In our setup, the IT department at UCI is the trusted infrastructure deployer. It also plays the role of the trusted notifier (notifying users over emailing lists). At UCI, 490 WiFi sensors, installed over 30 buildings, send data to a controller that forwards data to the cloud server, where IoT Notary is installed. The cloud keeps cryptographic log digests that are transmitted to the verifier, while sensor data, qualifies data-capture rules, is ingested into realtime applications supported by TIPPERS. We allow users to verify the data collected over the last 30minutes (min).

Dataset size. Though IoT Notary deals with live WiFi data, we report results for data processed by the system over 180 days during which time IoT Notary processed 13GB of WiFi data including 110 million WiFi events.

Data-capture rules. We set the following four data-capture rules: (*i*) *Time-based*: always retain data, except from t_i to t_j each day; (*ii*) *User-location-based*: do not store data about specified devices if they are in a specific building; (*iii*) *User-time-based*: do not capture data having a specific device-id from t_x to t_y ($x \neq i$, $y \neq j$) each day; and (*iv*) *Time-location-based*: do not store any data from a specific building from time t_x to t_y each day. The validity of these rules were 40-days. After each 40-days, variables i, j, x, y were changed.

Exp 1. Storage overhead at the cloud. We fix the size of each chunk to 5MB, and on average, each of them contains $\approx$ 37K sensor readings, covering around 30min data of 30 buildings in peak hours. Based on 5MB chunk size, we got 3291 chunks[5] for 180 days. For each chunk, the sealing function generates two types of logs: (*i*) for auditor verification that produced proof-of-integrity $\mathcal{PI}$ of size 512bytes, and (*ii*) for user verification that produces hashed values (see Figure 4) and proof-of-integrity for users $\mathcal{PU}$ of size 1.05MB. Figure 7 shows 180-days WiFi data size without having sealed logs (red color) and with sealed logs (green color).

Exp 2. Performance at the cloud. For each 5MB chunk, the sealing function took around 310ms to seal each chunk. This includes time to compute $\mathcal{PI}$, $\mathcal{PU}$ and encrypt them.

Exp 3. Auditor verification time. The auditor at our campus has a 7th-Gen quad-core i7CPU and 16GB RAM machine. It downloads the chunks from the cloud and executes auditor verification. The auditor varied the number of chunks from 1 to 3000; see Table I. Note that to attest one-day data across 30 buildings, the auditor needs to download at most 50 chunks, which took less than 1min to verify. Observe that as the number of chunks increases, the time also increases, due to executing the hash function on more data.

Number of Chunks	1	50	100	500	1000	3000
≈ duration (day)	30-60min	1-2	2-5	8-18	35-55	175
Time (seconds)	1	49	102	544	1160	4400

Table I: The auditor verification time. Duration varies due to different class schedules in buildings and working hours.

Exp 4: Verification at a resource-constrained user. To show the practicality of IoT Notary for resource-constrained users, we considered four types of users, differing on computational capabilities (*e.g.*, available main memory (1GB/2GB) and the number of cores (1 or 2 cores)). Each user verified 1-, 10-, and 20-days data; see Figure 8. Observe that verifying 1-day data, which is $\approx$ 50 blocks, at resource-constrained users took at most 30s. Further, as the number of blocks increases, the computational time also increases, where the maximum computational time to verify 20-days data was $<$ 10min. As the days increase, so does data transmitted to the user, which spills over to disk causing an increased latency. Also, we execute the same experiment on a powerful user having 4-core and 16GB machine. Observe that as the number of core and memory increase, it results in parallel processing and absence of disk data read. Hence, the computation time decreases (see user 5 in Figure 8).

[5] The reason of getting more chunks is that during non-peak hours 5MB chunk can store sensor readings for more than one hour. However, as per our assumption, we allow the user to verify the data collected over the last 30min. Hence, regardless of chunk is full or not, we compute the sealing function on each chunk after one hour.

Exp 5: Impact of communication. We measured the communication impact when a verifier downloaded the sensor data and/or sealed log for attestation. Consider a case when the verifier attests only one-hour/one-day data. The average size of one-hour (one-day) data in a peak hour was 14MB (250MB) having 103K (1.2M) connectivity events, while in a non-peak hour, it was 2MB (50MB) having 13.5K (320K) connectivity events. When using slow (100MB/s), medium (500MB/s), and fast (1GB/s) speed of data transmission, the data transmission time in case of 1-hour/1-day data was negligible.

Exp 6: Impact of parallelism. The processing time at each server can be reduced by parallelizing the computation. We investigated the impact of parallel processing to seal a 5MB chunk when having the number of threads 2 or 4 that took 322ms and 310ms, respectively. Increasing more threads did not provide speed-up, since the execution time increases due to thread maintenance overheads. Note that we only parallelized the hash function computation for $\mathcal{PU}$, (while $\mathcal{PI}$ cannot be computed in parallel, due to the formation of hash chains).

VII. Conclusion

This paper presented a framework, IoT Notary for sensor data attestation that embodies cryptographically enforced log-sealing mechanisms to produce immutable proofs, used for log verification. Our solution improves the naïve end-to-end encryption model, where retroactive verification is not provable. The service verification mechanism on failing at users allows them to revoke services of the concerned IoT space. Therefore, a user is not required to blindly trust in the IoT space, and we empower the users with the right-to-audit instead of right-to-own the data captured by sensors. IoT Notary is a part of a real IoT system (TIPPERS) and provides verification on live WiFi data with almost no overheads on users.

In addition, we are exploring several improvements to the current implementation of IoT Notary: (*i*) *Chunk-size determination*: In the current setting, we have fixed the chunk size to 5MB. Chunk size has an impact on both cryptographic operations (larger the chunk size, more amortized the overheads) and logging cost at SP (each chunk needs to be stored durably, so larger the chunk, lower the overhead). However, a large-sized chunk can increase service latency, and further, increase data access time during verification. (*ii*) *Supporting data-capturing rules where conditions depend on the value of other sensors*: The current implementation does not support context-dependent rules, where the context is determined based on data captured by other sensors. Extending the system to handle such data-capture rules, *e.g.*, "Do not capture my WiFi connectivity data if I am the only person connected to the access point," is a non-trivial challenge.

References

[1] D. Kozlov *et al.*, "Security and privacy threats in IoT architectures," in *BODYNETS*, 2012, pp. 256–262.

[2] S. Mehrotra *et al.*, "TIPPERS: A privacy cognizant IoT environment," in *PerCom Workshops*, 2016, pp. 1–6, http://tippersweb.ics.uci.edu/.

[3] E. Fernandes *et al.*, "Security implications of permission models in smart-home application frameworks," *IEEE Security & Privacy*, vol. 15, no. 2, pp. 24–30, 2017.

[4] A. Rao *et al.*, "Expecting the unexpected: Understanding mismatched privacy expectations online," in *SOUPS*, 2016, pp. 77–96.

[5] S. Madakam *et al.*, "Security mechanisms for connectivity of smart devices in the internet of things," 2016.

[6] H. Krawczyk, "Sigma: The 'SIGn-and-MAc' approach to authenticated diffie-hellman and its use in the IKE protocols," in *CRYPTO*, 2003.

[7] A. Ibrahim *et al.*, "AID: autonomous attestation of IoT devices," in *SRDS*, 2018.

[8] N. Asokan *et al.*, "Seda: Scalable embedded device attestation," in *CCS*, 2015, pp. 964–975.

[9] A. Ibrahim *et al.*, "Darpa: Device attestation resilient to physical attacks," in *WiSec*, 2016, pp. 171–182.

[10] M. Ambrosin *et al.*, "SANA: secure and scalable aggregate network attestation," in *CCS*, 2016, pp. 731–742.

[11] X. Carpent *et al.*, "Remote attestation of iot devices via smarm: Shuffled measurements against roving malware," in *HOST*, 2018, pp. 9–16.

[12] J. Wang *et al.*, "Enabling security-enhanced attestation with Intel SGX for remote terminal and iot," *IEEE TCDICS*, vol. 37, no. 1, pp. 88–96, 2018.

[13] D. C. G. Valadares *et al.*, "Achieving data dissemination with security using FIWARE and Intel software guard extensions (SGX)," in *ISCC*, 2018, pp. 1–7.

[14] J. Frankle *et al.*, "Practical accountability of secret processes," in *USENIX*, 2018, pp. 657–674.

[15] S. Eskandarian *et al.*, "Certificate transparency with privacy," *PoPETs*, vol. 2017, no. 4, pp. 329–344, 2017.

[16] S. A. Crosby *et al.*, "Efficient data structures for tamper-evident logging." in *USENIX*, 2009, pp. 317–334.

[17] S. Zawoad *et al.*, "Towards building forensics enabled cloud through secure logging-as-a-service," *IEEE TDSC*, vol. 13, pp. 148–162, 2016.

[18] E. Boyle *et al.*, "Function secret sharing," in *EUROCRYPT*, 2015.

[19] Y. Zhang *et al.*, "vSQL: Verifying arbitrary SQL queries over dynamic outsourced databases," in *IEEE SP*, 2017, pp. 863–880.

[20] V. Costan *et al.*, "Intel SGX explained," *IACR Cryptology ePrint Archive*, vol. 2016, p. 86, 2016.

[21] https://newsroom.intel.com/newsroom/wp-content/uploads/sites/11/2017/09/8th-gen-intel-core-product-brief.pdf.

[22] W. Wang *et al.*, "Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX," in *CCS*, 2017, pp. 2421–2434.

[23] S. Lee *et al.*, "Inferring fine-grained control flow inside SGX enclaves with branch shadowing," in *USENIX Security*, 2017, pp. 557–574.

[24] J. Wang *et al.*, "Interface-based side channel attack against intel SGX," *CoRR*, vol. abs/1811.05378, 2018.

[25] M. Shih *et al.*, "T-SGX: eradicating controlled-channel attacks against enclave programs," in *NDSS*, 2017.

[26] V. Costan *et al.*, "Sanctum: Minimal hardware extensions for strong software isolation," in *USENIX Security*, 2016, pp. 857–874.

[27] L. Lamport, "Password authentication with insecure communication," *Commun. ACM*, vol. 24, no. 11, pp. 770–772, 1981.

[28] M. S. P. M. M'Raihi, D. and J. Rydell, "TOTP: Time-based one-time password algorithm," in *RFC 6238 Internet Engineering Task Force*. [Online]. Available: https://doi.org/10.17487/RFC6238

[29] C. Priebe *et al.*, "Enclavedb: A secure database using SGX," in *IEEE SP*, 2018, pp. 264–278.

[30] W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Trans. Information Theory*, vol. 22, no. 6, pp. 644–654, 1976. [Online]. Available: https://doi.org/10.1109/TIT.1976.1055638

Appendix A
Future Password-based Notice Messages

The authorization and validation of data-capture rule declaration is based on Time-based One-Time Passwords (TOTP) [28] such that the users are pre-initialized to receive the notifications regarding any sensitive data collection until an upcoming reference time t_{max} happens in the future. All notifications from the time of initialization t_{init} until the future reference time are chained through exactly in the timeline sequence rounded over the equally distant interval of an epoch size $\mathcal{E}$. Once the chain of these pre-initialized notifications nears the end, *i.e.*, the current time t_{cur} is same as the future reference time t_{max} (as declared in the beginning); the service provider initializes a new chain for the next set of notifications until a new reference time t'_{max} in future such that

$$t'_{max} > t_{init} > t_{init}$$

Notification phase: The service provider selects an upcoming reference time t_{max} for which all registered users will be notified. The service provider computes a local secret passphrase such as x and compute $y_{init} = \mathcal{H}^{length}(x)$ where *length* is the length of the chain as $t_{max} - t_{init}/\mathcal{E}$ and $\mathcal{H}^{length}$ is the *length* successive iterations of hash function $\mathcal{H}$ over the value x. The registered users receive a tuple $\langle y_{init}, t_{init}, t_{max} \rangle$ to compute the hash-chain at each interval during $t_{max} - t_{init}$. Let us assume that the service provider dispatch a notification during ith epoch $\mathcal{E}_i$ in the hash-chain. The service provider computes $y_i = \mathcal{H}^{length-i}(x)$ and send it to all registered users along with the message contents and a HMAC over passphrase and the message contents.

$$\langle y_i, msg, HMAC(y_i, msg) \rangle$$

Subsequently, all registered users retrieve the previously validated y_{i-1} (which is y_{init} at the beginning of the hash-chain) and use it to validate the current passphrase y_i. Note that the users must compute $\mathcal{H}(y_i)$ and compare it to previously validated passphrase y_{i-1}. If $\mathcal{H}(y_i) = y_{i-1}$ the user replace previously validated passphrase y_{i-1} with the currently validated passphrase y_i and use it for next passphrase validation. It must be noted that x represent the passphrase at the end of the hash-chain where $t_{cur} = t_{max}$.

It must be noted that the usage of same hash function for the passphrase computation is vulnerable to birthday attacks. In addition the passphrase (hence the notification) remain valid and irrevocable for an indefinite duration which makes it vulnerable to leakage attacks. Therefore, to allow ephemeral passphrases and the notifications that are exposed for a short duration only a time-based counter is required. In particular, a primary hash function and a time-based counter can be used here together to derive as many separate hash functions as the length of the hash-chain $length$. In that case the initial passphrase would be:

$$y_{init} = \mathcal{H}_{length}(\mathcal{H}_{length-1}(\ldots(\mathcal{H}_1(x))\ldots))$$

In addition, for any notification validation during $t \in (t_{init}, t_{max}]$ the service provider must yield $\langle y_t, msg, HMAC(y_t, msg) \rangle$ such that

$$y_t = \mathcal{H}_{t_{max}-t}(\mathcal{H}_{t_{max}-t-1}(\ldots(\mathcal{H}_1(x))\ldots))$$

In order to validate the passphrase and the notifications at time $t_{cur} > t_{prev}$, the registered users must compute the hash-chain from $(t_{max} - t_{prev})$ to $(t_{max} - t_{cur} + 1)$.

Appendix B
AKE Adaptations

Authenticated Key Exchange. In order to provide the secure log retrieval we incorporate an Authenticated Key Exchange (AKE) protocol. The AKE based communication enables secure end-to-end retrieval of logs and assures secure authentication without any active impersonation, namely, man-in-the-middle attacks. Our AKE scheme is a natural extension of SIGMA (SIGn-and-MAc) Authenticated Key Exchange (AKE) protocol [6] based on Diffie-Hellman key exchange as defined below.

Computational Diffie-Hellman [30] Let $\langle g \rangle$ be a cyclic group G of order q. There is no efficient probabilistic algorithm $\mathbb{A}_{CDH}$ that produces g^{xy} given (g, g^x, g^y), where x, y are randomly chosen group elements.

We further show that the proposed AKE can be used in two different modalities: (i) with the static public-private key pair on each sides of the AKE participants. (ii) with the static as well as an ephemeral public-private key pair on each sides of the AKE participants. The difference is that second method allows participants to choose an additional pair of ephemeral keys for each new session which enables forward secrecy in terms of protecting all previous sessions. Below we briefly describe the protocol steps for original SIGMA, TLS and the adaptation based on forward secrecy property.

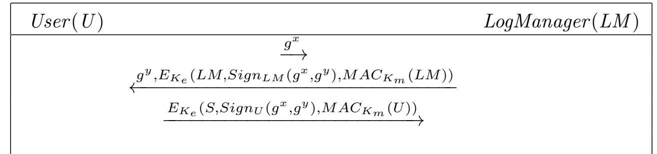

Figure 9: SIGMA protocol with sender identity protection.

SIGMA. A 3-round version of SIGMA protocol [6] provides sender identity protection while the key establishment is based on an authenticated DH protocol using the digital signatures. A session key K_s, an encryption key K_e and a message authentication key K_m are derived from g^{xy}. Note that keys K_s, K_e, and K_m are computationally independent from each other. Here, messages are decrypted using the key K_e and the identity is verified using the key K_m. There exists a 4-round version of SIGMA protocol in case receiver identity protection is a required feature. However, a 4-round version is prone to *reflection attacks* therefore, we chose to base our protocol construction on a 3-round version of the SIGMA protocol.

Transport Layer Security. We show that existing Internet security protocol such as TLS can adapt in our security model. As we can see in Figure 10, a mutually authenticated channel

requires multiple phases such as: negotiation phase (to establish cipher suit between user and the server), authentication or X.5.09 certificate exchange phase (to validate the identity and receive key exponents), and, a handshake complete phase (to exchange a complete set of individual messages exchanged during the handshake); and would impose scalability issues as compared to our proposed approach. Furthermore, connections can be resumed over TLS channels using 0-RTT mechanism and would scale a little better in case the users resume connections more frequently.

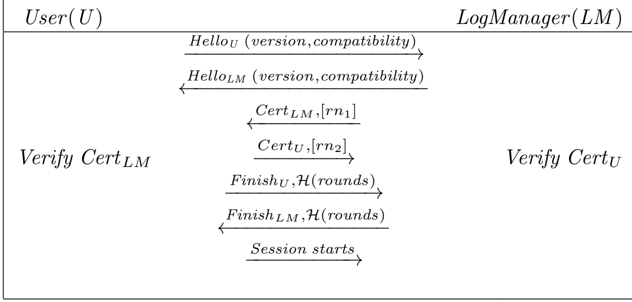

Figure 10: Transport Layer Security protocol.

Forward Secrecy. The Forward Secrecy (FS) provides security assurance regarding the session key used in past even if the long term key is compromised in the upcoming future. In case of key reveal attack all transcripts from past and future are vulnerable to manipulation, therefore, FS property provides an assurance that even if a long-term secret key (of one of the parties in communication) is revealed to the attacker still all keys derived during the past session remain secure and that the communication is secure in forward direction. Our solution can adapt to open a forward secret channel such that log exchange between the verifier and the SP remains confidential across the sessions.

The authenticated key establishment protocol would assume a static public-private key pair and an ephemeral public-private key pair at both parties, such that the session key would remain secret even if one of the long term static private key is revealed to an attacker.

- $PK_{D_i} = g^{SK_{D_i}}$ and $PK_{SP} = g^{SK_{SP}}$ as static public key for the verifying user U and log manager SP.
- $X = g^{\mathcal{H}(x,a)}$ and $Y = g^{\mathcal{H}(y,b)}$ as ephemeral public key where x and y are ephemeral secret keys and a, b are random values.
- The session key e is computed through $e = (\sigma, U, SP, X, Y)$ where $\sigma_U = (Y^{\mathcal{H}(x,a)})$ and $\sigma_{SP} = (X^{\mathcal{H}(y,b)})$.

The key derivation would provide the guarantee that the ephemeral key leakage or the static private key leakage would not reveal the session keys used in coming sessions. Therefore, all the sessions from past are secure to any key reveal attacks in upcoming sessions. The adversary might have stored the transcripts from previous sessions but it would not be useful without possessing the static secret and ephemeral secret key.

APPENDIX C
SECURITY ANALYSIS

In this section, we analyze the security claims regarding the validation of a bygone policy notification between user and service provider. The proposed solution highlights the possible scenario for policy violation, both, with and without the evidence of violation. In every possible scenario both properties (given below) must be satisfied regarding the policy infringement.

Theorem 1. *Let device D_i upload raw observations to SP and the corresponding log files A, B, C with the proof-of-order $\mathcal{P}$ are securely transmitted to SP. Assume that:*

- *integrity is retained across the channel, i.e., from devices to service provider, from service provider to the verifier, and on the untrusted storage at SP (see Lemma 1).*
- *logs are secure and the proof-of-order is intact (as shown in Theorem 2).*
- *policy adherence is accountable (see Lemma 2).*
- *policy violation is refutable (see Lemma 3).*

Our proposed solution ensures the integrity property regarding the logs at the channel and at the storage server. Other two properties, *i.e.*, accountability and refutability are inherently based on this integrity property. Therefore, integrity is the most significant and beyond that the first property is to provide the accountability feature for policy violation based on authentic logs; and second property is to provide disapproval feature for policy violation in the absence of authentic logs. In both cases, the attestation process boils down to a secure log verification mechanism that resolves the policy violation.

Lemma 1. *Integrity property is satisfied by using end-to-end encryption at log sealing unit (closer to the log origin itself, i.e., log generating sensors), while on channel, and, by proof-of-order validation, while not on channel.*

Proof: The log integrity is subject to malicious tampering either on the *channel* or at the destined *storage*. The integrity of logs on the storage at service provider is provable through a fusion between the successive hash-chains on log entries and the XOR-links between the successive hash-chains. The order on order guarantees that logs remain tamper-proof. In addition, the log orderliness is periodically signed by the sealing agent. Therefore, the log tampering requires signature forging on a fabricated order of logs. Therefore, our end-to-end encryption followed by a collision-resistant hash-chains provide integrity both at the channel and the end storage. ■

Lemma 2. *Service accountability property is satisfied by non-repudiation in the presence of a bidirectional communication between the log storage and the service auditor.*

Proof: The *non-repudiation* property yields *accountability* based on the conditionally available log evidences. The service accountability is based on periodic auditing of the proof-of-order. The auditing requires re-construction of successively binding hash-chains on log entries and the adjacency links

between these hash-chains. Moreover, the accountability is incomplete without correct retrieval of secured logs. Hence, a bidirectional verification is required here. Therefore, a correct authorization (with respect to the log verifier) is followed by an AKE between the verifier and service provider. ■

Lemma 3. *Refutability property is satisfied in the absence of a bidirectional communication between the log storage and the policy auditor.*

Proof: In the absence of a correct authorization (meaning the aforementioned accountability property is not fulfilled) there would not exist log based evidences from the log storage. Note that here only the logs already on the storage are required and the log integrity property regarding log generating device still prevails. However, the refutability property[6] is mutually exclusive with the accountability property. Hence, only a unidirectional verification is required here. Such that either a correct authorization yields log to the verifier or the log generating device resume refutability in the absence of any logs. ■

$\mathbf{Exp}^{\Sigma}_{\mathcal{A}}(1^{\lambda})$	**Oracle** $log(\mathcal{C})$
$a = Gen(1^{\lambda})$	$\mathcal{P} = log(\mathcal{C}, g^{a})$
$S_{eof} = g^{a}$	$Sign((\mathcal{P}), (g^{a}_{(i-1)}, g^{b}_{i}, g^{c}_{(i+1)}))$
`if` $\mathcal{P}' = \mathcal{A}(\mathcal{C}', g^{a'})$ `where` $\mathcal{P} = \mathcal{P}'$	
`then return` 1	
`else return` 0 `where` $\Pr[\mathbf{Exp}^{\Sigma}_{\mathcal{A}}(1^{\lambda}) = 1] \leq \epsilon^{\Sigma}$	

Figure 11: Security experiment for logging scheme

Proof Sketch. [Theorem 1.] The proof sketch for Theorem 1 is a direct implication of Lemma 1, 2, and 3. The polynomial advantage of an adversary $\mathcal{A}$ in a worst case attack scenario where an adversary is present on channel and the storage, simultaneously, is negligibly smaller. Therefore, at worst an adversary gain an advantage, either, by (*a*) attacking logs on storage $\Pr[Storage_{\mathcal{A}}] \leq \epsilon$ or, (*b*) attacking logs in transit $\Pr[Transit_{\mathcal{A}}] \leq \epsilon$. □

Theorem 2. *Let Σ be the log sealing scheme that generates $(\mathcal{P}, Header)$ as the verifiable proof-of-order with hash-chain $\mathcal{C}$ signed by SP and the advantage of an adversary is $\mathbf{Adv}(\mathcal{A})$.*

- *The advantage $\mathbf{Adv}(\mathcal{A})$ in preparing a proof-of-order $(\mathcal{P}', Header')$ such that $\mathcal{P}' = \mathcal{P}$ and $Header' = Header$ is negligible as the advantage $\mathbf{Adv}(\mathcal{A})$ in fabricating a proof $\mathcal{P}'$ does not accumulate over time.*

Proof: In order to prepare a proof $\mathcal{P}'$ same as $\mathcal{P}$ an adversary $\mathcal{A}$ must succeed in the following order.

- *Case 1:* An adversary must produce the exact sequence of each log entry in a file even before the proof is created for an original log sequence since the log files are append-only. However, it is difficult to predict the log entries in advance and to produce the identical hash-chain. Otherwise an adversary must record the original hash-chain and find the right sequence of matching preimage log entries that would generate the same hash output. However, assuming the underlying security of the hash function this preimage attack has a negligible probability of occurrence given the time-bound until next file is appended to the sequence.
- *Case 2:* Assume that an adversary have successfully predicted the sequence of log entries and have produced the identical hash-chain. Now, the adversary must *compute* the secret *eoc* string to completely fabricate the proof. However, if the adversary have successfully computed the secret string it can then be used as a subprocedure to break the Discrete Logarithm (DL) problem.
- *Case 3:* Assume that an adversary have successfully predicted the sequence of log entries and have produced the identical hash-chain. Now, the adversary *predicts* the secret *eoc* string. However, it requires to forge a signature on this fabricated proof without possessing the secret signing key. Therefore, the adversary is assumed to produce only random signatures for an existentially unforgeable signing scheme.

Therefore, an adversary has negligible advantage $\Pr[\mathbf{Exp}^{\Sigma}_{\mathcal{A}}(1^{\lambda}) = 1] \leq \epsilon^{\Sigma}$ to win the experiment in each of these cases (as shown in Figure 11). ■

[6]Our definition of refutability must not be mingled with the traditional *deniability* property. Since the conventional definition of deniability provides authorization with non-transferable credentials. In our definition the credential transfer is not a choice but mandatory, and therefore, the ability to deny is preserved with respect to the log generating device in question instead of log verifier (as it would have been with the deniability property).