

# UC Irvine

## UC Irvine Previously Published Works

### Title

IoT Notary: Sensor Data Attestation in Smart Environment

### Permalink

<https://escholarship.org/uc/item/1g9967nx>

### Authors

Panwar, Nisha

Sharma, Shantanu

Wang, Guoxi

et al.

### Publication Date

2019-01-26

### DOI

10.1109/nca.2019.8935047

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# 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 deliberately or IoT devices may transfer user data to a malicious third-party due to cyberattacks, 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,’ based on which a verifier can attest that captured sensor data 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, thereby users can verify their data of one day in less than two seconds.

## I. INTRODUCTION

While fine-grained continuous monitoring by IoT devices (*e.g.*, camera and WiFi access-points) offers numerous benefits and empowers existing systems with new capabilities, it also raises several privacy and security concerns (*e.g.*, smoking habits, gender, and religious belief). 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 [1], 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 the 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 contains access-point-id,  $s_i$ , user device MAC address,  $d_j$ , and the time of the association,  $t_k$ . All SNMP traps  $\langle s_i, d_j, t_k \rangle$  are sent to access-point’s controllers in

Accepted in IEEE International Symposium on Network Computing and Applications (NCA), 2019. For the final version, please refer to the conference proceeding.

This work is based on research sponsored by DARPA under agreement number FA8750-16-2-0021 and partially supported by NSF grants 1527536 and 1545071. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

realtime. The access-point controller anonymizes device MAC addresses (to preserve the privacy of 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. 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 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. Thus, mechanisms 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 are used to control data access, *i.e.*, which applications have access to which data generated at the mobile device (*e.g.*, location and contact list) for which purpose and in which context. However, in IoT settings, the data-capture framework differs from that in the mobile settings, in two important ways:

- 1) Unlike the mobile setting, where applications can seek user’s permission at the time of installation, in IoT settings,

there are no obvious mechanisms/interfaces to seek users' preferences about the data being captured by sensors of the smart environment. Recent work [2] has begun to explore mechanisms using which environments can broadcast their data-capture rules to users and seek their explicit permissions.

- 2) Unlike the mobile setting, users cannot control sensors in IoT settings. 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.

We focus on the above-mentioned second scenario and determine ways to provide trustworthy sensing in an untrusted IoT environment. Thus, the users can verify their data captured by IoT environment based on pre-notified data-capture rules. Particularly, we deal 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 verify the sensor data against pre-notified data-capture rules by a user, without being heavily involved in the attestation process.

**Our contribution and outline of the paper.** We provide:

- A user-centric framework (§III) to ensure trustworthy data collection in untrusted IoT spaces, entitled IoT NOTARY.
- Two models to inform the user about the data-capture rules (§IV-A): notice-only model and notice-and-ACK model.
- A secure log-sealing mechanism (§IV-B) implemented by secure hardware that cryptographically retains logs, data-capture rules, sensors' state, and contextual information to generate a *proof-of-integrity* in an immutable fashion.
- A secure attestation mechanism (§IV-C), mixed with SIGMA protocol [3], 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 §V.

**Full version.** Due to space limitations, we could not describe several details about IoT NOTARY, which are given in the full version [4]. These include: future temporal password-based notification method, log retrieval at the verifier using SIGMA, details of the verification phase, throughput and communication cost experiments, and security proofs.

## II. 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  $s_1, s_2, \dots, s_p$ ), 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

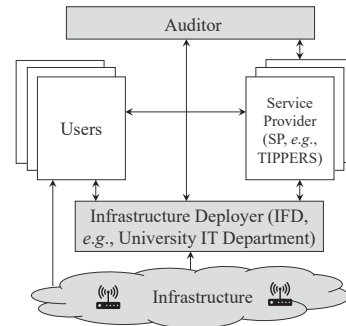


Figure 1: Entities in IoT NOTARY.

devices, especially WiFi access-points that also capture some user-related information in response to services. *E.g.*, 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); denoted by:  $\langle d_i, s_j, t_k, param \rangle$ , where  $d_i$  is the  $i^{th}$  user-device-id,  $s_j$  is the  $j^{th}$  sensor device,  $t_k$  is  $k^{th}$  time, and  $param$  is other parameters (we do not deal with  $param$  field and focus only the first three fields). All sensor data is collected at a controller (server) owned by IFD. The controller may keep sensor data in cleartext or in encrypted form; 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  $dc_1, dc_2, \dots, dc_q$ ). 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<sup>th</sup> 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 conditions on the user-id, as well as, on time and the 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, 0 (or passive). In the remaining paper, unless explicitly noted, opt-out is considered as the default rule, 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 Intel Software Guard eXtension (SGX) [5], 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*). The assumption of secure hardware at a machine is rational with the emerging system architectures, *e.g.*, Intel machines are equipped with SGX [6]. However, existing SGX architectures suffer from side-channel attacks, *e.g.*, cache-line, branch shadow, page-fault attacks [7], which are outside the scope of this paper.

**Users.** Let  $d_1, d_2, \dots, d_m$  be  $m$  (user) devices carried by  $u_1, u_2, \dots, u_{m'}$  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. Let  $d_i$  be the  $i^{\text{th}}$  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 opt-out was a default rule or when a user opted-out from an opt-in default). Users may also require SP to prove service integrity by storing all sensor data associated with the user (when users have opted-in into services), 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 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. Also, an adversarial user may *impersonate* a real user to retrieve the sensor data and secured-log during the verification phase. Thus, an adversarial user may reveal the privacy of the users by observing sensor data. Also, 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 rogue SP, as well as, SP can detect rogue users, and (ii) between SP and the verifier (auditor/user), before sending sensor data to the verifier to prevent any rogue 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 rogue sensor data by SP, as well as, preventing any rogue 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

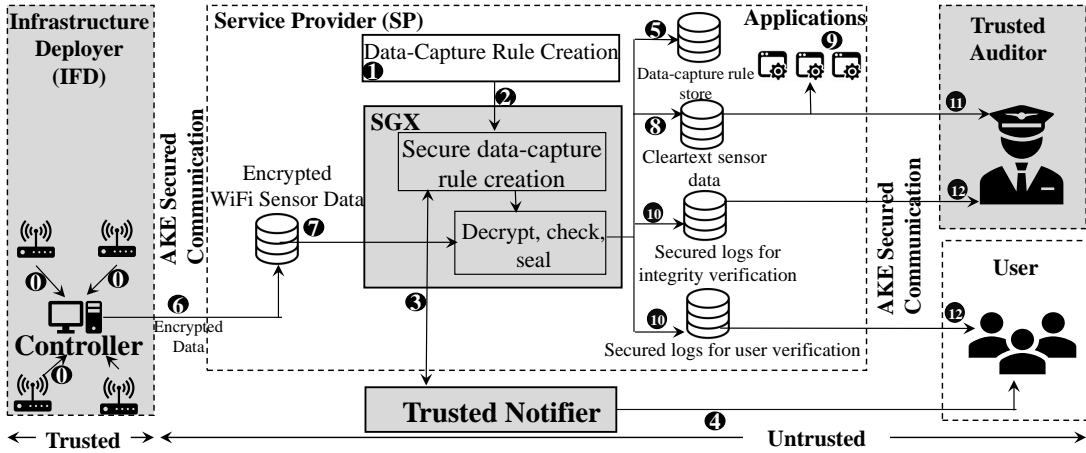


Figure 2: Dataflow and computation in the protocol. Trusted parts are shown in shaded boxes.

any additional information to the user about all the other users during the 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. Assumptions

This section presents assumptions, we made, as follows:

- 1) 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.
- 2) 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 at the controller at IFD, and it ensures that no rogue sensor device can generate the data on behalf of an authentic sensor. Further, we assume that an adversary cannot deduce any information from the dataflow between a sensor and the controller. Recall that in our setting the university IT department collects the entire sensor data from their owned and deployed sensors, before sending it to TIPPERS.
- 3) We assume the existence of an authentication protocol between the controller and SP, so that SP receives sensor data only from authenticated and desired controller.
- 4) 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.
- 5) 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.)

### III. IOT NOTARY

This section presents an overview of the three phases and dataflow among different entities and devices, see Figure 2.

**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 (1). Trusted hardware securely stores data-capture rules (2, 5) and informs the *trusted notifier* (3) that transmits the message to all users (4). Only notice messages need a trusted notifier to transmit the message (see §IV-A).

**Log-sealing phase: Sensor devices to SP messages.** Each sensor sends data to the controller (9). The controller receives the correct data, generated by the actual sensor, as per our assumptions (and settings of the university IT department). The controller sends encrypted data to SP (6) that authenticates the controller using any existing authentication protocol, before accepting data. Trusted hardware (Intel SGX) at SP reads the encrypted data in the enclave (7).

**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 adds a new field, entitled *sensor (device) states*. The sensor state of a sensor  $s_j$  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 (8) that provides services using this data (9). Note that the cryptographically sealed logs and cleartext sensor data are kept at untrusted storage of SP (8, 10).

**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 (8, 10). 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 (12). Note that using an underlying authentication technique (as per our assumptions), auditor/users and SP authenticate each other before transmitting data from SP to auditor/users.

#### IV. 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 user devices into the system. The trusted authority (which is the university IT department in our setup of TIPPERS) generates/renews/revokes keys used by the secure hardware enclave (denoted by  $\langle PK_E, PR_E \rangle$ ) and the notifier (denoted by  $\langle PK_N, PR_N \rangle$ ). The keys are provided to the enclave during the secure hardware registration process. Also,  $\langle PK_{di}, PR_{di} \rangle$  denotes keys of the  $i^{th}$  user device. *Usages of keys:* The controller uses  $PK_E$  to encrypt sensor readings before sending to SP.  $PR_E$  is also used by the enclave to write encrypted sensor logs and decrypt sensor readings.  $PK_N$  is used during the notification phase by SGX to send an encrypted message to the notifier. User device's keys are used during device registration, as given below.

We assume a registration process during which a user identifies 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 push messages to the user device). Such communication is used during the 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 application, see 9 in Figure 2). 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).<sup>1</sup> 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. Since 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 *the 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/some of the users; thus, an encrypted data-capture rule is not helpful. Thus, a trusted notifier ensures that the notice message is sent to all the registered users. 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.

<sup>1</sup>Due to the enclave's limited memory, it cannot keep all valid and non-valid data-capture rules, after a certain size. Thus, the enclave writes all non-valid data-capture rules on the disk, after computing a secured hash digest over all rules. Taking a hash over the rules is needed to maintain the integrity of all rules.

$\langle d_1, s_1, 1, t_1 \rangle$	$h_1 \leftarrow H(d_1    s_1    1    t_1    H(0))$	$o_1 \leftarrow H(d_1    t_1)$	$hu_1 \leftarrow H(o_1    1)$
$\langle d_2, s_2, 1, t_2 \rangle$	$h_2 \leftarrow H(d_2    s_2    1    t_2    h_1)$	$o_2 \leftarrow H(d_2    t_2)$	$hu_2 \leftarrow H(o_2    1)$
$\langle d_3, s_3, 1, t_3 \rangle$	$h_3 \leftarrow H(d_3    s_3    1    t_3    h_2)$	$o_3 \leftarrow H(d_3    t_3)$	$hu_3 \leftarrow H(o_3    1)$
$\langle d_4, s_4, 1, t_4 \rangle$	$h_4 \leftarrow H(d_4    s_4    1    t_4    h_3)$	$o_4 \leftarrow H(d_4    t_4)$	$hu_4 \leftarrow H(o_4    1)$
<b>Sensor data after passing the enclave</b>	$PI_{C_x} \leftarrow g^b, Sign_{PK_E}(h_4 \oplus S_{eoc}^x)$	$hu_{end} \leftarrow hu_1 \oplus hu_2 \oplus hu_3 \oplus hu_4$	
	<b>Sealing function execution for log-integrity</b>	$PU_{C_x} \leftarrow (g^b, Sign_{PK_E}(h_q^{end} \oplus S_{eoc}^x))$	
		<b>Sealing function execution for user's data/ query verification</b>	

Figure 3: Cryptographically sealing procedure executed on a chunk,  $C_x$ . 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 for a chunk,  $C_x$ .

### B. Log Sealing Phase

The second phase does cryptographically sealing of 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 3, 4), 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 (7) in Figure 2) 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 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 §IV-B1. Likewise, the protocol to deal with all sensor data having state one, a protocol can also deal with the case when some sensor data satisfies some data-capture rule, while remaining sensor data does not satisfy any rule (*i.e.*, the state of the remaining sensor data is zero). However, due to page limitations, we skip details of such a protocol.

**1) Sealing Entire Sensor Data:** The sealing operation contains 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,  $C_x$ , 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  $C_x$ , 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  $h_1^x$ . After processing the complete first sensor reading of the chunk  $C_x$ , the enclave writes cleartext first sensor reading of  $C_x$ , *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. It computes a hash function on the second sensor values, while mixing it with the hash digest of the first sensor reading, *i.e.*,  $H(d_2 || s_j || s_j.state || t_{k+1} || h_1^x)$  that results in a hash digest, say  $h_2^x$ . 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  $C_x$ . After processing the last sensor reading of the chunk  $C_x$ , the hash digest  $h_n^x$  is obtained.

**PHASE 3: Proof-of-integrity creation.** Since each sensor reading is written on disk, SP can alter sensor readings, to make 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, the sealing function prepares an immutable proof-of-integrity for each chunk, as follows:






Figure 4: PHASE 3: end of chunk,  $S_{eoc}$ , creation for three chunks. Observe that  $S_{eoc}^x = g^a \oplus g^b \oplus g^c$ .

For each chunk  $C_i$ , the sealing function generates a random string, denoted by  $g^j$ , where  $i \neq j$ . Let  $C_v$ ,  $C_x$ , and  $C_y$  be three consecutive chunks (see Figure 4), based on consecutive sensor readings. Let  $g^a$ ,  $g^b$ , and  $g^c$  be random strings for chunks  $C_v$ ,  $C_x$ , and  $C_y$ , respectively. The use of random strings will ensure that any of the consecutive chunks have not been deleted by SP (will be clear in §IV-C). Now, for producing the proof-of-integrity for the chunk  $C_x$ , the sealing function: (i) executes XOR operation on  $g^a$ ,  $g^b$ ,  $g^c$ , whose output is denoted by  $S_{eoc}^x$ , where  $eoc$  denotes the end-of-chunk; (ii) signs  $h_n^x$  XORed with  $S_{eoc}^x$  with the private key of the enclave; and (iii) writes the proof-of-integrity for log verification of the chunk  $C_x$  with the random string  $g^b$ , as follows:

$$\mathcal{PI}_{C_x} = (g^b, \text{Sign}_{PR_E}(h_n^x \oplus S_{eoc}^x))$$

**Note.** We do not generate the proof for each sensor reading. The enclave writes only the proof and the random string for each chunk to the disk, which is accessible by SP. Further, the sensor readings having the state one are written on the disk, based on which SP develops services.

**Example.** Please see Figure 3, 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 includes signed  $h_4 \oplus S_{eoc}^x$  and a 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. Thus, in order to initialize the secure binding for the first chunk, the seed value is used as a substitute random string.

## 2) 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 protocol presented so far requires entire cleartext data to be sent to the verifier to attest log integrity* (it will be clear soon in §IV-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<sup>2</sup> verification), we develop a new sealing method, consists of the three phases: (i) chunk creation, (ii)

<sup>2</sup>The users, who access services developed by SP (as mentioned in §I), may also wish to verify the query results, since SP may tamper with the data to show the wrong results.

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 §IV-B1. Below, we only describe PHASE 2 and PHASE 3.

**PHASE 2: Hash-generation.** Consider a chunk,  $C_x$ , that can have at most  $n$  sensor readings, each of them of the format:  $\langle d_i, s_j, s_j.state, t_k \rangle$ . 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  $hu_i^x$ . After processing the  $i^{th}$  sensor reading of the chunk  $C_x$ , the enclave writes  $o_i$  on the disk. After processing all the  $n$  sensor readings of the chunk  $C_x$ , the sealing function computes XOR operation on all hash digests,  $hu_i^x$ , where  $1 \leq i \leq n$ :  $hu_1^x \oplus hu_2^x \oplus \dots \oplus hu_n^x$ , whose output is denoted by  $hu_{end}^x$ . (Reason of computing  $hu_{end}^x$  will be clear in §IV-C).

**PHASE 3: Proof-of-integrity creation for the 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,  $\mathcal{PI}$  (§IV-B1), for each chunk, the sealing function obtains  $S_{eoc}$ ; refer to PHASE 3 in §IV-B1. Now, for producing  $\mathcal{PU}$  for the chunk  $C_x$ , the sealing function: (i) signs  $hu_{end}^x$  XORed with  $S_{eoc}^x$  with the private key of the enclave, and (ii) writes the signed output with the random string of the chunk,  $g^b$ , as  $\mathcal{PU}_{C_x}$ .

$$\mathcal{PU}_{C_x} = (g^b, \text{Sign}_{PR_E}(hu_{end}^x \oplus S_{eoc}^x))$$

**Note.** The enclave writes hash digests,  $o_i$  for each sensor reading, the proof for user verification, and the random string for each chunk on the disk. Of course, the sensor readings having the state one are written on the disk.

**Example.** Please see Figure 3, where the **last box** shows PHASE 2 execution on four sensor readings to obtain 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 SP to retrieve logs, and (ii) verification of the logs. Due to space restrictions, we skip the key establishment phase. Here, we briefly describe the verification process at the auditor and/or the 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  $\mathcal{C}_x$ ; see Figure 4. Hence, entire sensor data (the data written in first box of Figure 3) of the chunk  $\mathcal{C}_x$ , random strings  $g^a$ ,  $g^b$ , and  $g^c$  (corresponding to the previous and next chunks of  $\mathcal{C}_x$ ; see Figure 4), and proof-of-integrity  $\mathcal{PI}_{\mathcal{C}_x}$  are provided to the auditor. The auditor performs the same operation as in PHASE 2 of §IV-B2. Also, the auditor computes the end-of-chunk string  $S_{eoc}^x = g^a \oplus g^b \oplus g^c$ . At the end, the auditor matches the results of  $h_n^x \oplus S_{eoc}^x$  against the decrypted value of received  $\mathcal{PI}_{\mathcal{C}_x}$ , and if both the values are identical, then it shows that the entire chunk is unchanged.

Note that since SP transfers sensor readings of the chunk  $\mathcal{C}_x$ , random strings ( $g^a$ ,  $g^b$ , and  $g^c$ ) and  $\mathcal{PI}_{\mathcal{C}_x}$  to the user, SP can alter any transmitted data. However, SP cannot alter the signed  $Sign_{PR_E}(h_n^x \oplus S_{eoc}^x)$ , due to unavailability of the private key of the enclave,  $PR_E$ , which was generated and provided by the trusted authority to the enclave. Thus, by following the above-mentioned procedure on the sensor readings of  $\mathcal{C}_x$ , any inconsistency created by SP will be detected by the auditor.

**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 digests computed over device-id and time ( $o_i$ , see the last box in Figure 3), time, sensor state, random strings  $g^a$ ,  $g^b$ , and  $g^c$  (see Figure 4), 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. Also, to verify that no 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  $n$  is the number of sensor readings in  $\mathcal{C}_x$ ) and computes  $hu_{end}^x = h_1^x \oplus h_2^x \oplus \dots \oplus h_n^x$ . Finally, the user computes  $hu_{end}^x \oplus S_{eoc}^x$  and compares against the decrypted value of  $\mathcal{PU}$ . The correctness of this method can be argued in a similar manner to the correctness of the verification at the auditor.

## V. 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 use SHA-256 as the hashing algorithm and 256-bit length random strings in IOT NOTARY. We allow users to verify the data collected over the last 30minutes (min).

### Dataset size.

Although 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 having 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 was 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 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 3) and proof-of-integrity for users  $\mathcal{PU}$  of size 1.05MB. Figure 5 shows 180-days WiFi data size without having sealed logs (red color) and with sealed logs (green color).<sup>3</sup>

**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.

<sup>3</sup>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 30min.

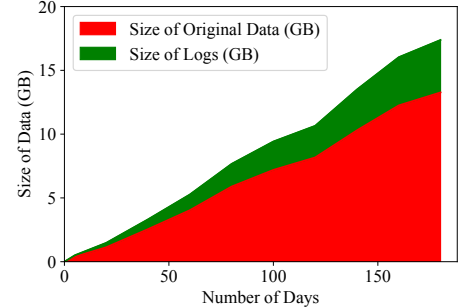


Figure 5: Exp 1: Storage overhead.




Figure 6: Exp 4: Verification time.

Number of Chunks	1	50	100	500	1000	3000
$\approx$ 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 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.

**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 1/10/20-days data; see Figure 6. Note that verifying 1-day data, which is  $\approx 50$  blocks, at resource-constrained users took at most 30s. As the number of blocks increases, the computational time also increases, where the maximum computational time to verify 20-days data was  $< 10$ min. 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. Note that as the number of core and memory increase, it results in parallel processing and absence of disk data read. Thus, the computation time decreases (see user 5 in Figure 6).

## VI. 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 [8] attests the internal state of neighboring devices through key exchange and proofs-of-non-absence. SEDA [9] attests embedded devices and provides the number of devices that pass attestation. Also, DARPA [10] and SANA [11] allow detection of physical attacks by using heartbeat messages and provide aggregate network attestation. In short, existing work cannot verify sensor data against the data-capture rules, except sensors’ internal state. In contrast, IOT NOTARY does not deal with the verification of the internal state of sensors, since in our case, (WiFi access-point) sensors deployed by a trusted entity (*e.g.*, the university IT department). Of course, cyberattacks are possible on sensors to maliciously record data and that can also be detected by IOT NOTARY.

**Attestation using secure hardware.** [12] provided SGX-based attestation method for physical attacks on the sensor. Fiware [13] provides secure key management through key vault running in SGX. However, [12], [13] cannot

verify any sensor data. Also, 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] proposed a privacy-preserving scheme based on zero-knowledge proofs to detect log-exclusion attacks. [15] proposed a Bloom tree that stores proof of logs at an untrusted cloud. vSQL [16] may be used for verifying cleartext query results. However, these techniques cannot detect log deletion and incur significant overheads. For example, vSQL takes more than 4000 seconds to verify a SQL query. In contrast, IOT NOTARY provides complete security to sensor data and realtime data attestation approach.

## VII. CONCLUSION

This paper presented a framework, IOT NOTARY for sensor data attestation through cryptographically enforced log-sealing mechanisms to produce immutable proofs, used for log verification. We improve the naïve end-to-end encryption model, where retroactive verification is not provable. The user-data verification mechanism allows users to revoke services of the concerned IoT space. Thus, 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.

## REFERENCES

- [1] S. Mehrotra *et al.*, “TIPPERS: A privacy cognizant IoT environment,” in *PerCom Workshops*, 2016, pp. 1–6, <http://tippersweb.ics.uci.edu/>.
- [2] A. Rao *et al.*, “Expecting the unexpected: Understanding mismatched privacy expectations online,” in *SOUPS*, 2016, pp. 77–96.
- [3] H. Krawczyk, “Sigma: The ‘SIGn-and-MAC’ approach to authenticated diffie-hellman and its use in the IKE protocols,” in *CRYPTO*, 2003.
- [4] Full version of the paper available at: <https://isg.ics.uci.edu/publications/>.
- [5] V. Costan *et al.*, “Intel SGX explained,” *IACR Cryptology ePrint Archive*, vol. 2016, p. 86, 2016.
- [6] <https://newsroom.intel.com/newsroom/wp-content/uploads/sites/11/2017/09/8th-gen-intel-core-product-brief.pdf>.
- [7] W. Wang *et al.*, “Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX,” in *CCS*, 2017, pp. 2421–2434.
- [8] A. Ibrahim *et al.*, “AID: autonomous attestation of IoT devices,” in *SRDS*, 2018.
- [9] N. Asokan *et al.*, “Seda: Scalable embedded device attestation,” in *CCS*, 2015, pp. 964–975.
- [10] A. Ibrahim *et al.*, “Darpa: Device attestation resilient to physical attacks,” in *WiSec*, 2016, pp. 171–182.
- [11] M. Ambrosin *et al.*, “SANA: secure and scalable aggregate network attestation,” in *CCS*, 2016, pp. 731–742.
- [12] J. Wang *et al.*, “Enabling security-enhanced attestation with Intel SGX for remote terminal and iot,” *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,” in *ISCC*, 2018.
- [14] J. Frankle *et al.*, “Practical accountability of secret processes,” in *USENIX*, 2018, pp. 657–674.
- [15] S. Zawoad *et al.*, “Towards building forensics enabled cloud through secure logging-as-a-service,” *IEEE TDSC*, vol. 13, pp. 148–162, 2016.
- [16] J. Zhang *et al.*, “vSQL: Verifying arbitrary SQL queries over dynamic outsourced databases,” in *IEEE SP*, 2017, pp. 863–880.