



PoTR: Accurate and Efficient Proof of Timely-Retrievability for Storage Systems

CLÁUDIO CORREIA, Computer Science, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Portugal

RITA PRATES, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Portugal

LUÍS FONSECA, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Portugal

MIGUEL CORREIA, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Portugal

LUÍS RODRIGUES, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Portugal

The use of remote storage has become prevalent both by organizations and individuals. By relying on third-party storage, such as cloud or peer-to-peer storage services, availability, fault tolerance, and low access latency can be attained in a cost-efficient manner. Unfortunately, storage providers may misbehave and violate Service-Level Agreements (SLAs). In this paper, we propose a new *Proof of Timely-Retrievability* (PoTR) that aims at assessing whether a provider is able to retrieve data objects with a latency lower than some SLA-specific threshold δ . We have implemented the PoTR and evaluated two distinct configurations of the proof, one tailored to estimate the average latency experience by clients and the other tailored to assess its variance. We leverage Trusted Execution Environments (e.g., Intel SGX) to ensure that the proof is produced by the node being audited and to reduce the communication between the auditor and the audited node. We have experimentally evaluated our prototypes considering a challenging edge computing setting, where storage services are provided by resource-constrained fog nodes, and the distance between the auditor and the audited node can be large. Despite the noise introduced by edge network delays, we show that the auditor is able to effectively detect SLA violations.

CCS Concepts: • **Security and privacy** → *Database and storage security; Distributed systems security.*

Additional Key Words and Phrases: edge storage, SLA violations, auditing, proofs of storage

1 Introduction

Today, there are many scenarios in which end users, or organizations, store data in machines run by third parties, either to ensure durability and availability, or to ensure that customers can access data with low latency. Relevant examples include cloud storage (e.g., Dropbox, iCloud, and Google Drive), peer-to-peer storage (e.g., Filecoin [21], IPFS [12], and Swarm [35]), content distribution networks (e.g., Akamai [3] and Cloudflare [16]), and, more recently, edge storage [3, 36]. In this emerging edge computing environment, the latency for data access will become a crucial constrain and challenge for any data placement algorithm, due to the highly distributed and zero-trust environment [10].

Despite significant advances in storage systems, trust in providers has remained unchanged [24]. Customers require mechanisms to verify that QoS (Quality of Service) is being respected. Relevant QoS aspects include the guarantee that the third party will not discard or corrupt the stored data, that the data is stored on multiple distinct

Authors' Contact Information: Cláudio Correia, Computer Science, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Lisboa, Portugal; e-mail: claudio.correia@tecnico.ulisboa.pt; Rita Prates, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Lisboa, Portugal; e-mail: rita.prates@tecnico.ulisboa.pt; Luís Fonseca, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Lisboa, Portugal; e-mail: luis.h.fonseca@tecnico.ulisboa.pt; Miguel Correia, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Lisboa, Portugal; e-mail: miguel.p.correia@tecnico.ulisboa.pt; Luís Rodrigues, Universidade de Lisboa, Instituto Superior Tecnico, Lisboa, Lisboa, Portugal; e-mail: ler@tecnico.ulisboa.pt.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 1433-299X/2024/8-ART

<https://doi.org/10.1145/3685935>

Systems	Data Locality	Efficient/Cheap Execution	Challenge Delegation Protection	Timer Delay Protection
<i>Single Remote Auditor</i>				
PoTR	✓	✓	✓	✓
PDP (2007) [9]	✗	✓	✗	✓
PoRet (2016) [8]	✗	✓	✗	✓
Benet et al. (2017) [13]	✗	✓	✗	✓
Li et al. (2020) [26]	✗	✓	✗	✓
Filecoin [21]	✗	✓	✗	✓
<i>Multiple Auditors</i>				
Benson et al. (2011) [14]	✓	✗	✗	✓
Gondree et al. (2013) [22]	✓	✗	✗	✓
<i>Local Auditor Relying on TEE Clocks</i>				
Dang et al. (2017) [18]*	✓	✓	✓	✗
EnclavePoSt (2022) [39]	✓	✓	✓	✗
<i>Multiple Auditors and Reliance on TEE Clocks</i>				
ReliableBox (2021) [25]*	✗	✗	✗	✗

* deprecated since Intel excluded trusted time from SGX Linux PSW [5].

Table 1. PoTR properties compared with the related work

machines, in specific geographic locations, and that users are served with some bounded delay. Unfortunately, a misbehaving provider may opt to avoid complying with the agreement if it can gain some benefits and pass unnoticed. For example, the provider may keep the data in fewer locations than agreed with the customer, assuming that it may be impossible for the customer to audit how many replicas are used or where these replicas are placed. This threat has motivated the development of auditing techniques that are capable of extracting storage proofs, that is, evidence that the third party is complying with (or violating) the defined quality of service [13, 14, 18, 22, 26].

In this paper, we present a mechanism that aims at assessing whether a given server node is able to retrieve data objects with a latency lower than some SLA-specific threshold δ . By estimating an upper bound on the data access latency, we can also determine if the data is placed where expected. Namely, when the SLA threshold δ is small, it is possible to verify if data is being stored at the audited node or elsewhere. Our new proof mechanism, named *Proof of Timely-Retrievability* (PoTR), takes advantage of the existence of Trusted Execution Environments (TEEs) [19, 28, 30] (concretely, Intel SGX enclaves[6]) to ensure that the challenge is executed by the node being audited [28], not by some other remote node. By using TEEs, we can also avoid prematurely revealing the data to be accessed during an audit, while keeping the communication between the auditor and the audited node to a single request-reply exchange. PoTR has been carefully designed to mitigate the noise introduced by this single message exchange. Our approach minimizes the network impact on our challenge accuracy, eliminating the need to rely on vulnerable and discontinued TEE clocks [5, 7], in contrast with recent related work, listed in Table 1.

Compared with previous work [18, 25, 39], we take a step forward and evaluate our proof in the highly challenging edge computing environment. When auditing edge services, the auditor may be located far away from the audited node and the communication network may exhibit large delays and jitter that can affect the accuracy of the proof. Edge computing relies on placing resources physically close to end users, such that applications can offer a latency lower than some SLA-specific threshold δ . By setting the SLA threshold δ to a small value that can only be satisfied by the provider if data is kept locally, we use PoTR to distinguish the case



Fig. 1. PoTR auditing scenarios.

where the edge node stores data locally (Figure 1a) from the case where it keeps the data in some remote node or in the cloud (Figure 1b). Although we illustrate PoTR in the context of edge computing, our proof is general and can be applied to other contexts.

We show that our PoTR can be configured and used in different ways to extract two relevant related metrics. The first configuration of PoTR, that we have called Proof of Average Timely-Retrievability (PoATR), allows to estimate accurately the average data access latency (this can be achieved by configuring appropriately the duration of the challenge). The second configuration of PoTR, that we have called Proof of Uniform Timely-Retrievability (PoUTR), allows to assess the variance in the data access latency (this can be achieved by running multiple, independent, shorter challenges). PoUTR is useful to determine if the audited node is able to serve data with uniform latency or if it serves data with latency spikes that cannot be revealed by PoATR alone.

We have experimentally evaluated the accuracy of both configurations, for average latency and variance, with nodes placed in two different university campuses and in the cloud. The results demonstrate that our PoTR can be effectively configured to distinguish between a fog node that respects the latency SLA and one that does not. We found that even in scenarios where a dishonest node stores the file in a nearby fog node (resulting in a latency difference of less than $1.5ms$), our PoTR can accurately pinpoint the misbehaviour. Finally, we have explored various alternatives to reduce the overhead on fog nodes during the execution of the proof, including varying the block size and using SGX switchless calls [37].

2 Background and Related Work

This section provides the essential background for our work.

2.1 Distributed Edge Data Storage

Resorting to third-party storage services provides clients with interesting benefits and challenges, especially in the highly demanding edge computing environment.

2.1.1 Edge Storage Benefits. Many applications strive to reduce their loading times (e.g., social media, entertainment, e-commerce, advertising, and gaming), knowing that even a small difference in milliseconds can significantly impact their revenue [2]. In recent years, a growing number of applications have relied on content delivery networks (CDN) [23], such as Akamai, Amazon, or Verizon, to store static content closer to end users, achieving lower access latencies. In recent years, edge computing has emerged to offer storage and computation even closer to end users, supporting applications for face or object recognition[34], real-time databases [3], and just-in-time video indexing[32], which demand response times below 5-30 milliseconds [31], something that cannot be guaranteed with cloud storage alone or CDNs. The intelligent and resource-aware edge computing environment is composed of a large number of resource-constrained servers known as *fog nodes* or *cloudlets*. Where the self-adapting data placement algorithms will need to consider both the limited resources and the access latency that edge clients suffer, within the distributed and zero-trust environment [10, 20, 27, 33].

2.1.2 Edge Storage Challenges. Fog nodes are managed by many local providers, and due to their limited resources, they cannot store copies of all objects stored in the cloud [17]. Instead, they typically keep copies of items that are required by local applications. Moreover, due to their limited storage capacity, edge storage providers may be tempted to oversell their storage and hide this behaviour by fetching, on-demand, data from the cloud or from other servers, instead of serving them from local (edge), resulting in increased latency for data users. One effective approach to address and detect such misbehaviour is through auditing mechanisms. These mechanisms can enhance the intelligent edge environment by providing real measurements to build trust autonomously between the local entities (e.g., through rating schemes) or to penalize non-compliant providers.

Filecoin storage [21] is a concrete example of an auditing mechanism, where a node that fails to prove its ability to store data as required will no longer receive financial rewards and may be expelled from the network. In cloud or edge storage, providers can be compelled to compensate their customers for violations of the defined contract, known as the SLA. It is essential to audit that the SLA is being delivered in edge applications such as: 1) Web, EdgeKV [4] (similar to CDNs) charge their clients to store data near clients and reduce latency; 2) Augmented reality apps, they need low-latency access to data to be shown to the user; 3) Autonomous vehicles, they need low-latency access to maps, directions, etc.

2.2 Auditing Third-Party Storage Services

When a storage provider is subject to an audit, it must provide a proof that it is applying the storage policies specified in the SLA. This proof is generically called a *proof of storage*[13]. Moreover, since an SLA may cover different aspects of storage implementation, such as the target number of replicas, the location of those replicas, or the latency observed by users when accessing data, it is possible to define different proofs.

2.2.1 Proofs of Storage. The literature is rich in techniques for obtaining proofs of storage, mostly focused on cloud storage, the most relevant ones are *Proofs of Data Possession* (PDP)[9] and *Proofs of Retrievability* (PoRet)[8] that aim to check if the storage provider keeps at least some copies of the stored data; *Proofs of Replication*[13, 26] that assess if the storage provider has n copies of a data item (they may be placed on the same machine); and *Proofs of Geographic Replication*[14, 22] that test if the provider keeps n data copies on different machines, in distinct geographic locations. Each of these proofs is issued as a response to a *challenge*[26], that is sent to the storage provider by one or more auditing entities. Typically, a challenge requires the provider to execute a set of reading operations over a subset of the stored data, and return on-time a value that proves the access of the correct data items (e.g. a cryptographic hash of the audited data).

2.2.2 Structure of a Challenge. A simple way to verify if a storage provider keeps a given data item would be to request that item and then check its integrity. Although this method could, in fact, offer a PoRet, it has several limitations.

First, this approach is inefficient, as it requires the auditor to consume a large amount of bandwidth to obtain the proof. Therefore, to save bandwidth, most challenges require the storage provider to read a subset of the stored data items and compute a cryptographic hash of them. Typically, the response has to be issued within a predefined deadline, determined by the auditor [13, 14, 18, 26]. Furthermore, when sending the challenge, the data items should be revealed interactively to prevent the storage provider from downloading missing items on demand [26]. If the provider cannot guess in advance the data items to be accessed, it has to keep all items within the constraint access time to build a correct and timely proof.

Second, a single request for data items cannot verify some of the service requirements. For example, it cannot assess whether the storage provider keeps just one or several replicas of the data items. A typical solution is to encode each file with a distinct secret key[13]. Unfortunately, this does not prevent a provider from keeping all replicas in the same machine, which may compromise availability. An option is to design the challenge so that it

is impossible to respond on time if all replicas are kept on the same machine, although it can be feasible if the proof is built in parallel[26].

These mechanisms are not enough to guarantee that the proof was generated by the target node. For this purpose, it is possible to leverage a TEE to interactively reveal the challenge and ensure that the operations essential to the proof construction are executed in a given machine[18].

2.3 Trusted Execution Environments

A processor with a TEE has two execution modes: *normal mode*, where the operating system and applications run, and *secure mode*, an isolated environment which ensures the integrity and confidentiality of data and code inside [19, 28, 30], even in the presence of an untrusted operating system.

There is a set of different TEE architectures[30], for example, Intel Software Guard Extensions (SGX), ARM TrustZone, AMD Secure Memory Encryption (SME) and AMD Secure Encrypted Virtualization (SEV). We resort to *Intel SGX* technology, as it is available in Intel systems (PCs and servers), and we use Intel machines in our experimental setup.

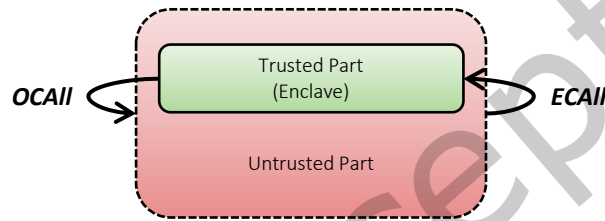


Fig. 2. *Intel SGX* components and execution mode switching calls.

Intel SGX splits the computing environment in two parts: the secure mode known as *enclaves*, and the normal mode known as the untrusted part. Furthermore, since the processor core can only execute one environment at a time, the exchange of environments occurs through hardware calls, more precisely *ECalls* and *OCalls*. In addition to code and data security, *Intel SGX* provides security guarantees regarding machine identification. Figure 2 provides an overview of *Intel SGX* components and the calls executed for the exchange of execution modes. In other words, an external machine that communicates with the enclave can get a guarantee that it is communicating with a real enclave, through the *attestation process*[17, 28].

2.4 Discussion

We now discuss the limitations of the related work in auditing storage systems, summarized in Table 1.

Efficient Execution. Both PDP[9] and PoRet[8] are designed to offer efficient cryptographic mechanisms to verify the integrity of cloud-stored data. Benet et al.[13] present a mechanism capable of verifying if there are a certain number of copies of a file, while Li et al.[26] audit if such copies are stored on different physical machines. These proofs depend on a single remote auditor, requiring low cost for deployment and execution. Filecoin [21] offers blockchain-based cooperative digital storage, which requires expensive cryptographic operations. These proofs cannot estimate data location or access latency, making them unsuitable for testing if a provider stores the required data in a specific node.

Data Locality. Triangulation mechanisms are one way to estimate data location. Gondree et al.[22] and Benson et al. [14] follow this approach by relying on multiple auditors/landmarks for the estimation. Unfortunately, triangulation can be expensive since it requires the intervention of multiple landmarks for each proof execution, and the accuracy of this mechanism depends on the proximity of the landmarks to the audited node. In this paper,

we take a different approach, where we decouple geolocation from data locality. An auditor can first geolocate a node (using techniques as in [22]) and then run our PoTR to check if the files are stored locally at that node. This approach has the advantage that the accuracy of the locality proof no longer depends on the landmarks, and exclusively relies on a single parameter that affects the proof duration, making it practical, particularly in edge environments.

Delegation Attack. Despite the use of triangulation, none of these systems is capable of enforcing the proof to be executed on a given machine; this is a critical obstacle when attempting to audit specific nodes, particularly in an edge storage environment. A storage provider, which controls the infrastructure, may capture the challenge request and delegate the production of the response to the remote storage: such behavior may be difficult to detect. Dang et al. [18], and EnclavePoS [39] resort to a TEE to execute their challenge on the correct machine to prevent this attack but, unfortunately, are still vulnerable to a clock delay attack.

Delay Attack. Using a TEE to trivially implement the auditor alone is not enough to ensure the generation of correct and honest proofs. These systems assume that the SGX clock can be trusted to measure the duration of the challenge at the audited node [18, 25, 39], but recent research has shown that the enclave cannot read an accurate and reliable system clock, due to the following issues [7]: (1) the storage node system clock is vulnerable to manipulations by the untrusted part; (2) the operation to read trusted timers, such as those provided by SGX or, TPM (*Trusted Platform Module*), has a large overhead, penalizing the proof accuracy; and (3) the untrusted part can maliciously delay trusted timer messages, to fetch remote data blocks on demand, and thus escape detection. Additionally, since 2020, Intel has excluded SGX trusted time from Linux PSW, leaving Dang et al. [18] and ReliableBox [25] deprecated. As a result, it is not practical to use the enclave for time measurements. Furthermore, as discussed below, ReliableBox focus on the triangulation task and is unable to verify if the audited node keeps the files locally or remotely.

Geolocation of the Node Being Audited. Some applications may require the geolocation of the machine being audited. The geolocation problem is orthogonal to the problem of timely retrievability and we advocate that these problems must be addressed by different, complementary, mechanisms. Geolocation typically requires the use of multiple auditors (namely, to perform triangulation) while, as we show in this paper, timely retrievability can be achieved using a single auditor. This separation allows performing the proof of geolocation only sporadically (for instance, when the machine boots) and then run a proof of locality more often.

One way to perform triangulation is by leveraging One-Way Delay (OWD) estimation. ReliableBox [25] measures the duration of the challenge both at a set of remote auditors and inside the enclave. It then computes the time difference between these two measurements to estimate the Round-trip time (RTT) between the auditors and the server, performing geolocation via triangulation. Unfortunately, ReliableBox only captures the network delays and is unable to assess accurately the duration of the proof construction at the audited node. This means that the response to the challenge can take an arbitrarily long duration (e.g. when data is stored in a remote node) without affecting the operation of ReliableBox: triangulation would still be accurate but the result of the challenge cannot guarantee data locality or any other properties. To provide both geolocation and timely retrievability, schemes such as ReliableBox need to be combined with schemes such as the one proposed in this paper.

3 System Architecture

We now present our system architecture in the context of edge computing that includes an auditor, a set of fog nodes, and a set of remote storage systems connected to the fog nodes (see Figure 3). Our proof, presented in the next section, is computed by a fog node to prove that it can access data with a latency lower than a given agreed threshold δ .

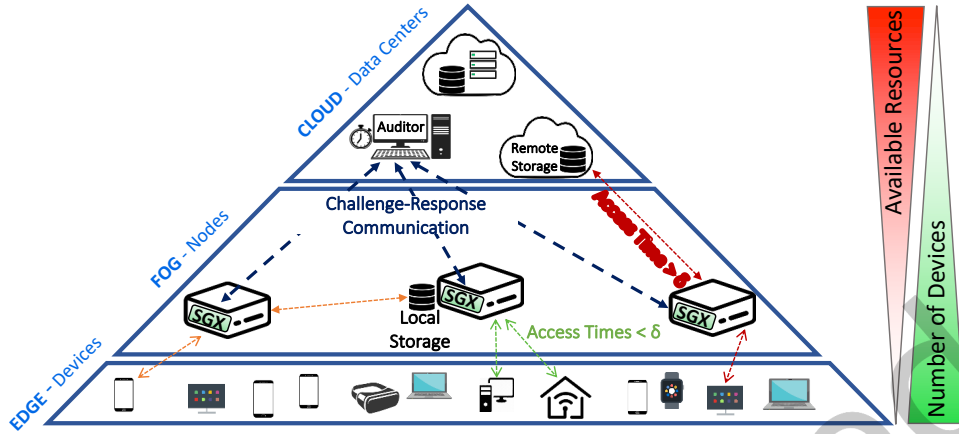


Fig. 3. System architecture.

3.1 Assumptions

The protocol for obtaining a proof is executed between two nodes: the auditor and the audited fog node. A fog node is said to be *correct* if it can retrieve the files assigned to it with a latency lower than some given threshold δ . Any fog node that cannot satisfy this requirement is denoted *faulty*. The value of δ is application specific, but can be small and, in some cases, can only be satisfied if data is stored in a storage device directly connected to the fog node. Providers should offer SLAs that define δ based on their capacity to maintain consistent performance under varying workloads. If a provider cannot meet δ latency under high loads, they should consider offering an SLA with higher latency.

Our proof requires the audited node to read a configurable number N of data blocks. N can be conservatively selected to mask worst-case errors that may result from the variability of access to local storage and the variability in round-trip times. We also show that knowledge of the distribution of network delays and the distribution of storage access delays can be used to optimize the value of N when auditing honest nodes, discussed in Section 6.3.1. Interestingly, these optimizations cannot be exploited by a rational provider to evade auditing.

We assume that each fog node has a processor with *Intel SGX*, as we rely on the guarantees provided by a TEE (our solution may be adapted to use other trusted environments, but the current prototype and evaluation is based on SGX). We assume that the auditor has the guarantee that it communicates with the expected enclave, due to the attestation process [17, 28], and also that the integrity and confidentiality of the data and code inside the enclave are guaranteed [17]. As explained in Section 2.4, the enclave cannot read a reliable system clock [5, 7]. However, the enclave can provide us with the guarantee that the proof is effectively built at the audited node, which is the property we leverage in the solution [28].

We assume that a correct provider keeps all data “locally”, i.e., in the local machine or in a nearby storage that allows it to serve data access requests with a latency smaller than the threshold δ defined by the SLA. A faulty provider may store parts of the data locally and parts (or all) of the data in one or more “remote” sites, i.e., in a remote machine or in a remote storage where it can only be served with a latency higher than δ , violating the SLA.

We assume that the auditor has previous knowledge of the average network delay between the auditor and the audited node but that it cannot measure the exact round-trip time of any particular interaction with the enclave, as the actual network delay is subject to variance. The variance of the network delay affects the configuration of our proof, as it will be explained later in the text. We do not require the exact distribution of the network delays

to be known, as the proof can be tuned to observed network variance in a calibration phase that will also be explained later in the text.

Finally, we assume that it is not feasible for a faulty edge storage provider to reallocate enough objects in less than δ at the beginning of the audit, from a remote storage location. As it will be seen, our audit executes quickly (under 500ms), limiting the number of files that can be downloaded in time, even if the edge node has a high bandwidth link.

3.2 Threat Model

In our system, we assume that the following two entities are trusted: the auditor and the code running within the enclave. When necessary, the auditor can establish a secure channel, via TLS, to the enclave. In our challenge, we leverage the SHA-256 hash function and AES in GCM mode with 128-bit keys, both of which are currently considered secure [11]. The choice of the hash function is due to its security guarantees, such as a high degree of difficulty in finding a collision, and its deterministic and unpredictable output. Hash functions based on SHA are very popular and heavily scrutinized, providing a high level of confidence that they can maintain these properties. Note that other alternatives, such as secure random generators, have similar complexity, as secure random generators typically also resort to secure hash functions or other cryptographic mechanisms.

The edge provider owns the fog node being audited and controls both the local storage device and the network at the audited nodes. This means the provider may attempt to tamper with the local data or change its location. Additionally, the provider can delay messages arriving at or leaving the fog node by adding delays to manipulate the network variance.

3.3 Fog Node Storage Organization

The edge storage provider is responsible for storing the files in the fog layer and ensuring that the files are stored in such a way that they can be retrieved with a latency lower than δ , the SLA defined threshold.

In each fog node with *Intel SGX*, local documents are kept in the untrusted part, as the storage capacity of the enclave is limited[17]. Thus, if the processor is running inside the enclave, there must be an exchange between execution environments to read a data item (from the enclave to the untrusted part). Even if the data item is remote, there is also an exchange of execution environments, as the untrusted part is responsible for communicating with remote machines[17].

The set of auditable files is known to both the auditor and audited node and sorted in a deterministic manner. Thus, both parties can use the index of a file in the sorted list as a mutually agreed short unique identifier for that file. We denote this index by *set index*. However, agreeing on the set of files and supporting file modifications is beyond the scope of this work. For simplicity, files assigned to a fog node, including their content, do not change. The modification of files can be trivially supported using versioning.

3.4 Enclave Geolocation

The goal of PoTR is to check whether a given node stores data locally. When an enclave is attested, it is possible to uniquely identify the specific enclave, but attestation does not reveal its location. Therefore, PoTR alone cannot ensure that data is stored at a given target location; for that, one also needs to geolocate the node. Several geolocation techniques have been proposed in the literature [1, 15, 38] and any of them can be combined with PoTR to ensure data locality *and* geolocation. Although geolocation proofs are orthogonal to our work, we briefly sketch two methods to discover the location of the enclave: 1) *Proximity Attestation* – the auditor physically launches the enclave in the local machine at the correct location and exchanges a certificate with that enclave for later authenticating the same specific machine/enclave; 2) *Triangulation* – any technique of triangulation can be

applied without requiring heavy cryptographic operations (the goal is to geolocate the machine and not data), so the accuracy is not compromised in any way by the storage proof.

4 Proof of Timely-Retrievability

In this work, we introduce the *Proof of Timely-Retrievability* (PoTR) mechanism: a storage proof that aims to assess whether a storage node can access stored data with a latency smaller than some specific threshold δ . The auditor can select the value of the δ parameter based on the specific requirements of a given application, but it is typically small and, in some cases, can only be satisfied if the audited (edge) node keeps the data locally. With the goal of applying the proof in edge computing scenarios, we assume small values of δ , in the order of the time required to read a data block from the local disk. By supporting such strict values of δ , PoTR is capable of distinguishing the case in which the fog node stores the data locally (Figure 1a) from the case it keeps it in some remote node or cloud (Figure 1b).

The PoTR is obtained by the storage node in response to an audit, so our solution requires a machine with auditing capabilities. We do not restrict the placement of these auditing machines, as there may be many geographically distributed fog nodes[29]. Therefore, our proof was designed to be obtainable from an audit machine anywhere on the Internet. This property offers a lot of flexibility regarding the deployment of the auditor and allows a single auditor to perform audits on a large number of fog nodes.

An alternative could be to ask clients to report the latency they observe and use this information to perform the audit. However, this approach raises many privacy challenges, an attacker could infer the client's location given its latency to a fog node. If a PoTR can be extracted independently of the auditor location, we avoid these vulnerabilities.

4.1 Challenges

In the design of our PoTR, we face some obstacles, namely: i) the timing information provided by the audited node cannot be trusted[7], thus the time to produce the proof must be measured by the auditor; ii) the network between the auditor and the audited node is subject to variance that introduces errors when estimating the time the audited node took to produce the proof; iii) storage/fog nodes are heterogeneous [29], and the time they require to perform computations and read data (even if the data is local) is not constant, so the proof should be based on average values from multiple readings; and iv) the audited node may attempt to delegate the generation of the proof to another node that has faster access to the data than the audited node itself, so it is required to ensure the proof is produced by the audited node, and not delegated.

4.2 Design of the Challenge

The challenge requires the untrusted part of the fog node to access a given number of data objects, in a certain sequence, and return, at the end, a value related to these data objects. *The delay the fog node takes to read these data blocks, and to compute the final value, is used by the auditor to estimate the reading delay observed at the audited node and to check if it matches the target threshold δ .*

Each challenge (implicitly) specifies a sequence of files that must be accessed by the audited node, and each file is uniquely identified by a *set index*. For efficiency reasons, the fog node for each file reads a data block of size s_b , instead of the entire file. In practice, the block size should be a multiple of the block size used by the fog node file system.

In each challenge c , the fog node has to read a pseudorandom and unpredictable sequence of N data blocks (each of size s_b), and return a cryptographic hash of the concatenation of all data blocks accessed. The number N of data blocks is a configuration parameter that influences the accuracy and efficiency of the challenge: the higher the value N , the more accurate but less efficient the proof. The approach to configure this parameter is

explained through PoATR in Section 5, and if the user is unable to configure s_b , it is still possible to execute our challenge defining only N , as discussed in Section 6.3.1.

The pseudorandom sequence of the data blocks is determined by a nonce η^c (unique per challenge, and generated by the auditor) and the content of the data blocks. For each challenge, the auditor sends the nonce (encrypted with a symmetric key), the number N of data blocks and the size s_b of a block to the enclave. To ensure that the nonce η^c is not disclosed to the untrusted part, the nonce never leaves the enclave. The untrusted part only has access to a cryptographic hash of the nonce, we denote this hash as $\#index$. With the $\#index$, the untrusted part is able to determine the set index of the first file to be accessed. The set index is determined by applying the modulo function (mod) to the hash with the total size of the set of files, i.e., the set index is the remainder of the division of the hash by the number of files.

Since the fog node has to read a data block of size s_b , and not the entire file, the auditor sends to the enclave a second nonce η_b^c that will determine a data block inside the file, with the computed set index. The enclave, in turn, computes the cryptographic hash of this second nonce and forwards the result to the untrusted part. Both nonces are hashed by the enclave, to ensure that the untrusted part does not know them in plain text. Otherwise, the untrusted part could maliciously predetermine the N data blocks. Now, with the hashed η_b^c , and the total size of the file to be accessed, the untrusted part can determine the data block, with size s_b , inside the file to be retrieved, by applying the modulo function.

When the untrusted part receives the first $\#index$ from the enclave, it proceeds to determine and read the corresponding initial data block. It then calculates a cryptographic hash of the data block's content, combined with the $\#index$, resulting in $result_hash$. Subsequently, the untrusted part returns the value of $result_hash$ to the enclave. The enclave, in turn, computes a new $\#index$ by hashing the response $result_hash$ together with the nonce η^c . This newly generated $\#index$ is then forwarded to the untrusted part, which uses it to identify the next file.

For each subsequent hash, the untrusted part repeats the execution of the modulo function to obtain the next set index and the index of the data block within the identified file. It reads the corresponding data block and calculates its cryptographic hash together with the file $\#index$, yielding $result_hash$ that is returned to the enclave. The enclave computes two new hashes (block and file index), and this process continues until all N data blocks are read. Note that each new $\#index$ is computed using the hash of the previously retrieved data blocks. Finally, the enclave computes a final hash using η^c and sends the result back to the auditor as the conclusive proof. The auditor, upon receiving the final hash, repeats all the aforementioned computations to construct a verifiable version and verify the correctness of the proof. If both computations match, the issued proof is deemed correct. In Figure 4, we present in detail all the steps of our communication protocol in the execution of our PoTR challenge. We separate the steps on the auditor's side between the untrusted part and the enclave, where the sensitive parts of our challenge, such as the nonce, are always securely stored inside the enclave.

Security Guarantees: Each challenge requires two unique nonces that will determine in a sequential and pseudo-randomly fashion which N data blocks the untrusted part must read. This interaction with the enclave ensures that the fog node is unable to retrieve the N data blocks in parallel from neighboring nodes. Instead, it must retrieve one block at a time to determine the next block it needs. Additionally, the dependency between data blocks, since the next block index depends on the content of the previous one, ensures that the fog node cannot reuse outputs from previous challenges, thus maintaining challenge freshness. The constant exchange between execution environments (enclave to untrusted part and vice versa) guarantees that the fog node cannot rely on a remote machine for the entire proof computation. The proof is computed solely on the expected machine [28]. It is important to note that through the verification of proof correctness, the auditor can assess the integrity of the stored documents. Any modification to a file will render the proof invalid.

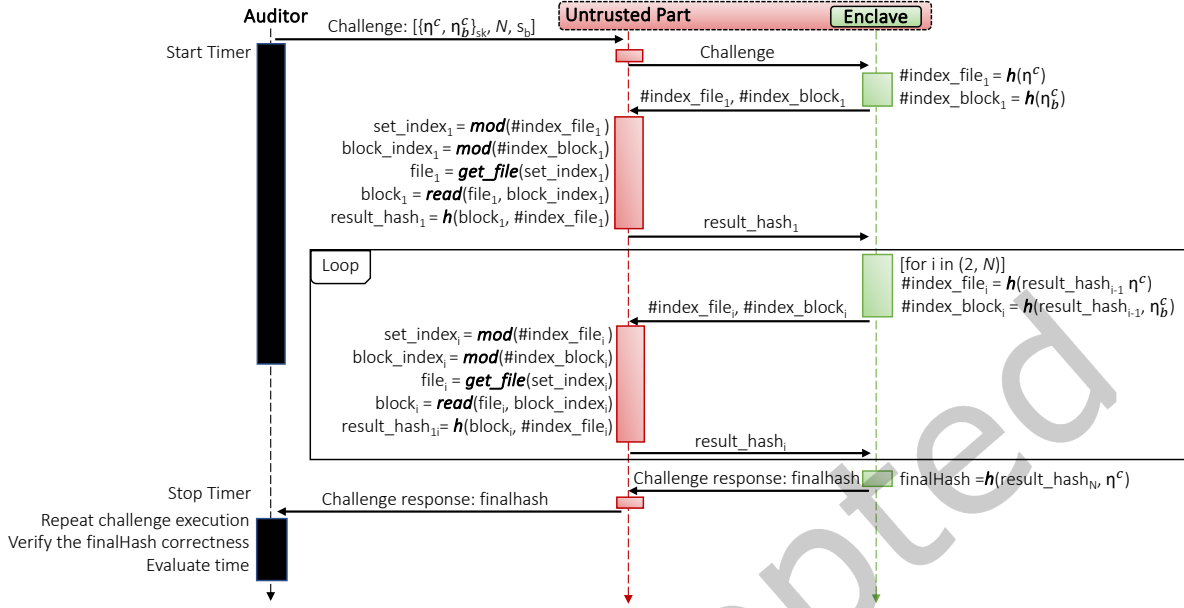


Fig. 4. PoTR challenge execution. set_index_i is the index of a file, and $block_index_i$ is the index of a data block with size s_b inside $file_i$. $block_i$ is the read data block. sk is a secret key between the auditor and the enclave. $\#index$ is a hash.

4.3 Reading Delay δ at the Audited Node

In addition to checking *proof correctness*, the auditor has to check the *proof timeliness*. When the challenge is sent to the enclave of the audited node, the auditor starts a timer that is stopped when the proof is received. A proof is valid if it is correct and the reading delay estimation, for a single data block, is acceptable to the auditor.

Having T_i as the time elapsed between the auditor sending the challenge i and receiving the response from the fog node, T_i can be decomposed into different factors, namely:

$$T_i = rtt_i + \alpha_i^1 + \dots + \alpha_i^N + \delta_i^1 + \dots + \delta_i^N + \epsilon \quad (1)$$

where rtt_i is the network round trip time for challenge-response messages, N the number of data blocks to be accessed, α_i^j the delay observed at the fog node to compute the cryptographic hash of a given data block j and compute the index of the next block, δ_i^j the delay observed to read a data block j , and a small constant delay ϵ to account for the cold start of the challenge. Note that for sufficiently large values of N , we will have the following:

$$\frac{\alpha_i^1 + \dots + \alpha_i^N}{N} = \bar{\alpha}_i \approx \bar{\alpha} \quad \frac{\delta_i^1 + \dots + \delta_i^N}{N} = \bar{\delta}_i \approx \bar{\delta}, \quad \frac{\epsilon}{N} \approx 0 \quad (2)$$

However, the auditor is unable to measure accurate values for rtt_i , $\bar{\alpha}_i$ and $\bar{\delta}_i$, as they are different and variable in each challenge i and the auditor is only able to measure the total delay T_i . Therefore, to estimate $\bar{\delta}$, i.e., the actual mean delay a fog node takes to read a data object, in our work we resort to mean values:

$$T_i = \overline{rtt} + N\bar{\alpha} + N\bar{\delta} \quad (3)$$

$$\bar{\delta}_{\text{estimate}} = \frac{T_i - \overline{rtt} - N\bar{\alpha}}{N} \quad (4)$$

Therefore, the estimate error for a given challenge i will depend on how far the observed values are from the mean values. As noted in Section 3.1, we assume that \overline{rtt} is known *a priori* by the auditor, based on the expected location of the audited node. Experimentally, we verified that the α values are subject to a negligibly small variance, whereby we assumed $\bar{\alpha}_i = \bar{\alpha}$, i.e., we ignore the error introduced by α sampling. Therefore, the error of the estimate $\bar{\delta}_{\text{estimate}}$ depends on two factors: i) the reading error ϵ^δ when estimating $\bar{\delta}$ (which depends on the sample size, i.e., the number of data blocks accessed) and ii) the variance between the observed network round-trip time (rtt_i) and the expected mean value (\overline{rtt}), divided by the number N of data blocks, as Equation 5 describes:

$$\epsilon^{\text{rtt}} = \frac{|rtt_i - \overline{rtt}|}{N} = \frac{\Delta^{\text{rtt}}}{N} \quad (5)$$

4.4 Variance of the Network Delay

Our proof is capable of finding the correct $\bar{\delta}_{\text{estimate}}$, as long as \overline{rtt} is known, even if delay distributions are unknown, by choosing worst-case default values that can be used conservatively (see Section 6.3.1). However, during a calibration stage, it is possible to leverage information regarding the observed network delay variance to select a more efficient N value, ensuring that the sum of both errors remains below a target value of $\epsilon_{\text{max}}^{\bar{\delta}}$. We demonstrate this in the concrete edge node scenario in Section 6.3.2. The choice of an efficient N , based on measuring network behavior, is safe and cannot be exploited by a malicious provider. Note that, during calibration, a faulty node could attempt to add delays to the messages used in the calibration to influence the variance observed by the auditor. We argue that artificially increasing or decreasing the variance observed by the auditor (with regard to the real network variance) brings no benefit to a faulty node:

- *Increasing the observed network variance:* If the attacker tries to inject delays to increase the network variance, this would result in much larger values for N (from Equation 5). As our evaluation shows, increasing N provides more accuracy in estimating $\bar{\delta}$ at the cost of a longer challenge. Therefore, this option offers no advantage to the attacker since the challenge would be tuned to tolerate the added variance at the cost of forcing the audited node to run more expensive challenges.
- *Reducing the observed network variance:* If the attacker injects well-precised network delays that negatively correlate with the real network variance, they could reduce the measured variance during calibration. This would result in a smaller N variable, meaning that challenges will be shorter and more likely to trigger false positives in face of the real network variance. This provides no advantage to the attacker, as it would increase the chances to be flagged, even when behaving correctly (i.e., while keeping all file locally).

5 Configuring the Proof of Timely-Retrievability

As previously mentioned, it is possible to deploy two different configurations of the PoTR, one aimed at assessing the average data access latency and the other aimed at assessing the data access latency variance, that we have named PoATR and PoUTR respectively. The PoATR is based on executing a single challenge and by setting the value of the N variable large enough to ensure that the challenge returns an accurate estimation of $\bar{\delta}$. The PoUTR is based on executing multiple challenges, each with a smaller value of the N variable, such that it can better capture the variance in accessing different sets of data objects. Both the PoATR and the PoUTR are instances of the general PoTR formulation described above, but they feature distinct configurations to fulfill specific objectives. The following section offers details on each configuration.



(a) False negative given by PoATR when just a small percent- age of data is maliciously in a remote node. (b) PoUTR detects access to the remote storage even for a small percentage of data.

Fig. 5. PoUTR auditing scenarios.

5.1 Proof of Average Timely-Retrievability

The PoATR challenge is strategically configured to control the error in estimating $\bar{\delta}$, thereby enhancing the accuracy of the average data access latency calculation. The challenge can be configured by providing two parameters: the maximum error $\epsilon_{\max}^{\bar{\delta}}$ that he is willing to tolerate, when estimating $\bar{\delta}$, and the reliability of the challenge ϕ , a parameter specified by the auditor that captures the probability of estimating $\bar{\delta}$ with an error lower than $\epsilon_{\max}^{\bar{\delta}}$ (in practice we use the 99.99 percentile). As noted before, the error of the estimate $\epsilon_{\max}^{\bar{\delta}}$ results from the variance of local reads experienced by the audited node ($\epsilon_{\max}^{\delta_i}$) and by the variance of the network delays ($\epsilon_{\max}^{\text{rtt}}$), i.e., $\epsilon_{\max}^{\bar{\delta}} = \epsilon_{\max}^{\delta_i} + \epsilon_{\max}^{\text{rtt}}$. For both sources, the error diminishes with the number of samples N , thus, N should be chosen so that $\epsilon_{\max}^{\text{rtt}} + \epsilon_{\max}^{\delta_i} < \epsilon_{\max}^{\bar{\delta}}$.

Let $\epsilon^{\delta_i^j}$ be the difference between the average access latency $\bar{\delta}$ and the latency observed when reading block j during the execution of challenge i . The error caused by the variance of storage access in a given challenge i is given by:

$$\epsilon^{\delta_i} = \frac{\sum_{j=1}^N \epsilon^{\delta_i^j}}{N}$$

Given ϕ , the expected value of $\epsilon_{\max}^{\bar{\delta}}$ can be derived from the distribution of the access delays. Similarly, with the Inverse Cumulative Distribution Function (ICDF) of the network distribution, we can also calculate the expected maximum difference between the observed value \overline{rtt}_i and the expected mean \overline{rtt} that matches ϕ . Section 6.3 evaluates experimentally the PoATR configuration.

5.2 Proof of Uniform Timely-Retrievability

A faulty storage node may opt to keep part of the data locally and part of the data remotely, as depicted in Figure 5. If the fraction of data that is stored remotely is small, this can pass undetected by the PoATR because the resulting average access delay can still be below the designated threshold δ . However, in this scenario, some clients will experience acceptable delays but others, that access data stored remotely, will experience delays that violate the SLA. The Proof of Uniform Timely-Retrievability is a configuration of the PoTR aimed at estimating if clients can assess all data with a uniform delay.

The PoUTR is based on analyzing the variance across several smaller challenges to identify if there's a portion of data stored remotely. This configuration employs the standard variance calculation method, utilizing multiple δ measurements to provide a more comprehensive analysis of the audited node behaviour. Namely, PoUTR computes σ^2

$$\sigma^2 = \frac{\sum_{i=1}^k (\delta_i - \bar{\delta})^2}{k - 1}; k > 1 \quad (6)$$

which is the sample variance derived from k samples, each executing our challenge to measure δ_i .

In Section 6.4.1, we show that the N required for each δ_i challenge can be small while still achieving sufficient accuracy to detect misbehaviour in the challenging scenarios where only a small fraction of the data is stored remotely. Note that reducing the value for N breaks the assumptions from Equation 2, impacting PoUTR ability to accurately estimate the latency $\bar{\delta}$. Therefore, PoATR should always be used for estimating averages, and PoUTR only for estimating variance.

6 Evaluation

We now present our evaluation of both PoTR configurations, PoATR and PoUTR. We begin by describing our experimental setup and the different benchmark scenarios. Subsequently, we experimentally evaluate and configure each configuration separately. For the PoATR configuration, we demonstrate that our proof is effective even without prior knowledge of the network distribution. We then discuss how to optimally configure the challenge, focusing on the block size and the number of N samples. Upon determining the appropriate configuration for PoATR, we assess its performance across different scenarios. Our experimental results show that PoATR can accurately distinguish between a misbehaving server and a compliant one in the context of edge computing. Additionally, we explore various alternatives to mitigate the impact of PoATR on the audited node. For the PoUTR, we begin by establishing the necessity for this second configuration of the proof by introducing scenarios where only parts of the data are placed remotely and showing that measuring the average access latency is not enough to detect this type of misbehaviour. We then discuss how to configure the PoUTR parameters, selecting a smaller number of N samples and evaluate the proof's performance in relation to the amount of samples taken by it. The results show that the variance of a challenge with only 40 samples can accurately identify nodes with multiple sources, even for remote percentages as low as 10%.

6.1 Experimental Setup

We resorted to two different types of machines: 1) Intel NUCs to represent fog nodes, and 2) virtual machines in the Windows Azure cloud to represent remote entities. For NUCs, we have resorted to the Intel NUC10i7FNB, it has an Intel i7-10710U CPU that supports Intel SGX, 16GB RAM, 240GB SSD M.2, and Ubuntu 20.04 LTS. We run the Intel SGX SDK Linux 2.13 Release and OpenSSL 1.1.1k. An Intel NUC is an example of what a fog node might be, as it possesses modest computational resources but is relatively inexpensive for large-scale deployments. For the cloud deployment in Azure, we resort to Standard D2s v3 with two vcpus, 8 GB RAM, 256 GB SSD, and Ubuntu 22.04 LTS in west Europe (Netherlands) the closest data center to our laboratory; our laboratory is located in a European capital. We leverage these two types of deployment to create different scenarios and evaluate our storage proof; we explain our scenarios in more detail in the next section.

For the stored data that our proof is responsible for auditing, we generated 150GB with random data divided into files of 1GB, our proof leverages a hash function to choose the next block to be accessed; this offers uniform distribution access over the 150GB of data. In our code implementation, we used the SHA-256 function as a cryptographic hash function and AES in GCM mode with 128 bit keys to cipher the nonces, both currently considered secure [11].

To simulate real world usage, all nodes run a local client that performs random read accesses on its stored data. We experimentally measure the required load for the read rate to plateau, and run the tests with 50% of such load on the background. On some occasions, results of tests performed without load will be presented for comparison. Our code is publicly available at: <https://github.com/claudio-correia/potr>




Fig. 6. Overview of PoATR δ estimative without network knowledge, picking a large N value of 3000, 64KB block size.

6.2 Evaluation Scenarios

In our experiments, we leverage three different entities to evaluate our PoTR storage proof: an auditor; the audited node, which represents a fog node that stores data; and the remaining one is the remote storage node, which is used by the audited node to access data in configurations where it does not store the files locally. In the first part of the evaluation (Section 6.3) we consider the case where the audited node stores all data at a single location (that can be either local or remote). In this case we focus on the performance of PoATR in accurately extracting the observed average access delay. In the second part of the evaluation (Section 6.4) we consider the case where the audited node stores part of the data locally and part of the data remotely. In this case we focus on the performance of PoUTR in detecting this scenario using the variance in the access delay.

We deployed four different scenarios to evaluate PoTR; the difference between the scenarios is the location of the remote storage node, since the closer the remote storage is to the audited node, the more difficult it is to detect malicious behavior. The first scenario is represented in Figure 1a: the fog node has honest behavior and stores all data locally. Figure 1b presents the malicious behavior that the edge provider can adopt by resorting to remote storage. Figure 1c represents the three remaining scenarios in which we vary the location of the remote storage from: a) a virtual machine on the Azure cloud; b) a NUC on a different campus of our university; c) a NUC located in our laboratory and connected through a switch to the audited node, with a mean network delay of $0.1ms$. Regardless of the scenario, both the auditor and the fog node/audited node do not change their location, the auditor is running in the Azure cloud, and the audited node in an NUC at our laboratory. By deploying different machines at different geographic locations, we capture the different variations and delays of wide-area links.

6.3 PoATR

We now present the experimental configuration of PoATR, where we configure PoTR to accurately estimate δ while keeping N small to reduce the cost of the audit procedure.

6.3.1 PoATR Without Network Knowledge. As stated in Section 3.1, the efficiency of PoTR can be improved given knowledge on the distribution of network delays and access delays. When executing PoATR without this knowledge, a large default value for N should be used, enough to compensate worst-case network variance; we show that a default value $N = 3000$ can be safely used in most scenarios. We have selected this default value based on latency measurements from our laboratory to the farthest Azure datacenter¹, the Australia Central 2. We have observed an average RTT of $\approx 286ms$ and a maximum latency of $\approx 1143ms$. When considering a network variance with this order of magnitude in Equation 5, by setting $N = 3000$ we can reduce ϵ^{rtt} to $\approx 0.28ms$. In this scenario, the variance in storage access delay is negligible compared to the network variance.

Figure 6 shows the results of running PoATR using the default value of $N = 3000$ in different scenarios. We can observe that it is possible to set the SLA threshold to a small value such as $\delta = 2.5ms$, to ensure that the audited

¹We leveraged this website: <https://cloudpingtest.com/azure>

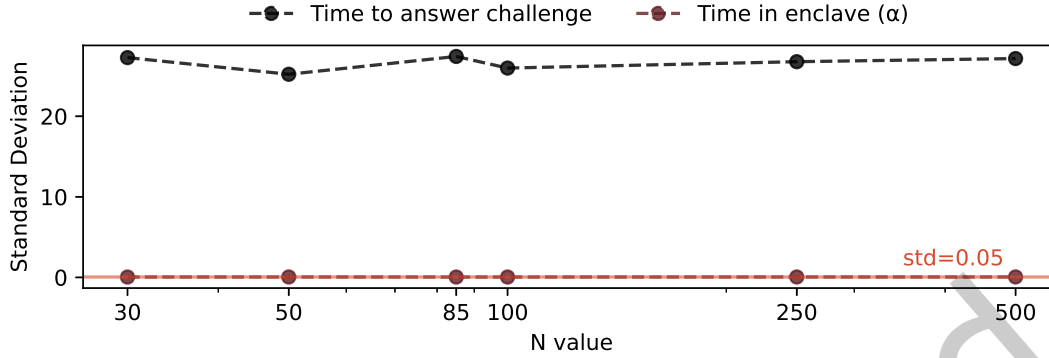

Fig. 7. Standard deviation of the α samples.

Fig. 8. PoATR experimental error for different scenarios and block size, EOE is Experimentally Observed Error.

node has the data stored locally (the local delay to read a data block is $\approx 1.7ms$), and that our $\bar{\delta}_{estimate}$ achieves high accuracy in both scenarios.

However, with $N = 3000$ our challenge takes approximately 5.2s to execute. When the challenge is executed using more stable networks, it is possible to fine-tune the value of N to achieve similar accuracy with a lower cost for the audited node. In the next section, we demonstrate results for various values of N . For instance, with $N = 500$, we achieve accurate results with a challenge executed in just 0.9s.

6.3.2 PoATR Configuration. We now derive, based on experimental results, the best configuration for our storage proof, by setting the block size and the N value for PoATR. Note that in a situation where the user is unable to configure the block size, they can simply use a high value for N , as demonstrated in the previous Section 6.3.1.

- *Negligible impact from α .* In Figure 7, we demonstrate experimentally that the α variance has a negligible effect on our challenge, which is why we ignore the error introduced by α sampling in Equation 5. Figure 7 shows that the standard deviation of the α samples is just 0.05, even for $N = 500$. This is minimal compared to the standard deviation of more than 20 for our challenge.

- *Selecting the appropriate Block Size.* To find the appropriate value for the size of the block to use in our challenge, we evaluated the performance of the challenge with different block sizes and different N values. The results are presented in Figure 8. We executed our challenge to estimate the δ value and calculated the error relative to the real reading delay in the fog node, we present the results for the honest fog node scenario in

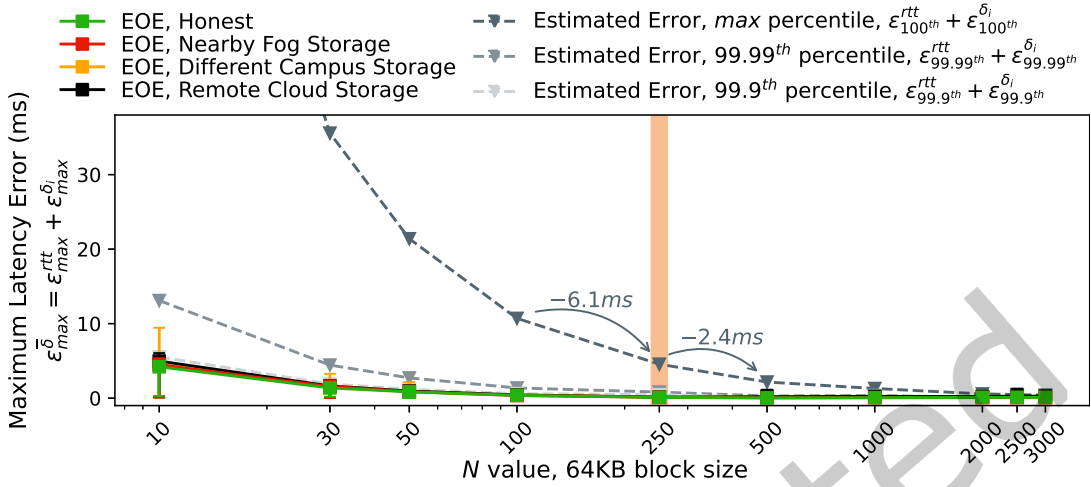


Fig. 9. Theoretical and experimental error for the PoATR challenge (EOE means Experimentally Observed Error).

Figure 8(a) and for the nearby fog storage scenario in Figure 8(b). We observe that for small values of N the results obtained exhibit a large variance and only by increasing N do we obtain more accurate results. In both scenarios, we observe that the size of the block has a relatively smaller impact on the accuracy of the result than the value of N . Still, it is possible to observe that, when using larger block sizes one can obtain more accurate results. This is clearly noticeable in the scenario where files are stored in a nearby fog node, where the configuration using blocks of 4KB takes longer to converge than configurations with larger block sizes. Furthermore, in Figure 8(b), it is possible to observe that when the N value goes above 100, blocks of 64KB offers a slightly smaller error than blocks of 16KB. For these reasons, we opted to set our challenge with blocks of 64KB.

- *Finding the appropriate value for N .* As described in Section 5, the value N can be determined by providing $\bar{\epsilon}_{\max}^{\delta}$, which depends both on the network error and the reading error, this network knowledge allows us to find an optimal value for N . Figure 9 presents the estimated error for our challenge under different percentile values, of the reliability ϕ , for each error ϵ^{rtt} and ϵ^{δ_i} (these two were measured experimentally). Note that $\bar{\epsilon}_{\max}^{\delta}$ decreases as we increase the value N , this results from the impact that N has on ϵ^{rtt} reducing its effect on the challenge, on the other hand, ϵ^{δ_i} has a negligible impact on $\bar{\epsilon}_{\max}^{\delta}$, ϵ^{δ_i} is mostly below 1ms reaching at most 5ms. Figure 9 also presents the measured error of our challenge in real experiments in the four different scenarios presented earlier. We observe that the error increases when we also increase the distance between the data and the audited node.

Both ϕ values of 100th and 99.99th percentiles are able to predominantly capture the real observed error of our challenge. We use the estimated error to choose a desirable value for N ; in this case, we chose $N = 250$ since this is where we observe the last most significant reduction in the estimated error of $-6.1ms$, versus the increase in the value of N , of 150. As expected, increasing the value of N also improves the accuracy of our challenge, for $N = 10$ we obtain an error of $\approx 5ms$, which can be undesirable for edge environments, while for $N = 250$ the observed error is only $\approx 0.1ms$, such a small error can provide a $\bar{\delta}_{\text{estimate}}$ with high reliability.

6.3.3 PoATR Accuracy. Using the configuration parameters derived from the previous analysis, i.e., using a block size of 64KB and by setting the number of samples N to 250, we have executed our challenge in multiple scenarios and captured the distribution of the data access latency estimated by our challenge. We compare the estimated

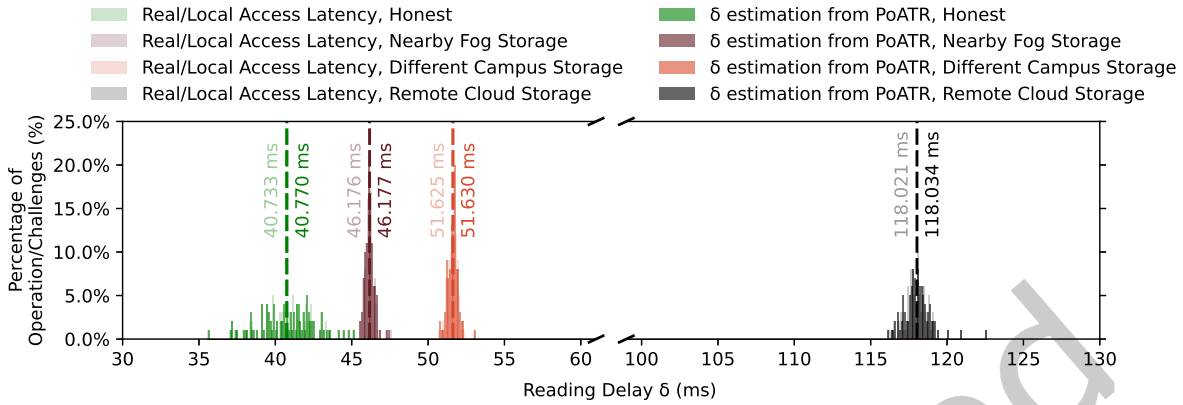


Fig. 10. Distribution and correspondent averages for reading delay δ , in milliseconds, for different scenarios where PoATR is configured with $N = 250$ and 64KB block size.




Fig. 11. Distribution and correspondent averages for reading delay δ , in milliseconds, for different scenarios where PoATR is configured with $N = 250$ and 64KB block size, without background load on audited node.

values with the real values, observed at the audited node itself. The results are depicted in Figure 10. Each point in the figure represents either the real latency to access a data block or the $\bar{\delta}_{\text{estimate}}$ to access a block resulting from the execution of our challenge.

It is possible to observe that our challenge estimate closely approximates the real value with a small error, as discussed in Section 6.3.2. The error between the real value and $\bar{\delta}_{\text{estimate}}$ increases as the audited data move further from the audited node. This happens because the farther away the audited node is, the more likely the network delay exhibits a larger variance. Nonetheless, our challenge can still distinguish whether the data is local to the fog node or at a remote site. The accuracy of the PoATR is sufficient to differentiate between scenarios where files are stored locally or at a nearby fog node. Even with a latency difference of less than 1.5ms, our PoATR accurately identifies the configuration used by the audited node.




Fig. 12. Challenge impact on concurrent reads.

Looking at the results in Figure 10, our PoATR enables the establishment of a threshold to accurately differentiate between a correct node (storing all files locally) and a faulty node that stores files elsewhere (even if they are kept in nearby nodes). This threshold is determined by setting a strict value for δ in the SLA, such as below the first reading delays measured at the nearby fog storage node. This threshold, located at $45.2ms$, corresponds to 3% false positives and 0% false negatives (we define as positive the detection of malicious behavior of the audited node).

As a point of comparison, the experiment was repeated without background load generation. The results can be seen in Figure 11, severely reducing the variance of the honest audited node. With this configuration, the same threshold as before becomes able to perfectly flag all malicious nodes, with 0% false negatives or positives. In real world deployments, the performance would be in between the two.

6.3.4 Overhead Imposed by a PoATR Challenge. We now assess the overhead imposed on the audited node while it responds to a challenge. For this, we run a local client that performs read/write access to the data stored on the edge node, and then we measure the loss in throughput of that client during challenge execution. The results are shown in Figure 12.

The throughput of the concurrent client decreases by approximately 57% during the processing of the challenge by the audited node. This is expected as our resource-constrained edge nodes need to access multiple files and compute their digest to respond to the challenge. Despite this non-negligible overhead, we argue that its impact on the overall operation of an edge node is not significant in practice. The challenge typically takes less than $500ms$ to complete, and auditing occurs sporadically in most scenarios (e.g., once per day). Nonetheless, in the next section, we discuss and evaluate strategies to further reduce this overhead.

6.3.5 Strategies to Reduce PoATR Overhead. We considered three complementary avenues to attempt to reduce the overhead of running a challenge in an audited node. The first was to modify the implementation, taking advantage of switchless calls [37] that reduce the ECall/OCall overhead. The second was to use blocks size smaller than 64K. As discussed earlier, smaller sizes can reduce the accuracy of the test but can have an impact on the overhead. Finally, one can simply reduce the number of samples N . By reducing the number of samples, one would reduce the challenge duration and therefore, the period during which clients would be affected by the concurrent execution of a challenge. However, as we show in this section, using switchless calls or reducing the size of the blocks has a minor impact on the overhead.

Figure 13 shows how the use of switchless calls, different block sizes, and different values of N , affect the drop in throughput observed by clients during the execution of the challenge and also the duration of the challenge.

It is clear from Figure 13a that the drop in throughput during the execution of the challenge is roughly the same regardless of the block size used or the use of switchless calls. Switchless calls offered a small performance improvement, just a latency reduction of $\approx 1.6\%$ in the best case. When considering block size, the largest




Fig. 13. PoATR duration and throughput impact on concurrent reads, Blk is block, and Sw is switchless.




Fig. 14. False positive and negative rates of PoATR under different N values for honest and nearby fog node scenarios.

difference is observable from 256KB block size, with a throughput reduction of -55% , versus 4KB block size alternative, with a reduction of -38% , relative to the parallel client. This is not a large improvement considering that the 256KB block size is 64 times the 4KB block size. This is because, during the execution of our challenge, our disk access is continuous, independently of the block size, and therefore the impact on other clients will be similar; however, our challenge duration may vary. Additionally, as discussed in Section 6.3.2, having a very small block size can still affect our storage proof error.

These results suggest that the only effective way to reduce the overhead of the challenge is to reduce its duration. Figure 13b shows that, although block size has some impact on the duration of challenge executions, its impact is relatively minor compared to the effect of reducing the number of samples N . This indicates that the most effective strategy to reduce the overhead of the challenge is to reduce the number of samples, as this reduces the duration of the period in which clients are affected. Of course, as discussed before, reducing the value of N will degrade the accuracy of the test, increasing the chances of producing false positives and false negatives.

Figure 14 shows the effect of N on both false positive and false negative rates when using the threshold $\delta = 45.2ms$, as discussed in Section 6.3.3. In this analysis, we consider the scenario where a faulty node stores data at the nearby fog storage, as this is the scenario where it is most challenging to accurately detect a misbehavior, given the small difference between the latency offered by correct and faulty nodes. It is interesting to observe

that our PoATR offers a false negative rate close to 0% for most values of N . Therefore, even where the values of N are smaller than 250 (the required to obtain the best accuracy), our PoATR will detect a misbehaving client. Unfortunately, the true positive rate can increase significantly when using low values of N , risking identifying a correct node as faulty. More precisely, the false positive rate increases slowly as we decrease the number of samples from 250 to 100 (at this point, our PoATR offers a false positive rate of approximately 1%), but then increases sharply. For values of N smaller than 100 the challenge no longer provides acceptable values for the false positive rate.

In summary, by tolerating a small fraction of false positives, the auditor can reduce the challenge overhead by decreasing the number of samples from 250 to 100. This roughly halves the period during which clients are affected by our challenge.

6.3.6 Combining Multiple PoATR Configurations. As seen in the previous section, using a reduced number of samples, 100 instead of 250, our challenge imposes significantly less overhead while still obtaining a negligible false negative rate and a very small positive rate ($\approx 1\%$). Faced with these observations, we conclude that in some deployments the auditor should combine different configurations of the PoATR: one that is cheaper and works 99% of the time; another that is more expensive but provides extreme accurate results.

For example, the auditor might use $N = 100$ to challenge the audited node. If the node was flagged as malicious (a positive), the auditor could then launch a second challenge with $N = 250$, to filter out false positives. Note that in such a case, the combined challenge would read on average 102.5 ($0.01 \times 100 + 0.01 \times 250 + 0.99 \times 100$) blocks instead of the 250 in the normal auditor, a reduction of 41% on average for the number of blocks required to execute our challenge. Notice that, as hinted before, for $N = 85$ this strategy would no longer be worth it since it would require reading 257.5 ($0.69 \times 85 + 0.69 \times 250 + 0.31 \times 85$) blocks on average, which is larger than the 250 required by the normal auditor solution.

6.4 PoUTR

We now present the experimental configuration of PoUTR, where we configure PoTR to estimate the variance of the access delay δ in order to detect scenarios where the audited node stores parts of the data locally and parts of the data remotely.

To experimentally evaluate the PoUTR configuration we consider a scenario where a faulty providers divides its data into two slices, where one slice with $p\%$ of the data is stored remotely and the other slice with $(1 - p)\%$ of the data is stored locally. The data in each slice is selected at random (note that the randomized nature of the PoTR challenge voids any attempt to optimize placement to evade auditing). In the experiments reported in this section, the local slice is stored in the fog node (an Intel NUC deployed in our laboratory) and both the remote slice storage server and the auditor are hosted in a remote cloud, namely in the Windows Azure cloud virtual machines described in Section 6.1.

The first subplot of Figure 15 portrays the distribution of PoATR estimations of δ for nodes with all storage in the same location. The other two subplots showcase mixed configurations where small percentages of the data are stored remotely in the cloud. Naturally, when a fraction of the data is stored remotely, the average access time increases. However, the average δ , albeit larger, can still be below the threshold and, therefore, a misbehaved storage node cannot be detected using the PoATR alone.

This is illustrated by Table 2. When the SLA is fixed to $\delta = 44.8ms$, PoATR is able to flag all faulty storage nodes storing 25% of its data remotely. However, for a smaller share of 10%, PoATR alone can no longer fully detect the misbehaved nodes, since the average access delay is below the threshold for 10% of the PoATR executions. This motivates the need to use PoUTR.

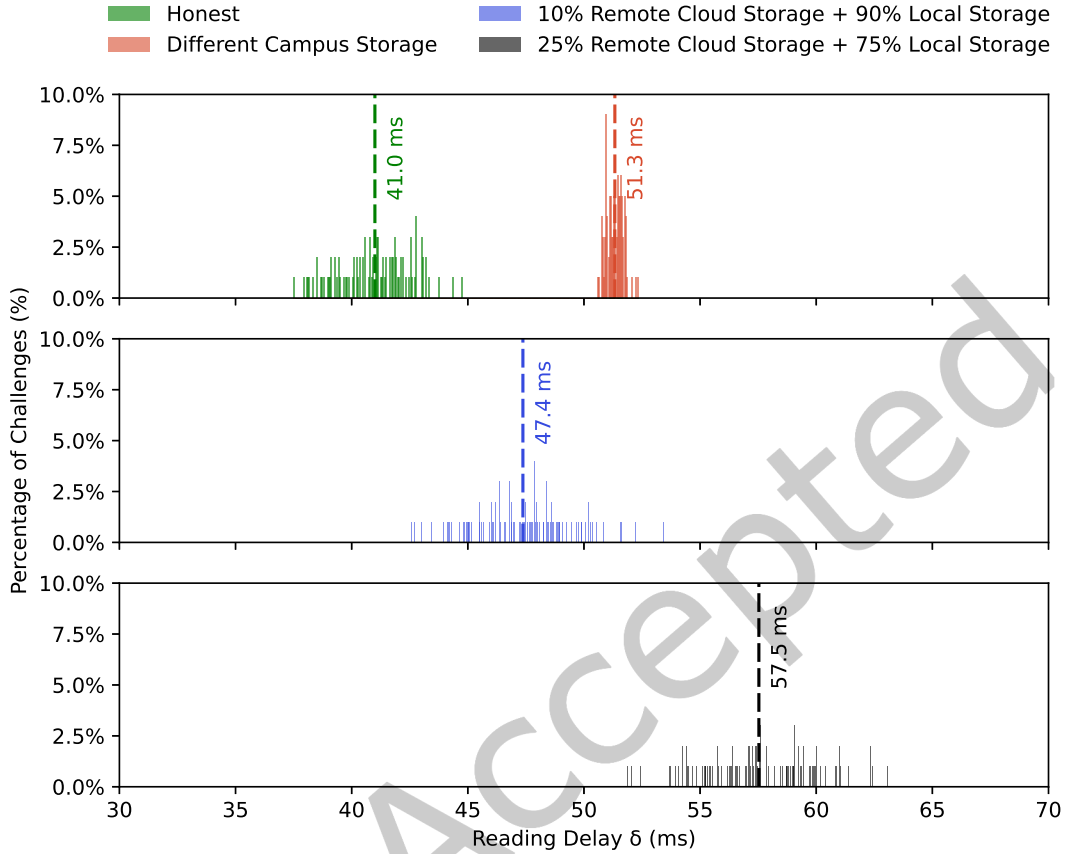


Fig. 15. Distribution and correspondent averages for reading delay δ , in milliseconds, for nodes with storage in a single location versus nodes with mixed local and remote cloud storage.

δ threshold	$p = 10\%$	$p = 25\%$
44.8ms	90%	100%

Table 2. PoATR detection rate for different percentages $p\%$ of remote data

6.4.1 PoUTR Configuration. When configuring the Proof of Uniform Timely-Retrievability, the focus shifts from the prior emphasis on minimizing the error ϵ_{\max}^{δ} in PoATR. Instead, our goal with PoUTR is to directly analyze the variance among multiple samples of the challenge.

As such, a new value for N must be experimentally determined to align with the specific goals of PoUTR. The experimental process involves varying the challenge size N and observing how this parameter affects the standard deviation across $k = 200$ executions of the challenge on nodes using different values of p . The results are presented in Figure 16. As expected, nodes configured with 50% remote storage exhibit the most pronounced standard deviation. Then it progressively diminishes as the configuration approaches either a single-source

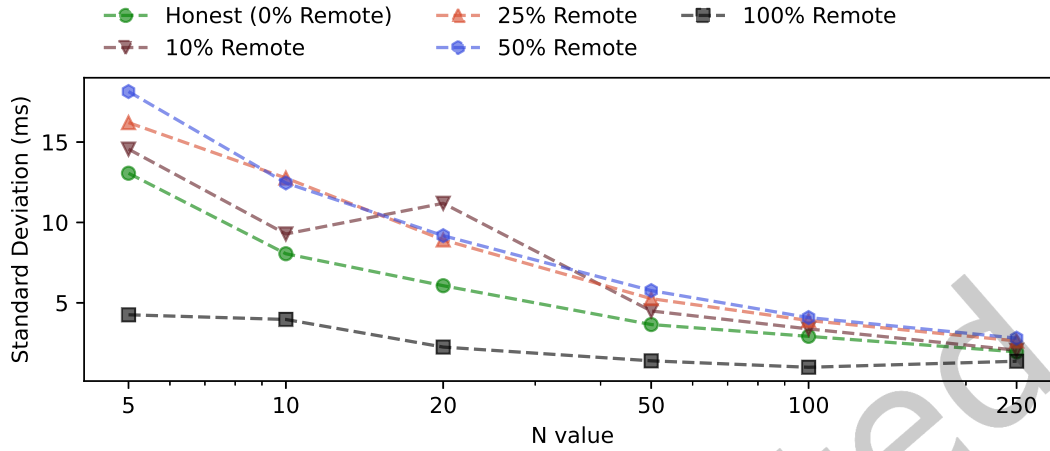


Fig. 16. Standard deviation (in ms) for δ on nodes with different values of $p\%$ using $k = 200$ challenges.

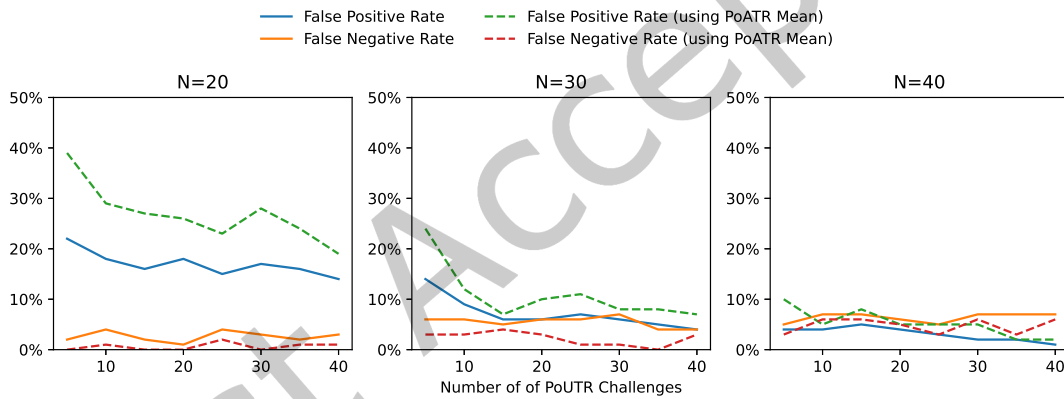


Fig. 17. Accuracy when using $\sigma = 6.8ms$ standard deviation threshold on nodes with 10% remote storage, for various values of N and number of challenges k .

scenario with fully local or remote storage. PoUTR can be configured to achieve good performance for nodes with as little as 10% and 25% remote storage. To discern between honest and malicious nodes at these minimal remote storage percentages, opting for smaller N values is necessary. We will assess the accuracy achieved when using values different values of N .

6.4.2 PoUTR Accuracy. To assess the accuracy of PoUTR, an evaluation encompassing 100 sets of challenges was conducted. These sets varied in size, ranging from 5 to 40 challenges per set. Across the range, the experiment was performed for N values of 20, 30 and 40. Furthermore, 100 PoATR challenges were ran to provide the $\bar{\delta}$ to be used when calculating the standard deviation. We have experimentally determined $\sigma = 6.8ms$ to be a good upper limit for the standard deviation of an honest node. The results of combining PoATR and PoUTR when

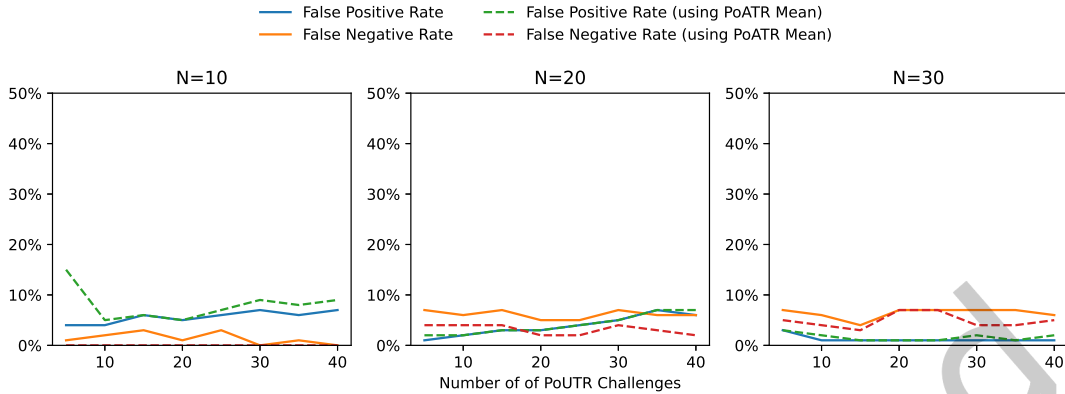


Fig. 18. Accuracy when using $\sigma = 6.8ms$ standard deviation threshold on nodes with 10% remote storage, for various values of N and number of challenges k , without background load on audited node.

using this value are presented in Figure 17, where it can be seen that a value of $N = 40$ outperforms the other two configurations, with 93% malicious nodes being detected with just $k = 35$ samples.

Since PoUTR is expected to be run in combination with PoATR, we leverage the results from the first test to improve PoUTR results, by using the more precise $\bar{\delta}$ calculation in the standard deviation formula. The improved results can be seen in the dashed lines from Figure 17, reducing false negatives from 7% to 3%, while keeping false positives at 2% in the $N = 40, k = 35$ configuration.

Therefore, combining PoATR with PoUTR provides a significant improvement over running PoATR alone, reducing the false negative rate from 10%, as presented in Table 2, to just 3%.

When the background load generation is disabled, similar results can be obtained for even smaller shares of remote data, as can be seen in Figure 18. In particular, with just 5% of the data remote, the lighter $N = 30, k = 15$ configuration can achieve 3% false negatives and 1% false positives.

7 Conclusions

We present a novel auditing mechanism that is capable of extracting a proof of timely-retrievability, that is, a proof that a given storage node is able to serve requests without violating some given data access latency constraint δ . We have implemented and evaluated two distinct configurations of the proof, one tailored to estimate the average latency experience by clients and the other tailored to assess its variance. The proof is designed in a way that, if the storage node does not store locally a significant fraction of the objects, it will be unable to respond in time. We enforce our proof to be executed in the audited node by leveraging SGX enclaves, for iteratively revealing the next data block to be read to the untrusted part. Our evaluation shows that our proof can accurately detect a node that cannot satisfy the target latency constraint δ under different edge computing scenarios, resulting in the detection of a misbehaving storage provider, including sophisticated attackers that only store parts of the data remotely.

Acknowledgments

This research partially supported by the Fundação para a Ciência e a Tecnologia (FCT) via scholarship 2020.05270.BD, the INESC-ID grant UIDB/50021/2020 (DOI:10.54499/UIDB/50021/2020) and the DACOMICO project (financed by the OE with ref. PTDC/CCI-COM- /2156/2021)."

References

- [1] AbdelRahman Abdou, Ashraf Matrawy, and Paul C van Oorschot. 2015. Accurate one-way delay estimation with reduced client trustworthiness. *IEEE Communications Letters* 19, 5 (2015), 735–738.
- [2] Akamai. [n. d.]. Online Retail Performance Report: Milliseconds Are Critical. <https://www.akamai.com/newsroom/press-release/akamai-releases-spring-2017-state-of-online-retail-performance-report>.
- [3] Akamai. [n. d.]. Real time key value store for edge computing. <https://www.akamai.com/products/edgekv>.
- [4] Akamai. [n. d.]. Supercharge your EdgeWorkers apps with a serverless key-value store. <https://www.akamai.com/products/edgekv>.
- [5] Fritz Alder, Gianluca Scoppelliti, Jo Van Bulck, and Jan Tobias Mühlberg. 2023. About Time: On the Challenges of Temporal Guarantees in Untrusted Environments. In *Proceedings of 6th Workshop on System Software for Trusted Execution (SysTEX 2023)*. Rome, Italy.
- [6] Ittai Anati, Shay Gueron, Simon Johnson, and Vincent Scarlata. 2013. Innovative Technology for CPU Based Attestation and Sealing. In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*. Tel-Aviv, Israel.
- [7] Fatima Anwar, Luis Garcia, Xi Han, and Mani Srivastava. 2019. Securing Time in Untrusted Operating Systems with TimeSeal. In *Proceedings of the 40th IEEE Real-Time Systems Symposium (RTSS)*. York, England.
- [8] Frederik Armknecht, Ludovic Barman, Jens Bohli, and Ghassan Karame. 2016. Mirror: Enabling proofs of data replication and retrievability in the cloud. In *Proceedings of the 25th USENIX Security Symposium*. Austin (TX), USA.
- [9] Giuseppe Ateniese, Randal Burns, Reza Curtmola, Joseph Herring, Lea Kissner, Zachary Peterson, and Dawn Song. 2007. Provable Data Possession at Untrusted Stores. In *ACM Conference on Computer and Communications Security*. Alexandria, Virginia, USA.
- [10] Autopoietic Cognitive Edge-cloud Services (ACES). [n. d.]. Cloud-edge service for data management. <https://www.aces-edge.eu/>.
- [11] Elaine Barker and Allen Roginsky. 2019. *Transitioning the use of cryptographic algorithms and key lengths*. NIST Special Publication 800-131A Revision 2. National Institute of Standards and Technology.
- [12] J. Benet. 2014. IPFS: Content Addressed, Versioned, P2P File System. *arXiv preprint arXiv:1407.3561* (2014).
- [13] J. Benet, D. Dalrymple, and N. Greco. 2017. Proof of replication. *Protocol Labs, July 27* (2017).
- [14] Karyn Benson, Rafael Dowsley, and Hovav Shacham. 2011. Do You Know Where Your Cloud Files Are?. In *Proceedings of the 3rd ACM Workshop on Cloud Computing Security Workshop*. Chicago (IL), USA.
- [15] Jin-Hee Choi and Chuck Yoo. 2005. One-way delay estimation and its application. *Computer Communications* 28, 7 (2005), 819–828.
- [16] Cloudflare. [n. d.]. CDN global network. <https://www.cloudflare.com/cdn/>.
- [17] Cláudio Correia, Miguel Correia, and Luís Rodrigues. 2020. Omega: a Secure Event Ordering Service for the Edge. In *Proceedings of the 50th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. València, Spain.
- [18] Hung Dang, Erick Purwanto, and Chien Chang. 2017. Proofs of Data Residency: Checking Whether Your Cloud Files Have Been Relocated. In *ACM on Asia Conference on Computer and Communications Security*. Abu Dhabi, United Arab Emirates.
- [19] J. Ekberg, K. Kostianen, and N. Asokan. 2014. The Untapped Potential of Trusted Execution Environments on Mobile Devices. *IEEE Security Privacy* 12, 4 (2014).
- [20] Leonardo Epifâneo, Cláudio Correia, and Luís Rodrigues. 2021. Cathode: A Consistency-Aware Data Placement Algorithm for the Edge. In *2021 IEEE 20th International Symposium on Network Computing and Applications (NCA)*. 1–10.
- [21] Filecoin. [n. d.]. A blockchain based storage network. <https://spec.filecoin.io/>.
- [22] M. Gondree and Z. Peterson. 2013. Geolocation of Data in the Cloud. In *Proceedings of the 3rd ACM Conference on Data and Application Security and Privacy (CODASPY)*. San Antonio (TX), USA.
- [23] Grand View Research. [n. d.]. CDN Market Report. <https://www.grandviewresearch.com/industry-analysis/content-delivery-networks-cnd-market>.
- [24] Haryadi Gunawi, Mingzhe Hao, Riza Suminto, Agung Laksono, Anang Satria, Jeffrey Adityatama, and Kurnia Eliazar. 2016. Why does the cloud stop computing? lessons from hundreds of service outages. In *Proceedings of the 7th ACM Symposium on Cloud Computing (SOCC)*. Santa Clara (CA), USA.
- [25] Tao Jiang, Wenjuan Meng, Xu Yuan, Liangmin Wang, Jianhua Ge, and Jianfeng Ma. 2021. ReliableBox: Secure and Verifiable Cloud Storage With Location-Aware Backup. *IEEE Transactions on Parallel and Distributed Systems* 32, 12 (2021).
- [26] L. Li and L. Lazos. 2020. Proofs of Physical Reliability for Cloud Storage Systems. *IEEE Transactions on Parallel and Distributed Systems* 31 (2020).
- [27] Ivan Lujic, Vincenzo De Maio, Srikumar Venugopal, and Ivona Brandic. 2022. SEA-LEAP: Self-Adaptive and Locality-Aware Edge Analytics Placement. *IEEE Transactions on Services Computing* 15, 2 (2022), 602–613.
- [28] Sinisa Matetic, Mansoor Ahmed, Kari Kostianen, Aritra Dhar, David Sommer, Arthur Gervais, Ari Juels, and Srdjan Capkun. 2017. ROTE: Rollback Protection for Trusted Execution. In *Proceedings of the 26th USENIX Security Symposium*. Vancouver, Canada.
- [29] Mithun Mukherjee, Rakesh Matam, Lei Shu, Leandros Maglaras, Mohamed Ferrag, Nikumani Choudhury, and Vikas Kumar. 2017. Security and Privacy in Fog Computing: Challenges. *IEEE Access* 5 (2017).
- [30] Zhenyu Ning, Jinghui Liao, Fengwei Zhang, and Weisong Shi. 2018. Preliminary Study of Trusted Execution Environments on Heterogeneous Edge Platforms. In *Proceedings of the IEEE/ACM Symposium on Edge Computing (SEC)*. Bellevue (WA), USA.

- [31] Gleen Ricart. 2017. A city edge cloud with its economic and technical considerations. In *Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*. Kona (HI), USA.
- [32] Mahadev Satyanarayanan, Phillip Gibbons, Lily Mummert, Padmanabhan Pillai, Pieter Simoens, and Rahul Sukthankar. 2017. Cloudlet-based Just-in-Time Indexing of IoT Video. In *Global Internet of Things Summit*. Geneva, Switzerland.
- [33] Zhong-Liang Shao, Cheng Huang, and Heng Li. 2021. Replica selection and placement techniques on the IoT and edge computing: a deep study. *Wireless Networks* 27, 7 (2021), 5039–5055.
- [34] Christopher Streiffer, Animesh Srivastava, Victor Orlikowski, Yesenia Velasco, Vincentius Martin, Nisarg Raval, Ashwin Machanavajjhala, and Landon Cox. 2017. ePrivateEye: To the Edge and Beyond!. In *Proceedings of the 2nd ACM/IEEE Symposium on Edge Computing (SEC)*. San Jose (CA), USA.
- [35] Swarm. [n. d.]. Distributed storage platform. <https://github.com/ethersphere/>.
- [36] Tarik Taleb, Konstantinos Samdanis, Badr Mada, Hannu Flinck, Sunny Dutta, and Dario Sabella. 2017. On Multi-Access Edge Computing: A Survey of the Emerging 5G Network Edge Cloud Architecture and Orchestration. *IEEE Communications Surveys Tutorials* 19, 3 (2017).
- [37] Hongliang Tian, Qiong Zhang, Shoumeng Yan, Alex Rudnitsky, Liron Shacham, Ron Yariv, and Noam Milshten. 2018. Switchless calls made practical in Intel SGX. In *Proceedings of the 3rd Workshop on System Software for Trusted Execution (SysTEX)*. Toronto, Canada.
- [38] Ahmad Vakili and Jean-Charles Gregoire. 2012. Accurate one-way delay estimation: Limitations and improvements. *IEEE Transactions on Instrumentation and Measurement* 61, 9 (2012), 2428–2435.
- [39] Yang Zhang, Weijing You, Shijie Jia, Limin Liu, Ziyi Li, and Wenfei Qian. 2022. EnclavePoSt: A Practical Proof of Storage-Time in Cloud via Intel SGX. *Security and Communication Networks* (2022).

Received 29 December 2023; revised 1 July 2024; accepted 22 July 2024