

Sharing Proofs of Retrievability across Tenants

Frederik Armknecht
University of Mannheim, Germany
armknecht@uni-mannheim.de

David Froelicher
NEC Labs Europe, Germany
david.froelicher@neclab.eu

Jens-Matthias Bohli
NEC Labs Europe, Germany
Hochschule Mannheim, Germany
JensMatthias.Bohli@neclab.eu

Ghassan O. Karame
NEC Labs Europe, Germany
ghassan@karame.org

ABSTRACT

Proofs of Retrievability (POR) are cryptographic proofs which provide assurance to a single tenant (who creates tags using his secret material) that his files can be retrieved in their entirety. However, POR schemes completely ignore storage-efficiency concepts, such as multi-tenancy and data deduplication, which are being widely utilized by existing cloud storage providers. Namely, in deduplicated storage systems, existing POR schemes would incur an additional overhead for storing tenants' tags which grows linearly with the number of users deduplicating the same file. This overhead clearly reduces the (economic) incentives of cloud providers to integrate existing POR/PDP solutions in their offerings.

In this paper, we propose a novel storage-efficient POR, dubbed SPORT, which transparently supports multi-tenancy and data deduplication. More specifically, SPORT enables tenants to securely share the same POR tags in order to verify the integrity of their deduplicated files. By doing so, SPORT considerably reduces the storage overhead borne by cloud providers when storing the tags of different tenants deduplicating the same content. We show that SPORT resists against malicious tenants/cloud providers (and against collusion among a subset of the tenants and the cloud). Finally, we implement a prototype based on SPORT, and evaluate its performance in a realistic cloud setting. Our evaluation results show that our proposal incurs tolerable computational overhead on the tenants and the cloud provider.

Keywords

Cloud security; Proofs of Retrievability; Deduplication

1. INTRODUCTION

The cloud is gaining several adopters among SMEs and large businesses that are mainly interested in minimizing the costs of both deployment and infrastructure management and maintenance.

Cost effectiveness is realized in the cloud through the integration of multi-tenancy and storage efficiency solutions with tailored distributed algorithms that ensure unprecedented levels of scalability and elasticity. The combination of multi-tenancy solutions with

storage efficiency techniques (e.g., data deduplication) promises drastic cost reductions. For instance, recent studies show that cross-user data deduplication can save storage costs by more than 50% in standard file systems [26, 27], and by up to 90-95% for back-up applications [26]. Moreover, nearly three quarters of these savings could also be obtained by means of straightforward whole-file deduplication [27].

The advent of cloud storage, however, introduces serious concerns with respect to the confidentiality, integrity, and availability of the outsourced data [21]. For instance, Google recently admitted permanent loss of customers' data in their storage systems due to a malfunction of local utility grid located near one of Google's data centers [4]. The literature features a number of solutions that enable users to verify the integrity and availability of their outsourced data [13, 15, 25, 30]. Solutions include Proofs of Retrievability (POR) [25, 30] which provide end-clients with the assurance that their data is retrievable, and Proofs of Data Possession (PDP) [7] which enable a client to verify that its stored data has not undergone any modifications. Although these solutions can be indeed effective in detecting data loss, they completely ignore storage-efficiency requirements, such as multi-tenancy and data deduplication, which are being widely utilized by existing cloud storage providers.

Namely, existing POR/PDP solutions assume a single trusted tenant (i.e., honest verifier) who pre-processes the files to create tags using secret material before outsourcing them to the cloud, and later regularly performs verifications on the pre-processed files and tags in order to react as early as possible in case of data loss.

A straightforward adaptation for the multi-tenant case would be that each tenant constructs and stores his own tags in the cloud. However, this approach threatens to cancel out the benefits of data deduplication over popular objects—which might reduce the (economic) incentives of cloud providers to integrate existing POR/PDP solutions in their offerings.

In Figure 1, we estimate the additional storage overhead incurred by such an instantiation using the existing privately-verifiable POR scheme of [30], the publicly-verifiable RSA-based POR/PDP schemes of [7, 30], and finally the publicly-verifiable BLS-based POR scheme of [30]. Here, we assume a block size of 8 KB. Notice that the smaller the block size is, the larger is the storage overhead due to the storage of the tenant tags (since the tags' sizes are fixed per block and are independent of the block size). Our results show that, even with block sizes as large as 8 KB, the metadata storage per tenant incurred by existing POR/PDP schemes are considerable. Clearly, such storage overhead reduces the profitability of existing storage-efficiency solutions and, in turn, reduces the incentives of providers to integrate POR/PDP schemes in their offerings. For example, in the 2048-bit RSA-based schemes of [7, 30], each POR tag requires around 1.5% additional storage overhead per tenant per

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

ASIA CCS '17, April 02-06, 2017, Abu Dhabi, United Arab Emirates

© 2017 ACM. ISBN 978-1-4503-4944-4/17/04...\$15.00

DOI: <http://dx.doi.org/10.1145/3052973.3052997>




Figure 1: Storage overhead due to the storage of individual tenant tags in (private and public) POR schemes of [30] and the public PDP scheme of [7]. Here, we assume 8 KB block sizes.

block; popular blocks, e.g., shared amongst 50 tenants, would then require 75% additional storage overhead.

Consequently, one way to minimize this additional storage overhead would require to use the same set of tags for each tenant. However, in practice, given that files are typically deduplicated across tenants, and different tenants do not tend to trust each other, tenants will be reluctant on sharing the secret material used to construct tags in POR/PDP. Notice that the leakage of the secret material invalidates the security of POR/PDP; for instance, a cloud provider which has access to the secret material can always construct correct responses in existing POR schemes—even if the outsourced data is deleted.

In this paper, we address this problem, and propose a novel solution (SPORT)—which goes one step beyond existing POR and transparently supports multi-tenancy and data deduplication. More specifically, SPORT enables different tenants to share the same POR tags in order to verify the integrity of their deduplicated files. By doing so, SPORT considerably reduces the storage overhead borne by cloud providers when storing the tags of different tenants deduplicating the same content. We show that SPORT resists against malicious tenants/cloud providers (and against collusion among a subset of the tenants and the cloud).

We argue that SPORT is technically and economically viable. Indeed, by reconciling functional (i.e., storage efficiency and multi-tenancy) and security requirements (i.e., data retrievability in the cloud), SPORT provides considerable incentives for cloud providers to issue security SLAs for their users accounting for data loss, in exchange, e.g., for additional service fees. While the main barriers of wider adoption of POR lie in the lack of compliance with functional requirements in the cloud, SPORT bridges these gaps and enables cloud providers to offer their customer extended POR services, without incurring additional bandwidth and storage overhead.

Our contributions can be summarized as follows:

Formal Framework. We propose the first formal framework and a security model for multi-tenant POR, dubbed MTPOR. Our framework extends the POR model outlined in [30] and addresses security risks that have not been covered so far in existing models.

Concrete Instantiation. We describe a concrete MTPOR scheme, dubbed SPORT, which builds upon publicly verifiable BLS POR scheme [30] and that is secure in the MTPOR model. SPORT deploys a novel mechanism which enables different tenants to (in-place) update the POR tags created by others corresponding to the same deduplicated file—thus saving considerable storage. We show that this process resists against

a malicious cloud provider, and malicious tenants (as well as against any collusion between tenants/providers).

Prototype Implementation. We implement and evaluate a prototype based on SPORT, and we show that our proposal incurs comparable overhead on the users and cloud providers, when compared to existing publicly-verifiable POR/PDP schemes.

The remainder of this paper is organized as follows. In Section 2, we introduce a novel framework for secure multi-tenant POR, MTPOR. In Section 3, we propose SPORT, an efficient instantiation of MTPOR, and analyze its security. In Section 4, we describe a prototype implementation and evaluation of SPORT and compare its performance to existing publicly-verifiable POR/PDP schemes. In Section 5, we overview related work in the area and we conclude the paper in Section 6.

2. MULTI-TENANT PROOFS OF RETRIEVABILITY

In this section, we introduce a formal model for multi-tenant POR (MTPOR). Since MTPOR extends POR, we first quickly recall the POR model from [30].

2.1 Single-Tenant Proofs of Retrievability

Proofs of Retrievability (POR) are interactive protocols that cryptographically *prove* the retrievability of outsourced data. More precisely, POR consider a model comprising of a *single* user (or tenant), and a service provider that stores a file pertaining to the user. Note that POR only provide a guarantee that a fraction p of the file can be retrieved. For that reason, POR are typically performed on a file which has been erasure-coded in such a way that the recovery of any fraction p of the stored data ensures the recovery of the file [6].

A POR scheme consists of four procedures [30], **setup**, **store**, **verify**, and **prove**. The latter two algorithms define a protocol for proving file retrievability. We refer to this protocol as the POR *protocol* (in contrast to a POR scheme that comprises all four procedures).

setup. This randomized algorithm generates the involved keys and distributes them to the parties. In case public keys are involved in the process, these are distributed amongst all parties.

store. This randomized algorithm takes as input the keys of the user and a file \tilde{M} . The file is processed and **store** outputs M which will be stored on the server. **store** also generates a file tag τ , which contains additional information about M shared by the user and the service provider and is used in the subsequent protocol.

prove, verify. The randomized proving and verifying algorithms define an interactive protocol for proving file retrievability. During protocol execution, both algorithms take as input the public key. The **prove** algorithm also takes as input the processed file M , and the **verify** algorithm takes as input the secret key. At the end of the protocol run, **verify** outputs TRUE if the verification succeeds, meaning that the file is being stored on the server, and FALSE otherwise.

2.2 Multi-Tenant POR (MTPOR)

In the sequel, we formally define the model for multi-tenant POR, dubbed MTPOR. Similar to the standard POR model, we distinguish between tenants and the service provider S . The main difference to the aforementioned POR model lies in the direct integration of the notion of multiple tenants (or users) that *can* upload the *same* file to be stored at S . In general, these tenants are independent of each other and do not necessarily trust each other. That is, we do

not assume the existence of shared secrets or direct communication between the different tenants.

Conforming with the operation of existing cloud providers, we assume that S stores duplicate data (either at the block level or the file level) uploaded by different users only once—thus tremendously saving storage costs. Recent studies show that cross-user data deduplication can save storage costs by more than 50% in standard file systems [26, 27], and by up to 90-95% for back-up applications [26]. Nonetheless, any user who uploaded a file should be able to execute PORs later on to verify that the file has been stored. Therefore, we assume that the provider stores for each uploaded file M the following:

1. The file M itself which is stored only *once*.
2. Additional metadata $\text{Data}(M)$ which contains e.g., a list of the users who uploaded M and information for conducting the POR on M .

If a provider stores several files, e.g., M_1, \dots, M_n , we assume that S maintains separate metadata $\text{Data}(M_i)$ for each file that he stores. Clearly, metadata corresponding to different files may contain duplicate information about the same user to account for the case where the user has uploaded different files. For simplicity, if we refer to a metadata $\text{Data}(M)$ where the file M has not been uploaded yet, we assume that the metadata is empty, e.g., $\text{Data}(M) = \perp$. Moreover, for ease of presentation, we simply write Data instead of $\text{Data}(M)$, whenever the considered file is clear from the context.

In the sequel, we assume that all communication is authenticated and encrypted (e.g., using the TLS protocol). Similar to the POR model outlined in Section 2.1, we consider three basic procedures: Setup, Store, and POR where POR is an interactive protocol that extends the notions of *verify* and *prove* discussed earlier.

The Setup Protocol.

The Setup protocol is used to establish the keys that are required by the users to instrument the POR; notice that these keys may be different from the keys used to authenticate and encrypt the communication. The input of this protocol is the security parameter λ and the outputs are a secret and a public key. In case such keys are not required in some settings, the respective values are simply \perp . Formally, it holds for a party running Setup that $(sk, pk) \leftarrow \text{Setup}(\lambda)$. Whenever it is necessary to refer to the keys of a specific user U , we denote the corresponding secret and public key by sk_U and pk_U , respectively.

The Store Protocol.

This randomized file-storing protocol takes as input the keys of the user, a file \tilde{M} from the user U that has to be stored by the service provider and the current metadata of the provider. During the protocol execution, it may happen that the file is processed prior to being stored—we denote this file by M . The output of Store is a verification key vk_U that allows the user to run the proof of retrievability. The verification key may be equal to (or depend on) the secret or public key that the user generated in the Setup procedure. On the provider's side, the (processed) file M is stored and the metadata $\text{Data}(M)$ is updated. Formally, it holds that:

$$\begin{aligned} \text{Store:} \quad & [U : sk_U, pk_U, \tilde{M}; S : \text{Data}(M)] \\ & \longrightarrow [U : vk_U; S : M, \text{Data}(M)]. \end{aligned}$$

Observe that unlike the single-tenant POR model outlined in Sec. 2.1, we do not utilize the notion of a file tag τ that needs to be known to both parties. Instead, we split this information into two parts: (i) the part that needs to be known to the user (or verifier)—captured by the notion of a verification key vk_U —and (ii) the part

that needs to be stored by the provider, which is part of the metadata $\text{Data}(M)$.

Observe that the Store protocol essentially requires a user to *join* the set of users who all uploaded the same file M . In practice, one may also prefer to have a protocol that allows users to *leave* this set. However, as we aim to achieve that the model for a multi-tenant POR falls back to the classical single-tenant POR when only one user is present, we do not integrate such a protocol into the model.

The POR Protocol.

The aim of the POR protocol is to allow a user (or a verifier) to check if the file is still entirely stored at the provider. To this end, the user uses his personal keys and the verification key to verify the response sent by the provider; the latter uses the public key of the user, the uploaded file, and the metadata to issue the response to the user. At the end, the user outputs a Boolean value $\text{dec} \in \{\text{TRUE}, \text{FALSE}\}$. Formally, we have:

$$\begin{aligned} \text{POR} \quad & : [U : sk_U, pk_U, vk_U; S : pk_U, M, \text{Data}(M)] \\ & \longrightarrow [U : vk_U, \text{dec}]. \end{aligned}$$

The protocol run is accepted by the user if and only if $\text{dec} = \text{TRUE}$. Unlike the Store protocol, we point out that metadata maintained by the provider is not modified in this protocol. However, it may be the case that the user has to update his verification key vk_U —which explains the reason why this is part of the local output.

Adopting common terminology, we say that an MTPOR is *publicly-verifiable* if the POR-protocol does not require the secret key sk_U of the user. In such cases, verification can be outsourced to a third party verifier V . Otherwise, an MTPOR is said to be *privately-verifiable*.

2.3 Correctness

We consider correctness from the perspective of the provider and the user. From the former perspective, it is necessary to ensure that if two honest users initiate Store with the same file \tilde{M} , then during the protocol run the same file M is outsourced. This notion ensures correct and effective whole-file deduplication.

From a user's perspective, correctness intuitively means that if an *honest* user U uploaded a file to an *honest* provider S and later on runs the POR protocol, the user should accept. For a formal treatment though, one needs to take into account that in the time lapse between the file upload by a user U and the execution of the POR procedure, the provider may have executed other Store protocol runs (i.e., with another user U') which may have changed the metadata. For example, while the file M would not be stored a second time, information about U' would be integrated into the metadata $\text{Data}(M)$. Still, the POR protocol executed by U afterwards should be successful although $\text{Data}(M)$ has changed from the initial Store performed by U . We stress that this property has to hold even if some tenants are malicious, e.g., deviate from the protocol, as long as both U and S are honest. We capture this formally by saying that the new content of $\text{Data}(M)$ *evolved* from the last known metadata that has been generated when U executed Store.

DEFINITION 1 (EVOLVABLE METADATA). For a key pair (sk, pk) , a file M , and two metadata $\text{Data}, \text{Data}'$, we denote by $\mathbb{E} = \mathbb{E}(sk, pk, M, \text{Data}, \text{Data}')$ the event that when running Store with inputs sk, pk, M on the user side and input Data on the provider side where the provider behaves honestly, the metadata with respect to M is updated to Data' . We say that Data' is *evolvable* from Data with respect to (sk, pk, M) if:

$$\Pr [\mathbb{E}(sk, pk, M, \text{Data}, \text{Data}')] > 0. \quad (1)$$

This captures that the metadata Data may be possibly updated to Data' while running Store for a file M with keys (sk, pk) . We express this by:

$$\text{Data} \xrightarrow{(\text{sk}, \text{pk}, M)} \text{Data}'. \quad (2)$$

Moreover, we extend this notion to the case where one metadata may be evolved from another metadata also after running Store multiple times. That is, we write $\text{Data} \xrightarrow{M} \text{Data}'$ if there exists metadata $\text{Data}_1, \dots, \text{Data}_\ell$ and key pairs $(\text{sk}_1, \text{pk}_1), \dots, (\text{sk}_\ell, \text{pk}_\ell)$ such that:

- $\text{Data} \xrightarrow{(\text{sk}_1, \text{pk}_1, M)} \text{Data}_1$
- $\text{Data}_i \xrightarrow{(\text{sk}_{i+1}, \text{pk}_{i+1}, M)} \text{Data}_{i+1}$ for $i = 1, \dots, \ell - 1$
- $\text{Data}_\ell \xrightarrow{(\text{sk}_{\ell+1}, \text{pk}_{\ell+1}, M)} \text{Data}'$

In this case, we likewise say that Data' is evolvable from Data . Finally, we say that a metadata Data is M -evolvable if $\perp \xrightarrow{M} \text{Data}$. Recall that the metadata is initially set to \perp . Hence, this notion expresses that Data could be a possible metadata resulting from multiple executions of Store for the same file M .

Note that the definition of evolvable metadata refers to the file M as seen by S (i.e., after the user has uploaded it to the cloud). We are now ready to give the formal definition for correctness of a MTPOR.

DEFINITION 2 (CORRECTNESS). An MTPOR scheme is correct if the following holds with overwhelming probability in the security parameter for any key pairs (sk, pk) , any file M , and any M -evolvable metadata Data .

Consider a Store execution with the above mentioned inputs and the outputs vk and Data' . Then, it holds for any metadata Data'' with $\text{Data}' \xrightarrow{M} \text{Data}''$, and any protocol execution of POR with inputs $\text{sk}, \text{pk}, \text{vk}$ (user side) and M, Data'' (provider side), the user accepts at the end of the protocol run.

2.4 Security

MTPOR exhibits similar security considerations as in the case of single-tenant POR. Recall that the purpose of executing a POR or an MTPOR is to allow users to verify whether the file is still entirely stored by the provider. Consequently, we do not consider in this paper other security goals besides data retrievability. For example, to ensure the confidentiality of the file \tilde{M} , the user can encrypt the file during the Store protocol by leveraging message-locked encryption using a key derived from the file contents such as [5, 9]. We also do not take into account any privacy leakage resulting from the use of deduplication [23]. More specifically, we assume that the cloud storage employs Proofs of Ownership [22] in order to ensure that tenants indeed possess the files.

In a nutshell, our attacker model concentrates on a dishonest provider A who aims to mislead the user by storing few (or no) parts of the file but still tries to pass the POR protocol. Hence, security can be captured by a similar notion as in single-tenant POR, i.e., using the notion of an extractor. Recall that an extractor algorithm expresses the notion that if a provider is able to convince a user within the POR protocol, one can extract the file from the provider.

However, as opposed to the case of single-tenant POR, we have to consider a stronger attacker model where A might collude with one or several tenants. We call such tenants *corrupted*. Corrupted tenants are under the full control of the attacker. In particular, the attacker knows all their keys, including the verification key.

To formalize security, we envision a game between an adversary A and an environment E . The task of the environment is to play the role of all honest users and to challenge the adversary. That is,

whenever we say that an honest user executes a certain protocol, we mean that the environment honestly executes the protocol on behalf of this user. However, we allow the attacker to initiate protocol runs by honest users at her wish. To keep track of all honest users, we assume that the environment maintains a set \mathcal{U} . The adversary can now interact with the environment using four types of queries: a Setup -query, a Store -query, a POR -query, and a Corrupt -query.

In a Setup -query, a new honest user U is created. That is, the environment generates a key pair $(\text{sk}_U, \text{pk}_U)$ and stores a quadruple $(U, \text{sk}_U, \text{pk}_U, \text{VK}_U)$ in \mathcal{U} . The first entry identifies the user, the following two entries give the keys of the user, while VK_U represents the set of verification keys stored by the user—each per file stored by U . Initially, VK_U is set to \emptyset . In addition, we assume that the environment hands the public key to the adversary.

In a Store -query, the adversary presents a user U and a file M to the environment E . The environment stores for each file M that has been part of a Store -request a set $\text{Users}(M)$ that contains the set of users that have executed Store with the attacker for the file M . Upon receiving a Store -query, the environment proceeds as follows:

- If the user has not been created before (in terms of a Setup -request) or if the user is corrupted, the environment aborts the game. This is equivalent to checking if some quadruple $(U, \text{sk}_U, \text{pk}_U, \text{VK}_U)$ is stored in \mathcal{U} .
- The environment executes on behalf of the user the Store protocol with the adversary, using the key pair sk_U, pk_U and the file M as input. Let vk_U denote the output for the user. If the set VK_U of verification keys stored for user U already contains a tuple (M, vk'_U) , this tuple is replaced by (M, vk_U) . This reflects the case that an honest user runs Store again for the same file. Otherwise, VK_U is updated to $\text{VK}_U \leftarrow \text{VK}_U \cup \{(M, \text{vk}_U)\}$ and the quadruple stored in \mathcal{U} is updated accordingly.
- Finally, the environment stores U in $\text{Users}(M)$ if this is not the case already.

Observe that this captures the situation where the adversary runs the Store protocol for a given file with an honest user. As he can simulate the protocol with a corrupted user on his own, we do not need a separate oracle query for this case. Similarly, corrupted users do not need to be stored in \mathcal{U} —see also below.

In a POR -query, the attacker A gives a user U and a file M to E . The environment first checks if U is recorded in $\text{Users}(M)$ and if a quadruple $(id, \text{sk}_U, \text{pk}_U, \text{VK}_U)$ is stored in \mathcal{U} . If this is not the case, E aborts. Otherwise, it executes on behalf of U the POR -protocol with A , using the values $\text{sk}_U, \text{pk}_U, \text{vk}_U$ as input such that $(M, \text{vk}_U) \in \text{VK}_U$. The purpose of this query is to allow the attacker to execute the POR -protocol with an honest user of its choice. Also here, POR executions with corrupted users can be simulated by A and hence do not require any interaction with the environment.

Finally, in a Corrupt -query, the attacker A hands a user U to E . If the user is not recorded in \mathcal{U} , it aborts the game. Otherwise, it locates the corresponding quadruple $(id, \text{sk}_U, \text{pk}_U, \text{VK}_U) \in \mathcal{U}$ and hands the secret key sk_U and the set of verification keys VK_U to the attacker. Furthermore, it removes the quadruple from \mathcal{U} , indicating that U is no longer honest.

The attacker is allowed to make a polynomial number of queries (i.e., polynomial in the security parameter λ). We say that the attacker is *static* if he makes no Corrupt -queries and *adaptive* otherwise. Notice that a static adversary can still collude with corrupted users but these are under the control of the attacker right from the start. That is, the attacker can create both honest and corrupted users but cannot corrupt honest users. Eventually, the attacker quits the game by selecting a file M , an honest user U that is recorded in $\text{Users}(M)$, and by providing the description of a cheating prover P .

Observe that the attacker cannot quit if no honest user is present at this point in time.

From this point on, the definitions are analogous to the standard POR definitions [30]. Namely, we say that a cheating prover P is ε -admissible if in an ε fraction of POR runs with U , this user accepts at the end of the protocol. Here, the probability is over the coins of the user and the prover.

DEFINITION 3 (SOUNDNESS). *We say a MTPOR is ε -sound if there exists an efficient extraction algorithm Extr such that, for every adversary A , whenever A executes the game explained above, outputs a ε -admissible cheating prover P for a file M and a user U , the extraction algorithm recovers M from P except with negligible probability. We say that it is strongly ε -sound if this property holds in the presence of an adaptive attacker.*

2.5 Tag Sharing

Within our model, we distinguish between two types of data that are stored by the provider: the uploaded files and associated metadata. The crucial aspect of multi-tenant POR lies in the deduplication of the files, meaning that duplicate files/objects are stored only once. For generality purposes, multi-tenant POR, however, does not make any specific assumptions about the storage of the associated metadata.

In fact, a trivial instantiation of multi-tenant POR would be that the metadata $\text{Data}(M)$ is simply a collection of the user data for all users $U \in \text{Users}(M)$ that are registered for this file M and the respective file tags τ_U as outlined for single-tenant PORs (cf. Section 2.1). Here, each tenant U constructs his own tags τ_U using his own secret material; that is, only the file is deduplicated, but the tags are stored separately for each user. While this approach would resist against malicious tenants, its storage overhead is essentially the number of tags times the number of users (cf. Figure 1).

In this section, we address this problem and introduce the definition of *tag sharing*. The overall goal of tag sharing is to enable storage efficiency with respect to the storage of that tags. To motivate the definition, observe that the size $|\text{Data}(M)|$ of the metadata can be expressed as a function of two parameters¹: the size $|M|$ of M , and the number k of users registered to the file. For example, a straightforward realization of MTPOR based on a single-tenant POR consists of storing the file only once but storing the tags separately for each user. In this case, it would hold that $|\text{Data}(M)| \in \mathcal{O}(k \cdot |M|)$. That is, the number of users and the file size are directly coupled in the storage complexity. The goal of tag sharing is to make the storage consumed for storing the tags *independent* of the number of users. We capture this as follows.

DEFINITION 4 (TAG SHARING MTPOR). *We say that an MTPOR achieves tag sharing if it holds for all files M (with their respective number of users k) and the maintained metadata, that its size can be expressed by*

$$|\text{Data}(M)| = f(|M|) + g(k) \quad (3)$$

for some function f and g . We denote by $\text{Data} = (\text{Data}_{\text{file}}, \text{Data}_{\text{users}})$ the splitting of the metadata such that $|\text{Data}_{\text{file}}| = f(|M|)$ and $|\text{Data}_{\text{users}}| = g(k)$.

We stress at this stage that the storage efforts induced by the number of users and the file size, respectively, are *decoupled*. Notice that we make no assumptions on the functions f and g —for instance, g could be a constant function.

¹We assume here that global values such as the security parameter are fixed.

Next, we highlight an intrinsic property of the Store protocol that holds in a tag sharing MTPOR. Recall that in a tag sharing MTPOR, the size of $\text{Data}_{\text{file}}$ is $f(|M|)$ and is therefore independent of the number of users. Hence, even if the number k of users that upload this file continuously increases, the size of $\text{Data}_{\text{file}}$ is bounded by a function that solely depends on $|M|$. This means that, after a certain number of uploads, the information stored in $\text{Data}_{\text{file}}$ may still change when Store is executed but the size will no longer increase. We denote this variant of Store as Store^* :

DEFINITION 5. *In a tag sharing MTPOR, we refer to the variant of the Store protocol where the $\text{Data}_{\text{file}}$ part of the metadata does not further increase as Store^* . It formally holds*

$$\begin{aligned} \text{Store}^* : & \quad [U : \text{sk}_U, \text{pk}_U, \widetilde{M}; S : \text{Data}_{\text{file}}, \text{Data}_{\text{users}}] \\ \rightarrow & \quad [U : \text{vk}_U; S : M, \text{Data}_{\text{file}}^*, \text{Data}_{\text{users}}^*]. \end{aligned}$$

with $|\text{Data}_{\text{file}}^*| \leq |\text{Data}_{\text{file}}|$.

Notice that Store^* updates the *metadata* associated to a file but not the file itself.

3. SPORT: SHARING POR ACROSS TENANTS

In this section, we detail SPORT, an efficient publicly verifiable instantiation of tag sharing MTPOR. We start by outlining the main intuition behind our scheme.

3.1 Overview

One straightforward approach to realize a tag sharing MTPOR from existing single-tenant POR schemes would consist of sharing the keys required to verify the POR among all tenants deduplicating the same file. However, this approach would not resist malicious tenants who can share the secret material used to construct the POR with the cloud provider. Namely, in case the secret material is leaked to the provider (e.g., by means of a malicious/colluding tenant), then the provider may be able to correctly answer all POR challenges, without the need to store the file. *This equally applies to existing public-verifiable POR schemes; if the original creator leaks the private keys used to construct the POR, the underlying security of the POR is invalidated. This also rules out the possibility to derive the secret material to instrument POR from the file content itself (e.g., in a way similar to message-locked encryption [9]).* In this case, the entropy of the secret material will be dependent on the predictability of messages; providers can acquire the material e.g., by means of brute-force search, and/or by colluding with malicious tenants.

To overcome these challenges, SPORT extends the publicly verifiable POR scheme of SW [30] based on BLS signatures. Here, tags were essentially signatures on a combination of message blocks. Unlike [30] which exploits the message-homomorphism of BLS signatures to allow for compact batch verification of multiple blocks, SPORT goes one step further and leverages key-homomorphism of BLS signatures to combine different tags (protected by different keying material) constructed by multiple untrusted tenants into one tag—thus resulting in tremendous storage savings. Key-homomorphism here ensures that the underlying secret material of the combined tag depends on the contribution of each of the users registered to the same content. As a result, SPORT supports public verifiability and provides security against a collusion between tenants/service provider as long as there is one single honest tenant storing the content (and not compromising his own secret material).

SPORT achieves this without requiring any trusted entity (e.g., an identity manager or a proxy) to mediate the *update* of tags shared by

different tenants. In SPORT, this is realized using the (decentralized) Store* protocol—a variant of the standard Store POR protocol—that is executed by tenants deduplicating the same content. Unlike Store, the Store* routine in SPORT only involves updating the tags; we show that the tags can be efficiently verified in a single batch (for the entire file) by the provider. Moreover, we show that tenants only need to verify the inclusion of their own private keys within the updated keys—while the correctness of the tags is implicitly ensured during the POR protocol. Observe that the service provider does not have to be trusted when creating or updating the tags. This is the main reason why cryptographic accumulators could not be used in our setting.² However, most efficient accumulator constructs assume a *trusted setup*, meaning that the party who setups the accumulator needs to be trusted. Although a number of accumulators can deal with untrusted setup case [20, 28], these proposals are either inefficient or exhibit other impeding drawbacks.

In SPORT, although the Store* protocol incurs computational load on the clients, we argue that SPORT offers tremendous economic benefits to the clients and the service provider. Namely, SPORT allows clients to trade the *permanent* storage costs associated with the additional storage of file tags, with the *one-time* use of their computational resources to update the tags—thus reducing the total costs borne on users. Moreover, SPORT enables the provider to offer differentiated services, such as proving the integrity of stored files, with minimal additional storage overhead.

In the following subsections, we go into greater detail on the various parts of SPORT, starting with the protocol specification, then moving on to the security analysis of our scheme.

3.2 Protocol Specification

We now detail the specifications for the procedures of SPORT.

Specification of the Setup protocol

SPORT makes use of the multi-signature scheme based on BLS signatures described in [29]. To this end, it deploys a group G of prime order p with generator g , and a computable efficient bilinear map $e : G \times G \rightarrow G_T$. In addition, we require three independent secure hash functions $h_0, h_1, h_2 : \{0, 1\}^* \rightarrow G$ from the set of bit strings into the group G (e.g., the BLS hash function). We refer the readers to [29] for more details.

Each user U that participates in the scheme is required to setup a key pair of signature key and verification key. More precisely, the user chooses a private key $sk_U \in \mathbb{Z}_p$, and the corresponding public key is then given by $pk_U = g^{sk_U} \in G$. Moreover, U also needs to provide a proof of possession $\text{pop} = \text{pop}(sk_U)$ to prove possession of the secret key sk_U . Proofs of possession were introduced in [11] and prevent the provider from replacing/creating tags using other keys whose secret-keys are unknown to him. In a nutshell, our proof of possession (adapted from [29]) is constructed as follows: $\text{pop}(sk_U) = h_0(\langle pk_U \rangle)^{sk_U}$ and $\langle pk_U \rangle$ denotes an encoding of pk_U as a bit string. As we show later, such a proof of possession is required to protect against a rogue-key attack, where an adversary forges the multi-signature of a group using an invalid key (i.e., a public key whose private key is unknown to the adversary).

²Recall that an accumulator is a cryptographic technique that allows to represent a set S of elements by a single value w , the witness, such that membership for any element $x \in S$ can be proved using w and an according membership proof.

Specification of the Store protocol

To store a new file \widetilde{M} , the user first has to encode the file with a deterministic erasure code as required by the utilized POR (in order to provide extractability guarantees). As mentioned earlier, we assume that prior to this step the file is encrypted with a deterministic file-key securely obtained, e.g., using the schemes of [5, 9], in order to support deduplication of encrypted content across different users. The result is denoted by M .

Next, the user creates a unique file identifier fid for M , e.g., by hashing the encrypted file, or by computing the root of the Merkle tree of the file [5]. The server checks if the corresponding file with identifier fid has already been stored by another tenant. Now, we distinguish between two cases: (i) the file has not been stored previously, and (ii) the file is already stored. In the latter case, the user will continue with the Store* protocol (see below).

We start by discussing the former case where the outsourced file has not been stored previously at S . Similar to [30], the file is interpreted as n blocks, each is s sectors long. A sector is an element of \mathbb{Z}_p with p being the order of G and is denoted by m_{ij} with $1 \leq i \leq n, 1 \leq j \leq s$. That is, the overall number of sectors in the file is $n \cdot s$. For each $1 \leq j \leq s$ a group element $u_j \in G$ is pseudo-randomly extracted from the file as $u_j = h_1(H(M)||j)$ where H denotes a cryptographic hash function.

Similar to the BLS scheme of [30], for each block index $i \in \{1, \dots, n\}$ a file tag σ_i is computed as follows:

$$\sigma_i = \left(h_2(H(M)||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} \right)^{sk_U}. \quad (4)$$

Observe that σ_i effectively represents a BLS signature under the key sk_U of the following value:

$$w_i = h_2(H(M)||i) \cdot \prod_{j=1}^s u_j^{m_{ij}}, \quad (5)$$

The user uploads the file M , the public key pk_U along with the proof of possession $\text{pop}(sk_U)$, and the file tags $\tau = (\sigma_1, \dots, \sigma_n)$ to the provider S . The user stores locally $pk_M = pk_U$ as the verification key for the file and the hash value $H(M)$ to be able to reconstruct the values $\{u_j\}_{j=1 \dots s}$ later on.

Upon receiving $(M, pk_U, \text{pop}(sk_U), \tau)$, the provider S first verifies the validity of U 's key using $\text{pop}(sk_U)$. Furthermore, S checks that the tags σ_i in τ are indeed valid for the key pk_U . A straightforward approach would be to check that each tag σ_i indeed represents a signature of w_i , incurring an effort to evaluate n pairings, which incurs considerable computational overhead on the provider. In SPORT, we exploit the fact that the signature scheme is homomorphic and verify all the signatures in a single batch verification. More precisely, S samples for each block index i a random exponent $r_i \in \mathbb{Z}_p^*$ and checks whether the following verification holds:

$$e \left(\prod_i \sigma_i^{r_i}, g \right) \stackrel{?}{=} e \left(\prod_i w_i^{r_i}, pk_U \right). \quad (6)$$

Observe that this reduces the effort from n bilinear pairings to n exponentiations and one pairing. In Section 4, we show that this verification considerably reduces the computational overhead borne by the provider in verifying the tags. Note that effort can be further reduced by exploiting the fact that the values w_i are all defined over the same basis $\{u_i\}$ (cf. (5)). That is, instead of computing all w_i first ($= n \cdot s$ exponentiations) and the exponentiations $w_i^{r_i}$ (another n exponentiations), one could first compute the exponents r_j^* for each u_j (being $\sum_i r_i \cdot m_{ij}$ modulo p where p is known) and then

compute the values $h_2(H(M)||i)^{r_i}$ (n exponentiations) and $u_j^{r_j^*}$ (n exponentiations), summing up to $2n$ exponentiations instead of $(s+1) \cdot n$ exponentiations.

If the key and tags are correct, S sets $pk_M := pk_U$, being the public key associated to a secret key $sk_M := sk_U$. Moreover, the provider creates a log file, $pklog$, that will provide the necessary information needed for any other user to verify that the keys are well formed. Here, $pklog$ is initialized with $\{(pk_U, pop(sk_U))\}$. Following the notation from Section 2.5, the provider stores the file M , the metadata $Data_{users} = (pk_M, pklog)$, and file metadata $Data_{file} = \tau$.

Specification of the Store* Protocol

In case the file M is already stored and will be deduplicated, the Store* protocol is executed in order to update the file tags τ on the fly, allowing the new user U to obtain guarantees from the POR protocol without the need to trust the correctness of the already stored tags (created by other tenants).

In Store*, the user proceeds analogously to Store in order to compute the file M , its hash value $H(M)$, and the file tags $\tau = \{\sigma_i\}_{i=1 \dots n}$. Here, U also computes M and the hash value $H(M)$ and stores the same values as in Store. The main difference is that the user does not upload M but only $(pk_U, pop(sk_U), \tau)$ to S .

Upon reception of $(pk_U, pop(sk_U), \tau)$, S checks the validity of pk_U using $pop(sk_U)$. Moreover, he verifies the correctness of the newly uploaded tags σ_i . In principle, one could apply the same verification given in Equation 6. However, this would require the provider to recompute or store the values w_i . To avoid this, we leverage the bilinear map to combine the two verifications (in the signature) and the fact that the already stored file tags σ_j^* are correct (as they have been previously validated by S).

Namely, let $\tau = (\sigma_i)_{1 \leq i \leq n}$ denote the file tags uploaded by the user and $\tau^* = (\sigma_i^*)_{1 \leq i \leq n}$ be the tags stored by the provider. SPORT applies cross-verification between the uploaded tags τ and the stored tags τ^* . More precisely, an uploaded tag σ_i is correct if and only if it holds that:

$$e(\sigma_i, pk_M) \stackrel{?}{=} e(\sigma_i^*, pk_U), \quad (7)$$

Similar to the Store procedure, S does a batch cross-verification by sampling for each block index i a random exponent $r_i \in \mathbb{Z}_p^*$ and checks whether Equation 8 holds:

$$e\left(\prod_i \sigma_i^{r_i}, pk_M\right) \stackrel{?}{=} e\left(\prod_i (\sigma_i^*)^{r_i}, pk_U\right) \quad (8)$$

If the verification succeeds, S updates pk_M with $pk_M \cdot pk_U$, and likewise the already stored tags σ_i^* by replacing them with $\sigma_i^* \cdot \sigma_i$. Due to the homomorphic properties of the BLS signature, the new tags will still be correct signatures under the updated key pk_M with respect to an updated secret key $sk_M \rightarrow sk_M + sk_U$. The metadata about the file, $Data_{file}$, is hence updated to the new tags τ^* —observe that the size of $Data_{file}$ remains unchanged. Moreover, the key pk_M is updated in $Data_{users}$ while the public key and the corresponding proof of possession of the new user $(pk_U, pop(sk_U))$ are appended to $pklog$.

Finally, the provider sends to the user the log file $pklog$. Then, U proceeds to check the validity of each pair $(pk', pop') \in pklog$ with $pk' \neq pk_U$. In case all these keys pass the verification, U sets $pk_M := \prod_{(pk', pop') \in pklog} pk'$.

The core idea here is that, at each point in time, the public key pk_M associated to a file M corresponds to a secret key $\sum_U sk_U$, which captures the sum of all secret keys belonging to the users who uploaded M (see also the discussion in Section 3.3.2). This ensures

that the secret key to pk_M is unknown to the provider as long as at least one user is honest (and hence kept his key secret). In addition, our construct easily allows a user to *unsubscribe* from M . To do so, a user with secret key sk_U executes Store* again using the secret key $-sk_U$ instead. This results into removing (or cancelling) sk_U from sk_M .

Specification of the POR Protocol

To conduct POR on M for a user U , a verifier V (who could be the user himself) downloads the current verification key pk_M^* from the provider and compares it with the locally stored key pk_M .³ If these are not equal, this signals that a new tenant has deduplicated the same content, and updated the corresponding file tags. In this case, V checks that pk_M^* has been correctly updated and includes information from his local key pk_M .

To this end, S sends to V all pairs $(pk', pop') \in pklog$ of public keys and proofs of possession for currently registered users that joined since the last time that V or U executed a protocol with S with respect to M . Here, we assume that the keys are sorted in the order of the executions of Store/Store* by the users so that these values can be easily read out from $pklog$. We denote the set of these pairs by \widehat{pklog} .

Clearly, this allows U to learn the number of users that uploaded the same file. We point out, however, that this information leakage can be prevented if, for example, S inserts a number of “bogus” keys within $pklog$; in this case, U can only learn an upper bound on the number of other users that store the same file. Observe that the knowledge of the number of users deduplicating the same file can prove to be useful in a number of use-cases; for example, a number of studies show that such a level of transparency might motivate new pricing models in the cloud by allowing users to collaboratively share the costs of storing the same file [5].

Additionally, we point out that U might pseudorandomly create different signature keys for each file in order to prevent any entity to associate the user’s public key with the stored files.

Starting from the locally stored verification key pk_M , the verifier V verifies $pk_M^* \stackrel{?}{=} pk_M \cdot \prod_{(pk', pop') \in \widehat{pklog}} pk'$. Furthermore, for each $(pk', pop') \in \widehat{pklog}$, V checks the validity of the key pk' by verifying the proof of possession pop' . If the verification succeeds, V accepts the downloaded key pk_M^* as the verification key for M , updates the verification key pk_M locally to this value, and proceeds with the actual POR protocol as follows.

Subsequently, V creates a random challenge C and sends it to S . Here, C contains a random ℓ -element set of tuples (i, ν_i) where $i \in \{1, \dots, n\}$ denotes a block index, and $\nu_i \xrightarrow{P} \mathbb{Z}_N$ is a randomly generated integer. These values will be used to verify the retrievability of the data. Upon reception of C , S computes:

$$\sigma \leftarrow \prod_{(i, \nu_i) \in C} \sigma_i^{\nu_i} \in G, \quad (9)$$

$$\mu_j \leftarrow \sum_{(i, \nu_i) \in C} \nu_i m_{ij} \in \mathbb{Z}_p, \quad 1 \leq j \leq s. \quad (10)$$

Finally, these values are sent back to V who reconstructs u_i and checks that:

$$e(\sigma, g) \stackrel{?}{=} e\left(\prod_{(i, \nu_i) \in C} h_2(H(M)||i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu_j}, pk_M\right). \quad (11)$$

³If the verifier is not the user, we assume that he received all values from the user before except of the secret key sk_U .

If this verification does not pass, \mathcal{V} is certain that his file has been modified, and takes actions, such as informing the user (if different from \mathcal{V}) or downloading and repairing the file.

3.3 Correctness and Security

In this section, we show that SPORT is correct according to Definition 2 and secure according to the soundness property of Definition 3.

3.3.1 Correctness

In the most simple case of exactly one user (which needs to be honest then), correctness follows directly from the correctness of the BLS-based POR of [30].

To address the multi-tenant case, we assume first that all involved users are honest. We say that a tag σ is *correctly formed* with respect to a key pair (sk, pk) if it fulfills Equation 4. Recall that correctly formed tags are essentially signatures of values that depend on message blocks, i.e., $\sigma_1 = w^{sk_1}$, with respect to some verification key $pk_1 = g^{sk_1}$. This signature scheme is key-homomorphic in the following sense: given a second signature $\sigma_2 = w^{sk_2}$ of the same value w with respect to $pk_2 = g^{sk_2}$, then $\sigma_1 \cdot \sigma_2 = w^{sk_1 + sk_2}$ is a valid signature of w with respect to $g^{sk_1 + sk_2} = pk_1 \cdot pk_2$. It follows by induction for any file M and for any sequence of key pairs (sk_i, pk_i) , $i = 1, \dots, \ell$, that: $\perp \xrightarrow{(sk_1, pk_1, M)} \dots \xrightarrow{(sk_\ell, pk_\ell, M)} \text{Data}$ implies that $\perp \xrightarrow{(sk_1 + \dots + sk_\ell, pk_1 \cdot \dots \cdot pk_\ell, M)} \text{Data}$. Hence, according to the correctness of BLS POR of [30], it follows that POR-protocol accepts with $vk = pk_1 \cdot \dots \cdot pk_\ell$ if the tags are correctly formed with respect to $(sk^*, pk^*) := (sk_1 + \dots + sk_\ell, pk_1 \cdot \dots \cdot pk_\ell)$. This shows correctness for the case that all users behaved honestly.

Next, assume that some of the users are malicious (i.e., they deviate from the protocols). Let (pk_i, pop_i, τ_i) denote the values uploaded by the users during Store/Store*. Observe that G is a cyclic group and hence any value pk_i can be expressed as g^{sk_i} for some value $sk_i \in \mathbb{Z}_p$. This shows that pk_i represents a valid verification key. Moreover, the tag verifications executed by the provider (see Equations 6 and 8) ensure that the tags stored in Data are correctly formed. To see why, assume that a user uploads malicious tags $\sigma_i^* \neq \sigma_i$ for $i \in I \subseteq \{1, \dots, n\}$. As the user does not know the values $\{r_i\}_{i \in I}$, Equations 6 and 8 would be violated with high probability and hence not be accepted by the honest provider. We now give an explanation for Equation 6 (which also holds for Equation 8). Assume that the following holds:

$$e\left(\prod_i (\sigma_i^*)^{r_i}, g\right) = e\left(\prod_i w_i^{r_i}, pk_v\right) = e\left(\prod_i (w_i^{sk_v})^{r_i}, g\right).$$

It follows that:

$$\prod_i (\sigma_i^*)^{r_i} = \prod_i (w_i^{sk_v})^{r_i} \Leftrightarrow \prod_{i \in I} \underbrace{\left(\sigma_i^* \cdot (w_i^{sk_v})^{-1}\right)}_{\neq 1}^{r_i} = 1.$$

Since the values r_i are sampled uniformly at random and independently, the probability for this event is at most $1/|G|$.

Finally, the correctness of the proofs of possession pop_i are each individually tested according to the specifications. This shows that all values uploaded by the users need to be correctly built as they would be otherwise rejected with high probability by the provider.

3.3.2 Security

We now prove that SPORT is ε -sound (cf. Definition 3). Recall that this expresses security with respect to a static attacker. This means that the attacker may collude with corrupted parties which

are corrupted right from the start. That is, whenever an attacker creates a new honest user by a Store-query during the challenge game (cf. Section 2.4), it will not be corrupted later on.

Our proof follows in principle the line of arguments given for the BLS POR in [31]. There, the proof was divided into three parts; the first part shows that if a verifier accepts a response to a query, this response was constructed correctly with overwhelming probability. The other two parts prove that if correctly built responses are given, one can reconstruct a constant fraction of the blocks of the file (part two) and that with the help of erasure codes, one can derive the full original file from these blocks (part three). As the structure of the response in SPORT is exactly the same as in BLS POR and as we likewise suggest the use of an appropriate erasure code, the arguments given for parts two and three in [31] apply directly to SPORT as well. Consequently, it is sufficient to show the first part, i.e., that if a verifier \mathcal{V} accepts a response, it has to be correctly built with overwhelming probability.

SPORT uses a proof of possession pop to prove knowledge of the secret key. pop emerges as an efficient instantiation (as opposed to non-interactive zero-knowledge proofs of knowledge). We prove the security of SPORT in the knowledge of secret key model (KOSK) that was introduced in [11, 29]. The KOSK model incurs one limitation on the adversary; namely, the adversary has to prove knowledge of secret keys during the public key registration. This is modeled by requesting the adversary to output the public and secret keys of corrupted users in the key generation algorithm.

Recall that the CDH refers to the problem of deriving the value h^a where a is unknown from a given triple $(g, g^a, h) \in G^3$ and is usually assumed to be hard. We prove the following theorem:

THEOREM 1. *If the computational Diffie-Hellman problem (CDH) is hard in the bilinear group G , then an adversary in the KOSK model and the random oracle model, where h_1 , h_2 , and H are represented by random oracles, can only cause a verifier \mathcal{V} to accept a POR-instance by responding with correctly computed values $\{\mu_j\}$ and σ —except with negligible probability.*

We conduct a sequence of three games adapted from [31]. We start with Game 0 which is simply the challenge game defined in Section 2.4. As we operate in the KOSK model, it holds that when the adversary eventually outputs a prover P for a user U and a file M , he has to hand to the environment a value sk_i such that $g^{sk_i} = pk_i$, for each pair (pk_i, pop_i) communicated to \mathcal{V} in the context of M . As the environment performs the POR-protocols honestly on behalf of the user \mathcal{V} (i.e., without using this additional knowledge), this does not impact the probability that the verifier accepts.

Game 1 is the same as Game 0 with the difference that the environment keeps a list of all messages communicated within the challenge game. If the adversary is successful in any of the POR instances (i.e., the verifier accepts, but the adversary's signature σ' is not equal to $\sigma = \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i}$, see Equation 9), the environment aborts the security game. Let M denote the file considered within this POR-instance and pk_M the used public verification key. As explained above, it is ensured that $pk_M = g^{a+b}$ for some secret value a (which is the secret key of the user represented by \mathcal{V}) and a known value b (being the aggregation of all remaining keys).

Similar to [31], one can now construct from the adversary a simulator that simulates the honest users and the random oracle for the adversary and that solves the CDH for a given triplet $(g, g^a, h) \in G^3$ by using g^a as the public key of simulated honest users. However, since we have to cater for multiple users, one cannot use the same value g^a each time as the public key. Here, we exploit the fact that the tags are signatures and that these signatures are homomorphic. More precisely, whenever a new honest user U needs to be created

(in terms of a Setup-query), the simulator chooses a random value $r_v \in \mathbb{Z}_p$ and outputs $\text{pk}_v = g^a \cdot g^{r_v} = g^{a+r_v}$.

Next, the simulator has to compute the following signatures:

$$\left(h_2(H(M)||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} \right)^{a+r_v}. \quad (12)$$

Since r_v is known to the simulator and given that the signatures can be homomorphically combined, this task boils down to computing $\left(h_2(H(M)||i) \cdot \prod_{j=1}^s u_j^{m_{ij}} \right)^a$ which needs to be performed only by the simulator. This is accomplished as follows. First, a random exponent $\rho_i \in \mathbb{Z}_p^*$ is chosen. Since the values u_j and m_{ij} are known and assuming the random oracle model, the simulator programs the random model such that:

$$h_2(H(M)||i) := g^{\rho_i} \cdot \left(\prod_{j=1}^s u_j^{m_{ij}} \right)^{-1}. \quad (13)$$

This implies that the signature is equal to $(g^a)^{\rho_i}$ which can be computed since ρ_i is known and g^a part of the given CDH triple.

Note that an attacker may request the evaluation of h_2 prior running Setup. To enable a consistent simulation, we make use of the fact that both, h_2 and H , are simulated. That is, whenever the adversary requests an evaluation of $h_2(x)$ for some input x , the simulator first parses x as $x = s||i$ and checks if s is the output of an H -query made before. If this is the case, the simulator (that also simulates H) knows an input M such that $s = H(M)$ and hence knows the values m_{ij} to compute (13).

A further change affects how the values u_j are computed. By definition, it holds that $u_j = h_1(H(M)||j)$. Again, the simulator exploits that he simulates the random oracle and sets $h_1(H(M)||j) := g^{\alpha_j} \cdot h^{\beta_j}$ for some randomly chosen values α_j, β_j . Here, the simulation of h_1 does not require to know M .

Now, if $(\sigma', \{\mu'_j\})$ denotes the response given by the adversary and $(\sigma, \{\mu_j\})$ the correctly-built response, then one can deduce from $\sigma \neq \sigma'$ that there exists at least one index j such that $\mu'_j \neq \mu_j$. Let $\Delta_j := \mu'_j - \mu_j$ for each j . We obtain $e(\sigma' \cdot \sigma^{-1}, g) = e(\prod_j u_j^{\Delta_j}, g)$, which can be rewritten to:

$$e(\sigma' \cdot \sigma^{-1} \cdot (g^a)^{-\sum_j \alpha_j \cdot \Delta_j}, g) = e(h, g^{a+b+r_v})^{\sum_j \beta_j \cdot \Delta_j}, \quad (14)$$

where g^{a+b+r_v} represents the current combined public key. Note that $\sum_j \beta_j \cdot \Delta_j \neq 0$ with high probability as the coefficients β_j have been chosen uniformly at random and independently. It follows that:

$$h^a = \left(\sigma' \cdot \sigma^{-1} \cdot (g^a)^{-\sum_j \alpha_j \cdot \Delta_j} \right)^{(\sum_j \beta_j \cdot \Delta_j)^{-1}} \cdot h^{-b-r_v}$$

is a solution to the given CDH instance. Recall that r_v has been chosen by the simulator and b is given to the simulator due to the KOSK model.

As an additional step in our scheme, the proof of possession pop for the public key $\text{pk}_v = g^{a+r_v}$ needs to be provided. As the exponent a is unknown, this cannot be directly computed by the simulator. However, recall that $\text{pop} = h_0(\langle \text{pk}_v \rangle)^{\text{sk}_v}$ where $\text{sk}_v = a + r_v$. Therefore, the simulator chooses a random integer $\rho \in \mathbb{Z}_p$ and computes g^ρ and pk_v^ρ . Obviously, the second value is a valid signature of the first. Hence, we re-use the trick of reprogramming the hash and simply set $h_0(\langle \text{pk}_v \rangle)$ to g^ρ and pop to pk_v^ρ .

The final Game 2 is the same as Game 1 with a single difference. Similar to before, the environment tracks all messages and observes all POR-instances. We know already with overwhelming probability, that for all responses $(\sigma, \{\mu_i\})$ which are

Parameter	Default Value
File size	64 MB
Group size lpl	224 bits
RSA modulus size	2048 bits
Elliptic Curve (BLS)	PBC Library Curve F
Challenge size ℓ	100

Table 1: Default parameters used in evaluation.

accepted, it holds that σ is correctly computed. Now, the environment checks for all accepted responses whether the aggregated messages μ_i are equal to correctly generated, i.e., are equal to $\sum_{(i,v_i) \in Q} v_i m_{ij} \in \mathbb{Z}_p$, $1 \leq j \leq s$ (see also Equation 10). If this is not the case, the environment aborts the security game. Similar to [31] and the explanations given for Game 1, one can show that this can be used to solve the discrete logarithm problem in G —which would contradict the assumption that this problem is hard in G .

The core idea is the following. Given $g, h \in G$, the simulator honestly generates all values with the exception of the values u_j that are again set to $g^{\alpha_j} \cdot h^{\beta_j}$. Now, assume that the adversary succeeds with a response $(\sigma', \{\mu'_j\})$ instead of the correct response $(\sigma, \{\mu_j\})$. We know already that $\sigma = \sigma'$ needs to hold with high probability. Assume that $\Delta_j := \mu'_j - \mu_j$ for all j . From the fact that both responses are accepted, it follows that $\prod_j u_j^{\mu'_j} = \prod_j u_j^{\mu_j}$, which implies the following:

$$1 = \prod_j u_j^{\Delta_j} = g^{\sum_j \alpha_j \cdot \Delta_j} \cdot h^{\sum_j \beta_j \cdot \Delta_j}. \quad (15)$$

Again, it holds that $\sum_j \beta_j \cdot \Delta_j \neq 0$ with high probability. In such case, we solve the discrete logarithm problem by

$$h = \left(g^{\sum_j \alpha_j \cdot \Delta_j} \right)^{(\sum_j \beta_j \cdot \Delta_j)^{-1}}. \quad (16)$$

Thus, we conclude that the existence of such an attacker would contradict the hardness assumption with respect to the discrete logarithm problem.

Summing up, a prover can only present responses that are accepted by the verifier if all values are correctly built, i.e., as displayed in Equations 9 and 10—thus concluding the proof of Theorem 1.

4. IMPLEMENTATION & EVALUATION

In this section, we evaluate an implementation of SPORT within a realistic cloud setting and we compare the performance of SPORT to the RSA-based and BLS-based POR schemes due to [30]. Note that all three POR schemes are publicly-verifiable.

4.1 Implementation Setup

We implemented a prototype of SPORT in Java. In our implementation, we relied on the Backblaze library [3] for constructing the erasure codes (instantiated using Reed-Solomon coding). For a baseline comparison, we also implemented the public POR (with its two BLS and RSA variants) schemes (see Appendix A for a description of these POR schemes). Here, we relied on SHA-256, the Java built-in random number generator, HMAC based on SHA-256, 2048-bit RSA modulus, and the JPBC library [2] (based on the PBC cryptographic library [1]) to implement BLS signatures. Table 1 summarizes the default parameters assumed in our setup.

We deployed our implementations on a private network consisting of two 24-core Intel Xeon E5-2640 with 32GB of RAM. In our network, the communication between various machines was bridged using a 100 Mbps switch. The storage server was running on one of the 24-core Xeon E5-2640 machine, whereas the clients were co-located on the second 24-core Xeon E5-2640 machine; this ensures




Figure 2: Performance evaluation of SPORT in comparison to the POR-SW schemes of [30]. Each data point in our plots is averaged over 10 independent runs. Notice that, due to their small size, 95% confidence intervals are omitted from these plots.

a fair comparison between the overhead incurred on the users and on the server for the different schemes.

To emulate a realistic Wide Area Network (WAN), we shape all traffic exchanged on the networking interfaces with a Pareto distribution with a mean of 20 ms and a variance of 4 ms [18]. In our setup, each client invokes an operation in a closed loop, i.e., a client may have at most one pending operation. Prior to the setup phase, each client disperses his files with a (9,12) code. Similar to [6, 7], we assume that clients query for the availability of $\ell = 100$ randomly selected blocks in the POR challenge phase.

When implementing SPORT, we spawned multiple threads on the cloud machine, each thread corresponding to a unique audit performed on behalf of a client. Each data point in our plots is averaged over 10 independent measurements; where appropriate, we include the corresponding 95% confidence intervals.

4.2 Evaluation Results

Before evaluating the performance of SPORT, we start by analyzing the impact of the block size on the latencies incurred in the verification of POR in the BLS POR scheme of [30]. Our results (cf. Figure 5 in the Appendix) show that modest block sizes of 8 KB yield the most balanced performance, on average, across all procedures. Throughout the rest of our evaluation, we therefore set the block size to 8 KB.

POR protocol performance: In Figure 2(a), we evaluate the time required by the user to verify a single POR in SPORT, when compared to the public SW schemes. As expected, our findings show that users in SPORT and the BLS SW scheme witness comparable performance since the (core) POR verification process in both these schemes are similar. One major difference between SPORT and the BLS SW scheme of [30] is that users need to seldomly update the private key in order to be able to verify the POR responses issued by the provider. This only occurs whenever a new tenant deduplicates the same file of interest in between two consecutive POR performed by users storing the same file (cf. Section 3.2). As shown in Figure 2(b), this process incurs negligible overhead on users; for example, users can update their private keys based on the contributions of 10 additional tenants in around 3 seconds.

Figure 2(a) also shows that the latency witnessed in the POR verification process in SPORT and the BLS-SW scheme is almost double when compared to that of the RSA-SW scheme. This is the case due to the fact that (i) pairing operations (to verify tags) are considerably more expensive than RSA exponentiations, and (ii) BLS-based schemes result in a larger number of sectors per block,

which incurs additional computations in the store, verify, and prove procedures.

That said, the fact that SPORT/BLS-SW exhibits more sectors per block allows the cloud provider to efficiently parallelize the computations required to answer a given POR challenge. In Figure 4, we evaluate the latency with respect to the throughput exhibited by the provider in SPORT when issuing POR responses (i.e., when computing Equation 10). We measure throughput as follows: we require that each client performs back to back POR verification operations; we then increase the number of clients in the system until the aggregated throughput attained by all clients is saturated. The peak throughput is then computed as the maximum aggregated number of POR responses that can be issued by the storage server within a period of time. Our results show that the provider can perform up to 3700 POR responses per SPORT within 100 seconds, resulting in a peak throughput 37 operations (POR responses) per second. In the RSA-SW scheme, a maximum of 500 POR responses can be performed within the same time lapse, resulting in a modest peak throughput of 10 operations per second. In other words, the provider can scale approximately 4 times better in SPORT when compared to the RSA-SW scheme of [30].

Recall that RSA-based tags are considerably larger in size when compared to their BLS counterparts (almost 5 times larger). Moreover, RSA does not exhibit key homomorphism and therefore does not allow tag sharing. As shown in Figure 1, this incurs considerable storage overhead on the provider when multiple tenants deduplicate the same file. For example, assuming a file size of 64 MB, storing the tags pertaining to 5 tenants in RSA-SW results in an additional overhead of 10 MB; this overhead reduces to only 0.2 MB in SPORT.

Store* protocol performance: In Figure 2(c), we evaluate the time incurred on the provider when verifying the tenant's tags in the store and update procedures. Here, we also recall the time incurred on the tenants in the store procedure for a baseline comparison. Our results show that the overhead of store (in verifying the first tags uploaded for a given file) incurred on the provider slightly improves that incurred on the tenants in spite of the fact that BLS signature generation (performed by clients) is considerably faster than the signature verification process (performed by the cloud). This is the case since SPORT leverages the homomorphism of the underlying signature scheme in order to verify all the tags with a single batch verification (cf. Section 3.2). As a result of this optimization, the overhead on the provider in this case is approximately 4% lower than that on the tenants.




Figure 3: Latency incurred in Store w.r.t. the file size.

Notice that this overhead is even further reduced in the Store* procedure whenever an additional tenant shares the same file/tags. In this case, the provider does not need to partially create the tags (i.e., to reconstruct w_i as given by Equation 5). Instead, S only needs to verify that Equation 8 holds. This verification incurs negligible overhead on S ; as shown in Figure 2(c), this process is almost 40 times faster than that required to construct the tags by the tenants.

Recall that the Store* protocol is performed only once by clients (upon upload of a file that is already stored). By doing so, SPORT trades the *permanent* storage costs associated with the additional storage of file tags, with the *one-time* use of their computational resources to update the tags—thus reducing the total costs borne on users.

Store protocol performance: In Figure 3, we measure the latency incurred in the store procedure in SPORT compared to the RSA-SW and BLS-SW schemes of [30]. As expected, the latency incurred in store increases almost linearly with the file size in all POR schemes. Our results show that the store latency incurred in SPORT is the same as that incurred in the BLS SW of [30]. Indeed, both schemes involve the same procedure for creating tags based on BLS signatures; in this respect, SPORT does not incur any additional overhead on the user when compared to the basic BLS SW store procedure. On the other hand, although BLS signatures are faster to generate by the user when compared to their RSA counterpart, BLS-based tags result in considerably smaller sector sizes than the RSA-based scheme, which results in a larger number of sectors per block, and thus in additional computations to create each block tag. This explains the reason why the store latency exhibited by the users of SPORT is almost twice as much larger than that of the RSA-SW protocol. We, however, point out that the store procedure is only performed once per user—after which users simply execute the POR protocol to verify the integrity of their outsourced data.

5. RELATED WORK

Single-tenant POR: Juels and Kaliski [25] introduce a single-tenant POR scheme, which relies on indistinguishable blocks, hidden among regular file blocks in order to detect data modification by the server. The authors also propose a Merkle-tree construction for constructing public POR, which can be verified by any external party without the need of a secret key. Bowers *et al.* [14] propose various improvements to the original single-tenant POR in [25], which tolerates a Byzantine adversarial model. Shacham and Waters [30] propose private-key-based and public-key-based (single-tenant) POR schemes which utilize homomorphic authenticators to yield compact proofs.

In [7], Ateniese *et al.* introduce a variant of POR called *proofs of data possession* (PDP). It supports an unbounded number of challenge queries and enables public verifiability of the PDP. Unlike




Figure 4: Latency incurred in POR w.r.t. the throughput as witnessed by the provider.

other POR schemes, this instantiation does not offer extractability guarantees. This proposal was later extended in [8] to address dynamic writes/updates from the clients. Cash *et al.* [16] propose a dynamic POR scheme which relies on oblivious RAM protocols. In [32], Shi *et al.* propose a dynamic POR scheme that considerably improves the performance of [16] by relying on a Merkle hash tree.

Secure Data De-duplication in Multi-tenant Settings: In [23], Harnik *et al.* describe a number of threats posed by client-side data deduplication, in which an adversary can learn if a file is already stored in a particular cloud by guessing the hashes of predictable messages. This leakage can be countered using Proofs of Ownership schemes (PoW) [17, 22], which enable a client to prove it possesses the file in its entirety. PoW are inspired by Proofs of Retrievability and Data Possession (POR/PDP) schemes [7, 30], with the difference that PoW do not have a pre-processing step at setup time.

Douceur *et al.* [19] introduced the notion of convergent encryption, a type of deterministic encryption in which a message is encrypted using a key derived from the plaintext itself. Convergent encryption is not semantically secure [10] and only offers confidentiality for messages whose content is unpredictable. To remedy this, a number of proposals introduce the notion of oblivious server-aided encryption to perform data deduplication scheme [5, 9]; here, the encryption key is obliviously computed based on the hash of the file and the private key of the assisting server.

In [33], Vasilopoulos *et al.* extend the scheme in [5] to allow the deduplication of the files and the related tags. In [34], Zheng and Xu propose the notion of Proof of Storage with Deduplication (POSD). Here, POSD is introduced as the combination of PDP/POR and PoW schemes. As we mention in this paper, such schemes do not resist collusion between malicious tenants and the cloud provider; namely, a cloud provider which has access to the secret material (e.g., leaked by the first user who created the tags) can always construct correct responses to the challenges issued by the clients—even if the cloud provider deletes the outsourced data. As far as we are aware, SPORT is the first secure POR instantiation which addresses multi-tenancy and data deduplication.

6. CONCLUSION

In this paper, we introduced the notion of multi-tenant proofs of retrievability (MTPOR), an extension of the traditional single-tenant POR concept, and proposed an efficient instantiation dubbed SPORT. We implemented a prototype based on SPORT, and evaluated its performance in a realistic cloud setting. Our results show that our proposal incurs minimal storage overhead on the cloud provider without deteriorating the performance witnessed by tenants and the cloud provider when compared to existing publicly-verifiable POR schemes. SPORT is provably secure in the random oracle model

under the computational Diffie-Hellman problem, assuming static corruptions. In this respect, we see the analysis of the security of SPORT under adaptive corruption as an interesting open question.

We argue that SPORT provides an important stepping stone to reconcile existing cloud integrity and security primitives with functional requirements (such as resource sharing and multi-tenancy) in the cloud. Namely, SPORT provides considerable incentives (i) for end-users to obtain guarantees about the retrievability of their files in the cloud, and (ii) for cloud providers to offer differentiated services while preserving the efficiency of their storage system.

7. ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers for their valuable feedback and comments. This work was partly supported by the TREDISEC project (G.A. no 644412), funded by the European Union (EU) under the Information and Communication Technologies (ICT) theme of the Horizon 2020 (H2020) research and innovation programme.

References

- [1] PBC Library. <http://crypto.stanford.edu/pbc/>, 2007.
- [2] JPBC:Java Pairing-Based Cryptography Library. <http://gas.dia.unisa.it/projects/jpbc/#.U3HBFfna5cY>, 2013.
- [3] Backblaze Open Sources Reed-Solomon Erasure Coding Source Code. <https://www.backblaze.com/blog/reed-solomon/>, 2015.
- [4] Google loses data after lightning strikes. <http://money.cnn.com/2015/08/19/technology/google-data-loss-lightning/>, 2015.
- [5] F. Armknecht, J. Bohli, G. O. Karame, and F. Youssef. Transparent data deduplication in the cloud. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, pages 886–900, 2015.
- [6] F. Armknecht, J.-M. Bohli, G. O. Karame, Z. Liu, and C. A. Reuter. Outsourced proofs of retrievability. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, CCS '14*, pages 831–843, New York, NY, USA, 2014. ACM.
- [7] G. Ateniese, R. C. Burns, R. Curtmola, J. Herring, L. Kissner, Z. N. J. Peterson, and D. X. Song. Provable data possession at untrusted stores. In *ACM Conference on Computer and Communications Security*, pages 598–609, 2007.
- [8] G. Ateniese, R. D. Pietro, L. V. Mancini, and G. Tsudik. Scalable and efficient provable data possession. *IACR Cryptology ePrint Archive*, 2008:114, 2008.
- [9] M. Bellare, S. Keelveedhi, and T. Ristenpart. DupLESS: Server-aided encryption for deduplicated storage. In *Proceedings of the 22nd USENIX Conference on Security, SEC'13*, pages 179–194, Berkeley, CA, USA, 2013. USENIX Association.
- [10] M. Bellare, S. Keelveedhi, and T. Ristenpart. Message-locked encryption and secure deduplication. In Johansson and Nguyen [24], pages 296–312.
- [11] A. Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-diffie-hellman-group signature scheme. In *International Workshop on Public Key Cryptography*, pages 31–46. Springer, 2003.
- [12] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the weil pairing. *J. Cryptology*, 17(4):297–319, 2004.
- [13] K. D. Bowers, A. Juels, and A. Oprea. HAIL: a high-availability and integrity layer for cloud storage. In *ACM Conference on Computer and Communications Security*, pages 187–198, 2009.
- [14] K. D. Bowers, A. Juels, and A. Oprea. Proofs of retrievability: theory and implementation. In *CCSW*, pages 43–54, 2009.
- [15] K. D. Bowers, M. van Dijk, A. Juels, A. Oprea, and R. L. Rivest. How to tell if your cloud files are vulnerable to drive crashes. In *ACM Conference on Computer and Communications Security*, pages 501–514, 2011.
- [16] D. Cash, A. K  p  , and D. Wichs. Dynamic Proofs of Retrievability via Oblivious RAM. In Johansson and Nguyen [24], pages 279–295.
- [17] R. Di Pietro and A. Sorniotti. Boosting efficiency and security in proof of ownership for deduplication. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security, ASIACCS '12*, pages 81–82, New York, NY, USA, 2012. ACM.
- [18] D. Dobre, G. Karame, W. Li, M. Majuntke, N. Suri, and M. Vukolić. Powerstore: Proofs of writing for efficient and robust storage. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, CCS '13*, pages 285–298, New York, NY, USA, 2013. ACM.
- [19] J. R. Douceur, A. Adya, W. J. Bolosky, D. Simon, and M. Theimer. Reclaiming space from duplicate files in a serverless distributed file system. In *ICDCS*, pages 617–624, 2002.
- [20] E. Ghosh, O. Ohrimenko, D. Papadopoulos, R. Tamassia, and N. Triandopoulos. Zero-knowledge accumulators and set operations. *Cryptology ePrint Archive*, Report 2015/404, 2015. <http://eprint.iacr.org/2015/404>.
- [21] T. T. W. Group. The notorious nine: Cloud computing top threats in 2013. Report, Cloud Security Alliance, February 2013.
- [22] S. Halevi, D. Harnik, B. Pinkas, and A. Shulman-Peleg. Proofs of ownership in remote storage systems. In *Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS '11*, pages 491–500, New York, NY, USA, 2011. ACM.
- [23] D. Harnik, B. Pinkas, and A. Shulman-Peleg. Side channels in cloud services: Deduplication in cloud storage. *IEEE Security & Privacy*, 8(6):40–47, 2010.
- [24] T. Johansson and P. Q. Nguyen, editors. *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*. Springer, 2013.
- [25] A. Juels and B. S. K. Jr. PORs: Proofs Of Retrievability for Large Files. In *ACM Conference on Computer and Communications Security*, pages 584–597, 2007.
- [26] D. T. Meyer and W. J. Bolosky. A study of practical deduplication. In *Proceedings of the 9th USENIX Conference on File and Storage Technologies, FAST'11*, pages 1–11, Berkeley, CA, USA, 2011. USENIX Association.
- [27] D. T. Meyer and W. J. Bolosky. A study of practical deduplication. *Trans. Storage*, 7(4):14:1–14:20, Feb. 2012.
- [28] L. Reyzin and S. Yakubov. Efficient asynchronous accumulators for distributed pki. *Cryptology ePrint Archive*, Report 2015/718, 2015. <http://eprint.iacr.org/2015/718>.
- [29] T. Ristenpart and S. Yilek. The power of proofs-of-possession: Securing multiparty signatures against rogue-key attacks. In *Advances in Cryptology - EUROCRYPT 2007, 26th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Barcelona, Spain, May 20-24, 2007, Proceedings*, volume 4515 of *Lecture Notes in Computer Science*, pages 228–245. Springer, 2007.
- [30] H. Shacham and B. Waters. Compact Proofs of Retrievability. In *ASIACRYPT*, pages 90–107, 2008.
- [31] H. Shacham and B. Waters. Compact Proofs of Retrievability. *Cryptology ePrint Archive*, Report 2008/073, 2008. <http://eprint.iacr.org/>.
- [32] E. Shi, E. Stefanov, and C. Papamanthou. Practical dynamic proofs of retrievability. In A.-R. Sadeghi, V. D. Gligor, and M. Yung, editors, *ACM Conference on Computer and Communications Security*, pages 325–336. ACM, 2013.
- [33] D. Vasilopoulos, M.   nen, K. Elkhyaoui, and R. Molva. Message-locked proofs of retrievability with secure deduplication. In *Proceedings of the 2016 ACM on Cloud Computing Security Workshop, CCSW '16*, pages 73–83, New York, NY, USA, 2016. ACM.
- [34] Q. Zheng and S. Xu. Secure and efficient proof of storage with deduplication. In *Proceedings of the Second ACM Conference on Data and Application Security and Privacy, CODASPY '12*, pages 1–12, New York, NY, USA, 2012. ACM.

APPENDIX

A. PUBLIC SW POR

To enable any entity which does not necessarily possess secrets to verify a POR, Shacham and Waters [30] propose two publicly verifiable POR schemes based on BLS signatures [12] and RSA, respectively.

Public BLS SW Scheme: The setup phase requires choosing a

group G with support \mathbb{Z}_p , and a computable bilinear map $e : G \times G \rightarrow G_T$. Additionally, the user chooses a private key $x \in \mathbb{Z}_p$, the corresponding public key $v = g^x \in G$ along with another generator $u \in G$. In the storage phase, a signature on each block i is computed $\sigma_i = \left(H(i) \prod_{j=1}^s u_j^{m_{ij}} \right)^x$. For verification, the challenge query Q is generated similarly to PSW and sent to the prover who computes:

$$\sigma \leftarrow \prod_{(i, \nu_i) \in Q} \sigma_i^{\nu_i} \in G, \quad \mu_j \leftarrow \sum_{(i, \nu_i) \in Q} \nu_i m_{ij} \in \mathbb{Z}_p.$$


These values are sent to the verifier who checks that:

$$e(\sigma, g) \stackrel{?}{=} e \left(\prod_{(i, \nu_i) \in Q} H(i)^{\nu_i} \cdot \prod_{j=1}^s u_j^{\mu_j}, v \right). \quad (17)$$


Public RSA SW Scheme: The public RSA-based SW scheme is similar to its public counterpart. Here, the block authenticator can be computed by $\sigma_i = (H(i)u^{m_i})^d \bmod N$, where d is the private key of the user. The cloud response is calculated similarly to the public BLS SW scheme. Given the public RSA key e , the verification unfolds as follows:

$$\sigma^e \stackrel{?}{=} \prod_{(i, \nu_i) \in Q} H(i)^{\nu_i} u^{\mu} \bmod N.$$


B. IMPACT OF BLOCK SIZE ON POR PERFORMANCE



(a) Latency in store of SPORT w.r.t. to the block size.



(b) User verification latency in the Proof Verification procedure of SPORT w.r.t. to the block size.



(c) Proof generation latency in SPORT.

Figure 5: Performance of SPORT with respect to the block size.