

An abridged version of this paper appears in the proceedings of CRYPTO 2024. This is the full version.

# A Formal Treatment of End-to-End Encrypted Cloud Storage

Matilda Backendal<sup>1</sup>, Hannah Davis<sup>2</sup>, Felix Günther<sup>3</sup>,  
Miro Haller<sup>4</sup>, and Kenneth G. Paterson<sup>1</sup>

<sup>1</sup>Department of Computer Science, ETH Zurich, Zurich, Switzerland

<sup>2</sup>Seagate Technology, Minneapolis, USA

<sup>3</sup>IBM Research Europe – Zurich, Zurich, Switzerland

<sup>4</sup>Department of Computer Science & Engineering, University of California San Diego, USA

## Abstract

Users increasingly store their data in the cloud, thereby benefiting from easy access, sharing, and redundancy. To additionally guarantee security of the outsourced data even against a server compromise, some service providers have started to offer end-to-end encrypted (E2EE) cloud storage. With this cryptographic protection, only legitimate owners can read or modify the data. However, recent attacks on the largest E2EE providers have highlighted the lack of solid foundations for this emerging type of service.

In this paper, we address this shortcoming by initiating the formal study of E2EE cloud storage. We give a formal syntax to capture the core functionality of a cloud storage system, capturing the real-world complexity of such a system’s constituent interactive protocols. We then define game-based security notions for confidentiality and integrity of a cloud storage system against a fully malicious server. We treat both selective and fully adaptive client compromises. Our notions are informed by recent attacks on E2EE cloud storage providers. In particular we show that our syntax is rich enough to capture the core functionality of MEGA and that recent attacks on it arise as violations of our security notions. Finally, we present an E2EE cloud storage system that provides all core functionalities and that is both efficient and provably secure with respect to our selective security notions. Along the way, we discuss challenges on the path towards bringing the security of cloud storage up to par with other end-to-end primitives, such as secure messaging and TLS.

**Keywords:** end-to-end encryption · cloud storage · provable security

# Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Our Contributions . . . . .	4
1.2	On Modeling Choices . . . . .	5
1.3	Related Work . . . . .	6
<b>2</b>	<b>Preliminaries</b>	<b>7</b>
<b>3</b>	<b>A Framework for Cloud Storage</b>	<b>8</b>
3.1	Syntax . . . . .	9
3.2	The Protocols . . . . .	10
3.3	Correctness . . . . .	12
<b>4</b>	<b>Client-to-Client Security of Cloud Storage</b>	<b>15</b>
4.1	Client-to-Client Confidentiality . . . . .	16
4.2	Client-to-Client Integrity . . . . .	20
4.3	Designing a Security Model for Cloud Storage: UC vs. Game-Based . . . . .	22
4.4	Extensions, Limitations and Rationale of the Model . . . . .	22
<b>5</b>	<b>Formally Breaking MEGA</b>	<b>25</b>
<b>6</b>	<b>A Provably Secure Cloud Storage Scheme</b>	<b>27</b>
6.1	Building Blocks . . . . .	27
6.2	Scheme Overview . . . . .	28
6.3	Detailed Description . . . . .	29
6.4	Proving CSS Secure . . . . .	30
6.5	Limitations and Extensions of CSS . . . . .	35
<b>7</b>	<b>Conclusion and Future Work</b>	<b>36</b>
	<b>References</b>	<b>37</b>
<b>A</b>	<b>Details of Formalizing (the Attacks on) MEGA</b>	<b>42</b>
A.1	Formalizing MEGA as a Cloud Storage Scheme . . . . .	42
A.2	Formal Attacks against MEGA . . . . .	47
<b>B</b>	<b>Building Block Details</b>	<b>50</b>
B.1	Standard Building Blocks . . . . .	50
B.2	OPRF Building Block . . . . .	51
<b>C</b>	<b>Proofs of Client-To-Client Security for CSS</b>	<b>52</b>
C.1	Proof of Theorem <a href="#">6.1</a> . . . . .	52
C.2	Proof of Theorem <a href="#">6.2</a> . . . . .	63

# 1 Introduction

The *cloud*, which outsources data storage and computation to third party services, is one of the prevalent computing paradigms today. Our focus in this paper is on *cloud storage*. Aimed at consumers and enterprises alike, cloud storage relieves users of worrying about backups and data availability, and allows the data to be accessed from anywhere. But what about other security properties, beyond availability? Users should be rightly concerned about the confidentiality and integrity of their data too. Today the largest service providers such as Google and Microsoft expect users to trust them to look after these aspects on their behalf; data can be cryptographically protected, but using keys under the control of the service provider. If the providers are ever compromised (or compelled to provide access by law enforcement agencies) then all security guarantees may be voided.

However, there is an alternative approach for cloud storage, often referred to as End-to-End Encryption (E2EE). Here, encryption is carried out by the users themselves – under keys that they control – before uploading files to the cloud. If properly executed, this results in cryptographic confidentiality and authenticity guarantees: The service provider may learn some metadata about files, but nothing about the data content, nor can they modify individual files in an undetectable manner, thanks to the integrity protection provided by the encryption. As a consequence, security of users’ data is guaranteed even in the event of a compromise of the service. Ideally, these security guarantees should continue to hold even against a fully malicious service provider, removing the need for users to trust the providers with anything but the availability of their data.

A particular challenge for E2EE cloud storage is that key management is delegated from the service provider to the individual user. But users, as humans, are not particularly good at memorizing the cryptographic key material necessary to access their data from anywhere. For this reason, security in such services is typically bootstrapped from a user-memorable secret such as a password. The reliance on passwords brings fundamental security limitations to E2EE cloud storage, since a malicious server can always carry out variants of dictionary attacks. This risk can be partially obviated with strong password policies, trading convenience for security.

The lack of trust in the service provider also brings other challenges for E2EE cloud storage services. A primary feature target is *file sharing*: the capacity for a file’s owner to share access to the file with selected other users. Services offering file sharing face another difficulty related to key management. For the recipient to access the shared file, they need the relevant decryption keys. But these keys cannot be sent unprotected via the service provider if the provider is not trusted. Consequently, file sharing is typically done by encrypting the file key under the authenticated public key of the recipient, or via some out-of-band mechanism.

Despite the challenges associated to combining strong security with usability, E2EE encrypted cloud storage is on the rise. Prominent providers include MEGA [40], Nextcloud [48] and recently also Apple, who added optional support for E2EE to iCloud in 2023 [4]. To illustrate their scale, MEGA claims to have 300 million users and store more than 150 billion files, corresponding to 1000 petabytes of data [40]. Smaller E2EE providers (such as icedrive [42], pCloud [1], Seafile [43] and Tresorit [52]) also exist in abundance, with tens of millions of users collectively.

This development should delight privacy-minded users. However, recent analyses have thrown doubt on the security claims of MEGA and Nextcloud [6, 30, 3, 2] – the two largest providers of E2EE cloud storage by default – finding multiple attacks on their E2EE security. While both providers patched their systems quickly against the specific attacks, they still rely on complex, *ad hoc* cryptographic designs that lack any formal security guarantees. The MEGA case is particularly interesting – with so much data encrypted under keys known only to the users, migrating to more secure cryptographic mechanisms proved to be infeasible, since re-encryption would require user

participation. Hence, MEGA initially adopted a minimal approach to patching after the initial attacks in [6], which was quickly found to be insufficient [3], leading to a second round of more complex patching. The consequences for Nextcloud were even more dire: they were forced to turn off their E2EE file sharing service altogether until a full redesign could be carried out, constituting a significant feature regression [2].

These examples highlight the value of starting with a formally analyzed cryptographic design. The security of the other above-mentioned providers is unknown, but all rely on proprietary designs that are yet to be analyzed. Overall, our confidence in the security of E2EE cloud storage lags far behind that of other E2EE applications such as messaging and browsing, which have seen significantly more design input and analysis from the cryptographic research community.

## 1.1 Our Contributions

In this paper, we initiate the formal study of E2EE cloud storage. Our goal is to construct a solid foundation on which we can build our understanding of its security, as a first step towards bringing it up to par with other E2EE applications. We are driven by practicality – we want E2EE cloud storage systems that are as efficient as those deployed today and with the same functionality. We hence accept passwords as the basis of security. At the same time, we want to be able to formally analyze the systems – we want security models that are rich enough to capture all the key features of E2EE cloud storage, whilst still being tractable enough to use for proofs. We offer four main contributions towards this goal:

1. We give a formal syntax capturing the core functionality of E2EE cloud storage. Our syntax encompasses user registration and authentication, as well as upload, download and sharing of files. Each of these operations is modeled as a stateful, potentially multi-round interactive protocol executed between a client and a server. For file sharing, we modularize the out-of-band component, allowing it to be instantiated in practice in a variety of different ways (e.g., by link sharing or the use of a PKI, if available).
2. We define game-based security notions for E2EE cloud storage, which we refer to as Client-to-Client (C2C) security to emphasize that the end points are (user) clients. Our notions hence focus on the setting where the adversary is a malicious service provider aiming to subvert data confidentiality or file integrity. Our C2C models allow the adversary to interact in a fine-grained manner with the clients of honest users. For example, it can execute a single step of a protocol at a time (without running it to completion), arbitrarily interleave protocol runs across users and functions, as well as deviate in any way from the honest behavior of the party it emulates. The model additionally supports compromise of honest users. In tackling this complexity, we took inspiration from the key exchange literature. We consider both selective and adaptive versions of our C2C security notions; in the former, the set of compromised users is specified up front by the adversary, while in the latter, the adversary can compromise users in a fully adaptive manner.
3. We put our model to the test and use it to formally capture some of the recent attacks on MEGA [6, 3, 30]. We thereby show that our framework is rich enough to model the complexity of a deployed system, and that our notions capture practical attacks.
4. We present a provably secure E2EE cloud storage system called CSS. Our construction makes use of standard components: an oblivious pseudorandom function (OPRF) called 2HashDH [32, 34] is used to convert a user’s password into a user-specific root key, a pseudo-random function is then used to derive separate keys for (key) encryption and authentication

from the root key. Finally, a nonce-based AEAD scheme is used to build a key hierarchy (which allows password-rotation without file re-encryption): the key encryption key is used to encrypt a user-specific master key, which in turn encrypts randomly sampled file keys, each of which protect a single file. We prove the confidentiality and integrity of CSS in our selective C2C model, achieving bounds that reflect the unavoidable brute-force password attacks by a malicious server. In principle, our CSS system, if implemented, would provide a fully functional E2EE cloud storage system with proven security guarantees.

Finally, we point to many open problems and topics for future work.

## 1.2 On Modeling Choices

Our model is a formal answer to the question “what security should an end-to-end encrypted cloud storage system provide?” At first glance, this seems to have a relatively simple solution, one that is outright stated in the name of the system. The term “end-to-end encryption” denotes confidentiality against a malicious service provider, and “cloud storage” suggests that users can access their files with a high degree of reliability, tamper-resistance, and portability.

However, the intuitive notion that first emerges from these constraints might look very different from our final definitions. In the abstract, a perfectly secure cloud storage system might register clients by their passwords, then take in client files and release them unchanged only to authorized clients, while revealing to the server only the number and length of files. We would like this perfect security definition to include strong guarantees, including a treatment of adaptive corruption and some protection against offline password-guessing attacks from external adversaries in the presence of an honest service provider. Without concern for undue complexity, we might even go further and also address advanced features like metadata hiding, key rotation, and forward security.

However, a good security definition also possesses two other properties: first, it should be achievable by at least one practical system, and second, it should accommodate convincing and legible proofs of security. It is in accomplishing these two goals that the “intuitive” notion described above breaks down irretrievably. Our security definitions are the result of many carefully weighed decisions and compromises that maximize feasibility, generality, and comprehensibility, while still preserving the spirit of that first intuition.

The first significant choice is our use of a game-based model. Many password-based cryptographic schemes prove security in simulation-based settings, most notably the universal composability (UC) framework [15]. Secure UC-realizable protocols emulate an “ideal functionality” that captures the intended behavior of a protocol, with only unavoidable deviations from this ideal. Because UC provides very strong security, not all ideal functionalities are realizable [18]. In particular, fully secure channels cannot be constructed without non-standard encryption primitives that are not likely to see practical adoption [17]. Cloud storage systems experience the same obstacle to UC realization as secure channels, and the UC technique used to build a “weakly-secure” channel from standard primitives does not apply due to the use of passwords as the root of trust.

A similar problem arises in the treatment of adaptive security. Because we empower our adversaries with decryption oracles, cloud storage, like encryption, requires attacks to identify their targeted ciphertexts in advance. In our CSS system, some of these ciphertexts appear in the user registration and authentication process to enable a practical key hierarchy, leading to so-called commitment issues when simulating these operations for honest users in the proof. Hence, in order for the reductions to hold, we need attackers to identify not only their target files but also the users that they will compromise in advance. Due to file sharing, the loss in security from the standard technique for avoiding this (guessing the choices of the adversary, also known as complexity

leveraging) would be exponential. The consequence is that we prove selective rather than adaptive security for CSS; efficiently achieving the latter remains an open challenge.

To minimize complexity in our game-based model, we restrict ourselves to the core set of features and security properties common to existing E2EE services. Even in the simplified setting, complications arise when we consider that real-world attacks may interleave the various multi-step subprotocols of a cloud storage system and attempt to exploit state confusion between simultaneous operations of many users on many files. This type of non-atomic interchange also appears to a lesser degree in many key exchange security models. A secure cloud storage system could in theory be used to instantiate a key exchange protocol (simply put the key in a shared file!), so it is not surprising that the complexity we encounter is (at least) as high as that of key exchange.

File sharing also adds significant complexity, as it allows interaction between honest and malicious users, allowing malicious users to learn some plaintext without using decryption oracles. Tracking this in the security model requires care. Furthermore, users may expect to be able to share files without relying on a trusted service provider to authenticate the recipient. We idealize this part of the system via an “out of band” (OOB) mechanism. This makes our model general enough to capture many different choices of implementing such a channel, such as link sharing (as in MEGA), using public key encryption (as in Nextcloud), or even sending keys over a secure messenger.

We do make one choice that increases complexity, in service of future results. In our C2C setting, a malicious server can eschew any steps a protocol takes to disambiguate clients. As a consequence of this, we could obtain equivalent levels of security for CSS by using a salted hash to convert user passwords into keys as we do by using a stronger primitive like an OPRF. But the OPRF and its security properties will come to the fore in the *external* adversary setting; there we expect CSS to limit password guesses to one per client login attempt (comparable to what one hopes to achieve in a PAKE protocol).

### 1.3 Related Work

The reader may wonder: is E2EE cloud storage not already a solved problem? The short answer is that while many schemes for cloud storage have been proposed in the literature, only some of them address the E2EE problem, and few of them come with formal security analysis.

Table 1 lists properties covered by representative work from different areas. All of these works propose cloud storage systems, in some cases with more advanced security features such as key rotation, forward security, master key recovery, or metadata hiding. However, none of them offer a formal security proof that models all core components of cloud storage. Many consider the server to be honest (but curious), do not model file sharing, or ignore the security impact of outsourcing key material to the server. One feature in particular emerges as unique to our work: the modeling of non-atomic, multi-step protocols, with adversary-controlled arbitrary interleaving of protocol steps in the security games. This adds significant complexity to our model, but in return captures the interactive nature of cloud storage operations in practice.

Companies like Apple and Meta have recently launched systems providing encrypted backups for user data with user-controlled encryption [4, 54]. There is no public specification or analysis of E2EE in iCloud, but the WhatsApp backup system was formally analyzed in [24] using a UC approach. The latter analysis does not model a fully-malicious server but relies on a (trusted) hardware security module on the WhatsApp server instead. Like our CSS scheme, the WhatsApp system uses an OPRF to derive keys from passwords. This usage prevents pre-computation attacks on user passwords by snapshot adversaries, as well as offline attacks by an external adversary, explaining its common appearance in their setting and ours. Backups by nature only pertain to a

	formal	share	pw	mal	m-sess	non-atomic	single	advanced
SiRiUs [27]	○	●	○	○	○	○	●	●
DepSky [13]	○	○	○	◐	○	○	○	●
Mylar [50]	◐	●	○	●	◐	○	●	●
Burnbox [53]	●	○	○	●	●	○	●	●
OKMS [36]	●	○	●	◐	○	○	○	●
PFS [5]	●	○	○	◐	○	○	●	●
Metal [22]	●	●	○	○	○	○	○	●
Titanium [21]	●	●	○	◐	○	○	○	●
DPaSE [23]	●	○	●	●	●	○	○	○
PBCS [20]	●	○	●	●	●	○	○	○
<i>This work</i>	●	●	●	●	●	●	●	○

Table 1: Representative work from different areas and the features they consider: formal security analysis (*formal*), file sharing (*share*), modeling password-based security (*pw*), considering malicious servers and users (*mal*), key management across sessions on multiple devices (*m-sess*), non-atomic operations and arbitrary interleavings (*non-atomic*), single server instead of requiring multiple servers, some of which are honest (*single*), and advanced security guarantees (*advanced*) such as hiding access patterns and metadata, forward security, or temporary access revocation. We denote missing, partially covered, and fully covered properties with ○, ◐, and ●, respectively.

single user, and hence do not need to address the complexity of file sharing.

The core idea of using an OPRF in these settings seems to date back to [25], which introduced the concept of a “PRF as a service”, and proposed various applications using Partially Oblivious PRFs (PO-PRFs), including enterprise password storage. Follow-up works propose password-hardened encryption and key management services [39, 38, 36, 20, 23]. These are related to our notion of E2EE cloud storage, but do not include the file sharing capability or key hierarchies to support efficient password changes. While some of them suggest cloud storage as an application, most do not specify in detail how files are encrypted. Moreover, none of these works seem to consider the malicious server setting. See also [32, 37] for the use of OPRFs in password-based authentication protocols. An excellent overview of OPRFs and their applications can be found in [19].

Closely related in (proof) techniques is the prior work on password-based encryption [9, 12]. While focused on simpler primitives, their techniques for capturing passwords in game-based security models assisted us in establishing the security achievable in password-based settings and in developing the proofs of our proposed system.

## 2 Preliminaries

**NOTATION.** If  $\mathbf{w}$  is a vector then  $|\mathbf{w}|$  is its length (the number of its coordinates) and  $\mathbf{w}[i]$  is its  $i$ -th coordinate. By  $\varepsilon$  we denote the empty string or vector. By  $x \parallel y$  we denote the concatenation of (bit-)strings  $x, y$ . Using a tuple  $(x, y)$  where a string is required indicates an unambiguous encoding of  $(x, y)$  into a string. If  $S$  is a finite set, then  $|S|$  denotes its size. We use  $x \leftarrow a$  for assigning the variable  $x$  to  $a$ , and let  $x \leftarrow^* S$  denote picking an element of  $S$  uniformly at random and assigning it to  $x$ . For conciseness, we use the bulk assignment  $x, y, z \leftarrow a$  as abbreviation for  $x \leftarrow a$ ,  $y \leftarrow a$ , and  $z \leftarrow a$ . Furthermore, we use pattern matching to assign variables to elements of tuples, e.g.,  $(x, (y, z)) \leftarrow (a, (b, c))$  assigns  $x, y$ , and  $z$  to  $a, b$ , and  $c$  respectively. For sets  $S_1$  and  $S_2$ , the

shorthand  $S_1 \stackrel{\cup}{\leftarrow} S_2$  denotes  $S_1 \leftarrow S_1 \cup S_2$ .

If  $A$  is an algorithm, we let  $y \leftarrow A^{O_1, \dots}(x_1, \dots)$  denote running  $A$  on inputs  $x_1, \dots$  with oracle access to  $O_1, \dots$ , and assigning the output to  $y$ . The output assignment of randomized algorithms is denoted with  $y \leftarrow_{\$} A^{O_1, \dots}(x_1, \dots)$ . Running time is worst case, which for an algorithm with access to oracles means across all possible replies from the oracles. We use  $\top$  (top) and  $\perp$  (bot) as special symbols, not in  $\{0, 1\}^*$ , to denote acceptance and rejection, respectively.

We use the notation  $T[k] \leftarrow v$  to denote storing the key-value pair  $(k, v)$  in table  $T$ . In case  $k = \perp$ , the lookup aborts. The keys are unique and reusing an existing key  $k$  overwrites the previous value stored in  $T[k]$ . Table entries are assumed to be implicitly initialized to  $\perp$ , such that the comparison  $T[k] = \perp$  is well-defined (and evaluates to true when no entry for  $k$  was stored). Furthermore, every table has the entry  $T[\varepsilon] = \varepsilon$  to map the empty index to an empty entry. We use  $k \in T$  as shorthand for checking whether  $T[k] \neq \perp$ . Let  $|T|$  denote the number of key-value pairs in the table. We use a dot (e.g.,  $T[(k, \cdot)]$ ) to denote a wildcard matching any element (here, any tuple where element one is  $k$ ).

We use the dot notation from object-oriented programming to access specific elements of a tuple when the context is clear. For instance, to store a password  $pw$  and an account ID  $aid$  in a state  $st$ , we write  $st \leftarrow (pw, aid)$  and access these values directly using  $st.pw$  and  $st.aid$ . Furthermore, we use this notation to set values, e.g.,  $st.sid \leftarrow sid$  appends a session ID  $sid$  to  $st$ . such that it now stores  $(pw, aid, sid)$ .

GAMES. We use the code-based game-playing framework of [11]. By  $\mathbf{G}_S^{\text{Sec-}b}(\mathcal{A})$ , we denote running the game for security experiment Sec, parameterized by scheme  $S$  and (optionally) bit  $b \in \{0, 1\}$ , with adversary  $\mathcal{A}$ . We denote the probability that the execution of game  $\mathbf{G}$  results in the output  $y$  as  $\Pr[\mathbf{G}(\mathcal{A}) \Rightarrow y]$ , and write just  $\Pr[\mathbf{G}]$  for  $\Pr[\mathbf{G}(\mathcal{A}) \Rightarrow 1]$  and  $\Pr[\mathbf{G}(\mathcal{A}) \Rightarrow \text{true}]$ . Our proofs are in the random oracle model [10]. In games, this is modeled via an oracle  $\text{RO}_i(v)$ , which implements lazy sampling for some random oracle  $H_i$  on input  $v$ .

In games, integer variables, set variables boolean variables and string variables are assumed initialized, respectively, to 0, the empty set  $\emptyset$ , the boolean false and  $\varepsilon$ . We use **global** to mark variables with global scope, i.e., accessible from all parts of the game, including oracles. The symbol  $\iota_{\text{mal}}$  is reserved for tracking actions by the adversary and is assumed not to be in  $\mathbb{N}$ .

### 3 A Framework for Cloud Storage

Our framework for cloud storage systems identifies seven core operations:

**Account registration (areg).** Register an account with the cloud storage server.

**Session authentication (auth).** Authenticate registered user and create a new (client) session to interact with the server over the following operations.

**File upload (put).** Upload a file to the server.

**File update (upd).** Replace an existing file on the server.

**File download (get).** Retrieve a file from the server.

**File share (share).** Prepare sharing an existing file with another user.

**Accept share (accept).** Accept the sharing of a file by another user.



Each of these operations generally involves an interaction between a client (session) and the cloud storage server, which we will capture as interactive, message-based protocols. In our framework, we explicitly treat the operations described above (and their expected security properties), but the notation is generic and can be extended to cover further operations one might be interested in when capturing systems that allow for, e.g., account deregistration, session deauthentication, or password recovery mechanisms. In the following, we will first introduce a generic syntax that allows us to capture all of the above protocols in a unified manner, then discuss the specific behavior and inputs of each protocol.

### 3.1 Syntax

Syntactically, we consider a cloud storage scheme  $\text{CS} = (\Pi_{\text{areg}}, \Pi_{\text{auth}}, \Pi_{\text{put}}, \Pi_{\text{upd}}, \Pi_{\text{get}}, \Pi_{\text{share}}, \Pi_{\text{accept}})$  to consist of the seven core protocols outlined above and described in detail in Section 3.2; see also Table 2 for their minimal input-output behavior. Each of these is a two-party interactive protocol executed between a client and a server, the formalities of which we define next.

**PROTOCOL STEPS.** The execution of protocol  $\Pi$  involves  $\Pi.\text{SN}$  protocol steps overall, alternating between client and server. We denote each round of the interaction as  $\Pi^{(C:i)}$  and  $\Pi^{(S:i)}$ , for step  $i$  of client and server, respectively. Syntactically, we write the  $i$ -th protocol step of a party  $P \in \{C, S\}$  as

$$(st_P, st_P^{\text{tmp}}, out_P, m_{out}) \leftarrow \$ \Pi^{(P:i)}(st_P, st_P^{\text{tmp}}, [in_P,] m_{in}),$$

where  $st_P$  and  $st_P^{\text{tmp}}$  is the local (persistent resp. temporary/protocol) state of party  $P$ ,  $in_P$  is the local input [present only in the first step,  $i = 1$ ],  $m_{in}$  and  $m_{out}$  is the incoming resp. outgoing message of the protocol,<sup>1</sup> and  $out_P$  is the local output of party  $P$ . As a convention, we set empty protocol inputs to  $\varepsilon$ .

For a first, rough example, a file download protocol  $\Pi_{\text{get}}$  might, in step  $\Pi^{(C:1)}$ , take a file identifier  $fid$  as client-side input  $in_C$  and send that identifier to the server as part its output message  $m_{out}$ . The server step  $\Pi^{(S:1)}$  then receives  $fid$  as message input  $m_{in}$ , might use it to look up the corresponding file  $f$  in its persistent state  $st_S$ , and send  $f$  back within its output message. In step  $\Pi^{(C:2)}$ , the client would then receive the file  $f$  from the server in the input message  $m_{in}$  and output that file to the user within  $out_C$ .

We use the convention that the client sends the first message and the server sends the last message. This restricts the syntax to client-initiated protocols, but we argue that it suffices to model all operations relevant for the cryptographic core of a cloud storage protocol. For example, our syntax cannot capture “pushing notifications” from the server to clients, say to inform the laptop of a user that their phone just uploaded a new file. However, in most cases, such functionality can be handled by the surrounding, non-cryptographic protocol: the cryptographic core functionality can simply assume that it is known which files can be requested. Alternatively, if necessary for security, a server-initiated operation can be modeled in our syntax by setting the first client message to the empty message (effectively turning a “push” mechanism into “polling”).

**FULL EXECUTION.** We write the full execution of the protocol  $\Pi$  on initial states  $st_C, st_S$  and inputs  $in_C, in_S$ , of client and server respectively,<sup>2</sup> as

$$(\mathbf{m}, st'_C, st'_S, out_C, out_S) \leftarrow \$ \mathbf{Exec}_{\Pi}(st_C, st_S, in_C, in_S).$$

<sup>1</sup>The input message  $m_{in}$  is empty, i.e.,  $m_{in} = \varepsilon$ , for the initiating party in step  $i = 1$ .

<sup>2</sup>Note that the inputs  $in_C$  and  $in_S$  are only arguments to the first step of their respective protocols. We discuss the reasons behind this choice and its implications in the general discussion of our cloud storage framework, Section 4.4.

**Exec $_{\Pi}$** ( $st_C, st_S, in_C, in_S$ ):

```

1 ( $st_C, st_C^{tmp}, out_C, \mathbf{m}[1]$ )  $\leftarrow$   $\Pi^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon)$ 
2 ( $st_S, st_S^{tmp}, out_S, \mathbf{m}[2]$ )  $\leftarrow$   $\Pi^{(S:1)}(st_S, \varepsilon, in_S, \mathbf{m}[1])$ 
3 for  $i = 3, \dots, \Pi.SN$ :
4    $r \leftarrow \lceil i/2 \rceil$  ; if  $i$  odd:  $P \leftarrow C$  else  $P \leftarrow S$ 
5   if  $out_P.dec = \perp$ : ( $st_P, st_P^{tmp}, out_P, \mathbf{m}[i]$ )  $\leftarrow$   $\Pi^{(P:r)}(st_P, st_P^{tmp}, \mathbf{m}[i-1])$ 
6 return ( $\mathbf{m}, st_C, st_S, out_C, out_S$ )

```

Figure 1: Execution **Exec $_{\Pi}$**  of a protocol  $\Pi$  between client  $C$  and server  $S$ .

Protocol		Inputs $in_C$	Outputs $out_C$
Account registration	$\Pi_{\text{areg}}$	$aid, pw$	—
Session authentication	$\Pi_{\text{auth}}$	$aid, pw$	— ( $st_C$ created)
File upload	$\Pi_{\text{put}}$	$fid, f$	—
File update	$\Pi_{\text{upd}}$	$fid, f$	—
File download	$\Pi_{\text{get}}$	$fid$	$f$
File share	$\Pi_{\text{share}}$	$raid, fid$	$oob$
Accept share	$\Pi_{\text{accept}}$	$fid, oob$	—

Table 2: Core protocols making up a cloud storage scheme CS, along with their mandatory *client-side* inputs and outputs; see the detailed explanation in Section 3.2.

The output of the execution is a transcript  $\mathbf{m}$  of the interaction, which is the sequence of messages exchanged, the new persistent states  $st'_C, st'_S$  of the parties, and their respective outputs  $out_C, out_S$ . As part of the output, the decision of party  $P \in \{C, S\}$  to accept or reject is encoded as  $out_P.dec \in \{\text{true}, \text{false}\}$ .

We specify the details of **Exec $_{\Pi}$**  in Figure 1: it calls the protocol for each party and relays messages as long as parties have not finalized their decisions  $out_P.dec$ . We will use this full (or “atomic”) execution of a protocol  $\Pi$  to define correctness, where we assume that operations are executed honestly and in sequential order (per client session). Later, when defining security, we allow the adversary to arbitrarily and maliciously interleave protocol steps.

**TERMINATION.** For more concise return statements in standard success or failure cases, we introduce shorthands **success $_P$**  and **fail $_P$**  for  $P \in \{C, S\}$  defined as

$$\begin{aligned}
\mathbf{success}_P(m_P): out_P.dec &\leftarrow \text{true}; \mathbf{return}(st_P, st_P^{tmp}, out_P, m_P), \\
\mathbf{fail}_P: out_P.dec &\leftarrow \text{false}; \mathbf{return}(st_P, st_P^{tmp}, out_P, \perp),
\end{aligned}$$

where by convention **success $_P$**  := **success $_P$** ( $\top$ ) returns the success symbol  $\top$ .

Failing procedures and operations – including **fail $_P$** , calls returning the error symbol  $\perp$ , operations on  $\perp$ , and using index  $\perp$  for tables – propagate this failure through calling procedures (comparable to program exceptions), causing the caller of the failing procedure to return **fail $_P$** .

### 3.2 The Protocols

Let us now describe the formalized core protocols of a cloud storage scheme CS in more detail. Table 2 lists them along with their mandatory inputs and outputs.

REGISTRATION AND AUTHENTICATION. The account registration protocol  $\Pi_{\text{reg}}$  creates a new user in the system, represented by an account ID (e.g., an email address) provided by the client. It expects that account ID  $aid$  as well as the user’s password  $pw$  as (minimal) client-side input within  $in_C$  and must be run before any other protocols can be run on behalf of that user.<sup>3</sup> The account registration protocol does not involve any persistent client session state yet (i.e.,  $st_C = st'_C = \varepsilon$ ), as such state pertains only to active sessions.

A registered user can then run the authentication protocol  $\Pi_{\text{auth}}$  to authenticate to the server and initiate a new client session, again using account ID  $aid$  and password  $pw$  as client-side input, and upon empty client state  $st_C = \varepsilon$ . This creates a persistent client session state  $st'_C$ , which is shared among all protocols run within that session. A user can initiate multiple client sessions in parallel (each holding their own state  $st_C$ ), which can concurrently access the user’s files in the cloud storage.

All remaining protocols operating on files can only be called with non-empty client state  $st_C$  (i.e., we implicitly require them to fail otherwise). Overall, this enforces that account registration must be run before authentication, and authentication must be run before any file operation protocol.

FILE UP- AND DOWNLOAD. Within an active session, a client can upload a new file using  $\Pi_{\text{put}}$  on input a (globally unique) file identifier  $fid$  and file content  $f$ . File retrieval is correspondingly done via  $\Pi_{\text{get}}$  on (client-side) input the file’s identifier  $fid$ .<sup>4</sup> We handle file updates and deletions through the  $\Pi_{\text{upd}}$  protocol, taking as client-side input an existing identifier  $fid$  and the updated file content  $f$ . File deletion can be interpreted as setting the file  $f$  in the client input to  $\perp$ .

FILE SHARING. We capture file sharing through two operations,  $\Pi_{\text{share}}$  and  $\Pi_{\text{accept}}$ . The sharing protocol  $\Pi_{\text{share}}$  allows a client to prepare sharing a file identified by  $fid$ ,<sup>5</sup> and notify the server about its intent to share said file with another user, identified by the receiver account ID  $raid$ . To capture that means of access (e.g., key material) might be exchanged between sender and receiver of a shared file through some out-of-band channel (e.g., messaging a sharing link), we allow the sharing protocol to output some dedicated out-of-band information  $out_C.oob$ . Apart from the OOB channel, there is no other client-to-client communication.

To accept the sharing of a file, the receiver client session runs  $\Pi_{\text{accept}}$  with the server. On input the file identifier  $fid$  of the shared file and the obtained out-of-band information  $oob$ , the  $\Pi_{\text{accept}}$  operation integrates the key material for the shared file into this client’s cloud storage for future access (through  $\Pi_{\text{get}}$ ,  $\Pi_{\text{upd}}$ , or  $\Pi_{\text{share}}$ ). In our notion, sharer and sharees all have full access to the *same* file with no distinction between the original file owner and share recipients.

While, in syntax, we do not put constraints on the out-of-band information, in practice, such out-of-band channels are highly restricted and hence realistic schemes will have to restrict their use of it.<sup>6</sup> In our proposed cloud storage scheme in Section 6, we use it only to exchange keys that are a few bytes long (embedded, e.g., in a URL or QR code) for providing access to shared files.

---

<sup>3</sup>We list here only the *minimal* required client-side inputs to protocols, needed to define basic functionality as well as security. A concrete cloud storage scheme might take further (client- or server-side) inputs.

<sup>4</sup>Clients can learn  $fid$  from communication outside of our modeled core protocols, e.g., by fetching a mapping from file names to file identifiers from the server.

<sup>5</sup>Some cloud storage protocols may use account-dependent globally unique file identifiers (e.g.,  $(aid, n)$  where  $n$  is a monotonic counter). In such a protocol, sharing files with user  $raid$  may require changing the file identifier (in the previous example, to  $(raid, n')$ ). This can still be modeled in with our framework, but the protocol description needs to introduce a level of indirection and map account-dependent identifiers to account-independent ones that can be used by multiple users simultaneously, and then translate them depending on context.

<sup>6</sup>There may be trivially secure schemes that make heavy use of the out-of-band channel, e.g., by using the channel to entirely bypass the service provider. While such protocols might match our definitions, they are not interesting in practice as they do not adhere to said real-world constraints.

There is no dedicated mechanism to revoke access to a shared file from a user. However, since all users share access to the same file, any one of them can delete and re-upload the file to create a version of it which is not shared. It is possible that by explicitly modeling revocation, one could capture further security guarantees associated to access rights management. For example, in the honest-server setting, it might be possible to ensure that a user who gets their access to a file revoked cannot download that file in the future, nor see updates or new versions of the file. We discuss extensions such as these to the functionality and security of our model in Section 4.4.

### 3.3 Correctness

We now define the expected functional behavior, i.e., correctness, of a cloud storage scheme CS in executions of honestly behaving clients and server. Due to the potential interleavings between protocol executions of different client sessions, we do this via the game  $\mathbf{G}_{\text{CS}}^{\text{corr}}$  given in Figure 2, with oracles defined in Figure 3.

The intuition of the correctness game is that if a user successfully uploaded a file to the server, or if a file was shared with them, then fetching that file should succeed and return the expected file content. Put differently, if the server behaves honestly (and in particular does not perform any availability attacks), then users can successfully recover their cloud storage files. Formally, this is captured by letting the adversary win the game (by setting a win flag) if a  $\Pi_{\text{get}}$  operation does not return the expected file content (see Figure 3, line 20) or if a  $\Pi_{\text{accept}}$  operation fails even though the file has been shared with the receiving user (Figure 3, line 35).

**Definition 3.1 (Cloud storage correctness)** *Let CS be a cloud storage scheme and  $\mathbf{G}_{\text{CS}}^{\text{corr}}$  be the correctness game for CS defined in Figures 2 and 3. We define the advantage of an adversary  $\mathcal{A}$  playing this game as*

$$\text{Adv}_{\text{CS}}^{\text{corr}}(\mathcal{A}) := \Pr[\mathbf{G}_{\text{CS}}^{\text{corr}}(\mathcal{A})],$$

and say that CS is correct if  $\text{Adv}_{\text{CS}}^{\text{corr}}(\mathcal{A}) = 0$  for all (even unbounded)  $\mathcal{A}$ .

Game  $\mathbf{G}_{\text{CS}}^{\text{corr}}(\mathcal{A})$ :

```

1 global  $S_A, n_c, st_S, T_{\text{stC}}, T_f, T_{\text{oob}}, S_{\text{shr}}, \text{win}$ 
2  $\text{win} \leftarrow \text{false}$  ;  $n_c \leftarrow 0$ 
3  $\mathcal{A}^{\text{AREG}, \text{AUTH}, \text{PUT}, \text{UPD}, \text{GET}, \text{SHARE}, \text{ACPT}}()$ 
4 return win

```

RunProt( $\Pi, i_c, in_C, in_S$ ):

```

5 pre:  $(T_{\text{stC}}[i_c] \neq \perp)$  or  $in_C.\{\text{aid}\}$ 
6 if  $T_{\text{stC}}[i_c] \neq \perp$ :  $(\text{aid}, st_C) \leftarrow T_{\text{stC}}[i_c]$  // Load client session state for PUT, GET, SHARE, ACPT
7 else:  $st_C \leftarrow \varepsilon$  // No client session state for AREG, AUTH
8  $(\mathbf{m}, st_C, st_S, out_C, out_S) \leftarrow \text{Exec}_{\Pi}(st_C, st_S, in_C, in_S)$ 
9 if  $(out_C.\text{dec} \wedge out_S.\text{dec})$ :
10   if  $\Pi = \Pi_{\text{areg}}$ :  $S_A \xleftarrow{\cup} \{\text{aid}\}$  // Register account upon successful AREG
11   if  $\Pi = \Pi_{\text{auth}}$ :  $T_{\text{stC}}[i_c] \leftarrow (\text{aid}, st_C)$  // Register client session state upon successful AUTH
12 if  $\Pi \notin \{\Pi_{\text{areg}}, \Pi_{\text{auth}}\}$ :  $T_{\text{stC}}[i_c] \leftarrow (\text{aid}, st_C)$ 
13 return  $(\mathbf{m}, st_C, st_S, out_C, out_S)$ 

```

Figure 2: Correctness game  $\mathbf{G}_{\text{CS}}^{\text{corr}}$  for cloud storage scheme CS, with oracles given in Figure 3, and helper function RunProt used by oracles.

```

AREG( $in_C, in_S$ ):
1 pre:  $in_C.aid \notin S_A, in_C.\{aid, pw\}$ 
2 return RunProt( $\Pi_{areg}, \perp, in_C, in_S$ )

AUTH( $in_C, in_S$ ):
3 pre:  $in_C.aid \in S_A, in_C.\{aid, pw\}$ 
4  $n_c \leftarrow n_c + 1$  // increment client session counter
5 return RunProt( $\Pi_{auth}, n_c, in_C, in_S$ )

PUT( $i_c, in_C, in_S$ ):
6 pre:  $i_c \in T_{stC}, in_C.\{fid, f\}$ 
7 if  $fid \in T_f$ : return  $\perp$  // not allowed to upload same file id twice
8  $(m, st_C, st_S, out_C, out_S) \leftarrow$  RunProt( $\Pi_{put}, i_c, in_C, in_S$ )
9 if  $out_S.dec$ :
10    $ow \leftarrow \{T_{stC}[i_c].aid\}; T_f[fid] \leftarrow (ow, f)$  // store owner and file in file table
11 return  $(m, st_C, st_S, out_C, out_S)$ 

UPD( $i_c, in_C, in_S$ ):
12 pre:  $i_c \in T_{stC}, in_C.\{fid, f\}$ 
13 if  $T_{stC}[i_c].aid \notin T_f[fid].ow$ : return  $\perp$  // only allow updated by file owners
14  $(m, st_C, st_S, out_C, out_S) \leftarrow$  RunProt( $\Pi_{upd}, i_c, in_C, in_S$ )
15 if  $out_S.dec$ :  $T_f[fid].f \leftarrow f$  // update file in file table
16 return  $(m, st_C, st_S, out_C, out_S)$ 

GET( $i_c, in_C, in_S$ ):
17 pre:  $i_c \in T_{stC}, in_C.\{fid\}$ 
18 if  $T_{stC}[i_c].aid \notin T_f[fid].ow$ : return  $\perp$  // only allow downloads by file owners
19  $(m, st_C, st_S, out_C, out_S) \leftarrow$  RunProt( $\Pi_{get}, i_c, in_C, in_S$ )
20 if  $(out_C.dec \text{ and } out_C.f \neq T_f[fid].f)$ : win  $\leftarrow$  true // retrieved unexpected file
21 if  $(\neg out_C.dec \text{ and } fid \in T_f[fid])$ : win  $\leftarrow$  true // fail to retrieve uploaded file
22 return  $(m, st_C, st_S, out_C, out_S)$ 

SHARE( $i_c, in_C, in_S$ ):
23 pre:  $i_c \in T_{stC}, in_C.\{raid, fid\}$ 
24 if  $T_{stC}[i_c].aid \notin T_f[fid].ow$ : return  $\perp$  // only allow sharing by file owners
25  $(m, st_C, st_S, out_C, out_S) \leftarrow$  RunProt( $\Pi_{share}, i_c, in_C, in_S$ )
26 if  $out_C.dec$ :
27    $S_{shr} \stackrel{\cup}{\leftarrow} \{(raid, fid)\}; T_{oob}[(raid, fid)] \leftarrow out_C.oob$  // store OOB info for receiver
28 return  $(m, st_C, st_S, out_C, out_S)$ 

ACPT( $i_c, in_C, in_S$ ):
29 pre:  $i_c \in T_{stC}, in_C.\{fid\}$ 
30 if  $(T_{oob}[(T_{stC}[i_c].aid, fid)] = \perp)$ : return  $\perp$  // only proceed if out-of-band info available
31  $(in_C.oob) \leftarrow T_{oob}[(T_{stC}[i_c].aid, fid)]$  // retrieve out-of-band info for receiver
32  $(m, st_C, st_S, out_C, out_S) \leftarrow$  RunProt( $\Pi_{accept}, i_c, in_C, in_S$ )
33 if  $out_S.dec$ :  $T_f[fid].ow \stackrel{\cup}{\leftarrow} \{T_{stC}[i_c].aid\}$  // mark receiver as file owner
34 if  $out_C.dec$ :  $T_{oob}[(T_{stC}[i_c].aid, fid)] \leftarrow \perp$  // remove out-of-band info for receiver
35 if  $\neg out_C.dec \wedge (T_{stC}[i_c].aid, fid) \in S_{shr}$ : win  $\leftarrow$  true // fail to accept shared file
36 return  $(m, st_C, st_S, out_C, out_S)$ 

```

Figure 3: Oracles for the correctness game  $\mathbf{G}_{CS}^{\text{corr}}$  for cloud storage scheme CS. The helper function RunProt, given in Figure 2, captures core execution steps the protocols.

Let us explain some further technical details of the correctness game.

ATOMIC EXECUTION. Recall that, for simplicity, correctness is defined w.r.t. “atomic” protocol executions where individual protocol steps are executed sequentially and cannot be interleaved. This simplifies tracking of variables and avoids handling race conditions (e.g., which file version should be expected to be downloaded after two concurrent file uploads using the same file id?). We emphasize that when defining security later, we capture arbitrary and malicious interleaving of individual protocol steps.

FOCUS ON CORE FUNCTIONALITY. Our correctness definition only captures the core functionality of a cloud storage scheme: a successfully uploaded or shared file should be retrievable by authorized clients. In particular, correctness does *not* demand that uploading or sharing operations have to be successful. This is intentional: it allows for different instantiations of these protocols. For example, a scheme might reject files that are too large or whose file identifiers are malformed – but we do not want to prescribe such rules for a generic scheme. As another example, one scheme may allow sharing a file only with existing accounts while another might allow sharing with non-existing accounts (e.g., to invite a new user to the system via their not-yet-registered email address) – hence, we do not prescribe when sharing ought to be successful. We note that this opens our correctness notion up to “correct” dummy schemes (e.g., where uploading or sharing is never successful); in practice such schemes are trivially uninteresting, and we accept this trade-off in favor of being able to express a rich variety of cloud storage schemes in our framework.

PRECONDITIONS AND OTHER NOTATION. We annotate game oracles with preconditions (written **pre:** ...) which specify the minimal set of input values required by the game for the oracle call of some protocol to be meaningful. An oracle may for example specify **pre:**  $in_C.\{fid, f\}$  to indicate that it requires the values  $in_C.fid$  and  $in_C.f$  to be set. For brevity, we refer to the values required in a precondition without the prefix  $in_C$  inside the code of the oracle (in the above example:  $fid$  instead of  $in_C.fid$ ). Unsatisfied preconditions cause the oracle to abort.

ACCOUNTS AND CLIENTS. Users are identified via their account ID  $aid$ , set as client input by the adversary in the oracles AREG for account registration and AUTH for session authentication. The game keeps a count of all authenticated client sessions; whenever authentication succeeds in AUTH on input account ID  $aid$ , it associates a new client session counter value  $i_c$  with  $aid$  and the session’s state  $st_C$  (see Figure 2, line 11).

GLOBAL GAME VARIABLES. The correctness game tracks several global variables across oracle calls. This includes the set  $S_A$  of all IDs for registered accounts, the counter  $n_c$  for client sessions, the server state  $st_S$ , and the table  $T_{stC}$  which tracks client session states. Uploaded files and their owners are tracked using table  $T_f$ , which is indexed via the (globally unique) file identifiers  $fid$  specified by the adversary. The game enforces globally unique file identifiers that are never repurposed by never resetting the owner set of a file (even when a file is “deleted” by overwriting it with  $\perp$ ). Out-of-band information exchanged between the sender and receiver of file shares is tracked using table  $T_{oob}$ .

The game oracles update most of these global variables only if operations succeed (indicated by  $out_C.dec$  and  $out_S.dec$  being true), encoding that most correctness guarantees only apply after a successful protocol run. For instance, the oracle  $\Pi_{put}$  updates the file table  $T_f$  after the protocol run succeeds on the server (see Figure 3, line 9). However, we do not require that both parties always succeed. Indeed, for all file-related protocols, we update the client session states in  $T_{stC}$  on line 12 of Figure 2 irrespective of the output decisions. This captures that even a correct scheme might be expected to fail, for instance, when trying to upload a file with a malformed file ID, but that in case of such a protocol failure, the scheme should still continue to function. (That is, the



failed operation should not lead to a corrupted state on either client or server.) In contrast, for registration and authentication, line 9 of Figure 2 ensures that  $T_{\text{stC}}$  is only updated when both parties succeed, capturing that these protocols can only be correct if both participants agree on their outcome. This is necessary, as later operations for the same user or in the same session rely on registration and authentication, respectively, to have been successful.

## 4 Client-to-Client Security of Cloud Storage

In our cloud storage framework, we are interested in guaranteeing end-to-end encryption security for the data of clients. We call this *client-to-client* (C2C) security. The aim is to protect honest clients against a fully malicious, adversarially-controlled cloud provider, giving both *confidentiality* (C2CConf) and *integrity* (C2CInt) guarantees. Our security model focuses on the operations of the cloud storage system itself. We hence assume clients communicate with the storage server over a server-authenticated and confidential channel (e.g., a TLS connection) and do not concern ourselves with capturing the setup or concrete security of that channel. We point out subtle reasons for further modeling choices where they arise, and expand on them in the discussion in Section 4.4. There, we also discuss potential extensions to our security model, including additional protection against weaker adversaries that do not control the cloud server, as well as the modeling of metadata. More future work is mentioned in Section 7.

**PASSWORD MODELING.** Users in our cloud storage framework are authenticated through passwords (see Section 3). All of our security games are parameterized by  $\mathcal{PW}_n$ , a distribution of password vectors of cardinality  $n$  (the maximal number of user accounts), and pre-sample a password vector  $\mathbf{pw} \xleftarrow{\vec{r}} \mathcal{PW}_n$  from this distribution. Security is then defined with respect to the distribution  $\mathcal{PW}_n$ . This approach is inspired by prior game-based treatments of password-based encryption [9, 12], modeling non-uniform password distributions and correlated or even identical passwords in game-based proofs – i.e., the choices users make when picking passwords.

Previous work [9, 12] assumed that  $\mathbf{pw}$  consists of unique passwords to “avoid trivial attacks.” We do not make this assumption, as we deem it impossible to enforce in practice even for an honest service provider. Instead, we assume that the *account IDs* *aid* are globally unique for honest users. This can be enforced in practice, e.g. by taking email addresses as account IDs, possibly verified via a token exchange. We hence demand that the cloud storage scheme achieves security even when multiple users have the same password. Note that we do assume that the password distribution is independent from the hash functions used to hash passwords in our proposed cloud storage scheme. This is required for our cloud storage proofs, as the existence of pathological distributions otherwise make the derivation of pseudorandom keys from passwords impossible.

Security obviously depends on the strength of the passwords chosen by users (as modeled by the distribution  $\mathcal{PW}_n$ ). More specifically, it depends on the chances of an adversary outright *guessing* a user’s password. Following [9, 12], we are then interested in the security of a cloud storage scheme beyond the “guessing probability” of passwords from the given distribution. Naturally, a heavily skewed password distribution where guessing user passwords is easy leads to poor security guarantees for such users; real-world systems (e.g., Nextcloud [48] and OnePassword [49]) employ according mitigations to assist users in picking secure passwords.

We capture the password-guessing probability through the game in Figure 4, where the adversary, given selective access to the passwords of compromised users, wins by guessing the password of an uncorrupted user.

**Definition 4.1 (Password guessing)** Let  $\mathbf{G}_{n, \mathcal{PW}_n}^{\text{PG}}$  be the password-guessing game defined in Fig-

Game $\mathbf{G}_{n, \mathcal{PW}_n}^{\text{PG}}(\mathcal{A})$ :	TEST( $i_u, pw$ ):
1 <b>global</b> $S_{\text{comp}}, \mathbf{pw}, \text{win}$	7 <b>if</b> $(pw = \mathbf{pw}[i_u]) \wedge (i_u \notin S_{\text{comp}})$ :
2 $\mathbf{pw} \xleftarrow{\vec{\cdot}} \mathcal{PW}_n$	8 $\text{win} \leftarrow \text{true}$
3 $S_{\text{comp}} \leftarrow \mathcal{A}()$	
4 <b>for</b> $i_u \in S_{\text{comp}}$ : $\mathbf{pw}_{\text{comp}}[i_u] \leftarrow \mathbf{pw}[i_u]$	
5 $\mathcal{A}^{\text{TEST}}(\mathbf{pw}_{\text{comp}})$	
6 <b>return</b> win	

Figure 4: Password guessing game with selective compromises against a password distribution  $\mathcal{PW}_n$  and at most  $n$  users.

ure 4, for at most  $n$  users and with respect to a password distribution  $\mathcal{PW}_n$ . We define the password-guessing advantage of an adversary  $\mathcal{A}$  playing this game as

$$\mathbf{Adv}_{n, \mathcal{PW}_n}^{\text{PG}}(\mathcal{A}) := \Pr [\mathbf{G}_{n, \mathcal{PW}_n}^{\text{PG}}].$$

**ADVERSARIAL INTERACTION.** Before introducing our confidentiality and integrity notions, we explain some common features of our games with respect to the interaction with the adversary. The first surrounds the fine-grained level at which we capture the operation of the cloud storage scheme: As mentioned in earlier sections, the protocols comprising a cloud storage scheme may consist of several, interactive steps. For instance, the file update protocol  $\Pi_{\text{upd}}$  might first fetch some information about the file to be updated from the storage server before uploading the updated file content. In practice, a dishonest server might maliciously interleave the executions of these protocol steps to its benefit. Our cloud storage security games hence allow step-by-step execution of the protocols, with arbitrary, adversary-controlled interleaving across protocols.

In the games (see Figures 5 and 6), this multi-step nature of protocols is captured by providing the adversary with oracles to initiate each protocol run (AREG1, AUTH1, etc.), which check potential preconditions, and a “stepping” oracle STEP through which the adversary can trigger the next client-side step of a protocol (identified by some process ID  $i_p$ ). The games use the table  $T_p$  to keep track of the state of each ongoing protocol execution, including the participants’ state.

On the client side, the adversary gets to pick relevant inputs (e.g., account and file identifiers, file contents) for honest clients. The server instead is embodied by the adversary, capturing the C2C guarantees against malicious cloud providers. In particular, this means that the adversary receives the honest clients’ messages and controls the server’s responses in each protocol step.

## 4.1 Client-to-Client Confidentiality

We first define client-to-client confidentiality, via the game  $\mathbf{G}_{\text{CS}, n, \mathcal{PW}_n}^{\text{C2CConfS}}$  given in the top panel of Figure 5, together with oracles and helper functions in Figures 5 and 6. (Our integrity game is defined through the same figures, as it heavily overlaps in the basic structure; it contains the code lines in `[dashed]` boxes, which are not present in the confidentiality game. These can hence safely be ignored for now.) The game is parameterized by the maximum number of user accounts  $n$  and a distribution  $\mathcal{PW}_n$  of password vectors of the same cardinality. We further distinguish two versions of the game, depending on whether the adversary is allowed to adaptively (C2CConf) or selectively (C2CConfS) compromise clients; the latter version includes the code in `[dotted]` boxes.



In brief, confidentiality is captured as left-or-right indistinguishability<sup>7</sup> (akin to classical IND-CCA security for encryption): The game samples a random challenge bit  $b$  and provides the adversary  $\mathcal{A}$  with a special challenge oracle CHALL which uploads  $f_b$  out of a pair of files  $(f_0, f_1)$ , via a put or upd operation of an adversary-selected client. The adversary’s task is to determine  $b$ , while acting as the malicious server and interacting with honest clients through the game oracles. To prevent trivial attacks, the adversary cannot compromise any user that owns a challenge file. We first state the formal security definition before we explain the game in more detail.

**Definition 4.2 (Client-to-client confidentiality)** *Let  $\text{CS}$  be a cloud storage scheme and let  $\mathbf{G}_{\text{CS},n,\text{PW}_n}^{\text{C2CConf},S}$  be the adaptive (*selective*) client-to-client confidentiality (C2CConf $\langle S \rangle$ ) game for  $\text{CS}$  defined in Figures 5 and 6, for at most  $n$  honest users and with respect to a password distribution  $\text{PW}_n$ . We define the advantage of an adversary  $\mathcal{A}$  playing this game as*

$$\text{Adv}_{\text{CS},n,\text{PW}_n}^{\text{C2CConf},S}(\mathcal{A}) := 2 \cdot \Pr \left[ \mathbf{G}_{\text{CS},n,\text{PW}_n}^{\text{C2CConf},S} \right] - 1.$$

ORACLES. In game  $\mathbf{G}_{\text{CS},n,\text{PW}_n}^{\text{C2CConf},S}$ , the adversary  $\mathcal{A}$  has access to the following oracles: one oracle of the form ORC1 for each  $\text{CS}$  protocol (used to initiate a protocol run), the oracle STEP to perform the next client step in a protocol, the challenge oracle CHALL for uploading a challenge file, and the oracle COMP to compromise (the entire knowledge) of a user. As in the correctness game (see Section 3.3), we annotate oracles with preconditions where, e.g.,  $\text{pre: } in_C.\{fid\}$  indicates that  $in_C.fid$  must be set.

Oracles of the form ORC1 capture basic protocol operations. Each of them simulates an honest client (session) initiating the corresponding  $\text{CS}$  protocol. The game tracks the honest clients’ accounts, passwords, as well as persistent session state ( $st_C$ ) and temporary, protocol-internal state ( $st_C^{\text{tmp}}$ ). Since the adversary plays the role of the (malicious) server, the game does not keep any server state, but only tracks client-related information relevant for checking trivial winning conditions, e.g., which files an honest user has uploaded or shared.

In more detail, the basic oracles work as follows:

- AREG1 initiates registering a new, honest account for a user with index  $i_u$  and account ID  $aid$ , setting the user password to  $\text{pw}[i_u]$  (hidden from  $\mathcal{A}$ ). The game allows repeated account registration attempts of honest clients, but enforces that the user index to account ID mapping is bijective.
- AUTH1 initiates a new session  $i_c$  of an honest user with given user index  $i_u$ , tracking the new session’s user index in  $T_u[i_c]$  and persistent session state in  $T_{stC}[i_c]$ .
- PUT1 and UPD1 let session  $i_c$  initiate an upload resp. update of a file  $f$  under ID  $fid$ , both adversary-chosen. The precondition of PUT1 enforces unique file IDs  $fid$ , which can be achieved in practice by prefixing file IDs by  $aid$ . For a newly put file, the user ID of the calling client is tracked as an owner in the set  $T_f[fid].U$ .
- GET1 lets session  $i_c$  initiate fetching a file with given ID  $fid$ . Successfully downloaded files are handed to  $\mathcal{A}$  in the helper function Check (serving as a decryption oracle), except for challenge files (to prevent trivial wins).
- SHARE1 and ACPT1 let session  $i_c$  initiate sharing a file resp. accepting a shared file. The game tracks the owners of shared files in  $T_f[fid].U$ , and of challenge files in  $S_{\text{challOwners}}$ , and keeps track of the out-of-band information in  $T_{\text{oob}}$ . Note that to prevent trivial attacks, the

<sup>7</sup>Our framework does not single out “ciphertext” outputs, leaving the cloud storage scheme maximal freedom in how it implements E2EE. Consequently left-or-right indistinguishability is the appropriate confidentiality notion, as an IND $\$$ -style notion would require such a distinguished ciphertext (in order to replace it by random).

<p>Game <math>\mathbf{G}_{\text{CS},n,\mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A})</math>:</p> <pre> 1 <b>global</b> <math>T_{\text{uaid}}, T_u, T_{\text{stC}}, T_p, T_{\text{oob}}, T_f, S_{\text{reg}}, S_{\text{chall}}, S_{\text{hon}}, S_{\text{comp}}, S_{\text{challOwners}}, \mathbf{pw}, b</math> 2 <math>\mathbf{pw} \xleftarrow{\mathcal{V}} \mathcal{PW}_n</math>; <math>n_c \leftarrow 0</math>; <math>n_p \leftarrow 0</math>; <math>b \leftarrow \{0, 1\}</math> 3 <math>S_{\text{comp}} \leftarrow \mathcal{A}()</math> 4 <math>b' \leftarrow \mathcal{A}^{\text{STEP}, \text{AREG1}, \text{AUTH1}, \text{PUT1}, \text{UPD1}, \text{GET1}, \text{SHARE1}, \text{ACPT1}, \text{CHALL}, \text{COMP}}()</math> 5 <math>\text{trivial} \leftarrow \exists i_u \in S_{\text{challOwners}} : i_u \in S_{\text{comp}} \vee i_u \notin S_{\text{hon}}</math> // <math>\exists</math> adversary-owned challenge file 6 <b>return</b> <math>\neg \text{trivial} \wedge (b' = b)</math> </pre>
<p>Game <math>\mathbf{G}_{\text{CS},n,\mathcal{PW}_n}^{\text{C2CIntS}}(\mathcal{A})</math>:</p> <pre> 7 <b>global</b> <math>T_{\text{uaid}}, T_u, T_{\text{stC}}, T_p, T_{\text{oob}}, T_f, S_{\text{reg}}, S_{\text{hon}}, S_{\text{comp}}, \mathbf{pw}, \text{win}</math> 8 <math>\mathbf{pw} \xleftarrow{\mathcal{V}} \mathcal{PW}_n</math>; <math>n_p \leftarrow 0</math>; <math>n_c \leftarrow 0</math>; <math>T_f[\cdot].F \leftarrow \emptyset</math>; <math>T_f[\cdot].U \leftarrow \emptyset</math>; <math>\text{win} \leftarrow \text{false}</math> 9 <math>S_{\text{comp}} \leftarrow \mathcal{A}()</math> 10 <math>\mathcal{A}^{\text{STEP}, \text{AREG1}, \text{AUTH1}, \text{PUT1}, \text{UPD1}, \text{GET1}, \text{SHARE1}, \text{ACPT1}, \text{COMP}}()</math> 11 <b>return</b> win </pre>
<p><u>Step1</u>(<math>\Pi, i_c, i_r, in_C</math>):</p> <pre> 12 <math>(T_{\text{stC}}[i_c], st_C^{\text{tmp}}, out_C, m') \leftarrow \Pi^{(C:1)}(T_{\text{stC}}[i_c], \varepsilon, in_C, \varepsilon)</math> 13 <math>i_p \leftarrow n_p++</math>; <math>T_p[(\Pi, i_p)] \leftarrow (i_c, i_r, st_C^{\text{tmp}}, in_C, 2)</math> 14 <b>return</b> <math>\text{Check}(\Pi, i_p, out_C, m')</math> </pre> <p><u>Check</u>(<math>\Pi, i_p, out_C, m'</math>):</p> <pre> 15 <math>(i_c, i_r, st_C^{\text{tmp}}, in_C, r) \leftarrow T_p[(\Pi, i_p)]</math>; <math>fid \leftarrow in_C.fid</math> 16 <b>if</b> <math>\Pi = \Pi_{\text{areg}} \wedge out_C.dec</math>: <math>S_{\text{reg}} \xleftarrow{\cup} \{T_u[i_c]\}</math> 17 <b>if</b> <math>\Pi = \Pi_{\text{get}}</math>: 18   <math>\text{if } T_u[i_c] \in S_{\text{hon}} \setminus S_{\text{comp}}</math>: 19     <math>\text{win} \leftarrow out_C.dec \wedge out_C.f \notin T_f[fid].F \wedge T_f[fid].U \subseteq S_{\text{hon}} \setminus S_{\text{comp}}</math> 20   <math>\text{if } fid \in S_{\text{chall}}</math>: <math>out_C.f \leftarrow \perp</math> 21 <b>return</b> <math>(out_C.f, m')</math> 22 <b>if</b> <math>(\Pi = \Pi_{\text{share}} \text{ and } out_C.dec)</math>: <math>T_{\text{oob}}[(i_r, fid)] \leftarrow out_C.oob</math> 23 <b>if</b> <math>(\Pi = \Pi_{\text{accept}} \text{ and } out_C.dec)</math>: <math>T_{\text{oob}}[(T_u[i_c], fid)] \leftarrow \perp</math> 24 <b>return</b> <math>m'</math> </pre> <p><u>STEP</u>(<math>\Pi, i_p, m</math>):</p> <pre> 25 <math>(i_c, st_C^{\text{tmp}}, in_C, r) \leftarrow T_p[(\Pi, i_p)]</math> 26 <b>if</b> <math>r &gt; \Pi.SN \vee (\Pi = \Pi_{\text{areg}} \wedge T_u[i_c] \in S_{\text{reg}})</math>: <b>return</b> <math>\perp</math> 27 <math>(T_{\text{stC}}[i_c], st_C^{\text{tmp}}, out_C, m') \leftarrow \Pi^{(C:r)}(T_{\text{stC}}[i_c], st_C^{\text{tmp}}, m)</math> 28 <math>T_p[(\Pi, i_p)] \leftarrow (i_c, st_C^{\text{tmp}}, in_C, r + 1)</math> 29 <b>return</b> <math>\text{Check}(\Pi, i_p, out_C, m')</math> </pre>

Figure 5: Client-to-client confidentiality  $\text{C2CConfS}$  (top) and integrity  $\text{C2CIntS}$  (middle) games for security against a fully malicious cloud storage server with fine-grained control, as well as helper functions **Step1** and **Check** and oracle **STEP** (bottom). Code only included in the  $\text{C2CConfS}$  or  $\text{C2CIntS}$  game is marked with boxed and dashed boxes, respectively. Moreover, dotted code is only included in the *selective security* version of the games; the adversary is implicitly assumed to be stateful in that case.

```

AREG1( $i_u, in_C$ ):
1 pre:  $in_C.\{aid\}, (\#i'_u \in T_{\text{uaid}} \setminus \{i_u\}): aid = T_{\text{uaid}}[i'_u], i_u \notin S_{\text{reg}}$ 
2  $in_C.pw \leftarrow pw[i_u]; S_{\text{hon}} \stackrel{\cup}{\leftarrow} \{i_u\}$ 
3 if  $i_u \in T_{\text{uaid}}: in_C.aid \leftarrow T_{\text{uaid}}[i_u]$  else:  $T_{\text{uaid}}[i_u] \leftarrow aid$ 
4 return Step1( $\Pi_{\text{areg}}, \varepsilon, \varepsilon, in_C$ )

AUTH1( $i_u, in_C$ ):
5  $in_C.aid \leftarrow T_{\text{uaid}}[i_u]; in_C.pw \leftarrow pw[i_u]; i_c \leftarrow n_c++ ; T_u[i_c] \leftarrow i_u ; T_{\text{stc}}[i_c] \leftarrow \varepsilon$ 
6 return Step1( $\Pi_{\text{auth}}, i_c, \varepsilon, in_C$ )

PUT1( $i_c, in_C$ ):
7 pre:  $in_C.\{fid, f\}, fid \notin T_f$ 
8  $U \leftarrow \{T_u[i_c]\}; \bar{F} \stackrel{\cup}{\leftarrow} \{f\}$ 
9  $T_f[fid] \leftarrow (U, F)$ 
10 return Step1( $\Pi_{\text{put}}, i_c, \varepsilon, in_C$ )

UPD1( $i_c, in_C$ ):
11 pre:  $in_C.\{fid, f\}$ 
12  $[\text{if } T_u[i_c] \in T_f[fid].U: T_f[fid].F \leftarrow T_f[fid].F \cup \{f\}]$ 
13 return Step1( $\Pi_{\text{upd}}, i_c, \varepsilon, in_C$ )

GET1( $i_c, in_C$ ):
14 pre:  $in_C.\{fid\}$ 
15 return Step1( $\Pi_{\text{get}}, i_c, \varepsilon, in_C$ )

SHARE1( $i_c, i_r, in_C$ ):
16 pre:  $in_C.\{fid\}$ 
17  $T_f[fid].U \stackrel{\cup}{\leftarrow} \{T_u[i_c]\}$ 
18 if  $i_r \in T_{\text{uaid}}: in_C.raid \leftarrow T_{\text{uaid}}[i_r]$ 
19 if  $fid \in S_{\text{chall}}: S_{\text{challOwners}} \stackrel{\cup}{\leftarrow} \{i_r, T_u[i_c]\}$ 
20 return Step1( $\Pi_{\text{share}}, i_c, i_r, in_C$ )

ACPT1( $i_c, in_C$ ):
21 pre:  $in_C.\{fid\}, in_C.oob \neq \perp \vee T_{\text{oob}}[(i_r, fid)] \neq \perp$ 
22  $i_r \leftarrow T_u[i_c]$ 
23  $T_f[fid].U \stackrel{\cup}{\leftarrow} \{i_r\}$ 
24 if  $in_C.oob = \perp: in_C.oob \leftarrow T_{\text{oob}}[(i_r, fid)]$ 
25 else:  $T_f[fid].U \stackrel{\cup}{\leftarrow} \{\iota_{\text{mal}}\}; [\text{if } fid \in S_{\text{chall}}: S_{\text{challOwners}} \stackrel{\cup}{\leftarrow} \{\iota_{\text{mal}}\}]$ 
26 return Step1( $\Pi_{\text{accept}}, i_c, \varepsilon, in_C$ )


CHALL( $i_c, in_C, f_0, f_1, \text{ORC}$ ):
27 pre:  $in_C.\{fid\}, |f_0| = |f_1|, \text{ORC} \in \{\text{PUT1}, \text{UPD1}\}$ 
28 if  $(\text{ORC} = \text{UPD1} \wedge fid \notin S_{\text{chall}}): \text{return } \perp$ 
29  $in_C.f \leftarrow f_b; S_{\text{chall}} \stackrel{\cup}{\leftarrow} \{fid\}; S_{\text{challOwners}} \stackrel{\cup}{\leftarrow} T_u[i_c] \cup T_f[fid].U$ 
30 return  $\text{ORC}(i_c, in_C)$ 


COMP( $i_u$ ):
31 if  $i_u \notin S_{\text{comp}} \text{ return } \perp$ 
32  $S_{\text{comp}} \leftarrow S_{\text{comp}} \cup \{i_u\}$ 
33 return  $(pw[i_u], \{T_{\text{stc}}[i_c] \mid \exists i_c : T_u[i_c] = i_u\}, \{T_{\text{oob}}[(i_u, \cdot)]\})$ 

```

Figure 6: Oracles for the client-to-client games (see Figure 5). Code only included in the C2CConf[S] or C2CInt[S] game is marked with boxed and dashed boxes, respectively. Moreover, dotted code is only included in the *selective security* version of the games.

sharer is added as an owner on lines 17 and 19 in oracle SHARE1. This avoids attacks from compromised users sharing files which they do not own with honest users, causing the receiver to use malicious file key material for future file updates. The game allows the adversary itself to share files with and accept files from honest users by specifying the out-of-band message  $in_C.oob$  in ACPT1 or using the special index  $\iota_{\text{mal}}$  as recipient in SHARE1, respectively.

All basic oracles use the common subroutine **Step1** (in Figure 5) to initialize a protocol run and game-internal process tracking. The adversary can then use **STEP** to execute the next client step in an already running protocol  $\Pi$ , identified by process index  $i_p$ , on input the next (adversary-chosen) server message  $m$ .

Finally, security is captured via the challenge (**CHALL**) and compromise (**COMP**) oracles:

- **CHALL** takes two files  $(f_0, f_1)$  of equal size and an oracle name  $\{\Pi_{\text{put}}, \Pi_{\text{upd}}\}$ . If  $\text{ORC} = \Pi_{\text{put}}$ , the oracle uploads  $f_b$ , depending on the game bit  $b$ . The adversary can subsequently update the challenge file by running **CHALL** again on the same file ID with  $\text{ORC} = \Pi_{\text{upd}}$ . The game allows for multiple challenge files, the IDs of which are tracked in the set  $S_{\text{chall}}$ .
- **COMP** compromises a given honest user  $i_u$ , yielding their password and entire (current) state to the adversary. It may be called at any time, and repeatedly, to capture a persistent user compromise. Compromised users are tracked via  $S_{\text{comp}}$ . We do not model partial user compromise. That is, the adversary can only compromise the full state of a user, including all of its active sessions.

**MALICIOUS VS COMPROMISED CLIENTS.** Beyond orchestrating actions of honest clients, the adversary can directly simulate any number of *malicious* clients locally, without the need for oracles, since it fully controls the server state. Furthermore, after compromising a client via the oracle **COMP**, the server can continue using the protocol oracles on that client, but can also use the compromised client state to compute operations outside the honest behavior of clients.

Client compromise in particular models that an adversary may learn information about the password distribution. For instance, a scheme that leaks which users have the same password is vulnerable to attacks compromising passwords.

**SELECTIVE SECURITY.** In the selective security version of the game, **C2CConfS**, the adversary has to commit to the set  $S_{\text{comp}}$  of users it is allowed to compromise at the outset of the game (see Figure 5, line 3). It can then still call **COMP** adaptively at any point and repeatedly in the game, but only on these selected users, to learn their current state.

**WINNING CONDITION.** The adversary wins the game if it correctly guesses the challenge bit  $b$  in a non-trivial way. Line 5 of Figure 5 encodes this trivial win check: files that the adversary called the challenge oracle on may only be owned by or shared with honest and non-compromised users  $i_u \in S_{\text{hon}} \setminus S_{\text{comp}}$ . (Note that this in particular encodes that challenge files may not be shared with fully malicious, adversary-controlled clients.)

## 4.2 Client-to-Client Integrity

We now turn to the second security goal, client-to-client integrity, captured by the game  $\mathbf{G}_{\text{CS},n,\mathcal{PW}_n}^{\text{C2CInt}}$  given in the middle panel of Figure 5, again with oracles and helper functions in Figures 5 and 6. The integrity game includes the dashed code and excludes that in solid boxes; in particular, the challenge oracle **CHALL** is not present in the integrity game.

The task of the adversary in the integrity game is to forge the content of a file, i.e., to make an honest user successfully retrieve a file with a content that differs from what they uploaded or received as a share. This is captured by setting a win flag in line 19 of Figure 5, conditioned on there being no trivial attacks. We define the resulting client-to-client integrity as follows.

**Definition 4.3 (Client-to-Client Integrity)** Let  $\mathcal{CS}$  be a cloud storage scheme and  $\mathbf{G}_{\mathcal{CS},n,\mathcal{PW}_n}^{\text{C2CInt},S}$  be the adaptive (*selective*) client-to-client integrity ( $\text{C2CInt},S$ ) game for  $\mathcal{CS}$  defined in Figures 5 and 6, for at most  $n$  honest users and with respect to a password distribution  $\mathcal{PW}_n$ . We define the advantage of an adversary  $\mathcal{A}$  playing this game as

$$\text{Adv}_{\mathcal{CS},n,\mathcal{PW}_n}^{\text{C2CInt},S}(\mathcal{A}) := \Pr \left[ \mathbf{G}_{\mathcal{CS},n,\mathcal{PW}_n}^{\text{C2CInt},S} \right].$$

WHAT COUNTS AS A FORGERY. To evaluate the winning condition (in Figure 5, line 19), we track the content of honestly uploaded files in set  $T_f.F$  and their owners in  $T_f.U$ , both of which are implicitly initialized to  $\emptyset$ .<sup>8</sup> As a valid forgery we count any successful get operation (i.e., for which  $\text{out}_C.\text{dec} = \text{true}$ ) for which either (1) the file ID does not exist, meaning that the file is not tracked in the game, or (2) the retrieved file content does not match any version of the file stored in  $T_f.F$ . For case (2), we also need to ensure the file is not owned by a compromised or malicious user; i.e., the file owners  $T_f.U$  must be a subset of  $S_{\text{hon}} \setminus S_{\text{comp}}$ .

Note that a file can be either entirely or partially untracked, but line 19 covers both cases. That is, if  $T_f[\text{fid}].F = \emptyset$ , then  $\text{out}_C.f \notin T_f[\text{fid}].F$  evaluates to true. The file then counts as forgery if  $T_f[\text{fid}].U \subseteq S_{\text{hon}} \setminus S_{\text{comp}}$  is true too, either because the file is entirely untracked and  $T_f[\text{fid}].U = \emptyset$ , or because the owner set only contains honest users who shared or accepted a file ID that was never uploaded through oracle PUT1.

To prevent trivial attacks, the win flag is only set to true if the get operation is performed for an honest user (line 18, Figure 5). That is, the user ID  $T_u[i_c]$  of the calling client must be in  $S_{\text{hon}} \setminus S_{\text{comp}}$ . Otherwise, the adversary could easily forge a file, since it knows the retrieving user's key material.

The tracking of file content and owners in table  $T_f$  is conservative, in three aspects.

- In PUT1 and UPD1, we set resp. add the to-be-uploaded file content immediately in the table, without waiting for successful termination of the protocol. The reasoning here is that we cannot know at which point in the execution of a generic upload protocol a file should be considered successfully stored on the server. In particular, the adversary acting as the malicious server might claim that an upload failed, when it actually did store the corresponding file information which would make a later download successful. However, we only update the tracked file content in oracle UPD1 if the calling user is already an owner of the file. This reflects that the “expected content” (for the purpose of what constitutes a forgery) is only modified if an owner performs the update.
- We never remove file contents from the set of uploaded files under some file ID. This captures that we do not aim to protect against rollback attacks, i.e., a server serving an old copy of a file to a client does not count as an integrity attack.
- As in the confidentiality game, we add sharers as owners of the files they share. This prevents a trivial attack where the adversary uses a compromised user to share a file that is not yet tracked by the game with an honest user. Since the adversary knows the key material of the compromised user, it can create a file header containing a malicious file key which the sharer will decrypt and put in the OOB channel. Upon accepting the shared file, the honest user would then be the (only) owner of a file for which the adversary knows the file key, making a forgery trivial. Adding the sharer as an owner ensures that the file is correctly tracked as owned by a compromised user.

---

<sup>8</sup>We combine our convention for tables and sets and consider the table entry uninitialized while the sets are implicitly initialized to empty. That is, we let  $\text{fid} \notin T_f$  evaluate to **true** while  $T_f[\text{fid}] = (\emptyset, \emptyset)$ .

### 4.3 Designing a Security Model for Cloud Storage: UC vs. Game-Based

Having introduced our security definitions in detail, let us take a moment to explain our decision to propose *game-based* definitions for cloud storage instead of simulation-based security notions in the universal composability (UC) framework [15].

Many protocols with similar aims, particularly those in the multi-server password-protected secret sharing line of research ([7], [35], [14] [33]), analyze their security in the UC framework. This approach has obvious advantages: only one proof is needed to capture all security properties, the results can build upon the framework’s powerful composition theorem, and a proof in this model can argue about the security of users’ files relative to their passwords without attempting to quantify password security directly.

When we attempt to capture single-server cloud storage security, however, a commitment problem arises. In the UC framework, all honest clients are controlled by the environment, which chooses their inputs to the protocol. If the environment instructs an honest client  $C$  to store a file  $f$  with a malicious server, the simulator must produce, without knowledge of  $f$ , a ciphertext indistinguishable from the one  $C$  would send in an honest protocol interaction. If  $C$  should later become compromised, the simulator will learn  $f$  and must reveal some key under which the prior ciphertext decrypts to  $f$ . This is only possible if the cloud storage scheme employs some form of non-committing encryption.

A similar problem appears in the UC treatment of secure channels, where a client can send a message that must be decryptable after compromise [17]. Canetti et al. proposed as a solution an auxiliary “non-information oracle”, which takes in client messages (or files, in our setting) and outputs information to the simulator which is computationally independent from the input messages. Effectively, the non-information oracle produces committing ciphertexts for the simulator, and the computational independence of the ciphertexts from their plaintexts demonstrates that no leakage occurs. Upon client compromise, the non-information oracle can forward its internal keys and assist in decryption. This technique achieves UC channel security from standard primitives, and similar methods are used to capture secure messaging in the UC model [16].

Unfortunately, the non-information oracle strategy breaks down when applied to cloud storage due to our reliance on passwords as a basis for security rather than strong cryptographic keys. No matter how many layers of keys a protocol introduces between clients’ passwords and the ciphertexts they upload, a client must always be able to recover files using only its password as input by running the `auth` and `get` protocols. Consequently the ciphertexts cannot possibly be computationally independent from their underlying files. To achieve a UC security proof, we would instead require some computational secrecy property that relates the degree of leakage from the non-information oracle to the probability of a successful password-guessing attack. By moving the bulk of the simulation into an oracle with knowledge of client files and then making a separate computational argument about the secrecy of this oracle’s outputs, we effectively introduce a game-based confidentiality model. The advantages of the UC framework if we must also have games are significantly less. Thus, we focus solely on game-based notions in the remainder of this work.

### 4.4 Extensions, Limitations and Rationale of the Model

In the process of defining our cloud storage framework and security notions, there were many choices to be made for how to abstract and formalize real-world cloud storage schemes. Here, we provide rationale for certain aspects of our framework and why we did or did not include certain real-world facets in the model, and discuss the limitations that these choices imply. We also list possible extensions to the security model that could be interesting for future work.

PROTOCOL INPUTS. In our syntax, the execution of a cloud storage protocol only takes external input from the client and server in the first step of the respective party. (See Section 3.1.)

This is a design decision that we made for two reasons. First, it reduces the complexity of games. If we were to allow inputs at all protocol steps, it would be substantially more complex to track the security critical variables. For instance, our correctness and security games need to track the file identifier *fid* in file operations to check (trivial) win conditions. A game that supports taking *fid* as input at any protocol step would require a precondition to check that *fid* is not set multiple times (otherwise the adversary could interfere with the game-internal tracking). Additionally, the tracking and checks could only be performed once *fid* has been provided, and this time-point would now be protocol specific, making generic tracking more difficult.

Second, all protocols that we are aware of can be expressed with a single, initial input, even if it may seem like they require communication at first sight. For instance, consider a protocol to get a file that first fetches a file tree, presents it to the user to let them pick a file name, and then downloads the file with the ID corresponding to the name specified by the user. It may seem like the file ID is needed as a separate input in the second client-side step. However, we can easily model this protocol in our single input syntax by encoding the user action of translating a file name to a file ID into the client-side protocol. A user calls the file retrieval algorithm with the name of the file it wants to download. The client downloads the file tree, resolves the file name to a file ID, and fetches the file.

It is more challenging to accurately model a protocol that uses an out-of-band channel between server and client. (Such channels are not captured by our model, which only supports the client-to-client OOB channel used for file sharing.) An example of such a server-to-client channel could be email or SMS, as used e.g. for two factor authentication. If values exchanged through this channel appear as inputs in intermediate protocol steps, then these inputs cannot be captured by omitting the interaction or reordering (as in the file fetching example above), since they might be influenced by earlier protocol steps. We can either get around this limitation by only modeling the authentication after out-of-band material was already received (see our modeling of email tokens for MEGA in Appendix A.1), or by including the OOB message exchange in the protocol messages over the normal client-server channel. Both approaches do not affect C2C security as the adversary playing the server anyway sees all exchanged messages. However, more accurate modeling would be required for security against external adversaries, as discussed later in this section.

CLIENT-SIDE STATE. We use a single persistent client-side session state  $st_C$  (beyond the temporary per-protocol execution state  $st_C^{\text{tmp}}$ ). Therefore, when modeling real-world schemes, our state handling might sometimes need to differ from their actual implementation. For instance, consider a cloud storage scheme that stores the user password on the device after the first login. This increases usability at the cost of security as the user does not need to enter their credentials for every login. However, if the device is corrupted, then the password is compromised, even if there is currently no active session on the device. In our model, we would either model such a protocol by storing the password in the session state, or by requiring the user to enter it for each operation. Both approaches lead to the same adversarial strength in our model with respect to compromise, as the adversary can only compromise users completely, learning both their password and any session state. In other words, we cannot capture that not storing the password offers better security against device compromise.

A more fine-grained model could distinguish between user, device, and session states, capturing settings where one device runs multiple sessions in parallel, keeping some information in device memory and other information in per-session memory. This could allow to capture more fine-grained compromise notions, e.g., the security effects of a compromise of an old, revoked user device



on later-uploaded files, or that of a compromised session on one device on another, uncompromised session on the same device. We leave capturing such aspects to future work.

OUT-OF-BAND CHANNEL. Our model modularizes the out-of-band channel used for file sharing. This allows it to capture a wide range of real-world instantiations, such as sharing via a PKI or through links with encoded passwords. However, this modularization also means that within our confidentiality and integrity notions, the security of the out-of-band channel is idealized; it is assumed to be perfectly secure and incorruptible. A full security analysis of a concrete cloud storage system would hence need to separately prove the security of the actual channel used, and additionally analyze the security of its composition with the cloud storage scheme.

SECURITY AGAINST EXTERNAL ADVERSARIES. An alternative security notion to the client-to-client security given above would be to consider security against an *external* adversary that does not control the server but only plays the role of malicious clients. While this at first seems strictly weaker, in this model one could capture security properties which are trivially broken by a malicious cloud provider in the C2C setting, e.g.:

- *Limited password guessing:* an external adversary must interact with the server to verify password guesses online. In particular, the protocol does not leak any information to unauthenticated users that would allow offline password guesses. (In the C2C setting, the malicious server can always perform offline password guesses against the stored files of honest users.)
- *No rollback attacks:* an external adversary should not be able to serve old versions of a file to honest users. (In the C2C setting, rollback attacks cannot be prevented without additional assumptions, such as MPC, third parties, or direct device-to-device communication: a malicious server can always roll back the full server state and serve a previous version to a user who is logging in from a fresh device. We are not aware of any currently deployed E2EE cloud storage system that attempts to protect against such rollback attacks.)
- *Availability:* an external adversary should not be able to prevent honest users from fetching files which are not owned or shared with compromised or malicious users. (In the C2C setting, the malicious server can always reject to serve, or garble, the data of honest users.)
- *OOB channels:* real systems often use multiple channels between the server and the client, beyond the main TLS connection. For instance, using email or SMS as an out-of-band channel for multi-factor authentication is common practice. While the malicious server trivially knows the messages exchanged in these channels, the benefit of multiple channels becomes visible against an external adversary who may compromise some, but not all, channels.

We leave formalizing security against external adversaries as future work, but note that we kept the above properties in mind when designing the cloud storage protocol in Section 6.

ADDITIONAL FUNCTIONALITY. The core functionality considered in this paper could be extended – at the cost of higher complexity – with extra features such as file deletion, revocation of shared files, password recovery, session deauthentication (and ensuing security notions for adversaries with access to expired sessions), account deregistration, and distinguishing between read and write access to files. Out of these, more advanced access management (e.g., supporting revocation of access to shared files, or providing distinct access classes such as “owner”, “editor” or “read-only”) is especially interesting in the external adversary setting, where the server can be (at least partially) trusted to help enforce access restrictions.

METADATA. Protection of metadata, such as file names, types and paths, can be crucial for the overall security of a cloud storage system. For instance, insufficient encryption of file names or paths would leak information about the uploaded file to the server. Moreover, not protecting the integrity of metadata may directly affect the file integrity (e.g., when a file length field is modified).



In our security notions, we implicitly model strong security guarantees for metadata by treating it as part of the “file” plaintext input to oracle PUT1. This implies that the metadata must be confidential and integrity-protected even from a malicious server. Note that protocols which encrypt metadata separately from files for efficiency can still be captured in this model (e.g., by treating the metadata as its own “file”).

For the sake of complexity, we leave the modeling of weaker security guarantees of metadata – such as “authentication-only” – and the impact thereof on file security to future work. Moreover, our CSS protocol does not aim to protect against usage metadata, such as file access patterns and share relationships. Constructing a protocol that achieves the metadata-hiding properties of Metal [22] and Titanium [21] while having the provable guarantees of our detailed security models – including password-based security, a single fully malicious server, non-atomic operations, and considering multiple devices – is an open problem.

CHANNEL ASSUMPTIONS. Additional work could explore different channel assumptions. For instance, only assuming an authenticated but not confidential communication channel would give the adversary access to the exchanged messages. This may be interesting as a standalone solution, or as defense-in-depth measure to protect against implementation failures that occasionally undermine confidentiality guarantees in practice. For instance, implementation failures of TLS session tickets lead to the loss of confidentiality of early data for some implementations [29].

Moreover, specific instantiations of the OOB channel, especially those operated by the malicious server, could be integrated into CSS and analyzed.

ADVANCED SECURITY GUARANTEES. It is an open question how post-compromise and forward security apply to the cloud storage setting, and what is achievable; key rotation has been considered in prior works [36], as has fine-grained forward security via puncturing techniques [5, 53], but not in combination with the full feature set we target. Furthermore, CSS does not attempt to protect against partial rollback attacks. For instance, future work could attempt to protect the integrity of the user’s cloud storage state to prevent a malicious server from selectively serving some old and some new files.

## 5 Formally Breaking MEGA

To put our new cloud storage framework to the test, we will give both a negative analysis, capturing some of the recent attacks against the MEGA cloud storage system from [6, 3, 30] in our framework (in this section), and a positive one, presenting a provably secure cloud storage scheme (in the Section 6).

MEGA provides end-to-end encrypted cloud storage since 2013 and is currently the largest provider with E2EE by default, with roughly 300 million users [40]. In addition to its popularity, MEGA is an interesting system to study because recent cryptanalysis [6, 3, 30] showed that clients released prior to June 2022 [44] were vulnerable to a range of attacks by a malicious server. In the following, we show that MEGA can be cast in our syntax, demonstrating the expressiveness of our framework, and that recent attacks [6, 3, 30] arise in our framework as formal violations of the C2C security notions (already for selective security, as the attacks do not rely on client compromise). While the attacks discussed here target the unpatched version of MEGA, the patched protocol still cannot be proven secure in our model due to various design decisions, including key reuse and the lack of binding between file IDs and files.

In this section, we only highlight the core, vulnerable protocol components and how the main attack ideas can be cast in our security framework. We defer the full details of formalizing the MEGA protocols, attacks, and barriers for proving the patched protocol secure to Appendix A.

<pre> : client: <math>M_{\text{auth}}^{(C:3)}(st_C, st_C^{\text{tmp}}, m_S)</math>: 1 <math>(c_m, (pk, c_{rsa}), c_{sid}) \leftarrow m_S</math> 2 <math>k_m \leftarrow \text{AES-ECB.DEC}(k_e, c_m)</math> 3 <math>sk \leftarrow \text{AES-ECB.DEC}(k_m, c_{rsa})</math> 4 <math>st_C.sid \leftarrow \text{RSA.DEC}(sk, c_{sid})</math> ; <math>st_C.k_m \leftarrow k_m</math> 5 <math>st_C.pk \leftarrow pk</math> ; <math>st_C.sk \leftarrow sk</math> 6 <b>return</b> <math>(st_C, st_C^{\text{tmp}}, out_C, st_C.sid)</math> <math>M_{\text{auth}}^{(C:4)}(st_C, st_C^{\text{tmp}}, m_S)</math>: 7 <b>if</b> <math>m_S = \perp</math>: <b>fail</b><sub>C</sub> <b>else</b>: <b>success</b><sub>C</sub> </pre>	<pre> : server: <math>M_{\text{auth}}^{(S:2)}(st_S, st_S^{\text{tmp}}, m_C)</math>: 8 <math>aid \leftarrow st_S^{\text{tmp}}.aid</math> 9 <math>(n_C, c_m, h_a, pk, c_{rsa}) \leftarrow st_S.acc[aid]</math> 10 ... // omitted 11 <math>c_{sid} \leftarrow \text{RSA.ENC}(pk, sid)</math> 12 <math>m_S \leftarrow (c_m, pk, c_{rsa}, c_{sid})</math> 13 <b>return</b> <math>(st_S, st_S^{\text{tmp}}, \varepsilon, m_S)</math> <math>M_{\text{auth}}^{(S:3)}(st_S, st_S^{\text{tmp}}, m_C)</math>: 14 <b>if</b> <math>m_C \neq sid</math>: <b>fail</b><sub>S</sub> 15 <math>st_S.sess[sid] \leftarrow st_S^{\text{tmp}}.aid</math> 16 <b>success</b><sub>S</sub> </pre>
---	---

Figure 7: The part of MEGA’s authentication protocol  $M_{\text{auth}}$  hosting the core vulnerabilities; see Figure 12 in Appendix A.1 for the full protocol.

THE VULNERABLE CORE OF MEGA, FORMALIZED. Highly simplified, in the account registration protocol  $M_{\text{reg}}$ , a client generates a uniform, symmetric master key  $k_m$  and an RSA key pair  $(pk, sk)$ . It then encrypts  $sk$  under  $k_m$ , and  $k_m$  in turn under a password-derived key  $k_e$ , using AES-ECB encryption for both, and uploads the resulting ciphertexts  $c_{rsa}$  and  $c_m$ , respectively, to the server.

During authentication ( $M_{\text{auth}}$ , see Figure 7), the server sends  $c_{rsa}$  and  $c_m$  back to the client, along with another ciphertext  $c_{sid}$  which encrypts a server-chosen session ID  $sid$  under the client’s RSA public key. The client decrypts, in order,  $c_m$ ,  $c_{rsa}$ , and  $c_{sid}$ , and sends the obtained session ID value back to the server to complete authentication.

To upload a new file ( $M_{\text{put}}$ ), the MEGA client generates a new file key  $k_f$  and encrypts the file using a custom-made authenticated encryption algorithm. For the attack, the key thing to note is that the file key is encrypted (alongside a nonce and an authentication tag) under the master key  $k_m$  using AES-ECB and stored on the server, just like the user’s RSA private key.

The authentication protocol hosts the core vulnerability; we hence show its relevant steps in Figure 7. The full details of all protocols are in Appendix A.1.

BREAKING CONFIDENTIALITY. The attacks from [3, 6, 30] now translate into formal violations of our security games for the MEGA cloud storage scheme. We summarize the attacks and their formalization in our framework here and provide the full pseudocode of the adversaries for attacks 1–3 from [6] (discussing aspects from [3, 30]) in Appendix A.2.

We begin by showing how the RSA key and plaintext recovery attacks by Backendal et al. [6] break C2CConfS confidentiality, by combining them into an adversary  $\mathcal{A}_{\text{conf}}$  that wins the  $\mathbf{G}_{M,n,PW_n}^{\text{C2CConfS}}$  game (see Definition 4.2) with high probability. The adversary starts by registering an honest user and lets them upload two distinct files as challenge. It then runs a simpler version of the RSA key recovery attack [6, Section III] by letting the user repeatedly authenticate, tampering with the AES-ECB ciphertext  $c_{rsa}$  and exploiting the returned session ID  $sid$  as a partial decryption oracle. For the latter, observe that  $\mathcal{A}_{\text{conf}}$ , acting as the malicious server in the  $M_{\text{auth}}$  protocol, can pick an arbitrary ciphertext (on line 12 of Figure 7) which the honest client will decrypt using the maliciously modified RSA secret key  $sk$ , and send back in the client message  $m_C$  of  $M_{\text{auth}}^{(S:3)}$ . This allows  $\mathcal{A}_{\text{conf}}$  to recover the user’s  $sk$  through a binary search.

Knowing the user’s RSA secret key, the adversary can then mount the AES-ECB decryption attack from [6, Section IV] to decrypt the challenge file. On a high level,  $\mathcal{A}_{\text{conf}}$  injects the encrypted challenge file key into the RSA ciphertext  $c_{rsa}$  sent in the authentication protocol. This

is possible because both the RSA key and file keys are encrypted with the master key  $k_m$  using AES-ECB, enabling the replacement of individual ciphertext blocks in  $c_{rsa}$ . This second attack step is probabilistic and fails with probability at most  $2^{-30}$  (see [6]). Adversary  $\mathcal{A}_{conf}$  wins with advantage  $\text{Adv}_{M,n,\mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A}_{conf}) = 1 - 2^{-30}$ , making 1023 calls to  $M_{auth}$ . (The number of oracle calls to  $M_{auth}$  needed to mount the attack was since reduced to 6 by Heninger and Ryan [30].)

**BREAKING INTEGRITY.** To violate C2CIntS integrity, we formalize the integrity attack from [6, Section V.B.2]. Instead of assuming the knowledge of a plaintext-ciphertext pair (which is required for their attack) we integrate the AES-ECB encryption oracle identified by Albrecht et al. [3, Section 2.2] to exercise the full scope of our framework. This allows the adversary to obtain the ciphertext for a chosen plaintext via the sharing protocol.

The C2CIntS adversary  $\mathcal{A}_{int}$  first honestly registers and authenticates a user, and then makes them accept a shared file chosen by  $\mathcal{A}_{int}$  (which, of course, does not count as win yet). The client will re-encrypt the shared file key under its master key using AES-ECB, at which point the adversary knows both the plaintext and ciphertext of the file key. Adversary  $\mathcal{A}_{int}$  can exploit that MEGA's custom decryption algorithm causes the client to derive an all-zero file key during file downloads when an AES-ECB block is duplicated in the file key ciphertext. Knowing the decryption of one ciphertext block and carefully crafting an encrypted file (see Appendix A.2 and [6, Section V.B.3]),  $\mathcal{A}_{int}$  succeeds in making the user decrypt a forged file and thus wins the integrity game. This adversary  $\mathcal{A}_{int}$  has advantage  $\text{Adv}_{M,n,\mathcal{PW}_n}^{\text{C2CIntS}}(\mathcal{A}_{int}) = 1$  with a total of only four oracle calls.

In conclusion, we observe that problems with the E2EE cloud storage protocol of MEGA show up naturally when it is written down in our syntax and tested against our security games.

## 6 A Provably Secure Cloud Storage Scheme

In this section, we showcase the instantiability of our framework by presenting a concrete cloud storage scheme CSS which enjoys provable client-to-client security. The scheme provides all functionality from the core protocols, including file sharing, and supports password rotation without file re-encryption. We first introduce the building blocks used in our construction, before explaining the scheme itself and then analyzing its security. Appendix 6.5 presents limitations and extensions of the scheme, including a discussion of selective vs. adaptive security and considerations in practice.

### 6.1 Building Blocks

**PRF, AEAD AND MAC.** We use the following basic building blocks, further specified in Appendix B.1, in our cloud storage scheme:

- a PRF  $F: \{0, 1\}^{kl} \times \{0, 1\}^* \rightarrow \{0, 1\}^{kl}$  taking a key  $k \in \{0, 1\}^{kl}$  and a variable-length label  $x \in \{0, 1\}^*$  as inputs and returning a string of length  $kl$ .
- a message authentication code  $\text{MAC} = (\text{Tag}, \text{Vrfy})$  with key space  $\{0, 1\}^{kl}$ .
- a nonce-based authenticated encryption scheme with associated data  $\text{AEAD} = (\text{Enc}, \text{Dec})$ , also with key space  $\{0, 1\}^{kl}$ . We write  $c \leftarrow \text{Enc}(k, n, m, ad)$  to encrypt message  $m \in \{0, 1\}^*$  using key  $k \in \{0, 1\}^{kl}$ , nonce  $n \in \{0, 1\}^{nl}$ , and associated data  $ad \in \{0, 1\}^*$ , and  $m \leftarrow \text{Dec}(k, n, c, ad)$  for the corresponding decryption of the ciphertext  $c \in \{0, 1\}^*$ .

Note that, for ease of presentation, we assume that the output length  $kl$  of the PRF is the key length of the MAC and AEAD scheme.

**2HASHDH OPRF.** Our construction further uses 2HashDH [32, 34], an oblivious pseudorandom function (OPRF) [47, 26], to derive a hardened secret from the user password. An OPRF is an

interactive protocol between two parties (henceforth called client and server), which allows the client, holding a secret input  $x$ , to compute a PRF evaluation  $y = F(k, x)$  with the help of the server, which holds the key  $k$ . In our use, we let the client sample the OPRF key  $k$  during account registration and subsequently send it to the server; this ensures that the initial client output is computed with an honestly generated key.

We briefly recap the 2HashDH OPRF here. Let  $G$  be a group of prime order  $q$ , and let  $H_1: \{0, 1\}^* \rightarrow G$  and  $H_2: \{0, 1\}^* \times G \rightarrow \{0, 1\}^{kl}$  be two hash functions. The 2HashDH OPRF interactively computes  $H_2(x, H_1(x)^k)$  for the client secret input  $x \in \{0, 1\}^*$  and server-side key  $k \in \mathbb{Z}_q$ . In the interactive protocol, the input  $x$  is hidden from the server with a random blinding value  $r \leftarrow \mathbb{Z}_q$ . The client sends  $\alpha \leftarrow H_1(x)^r$  to the server, which computes the blind evaluation  $\beta \leftarrow \alpha^k$ . The result is recovered on the client by unblinding  $\beta$  using  $r$ , computing  $y \leftarrow H_2(x, \beta^{1/r})$ . We give the full specification in our syntax in Appendix B.2.

The goal of using an OPRF in our cloud storage protocol is to provide a means for the client to derive a key from the user password, which is secure against a malicious server. This is a somewhat non-standard setting for an OPRF, prompting us to adapt 2HashDH in the following ways (further discussed in Section 6.3):

- The client secret  $x$  consists of  $(aid, pw)$ , i.e., the combination of the *public* account ID and the user password. Thanks to the assumption that account IDs are unique, this implies that client secrets are unique.
- The OPRF key is sampled by the client during registration, and subsequently sent to the server; this ensures that the initial client output is computed with an honestly generated key.
- $H_1$  and  $H_2$  should ideally be memory-hard and slow, to increase the work required for a password-guessing attack from a malicious server.

## 6.2 Scheme Overview

We specify the protocols that make up our cloud storage scheme  $CSS = (CSS_{\text{reg}}, CSS_{\text{auth}}, CSS_{\text{put}}, CSS_{\text{upd}}, CSS_{\text{get}}, CSS_{\text{share}}, CSS_{\text{accept}})$ , in pseudocode in Figures 8–10. Let us first explain the core design rationale behind our cloud storage scheme, before describing the algorithms in more detail.

REGISTRATION AND AUTHENTICATION. To register an account, a user first derives a hardened secret  $rw$  from their account ID and password using the 2HashDH OPRF. They then AEAD-encrypt a randomly generated master key under a key derived from  $rw$ . To complete registration, the client uploads the encrypted master key, the OPRF server key and a MAC key derived from  $rw$  to the server.

To authenticate, a user evaluates the 2HashDH OPRF on their account ID and password with the help of the server(-side key) to recover  $rw$ , and from it the master key and MAC key. The server assigns a fresh session ID ( $sid$ ) to the client and uses the MAC key to verify that the client can produce a valid tag for the  $sid$  (and interaction transcript) as a means of authentication. Subsequent requests in the same session are authenticated by sending along the session ID.

FILE HANDLING. Our scheme follows the established approach of using a key hierarchy, employing the master key as key encryption key to wrap file keys, which in turn encrypt single files. Both file keys and files are AEAD-encrypted, bound to the owner and file ID, respectively. Each encrypted file is stored along with a header containing the wrapped file key. Importantly, we use the independently sampled master key as the root of the key-wrapping hierarchy in our cloud storage scheme instead of directly using the raw secret  $rw$  derived from the password; this way, users can change their password without having to re-encrypt all their data. The server stores the encrypted files and headers and maintains a record of owners to perform access control.

SHARING AND UPDATES. To share access to a file with another registered user, an owner shares the file key with the recipient out-of-band and informs the server about the added owner. The newly added user can then wrap the file key under its own master key and add it to the file header, allowing regular future access to the file.

The separation between headers and encrypted file content allows lightweight updates of shared files: any owner can update the encrypted file content while preserving the file key, thereby retaining access for all other owners.

### 6.3 Detailed Description

NOTATIONAL CONVENTIONS. In the pseudocode description of the CSS scheme in Figures 8–10, we use the convention that subroutines returning  $\perp$  cause their callers to abort by running the appropriate choice of **fail<sub>C</sub>** or **fail<sub>S</sub>**.

ACCOUNT REGISTRATION. In the account registration protocol  $\text{CSS}_{\text{areg}}$ , shown at the top of Figure 8, the client samples a 2HashDH OPRF key  $k$  and locally computes the OPRF value  $rw$  on its account ID  $aid$  and password  $pw$ . Involving the unique account ID here ensures uniqueness of the OPRF input (even if passwords are reused across users), preventing a malicious server from serving encrypted files across users that it suspects to share the same password. In a sense, we use the pair  $(aid, pw)$  as what was the password input in prior usage of 2HashDH [32, 34]; this allows us to avoid assuming that passwords are unique – an assumption that was necessary in other password-based encryption work [12], but is problematic in practice – and only make the milder assumption that (honest) users use unique account IDs, such as email addresses.

The client then derives a key encryption key  $k_{kek}$  and MAC key  $k_{mac}$  from  $rw$  using a PRF  $F$ , samples a symmetric master key  $k_{mk}$  uniformly at random, and AEAD-encrypts  $k_{mk}$  under  $k_{kek}$  (bound to the account ID via the associated data field). It sends its  $aid$ , the OPRF key  $k$ , the MAC key  $k_{mac}$ , and the computed ciphertext and used nonce to the server for storage in its account registry  $acc$ .

AUTHENTICATION. In the authentication protocol  $\text{CSS}_{\text{auth}}$  (Figure 8 bottom), the client sends the blinded hash  $\alpha = H_1(aid, pw)^r$  to initiate a 2HashDH OPRF run. The server computes the 2HashDH evaluation  $\beta = \alpha^k$  and samples a random, unused session ID  $sid$  (of bit-length  $sidl$ ), sending both to the client. The client unblinds the OPRF value to obtain its raw secret  $rw$  and derives  $k_{kek}$  and  $k_{mac}$  from it. It then authenticates to the server by sending a MAC value over the first two protocol messages (in particular containing  $sid$  as random challenge), using key  $k_{mac}$ . Upon successful verification of the MAC, the server adds the client to its session registry  $sess$  and sends back the client’s encrypted master key. The client AEAD-decrypts its master key  $k_{mk}$  using  $k_{kek}$  (aborting implicitly upon failure) and saves the session state  $st_C$  containing  $aid$ ,  $sid$ , and  $k_{mk}$ .

CORE FILE OPERATIONS. The core operations for handling files in the protocols  $\text{CSS}_{\text{put}}$ ,  $\text{CSS}_{\text{upd}}$ ,  $\text{CSS}_{\text{get}}$ ,  $\text{CSS}_{\text{share}}$ , and  $\text{CSS}_{\text{accept}}$  (all in Figure 9) is captured through subroutines for getting/downloading and putting/uploading headers and contents of files, shown in Figure 10.

**putFile, putHeader.** In the core operation for putting a file or header, the client AEAD-encrypts the given file resp. file key under the file resp. master key, to be sent to the server. The server ensures that the client user is an owner of the file ( $aid \in \text{ows}[fid]$ ) and then stores the encryption in its file or file-key registry ( $fs$  and  $fks$ ), respectively.

**getFile, getHeader.** Conversely, for fetching a file or header, the client first sends the corresponding file ID  $fid$ . The server ensures that the client user owns the requested file and then returns the



encrypted header and (for `getFile`) file content from its registries. The client AEAD-decrypts the header ciphertext  $c_{fk}$  using its master key, and, for a file download, the file ciphertext  $c_f$  using the resulting file key.

In each of these subroutines, the client sends its session ID  $sid$  as authentication token along with its message. The server uses  $sid$  to associate the request to an authenticated user account  $aid \leftarrow sess[sid]$ , implicitly failing if  $sid \notin sess$ .

**UP- AND DOWNLOAD.** The complete upload of a file ( $CSS_{put}$ , first box in Figure 9) consists of the client sampling a fresh AEAD file key  $k_f \leftarrow_{\$} \{0,1\}^{kl}$  and uploading that key as header (`putHeader`) along with the file content (`putFile`); the server in turn ensures that this file ID is unused ( $fid \notin ows$ ) and then registers the client user as owner and stores the encrypted header and file content. To update a file ( $CSS_{upd}$ , second box in Figure 9), the client first fetches the file key  $k_f$  via `getHeader` and then uses it to upload the updated file content (`putFile`). Downloading a file ( $CSS_{get}$ , third box in Figure 9) uses `getFile` to retrieve the file content.

**SHARING.** Our CSS scheme allows sharing of files with other, registered users. To this end, a client invokes the  $CSS_{share}$  protocol (fourth box in Figure 9) on input a file ID  $fid$  and the account ID  $raid$  of the receiving user with whom that file shall be shared. The sharing client fetches the file key  $k_f$  via `getHeader`, upon whose success the server adds  $raid$  to the set of owners for file  $fid$ . The client can then share  $k_f$  out-of-band with the receiving user by setting  $out_C.oob \leftarrow k_f$ . We intentionally leave the details of the out-of-band channel open; we only assume that it is inaccessible to the malicious cloud adversary (except for compromised clients), and aim to minimize its usage. Possible instantiations include using a secure messaging app or an external PKI infrastructure.

To accept the sharing of a file, a client in the  $CSS_{accept}$  protocol (last box in Figure 9) uses the `putHeader` subroutine to upload the shared file key obtained through the out-of-band channel ( $in_C.oob$ ), encrypted under the user’s master key, as a new header for the shared file ID  $fid$ . That is, each file owner has its own file header (which all contain the same file key, wrapped under their individual master keys), but the actual file ciphertext is shared among all owners. Hence,  $CSS_{upd}$  reuses the file key to ensure that access to shared files is maintained across updates. We do not model an “unshare” operation; however, any owner can delete a file by overwriting it with  $\perp$ .

## 6.4 Proving CSS Secure

We are now ready to present our fully quantified bounds on the selective confidentiality and integrity of CSS. They are dominated by the probability of a successful offline dictionary attack on user passwords, which is inherent in the malicious-server setting due to unavoidable guessing attempts. That is, since the server possesses ciphertexts created by the user under a key derivable from their password, a malicious server can attempt to brute-force the password and use the ciphertexts to verify guesses without any client interaction. The only protection against this type of attack is hence to choose a strong password; this is reflected in our theorem bounds by the term corresponding to the password-guessing advantage against the password distribution from which users sample their passwords in our model.

As we outlined in the introduction, our proofs pertain to our selective security notions rather than the adaptive notions because of the commitment issues that arise with adaptive compromises. For CSS, in particular, the most challenging issues arise because the challenge files can be shared with multiple users. In the proof, this means that the reduction needs to replace the master keys of all owners of challenge files with random in order for us to be able to reduce to the security of the AEAD scheme used for file key encryption in the next step. This replacement needs to happen during registration, since the master key is generated in that step and then stored in encrypted form

$\text{CSS}_{\text{areg}}^{(C:1)}(\varepsilon, \varepsilon, in_C, \varepsilon) :$ 1 <b>pre:</b> $in_C.\{aid, pw\}$ 2 $k \leftarrow \mathbb{Z}_q$ 3 $rw \leftarrow H_2(aid, pw, H_1(aid, pw)^k)$ 4 $k_{kek} \leftarrow F(rw, \text{"kek"})$ 5 $k_{mk} \leftarrow \{0, 1\}^{kl}; n_{mk} \leftarrow \{0, 1\}^{nl}$ 6 $c_{mk} \leftarrow \text{Enc}(k_{kek}, n_{mk}, k_{mk}, (aid, \text{"mk"}))$ 7 $k_{mac} \leftarrow F(rw, \text{"mac"})$ 8 $m_C \leftarrow (aid, k, n_{mk}, c_{mk}, k_{mac})$ 9 <b>success</b> $_C(m_C)$	$\text{CSS}_{\text{areg}}^{(S:1)}(st_S, \varepsilon, in_S, m_C) :$ 10 <b>pre:</b> $st_S.\{acc\}$ 11 $(aid, k, n_{mk}, c_{mk}, k_{mac}) \leftarrow m_C$ 12 $acc[aid] \leftarrow (k, n_{mk}, c_{mk}, k_{mac})$ 13 <b>success</b> $_S$
$\text{CSS}_{\text{auth}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon) :$ 14 <b>pre:</b> $in_C.\{aid, pw\}$ 15 $r \leftarrow \mathbb{Z}_q; \alpha \leftarrow H_1(aid, pw)^r$ 16 $m_C \leftarrow (aid, \alpha)$ 17 $st_C^{\text{tmp}} \leftarrow (aid, pw, r, m_C)$ 18 <b>return</b> $(st_C, st_C^{\text{tmp}}, \varepsilon, m_C)$  $\text{CSS}_{\text{auth}}^{(C:2)}(st_C, st_C^{\text{tmp}}, m_S) :$ 19 $(st_C^{\text{tmp}}.sid, \beta) \leftarrow m_S$ 20 $(aid, pw, r, m'_C) \leftarrow st_C^{\text{tmp}}$ 21 <b>if</b> $\beta \notin G$ : <b>fail</b> $_C$ 22 $\gamma \leftarrow \beta^{1/r}; rw \leftarrow H_2(aid, pw, \gamma)$ 23 $st_C^{\text{tmp}}.k_{kek} \leftarrow F(rw, \text{"kek"})$ 24 $k_{mac} \leftarrow F(rw, \text{"mac"})$ 25 $m_C \leftarrow \text{Tag}(k_{mac}, (m'_C, m_S))$ 26 <b>return</b> $(st_C, st_C^{\text{tmp}}, \varepsilon, m_C)$  $\text{CSS}_{\text{auth}}^{(C:3)}(st_C, st_C^{\text{tmp}}, m_S) :$ 27 <b>pre:</b> $st_C^{\text{tmp}}.\{k_{kek}, aid, sid\}$ 28 $(n_{mk}, c_{mk}) \leftarrow m_S$ 29 $k_{mk} \leftarrow \text{Dec}(k_{kek}, n_{mk}, c_{mk}, (aid, \text{"mk"}))$ 30 $st_C \leftarrow (aid, sid, k_{mk})$ 31 <b>success</b> $_C$	$\text{CSS}_{\text{auth}}^{(S:1)}(st_S, \varepsilon, in_S, m_C) :$ 32 <b>pre:</b> $st_S.\{acc, sess\}$ 33 $(aid, \alpha) \leftarrow m_C$ 34 <b>if</b> $(aid \notin acc) \vee (\alpha \notin G)$ : <b>fail</b> $_S$ 35 $(k, n_{mk}, c_{mk}, k_{mac}) \leftarrow acc[aid]$ 36 $\beta \leftarrow \alpha^k$ 37 $sid \leftarrow \{0, 1\}^{sidl}$ s.t. $sid \notin sess$ 38 $m_S \leftarrow (sid, \beta)$ 39 $st_S^{\text{tmp}} \leftarrow (aid, k_{mac}, m_C, m_S)$ 40 <b>return</b> $(st_S, st_S^{\text{tmp}}, \varepsilon, m_S)$  $\text{CSS}_{\text{auth}}^{(S:2)}(st_S, st_S^{\text{tmp}}, m_C) :$ 41 <b>pre:</b> $st_S.\{sess\}$ 42 $(aid, k_{mac}, m'_C, m'_S) \leftarrow st_S^{\text{tmp}}$ 43 <b>if</b> $\neg \text{Vrfy}(k_{mac}, (m'_C, m'_S), m_C)$ : <b>fail</b> $_S$ 44 $sess[sid] \leftarrow aid$ 45 $m_S \leftarrow (n_{mk}, c_{mk})$ 46 <b>success</b> $_S(m_S)$

Figure 8: Account registration and authentication protocols of CSS.

on the server. But in the adaptive game, the reduction does not know at the time of registration whether a user will become the owner of a challenge file later, or be compromised. Hence we run into a commitment issue: the reduction needs to commit to a ciphertext corresponding to the user master key during registration, but it does not yet know whether it will later have to open this commitment (if the user is compromised) or if this ciphertext should be independent of the actual master key (if the user is made the owner of a challenge file).

We are not aware of any attacks on CSS in the adaptive setting. We hence hypothesize that this is a proof technique issue, rather than an issue with the protocol. Indeed, one could possibly

$\text{CSS}_{\text{put}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon):$ 1 <b>pre:</b> $in_C.\{fid, f\} \ st_C.\{sid\}$ 2 $k_f \leftarrow \{0, 1\}^{kl}$ 3 $m_{fk} \leftarrow \text{putHeader}^{(C:1)}(st_C, fid, k_f)$ 4 $m_f \leftarrow \text{putFile}^{(C:1)}(st_C, fid, k_f, f)$ 5 $m_C \leftarrow (fid, m_f, m_{fk})$ 6 <b>success</b> $_C(m_C)$	$\text{CSS}_{\text{put}}^{(S:1)}(st_S, \varepsilon, in_S, m_C):$ 7 <b>pre:</b> $st_S.\{ows, sess\}$ 8 $(fid, m_f, m_{fk}) \leftarrow m_C$ 9 <b>if</b> $fid \in ows$ : <b>fail</b> $_S$ 10 $ows[fid] \leftarrow \{sess[sid]\}$ 11 $\text{putHeader}^{(S:1)}(st_S, m_{fk})$ 12 $\text{putFile}^{(S:1)}(st_S, m_f)$ 13 <b>success</b> $_S$
$\text{CSS}_{\text{upd}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon):$ 14 <b>pre:</b> $in_C.\{fid, f\}, \ st_C.\{sid\}$ 15 $m_C \leftarrow \text{getHeader}^{(C:1)}(st_C, fid)$ 16 $st_C^{\text{tmp}}.\{fid, f\} \leftarrow \{fid, f\}$ 17 <b>return</b> $(st_C, st_C^{\text{tmp}}, \varepsilon, m_C)$  $\text{CSS}_{\text{upd}}^{(C:2)}(st_C, st_C^{\text{tmp}}, m_S):$ 18 <b>pre:</b> $st_C^{\text{tmp}}.\{fid, f\}$ 19 $k_f \leftarrow \text{getHeader}^{(C:2)}(st_C, fid, m_S)$ 20 $m_C \leftarrow \text{putFile}^{(C:1)}(st_C, fid, k_f, f)$ 21 <b>success</b> $_C(m_C)$	$\text{CSS}_{\text{upd}}^{(S:1)}(st_S, \varepsilon, in_S, m_C):$ 22 $m_S \leftarrow \text{getHeader}^{(S:1)}(st_S, m_C)$ 23 <b>return</b> $(st_S, \varepsilon, \varepsilon, m_S)$  $\text{CSS}_{\text{upd}}^{(S:2)}(st_S, st_S^{\text{tmp}}, m_C):$ 24 $\text{putFile}^{(S:1)}(st_S, m_C)$ 25 <b>success</b> $_S$
$\text{CSS}_{\text{get}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon):$ 26 <b>pre:</b> $in_C.\{fid\}, \ st_C.\{sid\}$ 27 $m_C \leftarrow \text{getFile}^{(C:1)}(st_C, fid)$ 28 $st_C^{\text{tmp}}.fid \leftarrow fid$ 29 <b>return</b> $(st_C, \varepsilon, \varepsilon, m_C)$  $\text{CSS}_{\text{get}}^{(C:2)}(st_C, st_C^{\text{tmp}}, m_S):$ 30 <b>pre:</b> $st_C.\{aid\}, \ st_C^{\text{tmp}}.\{fid\}$ 31 $f \leftarrow \text{getFile}^{(C:2)}(st_C, fid, m_S)$ 32 $out_C.f \leftarrow f$ 33 <b>success</b> $_C$	$\text{CSS}_{\text{get}}^{(S:1)}(st_S, \varepsilon, in_S, m_C):$ 34 $m_S \leftarrow \text{getFile}^{(S:1)}(st_S, m_C)$ 35 <b>success</b> $_S(m_S)$
$\text{CSS}_{\text{share}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon):$ 36 <b>pre:</b> $in_C.\{raid, fid\} \ st_C.\{sid\}$ 37 $m'_C \leftarrow \text{getHeader}^{(C:1)}(st_C, fid)$ 38 $st_C^{\text{tmp}}.fid \leftarrow fid$ 39 <b>return</b> $(st_C, st_C^{\text{tmp}}, \varepsilon, (m'_C, raid, fid))$  $\text{CSS}_{\text{share}}^{(C:2)}(st_C, st_C^{\text{tmp}}, m_S):$ 40 <b>pre:</b> $st_C^{\text{tmp}}.\{fid\}$ 41 $k_f \leftarrow \text{getHeader}^{(C:2)}(st_C, fid, m_S)$ 42 $out_C.oob \leftarrow k_f$ // share file key out-of-band 43 <b>success</b> $_C$	$\text{CSS}_{\text{share}}^{(S:1)}(st_S, \varepsilon, in_S, m_C):$ 44 <b>pre:</b> $st_S.\{ows\}$ 45 $(m'_C, raid, fid) \leftarrow m_C$ 46 $m_S \leftarrow \text{getHeader}^{(S:1)}(st_S, m'_C)$ 47 $ows[fid] \leftarrow \overset{\cup}{\cup} \{raid\}$ // authorize receiver 48 <b>success</b> $_S(m_S)$
$\text{CSS}_{\text{accept}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon):$ 49 <b>pre:</b> $in_C.\{fid, oob\} \ st_C.\{sid\}$ // $k_f$ received in oob 50 $m_C \leftarrow \text{putHeader}^{(C:1)}(st_C, fid, oob)$ 51 <b>success</b> $_C(m_C)$	$\text{CSS}_{\text{accept}}^{(S:1)}(st_S, \varepsilon, in_S, m_C):$ 52 $\text{putHeader}^{(S:1)}(st_S, m_C)$ 53 <b>success</b> $_S$

Figure 9: File download, upload, update, share, and accept protocols for CSS.



$\boxed{\text{putFile}} \boxed{\text{putHeader}}^{(C:1)}(st_C, fid, k_f, f):$ 1 $n \leftarrow \$\{0, 1\}^{nl}$ 2 $md \leftarrow (st_C.aid, fid, \text{"fk"})$ 3 $c_{fk} \leftarrow \text{Enc}(st_C.k_{mk}, n, k_f, md)$ 4 $c_f \leftarrow \text{Enc}(k_f, n, f, (fid, \text{"file"}))$ 5 $m_C \leftarrow (st_C.sid, fid, n, c_{fk}, c_f)$ 6 <b>return</b> $m_C$	$\boxed{\text{putFile}} \boxed{\text{putHeader}}^{(S:1)}(st_S, m_C):$ 7 <b>pre:</b> $st_S.\{sess, ows, fks, fs\}$ 8 $(sid, fid, n, c_{fk}, c_f) \leftarrow m_C$ 9 $aid \leftarrow sess[sid]$ 10 <b>if</b> $aid \notin ows[fid]$ : <b>return</b> $\perp$ 11 $fks[aid, fid] \leftarrow (n, c_{fk})$ 12 $fs[fid] \leftarrow (n, c_f)$
$\boxed{\text{getFile}} \boxed{\text{getHeader}}^{(C:2)}(st_C, fid):$ 13 $m_C \leftarrow (st_C.sid, fid)$ 14 <b>return</b> $m_C$  $\boxed{\text{getFile}} \boxed{\text{getHeader}}^{(C:2)}(st_C, fid, m_S):$ 15 $((n_{fk}, c_{fk}), (n_f, c_f)) \leftarrow m_S$ 16 $md \leftarrow (st_C.aid, fid, \text{"fk"})$ 17 $k_f \leftarrow \text{Dec}(st_C.k_{mk}, n_{fk}, c_{fk}, md)$ 18 <b>return</b> $k_f$ 19 $f \leftarrow \text{Dec}(k_f, n_f, c_f, (fid, \text{"file"}))$ 20 <b>return</b> $f$	$\boxed{\text{getFile}} \boxed{\text{getHeader}}^{(S:1)}(st_S, m_C):$ 21 <b>pre:</b> $st_S.\{sess, ows, fks, fs\}$ 22 $(sid, fid) \leftarrow m_C$ 23 $aid \leftarrow sess[sid]$ 24 <b>if</b> $aid \notin ows[fid]$ : <b>return</b> $\perp$ 25 $m_S \leftarrow (fks[aid, fid], fs[fid])$ 26 <b>return</b> $m_S$

Figure 10: Subroutines for putting and getting header and/or files for CSS.  $\boxed{\text{Solid-boxed}}$ ,  $\boxed{\text{dashed-boxed}}$  code is included in the correspondingly marked algorithms.

prove CSS adaptively secure by guessing the subset of users which will become compromised and/or the challenge file owners. However, the loss from this naïve “complexity leveraging” technique would be exponential, because the reduction needs to guess one subset in the full power set of the set of users. Developing other proof techniques to show that CSS is adaptively secure with less loss is an interesting question for future work.

An interesting proof artifact in the confidentiality bound is our reductions to two different AEAD security notions of indistinguishability. We could equally have reduced twice to AE security, which implies both IND-CCA and IND $\$$ ; we find that the extra granularity adds clarity to the proof. One AEAD security property we do not require is key robustness, which protects password-based cryptography against partitioning oracle attacks allowing an adversary to test several passwords in one interaction. In our setting, key derivation employs a random oracle and becomes interactive, so partitioning oracle attacks do not decrease adversarial interaction costs.

In the following, we let  $\mathcal{Q}_{\text{ORC}}(\mathcal{A})$  denote the number of queries to oracle ORC by adversary  $\mathcal{A}$ . Similarly, we let  $q_{\text{ORC}}(\mathcal{A})$  be the maximum number of queries *per user* made to oracle ORC by adversary  $\mathcal{A}$ . We use the convention that  $\mathcal{Q}_{\text{PUT1}}(\mathcal{A})$  and  $\mathcal{Q}_{\text{UPD1}}(\mathcal{A})$  include indirect queries made via oracle CHALL. To count only queries via oracle CHALL, we let CPUT1 and CUPD1 denote oracle CHALL with input ORC set to PUT1 and UPD1, respectively.

**Theorem 6.1 (C2CConfS of CSS)** *Let CSS be the cloud storage scheme defined in Figures 8–10, with building blocks specified in Section 6.1 and hash functions  $H_1, H_2$  modeled as random oracles. Let  $\mathcal{A}$  be an adversary against the C2CConfS security of CSS. Then we can construct adversaries*

$\mathcal{B}_{\text{pg}}, \mathcal{B}_{\text{prf}}, \mathcal{B}_{\text{aead}}^{12}$  and  $\mathcal{B}_{\text{aead}}^3$  such that

$$\begin{aligned} \mathbf{Adv}_{\text{CSS},n,\mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A}) &\leq 2 \cdot \left( \mathbf{Adv}_{n,\mathcal{PW}_n}^{\text{PG}}(\mathcal{B}_{\text{pg}}) + \mathbf{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B}_{\text{prf}}) + 2\mathbf{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^{12}) \right. \\ &\quad \left. + \frac{(\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + 1)^2 + (q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))^2}{2^{nl+1}} \right) + \mathbf{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^3). \end{aligned}$$

The query counts of adversaries the adversaries  $\mathcal{B}_{\text{pg}}, \mathcal{B}_{\text{prf}}, \mathcal{B}_{\text{aead}}^{12}$ , and  $\mathcal{B}_{\text{aead}}^3$  are bounded by the total query counts of adversary  $\mathcal{A}$ , as detailed in Appendix C.1. The running times of all adversaries are approximately that of  $\mathcal{A}$ .

**Proof sketch.** We give an overview over the main proof steps here; the full proof is in Appendix C.1. The proof proceeds through a sequence of games  $G_0$ – $G_7$ , which step by step replaces the keys in the key hierarchy of honest users ( $i_u \in (\mathcal{S}_{\text{hon}} \setminus \mathcal{S}_{\text{comp}})$ ) by keys generated independently at random, reducing to the relevant security properties of the respective parts of CSS. The advantage of  $\mathcal{A}$  in the final game then reduces to the left-or-right IND-CCA security of the AEAD scheme used for file encryption.

**Games  $G_0$ – $G_3$ : Replacing  $rw$  by random.** We begin with game  $G_0$ , which is equivalent to  $\mathbf{G}_{\text{CSS},n,\mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A})$ , except that whenever a protocol uses the hash functions  $H_1$  or  $H_2$ , the game instead queries the random oracles  $\text{RO}_1$  or  $\text{RO}_2$ , respectively. The effect of this change is only to model the hash functions as random oracles. Hence

$$\mathbf{Adv}_{\text{CSS},n,\mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A}) := 2 \cdot \Pr[G_0] - 1. \quad (1)$$

In the first sequence of games, the goal is to move to game  $G_3$ , where the client-side output  $rw \in \{0,1\}^{kl}$  from the OPRF flow in protocols  $\text{CSS}_{\text{areg}}$  and  $\text{CSS}_{\text{auth}}$  is replaced by an independent random string in  $\{0,1\}^{kl}$  for all honest (uncompromised) users.

We claim that the change from  $G_0$  to  $G_3$  is indistinguishable to an adversary, except with some small probability. Formally, we want to bound  $\Pr[G_0] - \Pr[G_3]$ . We do this through a sequence of game hops over two additional games,  $G_1$  and  $G_2$ , which allow us to show that  $\Pr[G_0] - \Pr[G_3] \leq \Pr[\text{Bad}]$ , where (informally)  $\text{Bad}$  is the event that adversary  $\mathcal{A}$  made a query  $\text{RO}_1(\text{aid}, pw)$  or  $\text{RO}_2(\text{aid}, pw, *)$  such that  $\text{aid}$  and  $pw$  are the account ID and password of an uncompromised user. We then construct an adversary  $\mathcal{B}_{\text{pg}}$  such that  $\Pr[\text{Bad}] \leq \mathbf{Adv}_{n,\mathcal{PW}_n}^{\text{PG}}(\mathcal{A})$ . Combined with standard equation rewriting, this gives

$$\Pr[G_0] \leq \mathbf{Adv}_{n,\mathcal{PW}_n}^{\text{PG}}(\mathcal{A}) + \Pr[G_3]. \quad (2)$$

**Games  $G_4$ – $G_7$ : Random keys and nonce collisions.** In the following hops, we apply PRF and AEAD security to ensure that keys in the key hierarchies of honest users are uniformly random and independent from anything the server sees. We also handle nonce collisions in AEAD encryptions.

First, in the hop to  $G_4$ , we replace  $k_{\text{kek}}$  and  $k_{\text{mac}}$  derived from  $rw$  with PRF  $\text{F}$  by random strings. This hop is bounded by

$$\Pr[G_3] - \Pr[G_4] \leq \mathbf{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B}_{\text{prf}}). \quad (3)$$

$G_5$  then replaces the encrypted master keys  $c_{mk}$  of honest clients by independent random strings, simultaneously ensuring that any deviation from the protocol by the server in authentication leads to a decryption failure when the client attempts to retrieve its master key. This hop is bounded by

$$\Pr[G_4] - \Pr[G_5] \leq \mathbf{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^1). \quad (4)$$

In  $G_6$  we handle nonce collisions in file and file key encryptions. We introduce a bad event if any two nonces for the same key collide, and in  $G_6$  sample a guaranteed fresh nonce as replacement. The bad event, and hence hop from  $G_5$  to  $G_6$  is bounded (through two birthday bounds) by

$$\Pr[G_5] - \Pr[G_6] \leq \frac{(\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + 1)^2 + (q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))^2}{2^{nl+1}}. \quad (5)$$

In the final game hop, to  $G_7$ , we replace the encrypted file keys  $c_{fk}$  of all honest users by random strings, again ensuring that only honest keys can be decrypted in get, update and share operations. This step is bounded by

$$\Pr[G_6] - \Pr[G_7] \leq \text{Adv}_{\text{AEAD}}^{\text{IND}\$}(\mathcal{B}_{\text{aead}}^2). \quad (6)$$

**Bounding  $G_7$ : IND-CCA security of file encryption.** In the last step of the proof, we bound the success probability of  $\mathcal{A}$  in game  $G_7$  by the advantage of an adversary  $\mathcal{B}_{\text{aead}}^3$  against the IND-CCA security of AEAD. This step requires careful checking of the trivial attack conditions, to ensure that  $\mathcal{B}_{\text{aead}}^3$  can properly simulate the game unless  $\mathcal{A}$  performs a trivial attack. We conclude that

$$2 \cdot \Pr[G_7] - 1 \leq \text{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^3). \quad (7)$$

Merging adversaries  $\mathcal{B}_{\text{aead}}^1$  and  $\mathcal{B}_{\text{aead}}^2$  into an adversary  $\mathcal{B}_{\text{aead}}^{12}$  and putting Equations (1)–(7) together yields the theorem statement.  $\square$

**Theorem 6.2 (C2CIntS of CSS)** *Let CSS be the cloud storage scheme defined in Figures 8–10, with building blocks specified in Section 6.1 and hash functions  $H_1, H_2$  modeled as random oracles. Let  $\mathcal{A}$  be an adversary against the C2CIntS security of CSS making at most  $p$  queries to oracle PUT1. We construct adversaries  $\mathcal{B}_{\text{pg}}, \mathcal{B}_{\text{prf}}, \mathcal{B}_{\text{aead}}^1, \mathcal{B}_{\text{aead}}^2, \mathcal{B}_{\text{aead}}^3$  and  $\mathcal{B}_{\text{aead}}^4$  such that*

$$\begin{aligned} \text{Adv}_{\text{CSS}, n, \mathcal{PW}_n}^{\text{C2CIntS}}(\mathcal{A}) &\leq \text{Adv}_{n, \mathcal{PW}_n}^{\text{PG}}(\mathcal{B}_{\text{pg}}) + \text{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B}_{\text{prf}}) + \text{Adv}_{\text{AEAD}}^{\text{IND}\$}(\mathcal{B}_{\text{aead}}^1) \\ &\quad + \frac{(\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + 1)^2 + (q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))^2}{2^{nl+1}} + \text{Adv}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{B}_{\text{aead}}^2) \\ &\quad + p \cdot \left( \text{Adv}_{\text{AEAD}}^{\text{IND}\$}(\mathcal{B}_{\text{aead}}^3) + \text{Adv}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{B}_{\text{aead}}^4) \right). \end{aligned}$$

*The query counts for adversaries  $\mathcal{B}_{\text{pg}}, \mathcal{B}_{\text{prf}}$  and  $\mathcal{B}_{\text{aead}}^1$  are as in the proof of Theorem 6.1, the others are described in Appendix C.2. The running times of all adversaries are approximately that of  $\mathcal{A}$ .*

The proof is analogous to that of Theorem 6.1, except for the last few steps. That is, the proof proceeds through a sequence of games  $G_0$ – $G_8$ , which step by step replaces the keys in the key hierarchy of honest users by keys generated independently at random, reducing to password-guessing, PRF and AEAD advantages along the way. In contrast to in the confidentiality proof, these hops do not incur a factor 2 in the bound, due to the advantage definition for C2CInt (Definition 4.3). Additionally, there is an intermediate hop to guess the file ID for which  $\mathcal{A}$  performs its forgery attempts, incurring a loss by a factor  $p$ . The advantage of  $\mathcal{A}$  in the final game reduces to the INT-CTXT security of the AEAD scheme used for file encryption. The full details of the proof are in Appendix C.2.

## 6.5 Limitations and Extensions of CSS

CSS is meant as a proof-of-concept cloud storage scheme, illustrating our formalism and security model. For the scope of this work, the design focuses on core cloud storage operations and security

mechanisms. We had to analyse CSS in the weaker selective security game to avoid commitment issues, leaving open the challenging task of achieving adaptive security for a scheme relying on standard cryptographic components. Moreover, in a real-world implementation, one would have to consider further functional requirements and additional security mechanisms, e.g., to strengthen the system’s behavior against external adversaries under an honest cloud service. In the following, we discuss limitations of our design in this regards and possible extensions.

ADAPTIVE SECURITY. Plausibly, the flexible password-based encryption (FPBE) approach by Bellare and Shea described in [12] could be used to achieve adaptive security when applied to cloud storage protocols. FPBE’s flat key hierarchy omits the key-encrypting ciphertexts that would cause issues in an adaptive analysis of CSS. However, there are concerns about the practicality of FPBE: Bellare and Shea omit any consideration of file-sharing. Without decoupling file keys from the password through key wrapping, it is unclear how to efficiently share files if all of a user’s keys are derived from their password (as in FPBE). Without naïve file replication, users sharing a file would need to derive the same key but from different passwords. Moreover, caching the password in clients (as it is needed for every operation) may be an operational security concern. Additionally, directly deriving file keys from user passwords makes password rotation very expensive: if a user wants to change their password, they have to re-encrypt all of their files.

However, FPBE is only one avenue towards achieving adaptive security. It might well be the case that CSS is adaptively secure; the challenge lies in constructing a proof which does not suffer an exponential loss (from naïve complexity leveraging) when addressing the commitment issues. We leave it as an open problem to prove stronger security of CSS, or to construct an adaptively-secure cloud storage scheme which avoids the commitment issues while still supporting all core functionality and not compromising on practicality.

ENFORCING UNIQUE IDENTIFIERS. In our security model, we assume unique account identifiers for honest users (e.g., email accounts) and file identifiers (e.g., by prefixing account IDs). This allows us to easily define what constitutes a “user” and a “file”: they correspond to an account ID and a file ID, respectively. As a consequence, we also rule out trivial collision attacks by a malicious cloud provider. Our CSS protocol hence, for reduced complexity in presentation, does not check for colliding identifiers (e.g., in  $\text{CSS}_{\text{areg}}$ ); in an actual implementation, an honest server would reasonably check for those to prevent malicious clients from overwriting data of other (honest) clients. Providing provably secure schemes under weaker assumptions on the uniqueness of account or file identifiers is an interesting direction for future work.

PASSWORD-AUTHENTICATED CHANNEL ESTABLISHMENT. Instead of relying on a secure channel protocol (like TLS), a protocol might expand the OPRF usage to a full password-based key exchange, e.g., employing the OPAQUE [37] protocol, possibly in combination with TLS [31].

## 7 Conclusion and Future Work

The deployment of E2EE systems began over a decade ago and brought major security improvements to areas like secure messaging (e.g., the Signal protocol) and browsing (TLS). Despite cloud storage seeing comparable adoption, these advancements have not translated to this area, as recent attacks on the largest E2EE cloud providers [6, 30, 3, 2] have shown. The vulnerabilities in the end-to-end security of these cloud storages are at least in part due to the complexity of the setting – which, when taking into consideration features like file sharing, is at least as intricate as key exchange – leaving it easy for providers to introduce subtle mistakes in protocols. These mistakes motivate the need for a strong cryptographic foundation for E2EE cloud storage.

In this paper, we initiated the formal study of E2EE cloud storage. As central contributions towards this, we introduced a syntax capturing the core functionality of cloud storage, as well as game-based security notions for the desired confidentiality and integrity against a malicious service provider. We put our model to test by both capturing recent attacks on MEGA as well as proving secure our own E2EE cloud storage scheme CSS.

In order to tame the complexity of the problem, we had to make some simplifying assumptions, leaving many opportunities for future work, including both model and protocol extensions. (See Sections 4.4 and 6.5, respectively.)

A natural next step is to consider Client-to-Server (C2S) security of CSS, i.e., security against external adversaries. We have sketched how to adapt our existing models to this C2S setting, and we expect the C2S security of CSS to follow smoothly. Working with an honest server and assuming a secure channel between clients and server allows C2S to avoid many complexities (e.g. malicious choices of the OPRF keys, offline password brute-forcing, etc).

Second, we had to analyse CSS in the weaker selective security setting to avoid commitment issues, leaving open the challenging task of achieving adaptive security for a scheme relying on standard cryptographic components. Plausibly, the flexible password-based encryption (FPBE) approach described in [12] could be used to achieve this, because FPBE’s flat key hierarchy omits the key-encrypting ciphertexts that would cause issues in an adaptive analysis of CSS. However, FPBE omits any consideration of file-sharing, and it not clear how to efficiently achieve this if all of a user’s keys are derived from their password (as in FPBE), other than by naive file replication.

Third, the core functionality considered in this paper could be extended – at the cost of higher complexity – with extra features such as file “unsharing” and deletion, password recovery, session deauthentication (and ensuing security notions for adversaries with access to expired sessions), and account deregistration.

Finally, many more advanced security properties could be considered. For example, it is an open question how post-compromise and forward security apply to the cloud storage setting, and how much is achievable; key rotation has been considered in prior works [36], as has fine-grained forward security via puncturing techniques [5, 53], but not in combination with the full feature set we target. Furthermore, CSS does not attempt to protect against roll-back attacks, nor provide advanced metadata-hiding, cf. [22, 21].

This paper introduces formally analyzed protocols instead of relying on *ad hoc* and/or proprietary designs. We hope that users will eventually enjoy the benefits of a single, well-analyzed scheme with provable security guarantees, akin to E2EE communication today. Standardization of such a scheme would resolve the long-standing issues that untrusted cloud providers need to be trusted with the design of a secure system and serving benign client code. A standardized scheme could have independent client implementations (i.e., not under the control of the potentially malicious provider) and even lead to interoperability between cloud storage providers. Our work is a first step towards bringing modern cryptographic guarantees to cloud storage.

## Acknowledgments

We thank Mihir Bellare for his ideas and contributions in the early phases of the project, in particular on formal syntax and definitions. We also thank Dennis Hofheinz for valuable discussions, as well as all anonymous reviewers for their comments and feedback.

The research of Miro Haller was supported by an Amazon Research Award Fall 2023. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of Amazon.

## References

- [1] pCloud International AG. pCloud – the most secure cloud storage. <https://www.pcloud.com/>. (Cited on page 3.)
- [2] Martin R. Albrecht, Matilda Backendal, Daniele Coppola, and Kenneth G. Paterson. Share with care: Breaking E2EE in Nextcloud. Cryptology ePrint Archive, Paper 2024/546, 2024. <https://eprint.iacr.org/2024/546>. (Cited on page 3, 4, 36.)
- [3] Martin R. Albrecht, Miro Haller, Lenka Mareková, and Kenneth G. Paterson. Caveat implementor! Key recovery attacks on MEGA. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 190–218. Springer, Heidelberg, April 2023. (Cited on page 3, 4, 25, 26, 27, 36, 42, 47, 48, 49, 50.)
- [4] Apple. Advanced data protection for iCloud. <https://support.apple.com/guide/security/advanced-data-protection-for-icloud-sec973254c5f/web>, Dec 2022. visited Feb 10, 2024. (Cited on page 3, 6.)
- [5] Matilda Backendal, Felix Günther, and Kenneth G. Paterson. Puncturable key wrapping and its applications. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part II*, volume 13792 of *LNCS*, pages 651–681. Springer, Heidelberg, December 2022. (Cited on page 7, 25, 37.)
- [6] Matilda Backendal, Miro Haller, and Kenneth G. Paterson. MEGA: Malleable encryption goes awry. In *2023 IEEE Symposium on Security and Privacy*, pages 146–163. IEEE Computer Society Press, May 2023. (Cited on page 3, 4, 25, 26, 27, 36, 42, 44, 46, 47, 48, 49, 50.)
- [7] Ali Bagherzandi, Stanislaw Jarecki, Nitesh Saxena, and Yanbin Lu. Password-protected secret sharing. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 2011*, pages 433–444. ACM Press, October 2011. (Cited on page 22.)
- [8] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th FOCS*, pages 514–523. IEEE Computer Society Press, October 1996. (Cited on page 50.)
- [9] Mihir Bellare, Thomas Ristenpart, and Stefano Tessaro. Multi-instance security and its application to password-based cryptography. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 312–329. Springer, Heidelberg, August 2012. (Cited on page 7, 15.)
- [10] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993. (Cited on page 8.)
- [11] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006. (Cited on page 8, 54, 58, 64.)
- [12] Mihir Bellare and Laura Shea. Flexible password-based encryption: Securing cloud storage and provably resisting partitioning-oracle attacks. In Mike Rosulek, editor, *CT-RSA 2023*,



- volume 13871 of *LNCS*, pages 594–621. Springer, Heidelberg, April 2023. (Cited on page [7](#), [15](#), [29](#), [36](#), [37](#).)
- [13] Alysson Bessani, Miguel Correia, Bruno Quaresma, Fernando André, and Paulo Sousa. Depsky: Dependable and secure storage in a cloud-of-clouds. *ACM Trans. Storage*, 9(4), nov 2013. (Cited on page [7](#).)
  - [14] Jan Camenisch, Anna Lysyanskaya, and Gregory Neven. Practical yet universally composable two-server password-authenticated secret sharing. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 525–536. ACM Press, October 2012. (Cited on page [22](#).)
  - [15] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001. (Cited on page [5](#), [22](#).)
  - [16] Ran Canetti, Palak Jain, Marika Swanberg, and Mayank Varia. Universally composable end-to-end secure messaging. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 3–33. Springer, Heidelberg, August 2022. (Cited on page [22](#).)
  - [17] Ran Canetti and Hugo Krawczyk. Universally composable notions of key exchange and secure channels. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 337–351. Springer, Heidelberg, April / May 2002. (Cited on page [5](#), [22](#).)
  - [18] Ran Canetti, Eyal Kushilevitz, and Yehuda Lindell. On the limitations of universally composable two-party computation without set-up assumptions. *Journal of Cryptology*, 19(2):135–167, April 2006. (Cited on page [5](#).)
  - [19] Silvia Casacuberta, Julia Hesse, and Anja Lehmann. SoK: Oblivious pseudorandom functions. In *7th IEEE European Symposium on Security and Privacy, EuroSecP 2022*, pages 625–646. IEEE, 2022. (Cited on page [7](#).)
  - [20] Long Chen, Ya-Nan Li, Qiang Tang, and Moti Yung. End-to-same-end encryption: Modularly augmenting an app with an efficient, portable, and blind cloud storage. In Kevin R. B. Butler and Kurt Thomas, editors, *USENIX Security 2022*, pages 2353–2370. USENIX Association, August 2022. (Cited on page [7](#).)
  - [21] Weikeng Chen, Thang Hoang, Jorge Guajardo, and Attila A. Yavuz. Titanium: A metadata-hiding file-sharing system with malicious security. In *29th Annual Network and Distributed System Security Symposium, NDSS 2022, San Diego, California, USA, April 24-28, 2022*. The Internet Society, 2022. (Cited on page [7](#), [25](#), [37](#).)
  - [22] Weikeng Chen and Raluca Ada Popa. Metal: A metadata-hiding file-sharing system. In *NDSS 2020*. The Internet Society, February 2020. (Cited on page [7](#), [25](#), [37](#).)
  - [23] Poulami Das, Julia Hesse, and Anja Lehmann. DPaSE: Distributed password-authenticated symmetric-key encryption, or how to get many keys from one password. In Yuji Suga, Kouichi Sakurai, Xuhua Ding, and Kazue Sako, editors, *ASIACCS 22*, pages 682–696. ACM Press, May / June 2022. (Cited on page [7](#).)
  - [24] Gareth T. Davies, Sebastian H. Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the WhatsApp end-to-end encrypted backup protocol.

- In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part IV*, volume 14084 of *LNCS*, pages 330–361. Springer, Heidelberg, August 2023. (Cited on page 6.)
- [25] Adam Everspaugh, Rahul Chatterjee, Samuel Scott, Ari Juels, and Thomas Ristenpart. The pythia PRF service. In Jaeyeon Jung and Thorsten Holz, editors, *USENIX Security 2015*, pages 547–562. USENIX Association, August 2015. (Cited on page 7.)
  - [26] Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 303–324. Springer, Heidelberg, February 2005. (Cited on page 27.)
  - [27] Eu-Jin Goh, Hovav Shacham, Nagendra Modadugu, and Dan Boneh. SiRiUS: Securing remote untrusted storage. In *NDSS 2003*. The Internet Society, February 2003. (Cited on page 7.)
  - [28] Miro Haller. Cloud storage systems: From bad practice to practical attacks. Master’s thesis, ETH Zurich, Jun 2022. (Cited on page 42.)
  - [29] Sven Hebrok, Simon Nachtigall, Marcel Maehren, Nurullah Erinola, Robert Merget, Juraj Somorovsky, and Jörg Schwenk. We really need to talk about session tickets: A Large-Scale analysis of cryptographic dangers with TLS session tickets. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 4877–4894, Anaheim, CA, August 2023. USENIX Association. (Cited on page 25.)
  - [30] Nadia Heninger and Keegan Ryan. The hidden number problem with small unknown multipliers: Cryptanalyzing MEGA in six queries and other applications. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part I*, volume 13940 of *LNCS*, pages 147–176. Springer, Heidelberg, May 2023. (Cited on page 3, 4, 25, 26, 27, 36, 47, 49.)
  - [31] Julia Hesse, Stanislaw Jarecki, Hugo Krawczyk, and Christopher Wood. Password-authenticated TLS via OPAQUE and post-handshake authentication. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part V*, volume 14008 of *LNCS*, pages 98–127. Springer, Heidelberg, April 2023. (Cited on page 36.)
  - [32] Stanislaw Jarecki, Aggelos Kiayias, and Hugo Krawczyk. Round-optimal password-protected secret sharing and T-PAKE in the password-only model. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 233–253. Springer, Heidelberg, December 2014. (Cited on page 4, 7, 27, 29.)
  - [33] Stanislaw Jarecki, Aggelos Kiayias, and Hugo Krawczyk. Round-optimal password-protected secret sharing and t-pake in the password-only model. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014*, pages 233–253, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg. (Cited on page 22.)
  - [34] Stanislaw Jarecki, Aggelos Kiayias, Hugo Krawczyk, and Jiayu Xu. Highly-efficient and composable password-protected secret sharing (or: How to protect your bitcoin wallet online). In *IEEE European Symposium on Security and Privacy, EuroS&P 2016*, pages 276–291. IEEE, 2016. (Cited on page 4, 27, 29.)
  - [35] Stanislaw Jarecki, Aggelos Kiayias, Hugo Krawczyk, and Jiayu Xu. TOPPSS: Cost-minimal password-protected secret sharing based on threshold OPRF. In Dieter Gollmann, Atsuko Miyaji, and Hiroaki Kikuchi, editors, *ACNS 17*, volume 10355 of *LNCS*, pages 39–58. Springer, Heidelberg, July 2017. (Cited on page 22.)



- [36] Stanislaw Jarecki, Hugo Krawczyk, and Jason K. Resch. Updatable oblivious key management for storage systems. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 379–393. ACM Press, November 2019. (Cited on page 7, 25, 37.)
- [37] Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 456–486. Springer, Heidelberg, April / May 2018. (Cited on page 7, 36.)
- [38] Russell W. F. Lai, Christoph Egger, Manuel Reinert, Sherman S. M. Chow, Matteo Maffei, and Dominique Schröder. Simple password-hardened encryption services. In William Enck and Adrienne Porter Felt, editors, *27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018*, pages 1405–1421. USENIX Association, 2018. (Cited on page 7.)
- [39] Russell W. F. Lai, Christoph Egger, Dominique Schröder, and Sherman S. M. Chow. Phoenix: Rebirth of a cryptographic password-hardening service. In Engin Kirda and Thomas Ristenpart, editors, *USENIX Security 2017*, pages 899–916. USENIX Association, August 2017. (Cited on page 7.)
- [40] MEGA Limited. MEGA – about us. <https://mega.io/about>. visited Jan 25, 2024. (Cited on page 3, 25.)
- [41] MEGA Limited. MEGA security white paper – third edition. <https://mega.nz/SecurityWhitepaper.pdf>, Jun 2022. (Cited on page 42.)
- [42] ID Cloud Services LTD. Icedrive – next-generation cloud storage. <https://icedrive.net/>. (Cited on page 3.)
- [43] Seafile Ltd. Seafile – open source file sync and share software. <https://www.seafile.com/en/home/>. (Cited on page 3.)
- [44] Team MEGA. MEGA security update. <https://blog.mega.io/mega-security-update>, Jun 2022. visited Feb 9, 2024. (Cited on page 25.)
- [45] Mega Limited. MEGA web client. <https://github.com/meganz/webclient>. visited Feb 6, 2024. (Cited on page 42.)
- [46] Kathleen Moriarty, Burt Kaliski, and Andreas Rusch. PKCS #5: Password-Based Cryptography Specification Version 2.1. RFC 8018, January 2017. (Cited on page 42.)
- [47] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudo-random functions. In *38th FOCS*, pages 458–467. IEEE Computer Society Press, October 1997. (Cited on page 27.)
- [48] Nextcloud. End-to-End Encryption Design. <https://nextcloud.com/c/uploads/2022/03/Nextcloud-end-to-end-encryption-Whitepaper.pdf>, September 2017. (Cited on page 3, 15.)
- [49] OnePassword. 1Password Security Design (Release v0.4.6). <https://1passwordstatic.com/files/security/1password-white-paper.pdf>, October 2023. (Cited on page 15.)

- [50] Raluca Ada Popa, Emily Stark, Steven Valdez, Jonas Helfer, Nickolai Zeldovich, and Hari Balakrishnan. Building web applications on top of encrypted data using mylar. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, pages 157–172, Seattle, WA, April 2014. USENIX Association. (Cited on page 7.)
- [51] Phillip Rogaway. Authenticated-encryption with associated-data. In Vijayalakshmi Atluri, editor, *ACM CCS 2002*, pages 98–107. ACM Press, November 2002. (Cited on page 50.)
- [52] Tresorit. End-to-end encrypted cloud storage for businesses. <https://tresorit.com/>. (Cited on page 3.)
- [53] Nirvan Tyagi, Muhammad Haris Mughees, Thomas Ristenpart, and Ian Miers. BurnBox: Self-revokable encryption in a world of compelled access. In William Enck and Adrienne Porter Felt, editors, *USENIX Security 2018*, pages 445–461. USENIX Association, August 2018. (Cited on page 7, 25, 37.)
- [54] WhatsApp. End-to-end encrypted backups on WhatsApp. <https://blog.whatsapp.com/end-to-end-encrypted-backups-on-whatsapp>. visited Feb 10, 2024. (Cited on page 6.)

## A Details of Formalizing (the Attacks on) MEGA

This section formalizes a simplified version of the cryptographic protocol implemented by MEGA in our syntax. We use MEGA’s whitepaper [41], their client source code [45], and previous analyses of their system [6, 28] as sources to aggregate a description of their protocol. The server code is closed-source and the official documentation only describes the server API. Thus, the server-side in our protocol is inferred from the cryptographic operations on the client and may differ from the actual implementation. For simplicity, we focus on cryptographic operations and omit various input validation and error messages. While this is sufficient to capture the attacks from [6] on MEGA, we remark that modeling the error messages *would be essential* for a full security proof. For instance, Albrecht et al. [3] exploited error messages in a now-patched client implementation of MEGA to build a plaintext recovery attack for a malicious server.

### A.1 Formalizing MEGA as a Cloud Storage Scheme

We define the MEGA cloud storage scheme  $M$  as a tuple of seven protocols  $M = (M_{\text{areg}}, M_{\text{auth}}, M_{\text{put}}, M_{\text{upd}}, M_{\text{get}}, M_{\text{share}}, M_{\text{accept}})$  to register accounts, authenticate clients, as well as upload, update, download, share, and accept files.

The following simplified descriptions of MEGA build on the symmetric encryption schemes AES-ECB, AES-CBC, and AES-CCM, as well as the asymmetric RSA encryption scheme. MEGA further uses the password-based key derivation function PBKDF [46] and a collision-resistant hash function  $H$ .

Our modeling focuses on the version of MEGA that was vulnerable to the attacks by Backendal et al. [6]. To this end, we omit functionality such as ephemeral accounts, two-factor authentication, account recovery, remote session deauthentication, and public sharing links (see MEGA whitepaper Sections 3.3 to 3.6 and 5 [41]). We describe the core protocols of MEGA in the remainder of this section, before discussing attacks.

ACCOUNT REGISTRATION  $M_{\text{areg}}$  (FIGURE 11). The client first picks a random master key  $k_m$  and client nonce  $n_C$ . It derives the encryption key  $k_e$  and authentication key  $k_a$ , from PBKDF applied to the concatenation of the user password  $pw$  and a salt  $s$ . The salt is the hash of the nonce  $n_C$

<p>client: <math>M_{\text{areg}}^{(C:1)}(\varepsilon, \varepsilon, in_C, \varepsilon)</math>:</p> <ol style="list-style-type: none"> <li>1 <b>pre:</b> <math>in_C.\{aid, pw, tk\}</math></li> <li>2 <math>k_m \leftarrow \{0, 1\}^{128}</math> ; <math>n_C \leftarrow \{0, 1\}^{128}</math></li> <li>3 <math>s \leftarrow H(const \parallel n_C)</math></li> <li>4 <math>k_e \parallel k_a \leftarrow \text{PBKDF}(s, pw)</math></li> <li>5 <math>h_a \leftarrow H(k_a)</math></li> <li>6 <math>c_m \leftarrow \text{AES-ECB.ENC}(k_e, k_m)</math></li> <li>7 <math>(pk, sk) \leftarrow \text{RSA.GEN}(2048)</math></li> <li>8 <math>c_{rsa} \leftarrow \text{AES-ECB.ENC}(k_m, sk)</math></li> <li>9 <math>m_C \leftarrow (aid, tk, n_C, c_m, h_a, pk, c_{rsa})</math></li> <li>10 <b>return</b> <math>(\varepsilon, \varepsilon, \varepsilon, m_C)</math></li> </ol> <p><math>M_{\text{areg}}^{(C:2)}(\varepsilon, st_C^{\text{tmp}}, m_S)</math>:</p> <ol style="list-style-type: none"> <li>11 <b>if</b> <math>m_S = \perp</math>: <b>fail</b><sub>C</sub> <b>else:</b> <b>success</b><sub>C</sub></li> </ol>	<p>server: <math>M_{\text{areg}}^{(S:1)}(st_S, \varepsilon, in_S, m_C)</math>:</p> <ol style="list-style-type: none"> <li>12 <b>pre:</b> <math>st_S.\{tokens, acc\}</math></li> <li>13 <math>(aid, tk, n_C, c_m, h_a, pk, c_{rsa}) \leftarrow m_C</math></li> <li>14 <b>if</b> <math>(aid, tk) \notin tokens</math>: <b>fail</b><sub>S</sub></li> <li>15 <math>acc[aid] \leftarrow (n_C, c_m, h_a, pk, c_{rsa})</math></li> <li>16 <b>success</b><sub>S</sub></li> </ol>
---	--

Figure 11: MEGA’s account registration protocol  $M_{\text{areg}}$ .

concatenated to a constant  $const$ . Furthermore, the client generates the RSA secret key  $sk$  and public key  $pk$ .<sup>9</sup> The client then sends  $n_C$ ,  $k_m$  encrypted under  $k_e$  using AES-ECB,  $pk$ ,  $c_{rsa}$ , and the hash of  $k_a$  to the server.

We do not model the preliminary account registration step where a user provides their email address  $aid$  and the server verifies that address by emailing a randomly-generated token  $tk$ .<sup>10</sup> The cryptographic account registration starts when the user already received  $tk$  – inputting it in  $in_C$  to the client – and returns to permanently store their encrypted key material. In  $M_{\text{areg}}^{(S:1)}$ , if the account ID  $aid$  and token  $tk$  sent by the client matches the ones stored in  $st_S.tokens$ , the server considers account registration to be successful and stores  $(n_C, c_m, h_a)$  in persistent state.

**AUTHENTICATION  $M_{\text{auth}}$  (FIGURE 12).** To authenticate, a client sends its email address  $aid$ , entered by the user in  $in_C$ , to the server. If  $aid$  is stored in the user table  $st_S.acc$ , the server answers with the correct salt. As a countermeasure against email enumeration, the server generates a random salt for non-existent users.

The client uses the salt to derive the encryption and authentication keys,  $k_e$  and  $k_a$ , respectively. It sends  $k_a$  in plaintext to the server. If the hash of  $k_a$  matches  $h_a$  (stored during registration), the server responds with the encrypted master key  $c_m$  and the (half-encrypted) RSA key pair  $(pk, c_{rsa})$ . Moreover, the server sends the randomized session ID  $in_S.sid$  (picked by the server in the bigger context of the protocol) encrypted for the RSA public key  $pk$  of the user that tries to authenticate.

The client recovers its key material ( $k_m$  and  $sk$ ) from the ciphertexts sent by the server and uses them to decrypt the session ID  $sid$ . It sends  $sid$  back to the server, which considers a correct value as proof that the client successfully recovered its key material. Additionally, the client stores  $k_m$ ,  $sk$ , and  $sid$  in the persistent client session state  $st_C$  to enable file operations.

**UPLOAD  $M_{\text{put}}$  (FIGURE 13).** For every file, the MEGA client picks a new file key  $k_f$  and nonce  $n_f$ . It then encrypts the file  $in_C.f$  with a custom authenticated encryption algorithm that we call MFILE

<sup>9</sup>We deviate slightly from MEGA’s protocol actual registration protocol here for the sake of simplicity. In the real system, the key material is only generated after the first successful login.

<sup>10</sup>Formalizing this would require a syntax extension to allow individual client-side inputs per protocol-step to capture the email out-of-band channel. (See “Protocol inputs” in Section 4.4.) Since the generated token in MEGA does not seem to depend on client input, we opt for a simplified representation taking the token as initial client input.

<p>client: <math>M_{\text{auth}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon)</math>:</p> <ol style="list-style-type: none"> <li>1 <b>pre:</b> <math>in_C.\{aid, pw\}</math></li> <li>2 <math>st_C^{\text{tmp}}.pw \leftarrow pw</math></li> <li>3 <b>return</b> <math>(st_C, st_C^{\text{tmp}}, \varepsilon, aid)</math></li> </ol> <p><math>M_{\text{auth}}^{(C:2)}(st_C, st_C^{\text{tmp}}, ms)</math>:</p> <ol style="list-style-type: none"> <li>4 <b>pre:</b> <math>st_C^{\text{tmp}}.\{pw\}</math></li> <li>5 <math>k_e \parallel k_a \leftarrow \text{PBKDF}(ms, pw)</math></li> <li>6 <math>st_C^{\text{tmp}}.k_e \leftarrow k_e</math></li> <li>7 <b>return</b> <math>(st_C, st_C^{\text{tmp}}, \varepsilon, k_a)</math></li> </ol> <p><math>M_{\text{auth}}^{(C:3)}(st_C, st_C^{\text{tmp}}, ms)</math>:</p> <ol style="list-style-type: none"> <li>8 <b>pre:</b> <math>st_C^{\text{tmp}}.\{k_e\}</math></li> <li>9 <math>(c_m, (pk, c_{rsa}), c_{sid}) \leftarrow ms</math></li> <li>10 <math>k_m \leftarrow \text{AES-ECB.DEC}(k_e, c_m)</math></li> <li>11 <math>sk \leftarrow \text{AES-ECB.DEC}(k_m, c_{rsa})</math></li> <li>12 <math>st_C.sid \leftarrow \text{RSA.DEC}(sk, c_{sid})</math></li> <li>13 <math>st_C.k_m \leftarrow k_m</math></li> <li>14 <math>st_C.pk \leftarrow pk</math> ; <math>st_C.sk \leftarrow sk</math></li> <li>15 <b>return</b> <math>(st_C, st_C^{\text{tmp}}, out_C, st_C.sid)</math></li> </ol> <p><math>M_{\text{auth}}^{(C:4)}(st_C, st_C^{\text{tmp}}, ms)</math>:</p> <ol style="list-style-type: none"> <li>16 <b>if</b> <math>ms = \perp</math>: <b>fail</b><sub>C</sub> <b>else:</b> <b>success</b><sub>C</sub></li> </ol>	<p>server: <math>M_{\text{auth}}^{(S:1)}(st_S, \varepsilon, in_S, mc)</math>:</p> <ol style="list-style-type: none"> <li>17 <b>pre:</b> <math>st_S.\{acc\}, in_S.\{sid\}</math></li> <li>18 <math>st_S^{\text{tmp}}.sid \leftarrow sid</math></li> <li>19 <b>if</b> <math>mc \in acc</math>:</li> <li>20 <math>st_S^{\text{tmp}}.aid \leftarrow mc</math></li> <li>21 <math>(n_C, c_m, h_a, pk, c_{rsa}) \leftarrow acc[mc]</math></li> <li>22 <math>s \leftarrow H(const \parallel n_C)</math></li> <li>23 <b>else:</b></li> <li>24 <math>s \leftarrow H(mc \parallel const' \parallel st_S.ns)</math></li> <li>25 <b>return</b> <math>(st_S, st_S^{\text{tmp}}, \varepsilon, s)</math></li> </ol> <p><math>M_{\text{auth}}^{(S:2)}(st_S, st_S^{\text{tmp}}, mc)</math>:</p> <ol style="list-style-type: none"> <li>26 <b>pre:</b> <math>st_S.\{acc\}, st_S^{\text{tmp}}.\{aid, sid\}</math></li> <li>27 <math>(n_C, c_m, h_a, pk, c_{rsa}) \leftarrow acc[aid]</math></li> <li>28 <b>if</b> <math>H(mc) \neq h_a</math>: <b>fail</b><sub>S</sub></li> <li>29 <math>c_{sid} \leftarrow \text{RSA.ENC}(pk, sid)</math></li> <li>30 <math>ms \leftarrow (c_m, pk, c_{rsa}, c_{sid})</math></li> <li>31 <b>return</b> <math>(st_S, st_S^{\text{tmp}}, \varepsilon, ms)</math></li> </ol> <p><math>M_{\text{auth}}^{(S:3)}(st_S, st_S^{\text{tmp}}, mc)</math>:</p> <ol style="list-style-type: none"> <li>32 <b>pre:</b> <math>st_S.\{sess\}, st_S^{\text{tmp}}.\{sid\}</math></li> <li>33 <b>if</b> <math>mc \neq sid</math>: <b>fail</b><sub>S</sub></li> <li>34 <math>sess[sid] \leftarrow st_S^{\text{tmp}}.aid</math></li> <li>35 <b>success</b><sub>S</sub></li> </ol>
--	--

Figure 12: MEGA's authentication protocol  $M_{\text{auth}}$ .

which returns the file ciphertext  $c_f$  and authentication tag  $\tau_f$ .<sup>11</sup> The client encrypts the file key material  $d_{fkm}$  – consisting of file key, nonce, and tag – with AES-ECB using the master key  $k_m$ . The client metadata  $d_{md}$ , including file size and parent handle for the file's position in the file tree, is encrypted with AES-CBC using the file key  $k_f$ . Finally, the client sends the session ID  $sid$ , file ID  $fid$ , and ciphertexts for the file key material, metadata, and file ( $c_{fkm}$ ,  $c_{md}$ , and  $c_f$ , respectively) to the server.

For a valid session ID, the server persistently stores the file ID  $fid$ , encrypted key material  $c_{fkm}$ , and metadata  $c_{md}$  in the file tree of the user (identified by their email  $aid$ ).<sup>12</sup> Furthermore, the file ciphertext  $c_f$  is also stored in  $st_S$  under the user's files.

**UPDATE  $M_{\text{upd}}$**  (FIGURE 13). Updating files is almost identical to  $M_{\text{put}}$ , except that it reuses the file key  $k_f$  and nonce  $n_f$  picked during the first upload. These values come from a file tree  $st_C.tree$  that is cached in the client state and updated on  $M_{\text{get}}$  operations (described next).

**DOWNLOAD  $M_{\text{get}}$**  (FIGURE 14). The client first downloads the file tree  $tree$  containing the encrypted file key material  $c_{fkm}$  and metadata  $c_{md}$  for every file and stores it in the client state  $st_C$  (e.g., to be used by  $M_{\text{upd}}$ ).<sup>13</sup> Note that fetching  $tree$  requires knowledge of the session token  $sid$ ; the

<sup>11</sup>We abstract away the details of the file encryption as they do not matter for our purposes. Figure 2–4 of Backendal et al. [6] describes the chunkwise file encryption using a variant of AES-CCM performed by MFILE.

<sup>12</sup>The server data structures are proprietary and, thus, unknown to us. Our abstractions  $st_S.tree$  and  $st_S.files$  are merely convenient to express MEGA's protocols.

<sup>13</sup>In practice, the client regularly polls for updates of this file tree, independent of  $M_{\text{get}}$ . For the purposes of our security analysis, this is equivalent to only fetching updates for the file tree on a file download request.

<p>client: <math>\overline{M_{\text{put}}^{(C:1)}} \mid \overline{M_{\text{upd}}^{(C:1)}}(st_C, \varepsilon, in_C, \varepsilon):</math></p> <pre> 1 <b>pre:</b> <math>in_C.\{fid, f\}in_C.\{md\}</math> 2 <math>\boxed{k_f \leftarrow \\$_{0,1}^{128}; n_f \leftarrow \\$_{0,1}^{64}}</math> 3 <math>\overline{(\cdot, (k_f, n_f, \cdot)) \leftarrow st_C.tree[fid]}</math> 4 <math>(c_f, \tau_f) \leftarrow E^{\text{MFILE}}(k_f, n_f, f)</math> 5 <math>d_{fkm} \leftarrow (k_f, n_f, \tau_f)</math> 6 <math>c_{fkm} \leftarrow \text{AES-ECB.ENC}(st_C.k_{mk}, d_{fkm})</math> 7 <math>c_{md} \leftarrow \text{AES-CBC.ENC}(k_f, md)</math> 8 <math>m_C \leftarrow (st_C.sid, fid, c_{fkm}, c_{md}, c_f)</math> 9 <b>return</b> <math>(st_C, \varepsilon, \varepsilon, m_C)</math> </pre> <p><math>\overline{M_{\text{put}}^{(C:2)}} \mid \overline{M_{\text{upd}}^{(C:2)}}(st_C, st_C^{\text{tmp}}, m_S):</math></p> <pre> 10 <b>if</b> <math>m_S = \perp</math>: <b>fail</b><math>_C</math> <b>else:</b> <b>success</b><math>_C</math> </pre>	<p>server: <math>\overline{M_{\text{put}}^{(S:1)}} \mid \overline{M_{\text{upd}}^{(S:1)}}(st_S, \varepsilon, in_S, m_C):</math></p> <pre> 11 <b>pre:</b> <math>st_S.\{sess, tree, files\}</math> 12 <math>(sid, fid, c_{fkm}, c_{md}, c_f) \leftarrow m_C</math> 13 <math>aid \leftarrow sess[sid]</math> 14 <math>tree[aid] \leftarrow \bigcup \{(fid, c_{fkm}, c_{md})\}</math> 15 <math>files[fid] \leftarrow c_f</math> 16 <b>success</b><math>_S</math> </pre>
--	---

Figure 13: MEGA’s file upload and update protocols  $M_{\text{put}}$  and  $M_{\text{upd}}$ . The  $\boxed{\phantom{x}}$  line is only included in  $M_{\text{put}}$ , the  $\overline{\phantom{x}}$  line only in  $M_{\text{upd}}$ .

<p>client: <math>\overline{M_{\text{get}}^{(C:1)}}(st_C, \varepsilon, in_C, \varepsilon):</math></p> <pre> 1 <b>pre:</b> <math>in_C.\{fid\}st_C.\{sid\}</math> 2 <math>st_C^{\text{tmp}}.fid \leftarrow fid</math> 3 <b>return</b> <math>(st_C, st_C^{\text{tmp}}, \varepsilon, sid)</math> </pre> <p><math>\overline{M_{\text{get}}^{(C:2)}}(st_C, st_C^{\text{tmp}}, m_S):</math></p> <pre> 4 <b>pre:</b> <math>st_C.\{k_{mk}\}, st_C^{\text{tmp}}.\{fid\}</math> 5 <b>for</b> <math>(fid', c_{fkm}, c_{md}) \in m_S</math>: 6   <math>d_{fkm} \leftarrow \text{AES-ECB.DEC}(k_{mk}, c_{fkm})</math> 7   <math>(k_f, n_f, \tau_f) \leftarrow d_{fkm}</math> 8   <math>d_{md} \leftarrow \text{AES-CBC.DEC}(k_f, c_{md})</math> 9   <math>st_C.tree[fid'] \leftarrow (d_{md}, d_{fkm})</math> 10 <b>return</b> <math>(st_C, st_C^{\text{tmp}}, out_C, fid)</math> </pre> <p><math>\overline{M_{\text{get}}^{(C:3)}}(st_C, st_C^{\text{tmp}}, m_S):</math></p> <pre> 11 <b>pre:</b> <math>st_C.\{tree\}, st_C^{\text{tmp}}.\{fid\}</math> 12 <math>(d_{md}, d_{fkm}) \leftarrow tree[fid]</math> 13 <math>(f, out_C.dec) \leftarrow D^{\text{MFILE}}(d_{fkm}, m_S)</math> 14 <b>if</b> <math>out_C.dec</math>: <math>out_C.f \leftarrow f</math> 15 <b>success</b><math>_C</math> </pre>	<p>server: <math>\overline{M_{\text{get}}^{(S:1)}}(st_S, \varepsilon, in_S, m_C):</math></p> <pre> 16 <b>pre:</b> <math>st_S.\{tree, sess\}</math> 17 <b>return</b> <math>(st_S, \varepsilon, \varepsilon, tree[sess[m_C]])</math> </pre> <p><math>\overline{M_{\text{get}}^{(S:2)}}(st_S, st_S^{\text{tmp}}, m_C):</math></p> <pre> 18 <b>pre:</b> <math>st_S.\{files\}</math> 19 <b>success</b><math>_S(files[m_C])</math> </pre>
--	---

Figure 14: MEGA’s file download protocol  $M_{\text{get}}$ .

server fails by returning  $\perp$  for non-existent  $sid$ . The client decrypts the file key material with the master key stored in the client session state  $st_C$ . It then requests the file with identifier  $fid$ <sup>14</sup> from the server and decrypts it using the key material in the file tree and MEGA’s custom decryption algorithm  $D^{\text{MFILE}}$ , returning the decrypted file  $f$  and a success boolean (stored directly in client protocol success decision  $out_C.dec$ ).

<sup>14</sup>Clients find the file identifier  $fid$  corresponding to a file of a particular name in the decrypted file tree. However, our games require specifying  $fid$  in  $in_C$  in the first round, which can be interpreted as using the identifier from the cached file tree of a previous  $M_{\text{get}}$  call (say, for a dummy  $fid$ ).

<p>client: <math>M_{\text{share}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon)</math>:</p> <ol style="list-style-type: none"> <li>1 <b>pre</b>: <math>in_C.\{raid, fid\} in_C.\{rpk\}</math></li> <li>2 <math>(d_{md}, d_{fkm}) \leftarrow st_C.tree[fid]</math></li> <li>3 <math>c_{fkm} \leftarrow \text{RSA.ENC}(rpk, d_{fkm})</math></li> <li>4 <math>m_C \leftarrow (st_C.sid, raid, fid, c_{fkm})</math></li> <li>5 <b>return</b> <math>(st_C, \varepsilon, \varepsilon, m_C)</math></li> </ol> <p><math>M_{\text{share}}^{(C:2)}(st_C, st_C^{\text{tmp}}, m_S)</math>:</p> <ol style="list-style-type: none"> <li>6 <b>if</b> <math>m_S = \perp</math>: <b>fail</b><sub>C</sub> <b>else</b>: <b>success</b><sub>C</sub></li> </ol>	<p>server: <math>M_{\text{share}}^{(S:1)}(st_S, \varepsilon, in_S, m_C)</math>:</p> <ol style="list-style-type: none"> <li>7 <b>pre</b>: <math>st_S.\{sess, acc, tree, shrtree\}</math></li> <li>8 <math>(sid, raid, fid, c_{fkm}) \leftarrow m_C</math></li> <li>9 <b>if</b> <math>sid \notin sess</math> <b>or</b> <math>raid \notin acc</math>: <b>fail</b><sub>S</sub></li> <li>10 <b>if</b> <math>(fid, \cdot, \cdot) \notin tree[sess[sid]]</math>:</li> <li>11 <b>fail</b><sub>S</sub></li> <li>12 <math>shrtree[raid][fid] \leftarrow c_{fkm}</math></li> <li>13 <b>success</b><sub>S</sub></li> </ol>
--	--

Figure 15: Simplified version of MEGA’s legacy file sharing protocol  $M_{\text{share}}$ .

<p>client: <math>M_{\text{accept}}^{(C:1)}(st_C, \varepsilon, in_C, \varepsilon)</math>:</p> <ol style="list-style-type: none"> <li>1 <b>pre</b>: <math>in_C.\{fid\}, st_C.\{sid\}, in_C.\{d_{md}\}</math></li> <li>2 <math>st_C^{\text{tmp}}.d_{md} \leftarrow d_{md}</math></li> <li>3 <b>return</b> <math>(st_C, \varepsilon, \varepsilon, (sid, fid))</math></li> </ol> <p><math>M_{\text{accept}}^{(C:2)}(st_C, st_C^{\text{tmp}}, m_S)</math>:</p> <ol style="list-style-type: none"> <li>4 <b>pre</b>: <math>st_C.\{sk, k_{mk}\}, st_C^{\text{tmp}}.\{d_{md}\}</math></li> <li>5 <math>d_{fkm} \leftarrow \text{RSA.DEC}(sk, m_S)</math></li> <li>6 <math>c_{fkm} \leftarrow \text{AES-ECB.ENC}(k_{mk}, d_{fkm})</math></li> <li>7 <math>(k_f, n_f, \tau_f) \leftarrow d_{fkm}</math></li> <li>8 <math>c_{md} \leftarrow \text{AES-CBC.ENC}(k_f, d_{md})</math></li> <li>9 <b>return</b> <math>(st_C, st_C^{\text{tmp}}, out_C, (c_{fkm}, c_{md}))</math></li> </ol> <p><math>M_{\text{accept}}^{(C:3)}(st_C, st_C^{\text{tmp}}, m_S)</math>:</p> <ol style="list-style-type: none"> <li>10 <b>if</b> <math>m_S = \perp</math>: <b>fail</b><sub>C</sub> <b>else</b>: <b>success</b><sub>C</sub></li> </ol>	<p>server: <math>M_{\text{accept}}^{(S:1)}(st_S, \varepsilon, in_S, m_C)</math>:</p> <ol style="list-style-type: none"> <li>11 <b>pre</b>: <math>st_S.\{sess, acc, tree, shrtree\}</math></li> <li>12 <math>(sid, st_S^{\text{tmp}}.fid) \leftarrow m_C</math></li> <li>13 <math>st_S^{\text{tmp}}.aid \leftarrow sess[sid]</math></li> <li>14 <math>m_S \leftarrow shrtree[st_S^{\text{tmp}}.aid][st_S^{\text{tmp}}.fid]</math></li> <li>15 <b>return</b> <math>(st_S, st_S^{\text{tmp}}, \varepsilon, m_S)</math></li> </ol> <p><math>M_{\text{accept}}^{(S:2)}(st_S, st_S^{\text{tmp}}, m_C)</math>:</p> <ol style="list-style-type: none"> <li>16 <b>pre</b>: <math>st_S.\{tree\}, st_S^{\text{tmp}}.\{aid, fid\}</math></li> <li>17 <math>c_{fkm}, c_{md} \leftarrow m_C</math></li> <li>18 <math>tree[aid] \leftarrow (fid, c_{fkm}, c_{md})</math></li> <li>19 <b>success</b><sub>S</sub></li> </ol>
---	---

Figure 16: Simplified version of MEGA’s legacy protocol to accept shared files  $M_{\text{accept}}$ .

SHARING  $M_{\text{share}}$  (FIGURE 15). Backendal et al. [6] exploited a deprecated code path that was still present in web clients before their disclosure. To show how their attack shows up in our security games, we describe a simplified version of this legacy file sharing protocol.

In MEGA’s system, users establish contact relationships and sign the public key material of their contacts on first use (Trust On First Use). Instead of modeling this complex setup, we assume that the client obtained the public key  $rpk$  of receiver with account ID  $raid$  over some out-of-band channel and provides it as input to  $M_{\text{share}}$  in  $in_C$ . The client encrypts the file key material  $d_{fkm}$  with the RSA public key of the recipient and sends all encrypted file material to the server. If the session  $sid$  and receiver account  $raid$  exist, and the caller owns the file to be shared (line 11 in Figure 15), the server temporarily stores the encrypted file keys  $c_{fkm}$  in the share tree  $shrtree$ .

RECEIVING  $M_{\text{accept}}$  (FIGURE 16). The recipient of a file gets notified that a file with ID  $fid$  was shared with them by polling the server (we do not model this mechanism). The user then calls  $M_{\text{accept}}$  to fetch the encrypted file key material and re-encrypt it using their RSA secret key  $sk$  (for decryption) and master key  $k_{mk}$  (for encryption), likely to avoid expensive public key operations in the future. Next, the new encrypted file key material  $c_{fkm}$  and metadata  $c_{md}$  are uploaded to the server and stored in the server’s persistent file tree data structure  $st_S.tree$ .<sup>15</sup>

<sup>15</sup>In practice, the metadata contains user-specific values such as a pointer to the parent folder of the file in a user’s storage. These values would need to be replaced when sharing a file, which we model as a user input for simplicity.



```

ExecHonORC,  $\Pi$ ( $in_S, st_{adv}, \mathbf{arg}$ ):
1 pre:  $ORC \neq \text{GET1}$ 
2  $m' \leftarrow ORC(\mathbf{arg})$ 
3  $(st_{adv}, st_{adv}^{\text{tmp}}, out_{adv}, m) \leftarrow \Pi^{(S:1)}(st_{adv}, st_{adv}^{\text{tmp}}, in_S, m')$ 
4  $st_{adv}^{\text{tmp}} \leftarrow \varepsilon$ ;  $i \leftarrow 2$ 
5 while ( $out_{adv}.dec \neq \varepsilon$  or  $out_C.dec \neq \varepsilon$ ):
6    $m' \leftarrow \text{STEP}(\Pi, i_p, m)$ 
7    $(st_{adv}, st_{adv}^{\text{tmp}}, out_{adv}, m) \leftarrow \Pi^{(S:i)}(st_{adv}, st_{adv}^{\text{tmp}}, m')$ 
8    $i \leftarrow i + 1$ 
9 return  $st_{adv}, out_{adv}$ 

```

Figure 17: Honest execution by using oracles and following the protocol steps.

## A.2 Formal Attacks against MEGA

We proceed to show how attacks 1–3 from [6] (incorporating results from [3, 30]) against the previously-described MEGA protocol  $\mathbf{M}$  lead to adversaries with non-negligible winning probabilities in games  $\mathbf{G}_{\mathbf{M}, n, \mathcal{PW}_n}^{\text{C2CConfS}}$  and  $\mathbf{G}_{\mathbf{M}, n, \mathcal{PW}_n}^{\text{C2CIntS}}$  for an arbitrary number of accounts  $n \geq 1$ . Moreover, we discuss a generic attack that more fundamentally prevents a security proof, even of the patched MEGA protocol.

The following attacks perform a number of operations honestly, for which we use the shared subroutine depicted in Figure 17.

**C2CConfS ATTACK.** The adversary  $\mathcal{A}_{\text{conf}}$ , described in Figure 18, combines the RSA key and plaintext recovery attacks from Backendal et al. [6] to win  $\mathbf{G}_{\mathbf{M}, n, \mathcal{PW}_n}^{\text{C2CConfS}}$  with non-negligible probability.

The adversary starts by honestly registering a user with email address  $aid$ . It picks two distinct files ( $f_0$  and  $f_1$ ) of the same length and uploads them as challenge. Next,  $\mathcal{A}_{\text{conf}}$  runs the RSA key recovery attack from [6, Section III] to recover the prime factors  $p$  and  $q$  from the RSA modulus  $N$ . In short, the attack exploits the lack of integrity protection for the RSA secret key ciphertext  $c_{rsa}$  and properties of MEGA’s RSA decryption  $\text{RSA.DEC}$  on the client to perform a binary search for the smaller factor of the modulus  $N$ . The adversary tampers with the AES-ECB ciphertext  $c_{rsa}$  on line 18 of Figure 18 to partially garble the secret key. Due to MEGA’s implementation of  $\text{RSA.DEC}$ , clients leak whether the encrypted challenge value  $p'$  is less than the smaller RSA prime factor  $\min(p, q)$  (by responding with  $m_C = 0$ ) or greater (resulting in  $m_C > 0$  with high probability).

With knowledge of the RSA secret key, adversary  $\mathcal{A}_{\text{conf}}$  runs the plaintext recovery attack [6, Section IV] to decrypt the challenge file. On a high level, the adversary injects the encrypted file key material  $c_{fkm}$  for the challenge file in the RSA secret key ciphertext  $c_{rsa}$  that is exchanged during the authentication protocol. This is possible because both the RSA key and file keys are encrypted with the master key and AES-ECB, enabling the replacement of individual ciphertext blocks in  $c_{rsa}$  and  $c_{fk}$ . Together with a specially crafted session ID (the value  $u' \cdot q'$ , using the recovered RSA secret key), the client inadvertently decrypts the injected file key material ( $k_f, n_f, \tau_f$ ) in the session ID that it sends back to the server in  $\mathbf{M}_{\text{auth}}^{(C:3)}$ . The attack is probabilistic and fails with probability at most  $2^{-30}$  (as explained by [6]). As a successful file decryption always results in the adversary winning the game  $\mathbf{G}_{\mathbf{M}, n, \mathcal{PW}_n}^{\text{C2CConfS}}$ , it has advantage  $\text{Adv}_{\mathbf{M}, n, \mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A}_{\text{conf}}) = 1 - 2^{-30}$  with 1023 calls to  $\mathbf{M}_{\text{auth}}$ . Implementing the optimization of [6], the adversary  $\mathcal{A}_{\text{conf}}$  can stop the binary search after 512 queries to  $\mathbf{M}_{\text{auth}}$  and recover the missing part of the secret key with lattice-based techniques. Moreover, Heninger and Ryan [30] proposed a tailored lattice-based RSA key recovery attack that further reduces the cost to 6 calls to  $\mathbf{M}_{\text{auth}}$ .



```

adversary  $\mathcal{A}_{conf}^{\text{AREG1,AUTH1,PUT1,UPD1,GET1,SHARE1,ACPT1,CHALL,COMP}}()$ :
  // Attack setup: honest account creation and challenge file upload
  1  $aid \leftarrow \text{"evil@email.com"} ; in_C.aid \leftarrow aid ; i_u \leftarrow 0 ; i_c \leftarrow 0 ; fid \leftarrow 0$ 
  2  $in_C.tk \leftarrow \{0, 1\}^{128} ; st_{adv}.tokens \leftarrow \{(aid, in_C.tk)\}$ 
  3  $(st_{adv}, out_{adv}) \leftarrow \text{ExecHon}_{\text{AREG1}, M_{\text{areg}}}(\varepsilon, st_{adv}, (i_u, in_C))$ 
  4  $sid \leftarrow \$S ; in_S.sid \leftarrow \$sid$ 
  5  $(st_{adv}, out_{adv}) \leftarrow \text{ExecHon}_{\text{AUTH1}, M_{\text{auth}}}(in_S, st_{adv}, (i_u, in_C))$ 
  6  $in_C.fid \leftarrow fid$ 
  7  $f_0 \leftarrow 0 ; f_1 \leftarrow 1$ 
  8  $(st_{adv}, out_{adv}) \leftarrow \text{ExecHon}_{\text{CHALL}, M_{\text{put}}}(\varepsilon, st_{adv}, (i_c, in_C, f_0, f_1, \text{PUT1}))$ 
  // Read target user ciphertexts from malicious server state
  9  $(n_C, c_m, h_a, (pk, c_{rsa})) \leftarrow st_{adv}.acc[sid]$ 
  10 for  $(fid', c'_{fkm}, c'_{md}) \in st_{adv}.tree[aid]$ : // Find key material ciphertext  $c_{fkm}$  for the newly uploaded file
  11   if  $fid = fid'$ :  $c_{fkm} \leftarrow c'_{fkm}$ 
  12  $c_f \leftarrow st_{adv}.files[fid]$ 
  // Recover RSA private key (see [6, Section III])
  13  $l \leftarrow 0 ; r \leftarrow \lfloor \sqrt{N} \rfloor ; i_p \leftarrow 4$ 
  14 do:
  15    $p' \leftarrow \lfloor \frac{r+l}{2} \rfloor ; in_C.aid \leftarrow aid$ 
  16    $m_C \leftarrow \text{AUTH1}(i_u, in_C)$ 
  17    $m_C \leftarrow \text{STEP}(M_{\text{auth}}, i_p, H(const \parallel n_C))$ 
  18    $m_{adv} \leftarrow (c_m, (pk, c_{rsa} \oplus 1), \text{RSA.ENC}(pk, p'))$ 
  19    $m_C \leftarrow \text{STEP}(M_{\text{auth}}, i_p, m_{adv})$ 
  20   if  $m_C = 0$ :  $r \leftarrow p'$  else:  $l \leftarrow p' + 1$ 
  21    $i_p \leftarrow i_p + 1$ 
  22 until  $p' \mid N$ 
  // Decrypt the challenge file (see [6, Section IV])
  23  $q' \leftarrow N/p' ; u' \leftarrow (p')^{-1} \bmod q'$ 
  24  $in_C.aid \leftarrow aid$ 
  25  $m_C \leftarrow \text{AUTH1}(i_u, in_C)$ 
  26  $m_C \leftarrow \text{STEP}(M_{\text{auth}}, i_p, H(const \parallel n_C))$ 
  27  $c'_{rsa} \leftarrow \lfloor c_{rsa} \cdot 2^{-8 \cdot 128} \rfloor \parallel c_{fkm} \parallel (c_{rsa} \bmod 2^{6 \cdot 128})$ 
  28  $m_{adv} \leftarrow (c_m, (pk, c'_{rsa}), \text{RSA.ENC}(pk, u' \cdot q'))$ 
  29  $m_C \leftarrow \text{STEP}(M_{\text{auth}}, i_p, m_{adv})$ 
  30 for  $y = 0, 1, \dots, 2^{16} - 1$ :
  31    $(k_f, n_f, \tau_f) \leftarrow \lfloor \frac{y \parallel m_C}{q} \cdot 256^{123} \rfloor \bmod 256^{32}$ 
  32    $(f, succ) \leftarrow D^{\text{MFILE}}(k_f, n_f, \tau_f, c_f)$ 
  33   if  $succ$ : return  $\mathbb{I}_{f=f_1}$ 
  34 return 0 // (attack failed; reached with probability less than  $2^{-30}$ )

```

Figure 18: Adversary running RSA key attack (recovering prime factors  $p, q$  of the modulus  $N = p \cdot q$ ) and the plaintext recovery attack from [6].

**C2CIntS ATTACK.** Figure 19 describes the adversary  $\mathcal{A}_{int}$ , running the integrity attack from Backendal et al. [6] to win  $\mathbf{G}_{M,n,\mathcal{PW}_n}^{\text{C2CIntS}}$  with non-negligible probability. The integrity attack from [6] requires knowing a plaintext-ciphertext pair. Instead of building on  $\mathcal{A}_{conf}$  and decrypting a randomly picked AES-ECB ciphertext (to chain this attack with the previous two as in [6]), adversary  $\mathcal{A}_{int}$  uses the AES-ECB encryption oracle from [3, Section 2.2] to obtain the ciphertext for a cho-

sen plaintext from the share accept protocol, as follows. The adversary first honestly registers and authenticates a user, and then makes them accept a “shared” file chosen by the adversary. The client decrypts the shared file key material  $d_{fkm}$  from  $c_{fkm}$  (using RSA) and re-encrypts it under the user’s master key (using AES-ECB), sending the resulting  $c'_{fkm}$  to the server.

Since the adversary picked the file key material  $d_{fkm}$  (which consists of two AES blocks, the first containing the file key  $k_f$ , and the second a nonce  $n_f$  and tag  $\tau_f$ ), it now knows both the plaintext blocks and the corresponding ciphertext blocks  $c'_{fkm}$ . It will use the second one of these known blocks to mount the integrity attack of [6].

Note that local file key ciphertexts are encrypted with AES-ECB, and therefore not integrity-protected. Hence, as [6] describe, the adversary can trivially create a forged file key ciphertext. For this attack, it will use the second of the ciphertext blocks from above for which it knows the corresponding plaintext (namely, the nonce  $n_f$  and tag  $\tau_f$ ), and create a ciphertext which consists of that block duplicated twice. That is, the forged ciphertext consists of two identical AES-ECB blocks, for which the adversary knows the plaintext value  $n_f \parallel \tau_f$ .

Due to the way  $d_{fkm}$  is parsed in MEGA’s custom decryption algorithm  $D^{\text{MFILE}}$ , the duplicate ciphertext blocks cause the client to derive the all-zeros file encryption key  $k_f'' = 0^{128}$  when it decrypts the malicious file key ciphertext. Hence, the adversary knows which key the user derives (despite that this key was never used before), and can use it to fabricate a file and metadata ciphertext. The nonce  $n_f$  and tag  $\tau_f$ , which are also derived using  $D^{\text{MFILE}}$  on the file key ciphertext, are unaffected by this vulnerability. Hence, in addition to knowing the file key, that adversary also knows the nonce  $n_f$  and tag  $\tau_f$  that will be used during file decryption. This is needed for the attack, since encrypting the fabricated file must produce the tag  $\tau_f$  when using nonce  $n_f$  in order to pass verification during decryption.

Figure 6 of [6] explains how it is possible to create such a file by changing only 128 bits of an arbitrary file; we call this algorithm  $E^{\text{MFILE}'}$ . As long as the file encrypted by  $E^{\text{MFILE}'}$  is not the file that was shared in the first step of the attack, this leads to a win in  $\mathbf{G}_{M,n,\mathcal{PW}_n}^{\text{C2CIntS}}$ . Therefore, adversary  $\mathcal{A}_{int}$  has advantage  $\mathbf{Adv}_{M,n,\mathcal{PW}_n}^{\text{C2CConf}}(\mathcal{A}_{int}) = 1$  with only a single call to each of  $M_{\text{areg}}$ ,  $M_{\text{auth}}$ ,  $M_{\text{accept}}$ , and  $M_{\text{get}}$ .

REMAINING OBSTACLES IN THE PATCHED MEGA PROTOCOL. Despite having mitigated the attacks from [6, 3, 30], MEGA still cannot be proven secure in our model. Even with the patches, the encryption scheme remains vulnerable to theoretic attacks which prevent the protocol from satisfying our security notions. Unfortunately, mitigating these attacks is infeasible for MEGA, as it would require all of their users to re-encrypt all of their files.

One example of a remaining issues shows up in Figure 19. In their patches, MEGA decided to retain the file sharing feature that exposes an AES-ECB encryption oracle unchanged, instead of switching to more expensive asymmetric encryption. Hence, an adversary could still run the attack until line 12 of Figure 19 and obtain file key material  $d_{fkm}$  that is encrypted with the target user’s master key. Instead of running the integrity attack from [6] to achieve a full break in practice, the adversary could also win  $\mathbf{G}_{M,n,\mathcal{PW}_n}^{\text{C2CIntS}}$  with a more theoretical attack by reusing the file key material  $d_{fkm}$  to encrypt a new file. This is possible because file keys are not bound to the file or identifier  $fid$ , hence allowing an adversary that knows any file key material to reuse it to encrypt other files. The new file is not tracked as shared by the adversary, and therefore constitutes a valid forgery.

```

adversary  $\mathcal{A}_{int}^{\text{AREG1,AUTH1,PUT1,UPD1,GET1,SHARE1,ACPT1,COMP}}():$ 
  // Attack setup: honest account creation
  1  $aid \leftarrow \text{"evil@email.com"} ; in_C.aid \leftarrow aid ; i_u \leftarrow 0 ; i_c \leftarrow 0$ 
  2  $in_C.tk \leftarrow \{0, 1\}^{128} ; st_{adv}.tokens \leftarrow \{(aid, in_C.tk)\}$ 
  3  $(st_{adv}, out_{adv}) \leftarrow \text{ExecHon}_{\text{AREG1}, M_{\text{areg}}}(\varepsilon, st_{adv}, (i_u, in_C))$ 
  4  $sid \leftarrow \$S ; in_S.sid \leftarrow \$sid$ 
  5  $(st_{adv}, out_{adv}) \leftarrow \text{ExecHon}_{\text{AUTH1}, M_{\text{auth}}}(in_S, st_{adv}, (i_u, in_C))$ 
  // Prepare adversary chosen "shared file"
  6  $(n_C, c_m, h_a, (pk, c_{rsa})) \leftarrow st_{adv}.acc[sid]$ 
  7  $(k_f, n_f) \leftarrow \$(\{0, 1\}^{128} \times \{0, 1\}^{64})$ 
  8  $(c_f, \tau_f) \leftarrow E^{\text{MFILE}}(k_f, n_f, 0)$  // Encrypt any file, for instance "0"
  9  $d_{fkm} \leftarrow (k_f, n_f, \tau_f) ; c_{fkm} \leftarrow \text{RSA.ENC}(pk, d_{fkm})$ 
  // Obtain valid plaintext-ciphertext pair (over file sharing, see [3, Section 2.2])
  10  $i_p \leftarrow 3 ; in_C.fid \leftarrow 0 ; in_C.d_{md} \leftarrow \varepsilon$ 
  11  $(st_{adv}, out_{adv}) \leftarrow \text{ACPT1}(i_c, in_C)$ 
  12  $(c'_{fkm}, c'_{md}) \leftarrow \text{STEP}(M_{\text{accept}}, i_p, c_{fkm})$  // Now,  $c'_{fkm} = \text{AES-ECB.ENC}(k_m, d_{fkm})$ 
  // Fabricate file ciphertext
  13  $c_0 \parallel c_1 \leftarrow c'_{fkm} ; c''_{fkm} \leftarrow c_0 \parallel c_0$  // for  $|c_0| = |c_1| = 128$  bits
  14  $k_f'' \leftarrow 0^{128}$  // Result of decrypting  $c'_{fkm}$ , an artifact of MEGA's encryption  $E^{\text{MFILE}}$ 
  15  $c_f'' \leftarrow E_{k_f''}^{\text{MFILE}'}(n_f, \tau_f)$  // Creating some file encryption that produces tag  $\tau_f$ , see [6, Figure 6].
  16  $c''_{md} \leftarrow \text{AES-CBC.ENC}(k_f'', \varepsilon)$ 
  // Serve fabricated file
  17  $i_p \leftarrow i_p + 1 ; in_C.fid \leftarrow 0$ 
  18  $(st_{adv}, out_{adv}) \leftarrow \text{GET1}(i_c, in_C)$ 
  19  $m_C \leftarrow \text{STEP}(M_{\text{get}}, i_p, \{(in_C.fid, c'_{fkm}, c''_{md})\})$ 
  20  $m_C \leftarrow \text{STEP}(M_{\text{get}}, i_p, c_f'')$  // client succeeds because  $c_f''$  was encrypted with  $k_f''$ 
  21 return

```

Figure 19: Adversary running the integrity attack from [6, Section V] with a valid plaintext-ciphertext pair from the sharing feature exploited in follow-up work by Albrecht et al. [3, Section 2.2].

## B Building Block Details

This section specifies the building blocks used in Section 6 to prove the security of our cloud storage scheme in more detail.

### B.1 Standard Building Blocks

We use the following standard building blocks:

**PRF.** Let  $F: \{0, 1\}^{kl} \times \{0, 1\}^* \rightarrow \{0, 1\}^{kl}$  be a function family. We define  $\mathbf{G}_F^{\text{PRF-1}}$  to be the multi-user [8] PRF game in which an adversary  $\mathcal{A}$  can generate new PRF keys through oracle **NEW** and through oracle  $\text{FN}(\cdot, \cdot)$ , on input  $i, x$ , gets back  $F(k_i, x)$ , where  $k_i$  is the  $i^{\text{th}}$  key generated by **NEW**. Game  $\mathbf{G}_F^{\text{PRF-0}}$  is identical, except that  $\text{FN}(i, x)$  returns a consistently sampled random string in  $\{0, 1\}^{kl}$ . We define the advantage of adversary  $\mathcal{A}$  against the PRF security of  $F$  to be  $\text{Adv}_F^{\text{PRF}}(\mathcal{A}) := \Pr[\mathbf{G}_F^{\text{PRF-1}}] - \Pr[\mathbf{G}_F^{\text{PRF-0}}]$ .

**AEAD.** We define a nonce-based authenticated encryption scheme with associated data [51] as  $\text{AEAD} = (\text{Enc}, \text{Dec})$ , where we write  $c \leftarrow \text{Enc}(k, n, m, ad)$  and  $m \leftarrow \text{Dec}(k, n, c, ad)$  for encryption resp. decryption of a message  $m \in \{0, 1\}^*$  resp. ciphertext  $c \in \{0, 1\}^*$  using key  $k \in \{0, 1\}^{kl}$ ,

<b>Game <math>\mathbf{G}_{\text{AEAD}}^{\text{IND\\$-}b}(\mathcal{A})</math>:</b> 1 <b>global</b> $u, S_n, T_{\text{AEAD}}$ 2 $b' \leftarrow \mathcal{A}^{\text{New,Enc,Dec}}()$ 3 <b>return</b> $b'$  <b>NEW():</b> 4 $u \leftarrow u + 1$ 5 $k_u \leftarrow \{0, 1\}^{kl}$  <b>ENC(<math>i, n, m, ad</math>):</b> 6 <b>if</b> $(i, n) \in S_n$ : 7 <b>return</b> $\perp$ 8 $S_n \leftarrow^{\cup} \{(i, n)\}$ 9 $c_1 \leftarrow \text{Enc}(k_i, n, m, ad)$ 10 $c_0 \leftarrow \{0, 1\}^{ c_1 }$ 11 $T_{\text{AEAD}}[i, n, c_b, ad] \leftarrow m$ 12 <b>return</b> $c_b$  <b>DEC(<math>i, n, c, ad</math>):</b> 13 $m_1 \leftarrow \text{Dec}(k_i, n, c, ad)$ 14 $m_0 \leftarrow T_{\text{AEAD}}[i, n, c, ad]$ 15 <b>return</b> $m_b$	<b>Game <math>\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-}b}(\mathcal{A})</math>:</b> 1 <b>global</b> $u, S_n, T_{\text{AEAD}}$ 2 $b' \leftarrow \mathcal{A}^{\text{New,Enc,Dec}}()$ 3 <b>return</b> $b'$  <b>NEW():</b> 4 $u \leftarrow u + 1$ 5 $k_u \leftarrow \{0, 1\}^{kl}$  <b>ENC(<math>i, n, m_0, m_1, ad</math>):</b> 6 <b>if</b> $(i, n) \in S_n \vee  m_0  \neq  m_1 $ : 7 <b>return</b> $\perp$ 8 $S_n \leftarrow^{\cup} \{(i, n)\}$ 9 $c_0 \leftarrow \text{Enc}(k_i, n, m_0, ad)$ 10 $c_1 \leftarrow \text{Enc}(k_i, n, m_1, ad)$ 11 $S_c \leftarrow^{\cup} \{(i, n, c_b, ad)\}$ 12 <b>return</b> $c_b$  <b>DEC(<math>i, n, c, ad</math>):</b> 13 <b>if</b> $(i, n, c, ad) \in S_c$ : 14 <b>return</b> $\perp$ 15 $m \leftarrow \text{Dec}(k_i, n, c, ad)$ 16 <b>return</b> $m$	<b>Game <math>\mathbf{G}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{A})</math>:</b> 17 <b>global</b> $\text{win}, u, S_n, S_c$ 18 $\mathcal{A}^{\text{New,Enc,Dec}}()$ 19 <b>return</b> $\text{win}$  <b>NEW():</b> 20 $u \leftarrow u + 1$ 21 $k_u \leftarrow \{0, 1\}^{kl}$  <b>ENC(<math>i, n, m, ad</math>):</b> 22 <b>if</b> $(i, n) \in S_n$ : 23 <b>return</b> $\perp$ 24 $S_n \leftarrow^{\cup} \{(i, n)\}$ 25 $c \leftarrow \text{Enc}(k_i, n, m, ad)$ 26 $S_c \leftarrow^{\cup} \{(i, n, c, ad)\}$ 27 <b>return</b> $c$  <b>DEC(<math>i, n, c, ad</math>):</b> 28 $m \leftarrow \text{Dec}(k_i, n, c, ad)$ 29 <b>if</b> $(i, n, c, ad) \notin S_c \wedge m \neq \perp$ : 30 $\text{win} \leftarrow \text{true}$ 31 <b>return</b> $m$
---	--	---

Figure 20: Games formalizing authenticated encryption (IND\$), confidentiality (IND-CCA) and integrity of ciphertexts (INT-CTXT) of an authenticated encryption scheme with associated data AEAD.

nonce  $n \in \{0, 1\}^{nl}$ , and associated data  $ad \in \{0, 1\}^*$ .

Below, we formally define the advantage of an adversary  $\mathcal{A}$  against the IND\$, IND-CCA and INT-CTXT security of AEAD in the multi-user setting denoted as  $\text{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{A})$ ,  $\text{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{A})$ , and  $\text{Adv}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{A})$ , respectively. The games formalizing IND\$, IND-CCA and INT-CTXT of an AEAD scheme are given in Figure 20. We define

$$\begin{aligned}
\text{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{A}) &= \Pr [\mathbf{G}_{\text{AEAD}}^{\text{IND\$-1}}] - \Pr [\mathbf{G}_{\text{AEAD}}^{\text{IND\$-0}}], \\
\text{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{A}) &= \Pr [\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-1}}] - \Pr [\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-0}}], \text{ and} \\
\text{Adv}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{A}) &= \Pr [\mathbf{G}_{\text{AEAD}}^{\text{INT-CTXT}}].
\end{aligned}$$

**MAC.** A message authentication code  $\text{MAC} = (\text{Tag}, \text{Vrfy})$  consists of a tag generation algorithm  $\text{Tag}: \{0, 1\}^{kl} \times \{0, 1\}^* \rightarrow \{0, 1\}^{tl}$  and a verification algorithm  $\text{Vrfy}: \{0, 1\}^{kl} \times \{0, 1\}^* \times \{0, 1\}^{tl} \rightarrow \{\perp, \top\}$ . Both take a key  $k$  of length  $kl$  and an arbitrary length message  $m$  as input,  $\text{Tag}$  produces a tag of length  $tl$ . Correctness is defined as  $\text{Vrfy}(k, m, \text{Tag}(k, m)) = \top$ .

## B.2 OPRF Building Block

We describe the 2HashDH OPRF in Figure 21, which we will use inlined in the protocols of CSS. It is based on a group  $G$  of prime order  $q$  and two hash functions  $H_1: \{0, 1\}^* \rightarrow G$  and  $H_2: \{0, 1\}^* \times G \rightarrow$

<u>2HashDH.Req(<math>x</math>):</u> 1 $r \leftarrow \mathbb{Z}_q$ ; $\alpha \leftarrow H_1(x)^r$ 2 <b>return</b> $((r, x), \alpha)$  <u>2HashDH.Finalize(<math>(r, x), \beta</math>):</u> 3 <b>if</b> $\beta \notin G$ : <b>return</b> $\perp$ 4 $\gamma \leftarrow \beta^{1/r}$ 5 $y \leftarrow H_2(x, \gamma)$ 6 <b>return</b> $y$	<u>2HashDH.KeyGen():</u> 7 $k \leftarrow \mathbb{Z}_q$ 8 <b>return</b> $k$  <u>2HashDH.BlindEv(<math>k, \alpha</math>):</u> 9 <b>if</b> $\alpha \notin G$ : <b>return</b> $\perp$ 10 $\beta \leftarrow \alpha^k$ 11 <b>return</b> $\beta$
--	--

Figure 21: The 2HashDH OPRF operations (client on the left, server on the right) for a group  $G$  and hash functions  $H_1: \{0, 1\}^* \rightarrow G$  and  $H_2: \{0, 1\}^* \rightarrow \{0, 1\}^{kl}$  which compute  $2\text{HashDH}_k(x) := H_2(x, H_1(x)^k)$  for a client-chosen value  $x$  and server key  $k$ . The colored values  $\alpha$  and  $\beta$  are the messages exchanged between client and server.

$\{0, 1\}^{kl}$ . The server-side key  $k \in \mathbb{Z}_q$  is generated using **KeyGen()**. The client computes a request via **Req( $x$ )**, which samples a random blinding value  $r \in \mathbb{Z}_q$ , computes the request to be sent to the server as  $\alpha \leftarrow H_1(x)^r$  and outputs it together with local state  $(r, x)$ . The server computes the blind evaluation **BlindEv( $k, \alpha$ )** as  $\beta \leftarrow \alpha^k$ . The client finally in **Finalize** unblinds  $\beta$  using  $r$  and outputs  $y \leftarrow H_2(x, \gamma)$ .

## C Proofs of Client-To-Client Security for CSS

### C.1 Proof of Theorem 6.1

This section provides the full details of the proof of the selective client-to-client confidentiality of CSS, as stated in Theorem 6.1. Recall that  $\mathcal{Q}_{\text{ORC}}(\mathcal{A})$  and  $q_{\text{ORC}}(\mathcal{A})$  denote the maximum number of queries to oracle ORC by adversary  $\mathcal{A}$  in total and per user, respectively. We use the convention that  $\mathcal{Q}_{\text{PUT1}}(\mathcal{A})$  and  $\mathcal{Q}_{\text{UPD1}}(\mathcal{A})$  include indirect queries made via oracle CHALL. To count only queries via oracle CHALL, we let  $\text{CPUT1}$  and  $\text{CUPD1}$  denote oracle CHALL with input ORC set to PUT1 and UPD1, respectively.

The proof proceeds through a sequence of games  $G_0 - G_7$ , which all provide the oracles from game  $\mathbf{G}_{\text{CSS}, n, \mathcal{PW}_n}^{\text{C2CConfS}}$ . Each game will begin by sampling a challenge bit  $b$ , which determines the file that is encrypted in oracle CHALL. The finalize procedure in each game is equivalent to that of  $\mathbf{G}_{\text{CSS}, n, \mathcal{PW}_n}^{\text{C2CConfS}}$ . Namely, the game runs  $\mathcal{A}$  to obtain  $b'$ . It returns  $\neg \text{trivial} \wedge (b' = b)$ , where **trivial** is the predicate (defined on line 5 of Figure 5) indicating if  $\mathcal{A}$  performed a trivial attack. Throughout, we let  $\Pr[G_i]$  denote the probability that  $G_i(\mathcal{A})$  returns **true**.

We model the hash functions  $H_1, H_2$  used in account registration and authentication as (programmable) random oracles, implemented through lazy sampling with consistency ensured by tables  $\text{T}_{H_1}[\cdot, \cdot]$  and  $\text{T}_{H_2}[\cdot, \cdot, \cdot]$ . The adversary can interact with the tables through oracles  $\text{RO}_1$  and  $\text{RO}_2$ .

**Games  $G_0 - G_3$ : Replacing  $rw$  by random.** We begin with game  $G_0$ , which is equivalent to  $\mathbf{G}_{\text{CSS}, n, \mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A})$ . In particular, the client-side output  $rw \in \{0, 1\}^{kl}$  from the OPRF flow in protocols  $\text{CSS}_{\text{areg}}$  and  $\text{CSS}_{\text{auth}}$  is computed using the random oracles  $\text{RO}_1$  and  $\text{RO}_2$  in place of  $H_1$  and  $H_2$ . For example, line 3 of Figure 8 is replaced by

<u><math>G_0: \text{CSS}_{\text{areg}}^{(C;1)}</math>:</u> 3 $rw \leftarrow \text{RO}_2(\text{aid}, pw, (\text{RO}_1(\text{aid}, pw))^k)$
--

where we only display the modified line.

The effect of this change is only to model the hash functions as random oracles. Hence

$$\mathbf{Adv}_{\text{CSS}_{n,\mathcal{PW}_n}}^{\text{C2CConfS}}(\mathcal{A}) := 2 \cdot \Pr[G_0] - 1. \quad (8)$$

In the first sequence of games, the goal is to move to game  $G_3$ , where  $rw$  is replaced by an independent random string in  $\{0,1\}^{kl}$  for all honest (uncompromised) users.

Specifically,  $G_3$  differs from  $G_0$  in the following ways. In protocol  $\text{CSS}_{\text{areg}}$ , when the adversary registers a user  $i_u \notin S_{\text{comp}}$  (by calling oracle  $\text{AREG1}$  for  $i_u$ ) the calls to  $\text{RO}_1$  and  $\text{RO}_2$  computing  $\gamma$  and  $rw$  are replaced with consistent random sampling from  $G$  and  $\{0,1\}^{kl}$ , respectively. That is, line 3 in  $G_0$  is replaced by the following lines

```

G3:  $\text{CSS}_{\text{areg}}^{(C:1)}$ :
3  if  $i_u \in S_{\text{comp}}$ :
     $\gamma \leftarrow (\text{RO}_1(\text{aid}, \text{pw}))^k$ ;  $rw \leftarrow \text{RO}_2(\text{aid}, \text{pw}, \gamma)$ 
4  else:
5     $T_1[\text{aid}, \text{pw}] \leftarrow_{\$} G$ ;  $\gamma \leftarrow (T_1[\text{aid}, \text{pw}])^k$ 
6     $T_{rw}[\text{aid}, \text{pw}, \gamma] \leftarrow_{\$} \{0,1\}^{kl}$ ;  $rw \leftarrow T_{rw}[\text{aid}, \text{pw}, \gamma]$ 

```

The same change is also applied in  $\text{CSS}_{\text{auth}}$  for all  $i_u \notin S_{\text{comp}}$ . That is,  $\alpha \leftarrow (\text{RO}_1(\text{aid}, \text{pw}))^r$  on line 15 of Figure 8 is replaced by

```

G3:  $\text{CSS}_{\text{auth}}^{(C:1)}$ :
15 if  $i_u \in S_{\text{comp}}$ :  $\alpha \leftarrow (\text{RO}_1(\text{aid}, \text{pw}))^r$  else:  $\alpha \leftarrow (T_1[\text{aid}, \text{pw}])^r$ 

```

and on line 22, computing  $rw \leftarrow \text{RO}_2(\text{aid}, \text{pw}, \gamma)$  is replaced by

```

G3:  $\text{CSS}_{\text{auth}}^{(C:2)}$ :
22 if  $i_u \in S_{\text{comp}}$ :  $rw \leftarrow \text{RO}_2(\text{aid}, \text{pw}, \gamma)$ 
23 else:
24   if  $T_{rw}[\text{aid}, \text{pw}, \gamma] = \perp$ :  $T_{rw}[\text{aid}, \text{pw}, \gamma] \leftarrow_{\$} \{0,1\}^{kl}$ 
25    $rw \leftarrow T_{rw}[\text{aid}, \text{pw}, \gamma]$ 

```

We claim that the change from  $G_0$  to  $G_3$  is indistinguishable to an adversary, except with some small probability. Formally, we want to bound  $\Pr[G_0] - \Pr[G_3]$ . To this end, we introduce two additional games,  $G_1$  and  $G_2$ . In the following, we focus on the places where these games differ from  $\mathbf{G}_{\text{CSS}}^{\text{C2CConfS}}$ . To this end, we let  $\text{AUTH2}(i_p, m) := \text{STEP}(\text{CSS}_{\text{auth}}, i_p, m)$  be an oracle which takes care of the query processing when  $\mathcal{A}$  runs the *second* client-side step of protocol  $\text{CSS}_{\text{auth}}$  through oracle  $\text{STEP}$ .

Games  $G_1$  and  $G_2$  function exactly like  $G_3$ , except in oracles  $\text{RO}_1$  and  $\text{RO}_2$ , and where  $rw$  is computed in oracles  $\text{AREG1}$  and  $\text{AUTH2}$ . There, as shown in Figure 22, a flag **bad** is set to true if oracle  $\text{RO}_1$  resp.  $\text{RO}_2$  is queried by the adversary on inputs  $(a, b)$  resp.  $(a, b, c)$  such that  $T_1[a, b] \neq \perp$ , resp.  $T_{rw}[a, b, c] \neq \perp$ . (That is, if the adversary queries the random oracles on inputs corresponding to the  $\text{aid}$  and  $\text{pw}$  of a non-compromised user that has been previously registered and/or authenticated.) Game  $G_1$ , which contains the boxed code, then overwrites  $T_{H1}[a, b]$  resp.  $T_{H2}[a, b, c]$  with  $T_1[a, b]$ , resp.  $T_{rw}[a, b, c]$ , whereas  $G_2$  does nothing. Additionally, **bad** is also set to true in  $\text{AREG1}$  and  $\text{AUTH2}$  if the adversary has already queried  $\text{RO}_1$  resp.  $\text{RO}_2$  on input  $(\text{aid}, \text{pw})$  resp.  $(\text{aid}, \text{pw}, \gamma)$ , such that  $T_{H1}[\text{aid}, \text{pw}]$  resp.  $T_{H2}[\text{aid}, \text{pw}, \gamma]$  are initiated. Game  $G_1$  then overwrites the affected entry in tables  $T_1$  and/or  $T_{rw}$  with these values, whereas  $G_2$  again does nothing.

It is clear that  $G_2$  is equivalent to  $G_3$  (the only difference being that the **bad** flag is set, which does not affect the adversary's chances of winning), and hence that  $\Pr[G_2] = \Pr[G_3]$ . Let **Bad** denote the event that  $G_2$  sets the **bad** flag.  $G_1$  and  $G_2$  are identical-until-**Bad**, so  $\Pr[G_1] - \Pr[G_2] \leq \Pr[\text{Bad}]$  by the fundamental lemma of game playing [11]. Finally, we claim that  $G_1$  is equivalent to  $G_0$ . To see this, note that in  $G_1$ , all  $rw$  values are computed in a way that is consistent with the random oracles – any potential inconsistency is fixed through the programming after the bad event. Hence,  $\Pr[G_0] = \Pr[G_1]$ .

Through standard equation rewriting, this gives

$$\Pr[G_0] = \Pr[\text{Bad}] + \Pr[G_3]. \quad (9)$$

**Games  $G_0$ – $G_3$ : Bounding **Bad**.** Next, we proceed to bound the probability that  $G_2$  sets **bad**. Let  $G'_2$  be identical to  $G_2$ , except that instead of returning  $\neg \text{trivial} \wedge (b = b')$  in the finalize procedure, it returns **bad**, such that  $\Pr[\text{Bad}] = \Pr[G'_2]$ . Next, we construct a game  $G''_2$ , which is identical to  $G'_2$ , except that it additionally keeps track of the queries made by adversary  $\mathcal{A}$ , as follows. All queries  $\text{RO}_1(a, b)$  resp.  $\text{RO}_2(a, b, c)$  by adversary  $\mathcal{A}$  (notably, not by the game itself!) are collected in sets  $S_{H1}$  and  $S_{H2}$ , respectively, through  $S_{H1} \stackrel{\cup}{\leftarrow} \{(a, b)\}$ ,  $S_{H2} \stackrel{\cup}{\leftarrow} \{(a, b, c)\}$ . Queries to oracle  $\text{AREG1}$  and  $\text{AUTH2}$  which initialize table entries  $T_I[\text{aid}, pw]$  resp.  $T_{rw}[\text{aid}, pw, \gamma]$  are collected in  $S_I \stackrel{\cup}{\leftarrow} \{(\text{aid}, pw)\}$  and  $S_{rw} \stackrel{\cup}{\leftarrow} \{(\text{aid}, pw, \gamma)\}$ . Furthermore, instead of returning **bad**,  $G''_2$  returns  $(S_I \cap S_{H1}) \cup (S_{rw} \cap S_{H2}) \neq \emptyset$ . Games  $G'_2$  and  $G''_2$  are given in Figure 23.

We claim that the two return statements are equivalent. The flag **bad** is set to **true** for one of two reasons: (1) when the adversary<sup>16</sup> makes a query  $\text{RO}_1(a, b)$  resp.  $\text{RO}_2(a, b, c)$  such that  $T_I[a, b]$  resp.  $T_{rw}[a, b, c]$  is already initialized. This implies that  $(a, b) \in S_{H1} \cap S_I$  or  $(a, b, c) \in S_{H2} \cap S_{rw}$ . Or (2)  $\mathcal{A}$  makes a query using user values  $\text{aid}, pw$  and  $\gamma$  such that it has previously queried  $\text{RO}_1(\text{aid}, pw)$  or  $\text{RO}_2(\text{aid}, pw, \gamma)$ . This implies that  $(\text{aid}, pw) \in S_{H1} \cap S_I$  or  $(\text{aid}, pw, \gamma) \in S_{H2} \cap S_{rw}$ . Hence,

$$\Pr[\text{Bad}] = \Pr[G'_2] = \Pr[G''_2]. \quad (10)$$

Next, we construct an adversary  $\mathcal{B}_{pg}$  against the password-guessing game (see Figure 4) for the distribution  $\mathcal{PW}_n$  such that

$$\text{Adv}_{n, \mathcal{PW}_n}^{\text{PG}}(\mathcal{B}_{pg}) \geq \Pr[G''_2]. \quad (11)$$

On a high level, adversary  $\mathcal{B}_{pg}$  will use the RO queries from  $\mathcal{A}$  to guess passwords: Whenever adversary  $\mathcal{A}$  makes a query  $\text{RO}_1(a, b)$  or  $\text{RO}_2(a, b, c)$ ,  $\mathcal{B}_{pg}$  lets  $S_H \stackrel{\cup}{\leftarrow} \{(a, b)\}$ . When  $\mathcal{A}$  halts,  $\mathcal{B}_{pg}$  checks for every element  $(a, b) \in S_H$  if there exists a user  $i_u$  s.t. the account ID of  $i_u$  is  $a$ . If so, it queries  $\text{TEST}(i_u, b)$ . This way,  $\mathcal{B}_{pg}$  makes a password guess for every element  $(a, b)$  resp.  $(a, b, c)$  in  $S_{H1}$  resp.  $S_{H2}$  of  $G''_2$  where  $a$  corresponds to an account ID of a registered user. In particular, it will query  $\text{TEST}(i_u, b)$  for each  $(a, b) \in S_I \cap S_{H1}$  and each  $(a, b, c) \in S_{rw} \cap S_{H2}$  such that  $T_{\text{uaid}}[i_u] = a$ . Since  $S_I$  and  $S_{rw}$  only contain account IDs and passwords of honestly registered users not in  $S_{\text{comp}}$ , adversary  $\mathcal{B}_{pg}$  wins game  $\mathbf{G}_{n, \mathcal{PW}_n}^{\text{PG}}$  if  $(S_I \cap S_{H1}) \cup (S_{rw} \cap S_{H2}) \neq \emptyset$ , proving Equation (11).

Formally, we also need to take care that adversary  $\mathcal{B}_{pg}$  can properly simulate game  $G''_2$  for  $\mathcal{A}$ . This is possible thanks to the selective compromises in the password guessing game, and the fact that for non-compromised users, the OPRF values  $\gamma$  and  $rw$  are independent from the user's password in game  $G''_2$ . The selective compromises are used to handle queries from  $\mathcal{A}$  on  $i_u \in S_{\text{comp}}$ : In the first stage of the simulation, adversary  $\mathcal{B}_{pg}$  runs  $\mathcal{A}$  to get  $S_{\text{comp}}$  and returns this to its own challenger, receiving the vector  $\mathbf{pw}_{\text{comp}}$  of compromised passwords back. In the second stage,  $\mathcal{B}_{pg}$  uses  $\mathbf{pw}_{\text{comp}}[i_u]$  in place of  $\mathbf{pw}[i_u]$  to handle queries from  $\mathcal{A}$  on  $i_u \in S_{\text{comp}}$ .

<sup>16</sup>Note that this relies on the assumption that *aids* are unique per user, so that there is no overlap between game queries to  $\text{RO}_1$  and  $\text{RO}_2$  for users in  $S_{\text{comp}}$  and queries from  $\mathcal{A}$ .



<p>Games <math>\underline{G_1}, G_2</math>:</p> <p><math>\vdots</math></p> <p><math>\underline{\text{AREG1}(i_u, in_C)}</math>:</p> <p><math>\vdots</math></p> <p>3 <b>if</b> <math>i_u \in S_{\text{comp}}</math>: <math>\gamma \leftarrow (\text{RO}_1(\text{aid}, pw))^k</math>; <math>rw \leftarrow \text{RO}_2(\text{aid}, pw, \gamma)</math></p> <p>4 <b>else</b>:</p> <p>5   <math>T_I[\text{aid}, pw] \leftarrow \\$ G</math> // <math>\text{aid}</math> is unique and <math>\text{AREG1}</math> can only be called once, so no resampling</p> <p>6   <b>if</b> <math>T_{H1}[\text{aid}, pw] \neq \perp</math>:</p> <p>      <math>\text{bad} \leftarrow \text{true}</math>; <math>\boxed{T_I[\text{aid}, pw] \leftarrow T_{H1}[\text{aid}, pw]}</math></p> <p>7   <math>\gamma \leftarrow (T_I[\text{aid}, pw])^k</math></p> <p>8   <math>T_{rw}[\text{aid}, pw, \gamma] \leftarrow \\$ \{0, 1\}^{kl}</math></p> <p>9   <b>if</b> <math>T_{H2}[\text{aid}, pw, \gamma] \neq \perp</math>:</p> <p>      <math>\text{bad} \leftarrow \text{true}</math>; <math>\boxed{T_{rw}[\text{aid}, pw, \gamma] \leftarrow T_{H2}[\text{aid}, pw, \gamma]}</math></p> <p>10   <math>rw \leftarrow T_{rw}[\text{aid}, pw, \gamma]</math></p> <p><math>\vdots</math></p> <p><math>\underline{\text{AUTH1}(i_u, in_C)}</math>:</p> <p><math>\vdots</math></p> <p>15 <b>if</b> <math>i_u \in S_{\text{comp}}</math>: <math>\alpha \leftarrow (\text{RO}_1(\text{aid}, pw))^r</math> <b>else</b>: <math>\alpha \leftarrow (T_I[\text{aid}, pw])^r</math></p> <p><math>\vdots</math></p> <p><math>\underline{\text{AUTH2}(i_p, m_S)}</math> // Oracle STEP for <math>\text{CSS}_{\text{auth}}^{(C:2)}</math>:</p> <p><math>\vdots</math></p> <p>22 <b>if</b> <math>i_u \in S_{\text{comp}}</math>: <math>rw \leftarrow \text{RO}_2(\text{aid}, pw, \gamma)</math></p> <p>23 <b>else</b>:</p> <p>24   <b>if</b> <math>T_{rw}[\text{aid}, pw, \gamma] = \perp</math>: <math>T_{rw}[\text{aid}, pw, \gamma] \leftarrow \\$ \{0, 1\}^{kl}</math></p> <p>25   <b>if</b> <math>T_{H2}[\text{aid}, pw, \gamma] \neq \perp</math>:</p> <p>      <math>\text{bad} \leftarrow \text{true}</math>; <math>\boxed{T_{rw}[\text{aid}, pw, \gamma] \leftarrow T_{H2}[\text{aid}, pw, \gamma]}</math></p> <p>26   <math>rw \leftarrow T_{rw}[\text{aid}, pw, \gamma]</math></p> <p><math>\vdots</math></p>	
<p>Game <math>G_0, G_3</math></p> <p><math>\underline{\text{RO}_1(a, b)}</math>:</p> <p>27 <b>if</b> <math>T_{H1}[a, b] = \perp</math>:</p> <p>28   <math>T_{H1}[a, b] \leftarrow \\$ G</math></p> <p>29 <b>return</b> <math>T_{H1}[a, b]</math></p> <p><math>\underline{\text{RO}_2(a, b, c)}</math>:</p> <p>30 <b>if</b> <math>T_{H2}[a, b, c] = \perp</math>:</p> <p>31   <math>T_{H2}[a, b, c] \leftarrow \\$ \{0, 1\}^{kl}</math></p> <p>32 <b>return</b> <math>T_{H2}[a, b, c]</math></p>	<p>Game <math>\underline{G_1}, G_2</math></p> <p><math>\underline{\text{RO}_1(a, b)}</math>:</p> <p>33 <b>if</b> <math>T_{H1}[a, b] = \perp</math>: <math>T_{H1}[a, b] \leftarrow \\$ G</math></p> <p>34 <b>if</b> <math>T_I[a, b] \neq \perp</math>:</p> <p>35   <math>\text{bad} \leftarrow \text{true}</math>; <math>\boxed{T_{H1}[a, b] \leftarrow T_I[a, b]}</math></p> <p>36 <b>return</b> <math>T_{H1}[a, b]</math></p> <p><math>\underline{\text{RO}_2(a, b, c)}</math>:</p> <p>37 <b>if</b> <math>T_{H2}[a, b, c] = \perp</math>: <math>T_{H2}[a, b, c] \leftarrow \\$ \{0, 1\}^{kl}</math></p> <p>38 <b>if</b> <math>T_{rw}[a, b, c] \neq \perp</math>:</p> <p>39   <math>\text{bad} \leftarrow \text{true}</math>; <math>\boxed{T_{H2}[a, b, c] \leftarrow T_{rw}[a, b, c]}</math></p> <p>40 <b>return</b> <math>T_{H2}[a, b, c]</math></p>

Figure 22: Games and additional oracles for the first game hops in proof of Theorem 6.1.  $G$  is a group of prime order  $q$  with generator  $g$ .

<u>Games <math>G'_2, G''_2</math>:</u> $\vdots$ 1 <b>return</b> bad <span style="float: right;">// <math>G'_2</math></span> 2 <b>return</b> $(S_I \cap S_{H1}) \cup (S_{rw} \cap S_{H2}) \neq \emptyset$ <span style="float: right;">// <math>G''_2</math></span>	
<u>AREG1(<math>i_u, in_C</math>):</u> $\vdots$ 3 $k \leftarrow \mathbb{Z}_q$ 4 <b>if</b> $i_u \in S_{comp}$ : 5 <b>if</b> $T_{H1}[aid, pw] = \perp$ : $T_{H1}[aid, pw] \leftarrow G$ 6 $\gamma \leftarrow (T_{H1}[aid, pw])^k$ 7 <b>if</b> $T_{H2}[aid, pw, \gamma] = \perp$ : $T_{H2}[aid, pw, \gamma] \leftarrow \{0, 1\}^{kl}$ 8 $rw \leftarrow T_{H2}[aid, pw, \gamma]$ 9 <b>else</b> : $T_I[aid, pw] \leftarrow G$ 10 <b>if</b> $T_{H1}[aid, pw] \neq \perp$ : <b>bad</b> $\leftarrow$ <b>true</b> // $(aid, pw) \in S_{H1}$ <span style="float: right;">// <math>G'_2</math></span> 11 $S_I \leftarrow \bigcup \{(aid, pw)\}$ <span style="float: right;">// <math>G''_2</math></span> 12 $\gamma \leftarrow (T_I[aid, pw])^k$ 13 $T_{rw}[aid, pw, \gamma] \leftarrow \{0, 1\}^{kl}$ 14 <b>if</b> $T_{H2}[aid, pw, \gamma] \neq \perp$ : <b>bad</b> $\leftarrow$ <b>true</b> ; // $(aid, pw, \gamma) \in S_{H2}$ <span style="float: right;">// <math>G'_2</math></span> 15 $S_{rw} \leftarrow \bigcup \{(aid, pw, \gamma)\}$ <span style="float: right;">// <math>G''_2</math></span> 16 $rw \leftarrow T_{rw}[aid, pw, \gamma]$ $\vdots$ <u>AUTH1(<math>i_u, in_C</math>):</u> $\vdots$ 17 <b>if</b> $i_u \in S_{comp}$ : $\alpha \leftarrow (T_{H1}[aid, pw])^r$ <b>else</b> : $\alpha \leftarrow (T_I[aid, pw])^r$ $\vdots$ <u>AUTH2(<math>i_p, m_S</math>):</u> $\vdots$ 18 <b>if</b> $i_u \in S_{comp}$ : 19 <b>if</b> $T_{H2}[aid, pw, \gamma] = \perp$ : $T_{H2}[aid, pw, \gamma] \leftarrow \{0, 1\}^{kl}$ 20 $rw \leftarrow T_{H2}[aid, pw, \gamma]$ 21 <b>else</b> : 22 <b>if</b> $T_{rw}[aid, pw, \gamma] = \perp$ : $T_{rw}[aid, pw, \gamma] \leftarrow \{0, 1\}^{kl}$ 23 <b>if</b> $T_{H2}[aid, pw, \gamma] \neq \perp$ : <b>bad</b> $\leftarrow$ <b>true</b> ; // $(aid, pw, \gamma) \in S_{H2}$ <span style="float: right;">// <math>G'_2</math></span> 24 $S_{rw} \leftarrow \bigcup \{(aid, pw, \gamma)\}$ <span style="float: right;">// <math>G''_2</math></span> 25 $rw \leftarrow T_{rw}[aid, pw, \gamma]$ $\vdots$	

  

<u>Games <math>G'_2, G''_2</math></u> <u>RO<sub>1</sub>(<math>a, b</math>):</u> 26 <b>if</b> $T_{H1}[a, b] = \perp$ : $T_{H1}[a, b] \leftarrow G$ 27 <b>if</b> $T_I[a, b] \neq \perp$ : <b>bad</b> $\leftarrow$ <b>true</b> // $(a, b) \in S_I$ // $G'_2$ 28 $S_{H1} \leftarrow \bigcup \{(a, b)\}$ // $G''_2$ 29 <b>return</b> $T_{H1}[a, b]$	<u>RO<sub>2</sub>(<math>a, b, c</math>):</u> 30 <b>if</b> $T_{H2}[a, b, c] = \perp$ : $T_{H2}[a, b, c] \leftarrow \{0, 1\}^{kl}$ 31 <b>if</b> $T_{rw}[a, b, c] \neq \perp$ : <b>bad</b> $\leftarrow$ <b>true</b> // $(a, b, c) \in S_{rw}$ // $G'_2$ 32 $S_{H2} \leftarrow \bigcup \{(a, b, c)\}$ // $G''_2$ 33 <b>return</b> $T_{H2}[a, b, c]$
---	---

Figure 23: Games  $G'_2$  and  $G''_2$  for the first game hops in proof of Theorem 6.1.  $G$  is a group of prime order  $q$  with generator  $g$ .

If  $i_u \in (S_{\text{hon}} \setminus S_{\text{comp}})$ ,  $\mathcal{B}_{\text{pg}}$  instead samples  $\gamma$  and/or  $rw$  (depending on the oracle) uniformly at random and stores them for consistency in  $T_I[i_u]$  and  $T_{rw}[i_u, \gamma]$ , respectively. Note here that swapping  $aid, pw$  for  $i_u$  in the key for the tables is sound thanks to the assumption that account IDs are globally unique and the fact that the game enforces a single  $aid$  per registered user. With this strategy, adversary  $\mathcal{B}_{\text{pg}}$  perfectly simulates game  $G_2''$  for  $\mathcal{A}$ , making at most  $\mathcal{Q}_{\text{RO}_1}(\mathcal{A}) + \mathcal{Q}_{\text{RO}_2}(\mathcal{A})$  queries to oracle TEST.

**Game  $G_4$ : Replace  $k_{\text{kek}}$  and  $k_{\text{mac}}$  with random.**

We begin with a change in bookkeeping. Let  $G_3'$  be identical to  $G_3$ , except that table  $T_{rw}$  tracking the  $rw$  values of honest clients is indexed by  $i_u, \gamma$  rather than  $aid, pw, \gamma$ . Thanks to the one-to-one mapping between  $aid$  and  $i_u$ , this has no effect on the functionality of the game, implying  $\Pr[G_3] = \Pr[G_3']$ .

Next,  $G_4$  replaces  $k_{\text{kek}}$  and  $k_{\text{mac}}$ , derived from  $rw$  using  $F$  on lines 4, 7 of  $\text{CSS}_{\text{areg}}$ , and 23, 24 of  $\text{CSS}_{\text{auth}}^{(C:2)}$  (Figure 8) with consistently sampled random strings of length  $kl$ . We construct an adversary  $\mathcal{B}_{\text{prf}}$  such that

$$\Pr[G_3'] - \Pr[G_4] \leq \text{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B}_{\text{prf}}). \quad (12)$$

Adversary  $\mathcal{B}_{\text{prf}}$  acts as the challenger in game  $G_3'$ . In particular, it samples a bit  $b$  and uses it to determine which file to encrypt in queries from  $\mathcal{A}$  to oracle CHALL. In contrast to the challenger in game  $G_3'$ , adversary  $\mathcal{B}_{\text{prf}}$  does not itself sample the table entries in  $T_{rw}$  used as keys for  $F$ . Instead, it calls oracle NEW for every new  $i_u, \gamma$  combination (where  $i_u \notin S_{\text{comp}}$ ) that it sees in  $\mathcal{A}$ 's queries to AREG1 and AUTH2, storing the index of the new key in place of the key itself in table  $T_{rw}[i_u, \gamma]$ . Furthermore, instead of computing  $F$  itself for uncompromised clients, it calls oracle  $\text{FN}(T_{rw}[i_u, \gamma], x)$  for  $x \in \{\text{"kek"}, \text{"mac"}\}$  to compute  $k_{\text{kek}}$  and  $k_{\text{mac}}$ . This way,  $\mathcal{B}_{\text{prf}}$  simulates game  $G_3'$  when playing  $\mathbf{G}_{\text{F}}^{\text{PRF-1}}$  and  $G_4$  when playing  $\mathbf{G}_{\text{F}}^{\text{PRF-0}}$ , making at most  $\mathcal{Q}_{\text{AREG1}}(\mathcal{A}) + \mathcal{Q}_{\text{AUTH1}}(\mathcal{A})$  queries to both oracles NEW and FN. (Note that  $\mathcal{Q}_{\text{AUTH2}}(\mathcal{A}) \leq \mathcal{Q}_{\text{AUTH1}}(\mathcal{A})$ , since the second step of a protocol can only be run if the first has been executed, allowing us to bound the number subsequent protocol steps by the initial queries.)

When  $\mathcal{A}$  halts and returns bit  $b'$ , adversary  $\mathcal{B}_{\text{prf}}$  halts and returns 1 if  $b' = b$ , else it returns 0. This gives Equation (12).

**Game  $G_5$ : Replace encrypted master key with random.**

Next, in game  $G_5$ , the goal is to apply AEAD security for the encryption of the client master key  $k_{mk}$  in  $\text{CSS}_{\text{areg}}^{(C:1)}$  for honest users. All of the following changes apply only if  $i_u \notin S_{\text{comp}}$ ; otherwise the game remains unchanged.

In game  $G_5$ , the encryption  $c_{mk}$  of the master key is replaced by a randomly sampled string  $c'_{mk}$  of length  $|c_{mk}|$  in line 6 of  $\text{CSS}_{\text{areg}}^{(C:1)}$ . Additionally, the game stores the master key  $k_{mk}$  in table  $T_{\text{AEAD}}[(i_u, \gamma), n_{mk}, c'_{mk}, (aid, \text{"mk"})]$ , and instead of using Dec to decrypt  $c_{mk}$  in line 29 of  $\text{CSS}_{\text{auth}}^{(C:3)}$ , the table entry corresponding to the inputs from the server is used. Note that this means that if the table entry has not been initialized in  $\text{CSS}_{\text{areg}}$ ,  $k_{mk}$  is set to  $\perp$  in  $\text{CSS}_{\text{auth}}$ .

We construct an adversary  $\mathcal{B}_{\text{aead}}^1$  against the IND\$ security of AEAD such that

$$\Pr[G_4] - \Pr[G_5] \leq \text{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^1). \quad (13)$$

Adversary  $\mathcal{B}_{\text{aead}}^1$  acts like the challenger in game  $G_4$ , with the following exceptions. First, instead of sampling  $k_{\text{kek}}$  at random for each  $i_u, \gamma$  pair s.t.  $i_u \notin S_{\text{comp}}$  in AREG1 and the second step of AUTH, adversary  $\mathcal{B}_{\text{aead}}^1$  queries oracle NEW and stores the index of the new key in  $T_{\text{F}}[i_u, \gamma, \text{"kek"}]$ .

This gives  $\mathcal{Q}_{\text{New}}(\mathcal{B}_{\text{aead}}^1) \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) + \mathcal{Q}_{\text{AUTH1}}(\mathcal{A})$ . Second, instead of encrypting  $k_{mk}$  itself in AREG1, adversary  $\mathcal{B}_{\text{aead}}^1$  queries oracle ENC in the IND\$ game with key index  $T_F[i_u, \gamma, \text{"kek"}]$  and message input  $k_{mk}$  and uses the resulting ciphertext as  $c_{mk}$ . Note that thanks to the precondition  $i_u \notin S_{\text{reg}}$  in AREG1 and the fact that  $\text{CSS}_{\text{areg}}$  always succeeds on the client,  $\mathcal{A}$  can query oracle AREG1 at most once per  $i_u$ . Hence adversary  $\mathcal{B}_{\text{aead}}^1$  will make at most one ENC query per key in the IND\$ game, ensuring that nonce collisions are not a problem. Last, instead of decrypting  $c_{mk}$  in the third step of AUTH,  $\mathcal{B}_{\text{aead}}^1$  queries oracle DEC in its own game with key index  $T_F[i_u, \gamma, \text{"kek"}]$ . This way,  $\mathcal{B}_{\text{aead}}^1$  simulates game  $G_4$  and  $G_5$  for  $\mathcal{A}$  when the hidden bit in the IND\$ game is 1 and 0, respectively. We have  $\mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^1) \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A})$  and  $\mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^1) \leq \mathcal{Q}_{\text{AUTH1}}(\mathcal{A})$ .

Let  $b$  denote the bit sampled by  $\mathcal{B}_{\text{aead}}^1$  at the start of the simulation and  $b_{\mathcal{B}}$  the bit sampled by the challenger in  $\mathbf{G}_{\text{AEAD}}^{\text{IND\$}}$ . When  $\mathcal{A}$  halts and returns  $b'$ ,  $\mathcal{B}_{\text{aead}}^1$  returns 1 if  $b' = b$ , otherwise it returns 0. This gives Equation (13).

### Game $G_6$ : Avoid nonce collisions.

In  $G_6$ , we take care of nonce collisions for AEAD encryptions. These occur with a probability that relates to the nonce length  $nl$  and the number of operations per user which sample nonces. In  $\text{CSS}$ , the nonces for file resp. file key encryptions are sampled in subroutines `putFile` resp. `putHeader`, which are called by protocols  $\text{CSS}_{\text{put}}$  (both) and  $\text{CSS}_{\text{upd}}$  resp.  $\text{CSS}_{\text{accept}}$ . Formally, we introduce a **bad** flag which is set to true whenever nonces collide in both  $G_5$  and  $G_6$ , and game  $G_6$  only differs from  $G_5$  in that if **bad** is set, the nonce is resampled such that it is guaranteed to be fresh. That is, line 1 of both `putFile` and `putHeader` in Figure 10 is replaced by

```

G5, [G6]: putFile, putHeader:
1   $n \leftarrow \{0, 1\}^{nl}$ 
2  if  $n \in \{S_{n,i_u}^{\text{pF}}, S_{n,i_u}^{\text{pH}}\}$ :
    bad  $\leftarrow$  true;  $n \leftarrow \{0, 1\}^{nl} \setminus \{S_{n,i_u}^{\text{pF}}, S_{n,i_u}^{\text{pH}}\}$ 
3   $S_{n,i_u}^{\text{pF}}, S_{n,i_u}^{\text{pH}} \leftarrow \cup \{n\}$ 

```

// Boxed code only in  $G_6$

Here  $i_u \leftarrow T_u[i_c]$  is the user index tracked by the game corresponding to the client index  $i_c$  input by  $\mathcal{A}$  to oracles PUT1 and ACPT1. Since the games are identical until the bad flag is set,  $\Pr[G_5] - \Pr[G_6] \leq \Pr[G_5 \text{ sets bad}]$  by the fundamental lemma of game-playing [11].

Let  $\text{Bad}_{\text{pF}}$  resp.  $\text{Bad}_{\text{pH}}$  be the event that  $G_5$  sets **bad** in `putFile` resp. `putHeader`. We have that  $\Pr[G_5 \text{ sets bad}] = \Pr[\text{Bad}_{\text{pF}} \vee \text{Bad}_{\text{pH}}] = \Pr[\text{Bad}_{\text{pF}}] + \Pr[\text{Bad}_{\text{pH}}]$ , since the two bad events are independent. We now proceed to bound  $\Pr[\text{Bad}_{\text{pF}}]$  and  $\Pr[\text{Bad}_{\text{pH}}]$ . Recall that  $q_{\text{PUT1}}(\mathcal{A})$ ,  $q_{\text{UPD1}}(\mathcal{A})$  and  $q_{\text{ACPT1}}(\mathcal{A})$  are the maximum number of queries made by  $\mathcal{A}$  to oracles PUT1 and ACPT1 *per user*  $i_u$ , respectively.  $\Pr[\text{Bad}_{\text{pH}}]$  is then bounded from above by the collision probability  $C(2^{nl}, q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))$ , which denotes the probability that from  $q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A})$  elements sampled uniformly at random out of a set of size  $2^{nl}$ , any two collide. By a birthday bound, this gives

$$\Pr[\text{Bad}_{\text{pH}}] \leq \frac{(q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))^2}{2^{nl+1}}.$$

For  $\text{Bad}_{\text{pF}}$ , the tightest upper bound would be achieved by considering the number of nonces sampled *per file*, which is at most the number of calls to UPD1 for that file (across all users), plus one (for the nonce sampled when the file is first put). However, since the number of update calls per file is at most the total number of update queries  $\mathcal{Q}_{\text{UPD1}}(\mathcal{A})$  (across all files), we bound  $\Pr[\text{Bad}_{\text{pF}}]$

by  $C(2^{nl}, \mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + 1)$ , for simplicity. This gives

$$\Pr[\text{Bad}_{\text{pF}}] \leq \frac{(\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + 1)^2}{2^{nl+1}},$$

again by a birthday bound. Together, this gives

$$\Pr[\text{G}_5] - \Pr[\text{G}_6] \leq \frac{(\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + 1)^2 + (q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))^2}{2^{nl+1}}. \quad (14)$$

### Game G<sub>7</sub>: Replace encrypted file keys with random.

Next, we apply AEAD security to the file key encryptions of honest users. Briefly, the aim of this step is to replace each file key ciphertext of honest users with a randomly sampled string of the same length, such that the key used to encrypt the file is independent from the ciphertext sent to the server.

Formally, game G<sub>7</sub> differs from G<sub>6</sub> in line 3 of `putHeader` (Figure 10), which in G<sub>7</sub> is replaced by

G<sub>7</sub>: putHeader:

3  $c_{fk} \leftarrow \text{Enc}(st_C.k_{mk}, n, k_f, md)$ ;  $c_{fk} \leftarrow \$ \{0, 1\}^{|c_{fk}|}$ ;  $T_{k_f}[i_u, n, c_{fk}, md] \leftarrow k_f$

if called from oracle PUT1 or ACPT1 on input a client index  $i_c$  s.t.  $T_u[i_c] \notin S_{\text{comp}}$ . Note that `putFile` remains unchanged. In particular, the file is still encrypted with the file key  $k_f$  given as input to the subroutine. To ensure that file decryption works, game G<sub>7</sub> also replaces line 17 of `getHeader` and `getFile` (Figure 10) with

G<sub>7</sub>: getFile, getHeader:

16  $k_f \leftarrow T_{k_f}[i_u, n, c_{fk}, md]$

if called from the oracles running procedures  $\text{CSS}_{\text{upd}}$ ,  $\text{CSS}_{\text{get}}$  and  $\text{CSS}_{\text{share}}$  with a client index  $i_c$  s.t.  $T_u[i_c] \notin S_{\text{comp}}$ . Note that this means that  $k_f$  is set to  $\perp$  if the table entry has not been initialized.

We construct an adversary  $\mathcal{B}_{\text{aead}}^2$  against the IND\$ security of AEAD such that

$$\Pr[\text{G}_6] - \Pr[\text{G}_7] \leq \text{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^2). \quad (15)$$

Adversary  $\mathcal{B}_{\text{aead}}^2$  acts like the challenger in game G<sub>6</sub>, with the following exceptions on operations for users  $i_u \notin S_{\text{comp}}$ . First, instead of sampling  $k_{mk}$  at random in `AREG1`, adversary  $\mathcal{B}_{\text{aead}}^2$  queries oracle `NEW` and stores the index of the new key in  $T_{\text{AEAD}}[(i_u, \gamma), n_{mk}, c'_{mk}, (aid, \text{"mk"})]$ , where  $\gamma$ ,  $n_{mk}$ ,  $c'_{mk}$  and  $aid$  are defined as in game G<sub>6</sub>. This gives  $\mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^2) \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A})$ . Second, when  $\mathcal{A}$  runs procedure  $\text{CSS}_{\text{auth}}$  via `AUTH1` and `STEP`, adversary  $\mathcal{B}_{\text{aead}}^2$  sets  $st_C.k_{mk} \leftarrow T_{\text{AEAD}}[(i_u, \gamma), n_{mk}, c'_{mk}, (aid, \text{"mk"})]$ . (If the table entry is  $\perp$ ,  $\mathcal{B}_{\text{aead}}^2$  will respond with  $\perp$  to all future queries within the session which uses  $st_C.k_{mk}$ .)

Next, when  $\mathcal{A}$  queries `PUT1` or `ACPT1`, adversary  $\mathcal{B}_{\text{aead}}^2$  queries oracle  $\text{ENC}(st_C.k_{mk}, n, k_f, md)$  to encrypt the newly sampled  $k_f$  and uses the resulting ciphertext as  $c_{fk}$ . Thanks to how nonce sampling in `putHeader` is handled in G<sub>6</sub>, these queries from  $\mathcal{B}_{\text{aead}}^2$  are guaranteed to be nonce-respecting. Finally, instead of decrypting  $c_{fk}$  when simulating oracles which call `getFile` and `getHeader`,  $\mathcal{B}_{\text{aead}}^2$  queries oracle `DEC` in its own game with key index  $st_C.k_{mk}$  on the inputs provided by  $\mathcal{A}$ . This way,  $\mathcal{B}_{\text{aead}}^2$  simulates game G<sub>6</sub> and G<sub>7</sub> for  $\mathcal{A}$  when the hidden bit in the IND\$ game is 1 and 0, respectively. It makes at most  $\mathcal{Q}_{\text{PUT1}}(\mathcal{A}) + \mathcal{Q}_{\text{ACPT1}}(\mathcal{A})$  queries to oracle `ENC` and at most  $\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + \mathcal{Q}_{\text{SHARE1}}(\mathcal{A}) + \mathcal{Q}_{\text{GET1}}(\mathcal{A})$  to oracle `DEC`.

Let  $b$  denote the bit sampled by  $\mathcal{B}_{\text{aead}}^2$  at the start of the simulation and  $b_{\mathcal{B}}$  the bit sampled by the challenger in  $\mathbf{G}_{\text{AEAD}}^{\text{IND\$}}$ . When  $\mathcal{A}$  halts and returns  $b'$ ,  $\mathcal{B}_{\text{aead}}^1$  returns 1 if  $b' = b$ , otherwise it returns 0. This gives Equation (15).

**File encryption: Bounding game  $G_7$ .** The final reduction in the proof is concerned with the encryption of challenge files in  $G_7$ . We construct an adversary  $\mathcal{B}_{\text{aead}}^3$  against the IND-CCA security of AEAD such that

$$2 \cdot \Pr[G_7] - 1 \leq \mathbf{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^3). \quad (16)$$

Adversary  $\mathcal{B}_{\text{aead}}^3$  acts as the challenger in  $G_7$ , with some exceptions. First,  $\mathcal{B}_{\text{aead}}^3$  does not sample a bit  $b$ . Second,  $\mathcal{B}_{\text{aead}}^3$  will use the oracles in its own game to handle some file operations, as follows. Let  $i_u$  be the user index of the client for which  $\mathcal{A}$  runs a protocol through its oracle queries. We consider three cases for operations on a file with file ID  $fid$ :

- (a)  $fid \notin S_{\text{chall}}$ : In this case,  $\mathcal{B}_{\text{aead}}^3$  will sample the file key itself when the file is first put, and perform subsequent file encryptions and decryptions exactly as the challenger in  $G_7$ .
- (b)  $(fid \in S_{\text{chall}}) \wedge (i_u \notin S_{\text{comp}} \cup \{\iota_{\text{mal}}\})$ : In this case,  $\mathcal{B}_{\text{aead}}^3$  will query oracle NEW to request a new file key when the file is first put, and then use oracles ENC and DEC in game  $\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA}}$  to encrypt and decrypt, respectively.
- (c)  $(fid \in S_{\text{chall}}) \wedge (i_u \in S_{\text{comp}} \cup \{\iota_{\text{mal}}\})$ : If adversary  $\mathcal{A}$  makes a query that falls into this case,  $\mathcal{B}_{\text{aead}}^3$  will either return  $\perp$  to  $\mathcal{A}$  (if that is what  $G_7$  would do), or halt and return 0 (if it cannot continue to simulate).

Next, we will go through the implications of this strategy, protocol by protocol, focusing on cases (b) and (c), since that is where  $\mathcal{B}_{\text{aead}}^3$  differs from the challenger in  $G_7$ . Hence, the following applies only when  $fid \in S_{\text{chall}}$ .

Challenge put. In case (b), adversary  $\mathcal{B}_{\text{aead}}^3$  generates a new key via oracle NEW in game  $\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA}}$ , and stores the index  $i$  in  $T_{k_f}[i_u, n_{fk}, c_{fk}, md]$ . Here  $c_{fk}$  is a random string of appropriate length (as per game  $G_7$ ),  $n_{fk}$  is the file key nonce sampled in `putHeader` and  $md = (st_C.aid, fid, \text{"fk"})$  includes the file ID  $fid$ . The file ciphertext is generated as  $c_f \leftarrow \text{ENC}(i, n, f_0, f_1, (fid, \text{"file"}))$ , where  $n$  is the nonce sampled in `putFile`. Thanks to nonces being sampled without collisions (see  $G_6$ ), adversary  $\mathcal{B}_{\text{aead}}^3$  is guaranteed to be nonce-respecting. This correctly simulates game  $G_7$ .

In case (c), adversary  $\mathcal{B}_{\text{aead}}^3$  halts and returns 0. We note that in this case,  $i_u$  is added to  $S_{\text{challOwners}}$  on line 29 of Figure 6. Since  $i_u \in S_{\text{comp}} \cup \{\iota_{\text{mal}}\}$  in this case, `trivial` (defined on line 5 of Figure 5) will evaluate to `true`.

Challenge update. In case (b),  $i \leftarrow T_{k_f}[i_u, n_{fk}, c_{fk}, md]$ , retrieved by  $\mathcal{B}_{\text{aead}}^3$  in `getHeader`, either contains the index  $i$  of a key registered with oracle NEW (if the file has been put or its sharing accepted by user  $i_u$  before, and  $\mathcal{A}$  inputs the honest  $c_{fk}$  and  $n_{fk}$ ) or is  $\perp$  (if  $i_u$  is not an owner of this file or  $\mathcal{A}$  does not provide honest inputs). In either case,  $\mathcal{B}_{\text{aead}}^3$  attempts to encrypt the challenge files with  $c_f \leftarrow \text{ENC}(i, n, f_0, f_1, (fid, \text{"file"}))$  (returning  $\perp$  if  $i = \perp$ ). This correctly simulates game  $G_7$ .

In case (c),  $\mathcal{B}_{\text{aead}}^3$  returns  $\perp$  in place of  $c_f$  if  $T_{k_f}[i_u, n_{fk}, c_{fk}, md] = \perp$ . (This occurs if  $i_u$  is not an owner of  $fid$  from a prior put or accept operation, or if  $\mathcal{A}$  does not provide honest inputs.) This correctly simulates game  $G_7$ . Otherwise, if the table entry exists,  $\mathcal{B}_{\text{aead}}^3$  halts and returns 0. We note that in this case,  $i_u$  was either added to  $T_f[fid].U$  when it first put the file in oracle CHALL and is therefore added to  $S_{\text{challOwners}}$  on line 29, or, if the file was shared with  $i_u$  through the game, they

were added to  $S_{\text{challOwners}}$  as a recipient on line 19, both in Figure 6. In the case that  $i_u$  accepts the file as a share directly from the adversary, the adversary  $\iota_{\text{mal}}$  is added to  $S_{\text{challOwners}}$  on line 25 of Figure 6. In either case,  $\text{trivial} = \text{true}$ .

We note that in case (a),  $\mathcal{B}_{\text{aead}}^3$  will act as the challenger in the game and return  $\perp$  if  $\text{UPD1}$  is called through oracle  $\text{CHALL}$  on  $\text{fid} \notin S_{\text{chall}}$ . The reason that such queries are silenced is that  $\mathcal{B}_{\text{aead}}^3$  will have sampled the file key itself when the file was first put, and hence could not embed a challenge to its own game in this case. (That is, allowing this would encode a form of adaptive security, which is not present in the selective game, as it would lead to a commitment issue.)

Get. In case (b),  $\mathcal{B}_{\text{aead}}^3$  lets  $i \leftarrow T_{k_f}[i_u, n_{fk}, c_{fk}, md]$  and attempts to decrypt the file ciphertext with  $f \leftarrow \text{DEC}(i, n_f, c_f, (\text{fid}, \text{"file"}))$ . For the same reasons as in case (b) of Update, this correctly simulates game  $G_7$ .

In case (c),  $\mathcal{B}_{\text{aead}}^3$  returns  $\perp$  in place of the file. Thanks to line 20 of Figure 5, this correctly simulates game  $G_7$ .

Share. In case (b), the file is being shared by an honest user.  $\mathcal{B}_{\text{aead}}^3$  then acts differently depending on the identity of the receiver: If the receiver is also honest ( $i_r \notin S_{\text{comp}} \cup \{\iota_{\text{mal}}\}$ ),  $\mathcal{B}_{\text{aead}}^3$  lets  $i \leftarrow T_{k_f}[i_u, n_{fk}, c_{fk}, md]$  and sends  $i$  over the OOB channel. For the same reasons as in case (b) of Update, this correctly simulates game  $G_7$ . If instead  $i_r \in S_{\text{comp}} \cup \{\iota_{\text{mal}}\}$ , then  $\mathcal{B}_{\text{aead}}^3$  halts and returns 0. We note that in this case,  $i_r$  is added to  $S_{\text{challOwners}}$  on line 19, Figure 6. Hence  $\text{trivial}$  is true in this case. Note that as a consequence of this strategy, the OOB channel  $T_{\text{oob}}[(i_r, \cdot)]$  of a user  $i_r \in S_{\text{comp}}$  only ever contains file keys sampled by  $\mathcal{B}_{\text{aead}}^3$  itself (for non-challenge files).

In case (c),  $\mathcal{B}_{\text{aead}}^3$  halts and returns 0. We note that in this case,  $i_u$  was added to  $S_{\text{challOwners}}$  on line 19 of Figure 6, causing  $\text{trivial}$  to be true.

Accept. Note that we call the share-receiving user  $i_r$  here rather than  $i_u$ , to stay consistent with the game.

In case (b), if  $\text{in}_C.\text{oob} = \perp$ ,  $\mathcal{B}_{\text{aead}}^3$  samples a file key nonce  $n_{fk}$  and ciphertext  $c_{fk}$  independently at random, as would the challenger in  $G_7$ . It then lets  $T_{k_f}[i_u, n_{fk}, c_{fk}, md] \leftarrow T_{\text{oob}}[(i_r, \text{fid})]$ . This correctly simulates game  $G_7$ . If  $\text{in}_C.\text{oob} \neq \perp$ , then  $\mathcal{B}_{\text{aead}}^3$  halts and returns 0. We note that in this case,  $\iota_{\text{mal}}$  was added to  $S_{\text{challOwners}}$  on line 25 of Figure 6, causing  $\text{trivial}$  to be true.

In case (c),  $\mathcal{B}_{\text{aead}}^3$  checks if both  $\text{in}_C.\text{oob}$  and  $T_{\text{oob}}[(i_r, \text{fid})]$  are  $\perp$ . If so, it returns  $\perp$ . Otherwise, it halts and returns 0. We note that in this case,  $\iota_{\text{mal}}$  is either added to  $S_{\text{challOwners}}$  on line 25, Figure 6 (if  $\text{in}_C.\text{oob} \neq \perp$ ), or  $i_r$  was already added on line 19 (when  $T_{\text{oob}}[(i_r, \text{fid})]$  was initialized in  $\text{SHARE1}$ ). Since  $i_r \in S_{\text{comp}} \cup \{\iota_{\text{mal}}\}$  in this case, this implies that  $\text{trivial}$  is true.

If adversary  $\mathcal{A}$  makes any other query not covered above that sets  $\text{trivial}$  to true, adversary  $\mathcal{B}_{\text{aead}}^3$  halts and returns 0. Otherwise, adversary  $\mathcal{B}_{\text{aead}}^3$  simulates the game as above until adversary  $\mathcal{A}$  halts and returns  $b'$ , in which case  $\mathcal{B}_{\text{aead}}^3$  does the same. This means that  $\mathcal{B}_{\text{aead}}^3$  either returns the same bit guess as  $\mathcal{A}$  (if  $\neg \text{trivial}$ ) or 0 (if  $\text{trivial}$ ). In terms of query counts, we have  $\mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^3) \leq \mathcal{Q}_{\text{CPUT1}}(\mathcal{A})$ ,  $\mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^3) \leq \mathcal{Q}_{\text{CPUT1}}(\mathcal{A}) + \mathcal{Q}_{\text{CUPD1}}(\mathcal{A})$ ,  $\mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^3) \leq \mathcal{Q}_{\text{GET1}}(\mathcal{A})$ .

Let  $G_7^b$  denote  $G_7$  with hidden bit  $b$ . With this strategy, adversary  $\mathcal{B}_{\text{aead}}^3$  perfectly simulates  $G_7^b$  for  $\mathcal{A}$  when playing game  $\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-}b}$ , unless it halts before  $\mathcal{A}$  does. In particular,  $\mathcal{B}_{\text{aead}}^3$  is able to simulate queries to oracle  $\text{COMP}$  for  $i_u \in S_{\text{comp}}$ ; it samples  $\text{pw}$  and  $k_{mk}$  itself (like the challenger in  $G_7$ , thanks to prior game hops) and the OOB channel only contains file keys that  $\mathcal{B}_{\text{aead}}^3$  sampled itself, as noted in case (b) of Share.

Let  $T_b$  be the event that  $\text{trivial}$  is set to true by  $\mathcal{A}$  when adversary  $\mathcal{B}_{\text{aead}}^3$  plays game  $\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-}b}$ .



By applying the law of total probability to the IND-CCA advantage definition, we have

$$\begin{aligned}\mathbf{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^3) &= \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-1}}] - \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-0}}] \\ &= \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-1}} \mid \mathsf{T}_1] \cdot \Pr[\mathsf{T}_1] - \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-0}} \mid \mathsf{T}_0] \cdot \Pr[\mathsf{T}_0] \\ &\quad + \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-1}} \mid \neg\mathsf{T}_1] \cdot \Pr[\neg\mathsf{T}_1] - \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-0}} \mid \neg\mathsf{T}_0] \cdot \Pr[\neg\mathsf{T}_0].\end{aligned}$$

Furthermore, since  $\mathcal{B}_{\text{aead}}^3$  returns 0 if  $\text{trivial} = \text{true}$ ,

$$\Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-1}} \mid \mathsf{T}_1] \cdot \Pr[\mathsf{T}_1] - \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-0}} \mid \mathsf{T}_0] \cdot \Pr[\mathsf{T}_0] = 0.$$

Next, by the definition of conditional probability, and since  $\mathcal{B}_{\text{aead}}^3$  returns the same bit as  $\mathcal{A}$  in case  $\neg\text{trivial}$ ,

$$\begin{aligned}\Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-1}} \mid \neg\mathsf{T}_1] \cdot \Pr[\neg\mathsf{T}_1] - \Pr[\mathbf{G}_{\text{AEAD}}^{\text{IND-CCA-0}} \mid \neg\mathsf{T}_0] \cdot \Pr[\neg\mathsf{T}_0] \\ = \Pr[\mathcal{A} \Rightarrow 1 \text{ in } \mathsf{G}_7^1 \wedge \neg\mathsf{T}_1] - \Pr[\mathcal{A} \Rightarrow 1 \text{ in } \mathsf{G}_7^0 \wedge \neg\mathsf{T}_0] \\ = 2 \cdot \Pr[\mathsf{G}_7] - 1,\end{aligned}$$

where the last equality follows from standard advantage rewriting, conditioning on the value of  $b$ .

**Conclusion.** Combining Equations (8)–(16) gives

$$\begin{aligned}\mathbf{Adv}_{\text{CSS},n,\mathcal{PW}_n}^{\text{C2CConfS}}(\mathcal{A}) &\leq 2 \cdot \left( \mathbf{Adv}_{\mathcal{PW}}^{\text{PG}}(\mathcal{B}_{\text{pg}}) + \mathbf{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B}_{\text{prf}}) \right. \\ &\quad + \frac{(\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + 1)^2 + (q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))^2}{2^{nl+1}} \\ &\quad \left. + \mathbf{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^1) + \mathbf{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^2) \right) + \mathbf{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^3). \quad (17)\end{aligned}$$

We merge  $\mathcal{B}_{\text{aead}}^1$  and  $\mathcal{B}_{\text{aead}}^2$  into  $\mathcal{B}_{\text{aead}}^{12}$ , as follows.  $\mathcal{B}_{\text{aead}}^{12}$  samples a bit  $\beta \leftarrow_{\$} \{0, 1\}$  and runs  $\mathcal{B}_{\text{aead}}^{\beta+1}$ , forwarding any oracle queries to its own oracles and returning whatever  $\mathcal{B}_{\text{aead}}^{\beta+1}$  returns. This gives

$$\mathbf{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^{12}) = \frac{1}{2} \cdot \mathbf{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^1) + \frac{1}{2} \cdot \mathbf{Adv}_{\text{AEAD}}^{\text{IND-CCA}}(\mathcal{B}_{\text{aead}}^2), \quad (18)$$

with query counts  $\mathcal{Q}_{\text{ORC}}(\mathcal{B}_{\text{aead}}^{12}) \leq \max(\mathcal{Q}_{\text{ORC}}(\mathcal{B}_{\text{aead}}^1) + \mathcal{Q}_{\text{ORC}}(\mathcal{B}_{\text{aead}}^2))$  for  $\text{ORC} \in \{\text{NEW}, \text{ENC}, \text{DEC}\}$ . Combining Equations (17) and (18) gives the bound claimed in Theorem 6.1 with the following query counts in summary:

$$\begin{aligned}\mathcal{Q}_{\text{TEST}}(\mathcal{B}_{\text{pg}}) &\leq \mathcal{Q}_{\text{RO1}}(\mathcal{A}) + \mathcal{Q}_{\text{RO2}}(\mathcal{A}), \\ \mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{prf}}) &\leq \mathcal{Q}_{\text{FN}}(\mathcal{B}_{\text{prf}}) \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) + \mathcal{Q}_{\text{AUTH1}}(\mathcal{A}) \\ \begin{cases} \mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^{12}) &\leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) + \mathcal{Q}_{\text{AUTH1}}(\mathcal{A}), \\ \mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^{12}) &\leq \max(\mathcal{Q}_{\text{AREG1}}(\mathcal{A}), \mathcal{Q}_{\text{PUT1}}(\mathcal{A}) + \mathcal{Q}_{\text{ACPT1}}(\mathcal{A})), \\ \mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^{12}) &\leq \max(\mathcal{Q}_{\text{AUTH1}}(\mathcal{A}), \mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + \mathcal{Q}_{\text{SHARE1}}(\mathcal{A}) + \mathcal{Q}_{\text{GET1}}(\mathcal{A})), \end{cases} \\ \begin{cases} \mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^3) &\leq \mathcal{Q}_{\text{CPUT1}}(\mathcal{A}), \\ \mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^3) &\leq \mathcal{Q}_{\text{CPUT1}}(\mathcal{A}) + \mathcal{Q}_{\text{CUPD1}}(\mathcal{A}), \text{ and} \\ \mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^3) &\leq \mathcal{Q}_{\text{GET1}}(\mathcal{A}). \end{cases}\end{aligned}$$

## C.2 Proof of Theorem 6.2

This section provides the details of the last steps of the proof of the selective client-to-client integrity of CSS. The proof is analogous to the proof of Theorem 6.1 up until and including  $G_6$ , with the difference that games  $G_0$ – $G_6$  provide the oracles named in game  $\mathbf{G}_{\text{CSS},n,\mathcal{PW}_n}^{\text{C2CIntS}}$ , and also match the finalize procedure thereof. That is, the games run (the stateful) adversary  $\mathcal{A}$  with no oracle access to obtain  $S_{\text{comp}}$ , and then run  $\mathcal{A}$  again with access to all oracles until the adversary halts. They return win; a flag set to true (on line 19 of Figure 5) if  $\mathcal{A}$  makes a client successfully retrieve a file which has not been added via the PUT oracle, or which is only owned by honest users, but the file content does not match what is expected. Throughout, we let  $\Pr[G_i]$  denote the probability that  $G_i(\mathcal{A})$  returns true, and note that, by definition,

$$\mathbf{Adv}_{\text{CSS},n,\mathcal{PW}_n}^{\text{C2CIntS}}(\mathcal{A}) = \Pr[G_0]. \quad (19)$$

Furthermore, with only slight adaptations to the reductions provided in the proof of Theorem 6.1, we have

$$\Pr[G_0] - \Pr[G_3] \leq \mathbf{Adv}_{n,\mathcal{PW}_n}^{\text{PG}}(\mathcal{B}_{\text{pg}}), \quad (20)$$

$$\Pr[G_3] - \Pr[G_4] \leq \mathbf{Adv}_{\text{F}}^{\text{PRF}}(\mathcal{B}_{\text{prf}}), \quad (21)$$

$$\Pr[G_4] - \Pr[G_5] \leq \mathbf{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^1), \text{ and} \quad (22)$$

$$\Pr[G_5] - \Pr[G_6] \leq \frac{(Q_{\text{UPD1}}(\mathcal{A}) + 1)^2 + (q_{\text{PUT1}}(\mathcal{A}) + q_{\text{ACPT1}}(\mathcal{A}))^2}{2^{nl+1}}. \quad (23)$$

Adversary  $\mathcal{B}_{\text{pg}}$  works exactly as in the proof of Theorem 6.1. Adversaries  $\mathcal{B}_{\text{prf}}$  and  $\mathcal{B}_{\text{aead}}^1$  no longer sample a bit at the start of the game; instead, they run  $\mathcal{A}$  and return 1 in their respective games if  $\mathcal{A}$  sets win to true in the simulated game. This gives the bounds above.

We now provide the final steps in the proof of Theorem 6.2. Recall that in game  $G_6$  (as per the proof of Theorem 6.1), the OPRF client secret  $rw$ , the key encryption key  $k_{\text{kek}}$  derived therefrom and the registered master key  $k_{\text{mk}}$  of all honest clients are uniformly random and independent from anything that the server sees. The master key recovered by honest clients during authentication and used in the subsequent sessions is additionally ensured to be equal to their registered master key. That is, malicious inputs from the server during authentication of an honest user will result in the client failing to derive the master key. Furthermore, in game  $G_6$  all nonces sampled during file key and file encryption are guaranteed to be unique per encryption key.

**Game  $G_7$ : Rule out completely forged file headers.** In this step, we rule out the possibility of win being set to true for a file ID which has not been uploaded via oracle PUT1. To this end, we introduce a **bad** flag in  $G_6$ , which is set to true in procedure Check if either a get or share operation is successful, when run for an honest user, on a file ID which has no tracked content and no or only honest owners. That is, we cover both the case when the file ID is not tracked by the game at all, as well as when the file ID has not been uploaded through oracle PUT1, but shared or received, in which case the file has owners but no content.

In the game, we modify lines 18–20 and add lines 23–24 to subroutine Check in Figure 5, as shown below.

```

 $G_6, \boxed{G_7}$ :  $\text{Check}(\Pi, i_p, out_C, m')$ :
17 if  $\Pi = \Pi_{\text{get}}$ :
18   if  $T_u[i_c] \in S_{\text{hon}} \setminus S_{\text{comp}} \wedge out_C.dec \wedge T_f[fid].U \subseteq S_{\text{hon}} \setminus S_{\text{comp}}$ :
19     if  $T_f[fid].F = \emptyset$ :  $\text{bad} \leftarrow \text{true}$ ; return  $\perp$ 
20      $\text{win} \leftarrow out_C.f \notin T_f[fid].F$ 
21   return  $(out_C.f, m')$ 
22 if  $(\Pi = \Pi_{\text{share}} \text{ and } out_C.dec)$ :  $T_{\text{oob}}[(i_r, fid)] \leftarrow out_C.oob$ 
23   if  $T_u[i_c] \in S_{\text{hon}} \setminus S_{\text{comp}} \wedge T_f[fid].F = \emptyset \wedge T_f[fid].U \subseteq S_{\text{hon}} \setminus S_{\text{comp}}$ :
24      $\text{bad} \leftarrow \text{true}$ ;  $T_{\text{oob}}[(i_r, fid)] \leftarrow \perp$ ;  $T_f[fid].U \leftarrow T_f[fid].U \setminus \{T_u[i_c]\}$ ; return  $\perp$ 

```

We additionally modify the return statement on line 11 of Figure 5 to

```

 $G_6, \boxed{G_7}$ :
11 if  $\text{bad}$ :  $\text{win} \leftarrow \text{false}$ 
12 return  $\text{win}$ 

```

where, in both of the above, the boxed code is not present in game  $G_6$ . We note that neither of these changes affect the winning condition, hence the modified  $G_6$  is equivalent to the original and any adversary has the same probability of winning in both. We now create game  $G_7$ , which is identical to  $G_6$ , except that it contains the boxed code shown on lines 19, 24 and 11 above.

Let  $\text{Bad}$  denote the event that  $G_6$  sets the  $\text{bad}$  flag.  $G_6$  and  $G_7$  are identical-until- $\text{Bad}$ , so by the fundamental lemma of game playing [11],  $\Pr[G_6] - \Pr[G_7] \leq \Pr[\text{Bad}]$ .

We construct an adversary  $\mathcal{B}_{\text{aead}}^2$  against the INT-CTXT security of AEAD such that

$$\Pr[\text{Bad}] \leq \text{Adv}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{B}_{\text{aead}}^2). \quad (24)$$

Adversary  $\mathcal{B}_{\text{aead}}^2$  acts as the challenger in game  $G_6$ , with the following differences on operations for users  $i_u \notin S_{\text{comp}}$ . First, instead of sampling  $k_{mk}$  at random in  $\text{AREG1}$ , adversary  $\mathcal{B}_{\text{aead}}^2$  queries oracle  $\text{NEW}$  and stores the index of the new key in a table indexed by the relevant messages in the transcript of the registration protocol. (Cf. adversary  $\mathcal{B}_{\text{aead}}^2$  in the proof of Theorem 6.1, hop from  $G_6$  to  $G_7$ .) This gives  $\mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^2) \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A})$ . Second, when  $\mathcal{A}$  runs procedure  $\text{CSS}_{\text{auth}}$  via  $\text{AUTH1}$  and  $\text{STEP}$ , adversary  $\mathcal{B}_{\text{aead}}^2$  sets  $st_C.k_{mk}$  to the indexed stored in the table. (If the table entry is  $\perp$  due to malicious server messages during authentication,  $\mathcal{B}_{\text{aead}}^2$  will respond with  $\perp$  to all future queries within the session which uses  $st_C.k_{mk}$ .)

Next, when adversary  $\mathcal{A}$  queries oracles  $\text{PUT1}$  or  $\text{ACPT1}$ , adversary  $\mathcal{B}_{\text{aead}}^2$  queries oracle  $\text{ENC}(st_C.k_{mk}, n, k_f, md)$  to encrypt the newly sampled  $k_f$  or the key from  $in_C.oob$  or  $T_{\text{oob}}$ , respectively. It uses the resulting ciphertext as  $c_{fk}$ . Thanks to how nonce sampling in  $\text{putHeader}$  is handled in  $G_6$ , these queries from  $\mathcal{B}_{\text{aead}}^2$  are guaranteed to be nonce-respecting. Finally, instead of decrypting  $c_{fk}$  when simulating oracles which call  $\text{getFile}$  and  $\text{getHeader}$ ,  $\mathcal{B}_{\text{aead}}^2$  queries oracle  $\text{DEC}$  in its own game with key index  $st_C.k_{mk}$  on the inputs provided by  $\mathcal{A}$ .

Adversary  $\mathcal{B}_{\text{aead}}^2$  handles queries to oracle  $\text{COMP}$  as would the challenger in game  $G_6$ , since – due to the fact that  $\mathcal{A}$  announces all of its compromise queries in advance – adversary  $\mathcal{B}_{\text{aead}}^2$  knows all of the values to be returned in response to such a query. (It only does not know the master key of honest users, which are never revealed to the adversary in the selective security game.) This way,  $\mathcal{B}_{\text{aead}}^2$  simulates game  $G_6$  for  $\mathcal{A}$  and makes at most  $\mathcal{Q}_{\text{PUT1}}(\mathcal{A}) + \mathcal{Q}_{\text{ACPT1}}(\mathcal{A})$  queries to oracle  $\text{ENC}$  and at most  $\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + \mathcal{Q}_{\text{SHARE1}}(\mathcal{A}) + \mathcal{Q}_{\text{GET1}}(\mathcal{A})$  to oracle  $\text{DEC}$ .

We claim that  $\Pr[\text{Bad}] \leq \text{Adv}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{B}_{\text{aead}}^2)$ . To see this, observe that  $\text{bad}$  is set to  $\text{true}$  in  $G_6$  if the adversary has successfully run the  $\text{get}$  or  $\text{share}$  protocols on a file ID  $fid$  such that

$$(T_u[i_c] \in S_{\text{hon}} \setminus S_{\text{comp}}) \wedge (out_C.dec) \wedge (T_f[fid].F = \emptyset) \wedge (T_f[fid].U \subseteq S_{\text{hon}} \setminus S_{\text{comp}}).$$

In order, these conditions correspond to (1) the get or share operation is performed by an honest user, (2) the protocol concluded successfully on the client, (3) the file content is not tracked by the game, and (4) the file is only owned by honest users. We now claim that together, these conditions correspond to a file header forgery for the file ID  $fid$  which sets **bad**. To see this, note that point (3) implies that  $fid$  has not been queried to oracle PUT1 and consider the following two cases for point (4):

1.  $T_f[fid].U = \emptyset$ . (The set of file owners is empty.) In this case,  $fid$  has not been queried to oracle SHARE1 or ACPT1, and, hence, **putHeader** has not been performed for  $fid$ . However, a successful get or share operation means that protocol **getFile** or **getHeader**, respectively, has been run successfully for an honest user. This in turn means that adversary  $\mathcal{B}_{aead}^2$  queried oracle  $DEC(st_C.k_{mk}, n_{fk}, c_{fk}, md)$  on the master key index of the honest client, the file header  $(n_{fk}, c_{fk})$  supplied by adversary  $\mathcal{A}$  and the metadata  $md = (st_C.aid, fid, "fk")$  and that the decryption succeeded. Since no file header has been put for  $fid$ , and hence not queried to oracle ENC by adversary  $\mathcal{B}_{aead}^2$ , this is a valid forgery in the INT-CTXT game of AEAD. (Note that  $fid$  is part of the metadata  $md$ . Hence,  $fid$  being fresh implies that the decryption query does not count as a trivial attack on line 29 of game  $\mathbf{G}_{AEAD}^{INT-CTXT}$  in Figure 20.)
2.  $\emptyset \neq T_f[fid].U \subseteq S_{hon} \setminus S_{comp}$ . (The set of file owners is non-empty, but only contains honest users.) In this case, the file has been queried to oracle SHARE1 and/or ACPT1 (where the share query can be the one that sets **bad**) with only honest users as sharer and receiver, respectively. We identify the following sub-cases:
  - (a)  $fid$  has not been queried to oracle ACPT1. Then it is again the case that **putHeader** has not been performed for  $fid$ , as **putHeader** is only run in  $CSS_{put}$  and  $CSS_{accept}$ . Hence, by the same reasoning as in case 1, this implies that adversary  $\mathcal{B}_{aead}^2$  successfully forged a file header.
  - (b)  $fid$  has been queried to oracle ACPT1. Then, by the precondition for oracle ACPT1, on line 21 of Figure 6, either  $in_C.oob \neq \perp$ , or  $T_{oob}[i_r, fid] \neq \perp$ . If  $in_C.oob \neq \perp$ , then  $\iota_{mal}$  is added to the set of owners, and hence the file is no longer only owned by honest users, ruling out this possibility. If  $T_{oob}[i_r, fid] \neq \perp$ , then the OOB table entry must have been previously initialized by some user via a successful query to oracle SHARE1. This user must be honest (otherwise the file would be owned by the compromised sharer), implying that there was a point in time when  $fid$  had *not* yet been queried to oracle ACPT1, but was queried to oracle SHARE1 with all of the conditions necessary for setting **bad** being true. Hence, we are back in case 2(a).

Therefore, **Bad** implies that adversary  $\mathcal{B}_{aead}^2$  set win to true in game  $\mathbf{G}_{AEAD}^{INT-CTXT}$ , yielding Equation (24).

**Game G<sub>8</sub>: Guess a forgery attempt.** We now observe that in game G<sub>7</sub>, the file ID  $fid$  of any successful forgery attempt must have appeared in a query to oracle PUT1. To see this, note that, thanks to the modified return statement in game G<sub>7</sub>, the condition for setting the win flag can be changed to

```

G7: Check( $\Pi, i_p, out_C, m'$ ):
:
17 if  $\Pi = \Pi_{get}$ :
18   win  $\leftarrow (T_u[i_c] \in S_{hon} \setminus S_{comp}) \wedge (out_C.dec) \wedge (T_f[fid].F \neq \emptyset) \wedge (out_C.f \notin T_f[fid].F) \wedge (T_f[fid].U \subseteq S_{hon} \setminus S_{comp})$ 

```

without changing the outcome of the game. That is, adversary  $\mathcal{A}$  can only win if it successfully gets a file  $fid$  through the client of an honest user, such that the file content is tracked (meaning the file has been put or updated at least once), the decrypted file content does not match any tracked version of the file, and the file is owned only by honest users. Now consider the possible ways in which the adversary can initialize  $T_f[fid].F$ , while ensuring that  $fid$  is only owned by honest users.

$T_f[fid].F$  is only set in oracles PUT1 and UPD1. In the latter, it is only updated if the calling user is already an owner of  $fid$  tracked in  $T_F[fid].U$ .  $T_F[fid].U$ , in turn, is only updated in oracles PUT1, SHARE1 and ACPT1. This leads to the following cases:

1.  $fid$  was first uploaded via oracle PUT1.
2.  $fid$  was updated via oracle UPD1 without first having been put. Then it must be the case that the updating user is already an owner of the file, and hence either:
  - (a) The updating user was added as an owner in a query to oracle SHARE1 while  $T_f[fid].F = \emptyset$ . If the sharing user is compromised, the owner set no longer only contains honest users, so this case is ruled out. If instead the sharing client is an honest user, then by line 24 of game  $G_7$  (as shown in the previous game hop above) the sharer is not added to  $T_F[fid].U$ . Hence this case is also ruled out.
  - (b) The updating user was added as an owner while receiving  $fid$  as  $i_r$  in oracle ACPT1. By the precondition of oracle ACPT1, it is either the case that  $in_C.oob \neq \perp$ , or  $T_{oob}[i_r, fid] \neq \perp$ . In the former case,  $\iota_{mal}$  is added to the set of owners, and hence the file is no longer only owned by honest users, ruling out this possibility. In the latter case,  $T_{oob}[i_r, fid]$  must have been previously initialized by some user via a query to oracle SHARE1. But this is again ruled out by the same reasoning as in case 2(a), since either the sharer was compromised, in which case the file is no longer only owned by honest users, or  $T_{oob}[i_r, fid]$  has been set to  $\perp$  on line 24 of game  $G_7$ .

Hence, we are left only with case 1, that  $fid$  was first uploaded via oracle PUT1. We now proceed to guess among the files uploaded via oracle PUT1 one that will constitute a successful forgery.

Let  $S_f$  denote the set of file IDs of all files uploaded to the server by an honest user during the course of the game. That is,  $S_f = \{fid \mid fid \text{ was queried to oracle PUT1 with } T_u[i_c] \in S_{hon} \setminus S_{comp}\}$ . Furthermore, let  $S_f^* \subseteq S_f$  be the subset of uploaded file IDs for which  $\mathcal{A}$  performed a successful forgery. That is,  $S_f^*$  contains the file IDs of every file for which  $win$  evaluated to **true** during the course of game  $G_7$ . (As noted above,  $win$  can only be set to **true** for file IDs in  $S_f$ .) Assume without loss of generality that  $S_f^* \neq \emptyset$ . (Otherwise  $\mathcal{A}$  has 0 advantage.) As per the theorem statement, let  $p$  be an upper bound on the number of queries to PUT1 by adversary  $\mathcal{A}$ . By definition,  $p \geq |S_f| \geq |S_f^*|$ .

In game  $G_8$ , the challenger begins by sampling  $j \in \{1, \dots, p\}$ . When, during the game,  $\mathcal{A}$  makes the  $j$ -th call to PUT1, the challenger remembers the file ID  $fid^*$  it uses. Furthermore, in the Check procedure, the win condition is modified to

$G_8$ : Check( $\Pi, i_p, out_C, m'$ ):  
 $\vdots$   
17 **if**  $\Pi = \Pi_{get}$ :  
18      $win \leftarrow (T_u[i_c] \in S_{hon} \setminus S_{comp}) \wedge (out_C.dec) \wedge (T_f[fid].F \neq \perp) \wedge (out_C.f \notin T_f[fid].F)$   
            $\wedge (T_f[fid].U \subseteq S_{hon} \setminus S_{comp}) \wedge (fid = fid^*)$

That is,  $fid = fid^*$  is added as a condition for  $win$  to be set to **true**. Hence, in game  $G_8$ , adversary  $\mathcal{A}$  wins only if it performs a successful forgery of  $fid^*$ , the file ID uploaded in its  $j$ -th PUT1 query.

Let  $E$  be the event that  $fid^* \in S_f^*$ . Then  $\Pr[E] \geq 1/p$ , since by definition  $p \geq |S_f^*| > 0$  and  $j$  is sampled uniformly at random among the  $p$  queries to PUT1. Furthermore,  $\Pr[G_8|E] = \Pr[G_7|E] = \Pr[G_7]$ , since the outcome of the games is equivalent if  $E$  occurs, and the outcome of  $G_7$  is independent of  $E$ . Hence

$$\Pr[G_7] \leq p \cdot \Pr[G_8]. \quad (25)$$

**Game  $G_9$ : Replace the encrypted file key of  $fid^*$  with random.** Next, we apply AEAD security to the file key encryption of file  $fid^*$ , the file uploaded through the  $j$ -th PUT1 query by an honest user. As in the hop to game  $G_7$  in the confidentiality proof, the aim of this step is to replace the file key ciphertext with a randomly sampled string of the same length, such that the key used to encrypt the file is independent from the ciphertext sent to the server.

We construct game  $G_9$ , which replaces line 3 of putHeader (Figure 10) in  $G_8$  with

$G_9$ : putHeader:

3  $c_{fk} \leftarrow \text{Enc}(st_C.k_{mk}, n, k_f, md); c_{fk} \leftarrow \$ \{0, 1\}^{|c_{fk}|}; T_{k_f}[i_u, n, c_{fk}, md] \leftarrow k_f$

when called for the  $j$ -th time in oracle PUT1 for an honest user (where  $j$  is sampled by the challenger at the start of the game as per game  $G_8$  of this proof), or subsequently when called through oracle ACPT1 for any honest user on input file ID  $fid^*$ , as recorded in the  $j$ -th put query. Note that putFile remains unchanged. In particular, the file is still encrypted with the file key  $k_f$  given as input to the subroutine. To ensure that file decryption works, game  $G_9$  also replaces line 17 of getFile and getHeader (Figure 10) with

$G_9$ : getFile, getHeader:

16  $k_f \leftarrow T_{k_f}[i_u, n, c_{fk}, md]$

if called for an honest user from the oracles running procedures  $\text{CSS}_{\text{upd}}$ ,  $\text{CSS}_{\text{get}}$  and  $\text{CSS}_{\text{share}}$  on input file ID  $fid^*$ . Note that this means that  $k_f$  is set to  $\perp$  if the table entry has not been initialized.

We construct an adversary  $\mathcal{B}_{\text{aead}}^3$  against the IND\$ security of AEAD such that

$$\Pr[G_8] - \Pr[G_9] \leq \text{Adv}_{\text{AEAD}}^{\text{IND\$}}(\mathcal{B}_{\text{aead}}^3). \quad (26)$$

Adversary  $\mathcal{B}_{\text{aead}}^3$  acts like the challenger in game  $G_8$ , with the following exceptions on operations for users  $i_u \notin S_{\text{comp}}$ . First, instead of sampling  $k_{mk}$  at random in AREG1, adversary  $\mathcal{B}_{\text{aead}}^3$  queries oracle NEW and stores the index of the new key in table  $T_{\text{AEAD}}[(i_u, \gamma), n_{mk}, c'_{mk}, (aid, \text{"mk"})]$ , indexed by the relevant messages in the transcript of the registration protocol. (Cf. adversary  $\mathcal{B}_{\text{aead}}^2$  in the proof of Theorem 6.1, hop from  $G_6$  to  $G_7$ .) This gives  $\mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^3) \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A})$ . Second, when  $\mathcal{A}$  runs procedure  $\text{CSS}_{\text{auth}}$  via AUTH1 and STEP, adversary  $\mathcal{B}_{\text{aead}}^3$  sets  $st_C.k_{mk} \leftarrow T_{\text{AEAD}}[(i_u, \gamma), n_{mk}, c'_{mk}, (aid, \text{"mk"})]$ . (If the table entry is  $\perp$ ,  $\mathcal{B}_{\text{aead}}^3$  will respond with  $\perp$  to all future queries within the session which uses  $st_C.k_{mk}$ .)

Next, when  $\mathcal{A}$  makes the  $j$ -th query to oracle PUT1 for an honest user, adversary  $\mathcal{B}_{\text{aead}}^3$  queries oracle  $\text{ENC}(st_C.k_{mk}, n, k_f, md)$  to encrypt the newly sampled  $k_f$  and uses the resulting ciphertext as  $c_{fk}$ . It records the file ID of the new file as  $fid^*$ . If adversary  $\mathcal{A}$  makes subsequent queries to oracle ACPT1 for an honest client on input file ID  $fid^*$ , adversary  $\mathcal{B}_{\text{aead}}^3$  again queries oracle  $\text{ENC}(st_C.k_{mk}, n, k_f, md)$  to encrypt the key from  $T_{\text{oob}}$  and uses the resulting ciphertext as  $c_{fk}$ . Thanks to how nonce sampling in putHeader is handled from the hop to  $G_6$ , the encryption queries from  $\mathcal{B}_{\text{aead}}^3$  are guaranteed to be nonce-respecting.

Finally, instead of decrypting  $c_{fk}$  when simulating oracles which call getFile and getHeader for an honest client on input file ID  $fid^*$ ,  $\mathcal{B}_{\text{aead}}^3$  queries oracle DEC in its own game with key index



$st_C.k_{mk}$  on the inputs provided by  $\mathcal{A}$ . This way,  $\mathcal{B}_{\text{aead}}^3$  simulates game  $G_8$  and  $G_9$  for  $\mathcal{A}$  when the hidden bit in the IND\$ game is 1 and 0, respectively. It makes at most  $1 + \mathcal{Q}_{\text{ACPT1}}(\mathcal{A})$  queries to oracle ENC and at most  $\mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + \mathcal{Q}_{\text{SHARE1}}(\mathcal{A}) + \mathcal{Q}_{\text{GET1}}(\mathcal{A})$  to oracle DEC.

When  $\mathcal{A}$  halts and returns  $b'$ ,  $\mathcal{B}_{\text{aead}}^3$  returns 1 if adversary  $\mathcal{A}$  managed to set `win` to `true` (as defined in game  $G_9$ ), otherwise it returns 0. This gives Equation (26).

**File integrity: Bounding game  $G_9$ .** In the final step of the proof, we construct an adversary  $\mathcal{B}_{\text{aead}}^4$  such that

$$\Pr[G_9] \leq \mathbf{Adv}_{\text{AEAD}}^{\text{INT-CTXT}}(\mathcal{B}_{\text{aead}}^4). \quad (27)$$

Adversary  $\mathcal{B}_{\text{aead}}^4$  acts as the challenger in game  $G_9$ , with some exceptions, which only apply if  $fid^*$  first appears through a query to PUT1 for an honest client (i.e.,  $T_u[i_c] \in S_{\text{hon}} \setminus S_{\text{comp}}$ ). The first exception is that when  $fid^*$  is first uploaded through oracle PUT1, adversary  $\mathcal{B}_{\text{aead}}^4$  does not sample the file key used to encrypt file  $fid^*$  itself, but instead queries oracle NEW to set up a new key in the integrity game. Then, adversary  $\mathcal{B}_{\text{aead}}^4$  uses oracle ENC to encrypt the file under the new key. This encryption step is repeated (for the new file content) in case of subsequent updates to  $fid^*$ . Second, when adversary  $\mathcal{A}$  performs the second step of protocol get for  $fid^*$ , adversary  $\mathcal{B}_{\text{aead}}^4$  forwards the inputs provided by  $\mathcal{A}$  to oracle DEC to decrypt the file.

If at some point adversary  $\mathcal{A}$  shares  $fid^*$  with a recipient  $i_r \in S_{\text{comp}}$ , and subsequently attempts to accept the share, then adversary  $\mathcal{B}_{\text{aead}}^4$  halts and aborts the simulation. Note that at this point, adversary  $\mathcal{A}$  can no longer set the win flag to true in game  $G_9$ , since  $fid^*$  is owned by a compromised user, per line 23 of Figure 6.

With this strategy,  $\mathcal{B}_{\text{aead}}^4$  simulates  $G_9$  (until  $\mathcal{A}$  performs an operation which prevents `win*` from being set to `true`). Furthermore, if  $\mathcal{A}$  sets `win*` to `true`, then by definition it must have provided inputs to protocol  $\text{CSS}_{\text{get}}$  such that the decryption of  $fid^*$  is successful and does not result in any of the plaintexts provided as input to the encryption oracle by  $\mathcal{B}_{\text{aead}}^4$  in simulations of put and update operations. By correctness of the AEAD scheme, the file content being new means that the combination of nonce, file ciphertext or associated data input to decryption must be new. Hence  $\mathcal{B}_{\text{aead}}^4$  has performed a successful forgery in the INT-CTXT game, proving Equation (27).

Adversary  $\mathcal{B}_{\text{aead}}^4$  has query counts  $\mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^4) \leq 1$ ,  $\mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^4) \leq 1 + \mathcal{Q}_{\text{UPD1}}(\mathcal{A})$  and  $\mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^4) \leq \mathcal{Q}_{\text{GET1}}(\mathcal{A})$ .

**Conclusion.** Combining Equations (19) through (27) yields the bound from Theorem 6.2 with the following query counts in summary:

$$\begin{aligned} \mathcal{Q}_{\text{TEST}}(\mathcal{B}_{\text{pg}}) &\leq \mathcal{Q}_{\text{RO1}}(\mathcal{A}) + \mathcal{Q}_{\text{RO2}}(\mathcal{A}) , \\ \mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{prf}}) &\leq \mathcal{Q}_{\text{FN}}(\mathcal{B}_{\text{prf}}) \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) + \mathcal{Q}_{\text{AUTH1}}(\mathcal{A}) , \\ \begin{cases} \mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^1) &\leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) + \mathcal{Q}_{\text{AUTH1}}(\mathcal{A}) , \\ \mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^1) &\leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) , \\ \mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^1) &\leq \mathcal{Q}_{\text{AUTH1}}(\mathcal{A}) , \end{cases} \\ \begin{cases} \mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^2) &\leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) , \\ \mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^2) &\leq \mathcal{Q}_{\text{PUT1}}(\mathcal{A}) + \mathcal{Q}_{\text{ACPT1}}(\mathcal{A}) , \\ \mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^2) &\leq \mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + \mathcal{Q}_{\text{SHARE1}}(\mathcal{A}) + \mathcal{Q}_{\text{GET1}}(\mathcal{A}) , \end{cases} \end{aligned}$$



$$\begin{cases}
\mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^3) & \leq \mathcal{Q}_{\text{AREG1}}(\mathcal{A}) , \\
\mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^3) & \leq 1 + \mathcal{Q}_{\text{ACPT1}}(\mathcal{A}) , \\
\mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^3) & \leq \mathcal{Q}_{\text{UPD1}}(\mathcal{A}) + \mathcal{Q}_{\text{SHARE1}}(\mathcal{A}) + \mathcal{Q}_{\text{GET1}}(\mathcal{A}) ,
\end{cases}$$

$$\begin{cases}
\mathcal{Q}_{\text{NEW}}(\mathcal{B}_{\text{aead}}^4) & \leq 1 , \\
\mathcal{Q}_{\text{ENC}}(\mathcal{B}_{\text{aead}}^4) & \leq 1 + \mathcal{Q}_{\text{UPD1}}(\mathcal{A}) , \text{ and} \\
\mathcal{Q}_{\text{DEC}}(\mathcal{B}_{\text{aead}}^4) & \leq \mathcal{Q}_{\text{GET1}}(\mathcal{A}) .
\end{cases}$$