

Modern Family: A Revocable Hybrid Encryption Scheme Based on Attribute-Based Encryption, Symmetric Searchable Encryption and SGX

Alexandros Bakas and Antonis Michalas

Tampere University,
Tampere, Finland

{alexandros.bakas,antonios.michalas}@tuni.fi

Abstract. Secure cloud storage is considered as one of the most important issues that both businesses and end-users take into account before moving their private data to the cloud. Lately, we have seen some interesting approaches that are based either on the promising concept of Symmetric Searchable Encryption (SSE) or on the well-studied field of Attribute-Based Encryption (ABE). In this paper, we propose a hybrid encryption scheme that combines both SSE and ABE by utilizing the advantages of both these techniques. In contrast to many approaches, we design a revocation mechanism that is completely separated from the ABE scheme and solely based on the functionality offered by SGX.

Keywords: Access Control · Attribute-Based Encryption · Cloud Security · Hybrid Encryption · Policies · Storage Protection · Symmetric Searchable Encryption

1 Introduction

Cloud computing plays a significant role in our daily routine. From casual internet users, to big corporations, the cloud has become an integral part of our lives. However, using services that are hosted and controlled by third parties raises several security and privacy concerns. For example, in [12] it is stated that there has been a 300% increase in Microsoft cloud-based user's account attacks over the past couple of years. However, when considering a cloud-based environment, cyber-attacks performed by remote adversaries is only a part of the problem. More precisely, when we design cloud services we also need to take into consideration cases where the actual cloud service provider (CSP) acts maliciously.

To overcome this, both academia and big industrial players have started looking on how to build cloud-based services that will utilize Symmetric Searchable Encryption (SSE) [7, 4]. In such a scheme, whenever a user wishes to access her files, she can search directly over the encrypted data for specific keywords. Unfortunately, revocation cannot be implemented efficiently since sharing an encrypted file implies sharing the encryption key. As a result, if a data owner wishes to revoke a user, then all files that are encrypted with the same key

must be decrypted and then re-encrypted under a fresh key. Another promising technique that fits cloud-based services is Attribute-Based Encryption (ABE). In ABE schemes, all files are encrypted under a master public key but in contrast to traditional public key encryption, the generated ciphertext is bounded by a policy. Each user has a distinct secret key which is associated with specific attributes. This way a user's secret key can decrypt a ciphertext if and only if the her attributes satisfy the policy bound to the ciphertext. However, using an asymmetric encryption scheme to store data is rather inefficient.

Contribution: We propose a hybrid encryption scheme that combines SSE and ABE in a way that reduces the problem of multi-user data sharing to that of a single-user. We use the ABE scheme as a sharing mechanism and not as a revocation one to achieve better efficiency. To deal with the problem of revocation, we utilize the functionality offered by SGX. Furthermore, this work extends the protocol presented in [11].

Organization: In Section 2, we present important works that have been published and address the problem of secure cloud storage, data sharing and revocation. In Section 3, we define our system model while in Section 4, we present the cryptographic tools needed for the construction of our scheme. In Section 5, we give a formal construction of our scheme which is followed by the security analysis in Section 6. Finally, Section 7 concludes the paper.

2 Related Work

In [13] authors present a revocable hybrid encryption scheme while at the same time a key-rotation mechanism is used to prevent key-scraping attacks. The authors use Optimal Asymmetric Encryption Padding (OAEP) as an All-or-Nothing-Transformation (AONT) [2] to prevent revoked users from accessing stored data. This is due to the fact that reversing OAEP, requires to the entire output. Thus, changing random bits, renders the reversion infeasible. Hence, to decrypt a file, the changed bits need to be stored. However, this implies that with each re-encryption, the size of the ciphertext grows. Thus, decrypting a file that has been re-encrypted multiple times is an expensive operation. Moreover, to achieve better efficiency, authors suggest that the AONT could be applied by the server. However, this implies the existence of a fully trusted server.

A promising idea is presented in [5], where the authors present a protocol based on functional encryption, with the main functionalities running in isolated environments. The decryption of a file, and the application of a function f on the decrypted file both occur in SGX enclaves. Moreover, all enclaves can attest to each other and exchange data over secure communication channels. In our construction, even though we use the same hardware principles, we build a hybrid encryption scheme by combining SSE and ABE.

In [9] authors present a revocable ciphertext-policy attribute-based encryption scheme. The revocation mechanism is offered by a revocation list that is attached to the resulted ciphertexts. To avoid maintaining long revocation lists, a policy through which users' keys expire after a certain period of time is enforced.

As a result, the revocation list only includes keys that have been revoked before the expiration date. Another Hybrid encryption scheme is presented in [6], in which authors propose a scheme based on SSE and ABE. In the proposed scheme, data owners encrypt their files using SSE, but the resulted indexes are encrypted under ABE. This way, users can locally generate search tokens based on their attributes, that are then sent to the cloud. However promising, their scheme is static and as a result can only have very limited applications in real-life scenarios. Moreover, authors do not provide a revocation mechanism – a problem of paramount importance in cloud-based services.

In our construction, we overcome these issues by designing an efficient revocation mechanism that is utilizing the SGX functionality and it is separated from the ABE scheme.

3 Architecture

In this section, we introduce the system model by explicitly describing the main entities that participate in our protocol as well as their capabilities. The system model of our work is built on top of the model presented in [10] and it is enhanced with some important additions.

Cloud Service Provider (CSP): We consider a cloud computing environment similar to the one described in [14, 15]. Moreover, the CSP must support SGX since core entities will be running in a trusted execution environment offered by SGX.

Master Authority (MS): MS is responsible for setting up all the necessary public parameters for the proper run of the involved protocols. MS is responsible for generating and distributing ABE keys to the registered users. Finally, MS is SGX-enabled and is running in an enclave called the Master Enclave.

Key Tray (KT): KT is a key storage that stores ciphertexts of the symmetric keys that have been generated by various users and are needed to decrypt data. Registered users can directly contact KT and request access to the stored ciphertexts. KT is also SGX-enabled and runs in an enclave called the KT Enclave.

Revocation Authority (REV): REV is responsible for maintaining a revocation list (rl) with the unique identifiers of the revoked users. Similar to MS and KT, REV is also SGX-enabled and is running in an enclave called the Revocation Enclave. Finally, for the security of the stored revocation list, it is important to mention that rl is generated by the enclave (i.e. in an isolated environment) and never leaves its perimeter. Therefore, there is no need to encrypt rl .

SGX: Below we provide a brief presentation of the main SGX functionalities needed for our construction. A more detailed description can be found in [5, 3]

Isolation: Enclaves are located in a hardware guarded area of memory and they compromise a total memory of 128MB (only 90MB can be used by software). Intel SGX is based on memory isolation built into the processor itself along with strong cryptography. The processor tracks which parts of memory belong to which enclave, and ensures that *only* enclaves can access their own memory.

Attestation: One of the core contributions of SGX is the support for attestation between enclaves of the same (local attestation) and different platforms (remote attestation). In the case of local attestation, an enclave enc_i can verify another enclave enc_j as well as the program/software running in the latter. This is achieved through a report generated by enc_j containing information about the enclave itself and the program running in it. This report is signed with a secret key sk_{rpt} which is the same for all enclaves of the same platform. In remote attestation, enclaves of different platforms can attest each other through a signed quote. This is a report similar to the one used in local attestation. The difference is that instead of using sk_{rpt} to sign it, a special private key provided by Intel is used. Thus, verifying these quotes requires contacting Intel’s Attestation Server.

Sealing: Every SGX processor comes with a Root Seal Key with which, data is encrypted when stored in untrusted memory. Sealed data can be recovered even after an enclave is destroyed and rebooted on the same platform.

4 Cryptographic Primitives

In this section, we give a formal definition for the two main encryption schemes that the paper is based on. We proceed with the definition of a CP-ABE and SSE schemes as described in [1] and [7] respectively.

Definition 1 (Ciphertext-Policy ABE). *A revocable CP-ABE scheme is a tuple of the following five algorithms:*

- CPABE.Setup is a probabilistic algorithm that takes as input a security parameter λ and outputs a master public key MPK and a master secret key MSK. We denote this by $(\text{MPK}, \text{MSK}) \leftarrow \text{Setup}(1^\lambda)$.
- CPABE.Gen is a probabilistic algorithm that takes as input a master secret key, a set of attributes \mathcal{A} and the unique identifier of a user and outputs a secret key which is bind both to the corresponding list of attributes and the user. We denote this by $(\text{sk}_{\mathcal{A}, u_i}) \leftarrow \text{Gen}(\text{MSK}, \mathcal{A}, u_i)$.
- CPABE.Enc is a probabilistic algorithm that takes as input a master public key, a message m and a policy $P \in \mathcal{P}$. After a proper run, the algorithm outputs a ciphertext c_P which is associated to the policy P . We denote this by $c_P \leftarrow \text{Enc}(\text{MPK}, m, P)$.
- CPABE.Dec is a deterministic algorithm that takes as input a user’s secret key and a ciphertext and outputs the original message m iff the set of attributes \mathcal{A} that are associated with the underlying secret key satisfies the policy P that is associated with c_P . We denote this by $\text{Dec}(\text{sk}_{\mathcal{A}, u_i}, c_P) \rightarrow m$.

Definition 2 (Dynamic Index-based SSE). *A dynamic index-based symmetric searchable encryption scheme is a tuple of nine polynomial algorithms $\text{SSE} = (\text{Gen}, \text{Enc}, \text{SearchToken}, \text{AddToken}, \text{DeleteToken}, \text{Search}, \text{Add}, \text{Delete}, \text{Dec})$:*

- SSE.Gen is a probabilistic key-generation algorithm that takes as input a security parameter λ and outputs a secret key K .

- **SSE.Enc** is a probabilistic algorithm that takes as input a secret key \mathbf{K} and a collection of files \mathbf{f} and outputs an encrypted index γ and a sequence of ciphertexts \mathbf{c} .
- **SSE.SearchToken** is a (possibly probabilistic) algorithm that takes as input a secret key \mathbf{K} and a keyword w and outputs a search token $\tau_s(w)$.
- **SSE.AddToken** is a (possibly probabilistic) algorithm that takes as input a secret key \mathbf{K} and a file f and outputs an add token $\tau_a(f)$ and a ciphertext c_f .
- **SSE.DeleteToken** is a (possibly probabilistic) algorithm that takes as input a secret key \mathbf{K} and a file f and outputs a delete token $\tau_d(f)$.
- **SSE.Search** is a deterministic algorithm that takes as input an encrypted index γ , a sequence of ciphertexts \mathbf{c} and a search token $\tau_s(w)$ and outputs a sequence of file identifiers $\mathbf{I}_w \subset \mathbf{c}$.
- **SSE.Add** is a deterministic algorithm that takes as input an encrypted index γ , a sequence of ciphertexts \mathbf{c} , an add token $\tau_a(f)$ and a ciphertext c_f and outputs a new encrypted index γ' and a new sequence of ciphertexts \mathbf{c}' .
- **SSE.Delete** is a deterministic algorithm that takes as input an encrypted index γ , a sequence of ciphertexts \mathbf{c} and a delete token $\tau_d(f)$ and outputs a new encrypted index γ' and a new sequence of ciphertexts \mathbf{c}' .
- **SSE.Dec** is a deterministic algorithm that takes as input a secret key \mathbf{K} and a ciphertext c and outputs a file f .

The security of an SSE scheme is based on the existence of a simulator that is given as input information leaked during the execution of the protocol. In particular to define the security of SSE we make use of the leakage functions $\mathcal{L}_{in}, \mathcal{L}_s, \mathcal{L}_a, \mathcal{L}_d$ associated to index creation, search, add and delete operations [4].

5 Modern Family (MF)

In this section, we present Modern Family (MF) – the core of this paper’s contribution. We start by giving an overview of the SGX hardware functionalities used by the communicating parties as defined in [5]. and we continue with a formal construction.

Hardware:

- **HW.Setup**(1^λ): Takes as input a security parameter λ and produces the secret key sk_{rpt}^1 used to MAC the reports.
- **HW.Load**(Q): Takes as input a program Q . An enclave enc_i is created in which Q will be loaded. Moreover a handle hdl_{enc} is created that will be used as an identifier for the enclave.
- **HW.Run**(hdl, in): Takes as input a handle hdl and some input in . It runs the program in the enclave specified by hdl with in as input.

¹ sk_{rpt} is shared with every enclave on the same platform

- **HW.Run&Report**(hdl, in): Takes as input a handle hdl and some input in. It will output a report that is verifiable by any other enclave on the same platform. The report contains information about the underlying enclave signed with sk_{rpt} .
- **HW.ReportVerify**(hdl', rpt): Takes as input a handle hdl' and a report rpt. Uses sk_{rpt} generated by HW.Setup to verify the MAC of the report.

5.1 Formal Construction

MF is divided into a *Setup phase* and four main phases; *Initialization*, *Key Sharing*, *Editing* and *Revocation*. During the *Setup phase*, all the necessary enclaves are initialized by running the MF.Setup algorithm. In the rest of the phases, the user is interacting with the enclaves by running one of the following algorithms: MF.ABEUserKey, MF.Store, MF.KTStore, MF.KeyShare, MF.Search, MF.Update, MF.Delete and MF.Revoke as described below.

Setup Phase: In this phase MF.Setup runs. Each entity receives a public/private key pair (pk, sk) for a CCA2 secure public cryptosystem PKE. In addition to that, the entities running in enclaves generate a signing and a verification key pair. Finally, MS runs CPABE.Setup to acquire the master public/private key pair (MPK, MSK). An enclave is initialized as follows:

MF.Setup(“initialize”, 1^λ): Each enclave is initialized by generating a public/private and signing/verification key pairs. To do so, the program Q_{ID}^{init} is loaded:

Q_{ID}^{init}

- On input (“initialize”, 1^λ):
 1. Run $(pk, sk) \leftarrow PKE.KeyGen(1^\lambda)$.
 2. Output pk.
 Run $hdl \leftarrow HW.Load(Q_{ID}^{init})$.

Additionally, during the setup phase, the MS enclave loads a program Q_{MS}^{Setup} that outputs the master public/private key pair (MPK, MSK):

Q_{MS}^{Setup}

- On input (“initialize”, 1^λ):
 1. Run $(MPK, MSK) \leftarrow PKE.KeyGen(1^\lambda)$.
 2. Output MPK.
 Run $hdl_{MS} \leftarrow HW.Load(Q_{MS}^{Setup})$.

Initialization Phase: As a first step, a user u_i contacts the MS enclave and requests a secret CP-ABE key. Upon reception, MS authenticates u_i and checks if the user is eligible for receiving such a key. If so, MS generates a CP-ABE key sk_{A, u_i} , encrypts it under pk_i and sends it back to u_i . This is done by running the program Q_{MS}^{SKey} in the MS enclave as shown below:

MF.ABEUserKey(“KeyRequest”, MSK, u_i , $cred_i$, \mathcal{A}): The master enclave program Q_{MS}^{SKey} for generating users’ ABE keys is defined as follows:

\mathbf{Q}_{MS}^{SKey}

- On input (“KeyRequest”, MSK, u_i , $cred_i$, \mathcal{A}):
 1. Verify that u_i is registered. If not, output \perp .
 2. Use MSK and compute $sk_{\mathcal{A},u_i}$.
 3. Compute and output $c = \text{PKE.Enc}(pk_i, sk_{\mathcal{A},u_i})$.
 Run $c \leftarrow \text{HW.Run}(\text{hdl}_{MS}, (\text{“KeyRequest”}, \text{MSK}, i, cred_i, \mathcal{A}))$.

After u_i successfully received $sk_{\mathcal{A},u_i}$ she can start using the CSP to store files remotely. To do so, she first sends a store request *StoreReq* to the CSP. Specifically, u_i sends $m_{req} = \langle r_1, E_{pk_{CSP}}(cred_i), StoreReq, H(r_1 || cred_i || StoreReq) \rangle$ where r_1 is a random number. The CSP authenticates u_i as legitimate and sends back an authorization *Auth* as $m_{ver} = \langle r_2, (Auth), \sigma_{CSP}(H(r_2 || u_i || Auth)) \rangle$. At this point, u_i generates a symmetric key K_i to encrypt her files and sends $m_{store} = \langle r_3, E_{pk_{CSP}}(\gamma_i), c_i, H(r_3 || \gamma_i || c_i) \rangle$ to the CSP.

$\text{MF.Store}(\text{“Store”}, m_{req})$: The CSP enclave program $\mathbf{Q}_{CSP}^{\text{Store}}$ that is responsible for storing encrypted files is defined as follows:

 $\mathbf{Q}_{CSP}^{\text{Store}}$

- On input (“StoreReq”, m_{req}):
 1. Open m_{req} ; verify the message^a; if the verification fails, output \perp .
 2. Compute and output $m_{ver} = \langle r_2, (Auth), \sigma_{CSP}(H(r_2 || u_i || Auth)) \rangle$.
 Run $m_{ver} \leftarrow \text{HW.Run}(\text{hdl}_{CSP}, (\text{“StoreReq”}, m_4))$.
- On input (“store”, m_{store}):
 1. Open m_{store} ; verify the message; if the verification fails, output \perp .
 2. Store (c_i, γ_i) .
 Run $\text{HW.Run}(\text{hdl}_{CSP}, (\text{“store”}, m_{store}))$.

^a By this, we mean that the entity receiving the message verifies the freshness and the integrity of the message and it can also authenticate the sender.

Initialization phase concludes with MF.KTStore where u_i encrypts K_i under MPK to get $c_P^{K_i}$ and sends $m_{keystore} = \langle E_{pk_{KT}}(r_4), c_P^{K_i}, \sigma_i(H(r_4 || c_P^{K_i})) \rangle$ to KT.

Upon reception, KT generates a random number r_{K_i} that is stored next to $c_P^{K_i}$. $\text{MF.KTStore}(\text{“store”}, m_{keystore})$: The KT enclave program $\mathbf{Q}_{KT}^{\text{Store}}$ that stores a symmetric key K_i encrypted with MPK is defined as follows:

 $\mathbf{Q}_{KT}^{\text{Store}}$

- On input (“store”, $m_{keystore}$):
 1. Open $m_{keystore}$; verify the message. If the verification fails, output \perp .
 2. Generate a random number r_{K_i} .
 3. Compute $c = \text{PKE.Enc}(pk_{u_i}, r_{K_i})$.
 4. Store $(c_P^{K_i}, c)$.
 Run $(c_P^{K_i}, c) \leftarrow \text{HW.Run}(\text{hdl}_{KT}, (\text{“store”}, m_{keystore}))$.

Key Sharing Phase: This phase begins with u_j executing MF.KeyShare to prove that is not revoked. To this end, u_j sends $m_{verReq} = \langle r_5, E_{pk_{REV}}(u_j), \sigma_j(r_5 || u_j) \rangle$ to REV. Upon reception, REV verifies the message and checks whether $u_j \in rl$ or not. Assuming that $u_j \notin rl$ (i.e. she has not been revoked), REV replies

with $m_{token} = \langle r_6, E_{pk_{KT}}(u_j), E_{pk_{KT}}(\tau_{KS}), \sigma_{REV}(H(r_6 || u_j || \tau_{KS})) \rangle$. The user then simply forwards m_{token} to KT who verifies it. After the verification is complete KT sends $m_{key} = \langle (E_{pk_{CSP}}(u_j, t)), c_p^{K_i}, \sigma_{KT}(H(u_j || t)) \rangle$ back to u_j , where t is a timestamp declaring the time that u_j accessed $c_p^{K_i}$. If u_j already received K_i in the past, KT will only send back the first and last components of m_{key} .

MF.KeyShare(“share”, m_4) : REV and KT enclave programs (Q_{REV}^{Ver} , Q_{KT}^{Share}) that are responsible for sharing $c_p^{K_i}$ are defined as follows:

Q_{REV}^{Ver}

- On input (“share”, m_{verReq}):
 1. Open m_{verReq} ; verify the message; if the verification fails, output \perp .
 2. Check if $u_j \in rl$; if so, output \perp .
 3. Generate τ_{KS} .
 4. Compute and output m_{token} .
 Run $m_{token} \leftarrow HW.Run(hdl_{REV}, (“share”, m_{verReq})).$

Q_{KT}^{share}

- On input (“share”, m_{token}):
 1. Open m_{token} ; verify the message; if the verification fails, output \perp .
 2. Decrypt $PKE.Enc(pk_{KT}, u_j)$ and $PKE.Enc(pk_{KT}, \tau_{KS})$.
 3. Compute and output m_{key} .
 Run $m_{key} \leftarrow HW.Run(hdl_{KT}, (“share”, m_{ver})).$

User u_j can now run MF.Search to access certain files that are stored in the CSP. To do so, she locally runs SSE.SearchToken to generate $\tau_s(w)$ and then sends $m_{search} = \langle E_{pk_{CSP}}(u_j, t, \tau_s(w)), \sigma_i(H(u_j || t || \tau_s(w))), \sigma_{KT}(H(u_j || t)) \rangle$ to the CSP². Upon reception, CSP runs SSE.Search.

MF.Search(“search”, m_{search} ,) : The CSP enclave program Q_{CSP}^{Search} that is responsible for searching over the encrypted data is defined as follows:

Q_{CSP}^{Search}

- On input (“search”, m_{search}):
 1. Open m_{search} ; verify the message; if the verification fails, output \perp .
 2. Run $SSE.Search(\gamma_i, c_i, \tau_s(w)) \rightarrow \mathbf{I}_w$
 3. Output \mathbf{I}_w .
 Run $HW.Run(hdl_{CSP}, (“search”, m_{search}), which internally runs $SSE.Search \rightarrow \mathbf{I}_w$.$

Editing Phase: In this phase³, registered users can add files to the database and data owners can also delete files. To do so, u_i executes MF.Update and MF.Delete. To update the database, u_i first generates an add token by running $(\tau_\alpha(f), c_f) \leftarrow SSE.AddToken(K_i, f)$. This token is sent to the CSP via $m_{add} = \langle E_{pk_{CSP}}(u_i, t, \tau_\alpha(f), c_i, \gamma_i), \sigma_i(H(u_i || t || \tau_\alpha(f) || c_i || \gamma_i)), \sigma_{KT}(u_i || t) \rangle$. Finally, the CSP verifies the message and its freshness and executes $SSE.Add(\gamma_i, c_i, \tau_\alpha(f), c_f) \rightarrow (\gamma'_i, c'_i)$.

² The user simply forwards the components of m_{key} to the CSP along with a search token $\tau_s(w)$.

³ One could completely ignore the *Editing Phase* and the result would be a static MF.

MF.Update(“update”, m_{add}) : The CSP enclave program Q_{CSP}^{Up} for adding files to the database is defined as follows:

Q_{CSP}^{Up}

- On input (“update”, m_{add}):
 1. Verify the message. If the verification fails, output \perp .
 2. Run $SSE.Add(\gamma_i, c_i, \tau_\alpha(f), c_f) \rightarrow (\gamma'_i, c'_i)$.
 Run $HW.Run(hdl_{CSP}, (“update”, m_{add}), which internally runs $SSE.Add(\gamma_i, c_i, \tau_\alpha(f), c_f) \rightarrow (\gamma'_i, c'_i)$.$

Deletion of a file is a more complicated task. This is due to the fact that we only allow the data owner to delete files. To achieve this, u_i needs to prove her ownership over K_i . This can be done by requesting the random number r_{K_i} from KT. After u_i receives r_{K_i} , she signs it, runs $\tau_d \leftarrow SSE.DeleteToken(K_i, f)$ and replies to KT with: $m_{delete} = \langle E_{pk_{CSP}}(u_i, t, \tau_d(f), \gamma'_i), \sigma_i(H(u_i || \tau_d(f)) || \gamma'_i || r_{K_i}) \rangle$. KT verifies the message and is convinced that u_i is the owner of K_i . Finally, KT generates a report (rpt) containing the delete token. This is sent to the CSP who proceeds with the deletion of the specified files.

MF.Delete(“request”, $\sigma_i(u_i || t), c_p^{K_i}$): The enclave programs Q_{CSP}^{Del} , Q_{KT}^{Del} that are responsible for deleting files from the database are defined as follows:

Q_{KT}^{Del}

- On input (“request”, $\sigma_i(u_i || t), c_p^{K_i}$):
 1. Verify the signature. If the verification fails, output \perp .
 2. Get r_{K_i} and compute $c = PKE.Enc(pk_{u_i}, r_{K_i})$.
 3. Output c .
 Run $c \leftarrow HW.Run(hdl_{KT}, (“request”, $\sigma_i(u_i || t), c_p^{K_i}$)).$
- On input (“delete”, m_{delete}):
 1. Open m_{delete} ; verify the message and authenticate u_i as the owner of K_i . If the verification or the authentication fail, output \perp .
 2. Generate and output rpt.
 Run $HW.Run(hdl_{KT}, (“delete”, m_{delete})) and then $rpt \leftarrow HW.RunReport(hdl_{KT}, (“delete”, m_{delete})).$$

Q_{CSP}^{Del}

- On input (“delete”, rpt):
 1. Verify rpt. If the verification fails, output \perp .
 2. Run $SSE.Delete(\gamma'_i, c'_i, \tau_d(f)) \rightarrow (\gamma''_i, c''_i)$.
 Run $HW.Run(hdl_{CSP}, (“delete”, rpt))$ who will internally run $HW.ReportVerify(hdl_{CSP}, rpt)$ and $SSE.Delete(\gamma_i, c_i, \tau_d(f)) \rightarrow (\gamma''_i, c''_i)$.

Revocation Phase: To successfully run **MF.Revoke**, u_i first needs to prove ownership over K_i by following the same steps as in **MF.Delete**. When u_i signs r_{K_i} , she sends $m_{revoke} = \langle r_{10}, E_{pk_{KT}}(u_i, u_j, c_p^{K_i}), \sigma_i(H(u_i || u_j || c_p^{K_i}) || r_{K_i}) \rangle$ to KT. Now that KT is convinced that u_i is the owner of K_i , it generates rpt containing u_j 's identity, which is then sent to REV, who adds u_j to rl .

MF.Revoke(“request”, $\sigma_i(u_i || t), c_p^{K_i}$) : The enclave programs Q_{KT}^{Rev} , Q_{REV}^{Rev} that are responsible for revoking users are defined as follows:

Q_{KT}^{Rev}

- On input (“request”, $\sigma_i(u_i || t), c_p^{K_i}$):
 1. Verify the signature. If the verification fails, output \perp .
 2. Get r_{K_i} and compute $c = \text{PKE.Enc}(\text{pk}_{u_i}, r_{K_i})$.
 3. Output c .
- Run $c \leftarrow \text{HW.Run}(\text{hdl}_{\text{KT}}, (\text{“request”}, u_i, c_p^{K_i}))$.
- On input (“revoke”, m_{revoke}):
 1. Open m_{revoke} ; verify the message and authenticate u_i as the owner of K_i . If the verification or the authentication fails, output \perp .
 2. Generate $r_{K_i'}$ and replace it with r_{K_i} .
 3. Generate and output rpt .
- Run $\text{HW.Run}(\text{hdl}_{\text{KT}}, (\text{“revoke”}, m_{\text{revoke}}))$ and then $\text{rpt} \leftarrow \text{HW.RunReport}(\text{hdl}_{\text{KT}}, (\text{“revoke”}, m_{\text{report}}))$.

Q_{REV}^{Rev}

- On input (“revoke”, rpt):
 1. Verify rpt . If the verification fails, output \perp .
 2. Add u_j to the revocation list rl .
- Run $\text{HW.Run}(\text{hdl}_{\text{REV}}, (\text{“revoke”}, \text{rpt}))$ who will internally run $\text{HW.ReportVerify}(\text{hdl}_{\text{REV}}, \text{report})$.

6 Security Analysis

We construct a simulator \mathcal{S} that simulates the algorithms of the real protocol in such a way that any polynomial time adversary \mathcal{ADV} will not be able to distinguish between the real protocol and \mathcal{S} . \mathcal{S} intercepts \mathcal{ADV} 's communication with the real protocol and replies with simulated outputs.

Definition 3. (*Sim-Security*). We consider the following experiments. In the real experiment, all algorithms run as defined in our construction while in the ideal one, \mathcal{S} intercepts \mathcal{ADV} 's queries and replies with simulated responses.

Real Experiment

1. $\text{EXP}_{MF}^{\text{real}}(1^\lambda)$:
2. $(\text{MPK}, \text{MSK}) \leftarrow \text{MF.Setup}(1^\lambda)$
3. $\text{sk}_{A, u_i} \leftarrow \mathcal{ADV}^{\text{MF.ABEUserKey}}(\text{MSK}, A)$
4. $ct \leftarrow \text{CPABE.Enc}(\text{mpk}, m)$
5. $(\gamma, c) \leftarrow \mathcal{ADV}^{\text{SSE}}(\text{Enc}(K, \mathbf{f}))$
6. $\text{MF.Search}(\text{“search”}, m_s) \rightarrow \mathbf{I}_w$
7. $\text{MF.Update}(\text{“update”}, m_{\text{add}}) \rightarrow (\gamma', c')$
8. $\text{MF.Delete}(\text{“delete”}, m_{\text{delete}}) \rightarrow (\gamma', c')$
9. Output b

Ideal Experiment

1. $\text{EXP}_{MF}^{\text{ideal}}(1^\lambda)$:
2. $(\text{MPK}) \leftarrow \mathcal{S}(1^\lambda)$
3. $\text{sk}_{A, u_i} \leftarrow \mathcal{ADV}^{\mathcal{S}}(1^\lambda)$
4. $ct \leftarrow \mathcal{S}(1^\lambda, 1^{|m|})$
5. $(\gamma, c) \leftarrow \mathcal{ADV}^{\mathcal{S}}(\mathcal{L}_{\text{in}}(\mathbf{f}))$
6. $\mathcal{S}(\text{“search”}, m_s) \rightarrow \mathbf{I}_w$
7. $\mathcal{S}(\text{“update”}, m_{\text{add}}) \rightarrow (\gamma', c')$
8. $\mathcal{S}(\text{“delete”}, m_{\text{delete}}) \rightarrow (\gamma', c')$
9. Output b'

We say that MF is sim-secure if for all PPT adversaries \mathcal{ADV} :

$$\text{EXP}_{MF}^{\text{real}}(1^\lambda) \approx \text{EXP}_{MF}^{\text{ideal}}(1^\lambda)$$

Everything \mathcal{ADV} observes in the real experiment can be simulated by \mathcal{S} . Moreover, we use an IND-CCA2 public key encryption scheme. If \mathcal{ADV} can distinguish between real and ideal answers, she can also break the IND-CCA2 security. Finally, we let \mathcal{ADV} can load different programs in the enclaves and record the output. This assumption significantly strengthens \mathcal{ADV} since we need to ensure that only honest attested programs will be executed in the enclaves.

Theorem 1. *Assuming that PKE is an IND-CCA2 secure public key cryptosystem and Sign is an EUF-CMA secure signature scheme then MF is a sim-secure protocol according to Definition 3.*

Proof. We start by defining the algorithms used by the simulator. Then, we will replace them with the real algorithms. Finally, the help of a Hybrid Argument we will prove that the two distributions are indistinguishable.

- MF.Setup*: Will only generate MPK that will be given to \mathcal{ADV} .
- MF.ABEUserKey*: Will generate a random key to be sent to the adversary. That is, when \mathcal{ADV} makes a key generation query, \mathcal{S} will simulate CPABE.KeyGen and it will output $\text{sk}_{\mathcal{A},u}^*$. This key is a random string that has the same length as the output of the real MF.ABEUserKey*. The key will be given to \mathcal{ADV} .
- MF.KeyShare*: In the ideal experiment, after \mathcal{ADV} requests a secret key, \mathcal{S} will encrypt a sequence of bits based on \mathcal{L}_{in} , under MPK. The ciphertext will be returned to \mathcal{ADV} .
- MF.Search*: When \mathcal{ADV} generates a search token $\tau_s(w)$, \mathcal{S} gets as input the leakage function \mathcal{L}_s and outputs a simulated response. When \mathcal{ADV} makes a search query, \mathcal{S} will once again generate a simulated \mathbf{I}_w^* which will be sent back to her.
- MF.Update*: When \mathcal{ADV} generates an add token $\tau_\alpha(f)$, \mathcal{S} gets as input the leakage function \mathcal{L}_a and outputs a simulated response. \mathcal{S} will simulate the add token, the ciphertext to be added to the database and will also update the encrypted index.
- MF.Delete*: When \mathcal{S} generates a delete token, \mathcal{S} gets as input the leakage function \mathcal{L}_d and outputs a simulated response. Apart from $\tau_d(f)$, \mathcal{S} will also update the encrypted index.
- MF.Revoke*: The system does not revoke any user.

In the pre-processing phase, \mathcal{S} runs $\text{HW.Setup}(1^\lambda)$, just as in the real experiment, in order to acquire sk_{rpt} . Moreover, the challenger \mathcal{C} generates a symmetric key K_i , that will be needed in order to reply to search, add and delete queries. We will now use a hybrid argument to prove that \mathcal{ADV} cannot distinguish between the real and the ideal experiments.

Hybrid 0 MF runs normally.

Hybrid 1 Everything runs like in Hybrid 0, but we replace MF.Setup with MF.Setup*.

These algorithms are identical from \mathcal{ADV} 's perspective and as a result the hybrids are indistinguishable.

Hybrid 2 Everything runs like in Hybrid 1, but MF.ABEUserKey^* runs instead of MF.ABEUserKey .

Hybrid 2 is indistinguishable from Hybrid 1 because nothing changes from \mathcal{ADV} 's point of view.

After Hybrid 2, we have ensured that \mathcal{ADV} has followed all the required steps in order to ask for K_i . We are now ready to replace MF.KeyShare with MF.KeyShare^* .

Hybrid 3 Like Hybrid 2, but MF.KeyShare^* runs instead of MF.KeyShare . Also, the algorithm outputs \perp if HW.Run is queried with $(\text{hdl}_{KT}, ("share", m_{token}))$ but \mathcal{ADV} never contacts REV .

Lemma 1. *Hybrid 3 is indistinguishable from Hybrid 2.*

Proof. Replacing the two algorithms, does not change from \mathcal{ADV} 's perspective. If \mathcal{ADV} can generate m_{token} , then she can forge REV 's signature. Given the security of the signature scheme, this can only happen with negligible probability. So \mathcal{ADV} can distinguish between the Hybrids with negligible probability. \square

At this point, \mathcal{ADV} has received what she thinks is a valid K_i . The simulator now gets access to all leakage functions \mathcal{L} from the SSE scheme.

Hybrid 4 Like Hybrid 3, but when HW.Run is queried with $(\text{hdl}_{\text{CSP}}, ("search", m_{search}))$, \mathcal{S} is given the leakage function \mathcal{L}_S and generates \mathbf{I}_w^* which is then sent to the user.

Lemma 2. *Hybrid 4 is indistinguishable from Hybrid 3.*

Proof. Assuming the \mathcal{L}_i - security of the SSE scheme, the token sent by \mathcal{ADV} to the CSP, as part of m_{search} , is generated by \mathcal{S} with \mathcal{L}_s as input. As a result when \mathcal{S} receives m_{search} , it will generate a sequence of file identifiers \mathbf{I}_w^* that will be send back to \mathcal{ADV} . \mathcal{ADV} cannot distinguish between the real and the ideal experiment since she receives a sequence of files corresponding to a search token that was also simulated by \mathcal{S} . Moreover, if \mathcal{ADV} manages to generate m_{search} without having contacted KT earlier, then she can also forge KT 's signature. However, this can only happen with negligible probability, and as a result \mathcal{ADV} can only distinguish between hybrids 4 and 3 with negligible probability. \square

Hybrid 5 Like Hybrid 4, but when HW.Run is queried with $(\text{hdl}_{\text{CSP}}, ("update", m_{add}))$, \mathcal{S} is given the leakage function \mathcal{L}_a and tricks \mathcal{ADV} into thinking that she updated the database.

Lemma 3. *Hybrid 5 is indistinguishable form Hybrid 4.*

Proof. By assuming the \mathcal{L}_i - security of the SSE scheme, we know that \mathcal{ADV} will not be able to distinguish between the real add token and the simulated one.

Moreover, similar to the previous Hybrid, if \mathcal{ADV} can generate m_{add} without having contacted KT, then she can also forge KT’s signature – which can only happen with negligible probability. Hence, \mathcal{ADV} can only distinguish between hybrids 5 and 4 with negligible probability. \square

Hybrid 6 Like Hybrid 5, but when HW.Run is queried with $(\text{hdl}_{\text{KT}}, ("delete", m_{del}))$, \mathcal{S} is given the leakage function \mathcal{L}_d and tricks \mathcal{ADV} into thinking that she deleted a certain file from the database. Moreover, \mathcal{S} outputs \perp , if ReportVerify is queried with $(\text{hdl}_{\text{CSP}}, \text{rpt})$ for a report that was not generated by executing $\text{HW.RunReport}(\text{hdl}_{\text{KT}}, ("delete", m_{delete}))$.

Proof. By assuming the \mathcal{L}_i – security of the SSE scheme, we know that \mathcal{ADV} will not be able to distinguish between the real delete token and the simulated one. Moreover, if \mathcal{ADV} can query HW.ReportVerify with $(\text{hdl}_{\text{CSP}}, \text{rpt})$, for a rpt that was not generated by KT, then \mathcal{ADV} can produce a valid MAC which can only happen with negligible probability since she does not know sk_{rpt} . Thus, \mathcal{ADV} can only distinguish between Hybrids 5 and 6 with negligible probability. \square

Hybrid 7 Like Hybrid 6 but instead of MF.Revoke , \mathcal{S} executes MF.Revoke^* .

The hybrids are indistinguishable since no one can access the content of the revocation list and as a result nothing changes from \mathcal{ADV} ’s point of view.

With this Hybrid our proof is complete. We managed to replace the expected outputs with simulated responses in a way that \mathcal{ADV} cannot distinguish between the real and the ideal experiment. \square

6.1 SGX Security

Recent works [3, 17, 8, 16] have shown that SGX is vulnerable to software attacks. However, according to [5], these attacks can be prevented if the programs running in the enclaves are data-oblivious. Thus, leakage can be avoided if the programs do not have memory access patterns or control flow branches that depend on the values of sensitive data. In our construction, no sensitive data are used by the enclaves. KT acts as a storage space for the symmetric keys and does not perform any computation on them. Hence, all the $c_p^{K_i}$ are data-oblivious. Moreover, rl is stored in plaintext and every entry in the list is padded to achieve same length.

7 Conclusion

In this paper, we proposed MF, a hybrid encryption scheme that combines *both* SSE and ABE in a way that the main advantages of each encryption technique are used. The proposed scheme enables clients to search over encrypted data by using an SSE scheme, while the symmetric key required for the decryption is protected via a Ciphertext-Policy Attribute-Based Encryption scheme. Moreover, our construction supports the revocation of users by utilizing the functionality provided by SGX. In contrast to recent works, the revocation mechanism has been separated from the actual ABE scheme and is exclusively based on the utilization of trusted SGX enclaves.

References

1. Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: Proceedings of the 2007 IEEE Symposium on Security and Privacy. pp. 321–334. SP '07, IEEE Computer Society, Washington, DC, USA (2007)
2. Boyko, V.: On the security properties of oaep as an all-or-nothing transform. In: Wiener, M. (ed.) Advances in Cryptology — CRYPTO' 99. pp. 503–518. Springer Berlin Heidelberg, Berlin, Heidelberg (1999)
3. Costan, V., Devadas, S.: Intel sgx explained. Cryptology ePrint Archive, Report 2016/086 (2016), <https://eprint.iacr.org/2016/086>
4. Dowsley, R., Michalas, A., Nagel, M., Paladi, N.: A survey on design and implementation of protected searchable data in the cloud. Computer Science Review (2017)
5. Fisch, B., Vinayagamurthy, D., Boneh, D., Gorbunov, S.: Iron: Functional encryption using intel sgx. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. pp. 765–782. CCS '17, ACM (2017)
6. Guo, W., Dong, X., Cao, Z., Shen, J.: Efficient attribute-based searchable encryption on cloud storage. Journal of Physics: Conference Series (2018)
7. Kamara, S., Papamanthou, C., Roeder, T.: Dynamic searchable symmetric encryption. pp. 965–976 (2012)
8. Lee, S., Shih, M., Gera, P., Kim, T., Kim, H., Peinado, M.: Inferring fine-grained control flow inside SGX enclaves with branch shadowing. In: 26th USENIX Security Symposium, BC, Canada, August 16–18, 2017. pp. 557–574 (2017)
9. Liu, J.K., Yuen, T.H., Zhang, P., Liang, K.: Time-based direct revocable ciphertext-policy attribute-based encryption with short revocation list. Cryptology ePrint Archive, Report 2018/330 (2018), <https://eprint.iacr.org/2018/330>
10. Michalas, A.: Sharing in the rain: Secure and efficient data sharing for the cloud. In: Proceedings of the 11th IEEE International Conference for Internet Technology and Secured Transactions (ICITST-2016). IEEE (2016)
11. Michalas, A.: The lord of the shares: Combining attribute-based encryption and searchable encryption for flexible data sharing. In: Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing. pp. 146–155. SAC '19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3297280.3297297>, <http://doi.acm.org/10.1145/3297280.3297297>
12. Microsoft: Microsoft Security Intelligence Report (2017)
13. Myers, S., Shull, A.: Practical revocation and key rotation. In: Smart, N.P. (ed.) Topics in Cryptology – CT-RSA 2018. pp. 157–178. Springer, Cham (2018)
14. Paladi, N., Gehrman, C., Michalas, A.: Providing user security guarantees in public infrastructure clouds. IEEE Transactions on Cloud Computing **5**(3), 405–419 (July 2017). <https://doi.org/10.1109/TCC.2016.2525991>
15. Paladi, N., Michalas, A., Gehrman, C.: Domain based storage protection with secure access control for the cloud. In: Proceedings of the 2014 International Workshop on Security in Cloud Computing. ASIACCS '14, ACM, New York, NY, USA (2014)
16. Weichbrodt, N., Kurmus, A., Pietzuch, P.R., Kapitza, R.: Asyncshock: Exploiting synchronisation bugs in intel SGX enclaves. In: Computer Security - ESORICS 2016 - 21st European Symposium on Research in Computer Security, Heraklion, Greece, September 26–30, 2016, Proceedings, Part I. pp. 440–457 (2016)
17. Xu, Y., Cui, W., Peinado, M.: Controlled-channel attacks: Deterministic side channels for untrusted operating systems. In: Proceedings of the 36th IEEE Symposium on Security and Privacy (Oakland). IEEE (May 2015)