

State-free End-to-End Encrypted Storage and Chat Systems based on Searchable Encryption

Keita Emura[§], Ryoma Ito[§], Sachiko Kanamori[§], Ryo Nojima[§], and Yohei Watanabe^{†, §}

[§]National Institute of Information and Communications Technology (NICT), Japan.

[†]The University of Electro-Communications, Japan.

July 14, 2021

Abstract

Searchable symmetric encryption (SSE) has attracted significant attention because it can prevent data leakage from external devices, e.g., clouds. SSE appears to be effective to construct such a secure system; however, it is not trivial to construct such a system from SSE in practice because other parts must be designed, e.g., user login management, defining the keyword space, and sharing secret keys among multiples users who usually do not have public key certificates. In this paper, we describe the implementation of two systems from the state-free dynamic SSE (DSSE) (Watanabe et al., ePrint 2021), i.e., a secure storage system (for a single user) and a chat system (for multiple users). In addition to the DSSE protocol, we employ a secure multipath key exchange (SXKEX) protocol (Costea et al., CCS 2018), which is secure against some classes of unsynchronized active attackers. It allows the chat system users without certificates to share a secret key of the DSSE protocol in a secure manner. To realize end-to-end encryption, the shared key must be kept secret; thus, we must consider how to preserve the secret on, for example, a user’s local device. However, this requires additional security assumptions, e.g., tamper resistance, and it seems difficult to assume that all users have such devices. Thus, we propose a secure key agreement protocol by combing the SXKEX and login information (password) that does not require an additional tamper-resistant device. Combining the proposed key agreement protocol and the underlying state-free DSSE protocol allow users who know the password to use the systems on multiple devices.

1 Introduction

1.1 Searchable Symmetric Encryption

Searchable symmetric encryption (SSE) [18, 51] provides search functionality against encrypted documents, and dynamic SSE (DSSE) [9–11, 14, 20, 30, 31, 34, 37, 40, 42, 50, 52, 54, 57, 59] allows us to update encrypted databases. For example, in practical applications, when encrypted storage is constructed, the database is updated frequently; thus, DSSE is employed. As a fundamental security of DSSE, Stefanov et al. [52] defined forward privacy, which guarantees that even if some data are added, information about whether the data contain keywords that have been searched previously is not revealed.

DSSE prevents data leakage from external storages, e.g., clouds, because all stored data are encrypted. Such a DSSE-based storage system is described as follows. A user selects a key k that is kept secret, and an identifier id is associated with each file f_{id} . Here, assume that no information

of f_{id} is revealed from id . A storage server manages an encrypted database that comprises the pair (id, c_{id}) , where c_{id} is the ciphertext of f_{id} . Let \mathcal{W}_{id} be a set of keywords of file f_{id} with identifier id . The user computes a search query from k , a keyword to be searched $\omega \in \mathcal{W}_{id}$, and the state information, sends the query to the server, and then obtains c_{id} in which the corresponding f_{id} (i.e., the decryption result of c_{id} using k) contains ω . No information of ω is revealed from the query (more precisely, a leakage function is defined, and no information of ω is revealed besides this function). Finally, the user obtains f_{id} by decrypting c_{id} using k .

A secure storage system can be constructed easily from DSSE; however, many issues must be considered if such a system is launched in practice. For example, a set of keywords \mathcal{W}_{id} , which is assumed to be given in advance in DSSE, must be defined. Although a promising approach is to employ a morphological analysis tool, this approach introduces other issues, e.g., selecting the most appropriate tool. In addition, in a storage system, an independent area is assigned to each user; thus, authentication and the user login process must also be considered. Moreover, DSSE attempts to prevent information leakage against the server; thus, we must also consider cases where search queries sent from users are modified in the communication channel. We must also consider the case where multiple users share a key, e.g., a secure chat system, where the ciphertexts of chat history are preserved on the server, each user can search chat messages owing to DSSE, and each user reads them in a plaintext manner by locally decrypting the ciphertexts. Here, we must consider how to share a secret key among multiple users, and, if the state information (which is updated periodically and used to generate search queries) must be managed, then it must also be shared among users, which represents an additional synchronization problem. Note that users do not possess public key certificates in many cases, e.g., smartphones; thus, man-in-the-middle attacks can be made by an active adversary that controls the communication channel among users. In summary, it would be beneficial to address these issues (in addition to DSSE) in secure systems.

1.2 Our Contribution

In this paper, we implement two systems from the DSSE scheme proposed by Watanabe et al. [57], i.e., a secure storage system (for a single user) and a chat system (for multiple users). Watanabe’s DSSE protocol is state-free, which means that if a user knows a (stateless) secret key, then no other state information is required. This allows us to consider multiple users easily we only need to handle key agreement. In other words, we do not have to consider the synchronization of state information.¹ By combining a key agreement protocol (which is explained later) and the Watanabe DSSE protocol, these two systems are state-free; thus, the user can use the systems via a web browser (without considering devices) if they know the appropriate login information (i.e., the user ID and password).

Security Model. In our system, we prepared two (semi-honest) servers, i.e., an authentication server (to manage login information) and an application server (that preserves encrypted data and responds to the users’ search queries). We considered a realistic situation where two servers have public key certificates via a public key infrastructure (PKI), and the users do not have certificates. Here, we pursue end-to-end encryption (E2EE), i.e., only the corresponding users have a secret key, and no server can observe the plaintext data (even two servers collude with each other). Thus, we considered a relaxed security model, i.e., unsynchronized active adversaries, presented by Costea et al. [17]. Costea et al. proposed the secure multipath key exchange (SXKEX) protocol, which is secure against unsynchronized active adversaries. The SMKEX protocol allows chat users to share

¹To the best of our knowledge, Watanabe’s DSSE protocol is the first state-free construction with forward privacy; thus, we employed this protocol in this paper.

a secret key without assuming a PKI.

Our Key Agreement Protocol. To realize E2EE, the shared key must be kept secret; thus, we must consider how to preserve key secrecy on, for example, a local user device. However, this requires additional security assumptions, e.g., tamper resistance, and it seems difficult to assume that all users have such a device, as in certificates. In addition, it would be beneficial to access the systems via multiple devices without synchronization; thus, we propose a secure key agreement protocol that combines the SXKEX protocol and login information (password). The proposed secure key agreement protocol does not require additional (tamper-resistant) devices. Here, a DSSE secret key is defined by the password and a random value preserved in the application server. Then, when a user logs into the system, they obtain the random and compute the DSSE secret key locally. Although this is similar to password-based authenticated key exchange (PAKE) [32], no secret value shared in advance is required in the proposed protocol (under relaxed security). By combining the key management protocol and Watanabe’s state-free DSSE protocol, users can access the systems on multiple devices, and state-free E2EE storage and chat systems can be constructed.

Concierge Functionality. We also consider the explainability of the system. Typically, general users do not aware SSE; thus, such secure systems should be used without recognizing the underlying cryptographic tools. Even for general users, it is highly desirable to easily explain how data are encrypted, how encrypted data are preserved on external storage devices, and so on. Thus, we also implemented a concierge functionality where DSSE-related data processing can be viewed (see Appendix).

1.3 Related Work

CryptDB [44, 45, 48] is a popular encrypted database system in which SQL queries are executed on encrypted data. As discussed in the literature [46], the application server obtains access to the unencrypted data and receives each user’s key when a user logs in. In this sense, it is not an E2EE system. Popa et al. [46] proposed Mylar, , which is a platform to build web applications using a multi-key DSSE protocol [47]. They also published the kChat chat service, which is based on Mylar. Although they insisted that Mylar protects data confidentiality against attackers who have full access to the servers, Grubbs et al. [24] demonstrated that Mylar is vulnerable against active adversarial servers that modify the encryption algorithm. Here, we assume that the two servers in our systems are semi-honest; thus, Mayer might be employed. However, this is not dynamic and requires paring groups; thus, we employ Watanabe’s DSSE protocol.

Recently, secure enclave-assisted constructions have been proposed, e.g., [6, 35, 41, 49, 56]. This direction is interesting but it additionally employs trusted execution environments (TEEs) such as Intel SGX. Moreover, enhanced security besides forward privacy, e.g., [5, 8, 19, 23, 29, 33, 43, 60] or enhanced search functionality, e.g., [12, 13, 21, 25, 28, 55], also have been proposed. Since our key agreement protocol is independent to the underlying (D)SSE, although it has a good compatibility to the Watanabe DSSE protocol in terms of state-freeness, these constructions might be employed instead of the Watanabe DSSE protocol.

Secure messaging protocols have been widely researched, to name a few [3, 4, 15, 16, 26, 39]. This is an E2EE protocol because all messages through on the communication channel are encrypted and nobody (except users) can decrypt them even by the service provider, e.g., Signal and WhatsApp. Unlike to our E2EE systems, they do not consider the search functionality over encrypted data. Crypto-chat [1] was established for secure messaging, where users share passwords, and encrypted messages are decrypted on the device only. To the best of our knowledge, no search functionality against encrypted data is supported.

2 Preliminaries

2.1 Watanabe et al. DSSE

In this section, we introduce the Watanabe DSSE protocol [57]. Let $\pi = \{\pi_k : \{0, 1\}^* \rightarrow \{0, 1\}^{\lambda+\ell}\}_{k \in \{0, 1\}^\kappa}$ be a variable input-length pseudorandom function, where λ is the keyword length, ℓ is the identity length, and κ is the key length, which are all polynomial of the security parameter. Typically, the DSSE protocol does not explicitly consider data encryption; however, here we consider it explicitly because the search result is a ciphertext in the storage and chat systems.

Setup: A user selects a secret key $k \in \{0, 1\}^\kappa$. For the simplicity, we assume that k is also used for data encryption.

Update: When data are preserved on the server, the user computes $\pi_k(\omega, \text{id})$ for all $\omega \in \mathcal{W}_{\text{id}}$. Here, \mathcal{W}_{id} is a set of keywords in the file f_{id} with the identifier id . The set of identifiers \mathcal{I} is considered to be the state information, which is updated periodically according to the current database. The user encrypts f_{id} using k and sends id , $\pi_k(\omega, \text{id})$, and the ciphertext c_{id} to the server. Then, the server preserves $(\text{id}, c_{\text{id}})$ on the address $\pi_k(\omega, \text{id})$. When data are removed, the user sends the id of the removed data to the server, and the server removes $(\text{id}, c_{\text{id}})$.

Search: If the user searches files containing keyword ω , the user computes a trapdoor $\pi_k(\omega, \text{id})$ for all $\text{id} \in \mathcal{I}$ and sends a search query $\{\pi_k(\omega, \text{id})\}_{\text{id} \in \mathcal{I}}$. The server sends $(\text{id}, c_{\text{id}})$ preserved on the address $\pi_k(\omega, \text{id})$. Finally, the user decrypts c_{id} using k and obtains f_{id} .

In the Watanabe DSSE protocol, the server is modeled as semi-honest, i.e., it always follows the protocol procedure but may extract information. Assume that id does not reveal any information of f_{id} . Then state information $\mathcal{I} = \{\text{id}\}$ can be publicly available, and simply the server preserves \mathcal{I} and sends it to the user before the user searches. The server knows ciphertexts c_{id} and pseudorandom numbers $\pi_k(\omega, \text{id})$. Moreover, queries $\{\pi_k(\omega, \text{id})\}_{\text{id} \in \mathcal{I}}$ are computed for the current database. Thus, the Watanabe DSSE protocol supports forward privacy and is state-free.

2.2 SMKEX and Unsynchronized Adversaries

In this section, we introduce SMKEX proposed by Costea et al. [17] and its security model. In the two adversaries case which we also employed, unsynchronized adversaries are defined as follows:

Definition 2 [17] Two adversaries X1 and X2 are said to be unsynchronized (written X1/X2) if they can only exchange messages before the start and after the end of a specific protocol session

For example, let two active adversaries be considered and two paths be prepared. Then, one active adversary can observe and modify data on the first path, and the other active adversary can also observe and modify the data on the second path; however, these adversaries cannot communicate with each other. Costea et al. proposed the SXKEX protocol, which is secure against the adversaries.

The SXKEX protocol is described as follows. Essentially, it is a simple Diffie-Hellman (DH)-type key exchange with an additional confirmation phase. Here, let \mathbb{G} be a group with prime order p and let $g \in \mathbb{G}$ be a generator. Two users, i.e., Alice and Bob, would like to share a key. Then, Alice (resp. Bob) selects $x \xleftarrow{\$} \mathbb{Z}_p$ (resp. $y \xleftarrow{\$} \mathbb{Z}_p$) and computes g^x (resp. g^y). Note that g^x and g^y are not long-lived keys, and they need to choose them for each key exchange. Through Pass 1,

Alice sends g^x to Bob, and Bob sends g^y to Alice. Note that these values may be modified because the adversary is active. Alice further selects a nonce N_A and sends it to Bob in Pass 2. Then, Bob selects a nonce N_B , computes $\text{hsess} = \text{Hash}(N_A, g^x, N_B, g^y)$, and sends N_B and hsess to Alice in Pass 2. Alice then checks whether $\text{hsess} = \text{Hash}(N_A, g^x, N_B, g^y)$ holds. Since the adversaries are unsynchronized, even if one adversary observes g^y in Pass 1, the adversary cannot compute $\text{Hash}(N_A, g^x, N_B, g^y)$ and send it to Alice in Pass 2. Here, the actual shared key is computed according to RFC5869 [36], where a negotiated secret string is computed from the DH key g^{xy} with a 0 seed via HKDF-extract, and the actual shared key is the HKDF-expand value of the string.

3 Proposed Systems

In this section, we present our storage and chat systems.

3.1 Common Part

DSSE Library. We implemented our DSSE library in the C programming language. Here, we defined the APIs by following the DSSE syntax (**Setup**, **Update**, **Search**). When data are added, a trapdoor is computed for the data. In addition, when data are removed, the user sends the corresponding **id**, and the server removes $(\text{id}, c_{\text{id}})$. In other words, no cryptographic operations are required. Thus, we implemented the **Add** API as **Update** and did not implement the **Delete** API in the library. We employed OpenSSL (1.1.1h) to select k randomly, and HMAC-SHA256 as π_k .² We also employed the WebCrypto API, which is a JavaScript API, to implement the encryption functionality as a web application. For encryption, pseudo-randomness against chosen plaintext attack (PCPA) security [18] is required.³, thus, we employed AES-CTR with a 256-bit key. In addition, we employed MeCab [2] as the underlying morphological analysis tool. Note that we used the wasm MeCab library (v 0.996; ipadic dictionary).⁴ After executing the morphological analysis tool, trapdoors are generated using the **Add** API. To the best of our knowledge, “pneumonoultramicroscopicsilicovolcanokoniosis” (containing 45 characters) is the longest English word; thus, we set $\lambda = 45$.

System Architecture. We prepared an authentication server to manage user login information and application server to preserve the encrypted data and respond to the users’ search queries. A user can use the DSSE library via a web browser (WebAssembly). In this implementation, we employed Amazon Elastic Compute Cloud (Amazon EC2)⁵ and assumed that the two servers have public key certificates. We also assumed that the communication channel between the user and the application server is secured via transport layer security (TLS). The system architecture is shown in Fig. 1

Login Interface. In this implementation, the login interface is common to both systems, and the user selects the storage or chat system (Fig. 2). Here, we employed a simple login system, where each user a username **uname** and password **pw**. The authentication server preserves $\text{Hash}(\text{pw})$ with **uname**, and the user sends $(\text{uname}, \text{Hash}(\text{pw}))$ via TLS to the server. The generation of the DSSE

²The Watanabe DSSE protocol requires that (1) $\pi_k(\omega, \text{id})$ is pseudorandom and (2) the probability that a probabilistic polynomial-time adversary finds two distinct inputs $(\omega, \text{id}) \neq (\omega', \text{id}')$ where $\pi_k(\omega, \text{id}) = \pi_k(\omega', \text{id}')$ holds is negligible for the security parameter. Thus, we employed HMAC-SHA256 in our implementation.

³The reason behind is that the simulator just responds a random value in the security proof. Thus, a standard CPA security is also enough owing to the indistinguishability of ciphertext of 0.

⁴<https://github.com/fasiha/mecab-emscripten>

⁵<https://aws.amazon.com/jp/ec2/>

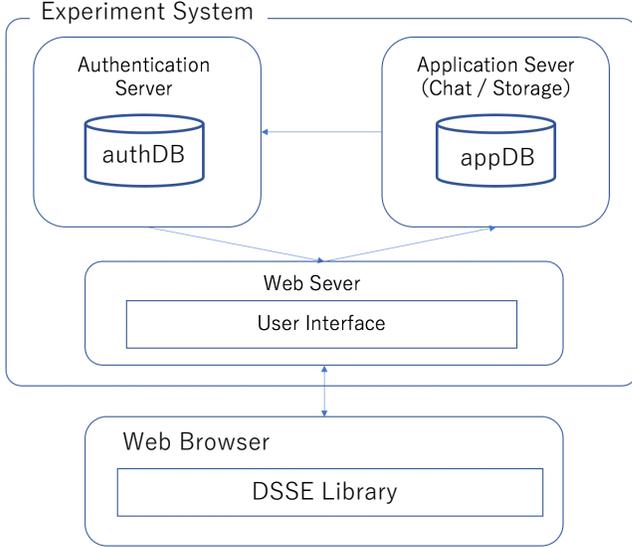


Figure 1: System Architecture

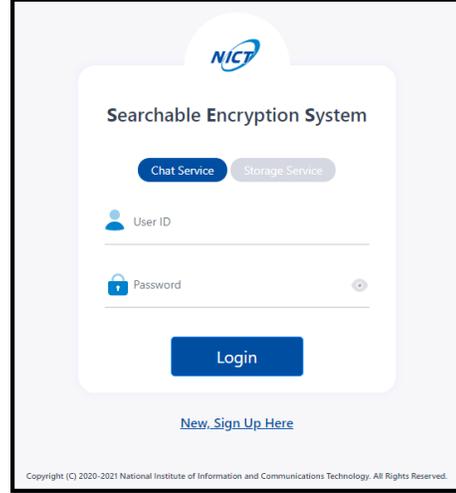


Figure 2: Login

secret key k is explained later. The application server preserves (id, c_{id}) as mentioned in Watanabe’s DSSE protocol. Here, id is generated by the universally unique identifier (UUID) version 4 [38]. It does not take file information as input; therefore, the requirement is satisfied, i.e., id does not reveal any information of f_{id} . The application server also preserves the state information $\mathcal{I} = \{id\}$ for each user.

3.2 Our Storage System

In this section, we give our storage system.

DSSE Key Generation. A user generates a DSSE key k as follows. First, the user selects a random value $R \in \{0, 1\}^\kappa$, where κ is the security parameter, and we set $\kappa = 256$. In the user registration phase, the user selects two different passwords. From a usability and practicality perspective, we assume that the user selects one password PW , and the system separates it such as $PW = pw||pw'$.⁶ The user sends R and $(uname, Hash(pw))$ via TLS to the authentication server, and the server preserves R in addition to $(uname, Hash(pw))$ where $Hash$ is SHA256.⁷ Then, a DSSE secret key is defined as

$$k = R \oplus Hash(pw')$$

where \oplus is a bitwise exclusive OR. In the login phase, the user sends $uname$ and $Hash(pw)$ to the application server via TLS, and the server returns R if $Hash(pw)$ is preserved with $uname$. This structure allows the user to generate the DSSE secret key k without requiring additional information (besides $uname$, pw , and pw'). Briefly, R is random, and no information of k is revealed from R . Even if the authentication server recovers pw from $Hash(pw)$ via an offline dictionary attack, no

⁶E.g., the first half and the second half, or more generally, PW is divided into $pw||pw'$ where $|pw| = \text{floor}(|PW|/2)$ and $|pw'| = \text{ceiling}(|PW|/2)$.

⁷We can employ some zero-knowledge proof system to demonstrate that the user actually knows pw , e.g., zk-SNARK [22]. Here, the communication channel is secure (TLS), and there is no intermediate adversary that can observe or modify $Hash(pw)$; thus, we did not further consider it in this implementation. However, the system can be extended easily in this sense.

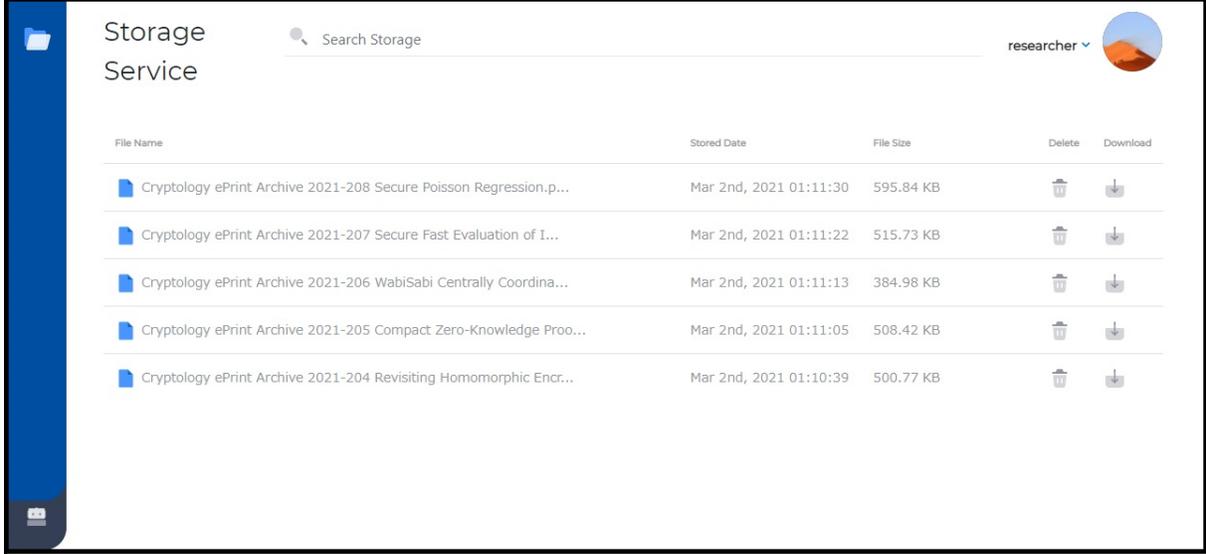


Figure 3: Storage System (Normal)

information of k is revealed because pw' is only used locally by the user. As a potential attack, if the authentication server obtains a ciphertext c_{id} , then the server performs an offline dictionary attack where choose pw' , compute $k = R \oplus \text{Hash}(\text{pw}')$, and check whether the decryption result of c_{id} using k is meaningful, e.g., whether a readable file is recovered or not. Note that c_{id} is sent from the user to the application server via TLS, which means that the authentication server does not perform this attack unless the authentication and application servers collude.

Secure Storage. When the user stores a file on the application server, the file is encrypted automatically. When a user downloads a file to the application server, the file is decrypted automatically. Although the file names are encrypted, they are also decrypted automatically and displayed as usual. Thus, users are not required to aware DSSE. The storage system is shown in Fig. 3, where the user name is researcher, and the preserved data are PDF files from the Cryptology ePrint Archive (<https://eprint.iacr.org/>).

3.3 Our Chat System

Here, we describe the chat system. The main difference from the storage system is the preparation of a random value R for each room in the chat system. In addition, a DSSE key is shared to users belonging to the room. Here, we assume that Alice creates a room and invites Bob to the room, and then both Alice and Bob are registered in the system (i.e., they have their own storage). Let pw_A and pw'_A (pw_B and pw'_B) be Alice's (Bob's) two passwords. We assume that there are two different communication passes as in the SMKEX protocol. Concretely, we consider the following.

Pass 1: Alice \leftrightarrow the authentication server \leftrightarrow Bob which are secure due to TLS.

Pass 2: Alice \leftrightarrow Bob which is different from Pass 1 and we simply assume an e-mail system.

In other words, the system is secure if the authentication server cannot read e-mails sent from Alice to Bob and from Bob to Alice, which is a realistic assumption. Finally, the authentication server preserves the random values R_A and R_B for Alice and Bob, respectively. Then, the room key k is

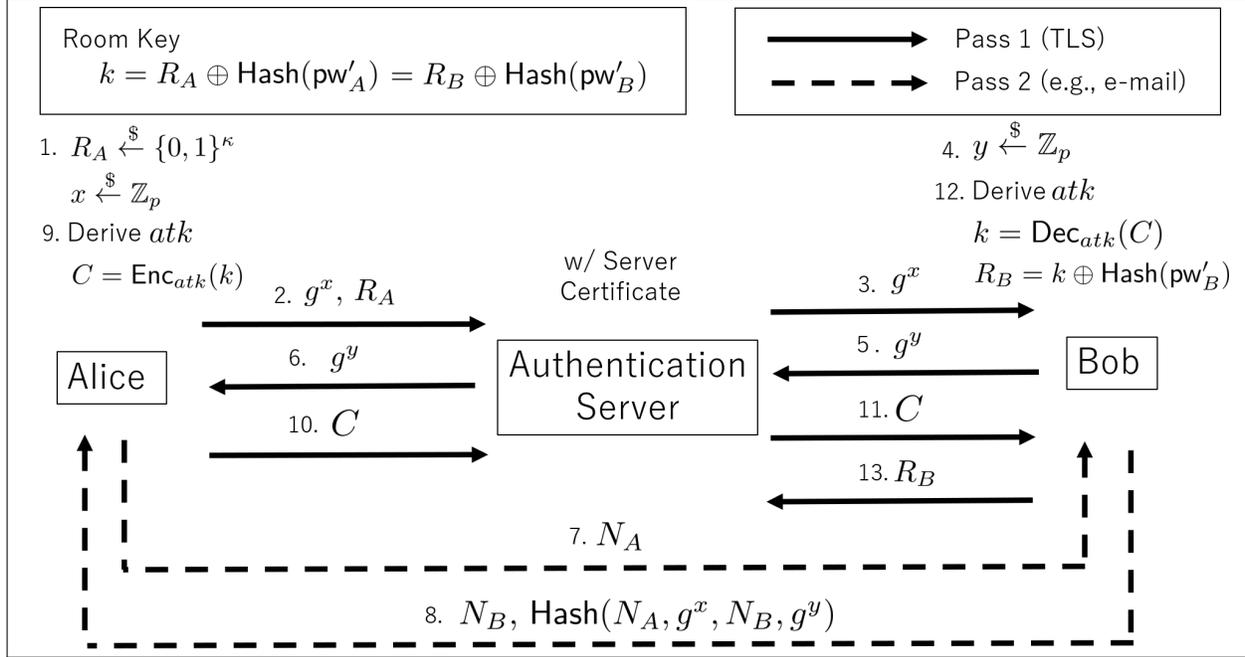


Figure 4: Our Key Agreement Protocol Based on SMKEX

defined as $k = R_A \oplus \text{Hash}(\text{pw}'_A) = R_B \oplus \text{Hash}(\text{pw}'_B)$. Our main idea is to encrypt the DSSE key by using a SMKEX key atk , and Alice sends the ciphertext to Bob. Then, Bob can obtain k and define R_B such that $R_B = k \oplus \text{Hash}(\text{pw}'_B)$. This protocol allows Alice and Bob to log into the chat system (similar to the storage system). The actual key agreement is described as follows (Fig. 4).

Alice: Choose a random value $R_A \in \{0, 1\}^\kappa$. Set the DSSE key for the room $k = R_A \oplus \text{Hash}(\text{pw}'_A)$.

Choose $x \xleftarrow{\$} \mathbb{Z}_p$ and compute a SMKEX public key g^x . Send g^x and R_A to the authentication server (via Pass 1).

Authentication Server: Preserve R_A with the user name Alice. Send g^x and R'_B to Bob (via Pass 1).

Bob: Choose $y \xleftarrow{\$} \mathbb{Z}_p$ and compute a SMKEX public key g^y . Send g^y to the authentication server (via Pass 1).

Authentication Server: Forward g^y to Alice (via Pass 1).

Alice: Choose a nonce N_A and send it to Bob (via Pass 2).

Bob: Choose a nonce N_B , compute $\text{hss} = \text{Hash}(N_A, g^x, N_B, g^y)$, and send N_B and hss to Alice (via Pass 2).

Alice: Compute $\text{Hash}(N_A, g^x, N_B, g^y)$ and if it is the same as hss , then derive atk and encrypt k using atk . We denote the ciphertext $C = \text{Enc}_{atk}(k)$ and assume AES-GCM256. Send C to the authentication server (via Pass 1).

Authentication Server: Forward C to Bob (via Pass 1).

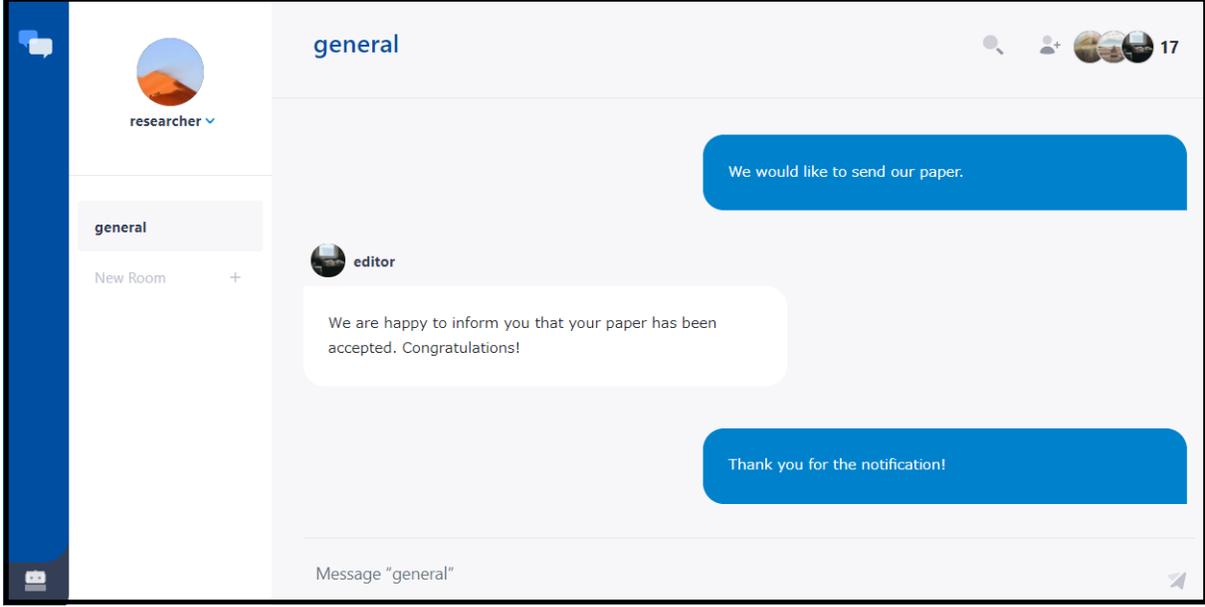


Figure 5: Chat System (Normal)

Bob: Derive atk , decrypt C using atk , and obtain k . Define $R_B = k \oplus \text{Hash}(\text{pw}'_B)$ and send R_B to the authentication server (via Pass 1).

Authentication Server: Preserve R_B with the user name Bob.

Here, R_A is chosen independently from k , pw_A , and pw'_A . Thus no information of them is revealed from R_A directly. Moreover, k is encrypted by atk and due to the security of SMKEX, only Alice and Bob know atk . Thus, no information of k is revealed from C . Finally, the authentication server knows R_A and R_B ; however, as in the storage system, the authentication server does not know pw'_A and pw'_B . Therefore, the authentication server cannot obtain k . Although Alice knows k , she does not know R_B because it is sent via a TLS communication between the authentication server and Bob. In other words, Alice cannot extract $\text{Hash}(\text{pw}'_B)$ from k . However, if Alice and the authentication server collude, then $\text{Hash}(\text{pw}'_B)$ can be extracted from k and R_B that allows they can observe Bob's storage and his chat messages sent in other room. Thus, we assume that the authentication server does not collude with any user.

Secure Chat. When a user posts a message to the application server, the message is encrypted automatically, and when a user displays a message, the message is decrypted automatically. Thus, users are not required to aware DSSE. The chat system is shown in Fig. 5, where the room name is general.

4 Performance Analysis

We employed AWS EC2 (t2.micro (vCPU1, 1GiB memory), OS: Ubuntu 20.04, CPU: Intel(R) Xeon(R) CPU E5-2676 v32.40GHz) as the authentication and application servers, and OS: Windows 10 Pro, CPU: Intel(R) Core™ i7-8565U CPU 1.80GHz as a user. We compared our system to a non-DSSE system. In this non-DSSE case, we employed a classical inverted index method as a searching

method for the storage system, and SELECT supported by PostgreSQL⁸ for the chat system.

Storage System. We used copyright-free Japanese books published by Aozora Bunko.⁹ For example, Botchan (Soseki Natsume) contains approximately 100,000 characters. When a le was uploaded, the runtime was 13.7 s in the non-DSSE case and 13.5 s in the DSSE case. We consider that the DSSE case was more efficient because indexes are generated in the non-DSSE case, whereas this procedure is not required in the DSSE case. When a keyword is searched, we gave the case when a keyword is found (Search Hit) in Table 1, and the case when a keyword is not found (Search does not Hit) in Table 2, respectively. Due to the AWS environment, it appears that computation resources are not always guaranteed; thus, there were fluctuations in run times; however, we found that the run time is generally linearly dependent on the number of files.

Table 1: Storage System: Search Hit (msec)

Cases \# Files	1	3	5	10
Non-DSSE(A)	54.0	55.5	62.5	53.5
DSSE(B)	82.5	89.0	112.0	103.5
(A)-(B)	28.5	33.5	49.5	50.0

Table 2: Storage System: Search does not Hit (msec)

Cases \# Files	1	3	5	10
Non-DSSE(A)	54.0	56.0	49.0	55.0
DSSE(B)	59.5	75.0	56.5	70.5
(A)-(B)	5.5	19.0	7.5	15.5

Chat System. We used Tweet data posted by NICT official publicity.¹⁰ When a message was posted, the running time was 35.0 ms in the non-DSSE case and 66.4 ms in the DSSE case. Although it becomes worse almost twice, it seems acceptable for practice application. When a keyword w searched, we gave the case when a keyword is found (Search Hit) in Table 3, and the case when a keyword is not found (Search does not Hit) in Table 4, respectively. Note that the difference (A)-(B) increased as number of messages increased due to the DSSE.

Table 3: Chat System: Search Hit (msec)

Case \# Messages	1	20	40	60	80
Non-DSSE(A)	35.6	43.8	41.0	61.2	54.8
DSSE(B)	65.7	114.6	129.4	158.2	211.4
(B)-(A)	30.1	70.8	88.4	97.0	156.6

⁸<https://www.postgresql.org/>

⁹<https://www.aozora.gr.jp/>

¹⁰https://twitter.com/NICT_Publicity

Table 4: Chat System: Search does not Hit (msec)

Case \ # Messages	1	20	40	60	80
Non-DSSE(A)	34.7	32.4	39.4	46.4	58.0
DSSE(B)	61.9	68.2	103.8	131.2	207.0
(B)-(A)	27.2	35.8	64.4	84.8	149.0

Acknowledgment: The authors thank Mr. Satoru Kanno (Infours Inc.), Dr. Yumi Sakemi (Infours Inc.), Prof. Tetsu Iwata (Nagoya University), and Prof. Shuichi Hirose (University of Fukui) for their invaluable comments and suggestions. This work was supported by JSPS KAKENHI Grant Numbers JP21K11897, 21H03441, 18K11293, and 18H05289.

References

- [1] Crypto-chat. <http://www.crypto-chat.com/>.
- [2] MeCab: Yet another part-of-speech and morphological analyzer. <https://sourceforge.net/projects/mecab/>.
- [3] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In *EUROCRYPT*, pages 129–158, 2019.
- [4] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. In *CRYPTO*, pages 248–277, 2020.
- [5] Ghouse Amjad, Sarvar Patel, Giuseppe Persiano, Kevin Yeo, and Moti Yung. Dynamic volume-hiding encrypted multi-maps with applications to searchable encryption. Cryptology ePrint Archive, Report 2021/765, 2021. <https://eprint.iacr.org/2021/765>.
- [6] Panagiotis Antonopoulos, Arvind Arasu, Kunal D. Singh, Ken Eguro, Nitish Gupta, Rajat Jain, Raghav Kaushik, Hanuma Kodavalla, Donald Kossmann, Nikolas Ogg, Ravi Ramamurthy, Jakub Szymaszek, Jeffrey Trimmer, Kapil Vaswani, Ramarathnam Venkatesan, and Mike Zwillig. Azure SQL database always encrypted. In *ACM SIGMOD*, pages 1511–1525, 2020.
- [7] Wei Bai, Michael Pearson, Patrick Gage Kelley, and Michelle L. Mazurek. Improving non-experts’ understanding of End-to-End encryption: An exploratory study. In *IEEE EuroS&P Workshops*, pages 210–219, 2020.
- [8] Laura Blackstone, Seny Kamara, and Tarik Moataz. Revisiting leakage abuse attacks. In *NDSS*, 2020.
- [9] Raphael Bost. $\sum\text{o}\varphi\text{o}\varsigma$: Forward secure searchable encryption. In *ACM CCS*, pages 1143–1154, 2016.
- [10] Raphaël Bost, Brice Minaud, and Olga Ohrimenko. Forward and backward private searchable encryption from constrained cryptographic primitives. In *ACM CCS*, pages 1465–1482, 2017.

- [11] David Cash, Joseph Jaeger, Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Catalin Rosu, and Michael Steiner. Dynamic searchable encryption in very-large databases: Data structures and implementation. In *NDSS*, 2014.
- [12] David Cash, Stanislaw Jarecki, Charanjit S. Jutla, Hugo Krawczyk, Marcel-Catalin Rosu, and Michael Steiner. Highly-scalable searchable symmetric encryption with support for boolean queries. In *CRYPTO*, pages 353–373, 2013.
- [13] David Cash, Ruth Ng, and Adam Rivkin. Improved structured encryption for SQL databases via hybrid indexing. In *Applied Cryptography and Network Security*, pages 480–510, 2021.
- [14] Javad Ghareh Chamani, Dimitrios Papadopoulos, Charalampos Papamanthou, and Rasool Jalili. New constructions for forward and backward private symmetric searchable encryption. In *ACM CCS*, pages 1038–1055, 2018.
- [15] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In *ACM CCS*, pages 1802–1819, 2018.
- [16] Katriel Cohn-Gordon, Cas J. F. Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. In *IEEE EuroS&P*, pages 451–466, 2017.
- [17] Sergiu Costea, Marios O. Choudary, Doru Gucea, Björn Tackmann, and Costin Raiciu. Secure opportunistic multipath key exchange. In *ACM CCS*, pages 2077–2094, 2018.
- [18] Reza Curtmola, Juan A. Garay, Seny Kamara, and Rafail Ostrovsky. Searchable symmetric encryption: improved definitions and efficient constructions. In *ACM CCS*, pages 79–88, 2006.
- [19] Saba Eskandarian and Matei Zaharia. ObliDB: Oblivious query processing for secure databases. *Proc. VLDB Endow.*, 13(2):169–183, 2019.
- [20] Mohammad Etemad, Alptekin Küpçü, Charalampos Papamanthou, and David Evans. Efficient dynamic searchable encryption with forward privacy. *Privacy Enhancing Technologies*, 2018(1):5–20, 2018.
- [21] Sky Faber, Stanislaw Jarecki, Hugo Krawczyk, Quan Nguyen, Marcel-Catalin Rosu, and Michael Steiner. Rich queries on encrypted data: Beyond exact matches. In *ESORICS*, pages 123–145, 2015.
- [22] Jens Groth. On the size of pairing-based non-interactive arguments. In *EUROCRYPT*, pages 305–326, 2016.
- [23] Paul Grubbs, Anurag Khandelwal, Marie-Sarah Lacharité, Lloyd Brown, Lucy Li, Rachit Agarwal, and Thomas Ristenpart. Pancake: Frequency smoothing for encrypted data stores. In *USENIX Security*, pages 2451–2468, 2020.
- [24] Paul Grubbs, Richard McPherson, Muhammad Naveed, Thomas Ristenpart, and Vitaly Shmatikov. Breaking web applications built on top of encrypted data. In *ACM CCS*, pages 1353–1364, 2016.
- [25] Hakan Hacigümüs, Balakrishna R. Iyer, Chen Li, and Sharad Mehrotra. Executing SQL over encrypted data in the database-service-provider model. In *ACM SIGMOD*, pages 216–227, 2002.

- [26] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. An efficient and generic construction for Signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable. In *Public-Key Cryptography*, pages 3–33, 2021.
- [27] Amir Herzberg and Hemi Leibowitz. Can johnny finally encrypt?: evaluating E2E-encryption in popular IM applications. In *ACM STAST*, pages 17–28, 2016.
- [28] Seny Kamara and Tarik Moataz. SQL on structurally-encrypted databases. In *ASIACRYPT*, pages 149–180, 2018.
- [29] Seny Kamara and Tarik Moataz. Computationally volume-hiding structured encryption. In *EUROCRYPT*, pages 183–213, 2019.
- [30] Seny Kamara and Charalampos Papamanthou. Parallel and dynamic searchable symmetric encryption. In *Financial Cryptography and Data Security*, pages 258–274, 2013.
- [31] Seny Kamara, Charalampos Papamanthou, and Tom Roeder. Dynamic searchable symmetric encryption. In *ACM CCS*, pages 965–976, 2012.
- [32] Jonathan Katz, Rafail Ostrovsky, and Moti Yung. Efficient password-authenticated key exchange using human-memorable passwords. In *EUROCRYPT*, pages 475–494, 2001.
- [33] Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O’Neill. Generic attacks on secure outsourced databases. In *ACM CCS*, pages 1329–1340, 2016.
- [34] Kee Sung Kim, Minkyu Kim, Dongsoo Lee, Je Hong Park, and Woo-Hwan Kim. Forward secure dynamic searchable symmetric encryption with efficient updates. In *ACM CCS*, pages 1449–1463, 2017.
- [35] Taehoon Kim, Joongun Park, Jaewook Woo, Seungheun Jeon, and Jaehyuk Huh. ShieldStore: Shielded in-memory key-value storage with SGX. In *ACM EuroSys*, pages 14:1–14:15, 2019.
- [36] H. Krawczyk and P. Eronen. HMAC-based extract-and-expand key derivation function (HKDF). <https://datatracker.ietf.org/doc/html/rfc5869>.
- [37] Russell W. F. Lai and Sherman S. M. Chow. Forward-secure searchable encryption on labeled bipartite graphs. In *Applied Cryptography and Network Security*, pages 478–497, 2017.
- [38] P. Leach, M. Mealling, and R. Salz. A Universally Unique Identifier (UUID) URN Namespace. <https://tools.ietf.org/html/rfc4122>, July 2005.
- [39] Moxie Marlinspike and Trevor Perrin. The X3DH key agreement protocol. <https://signal.org/docs/specifications/x3dh/>, November 2016.
- [40] Ian Miers and Payman Mohassel. IO-DSSE: scaling dynamic searchable encryption to millions of indexes by improving locality. In *NDSS*, 2017.
- [41] Pratyush Mishra, Rishabh Poddar, Jerry Chen, Alessandro Chiesa, and Raluca Ada Popa. Oblix: An efficient oblivious search index. In *IEEE Symposium on Security and Privacy*, pages 279–296, 2018.
- [42] Muhammad Naveed, Manoj Prabhakaran, and Carl A. Gunter. Dynamic searchable encryption via blind storage. In *IEEE Symposium on Security and Privacy*, pages 639–654, 2014.

- [43] Sarvar Patel, Giuseppe Persiano, Kevin Yeo, and Moti Yung. Mitigating leakage in secure cloud-hosted data structures: Volume-hiding for multi-maps via hashing. In *ACM CCS*, pages 79–93, 2019.
- [44] Raluca Ada Popa, Catherine M. S. Redfield, Nikolai Zeldovich, and Hari Balakrishnan. CryptDB: protecting confidentiality with encrypted query processing. In *ACM SOSP*, pages 85–100, 2011.
- [45] Raluca Ada Popa, Catherine M. S. Redfield, Nikolai Zeldovich, and Hari Balakrishnan. CryptDB: processing queries on an encrypted database. *Commun. ACM*, 55(9):103–111, 2012.
- [46] Raluca Ada Popa, Emily Stark, Steven Valdez, Jonas Helfer, Nikolai Zeldovich, and Hari Balakrishnan. Building web applications on top of encrypted data using Mylar. In *USENIX NSDI*, pages 157–172, 2014.
- [47] Raluca Ada Popa and Nikolai Zeldovich. Multi-key searchable encryption. *IACR Cryptol. ePrint Arch.*, 2013:508, 2013.
- [48] Raluca Ada Popa, Nikolai Zeldovich, and Hari Balakrishnan. Guidelines for using the CryptDB system securely. *IACR Cryptol. ePrint Arch.*, 2015:979, 2015.
- [49] Kui Ren, Yu Guo, Jiaqi Li, Xiaohua Jia, Cong Wang, Yajin Zhou, Sheng Wang, Ning Cao, and Feifei Li. HybriDX: New hybrid index for volume-hiding range queries in data outsourcing services. In *IEEE ICDCS*, pages 23–33, 2020.
- [50] Toshiya Shibata and Kazuki Yoneyama. Universally composable forward secure dynamic searchable symmetric encryption. In *ACM ASIA Public-Key Cryptography Workshop*, pages 41–50, 2021.
- [51] Dawn Xiaodong Song, David A. Wagner, and Adrian Perrig. Practical techniques for searches on encrypted data. In *IEEE Symposium on Security and Privacy*, pages 44–55, 2000.
- [52] Emil Stefanov, Charalampos Papamanthou, and Elaine Shi. Practical dynamic searchable encryption with small leakage. In *NDSS*, 2014.
- [53] Christian Stransky, Dominik Wermke, Johanna Schrader, Nicolas Huaman, Yasemin Acar, Anna Lena Fehlhauer, Miranda Wei, Blase Ur, and Sascha Fahl. On the limited impact of visualizing encryption: Perceptions of E2E messaging security. In *Symposium on Usable Privacy and Security*, 2021.
- [54] Shifeng Sun, Xingliang Yuan, Joseph K. Liu, Ron Steinfeld, Amin Sakzad, Viet Vo, and Surya Nepal. Practical backward-secure searchable encryption from symmetric puncturable encryption. In *ACM CCS*, pages 763–780, 2018.
- [55] Stephen Tu, M. Frans Kaashoek, Samuel Madden, and Nikolai Zeldovich. Processing analytical queries over encrypted data. *Proc. VLDB Endow.*, 6(5):289–300, 2013.
- [56] Viet Vo, Shangqi Lai, Xingliang Yuan, Surya Nepal, and Joseph K. Liu. Towards efficient and strong backward private searchable encryption with secure enclaves. In *Applied Cryptography and Network Security*, pages 50–75, 2021.

- [57] Yohei Watanabe, Takeshi Nakai, Kazuma Ohara, Takuya Nojima, Yexuan Liu, Mitsugu Iwamoto, and Kazuo Ohta. How to make a secure index for searchable symmetric encryption, revisited. Cryptology ePrint Archive, Report 2021/948, 2021. <https://eprint.iacr.org/2021/948>.
- [58] Alma Whitten and J. Doug Tygar. Why johnny can't encrypt: A usability evaluation of PGP 5.0. In *USENIX Security Symposium*, 1999.
- [59] Kazuki Yoneyama and Shogo Kimura. Verifiable and forward secure dynamic searchable symmetric encryption with storage efficiency. In *ICICS*, pages 489–501, 2017.
- [60] Yongjun Zhao, Huaxiong Wang, and Kwok-Yan Lam. Volume-hiding dynamic searchable symmetric encryption with forward and backward privacy. Cryptology ePrint Archive, Report 2021/786, 2021. <https://eprint.iacr.org/2021/786>.

Appendix: Concierge Functionality

Here, we introduce the concierge functionality, which is used to view DSSE-related data processing. In the storage system (concierge mode), when a file is uploaded,¹¹ the file identifier id and number of keywords (which are extracted by MeCab from the file name and file content) are displayed (Fig. 6). As shown in Fig. 3, file names are typically displayed. In concierge mode, the encrypted file names are displayed when they are moused over, which shows the application server's point of view (Fig. 7). When the file content is moused over, the corresponding ciphertext is displayed, which shows the application server's point of view (Fig. 8). In addition, when a keyword is searched, pairs of the keyword and a file identifier are displayed (Fig. 9). Here, the keyword to be searched is NIZK. When a keyword and file identifier pair is moused over, the corresponding trapdoor, i.e., $\pi_k(\text{NIZK}, \text{id})$, is displayed, which shows the application server's point of view (Fig. 10). Note that the chat system also supports concierge mode. We need to evaluate our concierge functionality from the usable-security point of view, e.g., [7, 27, 53, 58] and we left it as a future work of this paper.



Figure 6: Storage System (Concierge Mode: File Uploading)

¹¹In this case, the file name is “Cryptology ePrint Archive Report 2019609 CPA-to-CCA Transformation for KDM Security.pdf. Note that the second “-” is removed, and “CCA” is displayed immediately after “to” because a set of keywords is defined here.

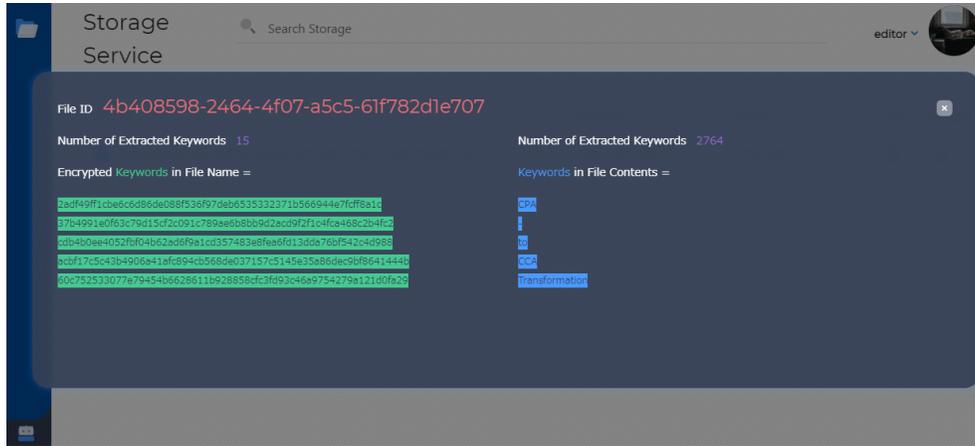


Figure 7: Storage System (Concierge Mode: File Name)

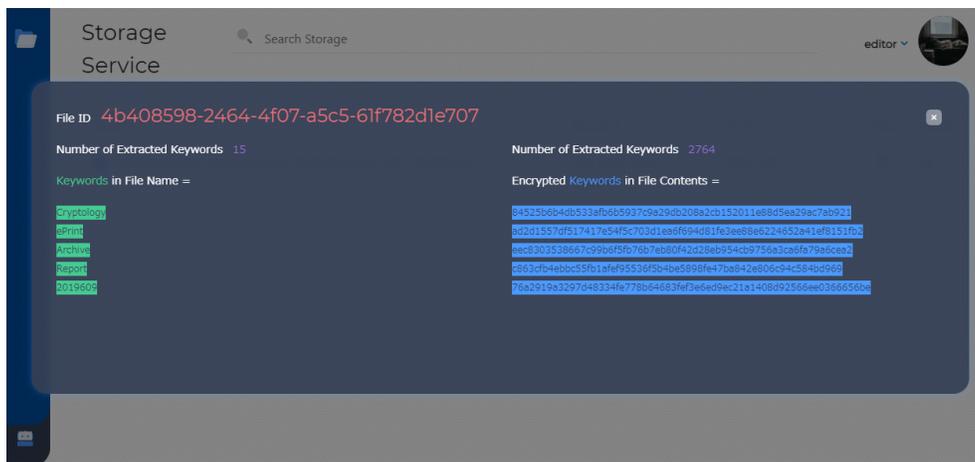


Figure 8: Storage System (Concierge Mode: Ciphertext)

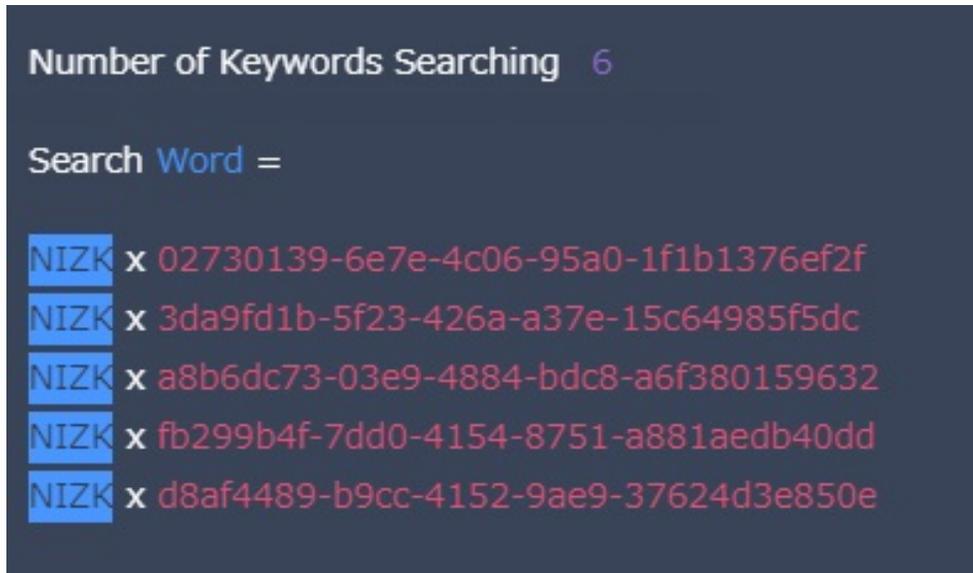


Figure 9: Storage System (Concierge Mode: Searching)

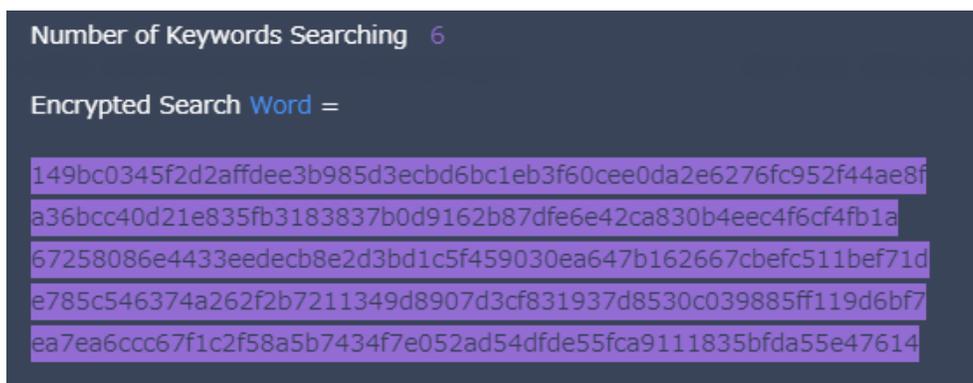


Figure 10: Storage System (Concierge Mode: Trapdoor)