



# **Findex**

Full Documentation

Version 1.0

Date: 16/12/2022



## Table Of Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
<b>2</b>	<b>Example - The Searchable Directory</b>	<b>4</b>
<b>3</b>	<b>Full Process</b>	<b>6</b>
3.1	Overview . . . . .	6
3.1.1	Symmetric Searchable Encryption . . . . .	6
3.1.2	Notations . . . . .	6
3.1.3	Index Tables . . . . .	7
3.1.4	Search Query . . . . .	8
3.2	Chain Table . . . . .	8
3.2.1	Chain Table Value . . . . .	9
3.2.2	Chain Table UID . . . . .	9
3.2.3	Size . . . . .	10
3.3	Entry Table . . . . .	11
3.3.1	Size . . . . .	11
3.4	Search Query Process . . . . .	12
<b>4</b>	<b>Update Process</b>	<b>15</b>
4.1	Overview . . . . .	15
4.1.1	Impact on the Efficiency . . . . .	15
4.2	Change in DB Table . . . . .	15
4.2.1	Delete Line . . . . .	15
4.2.2	Add Line . . . . .	16
4.2.3	Modify Line . . . . .	16
4.3	Change in Index Tables . . . . .	16
4.3.1	Delete Keyword . . . . .	16
4.3.2	Add Keyword . . . . .	17
4.4	ReIndexing . . . . .	17
4.4.1	Why . . . . .	17
4.4.2	How . . . . .	17
<b>5</b>	<b>Security</b>	<b>20</b>
5.1	Key . . . . .	20
5.2	Server Storage . . . . .	20
5.3	Client - Server Communication . . . . .	20
5.4	ReIndexing . . . . .	21



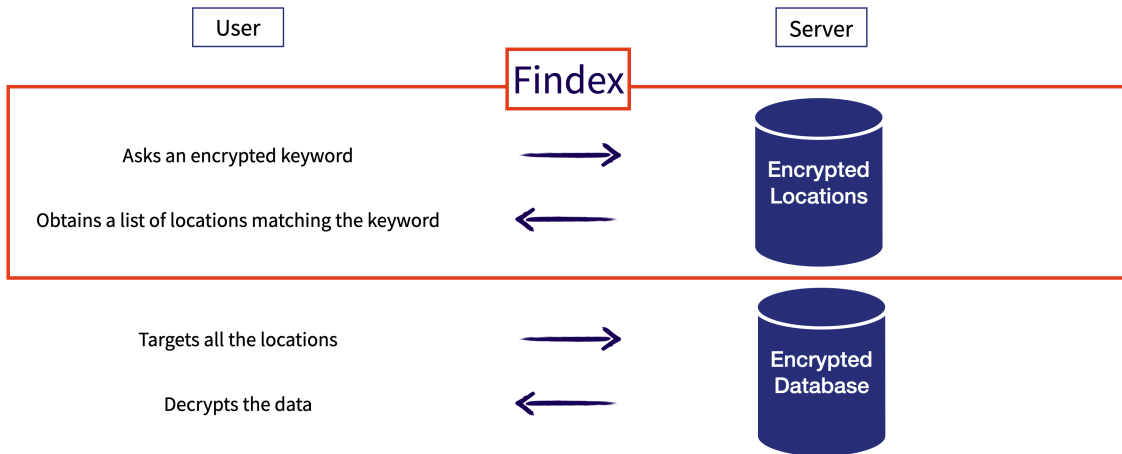
5.5	Dynamic Symmetric Searchable Encryption . . . . .	21
<b>6</b>	<b>Appendix</b>	<b>23</b>
6.1	Cryptographic Algorithms . . . . .	23
6.2	Keys . . . . .	23
	<b>References</b>	<b>24</b>



# 1 Introduction

**Findex** is a part of Cloudproof Encryption and helps to securely make search queries on outsourced encrypted data.

This documentation shows its running and explains the cryptographic details.






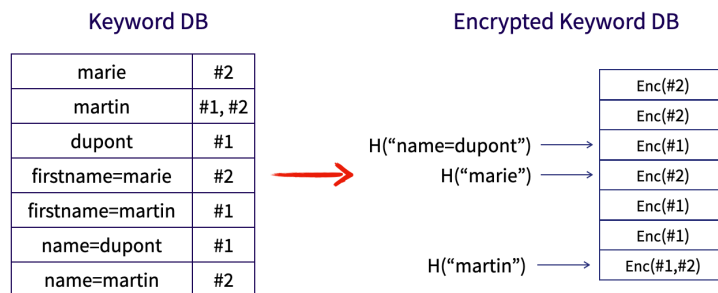
## 2 Example - The Searchable Directory

To understand the general idea behind Findex, let us assume one wants to outsource a directory while being able to securely make search queries on it.

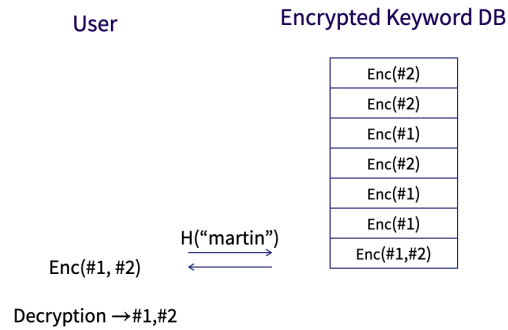
Here the directory is composed of two users. The first step consists of building a Keyword Database:



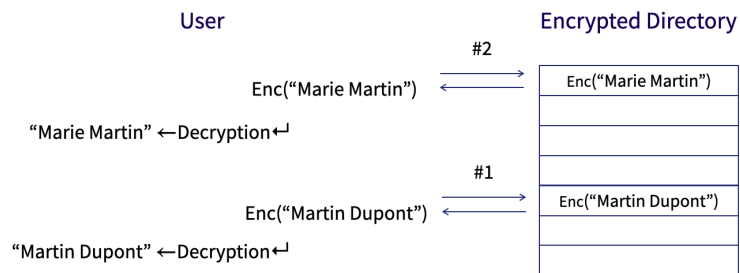
Then, the table can be encrypted and sent by an authenticated administrator to a first server. The lines are randomized not to be able to retrieve a keyword from a position in the encrypted database:



Now, the encrypted keyword database exists, a user can build requests. The user hashes the keyword "Martin" and asks for it to the server having the encrypted keyword database:



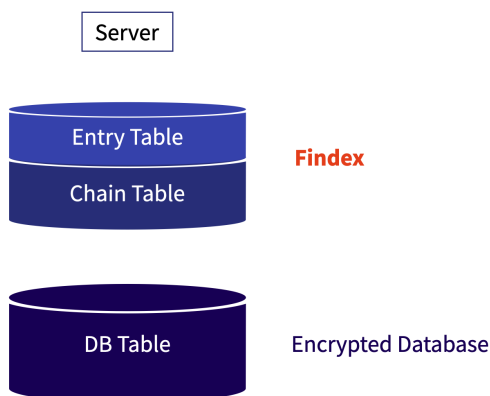
The user receives an encrypted message containing the location of all the matching queries: #1 , #2.  
The user can then requests the server hosting the encrypted directory for these two locations:





### 3 Full Process

#### 3.1 Overview



Findex relies on two server-side tables, *Entry Table* and *Chain Table*, to solve the following search problem:

How to *securely* recover the UIDs of DB Table to obtain the matching lines from a given keyword?

This solution is on top of an encrypted database, for consistency called DB Table, that actually stores the content to be requested.

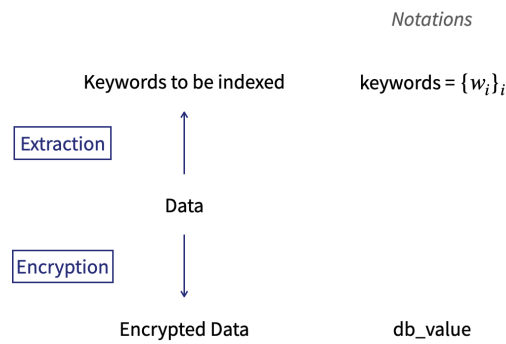
##### 3.1.1 Symmetric Searchable Encryption

To make efficient search queries on an untrusted cloud server, one needs to use an advanced cryptographic primitive called Symmetric Searchable Encryption (SSE). The security of SSE offers precise guarantees regarding the privacy of the user’s data and queries with respect to the host server.

##### 3.1.2 Notations

We assume that each line of DB Table is encrypted but at the time of encryption, some keywords  $\{w_i\}_i$  have been extracted to be stored with **Findex**.

In one line many keywords can be extracted and the same keyword can be extracted from several lines.

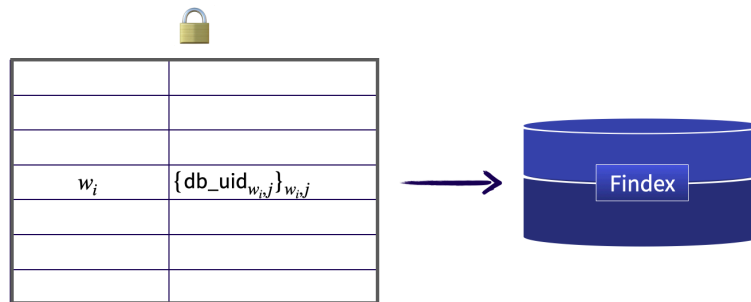


In this document, *key* refers to a cryptographic key and the databases are represented by a list of  $(uid_i, value_i)$ .

Hence,

- $(db\_uid_i, db\_value_i) = (uid_i, value_i)$  of DB Table,
- $(entry\_uid_i, entry\_value_i) = (uid_i, value_i)$  of Entry Table,
- $(chain\_uid_i, chain\_value_i) = (uid_i, value_i)$  of Chain Table.

### 3.1.3 Index Tables



**Figure 1:** Findex Input

After the extraction, each keyword  $w_i$  can be associated to a list  $DB[w_i] = \{db\_uid_{w_i,j}\}_j$  of db\_uids matching the keyword. Let  $DB = \{DB[w_i]\}_{w_i}$ .

**Chain Table:** *securely* stores DB, thus all the lists  $DB[w_i]$ .






*Example:* The keyword “Martin” is present in lines 3, 5, and 10 of a cleartext directory. These lines correspond to the db\_uids: db\_uid<sub>a</sub>, db\_uid<sub>b</sub>, and db\_uid<sub>c</sub> of the DB Table (i.e. the encrypted directory). The Chain Table will securely store {db\_uid<sub>a</sub>, db\_uid<sub>b</sub>, db\_uid<sub>c</sub>}.

**Entry Table:** provides the mandatory values to access the Chain Table.

### 3.1.4 Search Query



**Figure 2:** Search Query

**Search Query:** takes as input a bulk of keywords  $\{w_i\}_i$  and outputs the bulk of the lists  $\{DB[w_i]\}_i$ .

Findex considers search queries restricted to a single keyword  $w$ . To handle queries with several keywords (OR of keywords), several requests are made to the server, and possibly a combination can be done on the client’s side to deal with ANDs.

### 3.2 Chain Table



Let us see the content of the Chain Table.



### 3.2.1 Chain Table Value

We denote by  $DB[w_i]$  the list of the UIDs<sup>1</sup> of DB Table matching the keyword  $w_i$  and  $|DB[w_i]|$  the number of such elements.

To hide  $|DB[w_i]|^2$ ,  $DB[w_i]$  will be divided into blocks of equal size  $B^3$ :

$$DB[w_i] = \{DB[w_i]_1, \dots, DB[w_i]_{L_i}\}$$

Potentially, the last block is not full and will be padded.

Then, all the blocks are encrypted with a symmetric encryption scheme<sup>4</sup> under a key  $K_{w_i, \text{value}}$ <sup>5</sup> that is specific to the keyword  $w_i$ :

$$\text{Enc}_{\text{Sym}}(K_{w_i, \text{value}}, DB[w_i]_1), \dots, \text{Enc}_{\text{Sym}}(K_{w_i, \text{value}}, DB[w_i]_{L_i})$$

### 3.2.2 Chain Table UID

Before storing the chain values in the Chain Table, one needs to create their respective UIDs.

It would be possible to store all of them with random numbers but it would imply transmitting later all these chain\_uids to the user, as they correspond to the matching entries of its search request. The number of these UIDs is exactly the number  $L_i$  of blocks, and thus, is strictly smaller than the number of UIDs in the original list  $DB[w_i]$ . However, for the same reason that the length  $|DB[w_i]|$  must be hidden,  $L_i$  must be hidden too. The solution would be to repeat the same process by adding a new Index Table, and so on, but it is not practical.

To avoid that, our solution will exploit linked lists to create the UIDs of Chain Table:

- from the key  $K_{w_i, \text{uid}}$ , one can compute:  $\mathcal{H}(K_{w_i, \text{uid}}, \mathcal{H}(w_i)) \rightarrow \text{chain\_uid}_1$
- then, from  $\text{chain\_uid}_1$  and the key  $K_{w_i, \text{uid}}$ , one can compute:  $\mathcal{H}(K_{w_i, \text{uid}}, \text{chain\_uid}_1) \rightarrow \text{chain\_uid}_2$
- then, from  $\text{chain\_uid}_2$  and the key  $K_{w_i, \text{uid}}$ , one can compute:  $\mathcal{H}(K_{w_i, \text{uid}}, \text{chain\_uid}_2) \rightarrow \text{chain\_uid}_3$

---

<sup>1</sup>In the implementation,  $DB[w_i]$  is a list of locations where a location can be a UID or something else such as another keyword. The size of a location is not fixed but, in this case, the location is first divided into blocks of fixed size.

<sup>2</sup>See in Section Security to learn more about the importance of not only hiding the keywords and the db\_uids but also the length of the result.

<sup>3</sup>The size  $B$  is constant for a given Findex index. However, it may vary between use cases since the average length of the searchable keywords can vary.

<sup>4</sup>See Appendix for the description of the symmetric encryption scheme

<sup>5</sup>The two keys  $K_{w_i, \text{uid}}$  and  $K_{w_i, \text{value}}$  are derived from  $K_{w_i}$  randomly chosen by the Index Authority.



- and so on, until the  $L_i$  values have been generated.

In the end, instead of having  $L_i$  values to transmit to the user, we only have the last value of the linked list to know the stop criterion (and the key  $K_{w_i}$  to be able to derive  $K_{w_i,uid}$  and  $K_{w_i,value}$ ). This will be the goal of the Entry Table.

*Example:* If  $L_i = 3$ , the Chain Table could be:

UID	Value
...	...
chain_uid <sub>2</sub>	Enc <sub>Sym</sub> ( $K_{w_i,value}$ , DB[ $w_i$ ] <sub>2</sub> )
...	...
chain_uid <sub>1</sub>	Enc <sub>Sym</sub> ( $K_{w_i,value}$ , DB[ $w_i$ ] <sub>1</sub> )
...	...
chain_uid <sub>3</sub>	Enc <sub>Sym</sub> ( $K_{w_i,value}$ , DB[ $w_i$ ] <sub>3</sub> )
...	...

with in the other lines, blocks of  $B$  DB Table UIDs matching other keywords.

### 3.2.3 Size

About the size of Chain Table,

- the number of lines depends on the number of searchable keywords *and* on the number of lines needed to store all the locations indexed by these keywords
- in our implementation, a line is composed of:
  - **UID**: 32 bytes
  - **Value**:

	Nonce	AES-GCM encrypted data	MAC
Size (bytes)	12	$32 \times B$	16



### 3.3 Entry Table



The role of the Entry Table is to store for each  $w_i$ : the last value of its associated linked list, the key  $K_{w_i}$  used in this list and  $\mathcal{H}(w_i)$ .

First, for each keyword, a UID is computed from a common secret key  $K_{uid}$ . Then the data is symmetrically encrypted under a common secret key  $K_{value}$ .

*Example:* For a keyword  $w_i$ , the entry\_uid is computed:

$$\text{entry\_uid}_i = \mathcal{H}(K_{uid}, \mathcal{H}(w_i), T)$$

Then, if  $\text{chain\_uid}_3$  is the last value of the linked list and  $K_{w_i}$  is the key used in Chain Table,  $(\text{chain\_uid}_3, K_{w_i}, \mathcal{H}(w_i))$  is encrypted under  $K_{value}$ :

$$\text{entry\_value}_i = \text{Enc}_{\text{Sym}}(K_{value}, (\text{chain\_uid}_3, K_{w_i}, \mathcal{H}(w_i)))$$

*Note:* The two keys  $K_{uid}$  and  $K_{value}$  are derived from a secret key  $K$  known by all the authorized entities (i.e. the Index Authority and all the users).

Finally, the Entry Table looks like this:

UID	Value
...	...
$\mathcal{H}(K_{uid}, \mathcal{H}(w_i), T)$	$\text{Enc}_{\text{Sym}}(K_{value}, (\text{chain\_uid}_{L_i}, K_{w_i}, \mathcal{H}(w_i)))$
...	...

#### 3.3.1 Size

About the size of the Entry Table,

- the number of lines only depends on the number of searchable keywords
- in our implementation, a line is composed of:



- UID: 32 bytes
- Value:

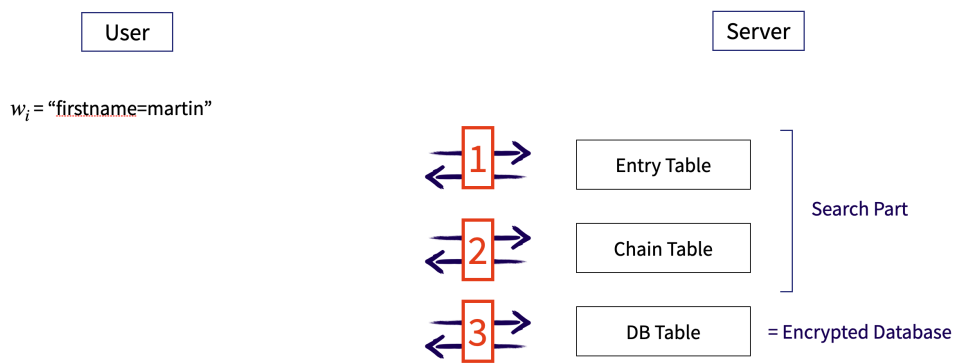
	Nonce	AES-GCM encrypted data	MAC
Size (bytes)	12	$UID.len + key.len + hash.len$	16

where:

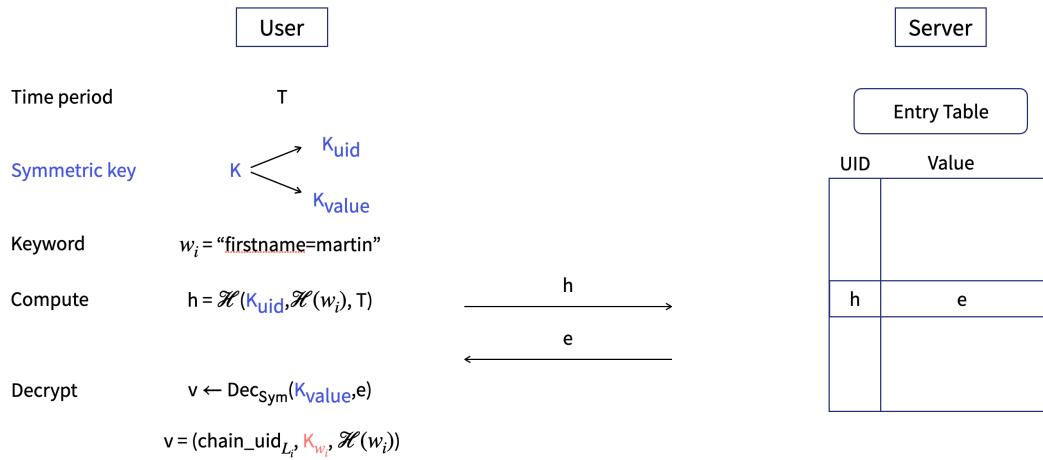
- $UID.len$ : 32 bytes
- $key.len$ : 16 bytes
- $hash.len$ : 32 bytes

### 3.4 Search Query Process

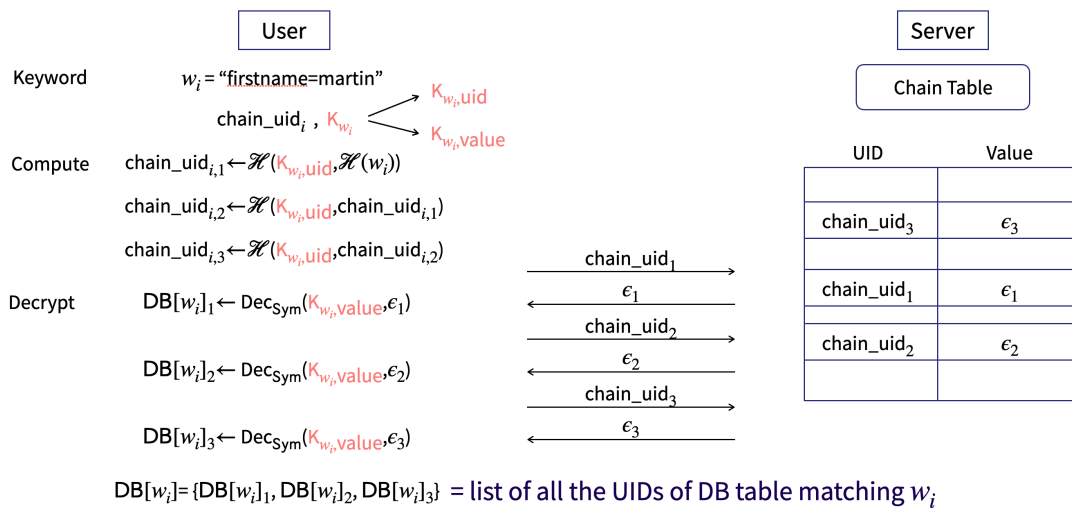
We recall the scenario: a user wants to make a search query on an encrypted database hosted in an external (untrusted) server. Now, we will see the interactions between this user and the server. Basically, the search query process travels along the tables in the opposite order as the one described before.



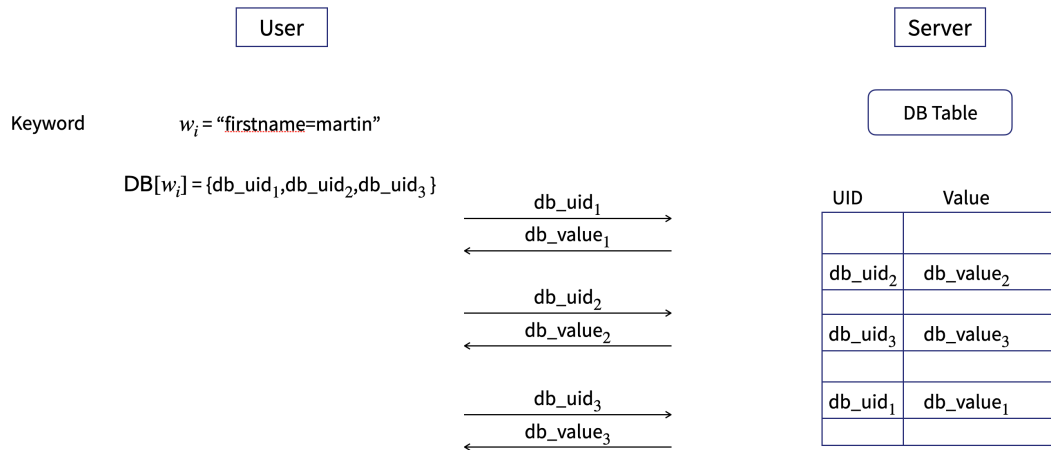
In the figure below, the user queries the Entry Table:



Then, the user computes  $\text{chain\_uid}_{i+1}$  until  $\text{chain\_uid}_{L_i}$  is found and queries the Chain Table:



Finally, the user queries the DB Table:



and tries to decrypt the results.



## 4 Update Process

### 4.1 Overview

Now, one can explain how to update the indexes to update a search request.

Overall, the changes can be:

- In DB Table:
  - to delete a line,
  - to add a new line,
  - to modify a line (the content or the right access).
- In Index Tables:
  - to delete already indexed keywords,
  - to add new keywords for data already encrypted in DB Table.

*Remark:* The changes in the Index Tables are supposed to be made by the Index Authority.

#### 4.1.1 Impact on the Efficiency

If you want to apply changes in order to improve efficiency, here are some remarks:

- **Changing the content of DB Table** can increase the size of the index tables (as new keywords may be indexed).
- **Reducing the number of searchable keywords** also reduces the size of the Index Tables but does not change the size of the DB Table.
- **Increasing the number of searchable keywords** does not affect the efficiency of a search request: the efficiency of a search request is *independent* of the total number of keywords.

## 4.2 Change in DB Table

### 4.2.1 Delete Line

Given the  $db\_uid_i$  of the line that must be deleted, delete the line  $db\_line_i$  in the DB Table.

The index tables do not change and thus, continue to refer to deleted contents. They are cleaned during a reindexing phase<sup>6</sup>.

---

<sup>6</sup>See Section ReIndexing





#### 4.2.2 Add Line

The addition of a line simply follows the description presented in the full process:

- extraction of the keywords,
- encryption of the line and insertion in DB Table,
- encryption of the chain of keywords with insertions in the relevant index tables.

Here, the three tables change.

#### 4.2.3 Modify Line

In our solution, a modification of the content (or possibly of the right access) will simply be the deletion of the old line followed by the addition of the new one containing the change.

A change will be:

- the deletion of the corresponding encrypted line in DB Table,
- the addition of a new line in the DB Table containing the change one wanted to apply.

Hence, it requires knowing which db\_uid must be modified for the deletion.

Here, the three tables change.

### 4.3 Change in Index Tables

In our solution, the choice of the indexed terms is completely free, they can be extracted without any structure as well as coming from the database structure. For example, the first two columns can be indexed and thus be 'searchable' while the other columns are not. Hence, it could make sense to modify the set of keywords even without changing the content of the encrypted database.

#### 4.3.1 Delete Keyword

It is possible to delete an indexed keyword without changing the DB Table by requesting the keyword and deleting all the corresponding lines in both Index Entry Table and Index Chain Table.

To delete an entire column of the searchable set, one needs to delete all the keywords present in it.



### 4.3.2 Add Keyword

To add new types of keywords (for example to index a new column of the database):

- For all the new entries (of DB Table), one simply considers all the keywords that must be indexed meaning the old ones and the new ones.
- For the already existing entries (of DB Table), all the concerned lines must be re-added.

## 4.4 ReIndexing

### 4.4.1 Why

With the changes presented in the previous sections, the Entry Table and Chain Table grow indefinitely. After the deletion of lines in the DB Table, some search requests still answer the deleted lines. Also, the server can link the modified lines to the changes.

The *ReIndexing* reduces the size of the index tables and blurs the information an attacker may have learned.

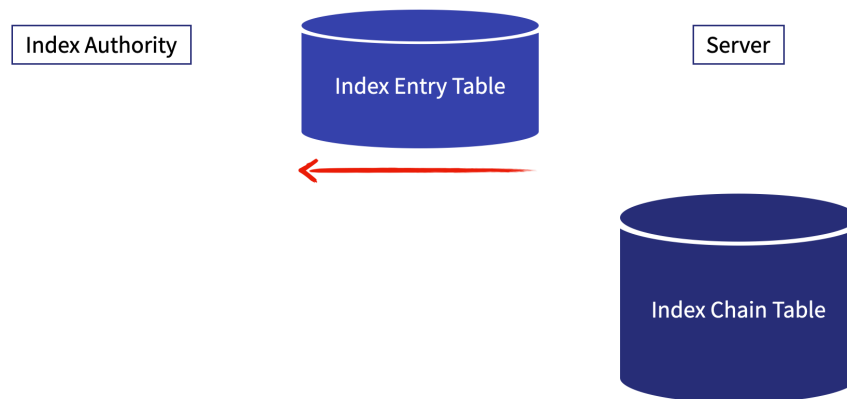
### 4.4.2 How

A reindexing:

- completely changes the Entry Table,
- cleans the search results of a fraction of keywords.

### Entry Table

The first step of the Index Authority is to download the entire *Entry Table* and to delete all the out-sourced existing lines:



Notation: -  $old_{entry\_table} = \{(entry\_uid_i, entry\_value_i)\}$  (the downloaded Entry Table) -  
 $entry\_uid_i = \mathcal{H}(K_{uid}, \mathcal{H}(w_i), T)$  -  $entry\_value_i = Enc_{Sym}(K_{value}, (chain\_uid_{L_i}, K_{w_i}, \mathcal{H}(w_i)))$   
 (AES-256-GCM encryption with nonce<sub>*i*</sub>) -  $new_{entry\_table} = \{(entry\_uid'_i, entry\_value'_i)\}$

Then, the Index Authority increments the (public) label  $T \rightarrow T'$  and, for all the Entry Table lines:

- computes  $entry\_uid'_i = \mathcal{H}(K'_{uid}, \mathcal{H}(w_i), T')$ ,
- changes nonce<sub>*i*</sub> used in the AES-256-GCM encryption into nonce'<sub>*i*</sub> to create  $entry\_value'_i$  an AES-256-GCM encryption with nonce'<sub>*i*</sub>.

Hence,  $entry\_value'_i$  is a randomization of  $entry\_value_i$ .

### Cleaning

In the second step, the Index Authority will clean the search part for a fraction of keywords.

For that, the Index Authority:

- implicitly chooses a fraction  $X$  of the keywords (of size  $x^7$ ) by choosing a random fraction of the Entry Table and decrypting it,
- for each of them,
  - makes the search request (i.e. obtains the list of all the positions matching the keywords) and deletes all the corresponding outsourced lines,
  - gets the result in the DB Tables (i.e. some of the positions will no longer exist),
  - for each non-existing value, deletes the db\_uid from the non-encrypted list of positions matching the keyword.

<sup>7</sup> $x = \lceil n \cdot (\log n + 0.58) \rceil / t$  is the required number of keywords per reindexing needed to reach the full set of keywords (#keywords =  $n$ ) in  $t$  reindexing phases. (cf. Coupon collector's problem).

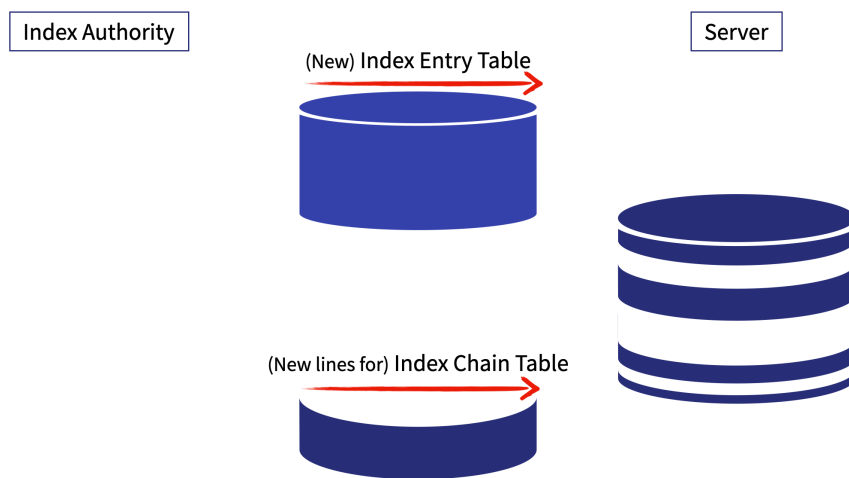


For each modified list of results, the Index Authority generates a completely new encryption of the indexes (i.e. create a new linked list of UID with the encryption of the db\_uids). This creates a new ephemeral key  $K_{w_i} \rightarrow K'_{w_i}$  and a new UID value  $chain\_uid_i \rightarrow chain\_uid'_i$  to be stored in Entry Table to be able to reconstruct the linked list during a request.

*Remark:* We presented the two steps of the reindexing one by one for readability but optimizations are made in the code.

Finally, the Index Authority can push all the computed lines in a *randomized* order:

- inserts the new Entry Table,
- insert the new lines of the Chain Table.





## 5 Security

Findex solves the following search problem:

How to *securely* recover the UIDs of DB Table to obtain the matching lines from a given keyword?

In this section, we explore the security details of Findex.

### 5.1 Key

The key  $K$  in Findex allows the authorized users to make search queries and the Index Authority to make update queries.

There is no security guarantee if the key is stolen. However, the key can be changed and distributed to only a subset of authorized users to revoke some of them. In that case, the content of the tables must be reencrypted.

### 5.2 Server Storage

The server stores two tables: Entry Table and Chain Table. Each table is composed of several lines containing (uid, value). The uid looks like a random number as it is the output of a Hash Function  $\mathcal{H}$  while value is the symmetric encryption of some elements. Hence, an adversary receiving a copy of the Entry Table and/or Chain Table cannot learn information on the data.

### 5.3 Client - Server Communication

The requests consist of hashed values and the answers to encrypted messages. Hence, for an adversary, the learnable information is limited to the frequencies of the requests. In particular, the server learns if a user asks two times the same keyword.

*Example - The Searchable Directory:* If **Dupont** is a frequent family name, the number of results matching the keyword will be large. On the contrary, the number of matching results of an uncommon family name will be small. If the attacker can see the length of the results, it can exploit the frequency of the keywords to retrieve the search query of an honest user which must be forbidden.

Moreover, if there are not enough requests for different keywords then, the server can “see” the interactions related to a keyword. This is true for the three tables: the Entry Table, the Chain Table and the DB Table.

One way to avoid that would be to use Oblivious Random Access Machine (ORAM) but it is not practical because not efficient. Hence, to avoid that, the simplest way is to generate fake requests to scramble communications.

## 5.4 ReIndexing

To make statistical analyses even more complicated, reindexing operations completely change the Entry Table and a fraction of the Chain Table.

## 5.5 Dynamic Symmetric Searchable Encryption

The classical version of a Symmetric Searchable Encryption (SSE) scheme considers a static database meaning the outsourced data can not be updated or new records can not be sent later. To overcome this issue a Dynamic version has been designed in [1] and security formalized in [2]. Such a scheme handles dynamic file collection.

**Findex** is a *Dynamic Symmetric Searchable Encryption*<sup>8</sup> scheme meaning one can define a triple (Setup, Search, Update) consisting of one algorithm and two protocols (between a client and a server) as defined below:

*Notation:* Let  $DB = \{(w_i, DB[w_i])\}_i$  be the (non-encrypted) table associating each keyword  $w_i$  to its matching results<sup>9</sup>.

- $Setup(DB) \rightarrow (EDB, K)$  : takes as input DB and outputs EDB an encrypted version of DB together with  $K$  a secret key.
- $Search(K, q; EDB) = (SearchClient(K, q), SearchServer(EDB))$  : is a protocol between a client with input the secret key  $K$  and a search query  $q$  and a server with input the encrypted table EDB. It outputs a list  $R$  of results to the client.
- $Update(K, q; EDB) = (UpdateClient(K, q), UpdateServer(EDB))$  : is a protocol between a client with input the secret key  $K$  and an update query  $q = (op, in)$  where the operation can be add or delete and in is the content to be added or deleted. It outputs a new  $EDB'$  to the server.

Findex considers search queries restricted to a *single* keyword  $w$ <sup>10</sup>.

An SSE scheme is said to be *correct* if the search protocol returns the correct result for every query:  $\forall w_i, R \leftarrow Search(K, w_i; EDB)$  is equal to  $DB[w_i]$ .

### Forward Secrecy

---

<sup>8</sup>See [3] for a thesis with a good introduction to SSE and the formal security definitions.

<sup>9</sup>See Overview for a presentation of  $DB[w_i]$ .

<sup>10</sup>Few SSE schemes deal with multiple keyword queries.



In Symmetric Searchable Encryption, a security notion called Forward Secrecy ([4],[5]) guarantees that updates do not reveal any information a priori about the modifications they carry out.

In our solution, there is no Forward Secrecy but it is possible to scramble the information during changes on the Index Tables by adding fake modifications.



## 6 Appendix

### 6.1 Cryptographic Algorithms

#### Hash Function

- $\mathcal{H}(\text{key}, m)$ : Message Authentication Code of  $m$  under the key  $\text{key}$ .

The implementation uses the KMAC128 scheme with a key of size 32 bytes.

#### Symmetric Scheme

- $\text{Enc}_{\text{Sym}}(\text{key}, m)$ : Symmetric Encryption of  $m$  under the key  $\text{key}$ .
- $\text{Dec}_{\text{Sym}}(\text{key}, m)$ : Symmetric Decryption of  $m$  under the key  $\text{key}$ . This algorithm “reverses” the  $\text{Enc}_{\text{Sym}}(\text{key}, m)$  function.

The implementation uses the AES256-GCM scheme with a key of size 32 bytes.

### 6.2 Keys

Key	Size (bytes)	Obtained from	Used in	Known by
$K$	16	random		All
$K_{\text{uid}}$	32	derivation of $K$	KMAC128	
$K_{\text{value}}$	32	derivation of $K$	AES256-GCM	
$K_{w_i}$	16	random		
$K_{w_i, \text{uid}}$	32	derivation of $K_{w_i}$	KMAC128	
$K_{w_i, \text{value}}$	32	derivation of $K_{w_i}$	AES256-GCM	

where ‘All’ means all the authorized users and the Index Authority.





## References

- [1] S. Kamara and C. Papamanthou, 'Parallel and dynamic searchable symmetric encryption', in *Financial cryptography and data security*, 2013, pp. 258–274.
- [2] S. J. D. Cash J. Jaeger, 'Dynamic searchable encryption in very-large databases: Data structures and implementation.', 2014.
- [3] R. Bost, 'Searchable encryption new constructions of encrypted databases', 2018, [Online]. Available: [https://raphael.bost.fyi/phd\\_docs/R\\_BOST\\_PhD\\_Thesis.pdf](https://raphael.bost.fyi/phd_docs/R_BOST_PhD_Thesis.pdf).
- [4] E. S. Emil Stefanov Charalampos Papamanthou, 'Practical dynamic searchable encryption with small leakage', 2014, [Online]. Available: <https://www.ndss-symposium.org/ndss2014/programme/practical-dynamic-searchable-encryption-small-leakage/>.
- [5] R. Bost, 'ΣΟΦΟΣ: Forward secure searchable encryption', in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, 2016, pp. 1143–1154, doi: 10.1145/2976749.2978303.