

Securing Confidential Data For Distributed Software Development Teams: Encrypted Container File

Tobias J. Bauer 

Fraunhofer Institute for Applied and Integrated Security
Weiden, Germany
e-mail: tobias.bauer@aisec.fraunhofer.de

Andreas Aßmuth 

Ostbayerische Technische Hochschule Amberg-Weiden
Amberg, Germany
e-mail: a.assmuth@oth-aw.de

Abstract—In the context of modern software engineering, there is a trend towards Cloud-native software development involving international teams with members from all over the world. Cloud-based version management services like GitHub are commonly used for source code and other files. However, a challenge arises when developers from different companies or organizations share the platform, as sensitive data should be encrypted to restrict access to certain developers only. This paper discusses existing tools addressing this issue, highlighting their shortcomings. The authors propose their own solution, Encrypted Container Files, designed to overcome the deficiencies observed in other tools.

Keywords—Cloud-based software development; hybrid encryption; agile software engineering.

I. INTRODUCTION

Modern software development processes are agile, span organizational boundaries, and are carried out by teams whose members are not confined to one location. They produce Cloud-native software components that communicate with each other. However, there is strong interest in ensuring the confidentiality of the transmitted data. Usually, encryption algorithms fulfill this task, which is also true for storing confidential data, e.g., in a database. Often during the development of Cloud-native software components, passwords, private certificate keys, and symmetric encryption keys arise. On the one hand, they must be kept secret, but on the other hand, they must be managed in a sensible way and some authorized team members have to have access to these secrets. In the following, we refer to these types of data as *confidential data*. To address the emerging conflict of these two opposing objectives, this paper aims to extend our solution we presented in [1].

The practice of Continuous Integration (cf. [2]) is a long-lasting trend in software development [3][4]. The developed applications are then distributed and deployed automatically via Continuous Delivery and Continuous Deployment [3]. However, designing a secure deployment process or pipeline is difficult and requires special tooling as hinted in [5][6]. As a positive aspect, the awareness to build these automated processes in a secure way has grown over the past two years (cf. [4][7]). This leads to incorporating the security aspect into the development and operation or *DevOps* cycle and stressing this new component by naming the new approach *DevSecOps* [4]. Although best practice guides stress the need for securing confidential data, e.g., in [6] and [8], they lack providing a pipeline vendor-independent solution (cf. [6][9]) or advertise third party (cloud) services (cf. [8]). Some implementations have been found to expose secrets [10].

Software developing teams often use version control systems, e.g., *git* [11], to manage both the source code and

other artifacts. The aforementioned practice of Continuous Integration requires the version control system to contain *all* artifacts, which – by definition – includes secrets we described earlier. However, storing confidential data without protective measures against unauthorized access would be grossly negligent and poses a serious hazard. There are tools available that offer the encryption of such confidential data before storing it within the version control system. However, we observed multiple drawbacks and shortcomings when using these tools. We present Encrypted Container File (ECF), our solution for a Cloud-based, hybrid-encrypted data storage aimed at software development teams. With ECF we extend the functionality of currently available tools and simultaneously overcome the found deficiencies.

The remainder of this paper is structured as follows: in Section II, we take a look at three existing tools and discuss their benefits and shortcomings. Next, we provide a high-level workflow description of the usage of ECF in Section III. This leads to the requirements and the file format description presented in Section IV. In Section V, we describe the operations of an ECF in detail. Since our initial presentation of ECF, we extended our work in various ways, which are described in Section VI. Finally, we conclude this paper in Section VII.

II. RELATED WORK

Whilst researching available solutions for the problem described above, we came across two projects: *jak* and *git-crypt*. In our opinion, however, both have some disadvantages, which we will briefly describe below. Lastly, we argue that Pretty Good Privacy (PGP) or GNU Privacy Guard (GPG) can solve the problem but the usage is inconvenient and error prone.

A. *jak*

jak [12], a tool written in Python, allows the encryption of files with symmetric keys using the Advanced Encryption Standard (AES). With *jak*, keys can be generated and stored in a key file, which itself is not encrypted. If a special text file containing a list of the file names of the corresponding files is added to the repository, *jak* can perform automatic encryption and decryption with a single command. As is well known, the so-called key exchange problem cannot be solved using encryption with symmetric keys only. This problem also exists with *jak*, which relies exclusively on AES as stated before. The tool therefore cannot offer a solution to the key exchange problem, but leaves it up to the developers/users to take care of it themselves. The distribution of confidential data (keys) becomes increasingly challenging, especially as

the size of the development team increases, which is why we believe that the practical use of *jak* is limited. Another issue with *jak* is that the content of confidential files is stored unencrypted on the developers' computers. The reason for this is that *jak* decrypts these files during checkout and re-encrypts them before committing. A prerequisite for the secure use of *jak* is therefore that external parties may only have read-only access to the repository. In addition, third parties must be prevented from accessing the developers' computers. If an attacker manages to gain access to such a computer, they can read the confidential information.

For further information on *jak*, please refer to the project documentation [12].

B. *git-crypt*

git-crypt [13] allows the encryption of files within a *git* repository with AES as well. This results in the same restrictions as with *jak* with regard to the necessary access restrictions for third parties. However, *git-crypt* offers a solution to the key exchange problem by using GPG. For this purpose, the public GPG keys of the recipients are added to the repository. When encrypting confidential files within the repository, *git-crypt* generates an asymmetrically encrypted key file for each recipient. Each authorized recipient can use this to access the symmetric key, which enables each of these recipients to decrypt the confidential files in this repository. *git-crypt* is implemented in a way that all confidential files are encrypted with the same symmetric key and this key must therefore be communicated to all recipients that are added to the repository. Unfortunately, this means that it is not possible to differentiate between different authorized recipients when accessing different confidential files. In our opinion, such a feature would be absolutely desirable. For example, confidential information about the production environment should only be accessible by the production team. However, there may also be other confidential files to which not only (some members of) the production team can have access. But this cannot be realized with *git-crypt*. As with *jak*, the contents of confidential files that are stored locally on the developer's computer are not secured. The reason for this is that *git-crypt* decrypts the confidential data during checkout (same as with *jak*). The integration of *git-crypt* into the mechanisms of *git* is optional, but recommended. With *git-crypt*, we also miss an option of being able to remove authorized users from the confidential files at a later date, for example if one of these developers leaves the project. In such a case, he or she should no longer have access to the confidential files. [13] justifies this by arguing that when using a version control system, a remote recipient can still access old versions of the repository and thus the confidential data stored in it. This is true, of course. However, it should also be borne in mind that a former developer should not have access to new certificates, passwords, etc. once the old ones have expired. Against the background of such changes of personnel, we consider this feature to be desirable. In addition, we recommend also changing the symmetric keys in this scenario.

C. *Pretty Good Privacy*

Another way to solve the problem would be to use PGP or GPG directly. To do this, one of the developers would generate a random session key to encrypt the confidential data. The random session key would then be encrypted with the public key of another developer. The file consisting of the encrypted session key and the symmetrically encrypted ciphertext could then be stored in the online repository. Only the aforementioned second developer would be able to access the confidential data.

As this brief description of PGP/GPG suggests, this solution does not scale very well. If the confidential data is required to be shared with multiple recipients, one such file per recipient must be created and stored in the repository. Of course, this is feasible in principle. However, it is important to note that each recipient must know which of these files is intended for them. In addition, it is absolutely impure that the same confidential data has to be stored in encrypted form in several instances. Changes to the content of the confidential data naturally require all these recipient-specific files to be recreated and stored. This makes this solution very challenging to use, and there is also the problem of keeping the content of all the confidential data's ciphertext files consistent. Just imagine that the confidential data of a project is spread over multiple files and each for different groups of recipients. Although handling with this solution is feasible in principle, it demands a great deal from everyone involved and is anything but convenient.

III. APPLICATION

In this section, we present a high-level description of multiple parties using an ECF. In order to illustrate the basic operations, we follow a common workflow involving various participants representing the different parties. The generic users Alice, Bob, and Charlie use a command line-based tool to interact with an ECF. This tool is also available as a part of this paper's accompanying source code¹. Furthermore, all basic operations are depicted in Figure 1. Each subfigure corresponds to its analogously named subsection, e.g., Figure 1a visualizes Subsection III-A.

This use case-oriented view serves as a foundation for the distillation of requirements of the ECF format presented in Section IV. We describe the technical details in Section V, which follows the workflow presented in this section.

A. *Creating an ECF*

Alice is the head of an operations team working for a company. Her task is to securely store secrets, such as server certificates for the company's website. Furthermore, a subset of her coworkers as well as an automated deployment job should get access to these secrets. Third parties, i.e., all other coworkers as well as any other person, must not get access to the sensitive data. Alice therefore chooses the ECF format to protect and store the secret data.

¹<https://github.com/Hirmmoder/ECF>

As a first step, Alice wants to create an ECF using the aforementioned command line tool (hereinafter referred to as the tool). To do so, Alice needs to generate a key pair consisting of a private and a public key. She may use this key pair for as many different ECFs as she likes, but she can also generate as many different key pairs as she likes. In order to keep this workflow description basic, we assume that Alice generates a single key pair using the tool. Since the private key must not be accessible to any person other than Alice, storing it unencrypted is not an option. This is why the tool asks for a self-chosen password which is then used to encrypt Alice's private key before storing it to disk.

In Figure 1a and all other subfigures in Figure 1, we depict a key pair as a keyring with two keys attached. Additionally, we always use the same color for a user as well as their keys and attributes, which is red for Alice.

After this initial step, Alice uses the tool to create an empty ECF. The tool then asks for her self-chosen password to get access to her private and public key. This is needed in order to add Alice as a recipient to the newly created ECF and to store its initial content, i.e., the secret server certificate, inside. To indicate that Alice is a recipient of that newly created ECF, we add an appropriately colored rectangle to it in Figure 1a.

When using a version control system, such as *git* [11], Alice may now add the ECF to the repository, commit the changes, and push it to a remote server. Since the sensitive data is encrypted within the ECF and no recipient but Alice was added to the ECF, only Alice is able to decrypt and access the secrets stored in that ECF.

B. Decrypting and Using an ECF

In order to access the data stored inside an ECF, Alice needs to be a recipient of the ECF in question. According to our workflow, Alice was added as a recipient to the ECF during its creation. She, thus, can decrypt and view the content of the ECF. In addition to that, Alice is also able to alter the recipient set as well as the sensitive data. Both of these actions are described later in Subsections III-C to III-E. This subsection is focused solely on the decryption and usage of an ECF.

Alice uses the tool to access the encrypted content. The tool first asks for Alice's self-chosen password and then decrypts the ECF using her private key. It then displays the decrypted content of the file for Alice to view and use. This operation is visualized in Figure 1b.

Usually, manual inspection of the content of an ECF is rare. It is much more likely that other processes, e.g., an automated deployment job, needs to use the secrets and therefore needs to access them. Such a process should run in a trusted environment, e.g., on a self-hosted deployment server. For the ECF, the process looks like any other recipient and, thus, Alice has to first add the process as a recipient to the ECF in order for it to be able to access the stored secret data. The procedure for adding a recipient is described in Subsection III-C. The process then can read (and possibly modify) the encrypted secrets inside the ECF and perform its task, e.g., deploying the company's website with the server certificate.


C. Adding Recipients

We now assume the following situation: Alice has two coworkers that also need access to the secrets, e.g., as a substitute when she is ill. Furthermore, the aforementioned automated deployment job also needs access to the secrets in order to fulfill its task.


As a preparatory step, her coworkers Bob and Charlie need to generate a key pair each, analogous to what Alice did in Subsection III-A. Then, Bob and Charlie export their public keys using the tool and sending those to Alice, e.g., via email. Bob and Charlie can perform these steps completely independently of each other and may use different computers and operating systems. Since Bob and Charlie keep their private keys private and only share their public keys, we use a single-key symbol for them in Figure 1c.

Since the automated deployment job is a process and not a person, Alice (or someone else she trusts for that matter) can generate a key pair for the process. This key pair can now be used to export the public key using the tool. The private key must be made available to the process while still keeping it safe from unauthorized access. This can be done by restricting access to the server the process is running on. This is also why the process should run within a trusted environment as stated in Subsection III-B.


Still, nobody but Alice has currently access to the content stored in the ECF Alice created in Subsection III-A. To change this, Alice now receives and checks the exported public keys of the recipients-to-be. She has to make sure that the received public keys were not tampered with during the transmission. A way to ensure authenticity is to compare, e.g., fingerprints of these public keys via a second channel. For example, we assume that Alice and Bob work in the same facility. They now can verify that the transmission of Bob's public key did not alter it by displaying the fingerprints on their laptop screens and comparing it visually.



(a) Alice creates an Encrypted Container File (cf. Subsection III-A).



(b) Alice decrypts and uses an Encrypted Container File (cf. Subsection III-B).



(c) Alice adds recipients Bob, Charlie and the automated deployment job (cf. Subsection III-C).

Figure 1. Overview of ECF operations according to Subsections III-A to III-C.

After verifying the received public keys, Alice can now add Bob, Charlie, and the automated deployment job as recipients to the ECF. She uses the tool for each recipient addition. Since modifying the ECF also involves decrypting it, Alice is prompted for her self-chosen password by the tool in order to access her private key. The order in which Alice adds the recipients does not matter. Afterwards, she may delete the exported public keys since they are no longer of use.

As a last step Alice must make the updated ECF available to the other users. She may commit the changes and push them to a remote server using *git*. Bob, Charlie, and the automated deployment job now gained access to the ECF using their respective private keys and, thus, can perform the same actions Alice can, i.e., decrypting and using the ECF, add or remove recipients, and change the encrypted content. In Figure 1c, we indicate that Bob, Charlie and the automated deployment job are now recipients by appropriately colored rectangles.

D. Removing Recipients


So far, we have tackled the recipient addition procedure. We now assume that Bob leaves the company and should, therefore, no longer have access to the content of the ECF. In this case, both, Alice and Charlie, can remove Bob from the recipient set – we opted for Alice to perform this step as she is the head of the team. This operation is also visualized in Figure 1d. Alice uses the tool to identify and remove Bob from the recipient set of the ECF. As with any operation, the tool asks for her self-chosen password beforehand. Alice then commits the changes and pushes them to a remote server.

Bob now loses access to the current and future versions of the ECF. However, Bob still has access to older versions. It is futile to withdraw access to older versions because (a malicious) Bob could have copied the secret content anyways when he still had access to. An ECF is designed to securely store secrets, such as access keys, private certificate keys, passwords, asymmetric and symmetric encryption keys, etc. As it is a good practice to rotate keys as described in [14], all secrets Bob once had access to will eventually expire. This is especially true for short-lived TLS certificates.


Nevertheless, it is advisable to only add recipients to an ECF if it is really necessary in the first place [15]. In fact, since each ECF is independent of other ECFs, the format allows for different recipient sets per file. Let us assume that Alice has two different secrets to share with different groups of people. Instead of creating a single ECF and adding all recipients to it, Alice should rather create two ECFs, one for each secret and only add the respective intended recipients to each.

E. Modifying the Content

The original ECF Alice created contains the server certificates for the company's website. Since these certificates are usually short-lived, they have to be renewed regularly. Alice's coworker Charlie – who got access to the ECF as described in Subsection III-C – now has the task of modifying the content of the ECF, which is depicted in Figure 1e.



(d) Alice removes recipient Bob (cf. Subsection III-D).



(e) Charlie modifies the content (cf. Subsection III-E).

Figure 1. (Continued.) Overview of ECF operations according to Subsections III-D and III-E.

Charlie uses the tool to modify the ECF and therefore is prompted to enter her self-chosen password. Then she can view and edit the content of the ECF. In order to fulfill her task, she removes the expired server certificate from the ECF and inserts the renewed one. In Figure 1e, we label the original content as x and the modified one as y . Last, she commits her changes using *git* and pushes them to a remote server to make the new version of that ECF available to all other recipients.

Now Alice, Charlie, and the automated deployment job can access the renewed certificate stored inside the ECF. Bob, on the other hand, cannot decrypt this ECF anymore as he was removed from the recipient set in Subsection III-D.

IV. REQUIREMENTS ENGINEERING AND INTERNAL FILE STRUCTURE

Based on the workflow described in Section III, we define requirements for the ECF format. Then, we use these requirements to formulate design goals for the definition of the file format. In the following subsections, we describe the file structure and used data types in detail. A technical description of the operations on the ECF format is given in Section V.

A. Use Case Requirements

We derive some requirements from the high-level workflow description given earlier. Furthermore, we take the features and weaknesses of the *jak* and *git-crypt* tools as well as the plain-PGP use into consideration. We have distilled the following requirements for the ECF format, cf. [1]:

- 1) The encryption of confidential data is mandatory,
- 2) possibility to modify the confidential data, i.e., the content is writable,
- 3) key distribution is no prerequisite,
- 4) decryption occurs not during checkout but on demand,
- 5) support for many recipients,
- 6) addition and removal of recipients,
- 7) minimal information gain for external parties,
- 8) a customizable set of recipients per file, and
- 9) be reasonably fast to be used in production.

B. Design Goals

The requirements in mind, we have decided to define the following design goals for our ECF format, cf. [1]. The numbers in brackets refer to the enumeration in Subsection IV-A.

- Use of hybrid encryption (Items 1, 3 and 5),
- inclusion of recipient information to allow re-encryption on changes (Items 2, 5, 6 and 8),
- minimal information about recipients publicly available (Item 7),
- obfuscation of recipient information for respective external parties (Items 7 and 8),
- delivery of the associated software as a library for embedding into existing applications (Item 4), and
- optimization of the encryption and decryption procedure to take less than 100 ms on modern hardware for common confidential data types (Item 9), while also
- be flexible enough to support different cipher suites for future extensions (Items 1, 3 and 9).

C. Structure, Data Types, Storage Format, and Notation

Each ECF comprises a public header, a private body, and a public footer. The general structure of the ECF format is depicted in Figure 2. The public fields are unencrypted and therefore readable by anyone. The private body, however, is fully encrypted using a symmetric encryption scheme and can

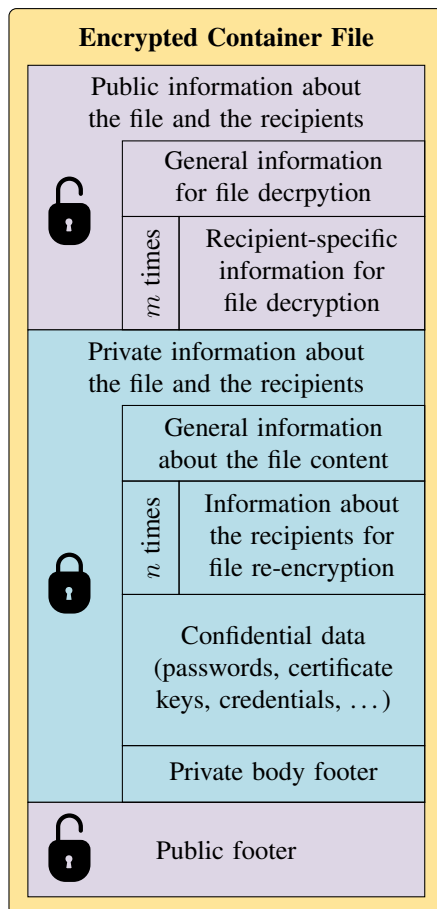


Figure 2. General structure of an Encrypted Container File, cf. [1].

only be decrypted by the recipients of the ECF. Therefore, external parties do not get access to the confidential data stored within the body of the ECF. We describe the public and the private components of an ECF in more detail in Subsections IV-D and IV-E, respectively.

Throughout the following descriptions of the public and private fields of an ECF, we use this notation to denote the data types of each field:

- `u32le` Unsigned 32 bit Integer, Little Endian
- `u8` Unsigned 8 bit Integer (byte)
- `s` String: length in bytes as `u32le`, then UTF-8 bytes without byte order mark (BOM)
- `x [n]` Array of type `x` with `n` entries, e.g., `u8 [16]` denotes an array of 16 bytes.

All variables referenced in this subsection are listed in Table I in Appendix B. Our solution employs four types of cryptographic algorithms to achieve the aforementioned goals:

- Diffie-Hellman-like Key Agreement/Exchange Algorithm
- Matching Signature Scheme (cf. Appendix A)
- Symmetric Encryption Scheme
- Suitable Hash Function

Each set of these algorithms is called a cipher suite. In our Proof of Concept (PoC) implementation, we define four cipher suites listed below. We designed the ECF format to be flexible with regards to the used cipher suite and in order to determine the cipher suite of a specific ECF, we assigned a unique identifier to each. In comparison to our previous work in [1], Cipher Suites III and IV were added to the PoC implementation with little effort, proving the flexibility of the ECF format. The four cipher suites are:

- I) Key Exchange: X25519 [16]
Signature: Ed25519 [17]
Symmetric Encryption: AES-256-GCM [18][19]
Hash Function: SHA-256 [20]
Identifier: 0x01010101
- II) Key Exchange: X25519 [16]
Signature: Ed25519 [17]
Symmetric Encryption: AES-256-GCM [18][19]
Hash Function: SHA-512 [20]
Identifier: 0x01010102
- III) Key Exchange: X25519 [16]
Signature: Ed25519 [17]
Symmetric Encryption: AEGIS-256 [21]
Hash Function: SHA-256 [20]
Identifier: 0x01010201
- IV) Key Exchange: X25519 [16]
Signature: Ed25519 [17]
Symmetric Encryption: AEGIS-256 [21]
Hash Function: SHA-512 [20]
Identifier: 0x01010202

The provided PoC contains a command line tool, which uses Cipher Suite II by default. When creating an ECF using the command line tool, the user can override the cipher suite to use. There is currently no built-in option to change the

cipher suite afterwards. This is because some combinations of cipher suites – in particular the keys for key exchange and signing – may be incompatible. This problem does not affect the implemented ones, but may arise due to future extensions.

D. Unencrypted, Public Header and Footer Fields

All authorized recipients must be able to decrypt the file and, thus, must be able to acquire the symmetric encryption key used to encrypt the private body. Simultaneously, in order to achieve the goal of minimal public information about the recipients, the public header fields contain as little identifying information as possible. Our solution allows us to separate the information needed for decrypting and the information needed for encrypting. Whereas the former must be made public in order for the whole scheme to work, the latter can be included in the encrypted body and therefore be protected against unauthorized access.

The public header of an ECF comprises a general, recipient-independent part followed by m identically constructed recipient-specific blocks. Figure 3 shows all public header fields in their respective order. The recipient-independent fields contain general information about the ECF:

- CONTAINER VERSION of type `u32le`
ECF format version; intended for future extensions; currently set to `0x00010000` for “Version 1.0”
- CIPHER SUITE of type `u32le`
Identifier for chosen cipher suite; cf. enumeration in Subsection IV-C
- PUBLIC LENGTH of type `u32le`
Length h of the public header in bytes
- PRIVATE LENGTH of type `u32le`
Length b of the encrypted private body in bytes
- RECIPIENT COUNT of type `u32le`
Number of recipients m in the public header
- SALT of type `u8 [16]`
Salt value; cf. Subsections V-A and V-B
- SYMMETRIC NONCE of type `u8 [c]`
Symmetric nonce value; cf. Subsections V-A and V-B; Length dependent on chosen cipher suite, usually $c = 12$

0	4	8	12
16			
CONTAINER VERSION	CIPHER SUITE	PUBLIC LENGTH h	PRIVATE LENGTH b
RECIPIENT COUNT m	u8 [16] SALT ...		
... SALT	u8 [c] SYMMETRIC NONCE		
dblock [m] RECIPIENT-SPECIFIC BLOCKS ...			
...			
...			
...			
... RECIPIENT-SPECIFIC BLOCKS			
h			

Figure 3. Public header fields of an Encrypted Container File.

All 4 byte-long fields are of type `u32le`. While the length of SALT is fixed to 16 bytes, the length of SYMMETRIC NONCE is dependent on the chosen cipher suite. However, since all implemented cipher suites use an Authenticated Encryption with Associated Data (AEAD) symmetric encryption scheme, we follow the recommendations in RFC 5116 [22] and set it to $c = 12$ bytes as shown in Figure 3 for the Galois/Counter Mode (GCM) [18] of AES. The AEGIS specification [21], however, requires $c = y$.

There are exactly m recipient-specific decryption blocks or `dblocks` following the recipient-independent part. The structure and length of these blocks is also highly dependent on the cipher suite. Figure 4 shows the general, three-part structure of these blocks. During the decryption procedure, recipients have to find their respective block by matching a unique identification number. This is why IDENTIFICATION TAG is the first entry in each block and has a fixed length of 16 bytes. It allows for up to $(2^8)^{16} = 2^{128}$, i.e., practically unlimited number of recipients while saving space compared to a full-length hash value. For illustration purposes, we decided to use the array lengths of the default Cipher Suite II. In this case, KEY AGREEMENT INFORMATION contains an ephemeral X25519 public key of length $a = 32$ bytes and SYMMETRIC PRE KEY 1 has, because of AES-256, a length of $y = 32$ bytes.

In Figure 4, we use o_i to denote the offset of a recipient-specific block $0 \leq i < m$ in bytes measured from the start of the file. It can be calculated as $o_i = i \cdot (16 + a + y) + o_0$. Since the blocks follow the recipient-independent part, $o_0 = 36 + c$. Furthermore, $o_m = h$ must hold in order for an ECF to be considered valid.

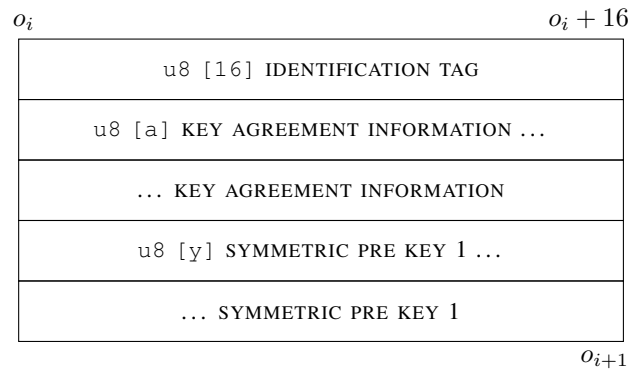


Figure 4. Recipient-specific block structure within the public header of an Encrypted Container File.

As shown in Figure 2, an ECF also contains a public footer. It consists of a single field: FILE HASH of type `u8 [d]`. Its length is dependent on the hash function of the chosen cipher suite. For example, $d = 64$ bytes applies to SHA-512. The field FILE HASH serves primarily to check the integrity of the file, e.g., after a copy or transfer action, and, thus, helps to detect accidental alterations. It does not increase the resilience against a malicious attacker as it can be recomputed easily after a deliberate alteration of other parts of the file.

The encryption procedure fills or generates all public header and footer fields and the decryption procedure uses them. Both operations are described in detail in Section V.

E. Encrypted, Private Body Fields

After retrieving the symmetric encryption key using the public header information, an authorized recipient is able to decrypt the private body of the ECF. External parties, who are not recipients of that ECF, are unable to obtain the encryption key and, thus, unable to access any information stored within the body of that ECF. The body of an ECF consists of four parts as shown in Figure 2. The fields within the ECF body, which are depicted in Figure 5, are composed as follows:

- CONTENT TYPE of type `u32le`
Describes the type of the confidential data
- PUBLIC HEADER HASH of type `u8 [d]`
Hash value of the public header
- RECIPIENT COUNT of type `u32le`
Number of true recipients $n \leq m$
- RECIPIENT INFORMATION BLOCKS of type `rblock [n]`
Information about the recipients of that ECF
- CONTENT LENGTH of type `u32le`
Length q of the confidential data in bytes
- CONTENT of type `u8 [q]`
Confidential data
- PRIVATE HASH of type `u8 [d]`
Hash value of the private body so far

The field `CONTENT TYPE` characterizes the type of confidential data stored within `CONTENT`. Our PoC implementation provides a built-in type of content, Binary Large Object (BLOB), using the identifier `0x00000001`. Since our implementation is flexible and extensible, future applications can and should define and handle their own types of content.

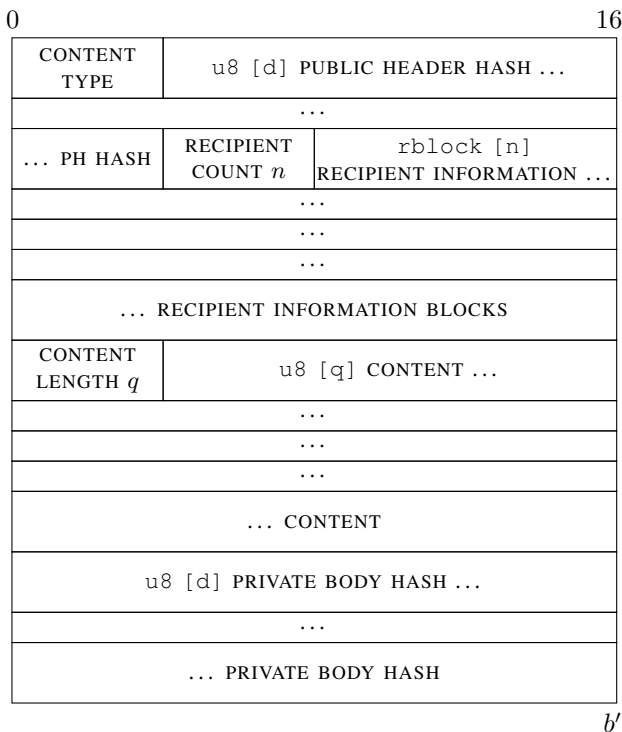


Figure 5. Private body fields of an Encrypted Container File.

Similar to the public footer described in Subsection IV-D, the lengths of the two fields `PUBLIC HEADER HASH` and `PRIVATE BODY HASH` is dependent on the chosen cipher suite. For Cipher Suite II, $d = 64$ bytes. Both fields aim to ensure integrity of the ECF and to detect unauthorized or unintended modifications of both the public header and the private body. For example, if one of the m public, recipient-specific blocks is altered preventing the corresponding recipient from obtaining the symmetric encryption key, all other recipients can detect this alteration. We opted to use fields for the hash values in order to allow for non-AEAD-based symmetric encryption schemes. This is also why we do not make use of the “associated data” option during encryption and decryption (cf. Section V).

Next, the private body contains information about the recipients in variable-sized recipient information blocks or `rblocks`. There are exactly n of these blocks present as denoted in Figure 5 and they contain the necessary information for file re-encryption. Allowing for $m \geq n$ is a design choice we made to hide the true number of recipients n from external parties, who can only read the public header field `RECIPIENT COUNT` and, thus, only obtain m . Furthermore, by storing the recipient-identifying information within the encrypted part of an ECF, we ensure that this information is only accessible by the recipients themselves and hidden from external parties, thus, pursuing the goal of minimal publicly available information about the recipients.

Each `rblock` contains three fields as shown in Figure 6:

- PUBLIC KEY of type `u8 [u]`
Public key of the recipient;
Length dependent on chosen cipher suite, e.g., $u = 32$
- NAME of type `s`
Self-chosen name of the recipient, usually email address or another organization-specific identifier
- NAME SIGNATURE of type `u8 [g]`
Signature over NAME;
Length dependent on chosen cipher suite, e.g., $g = 64$

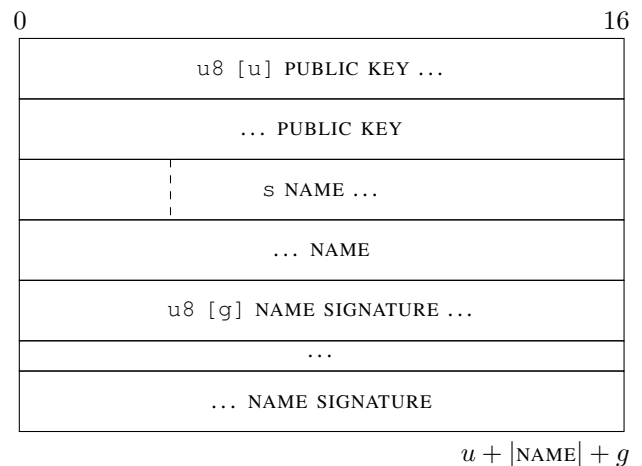


Figure 6. Recipient information block structure within the private body of an Encrypted Container File.

Since these blocks do not have a fixed size due to the variable-length self-chosen `NAME`, the offset measured from the start of the private body must be calculated by iterating through all previous blocks. The signature over the self-chosen name serves to assure that the person owning the associated private key has chosen the name and that no changes have been made to the name afterwards. The name is for information purposes only, e.g., when displaying the recipients or removing an existing recipient by name.

The main motivation for designing the ECF format was to securely store confidential data in a multi-recipient setting. The secrets, which may have any format the users want them to, are stored within the field `CONTENT`. Its theoretical maximum size $\max(q)$ in bytes is ultimately dependent on the number of true recipients n and the lengths of their self-chosen names as well as the chosen cipher suite. This is because the length of the encrypted private is limited to $\max(b) = 2^{32} - 1$ bytes ≈ 4 GiB. We can calculate the theoretical maximum size of the confidential data as follows:

$$\begin{aligned} \max(q) &= \max(b) - \bar{q} - (b - b') \\ \bar{q} &= b' - q \\ &= 4 + d + 4 + \sum_{i=1}^n (u + |\text{NAME}_i| + g) + 4 + d \\ &= 12 + 2d + n(u + g) + \sum_{i=1}^n |\text{NAME}_i| \end{aligned}$$

The total length of all private fields is denoted by $b' \leq b$, whereas b denotes the length of the encrypted private body. The public header field `PRIVATE LENGTH` holds the value b , which may be greater than b' depending on the used symmetric encryption scheme. This is definitely true for AEAD-based algorithms. Therefore, $b - b' \geq 0$ describes the number of bytes the encrypted private body exceeds the unencrypted one. The helper variable $\bar{q} = b' - q > 0$ denotes the length of all private body fields except `CONTENT`.

We can show that $\bar{q} + (b - b')$ is negligible compared to $\max(b)$, and therefore the amount of data that can be stored within an ECF is sufficient for passwords, certificate keys, and similar confidential data. We assume $n = 1000$ recipients with an average name length of $|\text{NAME}_i| = 500$ bytes and using the default Cipher Suite II:

$$\begin{aligned} \bar{q} + (b - b') &= 12 + 2 \cdot 64 + 1000 \cdot (32 + 64 + 500) + 16 \\ &= 596157 \approx 2^{19} \ll 2^{32} - 1 = \max(b) \\ \implies \max(q) &\approx 2^{32} - 2^{19} \approx 2^{32} \end{aligned}$$

In practice, both n and the average name length will be significantly less, supporting our argument even more. Our strong assumption is that, in practice, the number of recipients of a single ECF does typically not exceed 20. Furthermore, using email addresses as recipient identification in the field `NAME` as we suggested earlier yields far shorter names than in the assumption above.

V. ECF OPERATIONS

This section describes the basic operations on an ECF. We follow roughly the workflow presented in Section III and describe the encryption procedure in Subsection V-A, the decryption procedure in Subsection V-B, and the modification of the recipient set and confidential data in Subsections V-C and V-D, respectively.

Nomenclature. We use a common notation for the following subsections. Since the ECF format is flexible with regards to the used cipher suite, we opted to use variable and function names that are not tied to a specific algorithm. Due to readability concerns, we moved the notation to Appendix B, and specifically to Tables II and III.

A. Creating and Encrypting an ECF

As described in Subsection III-A, Alice first needs to generate a key pair consisting of a private and a public key. All cipher suites in our PoC implementation allow to convert an Ed25519 key pair to a X25519 key pair as described in [23][24]. Therefore, we generate the key pair for signing and then convert it to a key pair for key exchange:

$$\begin{aligned} (\text{sk}_A^S, \text{pk}_A^S) &\leftarrow \text{Gen}^S \\ (\text{sk}_A^X, \text{pk}_A^X) &= (\text{Convert}^{S \rightarrow X}(\text{sk}_A^S), \text{Convert}^{S \rightarrow X}(\text{pk}_A^S)) \end{aligned}$$

Two components are essential to an ECF: the confidential data $\Omega \in \{0, 1\}^{8q}$ and the set of recipients $\Psi = \{\psi_0, \dots, \psi_{n-1}\}$. When creating a new ECF, both components are empty. As a next step, Alice adds herself to the recipient set and, therefore, must generate a recipient entry $\psi_a = \text{pk}_A^S \parallel \text{NAME}_a \parallel t_a$. She may choose a name or use, for example, her email address as an identifier. Then she signs it:

$$t_a = \text{Sign}(\text{sk}_A^S, \text{UTF8}(\text{NAME}_a))$$

Alice now adds herself to the recipient set, i.e., $\Psi = \{\psi_a\}$. Following the basic workflow described earlier, Alice takes the byte representation of the confidential data in any format she likes and assigns it to Ω . As a final step, Alice encrypts the two components, thus, obtaining a valid ECF:

$$\mathcal{E} \leftarrow \text{Enc}^{\text{ECF}}(\Psi, \Omega)$$

The encryption procedure takes two inputs, the set of recipients Ψ and the confidential data Ω . It is important to note that parts of the cipher suite are selected implicitly by the construction of ψ_a . In the following, we assume that all cipher suite algorithms and their respective properties, e.g., the symmetric key length y and the nonce length c , are well defined according to the cipher suite Alice chose.

Additionally, the encryption procedure does not require any private key directly. However, the generation of ψ_a did require Alice's private key for signing, which is why our PoC implementation requests access to it. Yet, e.g., a pre-created recipient set Ψ , allows anyone to create an ECF without first generating a key pair. This operation, however, is not supported by our PoC implementation as the creator would not be a recipient and, thus, be unable to access the content of that ECF.

The encryption procedure is defined as follows:

- (1) Generate a symmetric key, a nonce, and a salt
 $k_{\text{final}} \leftarrow \text{Gen}^{\text{SYM}}, \alpha \leftarrow \text{Random}(c), \theta \leftarrow \text{Random}(16)$.
- (2) Choose a random integer $m \geq n = |\Psi|$.
- (3) For all $\psi_i \in \Psi, 0 \leq i < n$ generate r_i : ($R = \emptyset$ initially)
 - (a) Compute IDENTIFICATION TAG
 $\text{id_tag} = \text{H}(\text{pk}_{\psi_i}^S \parallel \theta) [0, \dots, 15]$.
 - (b) Generate an ephemeral key pair for key exchange
 $(\text{sk}_e^X, \text{pk}_e^X) \leftarrow \text{Gen}^X$.
 - (c) Convert the recipient's public key for signing to a public key for key exchange $\text{pk}_{\psi_i}^X = \text{Convert}^{\text{S} \rightarrow \text{X}}(\text{pk}_{\psi_i}^S)$.
 - (d) Execute the key exchange algorithm:
 $\text{ss} = \text{Kex}(\text{sk}_e^X, \text{pk}_{\psi_i}^X)$.
 - (e) Calculate $k_{\text{pre2}} = \text{H}(\text{ss} \parallel \text{pk}_{\psi_i}^X \parallel \text{pk}_e^X) [0, \dots, y-1]$.
 Using this hash function construction instead of the shared secret directly is recommended in [25].
 - (f) Calculate $k_{\text{pre1}} = k_{\text{final}} \oplus k_{\text{pre2}}$.
 - (g) Add $r_i = \text{id_tag} \parallel \text{pk}_e^X \parallel k_{\text{pre1}}$ to R .
- (4) Generate $m - n$ deception entries as described in Subsection V-E, add them to R and shuffle or sort the elements $r_i \in R, 0 \leq i < m$ by id_tag .
- (5) Write all public header fields into a temporary buffer \mathcal{H}^* , except for PRIVATE LENGTH as it is not known yet. Set its value b to $0_{\text{XECFFCODE}}$ (ECF format "code") and compute the hash value of the public header:

$$\begin{aligned} \mathcal{H}^* &= \text{CONTAINER VERSION} \parallel \text{CIPHER SUITE} \parallel h \parallel \\ &\quad 0_{\text{XECFFCODE}} \parallel m \parallel \theta \parallel \alpha \parallel r_1 \parallel r_2 \parallel \dots \parallel r_m \\ \beta_{\text{public}} &= \text{H}(\mathcal{H}^*). \end{aligned}$$

- (6) Write all private body fields unencrypted into a temporary buffer \mathcal{B}' with $q = |\Omega|$ in bytes and compute the hash value of the private body:

$$\begin{aligned} \psi_i &= \text{pk}_{\psi_i}^S \parallel \text{NAME}_i \parallel t_i \\ \mathcal{B}' &= \text{CONTENT TYPE} \parallel \beta_{\text{public}} \parallel n \parallel \\ &\quad \psi_1 \parallel \psi_2 \parallel \dots \parallel \psi_n \parallel q \parallel \Omega \\ \beta_{\text{private}} &= \text{H}(\mathcal{B}'). \end{aligned}$$

- (7) Append β_{private} to \mathcal{B}' and then encrypt the buffer:

$$\mathcal{B} \leftarrow \text{Enc}^{\text{SYM}}(k_{\text{final}}, \alpha, \mathcal{B}' \parallel \beta_{\text{private}}).$$

- (8) Update the previously written field PRIVATE LENGTH in \mathcal{H}^* with its actual value $b = |\mathcal{B}|$ in bytes to obtain the final header \mathcal{H} and compute the hash value over everything:

$$\begin{aligned} \mathcal{H} &= \text{CONTAINER VERSION} \parallel \text{CIPHER SUITE} \parallel h \parallel \\ &\quad b \parallel m \parallel \theta \parallel \alpha \parallel r_1 \parallel r_2 \parallel \dots \parallel r_m \\ \beta_{\text{all}} &= \text{H}(\mathcal{H} \parallel \mathcal{B}). \end{aligned}$$

- (9) The output \mathcal{E} of the ECF encryption procedure is

$$\mathcal{E} = \mathcal{H} \parallel \mathcal{B} \parallel \beta_{\text{all}}.$$

B. Decrypting an ECF

Only recipients of an ECF are able to decrypt it. In particular, having access to the private key of a recipient's key pair implies access to the encrypted content stored in the private body of an ECF. Alice has to use her private key for signing in order to decrypt the ECF and obtain the recipient set Ψ and the confidential data Ω :

$$(\Psi, \Omega) = \text{Dec}^{\text{ECF}}(\text{sk}_A^S, \mathcal{E})$$

The decryption procedure is defined as follows:

- (1) Load the public header fields and determine the container version as well as the cipher suite. Now h, b, m, θ , and α have values. Let \mathcal{H} be the public header, \mathcal{B} be the encrypted private body, and β_{all} the public footer.
- (2) Check file integrity. If $\text{Vrfy}^{\text{H}}(\mathcal{H} \parallel \mathcal{B}, \beta_{\text{all}}) \stackrel{?}{=} 0$, *Exit*.
- (3) Compute all necessary keys:
 $\text{sk}_A^X = \text{Convert}^{\text{S} \rightarrow \text{X}}(\text{sk}_A^S)$,
 $\text{pk}_A^S = \text{Convert}^{\text{sk} \rightarrow \text{pk}}(\text{sk}_A^S)$, $\text{pk}_A^X = \text{Convert}^{\text{sk} \rightarrow \text{pk}}(\text{sk}_A^X)$.
- (4) Compute $\text{id_tag} = \text{H}(\text{pk}_A^S \parallel \theta) [0, \dots, 15]$.
- (5) Search recipient-specific block $r_a = \text{id_tag} \parallel \text{pk}_e^X \parallel k_{\text{pre1}}$ with matching id_tag .
 If not found: Alice is not a recipient of this ECF. *Exit*.
- (6) Execute the key exchange algorithm:
 $\text{ss} = \text{Kex}(\text{sk}_A^X, \text{pk}_e^X)$.
- (7) Calculate $k_{\text{pre2}} = \text{H}(\text{ss} \parallel \text{pk}_A^X \parallel \text{pk}_e^X) [0, \dots, y-1]$.
- (8) Calculate $k_{\text{final}} = k_{\text{pre1}} \oplus k_{\text{pre2}}$.
- (9) Decrypt the private body \mathcal{B} using the computed symmetric key and the nonce α into a buffer \mathcal{B}^* :
 $\mathcal{B}^* = \text{Dec}^{\text{SYM}}(k_{\text{final}}, \alpha, \mathcal{B})$.
- (10) Deconstruct $\mathcal{B}^* = \mathcal{B}' \parallel \beta_{\text{private}}$ into the private body fields.
- (11) Compute the hash over the public header fields with b set to $0_{\text{XECFFCODE}}$ and verify it:
 $\mathcal{H}^* = \text{CONTAINER VERSION} \parallel \text{CIPHER SUITE} \parallel h \parallel$
 $0_{\text{XECFFCODE}} \parallel m \parallel \theta \parallel \alpha \parallel r_1 \parallel r_2 \parallel \dots \parallel r_m$
 If $\text{Vrfy}^{\text{H}}(\mathcal{H}^*, \beta_{\text{public}}) \stackrel{?}{=} 0$, *Exit*.
- (12) For each recipient $\psi_i \in \Psi$:
 - (a) Load and deconstruct ψ_i into $\text{pk}_{\psi_i}^S$, NAME_i , and t_i .
 - (b) Verify the signature:
 If $\text{Vrfy}^{\text{S}}(\text{pk}_{\psi_i}^S, \text{UTF8}(\text{NAME}_i), t_i) \stackrel{?}{=} 0$, *Exit*.
- (13) Compute the hash over the unencrypted private body excluding the field PRIVATE BODY HASH and verify it:
 $\mathcal{B}' = \mathcal{B}^* [0, \dots, b' - d]$
 $= \text{CONTENT TYPE} \parallel \beta_{\text{public}} \parallel n$
 $\psi_1 \parallel \psi_2 \parallel \dots \parallel \psi_n \parallel q \parallel \Omega$
 If $\text{Vrfy}^{\text{H}}(\mathcal{B}', \beta_{\text{private}}) \stackrel{?}{=} 0$, *Exit*.
- (14) The outputs of the ECF decryption procedure are the recipient set Ψ and the confidential data Ω .

Performing Steps 2, 11, 12 and 13 is optional but strongly recommended in order to detect unintended and/or malicious modifications to the ECF. However, Step 12 takes a considerable amount of time and therefore may be omitted in certain cases. For example, a Continuous Integration (CI) pipeline usually contains a build stage and a test stage [2]. We assume that the whole pipeline runs in a trusted environment, e.g., a self-hosted server. Then, it is advisable to check the recipient's signatures as a prerequisite to the build stage, while this step may be skipped during testing in order to speed up the pipeline execution. Nevertheless, the execution times in question are in the range of milliseconds, cf. Subsection VI-C.

C. Adding and Removing Recipients

Adding recipients to an ECF or removing recipients from an ECF implies modifying the recipient set Ψ . We follow the basic workflow described in Subsections III-C and III-D. In order to obtain Ψ (and the confidential data Ω), Alice is required to be a recipient of the ECF \mathcal{E} , i.e. $\psi_a \in \Psi$.

In the following, we describe the recipient addition procedure for a single recipient: Bob. Adding more recipients, such as Charlie and the automated deployment job, is performed analogously.

Before Alice can add Bob to \mathcal{E} , her coworker has to generate his recipient entry ψ_b and, thus, has to generate key pairs analogous to what Alice did in Subsection V-A:

$$\begin{aligned} (\text{sk}_B^S, \text{pk}_B^S) &\leftarrow \text{Gen}^S \\ (\text{sk}_B^X, \text{pk}_B^X) &= (\text{Convert}^{S \rightarrow X}(\text{sk}_B^S), \text{Convert}^{S \rightarrow X}(\text{pk}_B^S)) \\ t_b &= \text{Sign}(\text{sk}_B^S, \text{UTF8}(\text{NAME}_b)) \\ \psi_b &= \text{pk}_B^S \parallel \text{NAME}_b \parallel t_b \end{aligned}$$

Bob now sends ψ_b to Alice. Alice performs the following steps to add him as a recipient to \mathcal{E} :

(1) Verify signature:

$$\text{If } \text{Vrfy}^S(\text{pk}_{\psi_b}^S, \text{UTF8}(\text{NAME}_b), t_b) \stackrel{?}{=} 0, \text{Exit.}$$

(2) Decrypt the ECF: $(\Psi, \Omega) = \text{Dec}^{\text{ECF}}(\text{sk}_A^S, \mathcal{E})$.

(3) Check, if Bob is already a recipient by comparing the public keys for signing: If $\exists \psi_i \in \Psi : \text{pk}_{\psi_i}^S = \text{pk}_B^S$, *Exit*.

(4) Optionally check, if Bob's name is already present in the recipient set: If $\exists \psi_i \in \Psi : \text{NAME}_i = \text{NAME}_b$, *May exit*.

(5) Add Bob: $\Psi' = \Psi \cup \{\psi_b\} = \{\psi_0, \psi_1, \dots, \psi_{n-1}, \psi_b\}$.

(6) Encrypt the modified ECF: $\mathcal{E}' \leftarrow \text{Enc}^{\text{ECF}}(\Psi', \Omega)$.

The name of a recipient is for information purposes only, which results in Step 4 to be optional. However, in order to avoid (human) confusion about the recipients of an ECF, it is advisable to enforce the constraint that all recipient names must be unique. Our PoC implementation does not allow duplicate names by default but can be configured otherwise.

We now assume that Alice added her coworkers Bob and Charlie as well as the automated deployment job as recipients to the ECF, thus, following the basic workflow described in Subsection III-C.

Next, Alice wants to remove Bob from the recipient set as motivated in Subsection III-D. She has to identify Bob in the recipient set either by his public key or his name. If Alice uses Bob's name to identify him, his name must be unique within the ECF. Alice performs the following steps to remove Bob from the recipient set:

(1) Decrypt the ECF: $(\Psi, \Omega) = \text{Dec}^{\text{ECF}}(\text{sk}_A^S, \mathcal{E})$.

(2) Find $\psi_b \in \Psi$ based on pk_B^S and/or NAME_b .

If no $\psi_i \in \Psi$ matches, Bob is not a recipient of \mathcal{E} , *Exit*.

(3) Optionally check, if Alice tries to remove herself from the ECF, i.e., Alice and Bob are identical.

If $\text{pk}_B^S \stackrel{?}{=} \text{pk}_A^S = \text{Convert}^{\text{sk} \rightarrow \text{pk}}(\text{sk}_A^S)$, *May exit*.

(4) Remove Bob: $\Psi' = \Psi \setminus \{\psi_b\}$.

(5) Encrypt the modified ECF: $\mathcal{E}' \leftarrow \text{Enc}^{\text{ECF}}(\Psi', \Omega)$.

As noted in Subsection V-A, the final Step 5 does not require any private keys for the encryption procedure. This implies that recipients of an ECF are able to remove themselves. However, this is generally an undesired feature and, therefore, the implementation should prevent such action by executing the optional Step 3. Our PoC implementation of the command line tool does not allow any recipient self-removal.

Finally, it must be noted that the restriction explained in Subsection III-D still holds: All former recipients are always able to access old versions of an ECF. As stated earlier, this is unpreventable and should be mitigated by using short-lived secrets and by applying the principle of minimal privilege.

D. Changing the Content

Modifying the confidential data stored within an ECF is similar to modifying the recipient set. Following the workflow described in Subsection III-E, we describe how the recipient Charlie can change the content of an ECF. We denote the resulting bit string after modification as $\Omega' = \text{Modify}(\Omega)$ and Charlie's private key for signing as sk_C^S :

(1) Decrypt the ECF: $(\Psi, \Omega) = \text{Dec}^{\text{ECF}}(\text{sk}_C^S, \mathcal{E})$.

(2) Modify the confidential data: $\Omega' = \text{Modify}(\Omega)$.

(3) Encrypt the modified ECF: $\mathcal{E}' \leftarrow \text{Enc}^{\text{ECF}}(\Psi, \Omega')$.

Although possible, it is uncommon to change both the recipient set Ψ and the content Ω in a single operation. This is why our PoC command line tool offers both actions only separately.

E. Generating $m - n$ Deception Blocks

In order to achieve the goal of obfuscating the information about recipients in the public header of an ECF, we take two measures. First, the ECF format stores the names of recipients in the encrypted private body, thus, hiding them from external parties. Furthermore, if the public keys of users are not published to the open public, i.e., only visible to members of the same institution or company, external parties cannot determine if a given user is a recipient of an ECF. This is because without knowledge of the public key, one cannot compute the `id_tag` (cf. Step 4 in Subsection V-B) and therefore not search for a recipient-specific block within the public header of a given ECF.

Second, we opted to obfuscate the true number of recipients n in the public header. This is why $m \geq n$ gets chosen randomly in Step 2 in Subsection V-A. Our PoC implementation chooses m randomly depending on n , such that $\max\{8, 2n\} \geq m \geq n$. After generating n valid recipient-specific blocks in Step 3 in the encryption procedure, $m - n$ deception blocks must be generated to fill the public header and obtain a valid ECF.

The main purpose of these deception blocks is to obfuscate n . Therefore, these blocks must be generated in a way, such that an external party is not able to distinguish between real blocks and deception blocks. This is why the deception blocks should not be random bit strings because there is a possibility that the outputs of the used cryptographic algorithms used to generate the real blocks suffer from statistical biases.

To avoid distinguishability, we suggest a generation procedure as follows:

- (1) Generate a new deception key pair:
 $(sk_\phi^S, pk_\phi^S) \leftarrow \text{Gen}^S$, $pk_\phi^X = \text{Convert}^{S \rightarrow X}(pk_\phi^S)$.
- (2) Compute IDENTIFICATION TAG
 $id_tag = H(pk_\phi^S || \theta) [0, \dots, 15]$.
- (3) Generate an ephemeral key pair for key exchange
 $(sk_e^X, pk_e^X) \leftarrow \text{Gen}^X$.
- (4) Generate a random symmetric key $k_\phi \leftarrow \text{Gen}^{\text{SYM}}$.
- (5) Execute the key exchange algorithm:
 $ss = \text{Kex}(sk_e^X, pk_\phi^X)$.
- (6) Calculate $k_{pre2} = H(ss || pk_\phi^X || pk_e^X) [0, \dots, y - 1]$.
- (7) Calculate $k_{pre1} = k_\phi \oplus k_{pre2}$.
- (8) Add $r_\phi = id_tag || pk_e^X || k_{pre1}$ to R .

However, generating deception blocks using this procedure is computationally more expensive than generating real blocks. In order to simplify the procedure, decrease the computational load, and therefore increase runtime performance, our PoC implementation uses a different one. This is only possible if the used cryptographic hash function generates truly random looking bit strings, i.e., the outputs of Steps 2 and 6 could have been generated at random instead. This implies that the output of Step 7 also could have been generated at random, thus, eliminating most of the computationally expensive parts of the procedure. The assumption regarding the hash function is justified considering the lengths of the inputs, cf. [26]. The simplified version of the deception block generation procedure reads as follows:

- (1) Generate an ephemeral key pair for key exchange
 $(sk_e^X, pk_e^X) \leftarrow \text{Gen}^X$.
- (2) Generate a random IDENTIFICATION TAG:
 $id_tag \leftarrow \text{Random}(16)$.
- (3) Generate a random $k_{pre1} \leftarrow \text{Random}(y)$.
- (4) Add $r_\phi = id_tag || pk_e^X || k_{pre1}$ to R .

Of course, the generation of deception blocks does take some time, but the overhead is in the range of milliseconds as shown in Subsection VI-C.

VI. RECENT WORK AND IMPROVEMENTS

Since the original release of the ECF format and the PoC implementation in [1], we extended and improved our work in different regards. The changes, enhancements, and new experiments are presented in the following subsections.

A. Building on Linux and Dockerfile

The original version of our PoC implementation was targeted towards the Windows operating system. It could be cross-compiled for Linux, but building on Linux was not possible. However, enabling compilation on Linux was always a goal. The changes we made since the original release allow our PoC implementation to be compiled on both operating systems and cross-compiled to the respective other operating system. Furthermore, using the *single-file deployment* [27] feature of recent .NET [28] versions makes our PoC implementation portable and easy to distribute without the need of installing further dependencies. Additionally to the build on Linux, we provide a *Dockerfile* for compiling in a *Docker* [29] container.

Software and its dependencies need to be kept up-to-date, which is why we updated the targeted runtimes. This means that our PoC implementation now can be compiled for .NET 7 and .NET 8. We dropped support for .NET 6 in order to be able to use newer language features. The aforementioned *Dockerfile* uses the most recent runtime. Next, we updated all libraries used for our PoC implementation. This includes the *Sodium* library [30], which received new features, such as the support for the symmetric encryption scheme family AEGIS [21][31]. The wrapper library *NSec* [32] was also updated in order to be able to use the new algorithms in *Sodium*.

B. New Cipher Suite

As mentioned in Subsection IV-C, we introduced two new cipher suites to the PoC implementation. Namely, the symmetric encryption scheme AEGIS-256 [21][31] was added to the set of cipher suites. Although the specifications for this new algorithm family are very recent and still work-in-progress, the authors of *Sodium* recommend using AEGIS-256 over AES-256-GCM [33]. It is important to note that our use of AES-256-GCM never reaches the theoretical limit stated in [33], because the encryption procedure generates the symmetric encryption key every time anew (cf. Subsection V-A).

Because of the changing nature of the AEGIS specification, we opted to keep the default cipher suite in our PoC implementation. However, users of our command line tool may select the cipher suite that best suit their needs. Furthermore, our performance experiments in the following Subsection VI-C hint that Cipher Suite IV using AEGIS-256 as symmetric encryption scheme offers currently the best performance.

The flexibility of the ECF format allows for different cipher suites to be implemented and the source code of our PoC implementation is structured in a way to support the easy integration of more cryptographic algorithms. Therefore, implementing new cipher suites is straight-forward. This makes the ECF format suitable in the long term since algorithms that are considered weak or are cracked in the future can be easily replaced by secure ones.




Figure 7. Average encryption time for different content sizes, cipher suites, and runs.

C. Performance Analysis

Extending our previous work [1], we took performance measurements of our PoC implementation on two different machines and in total using two different runtimes. The raw data of all runs is published alongside the code repository on GitHub. For the charts and tables in this paper, we aggregate the data by using average values. We provide the numeric values used to generate the performance charts in Appendix C.

The first experiment encrypts ECFs with $n = 5$ recipients and different content sizes ranging from 1 MiB to 1000 MiB. The results are depicted in Figure 7. We chose a log-log scale plot as both the content sizes and the amount of time it takes to encrypt the ECF increase in orders of magnitudes. Although not directly visible in the plot, there is a linear relationship between the content size and the processing time, which can be validated using the numerical values in Table IV in Appendix C. The differences between the runtimes *.NET 7.0* and *.NET 8.0* (left vs. middle chart) are negligible while a stronger processor (right chart) results in better absolute values.

Similarly, the chosen cipher suite makes a difference in performance. While the two symmetric encryption algorithms implemented in our PoC perform almost identical, the hash function SHA-512 (Cipher Suites II and IV) outperforms its sibling SHA-256 (Cipher Suites I and III) regardless of content size, number of recipients, runtime, and processor in all of our experiments. This is to be expected as the throughput of SHA-512 is up to 1.6 times [34] as much as the throughput of SHA-256. Since our PoC implementation uses the cryptographic library *Sodium* [30], which focuses on portability and therefore does not use any specialized processor instructions for SHA-256, our observed performance difference of about a factor of 1.5 seems reasonable.

In a second experiment, we focus on the decryption speed and re-use the same parameters of $n = 5$ and the content sizes of our previous experiment. Figure 8 shows the results of this second experiment. We can see the same linear relationship between the content size and the processing time. The absolute decryption times are slightly lower than those for

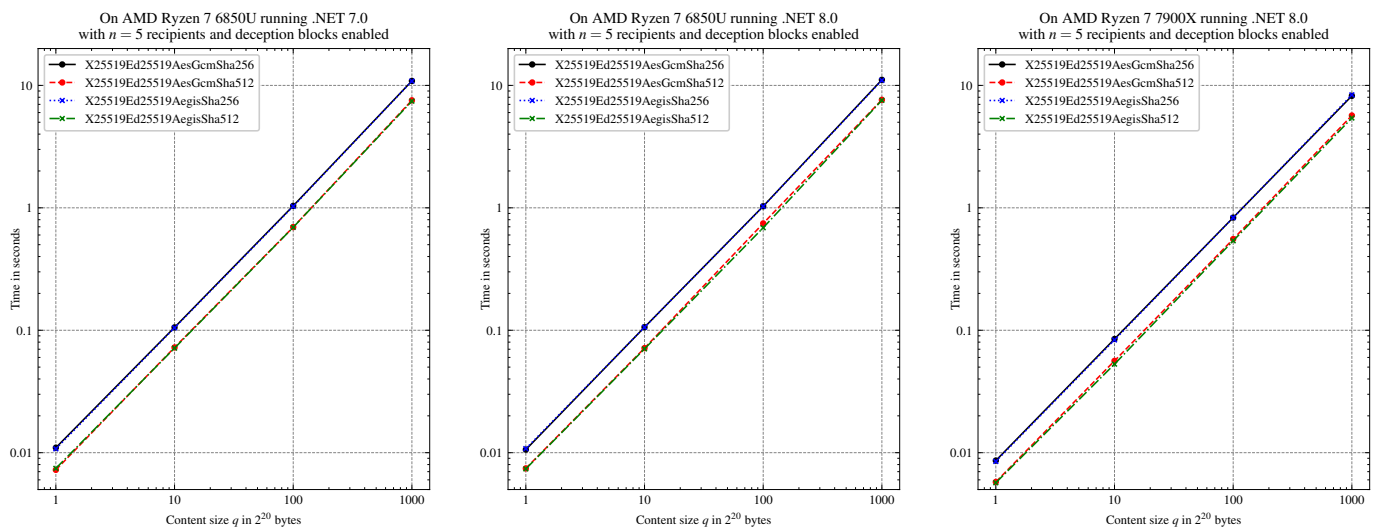


Figure 8. Average decryption time for different content sizes, cipher suites, and runs.




Figure 9. Average encryption and decryption time for different number of recipients, cipher suites, and whether deception blocks are enabled.

encryption, which must be due to the symmetric encryption and decryption algorithms. We can see the same speedup when using SHA-512 over SHA-256 as well as using a stronger processor and we observe little to no difference between the two symmetric encryption algorithms.

Our third experiment varies the number of recipients n from 5 to 1000 in order to make statements about the performance impact of generating the recipient-specific blocks in the public header. Furthermore, we can show that generating deception blocks increases the encryption time only slightly. We show the results in Figure 9, which are linear-linear plots and omit the samples for $n = 1000$ for clarity. The numerical values are listed in Table VI in Appendix C. Deception blocks have no impact on the decryption, which is why the two decryption charts differ only in run-to-run variance. All four charts show a linear relationship between the number of recipients and the average processing time.

In a fourth experiment, we focused on the performance impact of validating all recipient signatures during decryption. Similar to the previous experiment, the number of recipients was varied while the content size was fixed to 1 MiB. Activating the validation – essentially executing Step 12 in Subsection V-E – results in a linear performance hit as depicted in Figure 10. Deactivating the validation step does not remove the linear dependency completely since the recipients still have to be parsed as stated in Subsection IV-E. It is, however, drastically decreased, offering a performance benefit to certain applicable situations described earlier. In general, we recommend keeping the recipient signature validation enabled.




Figure 10. Average decryption time for different number of recipients, cipher suites, and whether recipient signature validation is enabled.

Accomplishing our goal of performing the encryption and decryption procedure in less than 100 ms (cf. Subsection IV-B) is therefore dependent on the amount of confidential data stored within an ECF as well as the number of recipients. Our experiments show that using a modern processor and choosing a cipher suite with SHA-512 as the hash function can yield to execution times below 100 ms for content sizes smaller than 10 MiB. As mentioned in Subsection III-D, an ECF is designed to store passwords, certificate keys, and other similar credentials for a small set of recipients. The size of these types of content is typically in the range of 2^5 to 2^{15} bytes. Therefore, even the smallest content size in our experiments, 1 MiB = 2^{20} bytes, is more than sufficient for typical ECF usage. Additionally, even though the ECF format supports far more than 1000 recipients, typically there are only a few recipients per file. Our experiments show that both the encryption and decryption procedure take less than 15 ms each with 1 MiB-sized content and 50 recipients. Therefore, we can safely state that the PoC implementation fulfills the objective.

D. Private Key Management

A major change was made to the management and storing of the private keys. While our original solution featured basic password-based encryption using a Key Derivation Function (KDF), we improved the implementation to support user-chosen parameters. Furthermore, the former static and pre-defined private key storage format was replaced by a more flexible one. This enables our PoC implementation to be extended in the future with cipher suites using different Diffie-Hellman-like key exchange and matching signature algorithms. We discuss the option of choosing the key exchange algorithm independently from the signature algorithm in Appendix A.

In order for the ECF format to work, three components must be selected to match: the cipher suite, a private key, and the recipient information in the private body of an ECF. This is why our PoC implementation provides three base types CipherSuite, ECFKey, and Recipient. Then, the three subtypes CSX25519Ed25519Base, EKX25519Ed25519, and RX25519Ed25519 are interlocked to form a group of matching components. Our implementation ensures this way that the correct type of key is used and the parsing of the recipient information blocks (cf. Figure 6) is done correctly.

The newly added, flexible format for private key storage is depicted in Figure 11. It starts similar to a public header with a VERSION to ensure future extensibility. The next three fields, KEY TYPE, SYMMETRIC ENCRYPTION TYPE, and KDF TYPE define which type of private key is stored, which symmetric encryption scheme was used to encrypt it, and which KDF was used to derive the symmetric encryption key from the user-chosen password. Next, a SALT that is used for key derivation is stored, followed by a SYMMETRIC NONCE used for encrypting and decrypting the private key.

Depending on the KDF, its configuration is stored next in the field KDF CONFIG. For the default algorithm Argon2id [35], three values are stored: the number of iterations, the memory size, and the degree of parallelism. These parameters can be chosen freely by the user as long as they meet the specification. We provide sensible defaults in our PoC command line tool.

Finally, the encrypted private key is appended last. Its length is dependent on the cipher suite and the chosen symmetric encryption algorithm. All implemented cipher suites use an Ed25519 private key as their basis, since it can be converted into an Ed25519 public key and furthermore into both a X25519 private key and a X25519 public key. Therefore, it is only necessary to store the Ed25519 private key. However, when using different cipher suites (cf. Appendix A), it might be necessary to store multiple keys, which is why the storage format is designed as flexible as possible.

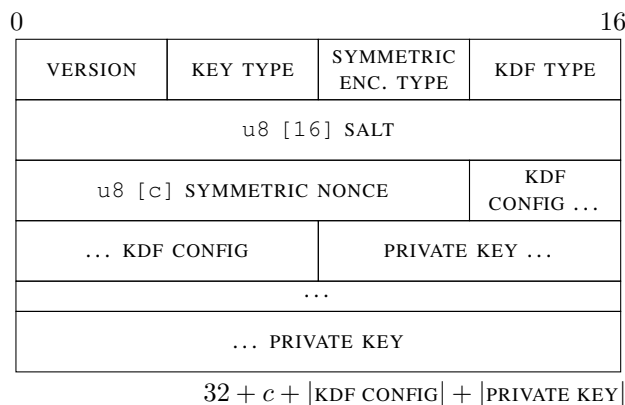


Figure 11. Private key storage format.

To keep the private key storage format simple, we opted to not include an encrypted hash value as a field. Instead, we utilize the “associated data” option of the AEAD-based symmetric encryption algorithm. All unencrypted fields are used as input to the “associated data” option, which means the PoC can detect modifications to any of the fields. This approach limits the choice of symmetric encryption schemes to AEAD-based ones. However, since both symmetric encryption schemes used within our PoC implementation are indeed AEAD-based, there is actually no restriction. The user can choose which symmetric encryption scheme and which KDF to use in order to protect the private key. Our PoC implementation uses AES-256-GCM [18][19] and Argon2id [35] by default.

VII. CONCLUSION AND FUTURE WORK

In this paper, we extended our previous work of Encrypted Container File in [1] and described the ECF format in more detail. Additionally, we presented an example use case workflow in Section III, which motivated the requirements, the design goals, and ultimately the ECF format presented in Section IV. To demonstrate the ECF format flexibility we claimed in our previous work, we implemented two new cipher suites in our Proof of Concept implementation and focused on a cipher suite-independent presentation of the ECF format and operations in Section V. Although all discussed cipher suites use a key exchange and a signature algorithm based on Curve25519 [16] [17], we provide guidance for implementing ECF with different algorithms in Appendix A.

Section VI describes more changes and improvements made to our PoC implementation, which is now native to Linux and sports a new and improved private key management. We also took performance measurements in order to validate our claim of providing a fast-enough implementation to be used in production. We are confident that our work is ready to be used in production.

The full code of the updated and extended PoC implementation as well as unit tests, performance tests, and performance analysis data for that code are available at:

<https://github.com/Hirmmoder/ECF>

The modular structure of our PoC implementation allows for an easy extension and the implementation of more features and functionalities. We are looking forward to the feedback from the community to further improve the ECF format and its implementation.

Future work may focus on finding further use cases and applications of ECF. For example, the subject of applying ECF or ECF-based solutions in diverse contexts, such as within public authorities, raises interesting research questions and may lead to fruitful improvements of the ECF format as well as our PoC implementation.

Another area of work could be the implementation of a Public Key Infrastructure (PKI)-based trust model amongst the recipients of an ECF. Such extension could help establishing common trust anchors in multi-organization settings. In addition to that, changes to the ECF format could be made to prove content modification ownership by employing digital signatures.

Furthermore, it may be possible to design a two-staged access control framework which allows a set of recipients to read the confidential data, but only a subset of those recipients to alter it. Such amendment could help distinguishing between different roles of entities accessing an ECF.

Finally, the authors would like to thank the International Academy, Research, and Industry Association for the opportunity to present the idea of Encrypted Container Files in more detail, supplemented by current developments of the project.

REFERENCES

- [1] Tobias J. Bauer and Andreas Aßmuth, "Encrypted Container File: Design and Implementation of a Hybrid-Encrypted Multi-Recipient File Structure," in *Proc of the 14th International Conference on Cloud Computing, GRIDs, and Virtualization (Cloud Computing 2023)*, Nice, France, Jun. 2023.
- [2] M. Fowler, "Continuous Integration," Jan. 18, 2024, [Online]. Available: <https://www.martinfowler.com/articles/continuousIntegration.html> (visited on 2024.05.27).
- [3] Intellectsoft, "Best Software Development Trends to Follow in 2024," Intellectsoft, Apr. 19, 2024, [Online]. Available: <https://www.intellectsoft.net/blog/software-development-trends/> (visited on 2024.05.27).
- [4] A. Lomas, "The Future of Software Development in 2024 and Beyond," Net Solutions, Mar. 10, 2023, [Online]. Available: <https://www.netsolutions.com/insights/the-future-of-software-development/> (visited on 2024.05.27).
- [5] P. Bajpai and A. Lewis, "Secure Development Workflows in CI/CD Pipelines," in *2022 IEEE Secure Development Conference (SecDev)*, 2022, pp. 65–66. DOI: 10.1109/SecDev53368.2022.00024.
- [6] J. Walsh, "Secure CI/CD Pipelines: Practices for Managing CI/CD Secrets," CyberArk, Feb. 15, 2023, [Online]. Available: <https://developer.cyberark.com/blog/secure-ci-cd-pipelines-best-practices-for-managing-ci-cd-secrets/> (visited on 2024.05.27).
- [7] N. Sakovich, "Top 10 Software Development Trends of 2024," SaM Solutions, Apr. 8, 2024, [Online]. Available: <https://www.sam-solutions.com/blog/software-development-trends/> (visited on 2024.05.27).
- [8] P. Belagatti, "Managing Secrets in Your DevOps Pipeline," The New Stack, Aug. 4, 2022, [Online]. Available: <https://thenewstack.io/managing-secrets-in-your-devops-pipeline/> (visited on 2024.05.27).
- [9] The DevSecOps Guides Authors, "Secure Pipeline," DevSecOps Guides, Jun. 2, 2023, [Online]. Available: <https://devsecopsguides.com/docs/code/secure-pipeline/> (visited on 2024.05.27).
- [10] I. Koishybayev *et al.*, "Characterizing the Security of Github CI Workflows," in *31st USENIX Security Symposium (USENIX Security 22)*, Boston, MA: USENIX Association, 2022, pp. 2747–2763, ISBN: 978-1-939133-31-1.
- [11] Git, "Git," 2024, [Online]. Available: <https://git-scm.com/> (visited on 2024.05.27).
- [12] Dispel LLC, "jak – simple git encryption," Dispel LLC, 2017, [Online]. Available: <https://jak.readthedocs.io/en/latest/> (visited on 2024.05.27).
- [13] A. Ayer, "git-crypt – transparent file encryption in git," 2023, [Online]. Available: <https://www.agwa.name/projects/git-crypt/> (visited on 2024.05.27).
- [14] E. Barker, *Recommendation for Key Management: Part 1 – General*. May 2020. DOI: 10.6028/nist.sp.800-57pt1r5.
- [15] J. Saltzer and M. Schroeder, "The protection of information in computer systems," *Proceedings of the IEEE*, vol. 63, no. 9, pp. 1278–1308, 1975. DOI: 10.1109/PROC.1975.9939.
- [16] D. J. Bernstein, "Curve25519: New diffie-hellman speed records," in *Public Key Cryptography - PKC 2006*, Springer Berlin Heidelberg, 2006, pp. 207–228.
- [17] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang, "High-speed high-security signatures," *Journal of Cryptographic Engineering*, vol. 2, no. 2, pp. 77–89, Aug. 2012.
- [18] D. A. McGrew and J. Viega, "The galois/counter mode of operation (GCM)," National Institute of Standards and Technology, May 31, 2005, [Online]. Available: <https://csrc.nist.rip/groups/ST/toolkit/BCM/documents/proposedmodes/gcm/gcm-revised-spec.pdf> (visited on 2024.05.27).
- [19] M. J. Dworkin, "Recommendation for block cipher modes of operation: Galois/counter mode (gcm) and gmac," Tech. Rep., Nov. 2007. DOI: 10.6028/nist.sp.800-38d.
- [20] Q. H. Dang, "Secure Hash Standard," Tech. Rep., Aug. 2015. DOI: 10.6028/nist.fips.180-4.
- [21] F. Denis and S. Lucas, "The AEGIS Family of Authenticated Encryption Algorithms," Internet Engineering Task Force, Internet-Draft draft-irtf-cfrg-aegis-aead-10, Jan. 20, 2024, Work in Progress.
- [22] D. McGrew, *An Interface and Algorithms for Authenticated Encryption*, RFC 5116, Jan. 2008. DOI: 10.17487/RFC5116.
- [23] "Squeamish Ossifrage", "Curve25519 over Ed25519 for key exchange? Why?" Crypto Stack Exchange, Mar. 19, 2019, [Online]. Available: <https://crypto.stackexchange.com/a/68129> (visited on 2024.05.27).
- [24] The Sodium Authors, "Ed25519 to curve25519 – libsodium," Oct. 27, 2023, [Online]. Available: <https://doc.libsodium.org/advanced/ed25519-curve25519> (visited on 2024.05.27).
- [25] The Sodium Authors, "Point*scalar multiplication – libsodium," Sep. 12, 2023, [Online]. Available: https://doc.libsodium.org/advanced/scalar_multiplication (visited on 2024.05.27).
- [26] "fgrieu", "Using SHA2 as random number generator?" Crypto Stack Exchange, Jun. 20, 2020, [Online]. Available: <https://crypto.stackexchange.com/a/81459> (visited on 2024.05.27).
- [27] Microsoft Learn, "Single-file deployment," Microsoft, Mar. 11, 2023, [Online]. Available: <https://learn.microsoft.com/en-us/dotnet/core/deploying/single-file/overview> (visited on 2024.05.27).
- [28] Microsoft, ".net – build. test. deploy.," Microsoft, 2024, [Online]. Available: <https://dotnet.microsoft.com/en-us/> (visited on 2024.05.27).
- [29] Docker Inc., "Docker: Accelerated Container Application Development," Docker Inc., 2024, [Online]. Available: <https://www.docker.com/> (visited on 2024.05.27).
- [30] The Sodium Authors, "Introduction – libsodium," May 25, 2024, [Online]. Available: <https://doc.libsodium.org/> (visited on 2024.05.27).
- [31] The Sodium Authors, "AEGIS-256 – libsodium," Nov. 1, 2023, [Online]. Available: https://doc.libsodium.org/secret-key_cryptography/aead/aegis-256 (visited on 2024.05.27).
- [32] K. Hartke, "NSec – Modern cryptography for .NET Core," 2024, [Online]. Available: <https://nsec.rocks/> (visited on 2024.05.27).
- [33] The Sodium Authors, "AES256-GCM – libsodium," Sep. 12, 2023, [Online]. Available: https://doc.libsodium.org/secret-key_cryptography/aead/aes-256-gem (visited on 2024.05.27).
- [34] R. Frame, "SHA-512 faster than SHA-256?" Crypto Stack Exchange, Jun. 17, 2015, [Online]. Available: <https://crypto.stackexchange.com/a/26351> (visited on 2024.05.27).
- [35] A. Biryukov, D. Dinu, and D. Khovratovich, "Argon2: The memory-hard function for password hashing and other applications," version 1.3, Mar. 24, 2017, [Online]. Available: <https://raw.githubusercontent.com/P-H-C/phc-winner-argon2/master/argon2-specs.pdf> (visited on 2024.05.27).

APPENDIX

A. Choosing the Key Exchange Algorithm Independently from the Signature Algorithm

The ECF format as well as our PoC implementation use a well-defined conversion method to convert an Ed25519 private or public key to its corresponding X25519 private or public key [23][24], respectively. This implies that storing, using, and having access to the Ed25519-version of a private or public key is enough to carry out all signing and key agreement operations described in this paper. In order to save memory and remove unnecessary redundancy, we opted to employ this conversion method throughout the ECF format and operations.

However, different key exchange and signature algorithms may not support such key conversion. In this case, the recipient information block stored inside the private body of an ECF (cf. Figure 6) must include not one but two public keys: one for the signature algorithm and one for the key exchange algorithm. We depict the altered structure in Figure 12 and assume for simplicity reasons that both public keys have the same length of u bytes. This is, however, not a requirement.

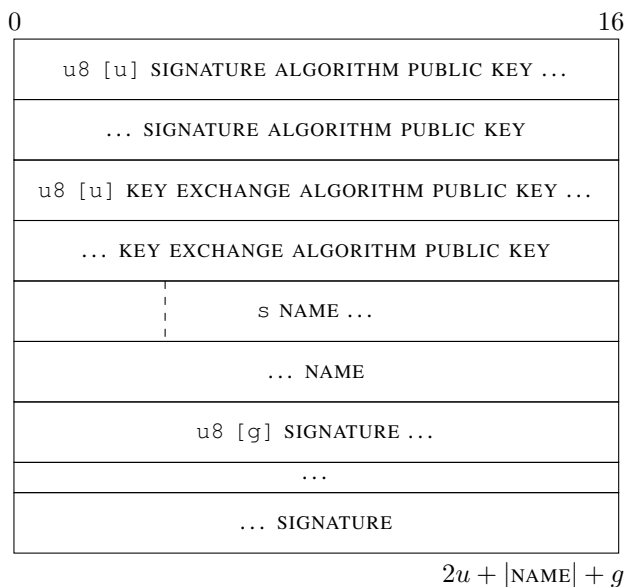


Figure 12. Recipient information block structure within the private body of an ECF when no key conversion method is available to convert between the keys of the key exchange and the signature algorithm.

The field SIGNATURE now contains a signature not only over the NAME, but over the combination of KEY EXCHANGE ALGORITHM PUBLIC KEY and NAME. This is compulsory in order to create a strong connection between the public key for key exchange and the human-readable name of a recipient.

The presented changes in data structure also need to be reflected in the ECF operations. The recipient entry is now constructed as follows:

$$\psi = \text{pk}_{\psi}^S \parallel \text{pk}_{\psi}^X \parallel \text{NAME} \parallel t$$

$$t = \text{Sign}(\text{sk}_{\psi}^S, \text{pk}_{\psi}^X \parallel \text{UTF8}(\text{NAME}))$$

During encryption, pk_{ψ}^X is already present and, therefore, Step 3c in Subsection V-A is omitted. In Step 6, the extended recipient information block structure is used.

The decryption procedure (cf. Subsection V-B) is extended and now needs three input parameters:

$$(\Psi, \Omega) = \text{Dec}^{\text{ECF}}(\text{sk}^S, \text{sk}^X, \mathcal{E})$$

Similarly to the encryption procedure, the first conversion in Step 3 in Subsection V-B is omitted. Next, the deconstruction in Step 12a produces four values $\text{pk}_{\psi_i}^S$, $\text{pk}_{\psi_i}^X$, NAME_i , and t_i and the verification in Step 12b needs to be adjusted, too:

$$\text{If } \text{Vrfy}^S(\text{pk}_{\psi_i}^S, \text{pk}_{\psi_i}^X \parallel \text{UTF8}(\text{NAME}_i), t_i) \stackrel{?}{=} 0, \text{Exit.}$$

Adding new recipients to an ECF requires them to generate their recipient entries as shown in Subsection V-C. The computations for t_b and ψ_b need to be adjusted accordingly and include the public key for key exchange as described above. This is also true for the verification in Step 1 in Subsection V-C. The mandatory check in Step 3 now needs to check for both pk_B^S and pk_B^X .

The optimized procedure for generating deception blocks, which is described in Subsection V-E, does not require any modifications. This is because the IDENTIFICATION TAG is generated randomly instead of being derived from a key pair and, thus, no key conversion is necessary.

B. Variables and Functions

Table I contains all simple variables used primarily to describe the ECF format. The encryption and decryption procedures make use of various cryptographic algorithms and their respective input and output parameters. These are defined and described in Table II. All used functions are listed in Table III. A randomized or probabilistic algorithm is indicated by the left arrow " \leftarrow ", whereas an equal sign " $=$ " indicates deterministic execution.

TABLE I. Variables.

Variable	Description
a	length of KEY AGREEMENT INFORMATION in bytes
b	length of the encrypted private body in bytes
b'	length of the unencrypted private body in bytes
c	length of SYMMETRIC NONCE in bytes
d	length of the output of the hash function in bytes
g	length of a signature in bytes
h	length of the public header in bytes
m	number of recipients in the public header
n	number of true recipients in the private body
o_i	offset of recipient-specific block i within the file in bytes
q	length of CONTENT in bytes
\bar{q}	length of all private body fields except CONTENT in bytes
u	length of a public key in bytes
v	length of a private key in bytes
y	length of the symmetric key in bytes

TABLE II. Component Variables.

Variable	Description
$pk_A^X \in \{0, 1\}^{8u}$	public <u>k</u> ey of <u>A</u> lice for key <u>e</u> xchange
$sk_B^X \in \{0, 1\}^{8v}$	private/ <u>s</u> ecret <u>k</u> ey of <u>B</u> ob for key <u>e</u> xchange
$k_{final} \in \{0, 1\}^{8y}$	final symmetric encryption <u>k</u> ey
$\alpha \in \{0, 1\}^{8c}$	symmetric nonce
$\theta \in \{0, 1\}^{8 \cdot 16}$	salt
$t \in \{0, 1\}^{8g}$	a signature t is g bytes long
$\beta \in \{0, 1\}^{8d}$	hash function's output is always d bytes
$r = id_tag pk_c^X k_{pre1}$	recipient in public header
$R = \{r_0, \dots, r_{m-1}\}$	recipient set in public header
$\psi = pk_\psi^S NAME t$	recipient in private body
$\Psi = \{\psi_0, \dots, \psi_{n-1}\}$	recipient set in private body
$\Omega \in \{0, 1\}^{8q}$	confidential data stored in an ECF, q bytes long
$\mathcal{E} = \mathcal{H} \mathcal{B} \beta_{all}$	an ECF consists of a public header \mathcal{H} , a private body \mathcal{B} , and a public footer β_{all}

TABLE III. Functions.

Function Definition	Description
$x = x_1 x_2$	concatenation of two bit strings
$x = x_1 \oplus x_2$	bitwise exclusive OR (XOR) operation on two same-length bit strings
$x' = x[0, \dots, j-1]$	truncation of the bit string $x \in \{0, 1\}^{l \geq 8j}$ to the first j bytes
$(sk^S, pk^S) \leftarrow Gen^S$	generates a key pair for <u>s</u> igning
$(sk^X, pk^X) \leftarrow Gen^X$	generates a key pair for key <u>e</u> xchange
$sk^X = Convert^{S \rightarrow X}(sk^S)$	converts a private/public key for signing to a private/public key for key exchange (*)
$pk^X = Convert^{S \rightarrow X}(pk^S)$	converts a private/public key for key exchange (*)
$pk^S = Convert^{sk \rightarrow pk}(sk^S)$	converts a private key to a public key
$pk^X = Convert^{sk \rightarrow pk}(sk^X)$	converts a private key to a public key
$k \leftarrow Gen^{SYM}$	generates a symmetric encryption key
$x = UTF8(s)$	converts the character string s without BOM into the UTF-8 bit string x
$x = Random(j)$	generates a random j -byte long bit string x
$ss = Kex(sk_A^X, pk_B^X)$ $= Kex(sk_B^X, pk_A^X)$	obtain a shared <u>s</u> ecret by performing the <u>k</u> ey <u>e</u> xchange with a secret and a public key
$t \leftarrow Sign(sk_A^S, x)$	signing the bit string x with <u>A</u> lice's private key for signing yields a signature t
$Vrfy^S(pk_A^S, x, t) \in \{0, 1\}$	returns 1, if and only if the signature t is valid for the given public key and the bit string x , otherwise 0
$\gamma \leftarrow Enc^{SYM}(k, \alpha, x)$	<u>e</u> ncrypting the bit string x with a symmetric encryption key k and a nonce α produces a ciphertext bit string γ
$x' = Dec^{SYM}(k, \alpha, \gamma)$	<u>d</u> ecrypting a ciphertext bit string γ with the correct symmetric encryption key k and the correct nonce α produces the bit string x'
$\beta = H(x)$	hash functions take any-length bit strings x and produce a fixed-sized output β
$Vrfy^H(x, \beta) \in \{0, 1\}$	returns 1, if and only if the hash value β matches the hash value of the bit string x , otherwise 0

(*) This behavior is specific for Ed25519 [17] and X25519 [16] using the conversion method described in [23][24]. However, if a cipher suite does not allow such conversion, cf. Appendix A.

C. Performance Analysis Value Tables

This subsection provides the numerical values of the performance charts in Subsection VI-C. All tables contain the aggregated data used to draw the performance charts. However, Tables VI and VII contain additional data points for $n = 1000$ recipients, showing that our PoC implementation is capable of handling even such unrealistic parameters efficiently. The raw data and the code to aggregate and plot the performance measurements is published on GitHub.

TABLE IV. Average encryption time (in seconds) for different content sizes and cipher suites. Visualized in Figure 7.

Content Size	Run	Cipher Suite			
		I	II	III	IV
1 MiB	R7 6850U (.NET 7.0)	0.012	0.010	0.014	0.010
	R7 6850U (.NET 8.0)	0.012	0.010	0.014	0.010
	R7 7900X (.NET 8.0)	0.010	0.008	0.010	0.008
10 MiB	R7 6850U (.NET 7.0)	0.136	0.100	0.144	0.103
	R7 6850U (.NET 8.0)	0.140	0.105	0.141	0.106
	R7 7900X (.NET 8.0)	0.103	0.075	0.104	0.074
100 MiB	R7 6850U (.NET 7.0)	1.210	0.883	1.295	0.983
	R7 6850U (.NET 8.0)	1.265	0.928	1.250	0.913
	R7 7900X (.NET 8.0)	0.934	0.683	0.950	0.680
1000 MiB	R7 6850U (.NET 7.0)	12.822	9.490	12.740	9.508
	R7 6850U (.NET 8.0)	12.711	9.665	12.726	9.650
	R7 7900X (.NET 8.0)	9.449	7.153	9.782	6.948

TABLE V. Average decryption time (in seconds) for different content sizes and cipher suites. Visualized in Figure 8.

Content Size q	Run	Cipher Suite			
		I	II	III	IV
1 MiB	R7 6850U (.NET 7.0)	0.011	0.007	0.011	0.007
	R7 6850U (.NET 8.0)	0.011	0.007	0.011	0.007
	R7 7900X (.NET 8.0)	0.009	0.006	0.008	0.006
10 MiB	R7 6850U (.NET 7.0)	0.106	0.072	0.105	0.071
	R7 6850U (.NET 8.0)	0.106	0.071	0.106	0.070
	R7 7900X (.NET 8.0)	0.085	0.056	0.083	0.053
100 MiB	R7 6850U (.NET 7.0)	1.038	0.696	1.028	0.693
	R7 6850U (.NET 8.0)	1.030	0.746	1.028	0.687
	R7 7900X (.NET 8.0)	0.832	0.555	0.831	0.538
1000 MiB	R7 6850U (.NET 7.0)	10.882	7.557	10.836	7.436
	R7 6850U (.NET 8.0)	11.113	7.626	10.992	7.559
	R7 7900X (.NET 8.0)	8.212	5.677	8.366	5.415

TABLE VI. Average encryption and decryption times (in milliseconds) for different number of recipients and cipher suites. Visualized in Figure 9.

n	Operation	Deception blocks	Cipher Suite			
			I	II	III	IV
5	Enc	✓	9.671	7.745	10.493	7.659
		✗	10.506	7.688	10.424	7.552
	Dec	✓	8.622	5.771	8.490	5.679
		✗	8.594	6.032	8.364	5.784
10	Enc	✓	10.050	7.630	10.253	8.061
		✗	10.363	7.593	10.205	7.211
	Dec	✓	8.676	6.172	8.771	6.011
		✗	8.675	6.111	8.859	5.858
20	Enc	✓	10.948	9.204	11.154	8.992
		✗	11.680	8.909	11.457	8.451
	Dec	✓	9.065	6.580	9.141	6.285
		✗	9.038	6.507	9.264	6.216
50	Enc	✓	15.023	12.688	14.358	12.424
		✗	14.664	12.076	14.706	11.646
	Dec	✓	10.204	7.713	10.459	7.485
		✗	10.164	7.670	10.414	7.315
100	Enc	✓	20.055	18.090	20.763	17.747
		✗	19.636	16.715	19.402	16.085
	Dec	✓	12.141	9.509	12.266	9.240
		✗	12.049	9.689	12.491	9.284
1000	Enc	✓	119.864	115.551	119.596	115.086
		✗	104.953	101.514	106.809	99.056
	Dec	✓	46.000	43.972	47.385	42.321
		✗	46.109	43.981	47.173	42.594

TABLE VII. Average decryption time (in milliseconds) for different number of recipients and whether recipient validation is enabled. Visualized in Figure 10.

n	Validation	Cipher Suite			
		I	II	III	IV
5	✓	8.622	5.771	8.490	5.679
	✗	8.244	5.639	8.573	5.547
10	✓	8.676	6.172	8.771	6.011
	✗	8.334	5.809	8.549	5.513
20	✓	9.065	6.580	9.141	6.285
	✗	8.357	5.833	8.414	5.578
50	✓	10.204	7.713	10.459	7.485
	✗	8.435	5.805	8.602	5.633
100	✓	12.141	9.509	12.266	9.240
	✗	8.644	5.914	8.802	5.813
1000	✓	46.000	43.972	47.385	42.321
	✗	11.239	8.300	11.366	7.840