



Three Lessons From Threema: Analysis of a Secure Messenger

Kenneth G. Paterson, Matteo Scarlata, and Kien Tuong Truong, *ETH Zurich*

<https://www.usenix.org/conference/usenixsecurity23/presentation/paterson>

**This paper is included in the Proceedings of the
32nd USENIX Security Symposium.**

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

**Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.**

Three Lessons From Threema: Analysis of a Secure Messenger

Kenneth G. Paterson
*Applied Cryptography Group,
ETH Zurich*

Matteo Scarlata
*Applied Cryptography Group,
ETH Zurich*

Kien Tuong Truong
*Applied Cryptography Group,
ETH Zurich*

Abstract

We provide an extensive cryptographic analysis of Threema, a Swiss-based encrypted messaging application with more than 10 million users and 7000 corporate customers. We present seven different attacks against the protocol in three different threat models. We discuss impact and remediations for our attacks, which have all been responsibly disclosed to Threema and patched. Finally, we draw wider lessons for developers of secure protocols.

1 Introduction

Threema is a Swiss encrypted messaging application. It has more than 10 million users [32] and more than 7,000 on-premise customers [35]. It is among the top Android apps in Switzerland, Germany, Austria, Canada, and Australia in the “paid for” category [8]. It has undergone independent security audits [38, 37]. Along with Signal and Telegram, Threema was widely advertised as a more secure alternative to WhatsApp [56, 24, 57] in the wake of Facebook’s acquisition of the company [24] and of the January 2021 change in WhatsApp Terms of Service [31, 61]. Threema is the messaging app of choice of the Swiss Government and of the Swiss Army [25], where its use is mandated for all official communications [59]. Threema is also used by German politicians, including the current Chancellor, Olaf Scholz [52]. The company behind Threema makes strong claims for its security [34] and advertises its Swiss, non-US jurisdiction as a virtue.¹

In general, secure messaging protocols are designed to provide end-to-end encrypted communication between two (or more) parties. Modern designs like Signal achieve not only confidentiality, integrity, origin authentication and verification of correct ordering of the exchanged messages, but also:

- **session key security:** compromise of a session key or ephemeral values in one session should not affect the security of keys in other, parallel sessions;

¹Quoting from [32]: “Threema is 100% Swiss Made, hosts its own servers in Switzerland, and, unlike US services [...] it is fully GDPR-compliant. . .

- **fine-grained perfect forward secrecy (PFS):** compromise of session keys in an on-going session should not affect security of earlier keys and exchanged messages in that session and compromise of long-term keys should not undermine the security of completed sessions; and
- **post-compromise security (PCS):** after a compromise of all key material, it should be possible to regain security after a protocol execution, provided the adversary is passive during that execution.

Such properties have been established for Signal for two-party messaging [14, 41, 15, 21].

At the same time, the current landscape of messaging apps includes a wide variety of solutions, with different trade-offs with respect to centralization and federation, company jurisdiction, ease of use, openness of design and code, and security guarantees provided by the underlying cryptographic protocols. The absence of a standard for secure messaging means that many of the contenders feature bespoke cryptographic protocols with unclear security guarantees [7, 4, 6]. Indeed, these security guarantees are often much weaker than those provided by Signal, which can be seen as providing a “gold standard” for secure messaging.

1.1 Our Contributions

We provide the first in-depth security analysis of the cryptography used in Threema, focussing on its protocols for securing end-to-end (E2E) and client-to-server (C2S) communications, for user registration, and for backing up users’ private keys. We also examine the interactions between these protocols.

Our analysis uncovers seven different attacks against Threema, in three different threat models (as depicted in Figure 1). One of the attacks was discovered before us by Krebs, disclosed to Threema and patched in December 2021 [33]. We include it here for completeness. We produced proof-of-concept implementations for most of the attacks to verify that they work. We summarise the threat models and attacks next.

In the “external actor” threat model we consider attacks by outsiders with no special access to Threema servers or clients. Here, we found two attacks:

1. An adversary who is able to compromise a single ephemeral value used by a client in the C2S protocol is able to permanently impersonate that client to the server, and through this, access the metadata for all E2E-protected messages. This strongly violates the expected security properties of a well-designed protocol.
2. An adversary is able to register the server’s public key as a user public key and, by tricking the victim into sending a single carefully crafted message in the E2E protocol, can obtain a value that enables it to again permanently impersonate the victim to the server in the C2S protocol.

In the “compromised Threema” threat model we consider attacks by an adversary who has gained access to Threema servers. Such access should not enable breaking the end-to-end security guarantees. Here, we describe three attacks:

3. Due to peculiarities in nonce-handling in the E2E protocol, the adversary is able to arbitrarily and undetectably reorder and/or delete messages sent from any client to any other client.
4. Again, due to how nonces are handled in the E2E protocol, Threema clients must keep a local database of all nonces used in sending and receiving messages. This database is deleted when a user reinstalls the app or changes devices. After such an action, the adversary can replay old messages. Because of poor key separation, the adversary can also send back *to* user Alice any old messages originally sent *by* Alice.
5. For registration, a client proves possession of its private key by encrypting a server-chosen message to a server-chosen public key using the E2E encryption mechanism. This process is invisible to the user. This allows a malicious server to create “Kompromat”: a potentially incriminating message that can be delivered at any later time to a target user. This is the attack that we rediscovered.

In the “compelled access” threat model, we consider attacks by actors with direct access to a user’s device, for example law enforcement officers or border guards. We present two attacks in this setting:

6. A trivial attack exploiting a Threema feature allowing users to easily export a backup of their long-term private key protected under a chosen password. Because of Threema’s design, this key can be used to clone a Threema account and, through this, silently access all the user’s future messages. In conjunction with a compromised Threema server, all old messages can be accessed too. This attack is only possible when the app and mobile device are both unlocked.

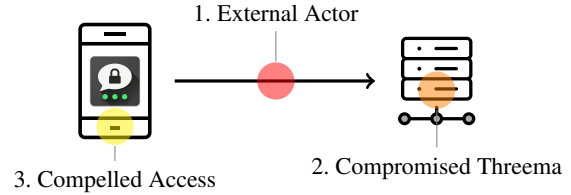


Figure 1: In the threat models we consider, the attacker can have access to network communication (1), the Threema servers (2) or the victim device itself (3).

7. A more sophisticated attack exploiting one of Threema’s backup mechanisms; this attack does not (by default) require an unlocked app (but does need an unlocked mobile device). It exploits Threema’s use of “compress-then-encrypt” when creating the backup, in combination with the ability of the adversary to inject chosen strings into the backup file through Threema nicknames. By observing the size of the backup file over many operations with chosen nicknames, the adversary can recover the user’s private key, with the same consequences as above.

Attacks 1 and 2 do not, by themselves, leak plaintext, but they target metadata that, under normal circumstances, only the server should be able to access. Furthermore, they break the PFS guarantees of Threema, allowing plaintext to be recovered if the long-term key of a user is ever exposed. Attacks 3, 4, and 5 show that the Threema server has the power to undetectably disrupt conversations, despite the E2E protocol (which, however, still guarantees message confidentiality). This requires users to put excessive trust in the server. Attacks 6 and 7 allow an attacker to obtain future plaintexts, as well as recover past plaintexts that are not protected by the C2S layer; however, they do require access to an end device. Our attacks are not directly caused by Threema’s lack of PFS and PCS, but they showcase how such a deficiency can amplify the impact of smaller attacks. The attacks range in practicality. For example, attack 6 is feasible for an unskilled attacker that has access to the device for a very short time, while attack 1 is more of theoretical interest. See Table 1 for a summary of the attacks’ prerequisites and impacts.

In totality, our attacks seriously undermine Threema’s security claims. In particular, we show that Threema’s approach to PFS (i.e. only providing it at the message transport layer) is insufficient even against a network attacker. All the attacks can be mitigated, but in some cases, a major redesign is needed. We discuss appropriate mitigations and the remediations undertaken by Threema, after presenting the attacks.

We also draw three wider lessons for developers of secure protocols: using secure libraries for cryptographic primitives does not on its own lead to a secure protocol design; beware of cross-protocol interactions; and adopt a proactive rather than reactive approach to secure protocol design.

1.2 Related Work

An independent early analysis of the Threema client-to-server protocol, dating back to the closed-source Threema version 1.3, was carried out by Ahrens [2], and was the basis for an open-source implementation by Berger [10]. We analyse the current version (version 4.831) of the Threema Android application, and base our analysis of the protocol on the 2020 open-source release of the Threema app [29] and the Threema security whitepaper [30].

Rösler et al. [54] considered the end-to-end security of group chats in various secure instant messaging services, including Threema. They observed the lack of forward secrecy and post-compromise security for Threema group chats, and described two attacks: a replay attack on group messages, and leakage of group membership information to non-members. Both attacks have since been patched by Threema.

Cremers et al. [18] investigated the security of Threema (and other apps) in a black-box post-compromise setting, by cloning a victim device and observing whether the clone retains access to exchanged messages. The authors note that Threema prevents the clone from accessing messages exchanged while the clone is offline. While this result holds in a black-box setting, the weaknesses we uncover in Threema’s client-to-server protocol enable even an adversary with limited technical proficiency to intercept messages if given access to the full cloned state of the victim.

Threema itself commissioned two security audits [38, 37]. Both audits focused on the Android and iOS apps for Threema, and do not touch on its cryptographic security beyond confirming that the implementation of the protocol matches the description provided in the Threema whitepaper [30].

Soatok [58] highlighted Threema’s lack of forward secrecy and transcript consistency in groups as deficiencies of the app. The lack of transcript consistency is not yet resolved by Threema, and, as [58] argues, can be combined with the lack of robustness of the XSalsa20-Poly1305 AEAD cipher to obtain “Invisible Salamander”-style attacks [19, 3]: an attacker can send attachments in group chats that decrypt to different plaintexts for different chat participants.

1.3 Ethical Considerations

We only used the description of Threema in [30] and the client-side source code to conduct our analysis. We did not attempt any reverse-engineering of the server behaviour other than what could be observed passively through normal client interaction. We performed experiments only with test accounts under our full control; no other user accounts were targeted. We reduced as much as possible the number of messages sent to Threema servers, and we avoided sending malformed messages to servers so as to avoid triggering DoS attacks.

We did purchase a “Threema Gateway” account (with Threema ID *LYTAAAS) and registered a public key under

that account identical to the Threema server public key (obviously, without knowing the corresponding private key). This step was necessary in order to validate one of our attacks (Section 3.1.2). It did not violate any terms or conditions of the service, so far as we could determine.

We disclosed the results in this paper to Threema on 03.10.2022, proposing a 90-day disclosure period. They acknowledged receipt also on 03.10.2022. We agreed on the details of the coordinated disclosure: a first batch of mitigations would be released in Q4 of 2022, while further mitigations would be released in Q1 of 2023, along with a post on Threema’s blog. The first batch of mitigations were included in Threema 5.0 for both Android and iOS² on 29.11.2022. To further remediate the issues we reported, Threema included their custom protocol, Ibex, in the same release. The new protocol aims to provide forward secrecy at the E2E layer.³ Threema also updated their whitepaper [30] on 15.12.2022 to add discussion of the forward security feature. We have not audited this new protocol. We agreed to publicly disclose our findings on 09.01.2023. The current publicly released patches address all of our attacks.

In one of our attacks, we leveraged a vulnerability in a library, Zip4j, used by Threema to create backup zip files. We disclosed our findings to the author of Zip4j on 08.10.2022, proposing a 60-day disclosure period. A CVE advisory for the bug (CVE-2023-22899) was published on 10.01.2023 [47]. A fix was released in version 2.11.3 on 26.01.2023.

1.4 Paper Structure

In Section 2 we introduce the Threema architecture and provide a description of Threema’s core cryptographic protocols and subsystems. Section 3 presents our attacks, grouped by threat model. We propose mitigations in Section 4. We discuss three wider lessons for secure protocol designers in Section 5.

2 Threema Architecture and Protocols

The design of Threema can be studied as the ensemble of several interoperable cryptographic protocols, all sharing the same user state and cryptographic primitives. The functional core of Threema consists of an end-to-end (E2E) protocol, which employs the long-term cryptographic secrets of the users to encrypt messages, and a client-to-server (C2S) protocol, a secure channel protocol that protects these messages in flight from clients to the Threema servers. Because of the composition of these two protocols, each message is wrapped in two layers of encryption: an inner end-to-end layer, and an outer client-to-server layer. We depict the composition in Fig. 2. Also relevant to our analysis and attacks are the registration protocol, data backup and contact discovery system.

²<https://threema.ch/en/whats-new>

³<https://threema.ch/en/blog/posts/ibex>

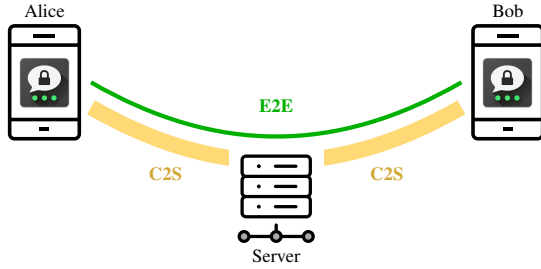


Figure 2: E2E and C2S protocol composition. Clients establish a secure channel with the server using the C2S protocol (in yellow) to send and receive E2EE messages from other users, which are relayed via the server (in green).

We first describe the cryptographic primitives used by Threema and establish our notation. Then we analyze the registration protocol, the E2E protocol, and the C2S protocol. Finally, we briefly cover the backup mechanisms. For completeness, we also include a description of the contact discovery mechanism in Appendix A.

2.1 Cryptographic Primitives and Notation

The protocols used by Threema rely on standard cryptographic primitives. The main tool used is the `crypto_box` abstraction provided by the Networking and Cryptography Library (NaCl) [11] which consists of a Curve25519 Diffie-Hellman key agreement, followed by encryption with an Authenticated Encryption with Additional Data (AEAD) algorithm. The AEAD algorithm used by NaCl is the nonce-based XSalsa20-Poly1305. As a convention, long-term keys will be indicated with the same letter as the symbol for the user, with the public key in uppercase and the private key in lowercase (e.g. user \mathcal{A} has private key a and public key A). We reserve (x, X) , (y, Y) , and (z, Z) for values used as ephemeral keys.

Let $X25519(\cdot)$ denote the NaCl function that takes a private key and a public key, and outputs a shared secret byte string. In NaCl, this string results from applying the `hsalsa20` key derivation function to the byte representation of the shared elliptic curve point and a fixed nonce. Internally, the `X25519` algorithm operates on the Curve25519 elliptic curve: the algorithm uses 32-byte scalars as private keys and computes the scalar multiplication of the chosen base point of the curve g with the private key to obtain a 32-byte public key. We will use multiplicative notation for such operations: e.g. secret key a has a corresponding public key $A = g^a$ for some fixed generator g . We denote by $(x, X) \leftarrow \text{KeyGen}()$ the key generation procedure of Curve25519, which outputs a private key x and a public key $X = g^x$. Let K be an AEAD key; we denote by $E_K(m; n)$ the encryption of message m under key K using nonce n , and by $D_K(m; n)$ the corresponding decryption operation. If the decryption fails, we assume that the algorithm returns a special value \perp . This may happen, for example, if the ciphertext has been tampered with, or if the wrong key

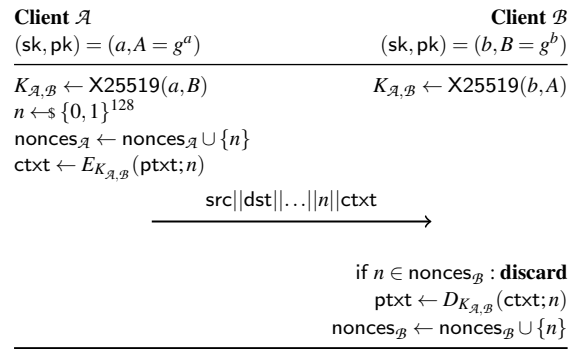


Figure 3: User \mathcal{A} sending ptxt to user \mathcal{B} with the E2E protocol. The set $\text{nonces}_{\mathcal{U}}$ represents the nonce database of user \mathcal{U} .

was used to decrypt.

We denote by $k \leftarrow \text{KDF}(K, \sigma, \tau)$ a Key Derivation Function taking some key material K , a salt σ , and a label τ as input, and producing key k as output. Threema uses a KDF based on the Blake2b hash function [55].

2.2 Registration Protocol

In order to create a new account and register it with the Threema server, each user generates a Curve25519 key pair (a, A) . The entropy for the generation of the key pair is taken from system randomness, as well as random user input which is processed and mixed with the former to obtain the 32-byte private key a . The client later runs a registration protocol with the server to prove knowledge of a . This is done via a challenge-response exchange run over a TLS-protected connection: the client sends A to the server. The server replies with an ephemeral public key X and a message m . The client computes $K = X25519(a, X)$ and AEAD-encrypts m using K and a fixed nonce (“createIdentity response.”) to create a response. If this response decrypts correctly and yields m as plaintext, the server accepts the registration and then issues a new Threema ID, storing it along with the public key of the user in the Threema database.

2.3 Threema End-to-End Protocol

The Threema E2E protocol, depicted in Figure 3, is concerned with guaranteeing end-to-end security of messages. Every user \mathcal{U}_i of the protocol has an alphanumeric 8-character identity $ID_{\mathcal{U}_i}$ (the *Threema ID*, or simply *ID*), a private key u_i , and a corresponding public key $U_i = g^{u_i}$. Messages sent from user \mathcal{U}_i to user \mathcal{U}_j are encrypted under $K_{\mathcal{U}_i, \mathcal{U}_j} = X25519(u_i, U_j) = X25519(u_j, U_i)$, a static Diffie-Hellman key. Note that the same key is derived by user \mathcal{U}_i for \mathcal{U}_j as the one from user \mathcal{U}_j for \mathcal{U}_i : the key $K_{\mathcal{U}_i, \mathcal{U}_j}$ is bidirectional.

A message m is first serialized into a byte string, and prefixed with a byte that indicates the message type. The specific serialization method depends on the message itself. For in-

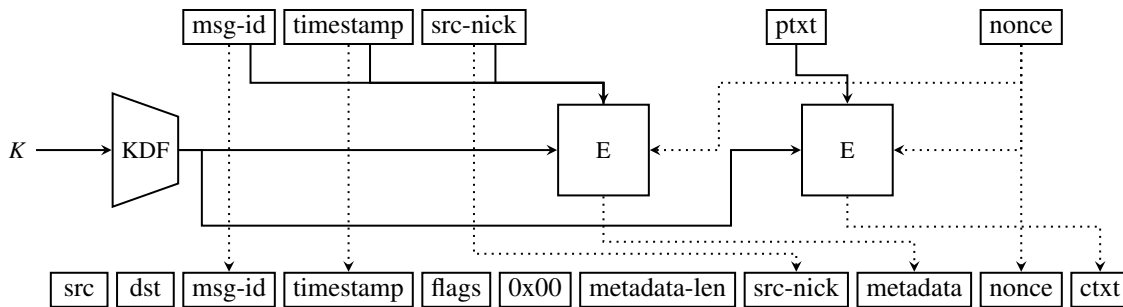


Figure 4: Structure of an E2E encrypted message. The input to the KDF is the key material K . The salt and label are implicitly input into the KDF, taking values “mm” and “3ma-csp”, respectively.

stance, a *text message* would be directly encoded as a byte string representing the UTF-8 string, and prefixed with $0x01$. A random amount of PKCS7 padding, between 1 and 254 bytes, is appended to the message. The resulting plaintext `ptxt` is then encrypted under K using the AEAD algorithm with a random nonce, `nonce`, to obtain a ciphertext `ctxt`.

To create the final message packet, the following information is prepended to the ciphertext: the Threema ID of the sender `src`, and of the receiver `dst`, a random 8-byte message ID `msg-id`, the timestamp at which the message was sent `timestamp`, the nickname of the sender `src-nick`, the nonce `nonce` and an *optional* metadata value. To construct this metadata, the app takes the message ID, the nickname of the sender and the timestamp, concatenates them, and encrypts the resulting plaintext under $KDF(K, \text{“mm”}, \text{“3ma-csp”})$, using the same nonce value. Since the presence of the metadata box is optional and it is not cryptographically bound to the ciphertext, it can always be stripped off and its length set to 0 by an adversary. In Fig. 4 we depict the structure of an E2E-encrypted packet in Threema.

When sending an E2E message, the associated random nonce is stored in the device’s local storage in order to avoid reusing it. If a generated nonce matches a nonce in the local storage, up to five attempts are made to sample a new nonce. When receiving a message, the app checks if the corresponding nonce has already been seen: if so, it rejects the message. If the nonce is otherwise fresh, the app tries to decrypt the message, verifying the MAC tag as part of the process. If successful, the app stores (a hashed representation of) the nonce. This storage of nonces is aimed at preventing trivial replay and reflection attacks (which would otherwise be possible since the key K_{u_i, u_j} is bidirectional) and hedges against the possibility of a faulty randomness source.

Remarkably, the E2E protocol does not provide any forward secrecy, since it only uses the long-term keys of the users, with no ephemeral values being exchanged. This means that a malicious server could take advantage of its position to execute a “store now, decrypt later” attack. This has supposedly been fixed with the introduction of the Ibex protocol.

2.4 Threema Client-to-Server Protocol

The Threema C2S protocol is concerned with establishing a secure channel between a client and the server to protect messages in flight, shielding metadata and the non-forward secret E2E protocol messages from a network adversary. Note that this protocol is only used for messaging: other functionalities such as registration, backup, and contact matching are implemented over TLS.

Threema’s C2S protocol is a novel secure channel protocol, broadly comparable in function and construction to the TLS protocol. It assumes a reliable transport, and can be decomposed as a *handshake subprotocol*, in which the server and the client interact to establish a session key, and a *transport subprotocol*, where the server and the client use the established session key to exchange messages.

The C2S protocol requires that each user connects to the server from only one device at any given moment. If two devices try to connect to the server at the same time, the older connection is dropped in favour of the newer one. Before terminating the connection, an error message is sent to the device, informing it that another device has connected with the same Threema ID.

2.4.1 Handshake Subprotocol

The handshake subprotocol runs between the server \mathcal{S} (with long-term key pair $(s, S = g^s)$) and a client \mathcal{U} (with long-term key pair $(u, U = g^u)$). Fig. 5 gives a pictorial representation of the protocol. We refer to the figure for the naming of the messages below. We can assume that the server and the client already know each other’s public keys: the client knows the server’s key as it is pinned in the app and the server receives the client’s key at registration time.

The client first generates an ephemeral key pair $(x, X = g^x)$ and a 16-byte value c called the *client cookie*, and sends (X, c) to the server (`client-hello`). After receiving the client’s ephemeral public key and the client cookie, the server generates its own ephemeral key pair $(y, Y = g^y)$ and a 16-byte value called the *server cookie*. It then computes the first shared symmetric key of the exchange $K_1 = X_{2519}(s, X)$, mixing

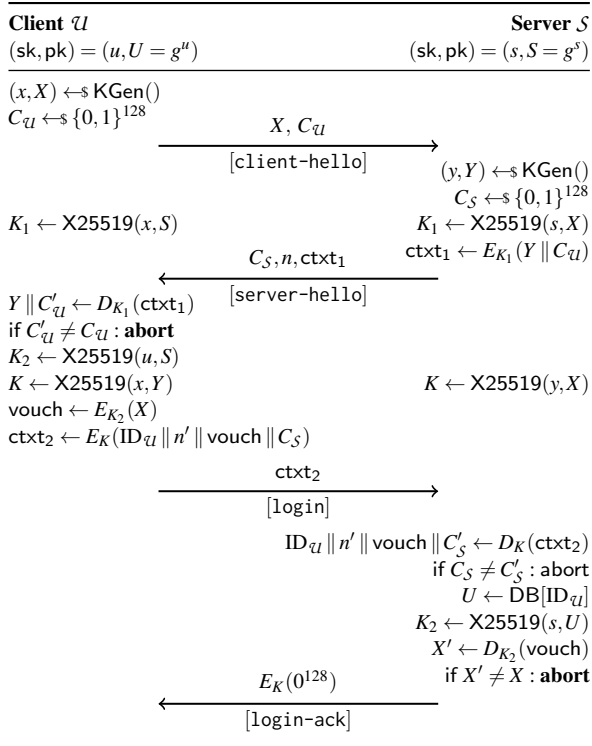


Figure 5: The Threema C2S Protocol (Handshake Subprotocol). For the sake of brevity, nonces are omitted from this diagram. The table DB is an abstraction for the database that maps Threema IDs to public keys. Message names are taken from the Android code and depicted in square brackets.

the long-term key of the server with the ephemeral key of the client, and uses it to encrypt its ephemeral public key concatenated with the client cookie, using a random nonce n . The server sends the resulting ciphertext, along with its own cookie (server-hello).

The client and server cookies prevent replay attacks and are used to initialize client and server *counters*. For each counter, the most significant 16 bytes of the counter are set to the respective cookie, while the lower 8 bytes function as a monotonically increasing value. Subsequent AEAD nonces for the exchanged messages are drawn from these counters.

When the client receives the message, it recomputes K_1 and checks if the ciphertext decrypts correctly and if the client cookie in the plaintext corresponds to the one that the client previously generated. If either check fails, it will abort the protocol. The client and the server derive two additional keys: K , obtained by combining the two ephemeral key pairs (x, X) and (y, Y) , and K_2 , obtained by combining the two long-term key pairs (u, U) and (s, S) . The key K is used in the transport subprotocol as the *session key*. The long-term symmetric key K_2 is used by the client to create the so-called *vouch box* $\text{vouch} \leftarrow E_{K_2}(X; n')$ using a random nonce n' . The *vouch box* is used by the server to authenticate the client, and to bind the

client's ephemeral key X to the client's identity.

The *vouch box* is concatenated with $\text{ID}_{\mathcal{U}}$ and the server cookie $C_{\mathcal{S}}$, and encrypted using the session key K (login). The server decrypts the message with K and tries to decrypt the *vouch box* with K_2 . If the decryption fails or the value contained in the *vouch box* does not correspond to the client's ephemeral key, the server will abort the protocol. Otherwise, the server sends a final confirmation message composed of 16 zero bytes, encrypted with the session key K (login-ack).

Taken as a standalone protocol, the C2S handshake could be formally analysed using, for example, the extended Canetti-Krawczyk security model [42] or the MSKE security model of [26] that was used in [14, 15] to analyse Signal. We will (informally) use these models when relating the attacks in Section 3 to the formal security properties they violate.

2.4.2 Transport Subprotocol

In the transport subprotocol, the client and the server exchange messages, encrypted using the session key K . The client and server nonces are derived from the same (respective) counters inherited from the handshake subprotocol, ensuring that handshake messages cannot be used as transport messages (and *vice versa*). If a decryption fails, the entire connection is dropped. This prevents message reordering and message deletion by a network adversary, since either attack would result in a party decrypting with a nonce different from the correct one, and then the AEAD integrity guarantees would ensure that an error is generated on decryption.

When the connection is dropped, the client attempts to reconnect, with up to five reconnection attempts being made. If these all fail, the app will display an error to the user and will require a restart to be able to connect to the server. If two devices are actively trying to connect to the server, they will keep interfering with each other until one of them relinquishes the connection. This is relevant for any attacker that wants to attempt an impersonation attack, since using all five attempts will alert the user that another device is trying to connect to the server with the same Threema ID.

While the connection is alive, the server will send messages that are meant for the user encrypted under the session key K . The client replies with an ACK in response to each message. If the client does not send an ACK, the server will try to send the message again when the client later reconnects.

2.5 Threema Safe

Threema allows a user to activate, at any point in time, the Threema Safe feature. This allows the user to backup their long-term private key, their nickname, the list of their contacts and additional account information on the Threema servers. Note that the backup does not store any message history.

When the Threema Safe is first set up, the user chooses a password for their backups. A minimum length of 8 characters

Table 1: Summary of the attacks, listing the prerequisites and their impacts. The last column focuses on the ability of an attacker to obtain plaintext after executing the attack.

	Network Attacker	C2S Eph. Key Leak	User Action	Unlocked Phone	No UI Lock	No Passphrase	Server Compromise	Leak in-metadata	Leak out-metadata	Break of PFS	C2S MiM	Permanent C2S Impersonation	Conversation Integrity Break	Long-term key leak	Recovers Plaintext
C2S Key Compromise	✓	✓	✗	✗	✗	✗	✗	✓	✓ ^a	✓	✓ ^a	✓	✓	✗	✗
Vouch Box Forgery	✗	✗	✓	✗	✗	✗	✗	✓	✗	✓	✓	✓	✓	✗	✗
Reordering/Deletion	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✓	✗	–
Replay/Reflection	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✓	✗	–
Kompromat	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗	✓	✓	✗	–
Threema ID Export	✗	✗	✗	✓	✓	✓	✗	✗	✗	✗	✗	✓	✓	✗	–
Backup Side-Channel	✓	✗	✗	✓	✗	✓	✗	✓	✗	✓	✓	✓	✓	✓	✓

^a As long as the client keeps using the revealed ephemeral key (i.e. up to one week)

is enforced, and the password is checked against an offline list of weak passwords. The app derives two 32-byte strings, the Backup ID BID and the Backup Key BK, using `scrypt` with the password as input and the Threema ID of the user as salt. For every backup, the app creates a JSON payload containing all the data to be backed up and serializes it into a byte string, which is then compressed using `gzip`. Finally, the app encrypts this compressed JSON string with `XSalsa20-Poly1305` under BK with a random nonce. The resulting ciphertext is sent to the server over a TLS connection.

The server will store ciphertexts indexed by BID, supposedly without any other information related to the user. In the Threema whitepaper [30], the developers claim that this process is designed to make it impossible for the Threema Safe server to distinguish which backup comes from whom by looking at the uploaded data. We note, however, that when a backup is created the app is likely to have an active C2S protocol session, which reveals the Threema ID of the user. By linking the two connections coming from the same IP, the server can learn the identity associated with a given backup.

A backup is scheduled to happen every day or, in case the previous backup failed to be uploaded, as soon as the app is restarted. This happens even if access to the app is protected by a PIN or biometric login.

2.6 Other Backup Methods

Threema provides two additional methods to save personal data: the first one is called a *Threema ID export* and the second one is called a *data backup*.

The Threema ID export requires that the user provide a password, which is used to derive a 32-byte symmetric key by using the `PBKDF2` algorithm with a random 8-byte salt. This key is used to encrypt the ID of the user, as well as their long-term private key using `XSalsa20` with a zero nonce (but with no integrity protection). The result is base32-encoded and given to the user in order to be saved outside of the app.

A data backup consists of an encrypted zip containing

various files. The identity file stores the Threema ID and the long-term key of the user; the contact file stores Threema IDs, nicknames, and public keys of contacts in CSV form, and for each chat, a file named with the Threema ID of the correspondent stores the chat history. The zip is encrypted following the WinZip AE-2 specification [16].

We note that, because WinZip does not hide the file names, this encrypted zip leaks information about with whom the user has been communicating. Due to the length of each file being known, an attacker may also infer the length of the conversation from the encrypted zip.

3 Attacks on Threema

We consider three threat models: first, an external actor who can monitor communications between a device running the app and the Threema servers; second, an actor that has compromised the Threema servers; third, an actor that can take control of the device for a short period of time.

The first of these models corresponds to the standard adversary for secure network protocols, often called the Dolev-Yao adversary. We refer to it as the *external actor* threat model.

Concerning the second model, any reasonable security analysis must consider the possibility that the messaging server is malicious: national security agencies and hacking groups are tasked with penetrating servers like those run by Threema in order to gain persistent access to concentrated amounts of sensitive data. It is then reasonable to ask what security guarantees remain for users in the event of such a breach. This is especially so for a system that claims to provide end-to-end security. We refer to this model as the *compromised Threema* threat model.

We call the third model the *compelled access* threat model. It is relevant, for example, in the case of border searches of electronic devices, when protesters are detained by police forces and searched for incriminating evidence, or in the setting of intimate partner violence [27]. We may subdivide this setting into two sub-cases: one where the app is locked, and

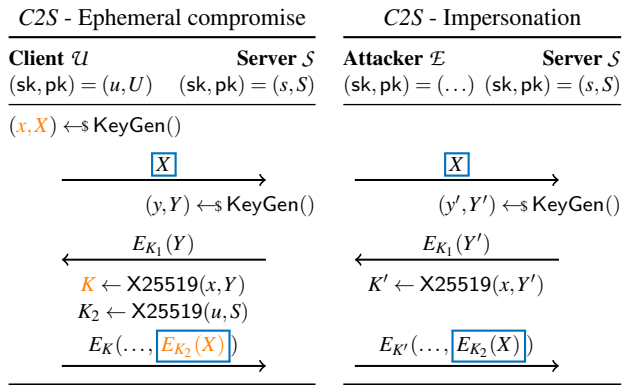


Figure 6: Attack 1: two C2S sessions. The attacker compromises an honest session by revealing the ephemeral private key x , from which they can compute the session key K and obtain the vouch box $E_{K_2}(X)$ (in orange, Left side). The values are replayed by the attacker in a new session (blue boxes, Right side). Note the lack of *session independence*.

the other where it is unlocked (we will always assume the device itself is unlocked).

We next introduce our attacks, grouped by threat model. Table 1 summarizes our findings.

3.1 External Actor Threat Model

Our first two attacks allow the attacker to impersonate a victim in the C2S protocol. The core of both attacks lies in the structure of the vouch box vouch. This is used in the protocol to authenticate clients but contains only client-chosen freshness. By either replaying (Attack 1) or forging (Attack 2) a user’s vouch box, the attacker can authenticate itself as that user to the server. These impersonation attacks are *permanent*, despite the fact that the attacker never learns the value of the victim’s long-term key.

Our attacks show that the weak notion of forward secrecy achieved by Threema is insufficient: by using either of our attacks, the adversary gains access to E2E-encrypted messages, which use long-term keys only. The reuse of ephemeral keys which we describe in the first attack further weakens the guarantees given by forward secrecy, since an attacker gains access to messages up to a week in the past. Comparing this design to Signal’s, we note that in the latter each E2E message is encrypted with a different key, providing a much stronger, fine-grained notion of forward secrecy.

3.1.1 Attack 1 (C2S Ephemeral Key Compromise)

Assumptions: We assume that the attacker has discovered one of the ephemeral keys x used by the client in a C2S handshake session. This capability is commonly referred to in formal security models as *ephemeral key reveal*, and it can arise in practice from a randomness failure or via a

side-channel attack. We further assume that the attacker has passively recorded the corresponding C2S session.

Execution: We now show that, with the knowledge of x and the C2S session transcript, the attacker can create new sessions with the server in which they impersonate the victim and can tamper with the flow of messages.

As highlighted in Figure 6, the attacker has all the material necessary to derive K_1 and decrypt the login message from the recorded C2S protocol run, from which it can recover the value of vouch. Now the attacker initiates a new run of the C2S protocol, playing the role of the victim user. In this run it replays the ephemeral public value X . A new session key K' known to the attacker is generated (the server should send a fresh Y' in its server-hello but the attacker still knows x). The attacker can then reuse vouch (without learning K_2) when constructing its response in the login message, which makes the handshake complete successfully. We have verified that the server does not detect such replays.

Impact: Since vouch can be reused repeatedly, this attack (involving the loss of just a single ephemeral value) allows the attacker to permanently impersonate the victim user to the server, enabling the attacker access to all E2E-encrypted messages meant for the victim. Despite the E2E layer of encryption, this still reveals the metadata information in all communications sent to the victim, allowing the attacker to learn with whom and when the victim has been communicating.

Most importantly, this attack leads to a break of the forward secrecy property that the C2S protocol is intended to provide. Indeed, the E2E-encrypted messages that the attacker now sees are encrypted only using long-term keys. Full decryption of those messages becomes possible if the victim’s long-term key is later revealed (for example through attacks 6 or 7).

In addition, the attacker gains the opportunity to selectively drop messages from the conversation by deciding which messages to ACK to the server. If a message is ACK-ed by the attacker, it will not be received by the victim when they next connect to the server, virtually deleting that message from the conversation on the victim’s end.

We stress that, as long as the attacker and the victim do not connect to the server at the same time, this attack is not detectable by the victim. If they were to connect at the same time, the attacker could observe error messages from the server to learn this. At this point, the attacker can relinquish the connection in order to avoid detection.

This attack shows that, for the C2S protocol, ephemeral keys are as important to protect as long-term keys. This is a very strong assumption to have to make in order to guarantee security for a key exchange protocol.

We note that a similar vulnerability was present in the Off-The-Record (OTR) protocol [51]. The authors of [51] highlight the consequent lack of the *session independence* property in OTR: “the exposure of ephemeral session-specific secrets should have no bearing on the security of other ses-

sions”. Threema’s C2S protocol suffers from the same issue.

Non-ephemeral Ephemerals: An additional problem with the C2S protocol lies in how the ephemeral keys are handled client-side: if the app is never restarted, then the same ephemeral key is used by the client for up to seven days before generating a new one. Leaking such a re-used ephemeral key then allows an attacker to also impersonate the server to the user by forging a server-hello using the server’s public key and the compromised ephemeral key (this is known as an *ephemeral Key Compromise Impersonation* attack in the literature). In turn, this implies that the attacker has the power to also read the outbound E2E-encrypted messages of the victim that are meant for the real Threema server. This reveals additional metadata, allows the attacker to drop arbitrary messages and allows a network adversary to execute attacks that would otherwise be in the “compromised Threema” threat model, such as the ones described in Sections 3.2.1 and 3.2.2. If ephemeral keys were to be freshly generated for each handshake, this problem would be confined to the single compromised session. However, by reusing the same ephemeral key for a long time, the effect of this attack is heavily amplified.

Finally, there is similar reuse of ephemeral keys on the server side. We noticed that whenever the same ephemeral key X is used multiple times by a user within a short span of time, the server will use the same ephemeral key Y as well. An attacker is thus able to force the same session key to be used for a long time by regularly connecting to the server with the same value X . We note that compromising said server ephemerals would allow the impersonation of multiple users to Threema for an indefinite period of time, as long as the attacker has recorded C2S handshakes involving the revealed ephemeral keys. In fact, it suffices for the attacker to replay the client-hello, inducing the server to reuse a previously revealed key y . As the attacker can compute the session key, they can create a new login message, which contains an old vouch box, thus completing the handshake.

In order to test all these weaknesses, we created a Python script which, given a valid transcript of a C2S handshake and the client ephemeral key used, can impersonate the user to the server. The script shows all the metadata of incoming messages and can drop arbitrary messages from the conversation.

3.1.2 Attack 2 (Vouch Box Forgery)

An alternative route to achieve the same goal as the previous attack is to forge a new vouch box containing an ephemeral key X for which the adversary knows the private value x . This is feasible due to a cross-protocol interaction between the E2E and the C2S protocols. The consequence is that the attacker again gains the ability to permanently impersonate the victim.

Assumptions: The attack requires two steps. In the first step, the attacker \mathcal{E} has the victim use the public key S of the C2S protocol server when communicating with \mathcal{E} in the E2E protocol. In the second step, the attacker, now acting as a

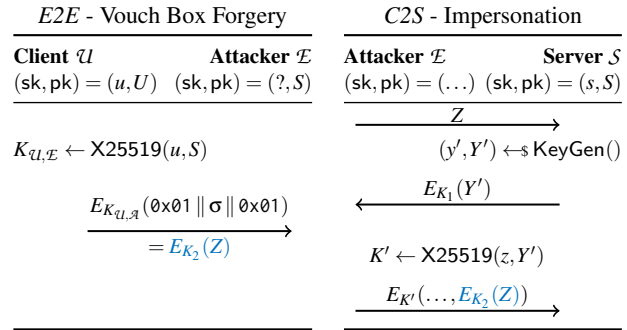


Figure 7: Attack 2: a cross-protocol interaction of an E2E and a C2S session. The attacker claims the public key of the server, and knows a keypair of the form $(z, Z = 0x01 \parallel \sigma \parallel 0x01)$. They convinced the victim \mathcal{U} to send σ to them as a E2E text message (in blue, Left side). The attacker can now start a session of the C2S protocol (Right side) where they use the “ephemeral” keypair (z, Z) and the corresponding vouch box $E_{K_2}(Z)$ (in blue) in order to authenticate as \mathcal{U} to the server.

Threema user, has the victim user send a specifically crafted text message to the attacker in the E2E protocol. Under normal circumstances, these actions should not lead to an attack: if the victim sends a message encrypted against the server’s public key, then the attacker should not be able to decrypt it, nor use it for other purposes. Furthermore, sending an E2E message should not compromise security of the C2S protocol.

We will show two different methods to implement the first step of the attack. In the first method, we use part of the infrastructure of Threema to create a new account which has the server’s public key S as its own public key. In the second method, we leverage a vulnerability in a dependent library used by the Android version of Threema.

Execution: The attacker’s objective is to forge a vouch box vouch for this key X that would allow the attacker to authenticate as the victim in the C2S protocol. Let (u, U) be the long-term key pair of the victim. Assume that the first step of the attack has already been executed, meaning that, whenever the victim wants to send a message to the attacker, they will use their private key u and the server’s public key S to encrypt their message. Assume also that the attacker manages to find a private key x such that the corresponding public key $X = g^x$ has the following form: a $0x01$ byte, followed by a string σ composed of 30 printable UTF-8 characters, followed by another $0x01$ byte. If the victim sends σ as a text message to the attacker in the E2E protocol (for example, via social engineering, e.g. sending a special code to win a prize), this would result in the victim deriving a key $K^* = \text{X25519}(u, S)$ and encrypting the message. Since σ is sent as a text message, the ciphertext will include $0x01$ as the first byte, followed by the string σ and then PKCS7 padding at the end. The attacker hopes to obtain a ciphertext $c = E_{K^*}(0x01 \parallel \sigma \parallel 0x01; n) = E_{K^*}(X; n)$ (for some nonce n). This has probability $1/254$ of working, since this is the probability of obtaining $0x01$ as a



Figure 8: Attack 2 in practice: on the left, a suitable keypair (base64 encoded). The public key bytes 1 to 31, also encoded in the QR code, all consist of printable UTF-8 characters. On the right, the *LYTAAAS Threema gateway account (since revoked), with the hijacked public key of the server. User \mathcal{U} sending the contents of the QR to *LYTAAAS as a message will allow *LYTAAAS to authenticate to Threema as \mathcal{U} .

PKCS7 padding. We note that K^* here is actually equal to K_2 , the key that would be derived by the victim during the C2S protocol and used in creating a vouch box. Because of this, by construction, the ciphertext c that was created and sent to the attacker is thus a valid vouch box vouch for X . Since the attacker knows the corresponding x , the attacker can now use (x, X) and vouch to authenticate to the Threema server as the victim. The attack is depicted in Figure 7.

Finding a good X : Creating a suitable key X is feasible, simply by randomly sampling x and computing $X = g^x$ until X has the desired form. We have empirically estimated the probability of each trial being successful in finding a good X as being $\approx 2^{-51}$. Thus, after 2^{51} trials, the success probability would be 0.63. We implemented this approach, using some batch inversion optimisations to reduce the cost of coordinate conversions. We found a key of the correct form after $2^{50.0}$ trials. In total, this required approximately 8100 core-days across a variety of computational clusters available to us. The results are shown in Figure 8. We used this key in our experimental environment to demonstrate that the attack works in practice. We note that this is a one-time cost: once a suitable key is found, it can be reused to target any user.

Claiming the server’s public key: We now describe two methods allowing an attacker to claim S as their long-term public key, the first step in the attack.

Our first method exploits the fact that Threema provides a paid API, the “Threema Gateway”, for companies to interface with the Threema messaging system automatically. When registering a new account, the user is required to provide a public key. The gateway, however, does not check (1) if the user has the corresponding private key and (2) that the key

does not correspond to the server’s public key S .⁴ This allows the creation of an account which has the public key of the server as its own. We created such a Threema account under the Threema ID “*LYTAAAS”.⁵

Our second method uses a vulnerability in one of the dependencies that Threema has on the Android version of the app for creating data backups (described in Section 2.6). The library used by Threema to create zip files is called Zip4j [43]. This library possesses a bug where the MAC is not checked when decrypting the zip file, if certain conditions are met. This is the case whenever a Threema backup is restored, meaning that any tampering with the zip would not be detected. In the context of Threema, this allows an attacker that has write access to the encrypted zip to modify the contacts file within it, allowing it to overwrite their public key with the server’s. For example, a user might choose to save their zip file in shared storage (e.g. a cloud service, or a folder on a computer which the attacker can access), believing it to be both encrypted and integrity-protected. While this second method does not exploit a vulnerability in Threema *per se*, a more robust design would have prevented this bug from being escalated to client impersonation in the C2S protocol.

Impact: This attack yields the same effects as the previous one. The attacker gains access to messages (still E2E-encrypted), along with their metadata. This puts the attacker in a position where they can store messages meant for the victim and decrypt them at a later point, when the long-term key of the victim is compromised. The attacker, again, has the ability to drop arbitrary messages from the conversation.

3.2 Compromised Threema Threat Model

The following attacks assume adversarial server infrastructure. Attack 5 requires access to the server’s secret cryptographic material. Attacks 3 and 4 only require being able to “strip away” the C2S security layer. This is trivial for the Threema server, but could also be achieved by exploiting C2S protocol weaknesses, like the one we present in Section 3.1.1.

On forward secrecy: We again stress that, from the point of view of an attacker located at the Threema server, the E2E-encrypted communications have no forward secrecy whatsoever, since it is only provided at the client-to-server level rather than the end-to-end level. This leaves the user completely vulnerable in the event of a long-term key compromise in the “compromised Threema” threat model. This is explicitly acknowledged in the Threema whitepaper [30].

In a messaging app, compromising the server should not give away significant data to the attacker: the server should

⁴Note that performing this check properly is non-trivial: we also successfully registered keys $T \neq S$ such that $T = S \cdot P$, where P is a Curve25519 point of order 8. When computing a Diffie-Hellman share using the private key a and the public key T , due to clamping $T^a = P^a \cdot S^a = S^a$, thus claiming T is equivalent for our attack to claiming S .

⁵“Lose Your Threema Account As A Service”

merely act as a message router, which cannot read messages and has no gain in storing the messages that it sees, or tampering with them. Due to the lack of forward secrecy on the E2E level, this expectation is not met by Threema.

3.2.1 Attack 3 (Message Reordering and Deletion)

The lack of an ordering mechanism in the E2E protocol combined with the lack of authentication of the metadata in that protocol allows the server to carry out trivial reordering attacks on E2E messages. This means that the server can change the order in which the messages are delivered to the receiver, thereby potentially changing the semantics of the conversation. Furthermore, the server can withhold messages for as long as it pleases and send them at a later time with an updated timestamp, in order to evade detection.

In addition to the above, we noticed that incoming messages are shown in the order in which they are received by the Threema app, rather than being ordered by timestamp.

3.2.2 Attack 4 (Message Replay and Reflection)

In Section 2.3 we explained how Threema tries to prevent replay and reflection attacks on the E2E protocol by storing nonces of both outgoing and incoming messages. Fundamentally, this requires that the nonce database is always kept updated and is never deleted in the client app. Unfortunately, this cannot be guaranteed whenever the app is reinstalled or when the user changes device: in either case⁶, the nonce database is reset and a compromised Threema server can replay messages that were received by the victim in the past or reflect messages that the victim has sent.

To make this attack more effective, we note that the Threema server can easily tell when a victim is likely to have reset their nonce database, as long as the victim has decided to use Threema Safe as their backup option. (A reasonable assumption, since the user is strongly encouraged at registration time to use it.) Whenever a user restores a backup from Threema Safe, it will connect to the server with the C2S protocol shortly after, most likely with the same IP address. Thus, not only can the server learn the identity associated with a given backup, but it will also know that, with high probability, the user's nonce database has been deleted, making that user a target for the aforementioned attacks.

To test Attacks 3 and 4, we created a simulated Threema server in Python to which clients can connect. The server provides control over the order of the messages and their inclusion in the conversation, and allows the attacker to save the messages in order to be later replayed or reflected. We experimentally verified all the behaviours observed above.

⁶Note that on iOS, migrating the app to a new device using *Quick Start* would preserve the database.

3.2.3 Attack 5 (Kompromat)

This attack provides another example of a dangerous cross-protocol interaction, in this case between the registration protocol and the E2E protocol.

Execution: Recall from Section 2.2 that, when a user \mathcal{A} registers to Threema, they must prove ownership of its public key to the server by use of a challenge-response protocol: the user combines their private key a with a public key provided by the server $X = g^x$ and encrypts a message m chosen by the server. When the contact discovery protocol is run (see Appendix A), the same challenge-response protocol is executed to prevent an excessive number of queries. The encryption method used is identical to that in the E2E protocol. In this protocol, the server does not necessarily have to own the private key x corresponding to the provided ephemeral public key, and they can present the long-term public key of any other Threema user \mathcal{B} instead. Thus, \mathcal{A} will be encrypting the challenge with the same key that they would use to communicate with \mathcal{B} in the E2E protocol. Hence, the resulting ciphertext is a valid E2E message that can be sent to either \mathcal{A} or \mathcal{B} (and which can be made to appear to come from the other party).

Impact: This attack targets the integrity of the E2E protocol in Threema. In the unpatched protocol, every time the contact discovery protocol is run (usually once per day) the server is able to forge a valid ciphertext. Each ciphertext can only be sent once per user due to the replay reflection provided by the nonces. This flaw was already fixed when we re-discovered it. The app now requires that the message m provided by the server has a type byte equal to `0xff`, which does not correspond to any valid message type, thus preventing it from being accepted as a valid message. While this does effectively prevent the attack, it does not tackle the other cross-protocol attacks that we present. In Section 4 we discuss mitigations that offer full protection.

3.3 Compelled Access Threat Model

The following attacks are in the strongest threat model of compelled access, which grants the attacker access to an unlocked device. These attacks allow the recovery of the long-term key of the victim. Due to the lack of forward secrecy and post-compromise security, this irreversibly forfeits all security. For instance, considering a combined attack setting of compelled access *and* a compromised Threema server, we can obtain an attack that can decrypt all past and future communications.

3.3.1 Attack 6 (Cloning via Threema ID Export)

Assumptions: Access to an unlocked phone, with no Threema App UI lock.

Execution: Among the backup methods provided, Threema allows the user to export the Threema ID (as explained in

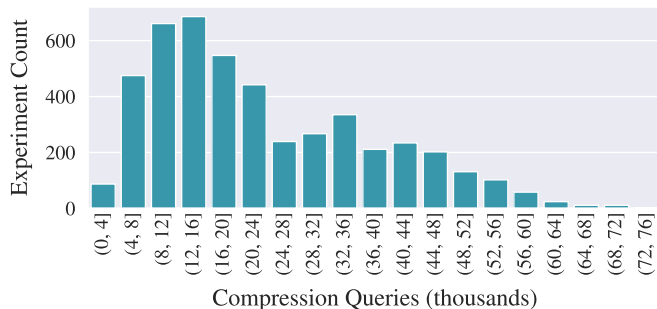


Figure 9: The distribution of the number of compression side-channel oracle queries in our successful key-recovery experiments. Each “query” is a backup attempt by the Threema app.

Section 2.6). However, this method directly exposes the long-term secret key of the user, and does not by default require any form of authentication, allowing a compelled access attacker to backup the private key with a password of their choice and later decrypt it on their own device.

This is an intentional feature in Threema, but one which we deem to be a grave security danger: all security in Threema relies on maintaining the secrecy of long-term private keys, and this attack allows an adversary to recover such keys by having access to the phone and app for a few seconds.

Impact: While similar mechanisms are present in other messengers, their security impacts are much smaller. For instance, Signal allows users to transfer their accounts to a new device, but it requires physical proximity of the new and the old devices, and disables the old device permanently after such a transfer is completed. This is in contrast to Threema, where a Threema ID export is undetectable on the victim device after it happened. Such a scenario is particularly impactful in the context of intimate partner violence [27] and due to the usage of Threema by high-ranking officials [52, 25].

The Threema app does provide some optional locking mechanisms. For example, the app can be locked with a PIN, a passphrase or by using biometrics, thus preventing this attack.

3.3.2 Attack 7 (Compression Side-Channel)

While the app locking mechanism does prevent a compelled access attacker from accessing the user’s data via the user interface, it does not prevent background processes from being executed. In particular, the app will attempt to make a new Threema Safe backup whenever the application is launched if the last successful backup is more than one day old. Moreover, incoming messages will be processed and may affect the contents of the backup. We exploit these facts in our final attack to recover a victim user’s private key.

Assumptions: Access to an unlocked phone, ability to observe the network.

Execution: Recall that a Threema Safe backup contains the private key of the user, and a list of the user’s contacts.

Before being sent to the server for storage, the backup is *compressed and then encrypted*. It is well-known that this combination of processing steps can be vulnerable to attack, especially if the attacker can control part of the plaintext before compression [40, 53]. Specifically, the length of the ciphertext may leak information about how much plaintext compression has been achieved, and this in turn can leak information about the degree of overlap between the attacker-controlled and the unknown parts of the plaintext. In Threema Safe, the XSalsa20 stream cipher is used for encryption, so a ciphertext leaks the exact length of the plaintext. Furthermore, usernames are handled locally in Threema. This means that if a user \mathcal{U} sends a message to victim \mathcal{V} , then \mathcal{U} ’s Threema ID and chosen nickname are updated in the local contact list of \mathcal{V} on receipt of the message. This happens even if \mathcal{V} ’s app is locked. This provides a means for an attacker to inject chosen plaintext into the backup of a victim. In this way, the conditions needed to build a compression-based side-channel attack are fulfilled. We provide a more detailed description of the attack in Appendix B.

Proof-of-Concept: We created a custom Threema server and instrumented the client code in order to redirect requests to our server rather than the actual Threema Safe server. This allows us to both see the length of the backups and to make backups fail by returning an HTTP error code. To automatically induce a new backup, we use Android debugging tools which can force an app to stop and restart. In order to measure the number of queries needed to complete the attack at scale, we also reimplemented the client backup mechanism, and simulated the attack locally. Figure 9 summarizes our experimental results. We successfully recovered the user’s private key in 4.7k out of 10k experiments; in these successful runs, we required a median of 19.4k backup attempts (mean 23.4k). In our PoC, each attempt requires around 2s, implying a total running time of around 11 hours for the attack.

The backup is normally transmitted to the server on a TLS connection, but with modern cipher suites (AES-GCM or ChaCha20-Poly1305), TLS does not hide the length of its payloads. So we consider the attack to remain viable even if the attacker can only see TLS-protected traffic as in our compelled access threat model.

Impact: Recovering the private key leads to a complete loss of security: after doing so the attacker can impersonate the user in any action and decrypt E2E protocol message.

4 Mitigations

We describe the mitigations we proposed to Threema, along with the ones that Threema employed in practice, and comment on what our attacks mean in practice for Threema users.

Forward Secrecy at the E2E level: We identify the lack of forward secrecy as one of the main weaknesses of Threema, since it accentuates the severity of many of our attacks. There

is no simple fix: a protocol redesign supporting the use of ephemeral keys is necessary. We recommend that Threema adopt the Signal protocol, whose cryptographic API is available through the libsignal library [44]. The Signal protocol has received extensive security analysis [28, 14, 41, 15] and provides forward secrecy and post-compromise security at the end-to-end layer. We note that Whatsapp has transitioned towards using the Signal protocol in the past [45], showing that such a change is feasible even for a messaging app with a huge user base.

Following the disclosure, Threema released Ibex, an additional protocol developed independently of our work. Ibex is layered on top of the E2E protocol, and aims to provide forward secrecy via an asynchronous, ephemeral key exchange. The keys derived from this exchange are used to instantiate a symmetric ratchet analogous to Signal's. The lack of an asymmetric ratchet means that Threema still lacks post-compromise security. Ibex remains opt-in at the time of writing. PFS of past conversations remains forsaken.

Strengthening the C2S Protocol: The least invasive modification to the C2S protocol that directly prevents our attack of Section 3.1.2 would be to include the server cookie inside the vouch box. This ensures that, as long as the server picks a fresh cookie every time the protocol is run, the vouch box cannot be trivially replayed. We recommend that Threema replaces its bespoke C2S protocol with the IK protocol of the Noise Protocol Framework [49], which allows for 0-RTT encrypted communications between the client and the server and provides authentication and forward secrecy. Furthermore, it prevents the attack of Section 3.1.1 (Impersonation by Randomness Failure). The Noise protocols have been analyzed and proven secure [22, 23], and Noise protocols have seen practical deployment in Wireguard [20] and Whatsapp [50].

Threema fixed attacks 1 and 2 via key separation using an additional KDF step with distinct labels, as well as by including the server cookie in the vouch box. Caching of ephemeral keys has been removed on the client and server. We do not believe these attacks affected Threema users at large, since we are not aware of practical mechanisms to recover ephemeral keys for (attack 1), and Threema reported that no Threema Gateway account was registered with the public key of their server in the past (attack 2).

Preventing Cross-Protocol Attacks: The attacks in Sections 3.1.2 (Vouch Box Forgery) and 3.2.3 (Kompromat Forgery) exploit the fact that the same “X25519-then-Encrypt” paradigm is used multiple times for different purposes but with the same keys. By using payloads created in one protocol in another, the adversary gains capabilities that they would not have when viewing the protocols separately. A secure application should cryptographically compartmentalise its protocols, taking a very conservative approach when using the same cryptographic material in different protocols (if doing so at all). A common way to prevent this class of attack is to use

proper *key separation*: the X25519 key material will then only be used to derive context-dependent keys using a KDF with proper labels. We stress that key material that is used to derive other keys should not be used itself for other purposes, since it is important to ensure that leaking one of the keys during its usage does not affect any other key, i.e. the keys should remain effectively *independent* of each other.

In correspondence, Threema indicated that they will make extended use of KDFs to ensure key separation.

Protecting Metadata: The simplest way to prevent the attacks of Sections 3.2.2 (Replay and Reflection Attacks) and 3.2.1 (Message Reordering) is to protect the integrity of the metadata contained in the E2E packet. This can be simply achieved by including the metadata in the (authenticated but not encrypted) “additional data” field of the AEAD scheme. This modification has zero message overhead and supersedes the usage of the metadata box, eliminating the need for an additional ciphertext and tag. Protecting the metadata in this way would prevent an adversary from swapping the source and destination of the message and thus prevent reflection attacks. To prevent message reordering attacks, the app could use a per-recipient counter for the AEAD nonces in the E2E protocol. This would also allow the recipient to detect lost or adversarially deleted messages, but would still leave the problem of how to avoid nonce repetitions when the Threema app is reinstalled or when the user changes device. We suggest that the app should re-run the registration protocol in this event, establishing a fresh key pair (a, A) .

Threema's introduction of Ibex fixes the reordering and replay attacks. The metadata box is now mandatory. Attacks 3 and 4 have limited implications if the underlying E2E plaintext is unknown.

Preventing Cloning via Threema ID Export: The ID export feature should be protected, possibly by requiring the user to provide a passphrase chosen at registration time, or at least by enforcing the use of the phone OS's authentication mechanisms such as biometrics and pin codes. It is up to Threema to strike a fair balance between usability and security, but we would err on the side of the latter.

The client now keeps track of ephemeral keys used in the C2S protocol and the user is alerted if an extraneous ephemeral key has been used in a C2S session. Attack 6 is the most likely to have been exploited in real world. This attack did not impact Threema users at large, but left at-risk individuals vulnerable to targeted attacks.

Mitigating the Compression Side-Channel: Threema should avoid using compression. To avoid an explosion in server-side storage requirements, a binary serialization format such as CBOR [13] could be used instead of JSON.

Threema disabled compression in backups. Attack 7 has serious implications, since observing the size of past Threema backups provides side channel information on a user's secret key, but does require targeting a specific user.

5 Conclusions

Threema is marketed as a secure app, suitable for both individuals and companies. It is used by government departments and high-profile politicians. On the other hand, the seven attacks we have presented highlight fundamental weaknesses in the design of Threema. Indeed, the Threema protocols lack basic properties that are nowadays considered *de rigueur* for a messenger app to be regarded as secure: forward secrecy with respect to a malicious server, and protection against replay, reflection, and reordering attacks. We believe that the cryptography in Threema has design flaws that need to be addressed in order to meet the security expectations of its users.

Stepping back from the details of our analysis, we believe that our work presents (at least) three useful lessons for developers when deploying complex cryptographic protocols.

First, *using modern, secure libraries for cryptographic primitives does not on its own lead to a secure protocol design*: It is possible to misuse libraries such as NaCl and libsodium when building a more complex protocol and developers must be wary not to be lulled into a false sense of security. A pertinent example is the C2S protocol in Threema, which has many issues with ephemeral keys. Another recent example is Bridgefy’s mis-integration of libsodium [6]. In some cases, it is the API of the libraries themselves that encourages bad cryptographic practice. For instance, NaCl’s authenticated encryption routines do not allow for associated data. This was the reason for the Threema’s bespoke metadata box construction. To their credit, Threema’s developers did (largely) avoid “rolling their own crypto” but this advice should be extended to “don’t roll your own cryptographic protocol”. Of course this advice is only useful if good alternatives are available. In the particular case of the C2S protocol, Threema could have used the Noise IK protocol or TLS.

Second, *beware of cross-protocol interactions*: A standalone protocol may appear to be secure, but cross-protocol interactions can undermine security guarantees, as we have shown in Threema with the vouch box and Kompromat attacks. So developers should be careful not to introduce such vulnerabilities. Every message should be cryptographically bound to its protocol and thereby rejected whenever it is used in a different protocol. This can be done by strictly following the key separation principle – use different keys for different purposes. Notably, a recent analysis of the secure messaging protocol Matrix also exploited cross-protocol interaction to violate the confidentiality of messages [5].

Third, *proactive, not reactive security*: Our inability to find an attack on a protocol does not imply it is secure: new attacks could be found at any moment and known attacks only get stronger over time. Often, software releases follow a *design-release-break-patch* process (a *reactive* approach [48]). This can be inconvenient for users and may also lead to the need to maintain backwards compatibility, opening the way to possible downgrade attacks [39, 12, 1, 9]. Developers should

adopt a *proactive* approach, where the protocol is formally analysed during the design stage. Ideally, different forms of analysis would be carried out, e.g. using symbolic approaches or by producing computational security proofs. While this cannot exclude every attack possibility, it can help improve confidence in the protocol’s security. Furthermore, protocols should be designed in such a way that this formal analysis can be more easily carried out. In the present setting, because the C2S protocol uses its session key directly in the handshake (to derive the metadata key), it is impossible to prove session key security, a standard property targeted in security proofs for key exchange protocols.

None of these lessons is fundamentally new, but given the results of our analysis of Threema, and the recent works in the field of secure protocol analysis [5, 6, 7], they apparently bear repeating.

References

- [1] David Adrian, Karthikeyan Bhargavan, Zakir Durumeric, Pierrick Gaudry, Matthew Green, J. Alex Halderman, Nadia Heninger, Drew Springall, Emmanuel Thomé, Luke Valenta, Benjamin VanderSloot, Eric Wustrow, Santiago Zanella Béguelin, and Paul Zimmermann. Imperfect forward secrecy: How Diffie-Hellman fails in practice. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-16, 2015*, pages 5–17. ACM, 2015.
- [2] Jan Ahrens. Threema protocol analysis. <https://blog.jan-ahrens.eu/files/threema-protocol-analysis.pdf>, March 2014.
- [3] Ange Albertini, Thai Duong, Shay Gueron, Stefan Kölbl, Atul Luykx, and Sophie Schmieg. How to abuse and fix authenticated encryption without key commitment. *IACR Cryptol. ePrint Arch.*, page 1456, 2020.
- [4] Martin R Albrecht, Jorge Blasco, Rikke Bjerg Jensen, and Lenka Mareková. Mesh messaging in large-scale protests: Breaking Bridgefy. In *Cryptographers’ Track at the RSA Conference*, pages 375–398. Springer, 2021.
- [5] Martin R. Albrecht, Sofía Celi, Benjamin Dowling, and Daniel Jones. Practically-exploitable cryptographic vulnerabilities in Matrix. <https://nebuchadnezzar-megolm.github.io>, September 2022.
- [6] Martin R Albrecht, Raphael Eikenberg, and Kenneth G Paterson. Breaking Bridgefy, again: Adopting libsodium is not enough. In *31st USENIX Security Symposium (USENIX Security 22)*, Boston, MA, August 2022. USENIX Association.

- [7] Martin R Albrecht, Lenka Mareková, Kenneth G Paterson, and Igors Stepanovs. Four attacks and a proof for Telegram. In *43rd IEEE Symposium on Security and Privacy (IEEE S&P 2022)*, 2022.
- [8] AppBrain. Threema - Android app on AppBrain. <https://www.appbrain.com/app/threema/ch.threema.app>, October 2022.
- [9] Nimrod Aviram, Sebastian Schinzel, Juraj Somorovsky, Nadia Heninger, Maik Dankel, Jens Steube, Luke Valenta, David Adrian, J. Alex Halderman, Viktor Dukhovni, Emilia Käsper, Shaanan Cohney, Susanne Engels, Christof Paar, and Yuval Shavitt. DROWN: breaking TLS using sslv2. In Thorsten Holz and Stefan Savage, editors, *25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016*, pages 689–706. USENIX Association, 2016.
- [10] J. Berger. openMittsu. <https://github.com/blizzard4591/openMittsu>, September 2016.
- [11] Daniel J Bernstein. Cryptography in NaCl. <https://cr.yp.to/highspeed/naclcrypto-20090310.pdf>, 2009.
- [12] Benjamin Beurdouche, Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, Pierre-Yves Strub, and Jean Karim Zinzindohoue. A messy state of the union: Taming the composite state machines of TLS. In *2015 IEEE Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*, pages 535–552. IEEE Computer Society, 2015.
- [13] Carsten Bormann and Paul E. Hoffman. Concise Binary Object Representation (CBOR). RFC 8949, December 2020.
- [14] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the Signal messaging protocol. In *2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26-28, 2017*, pages 451–466. IEEE, 2017.
- [15] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the Signal messaging protocol. *J. Cryptol.*, 33(4):1914–1983, 2020.
- [16] WinZip Computing. AES Encryption Information: Encryption Specification AE-1 and AE-2. <https://www.appbrain.com/app/threema/ch.threema.app>, 2009.
- [17] Graeme Connell. Technology deep dive: Building a faster ORAM layer for enclaves. <https://signal.org/blog/building-faster-oram>, August 2022.
- [18] Cas Cremers, Jaiden Fairoze, Benjamin Kiesl, and Aurora Naska. Clone Detection in Secure Messaging: Improving Post-Compromise Security in Practice. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, CCS '20*, pages 1481–1495, New York, NY, USA, October 2020. Association for Computing Machinery.
- [19] Yevgeniy Dodis, Paul Grubbs, Thomas Ristenpart, and Joanne Woodage. Fast Message Franking: From Invisible Salamanders to Encryptment. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Lecture Notes in Computer Science*, pages 155–186, Cham, 2018. Springer International Publishing.
- [20] Jason A. Donenfeld. WireGuard: Next generation kernel network tunnel. In *24th Annual Network and Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February 26 - March 1, 2017*. The Internet Society, 2017.
- [21] Benjamin Dowling, Felix Günther, and Alexandre Poirier. Continuous authentication in secure messaging. In Vijayalakshmi Atluri, Roberto Di Pietro, Christian Damsgaard Jensen, and Weizhi Meng, editors, *Computer Security - ESORICS 2022 - 27th European Symposium on Research in Computer Security, Copenhagen, Denmark, September 26-30, 2022, Proceedings, Part II*, volume 13555 of *Lecture Notes in Computer Science*, pages 361–381. Springer, 2022.
- [22] Benjamin Dowling and Kenneth G. Paterson. A cryptographic analysis of the WireGuard protocol. In Bart Preneel and Frederik Vercauteren, editors, *Applied Cryptography and Network Security - 16th International Conference, ACNS 2018, Leuven, Belgium, July 2-4, 2018, Proceedings*, volume 10892 of *Lecture Notes in Computer Science*, pages 3–21. Springer, 2018.
- [23] Benjamin Dowling, Paul Rösler, and Jörg Schwenk. Flexible Authenticated and Confidential Channel Establishment (fACCE): Analyzing the Noise Protocol Framework. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *Public-Key Cryptography - PKC 2020 - 23rd IACR International Conference on Practice and Theory of Public-Key Cryptography, Edinburgh, UK, May 4-7, 2020, Proceedings, Part I*, volume 12110 of *Lecture Notes in Computer Science*, pages 341–373. Springer, 2020.
- [24] Kit Eaton. These Apps Promise to Encrypt Your Smartphone Communications. <https://www.nytimes.com/2016/03/24/technology/personaltech/encryption-by-app-adds-security-to-smartphones.html>, March 2016.

- [25] Corin Faife. Swiss Army drops WhatsApp for homegrown messaging service, citing privacy concerns. <https://www.theverge.com/2022/1/7/22871881/swiss-army-whatsapp-messaging-threema-privacy-concerns-us-jurisdiction>, January 2022.
- [26] Marc Fischlin and Felix Günther. Multi-stage key exchange and the case of Google’s QUIC protocol. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 1193–1204. ACM, 2014.
- [27] Diana Freed, Jackeline Palmer, Diana Elizabeth Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. “A Stalker’s Paradise”: How intimate partner abusers exploit technology. In Regan L. Mandryk, Mark Hancock, Mark Perry, and Anna L. Cox, editors, *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, Montreal, QC, Canada, April 21-26, 2018*, page 667. ACM, 2018.
- [28] Tilman Frosch, Christian Mainka, Christoph Bader, Florian Bergsma, Jörg Schwenk, and Thorsten Holz. How secure is textsecure? In *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*, pages 457–472. IEEE, 2016.
- [29] Threema GmbH. GitHub - threema-ch/threema-android: Threema App for Android., December 2020.
- [30] Threema GmbH. Cryptography whitepaper. https://threema.ch/press-files/2_documentation/cryptography_whitepaper.pdf, November 2021.
- [31] Threema GmbH. Why Threema Instead of WhatsApp? <https://threema.ch/en/work/blog/posts/why-threema-instead-of-whatsapp>, May 2021.
- [32] Threema GmbH. About – Threema. <https://threema.ch/en/about>, 2022.
- [33] Threema GmbH. Threema changelog. <https://threema.ch/en/versionhistory>, 2022.
- [34] Threema GmbH. Threema provides maximum security and comprehensive privacy protection. <https://threema.ch/en/security>, 2022.
- [35] Threema GmbH. Threema.OnPrem. <https://threema.ch/en/onprem>, 2022.
- [36] Christoph Hagen, Christian Weinert, Christoph Sendner, Alexandra Dmitrienko, and Thomas Schneider. All the numbers are US: large-scale abuse of contact discovery in mobile messengers. In *28th Annual Network and Distributed System Security Symposium, NDSS 2021, virtually, February 21-25, 2021*. The Internet Society, 2021.
- [37] M. Heiderich, M. Wege, and C. Kean. Pentest and audit-report Threema mobile apps. https://threema.ch/press-files/2_documentation/security_audit_report_threema_2020.pdf, October 2020.
- [38] Fabian Ising, Damian Poddebniak, and Sebastian Schinzel. Security audit report Threema 2019. https://threema.ch/press-files/2_documentation/security_audit_report_threema_2019.pdf, March 2019.
- [39] Tibor Jager, Kenneth G. Paterson, and Juraj Somorovsky. One bad apple: Backwards compatibility attacks on state-of-the-art cryptography. In *20th Annual Network and Distributed System Security Symposium, NDSS 2013, San Diego, California, USA, February 24-27, 2013*. The Internet Society, 2013.
- [40] John Kelsey. Compression and information leakage of plaintext. In Joan Daemen and Vincent Rijmen, editors, *Fast Software Encryption, 9th International Workshop, FSE 2002, Leuven, Belgium, February 4-6, 2002, Revised Papers*, volume 2365 of *Lecture Notes in Computer Science*, pages 263–276. Springer, 2002.
- [41] Nadim Kobeissi, Karthikeyan Bhargavan, and Bruno Blanchet. Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach. In *2017 IEEE European Symposium on Security and Privacy, EuroS&P 2017, Paris, France, April 26-28, 2017*, pages 435–450. IEEE, 2017.
- [42] Brian A. LaMacchia, Kristin E. Lauter, and Anton Mityagin. Stronger security of authenticated key exchange. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *Provable Security, First International Conference, ProvSec 2007, Wollongong, Australia, November 1-2, 2007, Proceedings*, volume 4784 of *Lecture Notes in Computer Science*, pages 1–16. Springer, 2007.
- [43] Srikanth Lingala. Zip4j. <https://github.com/srikanth-lingala/zip4j>.
- [44] Signal Messenger LLC. libsignal. <https://github.com/signalapp/libsignal>.
- [45] Moxie Marlinspike. WhatsApp’s Signal protocol integration is now complete. <https://signal.org/blog/whatsapp-complete>, April 2016.

- [46] Moxie Marlinspike. Technology preview: Private contact discovery for Signal. <https://signal.org/blog/private-contact-discovery>, September 2017.
- [47] MITRE. CVE-2023-22899. <https://www.cve.org/CVERecord?id=CVE-2023-22899>, January 2023.
- [48] Kenneth G. Paterson and Thyla van der Merwe. Reactive and proactive standardisation of TLS. In Lidong Chen, David A. McGrew, and Chris J. Mitchell, editors, *Security Standardisation Research - Third International Conference, SSR 2016, Gaithersburg, MD, USA, December 5-6, 2016, Proceedings*, volume 10074 of *Lecture Notes in Computer Science*, pages 160–186. Springer, 2016.
- [49] Trevor Perrin. The Noise Protocol Framework. <http://www.noiseprotocol.org/noise.html>, July 2018.
- [50] Meta Platforms. WhatsApp encryption overview. <https://www.whatsapp.com/security/WhatsApp-Security-Whitepaper.pdf>, November 2021.
- [51] Mario Di Raimondo, Rosario Gennaro, and Hugo Krawczyk. Secure off-the-record messaging. In Vijay Atluri, Sabrina De Capitani di Vimercati, and Roger Dingledine, editors, *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society, WPES 2005, Alexandria, VA, USA, November 7, 2005*, pages 81–89. ACM, 2005.
- [52] Johannes Ritter. Was Russland stört, überzeugt Olaf Scholz. <https://m.faz.net/aktuell/wirtschaft/unternehmen/threema-was-russland-stoert-ueberzeugt-olaf-scholz-18248712.amp.html>, January 2022.
- [53] Juliano Rizzo and Thai Duong. The CRIME attack. https://docs.google.com/presentation/d/11eBmGiHbYcHR9gL5nDyZChu_-1Ca2Gizeu0faLU2H0U, 2012.
- [54] Paul Rösler, Christian Mainka, and Jörg Schwenk. More is less: On the end-to-end security of group chats in Signal, WhatsApp, and Threema. In *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018, London, United Kingdom, April 24-26, 2018*, pages 415–429. IEEE, 2018.
- [55] Markku-Juhani O. Saarinen and Jean-Philippe Aumasson. The BLAKE2 Cryptographic Hash and Message Authentication Code (MAC). RFC 7693, November 2015.
- [56] Michel Schreiner and Thomas Hess. Examining the role of privacy in virtual migration: The case of WhatsApp and Threema. In *21st Americas Conference on Information Systems, AMCIS 2015, Puerto Rico, August 13-15, 2015*. Association for Information Systems, 2015.
- [57] Jochen Siegle. Schutz beim Chatten: Threema wird noch sicherer. <https://www.nzz.ch/digital/whatsapp-alternative-threema-wird-noch-sicherer-ld.1498576>, July 2019.
- [58] Soatok. Threema: Three Strikes, You’re Out. <https://soatok.blog/2021/11/05/threema-three-strikes-youre-out/>, November 2021.
- [59] SRF. Threema setzt sich durch - Schweizer Armee verbietet WhatsApp und Co. <https://www.srf.ch/news/schweiz/threema-setzt-sich-durch-schweizer-armee-verbietet-whatsapp-und-co>, January 2022.
- [60] Paul C. van Oorschot and Michael J. Wiener. Parallel collision search with cryptanalytic applications. *J. Cryptol.*, 12(1):1–28, 1999.
- [61] WhatsApp. WhatsApp Help Center - Answering your questions about WhatsApp’s January 2021 Privacy Policy update. <https://faq.whatsapp.com/general/security-and-privacy/answering-your-questions-about-whatsapps-privacy-policy/?lang=en>, January 2021.

A Contact Discovery Protocol

If the user gives permission to do so, Threema can run a protocol to check which of the people in the user’s phone’s contact list are using Threema. To do so, it must compare the user’s contact list with the list of all registered Threema users, while at the same time preserving both the privacy of the user and the confidentiality of the list of registered Threema users.

Threema does this by hashing all the contacts of the user and sending the hashes to the Threema server over TLS. The server will reply with the hashes that correspond to users that have previously registered with Threema. The hash algorithm is implemented using HMAC-SHA256 with a fixed key. This is done to “ensure that hashes generated by Threema are unique and do not match those of any other app” [30]. This is similar to the process used by apps such as WhatsApp, Telegram and Signal [36]. Such protocols tend to raise privacy concerns [36] due to the information leaked to the server, including information about people who might not even be aware of the app’s existence. Hashing the information is insufficient, since, for example, phone numbers carry little entropy and can be brute-forced. This allows the service provider to build a partial social graph for an entire population. Signal tries to provide stronger guarantees by implementing their contact discovery protocol in an Intel SGX enclave [46], minimizing the information leak due to memory access patterns, and using techniques such as Oblivious RAM to minimize the information leaked to the server [17].

```

{
  "info": { ... },
  "user": {
    "privatekey": <private key>,
    "nickname": "user_nickname"
  },
  "contacts": [
    {
      "identity": "ABCDEFGH",
      "nickname": "other_user_nickname"
    }
  ],
  ...
}

```

Figure 10: Example of a backup, with the private key redacted.

To prevent abuse of the system, anyone that wants to run the contact discovery protocol must first obtain a *match token* by authenticating to the Threema server. To obtain one, the client must pass a challenge-response protocol which is essentially the same as the one done for the registration process, except for the fact that the nonce is now randomly chosen by the client, rather than being fixed.

B Compression Side Channel

From a high-level perspective, the attack leverages the fact that nicknames are handled client-side: when a message is received, the nickname contained in the metadata is set as the current nickname of the sender. Since nicknames are contained in the backup, this gives the attacker partial control of the plaintext in the victim’s backup. The objective of the attacker is to leverage this partial control in order to force a backlink to be created from the attacker’s nickname in the contacts field to the private key in the user field. We depict an example of a JSON backup, with only the relevant information present, in Fig. 10. If the attacker’s nickname and the private key are sufficiently close together, they will both fall within the same sliding window of the compression algorithm. This ensures that there is the possibility of a backlink being created. In order to increase the chances of a backlink being created, we include a string called a *canary* which we know is already included in the JSON backup, just before the string that we want to leak. This tricks the compression algorithm into finding redundancy between the nickname and the private key, creating the compression side-channel that we need.

For our experiments, we used “privatekey” as a canary for the beginning of the private key and “=” for the end of the private key. The equal sign always appears as the last base64 character of the key due to padding of the private key during encoding, since the key is 32 bytes long and must be padded to a multiple of 6 bits. We noticed that both canary strings were necessary. The reasons for this are opaque to us. Between the two canary strings, the attacker needs to place,

in order, its guess for the next character of the key, and the currently known characters of the key (initially none, thus an empty string). If the guess is correct, it should result in a shorter Threema Safe ciphertext. If the shortest ciphertext was induced by only one character among all guesses, then the currently known characters of the key are updated to include the current guess. Otherwise, two-character guesses can be attempted. This method effectively tries to guess characters of the base64 encoding of the private key from the last one to the first one.

In Fig. 9, we show the number of compression side-channel oracle queries needed to extract 31 characters of the private key, over 10,000 experiments; of these attempts 4727 were successful and 5273 failed to recover key material. Over the successful experiments, the median number of queries needed was 19,382 and the mean was 23,382. We remark that “failure” here means that our *automated* exploit did not successfully retrieve the key. If a targeted attack were to be run, one could manually guide the process. In fact, when a failure happens it is often because the exploit begins revealing characters of other parts of the backup, which can be easily spotted as the exploit is running.

To obtain the required queries, it suffices for an attacker to forcibly close the app and start it again. This induces the app into retrying to upload the backup. The forced closure can be done by either using debugging tools or by simulating user input on the unlocked phone. Then, in the setting of a compromised Threema Safe server, it is sufficient for the server to reply with an HTTP error code. Otherwise, a network attacker can forcibly close the connection, making the upload fail. Both of these techniques ensure that a new upload will be attempted at the next restart. At the pace of one restart every two seconds, an attacker can recover 31 base64 characters of the private key in a median time of about 11 hours (in a successful case).

Not all the bytes of the private key have to be recovered, since algorithms such as the van Oorschot-Wiener algorithm [60] can be used to recover the remaining bits. Indeed, this creates a trade-off between the number of queries and the amount of offline work required by the attacker. Since the latter does not require access to the device, it is convenient to execute just the sufficient number of queries required to make the offline attack feasible. For example, the van Oorschot-Wiener algorithm has time complexity $O(2^{n/2})$, where n is the bit-size of the set to be searched, and it is parallelisable. Recovering 31 base64 characters via the compression side channel (as above), leaves 13 characters to be recovered, equivalent to a search space of size 2^{78} after base64 decoding. This makes the offline part of the attack perfectly feasible.