

Threema. Cryptography Whitepaper

Threema uses modern cryptography based on open source components that strike an optimal balance between security, performance and message size.

In this whitepaper, the algorithms and design decisions behind the cryptography in Threema are explained.

Contents

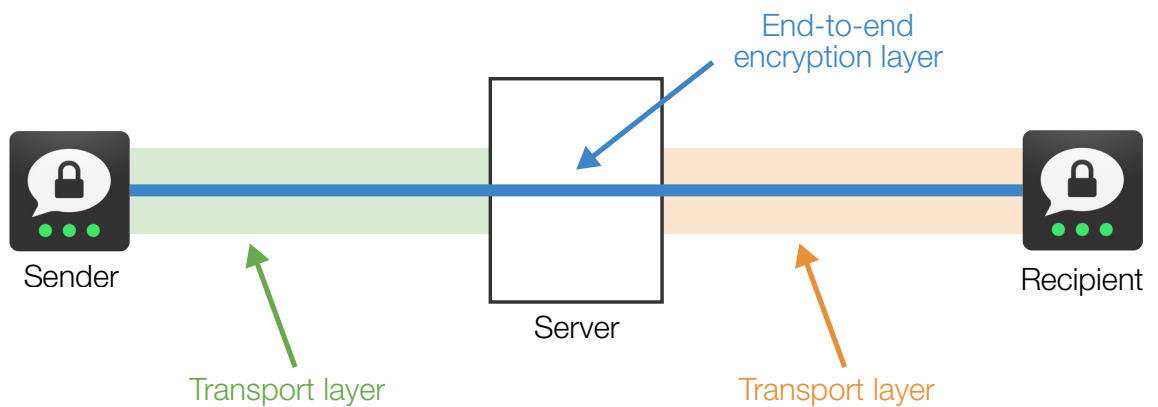
Overview	3
End-to-End Encryption	4
Key Generation and Registration	4
Key Distribution and Trust	4
Message Encryption	5
Group Messaging	6
Key Backup	6
Client-Server Protocol Description	7
Directory Access Protocol	7
Media Access Protocol	8
Cryptography Details	8
Key Lengths	9
Random Number Generation	9
Forward Secrecy	10
Repudiability	10
Padding	10
Replay Prevention	10
Local Data Encryption	11
iOS	11
Android	11
Key Storage	11
iOS	11
Android	12
Push Notifications	12
iOS	12
Android	12
Address Book Synchronization	13
Linking	13
ID Revocation	14
An Example	14
Profile Pictures	14
Web Client	15
Architecture	15
Connection Buildup	15
WebRTC Signaling	16
WebRTC Connection Buildup	16
Threema Calls	18
Signaling	18
Call Encryption	18
Audio Encoding	18
Video Encoding	18
Privacy / IP exposure	18
Threema Safe	19
Overview	19
Backup format	19
Encryption	19
Upload/Storage	20
Backup intervals	20
Restore/Decryption	20
Running a custom Threema Safe server	21
Security Recommendations	22

Overview

Threema uses two different encryption layers to protect messages between the sender and the recipient.

- **End-to-end encryption layer:** this layer is between the sender and the recipient.
- **Transport layer:** each end-to-end encrypted message is encrypted again for transport between the client and the server, in order to protect the header information.

These layers and other important aspects of the cryptography in Threema are described in detail in the following chapters. The crucial part is that the end-to-end encryption layer passes through the server uninterrupted; the server cannot remove the inner encryption layer.



End-to-End Encryption

In Threema, all messages (whether they are simple text messages, or contain media like images, videos or audio recordings) are end-to-end encrypted. For this purpose, each Threema user has a unique asymmetric key pair consisting of a public and a private key based on Elliptic Curve Cryptography. These two keys are mathematically related, but it is not technically feasible to calculate the private key given only the public key.

Key Generation and Registration

When a Threema user sets up the app for the first time, the following process is performed:




1. The app generates a new key pair by choosing a private key at random¹, storing it securely on the device, and calculating the corresponding public key over the Elliptic Curve (Curve25519).
2. The app sends the public key to the server.
3. The server stores the public key and assigns a new random Threema ID, consisting of 8 characters out of A-Z/0-9.
4. The app stores the received Threema ID along with the public and private key in secure storage on the device.

Key Distribution and Trust

The public key of each user is stored on the directory server, along with its permanently assigned Threema ID. Any user may obtain the public key for a given Threema ID by querying the directory server. The following input values can be used for this query:

- a full 8-character Threema ID
- a hash of a E.164 mobile phone number that is linked with a Threema ID
 - see section “Address Book Synchronization” for details on the hashing
- a hash of an email address that is linked with a Threema ID

The Threema app maintains a “verification level” indicator for each contact that it has stored. The following three levels are possible:

-  Red (level 1): The public key has been obtained from the server because a message has been received from this contact for the first time, or the ID has been entered manually. No matching contact was found in the user’s address book (by phone number or email), and therefore the user cannot be sure that the person is who they claim to be in their messages.
-  Orange (level 2): The ID has been matched with a contact in the user’s address book (by phone number or email). Since the server verifies phone numbers and email addresses, the user can be reasonably sure that the person is who they claim to be.
-  Green (level 3): The user has personally verified the ID and public key of the contact by scanning their 2D code. Assuming the contact’s device has not been hijacked, the user can be very sure that messages from this contact were really written by the person that they indicate.

To upgrade a contact from the red or orange to the green level, the user must scan that contact’s public QR code. This QR code is displayed in the “My ID” section of the app, and uses the following format:

```
3mid:<identity>,<publicKeyHex>
```

where <identity> is the 8-character Threema ID, and <publicKeyHex> is the hexadecimal (lowercase) representation of the user’s 32 byte public key. The app verifies that the scanned public key matches the one that was returned by the directory server.

¹ including entropy generated by user interaction; see section “Random Number Generation” for details

Key Fingerprints

The app displays a key fingerprint for each contact, and for the user's own identity. This can be used to compare public keys manually. A key fingerprint is created by hashing the raw public key with SHA-256, and taking the first 16 bytes of the resulting hash as the fingerprint.

Message Encryption

Threema uses the “Box” model of the [NaCl Networking and Cryptography Library](#) to encrypt and authenticate messages. For the purpose of this description, assume that Alice wants to send a message to Bob. Encryption of the message using NaCl works as follows:

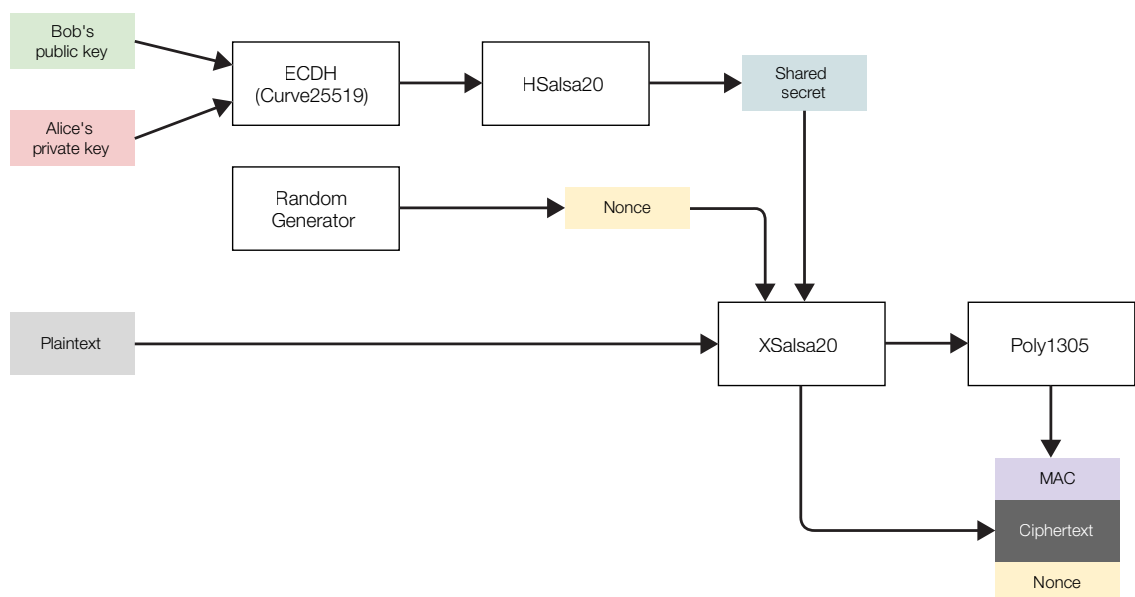
Preconditions

1. Alice and Bob have each generated a key pair consisting of a public and a private key.
2. Alice has obtained the public key from Bob over an authenticated channel.

Procedure to encrypt a message

1. Alice uses Elliptic Curve Diffie-Hellman (ECDH) over the curve Curve25519 and hashes the result with HSalsa20 to derive a shared secret from her own private key and Bob's public key.
 Note that due to the properties of the elliptic curve, this shared secret is the same if the keys are swapped (i.e. if Bob's private key and Alice's public key are used instead).
2. Alice generates a random nonce.
3. Alice uses the XSalsa20 stream cipher with the shared secret as the key and the random nonce to encrypt the plaintext.
4. Alice uses Poly1305 to compute a Message Authentication Code (MAC), and prepends it to the ciphertext. A portion of the key stream from XSalsa20 is used to form the MAC key.
5. Alice sends the MAC, ciphertext and nonce to Bob.

By reversing the above steps and using his own private key and Alice's public key, Bob can decrypt the message and verify its authenticity.



For further details, see [Cryptography in NaCl](#).

Group Messaging

In Threema, groups are managed without any involvement of the servers. That is, the servers do not know which groups exist and which users are members of which groups. When a user sends a message to a group, it is individually encrypted and sent to each other group member. This may appear wasteful, but given typical message sizes of 100-300 bytes, the extra traffic is insignificant. Media files (images, video, audio) are encrypted with a random symmetric key and uploaded only once. The same key, along with a reference to the uploaded file, is then distributed to all members of the group.

Key Backup

The user can back up their private key so that they are able to move their Threema ID to another device, or to restore it in case of loss/damage to the device. The app automatically reminds the user to do so, as without a backup, there is no way to recover a lost private key.

To generate a backup, the user must first choose a password (min. 8 characters).

The backup data is then generated as follows:

1. Calculate the SHA-256 hash of the following binary string:

`<identity><private key>`

where `<identity>` is the 8 character Threema ID, and `<private key>` is the 32 byte private key.

Keep only the first two bytes of the resulting hash. It is used during restoration to verify with reasonable confidence that the provided password was correct.

2. Choose a random 64 bit salt.

3. Derive a 256-bit encryption key from the given password and the random salt using PBKDF2 with HMAC-SHA256 and 100000 iterations:

$key_{enc} = \text{PBKDF2}(\text{HMAC-SHA256}, \text{password}, \text{salt}, 100000, 32)$

4. Use the XSalsa20 stream cipher with key_{enc} and an all-zero nonce to encrypt a binary string of the following format:

`<identity><private key><hash>`

where `<hash>` is the two byte truncated hash calculated in step 1. Note that this can be done using the [crypto_stream function of NaCl](#).

5. Prepend the salt to the encrypted string.

6. Base32 encode the result.

7. Split the Base32 encoded string into groups of four characters and separate the groups with dashes.

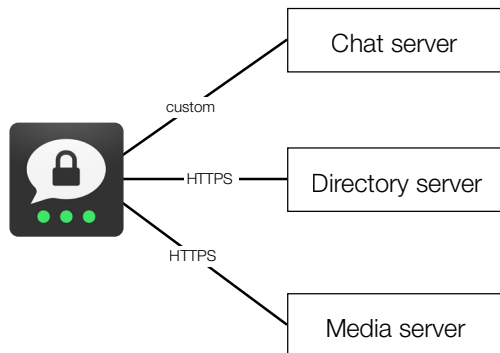
The result will look like this:

XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-
XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX

The encrypted backup data consists of 80 characters (A-Z, 2-7). It is also displayed in QR code form so that the user can more easily transfer it to another device.

Client-Server Protocol Description

Threema communicates with three different types of servers. To access the directory and to download/upload encrypted media files, HTTPS is used. To transport the actual chat messages, a custom protocol built on TCP is used.



Chat Protocol (Message Transport Layer)

The chat protocol is used to transport incoming and outgoing messages between the client (app) and the Threema servers. It is a custom binary protocol that uses the NaCl library to secure the connection on the application layer. Its properties are:

- Provides Forward Secrecy
 - New temporary keys are generated whenever the app restarts
 - Temporary private keys are never stored in permanent storage, but are only kept in volatile memory
- Optimized for minimal overhead
 - Small message headers
 - Minimum number of round-trips required for connection setup
- User is authenticated using his public key during connection setup
 - Ensures that a user can only log in if he is in possession of the private key for the Threema ID

Directory Access Protocol

The directory access protocol is used for the following purposes:

- Creating new Threema IDs
- Fetching the public key of another user
- Linking email addresses and mobile phone numbers
- Matching address book contacts (hashes of phone numbers and email addresses) to Threema IDs

Requests are authenticated with a challenge/response protocol based on the user's Threema ID and key pair.

HTTPS (HTTP with TLS) is used as the transport protocol. Strong TLS cipher suites with forward secrecy (ECDHE/DHE) and TLS v1.2 are supported. In order to preclude man-in-the-middle (MITM) attacks even if the system's trusted CA store has been tampered with, or if a trusted CA has illegally issued a certificate for a Threema domain name, the app uses public-key pinning with hardcoded pins to only accept specific, Threema-owned server certificates.

Media Access Protocol

The media servers are used for temporary storage of large media data (e.g. images, videos, audio recordings). Such media is not sent directly via the chat protocol. Instead, the following procedure is used:

1. The sender encrypts the media file with a random 256-bit symmetric key using XSalsa20 and adds a Poly1305 authenticator.
2. The sender uploads the encrypted media file to a media server via HTTPS.
3. The media server assigns a unique ID for this upload and returns it to the sender.
4. The sender sends an end-to-end encrypted message to the recipient, which contains the media ID and the symmetric key.
5. The recipient receives and decrypts the end-to-end encrypted message, obtaining the media ID and the symmetric key.
6. The recipient uses the media ID to download the encrypted media file from the media server.
7. The recipient decrypts the media file using the symmetric key.
8. The recipient signals the media server to delete the file, unless it was sent in a group conversation, in which case it remains on the media server for the maximum message lifetime (14 days), after which it is deleted automatically in any case.

The media servers use the same TLS configuration (cipher suites with forward secrecy, TLS v1.2) as the directory servers, and the app uses public-key pinning when accessing them.

Cryptography Details

As mentioned earlier, Threema uses the [NaCl Networking and Cryptography Library](#) for both the end-to-end encryption, and to secure the chat protocol at the transport level. This library uses a collection of algorithms to provide a simplified interface for protecting a message using what the authors call “public-key authenticated encryption” against eavesdropping, spoofing and tampering. By default, and as implemented in Threema, it uses the following algorithms:

- Key derivation: Elliptic Curve Diffie-Hellman (ECDH) over the curve [Curve25519](#)
- Symmetric encryption: XSalsa20
- Authentication and integrity protection: Poly1305-AES

It is worth mentioning that the ECC curve used by NaCl (and thus by Threema) is **not one of the NIST-recommended curves** that have been suspected of containing deliberately selected weakening constants in their specification.

For more details, see [Cryptography in NaCl](#).

The following implementations of the cryptographic algorithms are used:

	Curve25519	XSalsa20	Poly1305
iOS	Donna C implementation	reference C implementation	“53” C implementation
Android	reference C implementation (native code, fallback to jnacl)		

Key Lengths

The asymmetric keys used in Threema have a length of 256 bits, and their effective ECC strength is 255 bits.

The shared secrets, which are used as symmetric keys for end-to-end message encryption (derived from the sender's private key and the recipient's public key using ECDH, and combined with a 192 bit nonce), have a length of 256 bits.

The random symmetric keys used for media encryption are also 256 bits long.

The message authentication code (MAC) that is added to each message to detect tampering and forgery has a length of 128 bits.

Discussion of reasonable key lengths

According to [NIST Special Publication 800-57](#) (page 53), the security level of ECC based encryption at 255 bits can be compared to RSA at roughly 2048 to 3072 bits, or a symmetric security level of ~128 bits. The possibility of a successful brute force attack on a key with a 128 bit security level is considered extremely unlikely using current technology and knowledge, according to the judgment of reputable security researchers. A revolutionary breakthrough in mathematics or quantum computing would most possibly render keys breakable anyway, whether they are 128 or 256 bits in length.

- <http://cr.yt.to/talks/2005.09.19/slides.pdf> (Daniel J. Bernstein, author of Curve25519)
- <https://www.imperialviolet.org/2014/05/25/strengthmatching.html> (Adam Langley, Google security researcher)

Random Number Generation

Threema uses random numbers for the following purposes, listed by descending order of randomness quality required:

- Private key generation
- Symmetric encryption of media files
- Nonces
- Backup encryption salt
- Padding amount determination

Nonces and salts must never repeat, but they are not required to be hard to guess.

To obtain random numbers, Threema uses the system-provided random number generator (RNG) intended for cryptographically-secure purposes provided by the device's operating system. The exact implementation varies among operating systems:

Platform	Facility	Implementation
iOS	/dev/random	Yarrow (SHA1)
Android	/dev/random ²	Linux PRNG

Note that Threema does not use the flawed [Dual_EC_DRBG](#) random number generator.

² To avoid the flawed SecureRandom implementation in some Android versions, Threema uses its own implementation that directly accesses /dev/random (which is not affected by the SecureRandom implementation bug).

User generated entropy for private key generation

Due to the requirement for very high quality randomness when generating the long-term private key, the user is prompted to generate additional entropy by moving a finger on the screen. The movements (consisting of coordinate values and high-resolution timestamps) are continuously collected and hashed for several seconds. The resulting entropy is then mixed (XOR) with entropy obtained from the system's RNG.

Forward Secrecy

Due to the inherently asynchronous nature of mobile messengers, providing reliable Forward Secrecy on the end-to-end layer is difficult. Key negotiation for a new chat session would require the other party to be online before the first message can be sent. Experimental schemes like caching pre-generated temporary keys from the clients on the servers increase the server and protocol complexity, leading to lower reliability and more potential for mistakes that impact security. The user experience can also be diminished by events that are not under the control of the sender, for example when the recipient loses their phone's data, and along with it the ephemeral keys. Due to these and the following considerations, Threema has implemented Forward Secrecy on the transport layer only:

- Reliability is very important to ensure that users do not feel negatively impacted by the cryptography.
- The risk of eavesdropping on any path through the Internet between the sender and the server, or between the server and the recipient, is orders of magnitude greater than the risk of eavesdropping on the server itself.

With Threema's implementation, an attacker who has captured and stored the network traffic between a client and a Threema server will not be able to decrypt it even if he learns the private key of the server or the client afterwards.

Repudiability

In general, cryptographically signed messages also provide non-repudiation; i.e. the sender cannot deny having sent the message after it has been received. The NaCl library's box model uses so-called *public-key authenticators* instead, which guarantee repudiability (see <http://nacl.cryp.to/box.html>, "Security model"). Any recipient can forge a message that will look just like it was actually generated by the purported sender, so the recipient cannot convince a third party that the message was really generated by the sender and not forged by the recipient. However, the recipient is still protected against forgeries by third parties. The reason is that in order to perform such a forgery, the private key of the recipient is needed. Since the recipient himself will know whether or not he has used his own private key for such a forgery, he can be sure that no third party could have forged the message.

Padding

In order to thwart attempts to guess the content of short messages by looking at the amount of data, Threema adds a random amount of PKCS#7 padding to each message before end-to-end encryption.

Replay Prevention

The Threema app remembers the nonce of every message that has been sent in the past, and rejects messages with duplicate nonces. Since the server cannot successfully modify the nonce of a message without knowing the private key of one of the parties involved in the communication, this prevents a malicious server from replaying/duplicating previously sent messages.

Local Data Encryption

The Threema app stores local data (such as the history of incoming and outgoing messages, and the contact list) in encrypted form on the device. The way in which this data is encrypted varies among platforms.

iOS

On iOS, Threema stores local data in a Core Data database, which is backed by files in the app's private data directory. The iOS sandbox model ensures that no other apps can access this data directory. All files stored by Threema are protected using the iOS Data Protection feature using the Data Protection class `NSFileProtectionCompleteUntilFirstUserAuthentication` (see [iOS Security Guide](#), page 13). The encryption key used to protect the files is derived from the device's UID key and the user's passcode. A longer passcode will provide better security (on devices equipped with a Touch ID sensor, longer passcodes can be used without a loss of comfort as the passcode only needs to be entered after a reboot).

Note that the private key for the Threema ID is stored separately in the iOS Keychain.

Android

On Android, Threema stores local data in an SQLite database within the app's private data directory. Android ensures that no other apps can access this data directory (as long as the device is not rooted). The database itself is protected by SQLCipher with AES-256, and any media files (which are stored separately in the app's private directory within the file system) are also encrypted using AES. The key is randomly generated when the database is first created, and can be protected with a user-supplied passphrase (which is of course necessary if the user wishes to take advantage of this encryption). The passphrase is never written to permanent storage and therefore needs to be entered whenever the app process restarts (e.g. after a low-memory situation, or after the device has rebooted). Alternatively, the user may enable full-device encryption if supported by the device and Android version.

Key Storage

On the user's mobile device, the private key is stored in such a way as to prevent access by other apps on the same device, or by unauthorized users. The procedure differs between platforms.

iOS

- The private key is stored in the iOS Keychain with the attribute **kSecAttrAccessibleAfterFirstUnlock-ThisDeviceOnly**.
- Only the Threema app can access this keychain entry.
- While the entire iOS keychain is included in backups via iTunes or iCloud, the `kSecAttrAccessibleAfterFirstUnlockThisDeviceOnly` attribute causes the keychain entry of Threema to be encrypted with a device-specific key ("UID key") that cannot be read directly and is not known to Apple. Therefore, the keychain entry is only usable if the iTunes/iCloud backup is restored to the same device.
 - See [iOS Security Guide](#), page 14
- When the Threema app is deleted and reinstalled, the keychain entry persists and the Threema ID is not lost.

Android

- The private key is stored in a file in the app's private data directory.
- Other apps cannot access this file.
- AES-256-CBC encryption is applied to the private key before it's written to the file. The key for this encryption is stored separately and can be protected using a passphrase (see section "Local Data Encryption" for details).

Push Notifications

Threema uses the standard push notification service provided by the operating system. This is necessary to ensure reliable delivery of push notifications, as mobile operating systems place high importance on conserving battery power, and thus severely restrict background operation of apps (such as would be needed to maintain a separate, app-specific persistent connection for push notifications). Channelizing push notifications through the operating system's standard mechanism thus helps conserve battery power, and on iOS, it is the only means possible and accepted in the App Store.

iOS

The Threema-specific APNS payload consists of a JSON/Base64 encoded NaCl "[crypto_secretbox](#)" ciphertext and corresponding nonce. The symmetric key used by the Threema chat server to encrypt this payload is chosen at random by the client, and securely transmitted to the server beforehand using the chat protocol described above. The symmetric key is not known to Apple. When a push notification arrives, the Threema app decrypts the APNS payload using the symmetric key, and thus recovers the Threema ID and nickname of the sender of the corresponding message, as well as the message ID. With this information, the app can obtain the actual encrypted message by connecting to the chat server, and then decrypt the message and display a preview (if enabled) and the sender's contact name in the local notification.

Android

The GCM payload is empty; that is, the GCM message only serves to wake up the Threema app in the background. It then connects to the chat server to retrieve any pending messages, decrypts them, and shows local notifications for them.

Address Book Synchronization

Threema optionally lets the user discover other Threema contacts by synchronizing with the phone's address book. If the user chooses to do so, the following information is uploaded through a TLS connection to the directory server:

- HMAC-SHA256 hash of each email address found in the phone's address book
 - Key: 0x30a5500fed9701fa6defdb610841900febb8e430881f7ad816826264ec09bad7
- HMAC-SHA256 hash of the E.164 normalized form of each phone number found in the phone's address book
 - Key: 0x85adf8226953f3d96cfd5d09bf29555eb955fcd8aa5ec4f9fcd869e258370723

The directory server then compares the list of hashes from the user with the known email/phone hashes of Threema IDs that have been linked with an email address and/or phone number. Any matches are returned by the server as a tuple (Threema ID, hash). Only those hashes that have already been submitted by the user are returned (i.e. if the user submits an email hash which then matches a Threema ID, the server will only return the email hash of that ID and not the linked phone number's hash, even if one exists). After returning the matches to the client, the directory server discards the submitted hashes.

Use of HMAC keys

The reason why HMAC-SHA256 is used instead of plain SHA256 is not for obfuscation, but as a best practice to ensure that hashes generated by Threema are unique and do not match those of any other application. Obviously the keys need to be the same for all users, as random salting (such as is used when hashing passwords for storage) cannot be used here because the hashes of all users must agree so that matching contacts can be found.

A word about hashing phone numbers

Due to the relatively low number of possible phone number combinations, it is theoretically possible to crack hashes of phone numbers by trying all possibilities. This is due to the nature of hashes and phone numbers and cannot be solved differently (using salts like for hashing passwords does not work for this kind of data matching). Therefore, the servers treat phone number hashes with the same care as if they were raw/unhashed phone numbers. Specifically, they never store hashes uploaded during synchronization in persistent storage, but instead discard them as soon as the list of matching IDs has been returned to the client.

Linking

If a user chooses to link their Threema ID to an email address or a phone number, the directory server verifies that the user actually owns the email address or phone number in question.

- For email addresses, a verification email with a hyperlink is sent to the user. The user must open the hyperlink and confirm in the browser before the ID link is established.
- For phone numbers, the directory server sends an SMS message with a random 6 digit code. The user must enter that code in the app, which sends it back to the server to verify the phone number.
 - If the user cannot receive the SMS message, he/she may choose to receive an automated phone call in which the code is read out.

ID Revocation

Threema users may revoke their ID at any time by going to a website (<https://myid.threema.ch/revoke>) and entering their Threema ID and a revocation password that they have set beforehand. Revoked IDs can no longer log in, and they will be shown in strikethrough text on (or, depending on the users' settings, disappear from) the contact lists of other users within 24 hours. Messages cannot be sent to revoked IDs.

Revocation passwords are hashed with SHA256 on the client side, but only the first four bytes are sent to and stored/compared on the directory server. This mitigates the possibility of brute-force hash cracking on the server side to recover the user's password. Obviously, the downside is that the stored key has significantly less entropy than a reasonable user password and a large number of possible passwords correspond to each key. However, the server can easily limit the number of key revocation attempts over a period of time.

An Example

A user has set an eight character revocation password with a combination of characters consisting of the lower-case latin alphabet (a-z) and digits. This results in $36^8 \approx 2^{41.36}$ combinations.

Therefore, a given key can correspond with over 600 different passwords of this length and character set. On the other hand, the probability of randomly choosing the correct hash value without knowing anything about the password is $2^{-32} \approx 0.000000023\%$. Assuming that a user gets three attempts per day, the probability of finding the correct hash value within 10 years is $\approx 0.000255\%$.

Profile Pictures

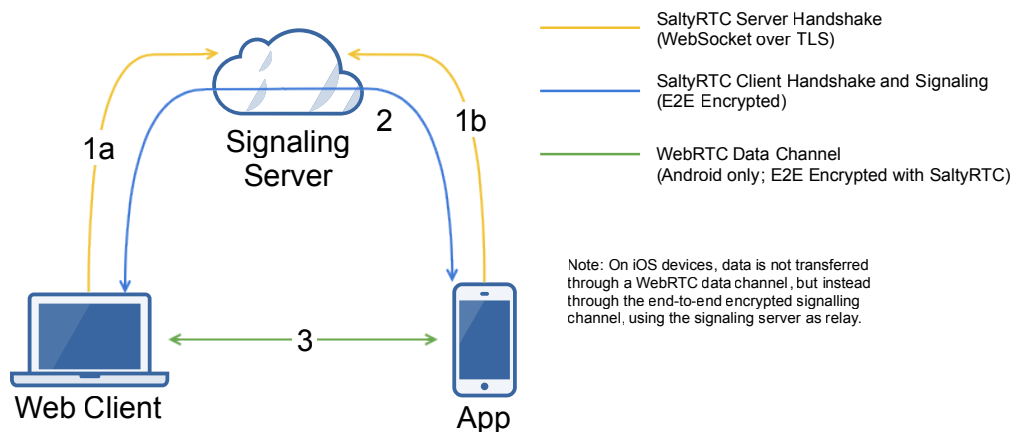
It is possible for a user to select a photo to be used as his/her "profile picture". The user can choose within the app whether he/she wants to share the profile picture with all their contacts, with only those contacts selected from a list, or with nobody. If he/she chose to share the profile picture and sends a message to one of the designated recipients, the newly set profile picture is encrypted with a random symmetric key within the app, and then uploaded to the media server like a regular image message (see "Media Access Protocol" on page 8). A reference to the encrypted picture on the media server, along with the symmetric key, is subsequently sent as an end-to-end encrypted message in the background whenever the user sends a regular message to one of the designated contacts who has not received the current profile picture yet. As the media server deletes files after two weeks, the profile picture is uploaded again once a week (if necessary), using a new symmetric key each time.

Web Client

The Threema web client communicates directly with the Threema app over [WebRTC](#) data channels (Android) and WebSockets (iOS). In addition to the standard WebRTC channel encryption (Android) and WebSocket TLS transport encryption (iOS), all packets are end-to-end encrypted using NaCl. The WebRTC signaling is also implemented with end-to-end encryption using the [SaltyRTC protocol](#).

The following paragraphs explain the general architecture of the Threema web client implementation as well as some aspects of the signaling protocol. If you want to learn more about the WebRTC signaling protocol, please refer to the [main specification](#) as well as the [WebRTC task specification](#).

Architecture



The web client and the app form a client-server relationship. All data about new events and old conversations is exchanged directly between the two through an end-to-end encrypted channel.

When the web client is started, it requests the initial data (conversations, contacts, avatars, etc.) from the app. The app responds with the requested data.

To send a message to a contact, the web client dispatches a message request to the app, which will then send the message to the recipient. When a new message arrives, the app notifies the web client. No messages are exchanged directly between the web client and the Threema chat server. The user's private key never leaves the device.

Connection Buildup

The full connection buildup happens in three stages.

1. SaltyRTC Server Handshake

First, the web client generates a new permanent key-pair and a random authentication token. It then connects to the SaltyRTC signaling server as initiator. The hex-encoded public permanent key is used as the WebSocket path.

The public permanent key of the web client as well as the authentication token and the protocol version are transferred to the app through a QR code. The app first generates its own permanent key-pair and then connects to the same WebSocket path as responder.

Both peers conduct the server handshake according to the SaltyRTC protocol.

2. SaltyRTC Client Handshake and Signaling

As soon as both peers have successfully finished the server handshake, they can start exchanging client handshake messages. Because the server has established a trusted connection with both the initiator and the responder, it can relay encrypted messages without knowing the content. This is how the session keys and the signaling data are transferred.

The two peers each generate a random session key-pair independent of the permanent key-pair. The public session keys are exchanged. Additionally, the peers exchange information on how they intend to communicate after the client handshake (the so-called SaltyRTC Task):

- Android devices request to communicate over [WebRTC data channels](#)
- iOS devices request to communicate over the [signaling channel](#)

The transfer of data between browser and app differs between Android and iOS, due to platform constraints:

2.1 Android

After the client handshake is successful and both peers know each other's public session key, they initiate a WebRTC PeerConnection. The necessary signaling information to build up the connection (like offer, answer and ICE candidate messages) are exchanged via the WebSocket connection, encrypted with the session keys.

As soon as the WebRTC PeerConnection is established, the two peers initiate the handover of the signaling channel. Future signaling messages (like additional ICE candidates when the network configuration changes) are exchanged over a secure signaling data channel. Once the handover is complete, the WebSocket connection to the server is closed.

The two peers now open a second WebRTC data channel in order to exchange application data. All packets transferred through that data channel are encrypted using the session keys.

2.2 iOS

After the client handshake is successful and both peers know each other's public session key, they can immediately start exchanging application data over the established WebSocket signaling connection. All packets transferred through that channel are encrypted using the session keys.

3. Data Channel (Android only)

All required web client data (conversations, contacts, messages, etc.) can now be freely exchanged through the encrypted WebRTC data channel.

WebRTC Signaling

The direct communication channel between the app and the Android web client is established using a WebRTC PeerConnection. In order to establish such a peer-to-peer connection, a signaling channel is inevitable. Common signaling server implementations often use WebSockets without any end-to-end encryption (often even without transport encryption), meaning that the server can read all (potentially sensitive) network information of the peers connecting. There is also the risk of a server manipulating the data being transmitted, opening up possibilities of MITM attacks.

In order to mitigate this risk as well as minimizing metadata exposure in general, we participated in the design and implementation of the SaltyRTC protocol, which offers end-to-end encryption of signaling data and does not require the clients to trust the server at all.

WebRTC Connection Buildup

Because networks in the real world don't always make it possible to establish a direct connection between two peers due to obstacles like firewalls, NATs, CGNs and the like, WebRTC connection buildup may have to resort to mechanisms like STUN and TURN. STUN (Session Traversal Utilities for NAT) is a protocol that allows to find each other's public IP despite the existence of NATs. TURN (Traversal Using Relays around NAT) is a protocol that provides relaying of data packets in case the connection buildup using STUN alone fails. Note, however,

that even though a TURN server relays packets between two peers, it cannot know anything about the content of these packets as they are both end-to-end encrypted by WebRTC using DTLS and end-to-end encrypted by SaltyRTC.

Threema runs its own STUN and TURN servers.

Trusted Keys/Stored Sessions

In order to prevent having to scan the QR code each time a connection needs to be established, the public permanent key of the peer can optionally be stored as a trusted key.

To be able to securely store this information in the browser, users must provide a password if they want to be able to restore the session at a later point in time. The public permanent key of the app as well as the private permanent key of the web client are then encrypted with the provided user password using authenticated NaCl secret key encryption (XSalsa20 + Poly1305) and stored in local browser storage.

On the app side, the public permanent key of the web client and the private permanent key of the app are stored in the encrypted app database.

When reconnecting to an existing session, instead of creating a new permanent key-pair, the peers restore the trusted key-pair before initiating the server and client handshake. Note, however, that new session keys are still created, thus offering perfect forward secrecy on a session level.

Push Service

On Android devices, if Google Play Services are installed on the device of the user, the FCM push token (an opaque string provided by Google) is transferred to the web client together with the initial data.

On iOS devices, the APNs push token (an opaque string provided by Apple) is transferred to the web client together with the initial data.

If the user requests to persist a session, the token is then stored in the local browser storage together with the trusted keys, encrypted with a key derived from the user-provided password.

When reconnecting a web client session (e.g. when the user requests to restore a previous session, or when automatically reconnecting due to connection loss) and if a push token is available, the browser sends this token via HTTPS to a push relay server provided by Threema, together with the SHA256 hash of the public permanent key of the initiator. The push relay server then sends the hash to the app as a FCM/APNs push notification. When the app receives such a notification, it first checks whether the web client is enabled. If it is enabled and if a web client session with the specified public key exists, that session is started and connects to the SaltyRTC server for the handshake and reconnection procedure.

Open Source

All components of the web client (except the code in the Threema app itself) are published under open source licenses:

Web client: <https://github.com/threema-ch/threema-web>

SaltyRTC server: <https://github.com/saltyrtc/saltyrtc-server-python>

Push relay: <https://github.com/threema-ch/push-relay>

The URL and public permanent key of the SaltyRTC server to be used are transferred through the QR code and stored together with the session information. This allows the user to entirely bypass the web server and SaltyRTC server operated by Threema, if desired. Instructions on how to set up a self-hosted Threema Web instance can be found in the [Threema Web repository](#).

The push relay cannot practically be run by an end user, since the FCM API key and the APNs certificates are not public, but the source code is provided on GitHub for review purposes.

Threema Calls

Threema Calls are based on [WebRTC](#), an open [IETF standard](#). WebRTC uses the ICE, STUN and TURN protocols to establish a secure peer-to-peer connection.

Signaling

Signaling data is sent through Threema messages, thus providing the same end-to-end security and trust level as a regular Threema conversation. If you have verified your contact, you can also trust your phone calls with that contact.

Call Encryption

The audio stream is encrypted with the SRTP protocol, with DTLS-SRTP being used for the key exchange. The certificates used for the DTLS session are cryptographically linked to the keys used for Threema's end-to-end encryption by means of including the certificate fingerprints in signaling messages. DTLS version 1.2 is enforced.

The DTLS ciphersuites offered by the Threema app are (in that order):

- TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256 (0xccca9)
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 (0xc02c)
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 (0xc02b)

The SRTP ciphersuites offered by the Threema app are (in that order):

- SRTP_AEAD_AES_256_GCM (0x0008)
- SRTP_AEAD_AES_128_GCM (0x0007)
- SRTP_AES128_CM_HMAC_SHA1_80 (0x0001)

RTP header extensions are only used if offered in an encrypted form.

Audio Encoding

Audio data is encoded using the state-of-the-art [Opus codec](#) at 48 kHz. To avoid potential privacy and security issues associated with variable bitrate audio codecs³, Threema Calls always enforce constant bitrate to encode the audio stream.

Video Encoding

Video data is encoded by VP9, VP8 or H.264.

Privacy / IP exposure

WebRTC connections can be established in two ways: Either a direct connection between the two devices is used (using STUN to help with NAT traversal), or one or both sides use a TURN server to relay the encrypted packets.

Relaying has the advantage of hiding the public IP address from the call partner, thus providing added privacy when the peer is not trusted. On the other hand it exposes some metadata towards the TURN server, and the connection quality (especially latency) may be worse than in a direct call. By default, Threema enforces relaying through TURN for all unverified contacts (with a single red dot as verification status). Optionally, users can enforce relaying for all calls through the "Always relay calls" option in the "Privacy" settings.

3 Schneier on Security: Eavesdropping on Encrypted Compressed Voice (https://www.schneier.com/blog/archives/2008/06/eavesdropping_o_2.html)

Threema Safe

Overview

Threema Safe is an encrypted server-based backup solution that encompasses the following user-specific data:

- Threema ID
- Private key
- Profile information
 - Nickname
 - Profile photo
 - Linked email address and phone number
- Contact list
 - Contact Threema ID
 - Public key
 - Name
 - Verification level
- Group definitions
 - Group ID
 - Creator
 - Group name
 - List of members
- Distribution lists
 - Name
 - List of members
- App settings

Messages and media data are not part of a Threema Safe backup.

The rationale behind Threema Safe is to give users a secure backup option that allows them to restore their Threema ID and important related data, or to move it to another device, using nothing but the knowledge of their Threema ID and the Threema Safe password that they have chosen.

The use of Threema Safe is optional; the user can choose to enable or disable it at any time. By default, encrypted backups are stored on Threema servers, but the user can use their own server as a backup store instead. A Threema Safe backup server cannot tell which backup belongs to which Threema user by looking at the uploaded data.

Backup format

The Threema app compiles the information mentioned in the previous section as a JSON document (binary values like the private key or profile photo are encoded in Base64). This JSON document is then compressed with gzip and encrypted as follows. The resulting encrypted backup typically has a size of only a few dozen kilobytes (depending mostly on the profile photo size, number of contacts and groups).

Encryption

In order to set up Threema Safe, the user must choose a password consisting of at least 8 characters. The app checks whether the password is on a list of the most frequently used passwords, and warns the user if that is the case.

From the password and the Threema ID of the user, a *Threema Safe Master Key* is derived as follows:

```
threemaSafeMasterKey = scrypt(P = Password, S = Threema ID, N = 65536,  
    r = 8, p = 1, dkLen = 64)
```

The `scrypt` parameters have been chosen to provide a reasonable trade-off between calculation time, memory usage and security without precluding the use on older/lower-end devices.

The `scrypt` output is then split into two parts of 32 bytes each as follows:

- `threemaSafeBackupId = threemaSafeMasterKey[0..31]`
- `threemaSafeEncryptionKey = threemaSafeMasterKey[32..63]`

The Threema Safe backup data in the form of a UTF-8 encoded JSON string (`plaintext`) is then encrypted as follows:

1. Compress using gzip: `gzippedPlaintext = gzip(plaintext)`
2. Generate random 24 byte nonce: `nonce = randombytes(24)`
3. `threemaSafeEncryptedBackup = nonce ||
 crypto_secretbox(gzippedPlaintext, threemaSafeEncryptionKey, nonce)
 || = concatenation`

The `crypto_secretbox` operation is described in the [NaCl documentation](#).

Upload/Storage

The encrypted data (`threemaSafeEncryptedBackup`) is uploaded to, or downloaded from, the Threema Safe server using simple WebDAV-compatible HTTPS requests (see below for details). The `threemaSafeBackupId` in lowercase hex encoding is used as a filename. If the user disables Threema Safe, a DELETE request is issued by the app.

Backup intervals

When Threema Safe is enabled and a backup password is set, the Threema app generates a backup once per day when the app is in use. If the backup data has changed since the last upload, or if more than half of the server-specified backup retention time has passed, then the backup is uploaded to the server, replacing any existing backup with the same ID that may already exist on the server.

Restore/Decryption

1. User enters their Threema ID and password in the setup screen.
2. Derive the master key as described above.
3. Fetch `threemaSafeEncryptedBackup` from the server using the `threemaSafeBackupId`.

4. `nonce = threemaSafeEncryptedBackup[0..23]`
5. `gzippedPlaintext =
 crypto_secretbox_open(threemaSafeEncryptedBackup[24..],
 nonce, threemaSafeEncryptionKey)`

The `crypto_secretbox` operation is described in the [NaCl documentation](#).

6. `plaintext = gunzip(gzippedPlaintext)`

The decrypted and decompressed JSON document (`plaintext`) is then used to restore the Threema ID and other information.

Running a custom Threema Safe server

When setting up Threema Safe, a custom backup server URL may be specified by the user. This URL must start with `https://` and may include username/password information (e.g. `https://username:password@my.server.example/path`). The server must use a TLS certificate trusted by the operating system.

The server must respond to the following requests (paths relative to the backup URL):

Request Server Configuration

```
GET config
Accept: application/json
```

Response Codes

Code	Description
200 OK	Response body contains the configuration, see example below.
400 Bad Request	Validation failed.

Request Server Configuration

```
{
  "maxBackupBytes": 524288,
  "retentionDays": 180
}
```

The app uses `maxBackupBytes` to determine whether a backup is too big to be uploaded (in which case the user is alerted). `retentionDays` is used to determine backup expiration: if `retentionDays/2` days have elapsed since the last backup, the backup is uploaded again even if the backup data has not changed, to keep the backup from expiring.

Create/Update Backup

```
PUT backups/<threemaSafeBackupId>
Content-Type: application/octet-stream
```

The `threemaSafeEncryptedBackup` is sent along unmodified in the PUT body.

Validation

- The `threemaSafeBackupId` MUST be a 64 character lowercase hex string
- The `Content-Type` header MUST be set to `application/octet-stream`

Response Codes

Code	Description
200 OK	Response body contains the configuration, see example below.
400 Bad Request	Validation failed.
404 Not Found	No backup was found at the specified URL.
429 Too Many Requests	Rate limit reached.

Delete Backup

```
DELETE backups/<threemaSafeBackupId>
```

Validation

- The `threemaSafeBackupId` MUST be a 64 character lowercase hex string

Response Codes

Code	Description
200 OK	Response body contains the configuration, see example below.
400 Bad Request	Validation failed.
404 Not Found	No backup was found at the specified URL.
429 Too Many Requests	Rate limit reached.

Security Recommendations

The following recommendations are not required for automated backups to work, but they improve security. Implementing these recommendations is especially important when offering a public service.

Throttling

To avoid misuse, throttling/rate limiting SHOULD be enforced in the server. It is up to the implementer whether to limit requests based on the peer's IP address, or the requested `threemaSafeBackupId`. If the limit has been reached, requests MUST be responded to with a HTTP 429 (Too Many Requests) status code.

If the server software itself does not support rate limiting, this can be implemented by using a Proxy like [Nginx](#) or [HAProxy](#).

User Agent Validation

All requests from a real Threema app contain a user agent header containing the string "Threema". If the user agent does not contain that string, HTTP 400 SHOULD be returned by the server.

Data Retention

Backups SHOULD be automatically deleted after a certain amount of time, as specified in the "retentionDays" configuration key.