



Messenger Secret Conversations

Technical Whitepaper

Version 2.0
May 18, 2017

Version History:

2.0 — May 18, 2017

1.0 — Jul 8, 2016

Contents

Introduction 4

Transport Protocol Overview 5

Abuse Reporting 11

Message Storage 13

Conclusion 15

Introduction

In this document we provide a brief technical overview of SECRET CONVERSATIONS — a specialized conversation mode within Facebook Messenger. SECRET CONVERSATIONS provides *end-to-end encryption for messages* using keys that are only available on users' devices.

SECRET CONVERSATIONS is a distinct conversation mode inside Facebook Messenger. Individual secret conversations are displayed as separate threads in Messenger and share many UX elements with regular Messenger conversations. However, SECRET CONVERSATIONS uses a different transport protocol, specialised on-device storage and separate back-end infrastructure.

The SECRET CONVERSATIONS threat model considers the compromise of server and networking infrastructure used by Messenger — Facebook's included. Attempts to obtain message plaintext or falsify messages by Facebook or network providers result in explicit warnings to the user. We assume however that clients are working as designed, e.g. that they are not infected with malware.

SECRET CONVERSATIONS relies upon the Signal Protocol. Messenger uses Signal Protocol's implementation as available in the open-source libsignal-protocol-java and libsignal-protocol-c libraries for Android and iOS respectively. SECRET CONVERSATIONS also incorporates abuse-reporting features which are not present in other platforms which use the Signal Protocol.

In this document we describe SECRET CONVERSATIONS, starting with the transport protocol. We then cover abuse reporting and close with how we handle on-device storage.

Transport Protocol Overview

SECRET CONVERSATIONS is a *device-to-device* conversation mode. Messages are only accessible from the devices which participate in a conversation when the conversation occurs. This differs from regular Messenger conversations which are stored server-side and are accessible from any device connected to the participating Facebook accounts; including present and future devices and browser instances. Instead, to use SECRET CONVERSATIONS users enable some subset of devices, such as their mobile phone and a tablet, upon which only their future secret conversations will be available.

Device Management

A user can enable or disable a device for SECRET CONVERSATIONS at any point, and can disable other devices remotely from any currently enabled device. SECRET CONVERSATIONS may also be enabled or disabled automatically. When a user logs into a compatible version of Messenger, the new device gets enabled for use with SECRET CONVERSATIONS. Facebook disables devices after they have been offline for a period of at least 30 days. Upon enabling a new device existing SECRET CONVERSATIONS messages and cryptographic keys are not transferred to the new device.

Whenever the list of devices on which SECRET CONVERSATIONS is enabled changes for an account, Messenger displays a warning that this occurred. When a user's own set of enabled devices changes, all of their other devices enabled for SECRET CONVERSATIONS proactively receive an update. People with pre-existing secret conversations with the user who changed their list of enabled devices receive the same warning when they return to these pre-existing conversations.

Facebook bounces messages sent to an incorrect list of participants, and includes the correct device list within that bounce. Messenger does not automatically resend bounced messages when new devices are added to a secret conversation – an explicit resend action from the user is required. Messenger, however, may automatically

resend messages without user interaction in case of device removal, as no device will receive the message that the sender has not been notified about.

Keys

Each device manages various cryptographic keys. All keys are generated or derived on-device. Private keys are never sent to Facebook.

Public keys All public key operations use Curve25519. Each device uses the following public-secret keypairs:

- The *Identity Key* keypair (IK_{pk}, IK_{sk}). This is a long-term keypair which is generated the first time Messenger runs.
- The *Signed Pre-Key* keypair (SPK_{pk}, SPK_{sk}). This is a medium-term keypair which is rotated periodically. It is signed by IK_{sk} .
- The *One-Time Pre-Key* keypairs ($OTPK_{pk}, OTPK_{sk}$). These keypairs are generated in batches by clients. They facilitate asynchronous conversation initiation.¹
- The *Ephemeral Key* keypairs (EK_{pk}, EK_{sk}). A new ephemeral keypair is generated for each round of communication within a secret conversation and is subsequently discarded.

¹ The client also generates a single *Last-Resort Pre-Key*. This is used like a *One-Time Pre-Key*, but is simply provided when the server has no *One-Time Pre-Keys* available for a given device.

Pairwise Session keys When starting a pairwise cryptographic channel the participating devices derive symmetric session keys. These are:

- The *Root Key* (RK) is a 256-bit key which is used to derive Chain Keys in the Signal Protocol ratchets.
- *Chain Keys* (CK) are each 256-bit values which are used to derive Message Keys.
- *Message Keys* (MK) are each 640-bit values which consist of 256 bits for an AES-256 key, 256 bits for an HMAC-SHA256 key, and 128 bits for an Initialization Vector (IV) for AES-CBC encryption.

Multicast Session keys In SECRET CONVERSATIONS involving more than 2 devices each device uses session keys for sending messages. These are:

- *Sender Chain Keys* (SCK) are 256-bit values used to derive Sender Message Keys.

- *Sender Message Keys* (SMK) are 384-bit values consisting of 256 bits for an AES-256 key, and 128 bits for an Initialization Vector (IV) for AES-CBC encryption.
- *Sender Signing Keys* (SSK) are key pairs used to sign multicast messages.

When Messenger initialises SECRET CONVERSATIONS it generates and then uploads to Facebook its permanent IK_{pk} and the current SPK_{pk} . It generates a batch of one-time pre-key keypairs and uploads their public parts to Facebook on demand.

Pairwise Channel Initiation

Each device in a secret conversation must have a pairwise channel with each other device before it can send messages. Each pairwise channel consists of two devices: one *Initiator* device and one *Responder* device (*I* and *R* respectively). Let HKDF be a secure hash-based key derivation function, and ECDH indicate the elliptic curve Diffie-Hellman function applied to a secret and public key. To create a new pairwise channel:

1. The Initiator obtains from Facebook IK_{pk}^R , SPK_{pk}^R and $OTPK_{pk}^R$ for an one-time pre-key keypair generated by the Responder device.
2. Facebook deletes $OTPK_{pk}^R$ from the Responder's list of available one-time pre-keys, so long as it is not the last-resort.
3. The Initiator generates a fresh ephemeral keypair (EK_{pk}^I, EK_{sk}^I) .
4. The Initiator now computes the first *root key* RK as follows:

$$\begin{aligned} a &= \text{ECDH}(IK_{sk}^I, SPK_{pk}^R), & b &= \text{ECDH}(EK_{sk}^I, IK_{pk}^R) \\ c &= \text{ECDH}(EK_{sk}^I, SPK_{pk}^R), & d &= \text{ECDH}(EK_{sk}^I, OTPK_{pk}^R) \\ \text{RK} &= \text{HKDF}(a||b||c||d) \end{aligned}$$

Using the RK the Initiator can calculate the first CK and MK (as described next) and use those to start sending messages.

5. The Initiator sends the first encrypted message to the Responder, including the fresh EK_{pk}^I it generated previously.

Upon receiving the first encrypted message the Responder:

1. Recomputes RK as above by performing the four ECDH operations using the other part of the same keypairs available locally.
2. Recomputes the first CK and MK, and decrypts the message.
3. Deletes $(OTPK_{pk}^R, OTPK_{sk}^R)$ from its local storage.

Pairwise Message Exchange

Each pairwise message is encrypted with AES·CBC and authenticated using HMAC·SHA256. The unique MK is derived from the current CK and RK. Their first value is:

$$\text{CK} = \text{RK}$$

$$\text{MK} = \text{HKDF}(\text{CK})$$

In each message exchange, the sender generates a fresh ephemeral keypair $(\text{EK}_{pk}^{\text{sender}}, \text{EK}_{sk}^{\text{sender}})$ and includes the public part in the outgoing message. The recipient calculates the current MK value using $\text{EK}_{pk}^{\text{sender}}$ and can decrypt the message. It too then generates a fresh ephemeral keypair $(\text{EK}_{pk}^{\text{receiver}}, \text{EK}_{sk}^{\text{receiver}})$ and derives new keys RK' , CK' and MK' for use with the next response by updating the previous symmetric key values as follows:

$$\text{RK}', \text{CK}' = \text{HKDF}(\text{ECDH}(\text{EK}_{sk}^{\text{receiver}}, \text{EK}_{pk}^{\text{sender}}))$$

$$\text{MK}' = \text{HKDF}(\text{CK}')$$

If a second message is sent before the opposing party responds, the sender uses a new chain key $\text{CK}'' = \text{HKDF}(\text{CK})$ and a corresponding $\text{MK}'' = \text{HKDF}(\text{CK}'')$.

When a secret conversation contains just two devices, we use the pairwise channel for transmitting messages directly. The Signal Protocol's implementation is open-source and available at <https://github.com/whispersystems/libsignal-protocol-java/>.

Multi-device Message Exchange

Secret conversations with more than two devices use the Signal Protocol's Group Messaging protocol.

All messages have a *Sender* and at least two *Receivers*. Prior to sending to a multi-device conversation *Sender* must have a pairwise channel with every *Receiver*. *Sender* also generates a SCK and SSK, and sends them along its pairwise channel to each *Receiver*.

Each multi-device message is encrypted with AES·CBC and signed by the SSK. The unique SMK is derived from the SCK similarly to the pairwise channel:

$$\text{SMK} = \text{HKDF}(\text{SCK})$$

For subsequent messages the *Sender* ratchets the SCK using HKDF

$$\text{SCK}' = \text{HKDF}(\text{SCK})$$

and uses each subsequent SCK' key to derive SMK and SSK. This approach provides forward secrecy, as earlier SCKs and SMKs cannot be calculated from the current state.

When a new device is added to a secret conversation, each other device will send the new device their latest SCK upon sending their next message. When a device is removed however, each device will generate a new SCK. This ensures that the necessary key material is restricted to the devices within the thread.

Attachments

If a message includes attachments, such as images or videos, they are encrypted and uploaded to Facebook. For each attachment, the sender:

1. Generates a pseudorandom 256-bit AES key K and a 96-bit pseudorandom Initialization Vector IV .
2. Encrypts the attachment using AES-GCM² encryption using K and IV . The IV is placed into the header (Associated Data) of the GCM encryption. The IV and the GCM authentication tag are attached to the resulting ciphertext to form the encrypted file.
3. Computes a SHA256 hash H of the resulting encrypted file and uploads the file to Facebook. The sender obtains a unique identifier for future retrieval.
4. Encodes the unique identifier, K , H , attachment metadata, and an optional thumbnail into a message. The message is encrypted and sent to the recipient who downloads the content, verifies the hash H , and decrypts the file using K .

² Messenger uses the AES-GCM implementation provided by OpenSSL. On Android, Messenger employs Conceal. Conceal is open-source and available at <https://facebook.github.io/conceal/>

Stickers

Stickers are images chosen from a set of collections hosted by Facebook. Each sticker is referenced by a unique identifier. The sender chooses a sticker to attach in a conversation from a set of sticker previews available in the message composer. The sender then sends the corresponding sticker identifier as an encrypted message to the recipient. Unless the sticker file is available locally, both the sender and the receiver submit the sticker identifier to Facebook and download the corresponding sticker file. Facebook may infer the use of individual stickers when they are first used on a device but devices cache sticker files and avoid repeat fetches if a sticker is later reused.

Conversation Metadata

During a secret conversation, devices generate metadata such as delivery and read receipts to indicate successful message transmission and display, respectively. Conversation metadata such as delivery and read receipts do not contain message plaintext and are not end-to-end encrypted.

Key Verification

For every secret conversation Messenger exposes in its interface the identity keys (i.e. IK_{pk}) for every participating device. Users may optionally verify these keys in order to ensure no man-in-the-middle attack is compromising their secret conversations. Messenger displays the 256-bit IK_{pk} values in hexadecimal format.

Abuse Reporting

A participant in a secret conversation may voluntarily notify Facebook of abusive content. Facebook uses such reports to identify users who violate Facebook’s terms of service.

The ability to report abuse does not represent a relaxation of the end-to-end encryption guarantees of SECRET CONVERSATIONS. Facebook will never have access to plaintext messages *unless one participant in a secret conversation voluntarily reports the conversation*.

SECRET CONVERSATIONS includes a mechanism to perform “franking” of messages sent through Facebook. This mechanism is analogous to placing a cryptographic stamp on the message, without learning the message content.

Franking

The franking mechanism must satisfy three main guarantees: *authenticity*, *confidentiality* and *third-party deniability*. The authenticity property ensures that if a user submits a report then the message must have legitimately originated from the sender’s device. The confidentiality property ensures that no outside party — including Facebook — should learn the content of a message unless a participant in a secret conversation voluntarily shares that information. Finally, the third-party deniability property ensures that no party outside of Facebook can cryptographically determine the validity of a report.

Franking tag Authenticity for messages in a secret conversation is provided by the *Franking tag* T_F . Senders must send the Franking tag along with each encrypted message. To compute T_F , the sender first generates a 256-bit random nonce N_F . N_F is added to the unencrypted message being transmitted. Next, the entire data structure is serialized into a string M , and T_F is computed as:

$$T_F = \text{HMAC}\cdot\text{SHA256}(N_F, M)$$

N_F remains within the serialised, encrypted data sent to the recipient. The sender destroys any other copies of N_F after transmission. T_F is

transmitted to Facebook along with each encrypted message.

Franking messages When Facebook receives T_F , it uses a Facebook key K_F to compute the *Reporting* tag R_F over T_F and conversation context (e.g., sender and recipient identifiers, timestamp) as:

$$R_F = \text{HMAC}\cdot\text{SHA256}(K_F, T_F \parallel \text{context})$$

Both T_F and R_F are sent to the recipient along with the encrypted message. The recipient decrypts the ciphertext, parses the resulting plaintext to obtain N_F , and verifies the structure of T_F prior to displaying the message. If T_F is not verified then the recipient discards the message without displaying it. The recipient stores the message M , N_F , T_F , R_F and context in its local storage.

Reporting abuse To report abuse, the recipient of a message submits to Facebook the full serialized message plaintext, R_F , N_F , and context. Upon receiving this message Facebook first recomputes T_F and then validates R_F using the provided information as well as its internal key K_F .

Security and Privacy The authenticity properties of the franking mechanism are based on reasonable assumptions about the collision-resistance of the SHA256 hash function and the unforgeability of HMAC·SHA256.

Authenticity In order to “forge” invalid content M' , a user must either (a) produce a forged HMAC tag under Facebook’s key K_F , or (b), identify a collision $N_F', M', \text{context}'$ such that the HMAC of these values is equal to the HMAC of a different valid message M sent through Facebook.

Confidentiality Similarly, under reasonable assumptions about HMAC·SHA256, the resulting tag reveals no information about the message to Facebook or to eavesdroppers.

Third-party deniability The guarantee holds under the assumption that HMAC·SHA256 is a pseudorandom function and that K_F is never publicly revealed. Facebook rotates K_F periodically.

Message Storage

SECRET CONVERSATIONS plaintext messages are stored permanently only on the devices that participate in each conversation. Plaintext messages are protected using on-device symmetric-key encryption and optional DISAPPEARING MESSAGES functionality.

On-device encryption ensures that messages stored permanently on a particular device are only accessible while a user is authenticated to Facebook. Messenger allows users to switch accounts. While a second user is logged in to a particular device messages of the first user should not be accessible. However, when the first user returns to the same device they should find their messages intact.

To fulfill these requirements, clients employ two encryption keys: K_{local} and K_{remote} . Both these keys are used for AES-GCM encryption. K_{local} is generated on-device and never leaves the device it was generated on. It is used to encrypt plaintext messages before these are stored permanently on a device. K_{remote} is a long-term, user-specific key held on Facebook and delivered to the device when a user authenticates. It is used to encrypt K_{local} in local storage. When Messenger switches accounts the device persists an encrypted version of K_{local} and erases K_{remote} . Upon successful reauthentication the device obtains K_{remote} from Facebook, uses it to decrypt K_{local} and gains access to messages.

On iOS, database files are protected using the filesystem protection level `FileProtectionComplete`. This ensures that database files can no longer be accessed shortly after the user locks their device. The key used for local storage, i.e. K_{local} , is stored in keychain using the `kSecAttrAccessibleWhenUnlockedThisDeviceOnly` attribute.

DISAPPEARING MESSAGES ensures that messages are no longer visible within a selected time after they are sent or received. In DISAPPEARING MESSAGES a *timeout* value is added to the message data structure before serialisation and encryption. Both devices automatically hide messages that specify a timeout once the message timeout has elapsed. The actual deletion of message plaintext from local storage occurs shortly after each message has expired in order to enable abuse reporting in the interim.

Video Attachments on iOS

As with other attachment types on iOS, video attachments are stored on-device encrypted. To play videos however, Messenger's video player on iOS must access them via a URL. As memory cannot be addressed by URL, this prevents videos from being playable directly from an in-memory decrypt. Messenger decrypts videos and stores them in a session-based least-recently-used filesystem cache when they need to be viewed. This cache is cleared when a user logs out. When a secret conversation is deleted, any associated decrypted video attachments in the cache are deleted.

Conclusion

SECRET CONVERSATIONS is a specialised conversation mode in Messenger. Messages in SECRET CONVERSATIONS are encrypted end-to-end between the sender and the recipient using the Signal Protocol and its open-source implementations. Third parties — Facebook included — do not have access to message plaintext and messages can only be decrypted by their intended recipient. Users may inspect the identity keys used for end-to-end encryption and verify the confidentiality and authenticity of their communications. Decrypted messages do not leave the devices that participate in the conversation. Users retain the ability to report abusive content to Facebook.