



Exposure Notification

Cryptography Specification

Preliminary — Subject to Modification and Extension

April 2020

v1.2

Contents

- Overview3
- External Functions4
- Key Schedule for Exposure Notification.....5
- Positive Diagnosis8
- Value Matching of Positive Users.....9
- Privacy Considerations.....10
- Test Vectors11
- Revision History12

Overview

This document provides the detailed technical specification for cryptographic key scheduling for a new privacy-preserving Bluetooth protocol to support Exposure Notification. Exposure Notification makes it possible to combat the spread of the coronavirus — the pathogen that causes COVID-19 — by alerting participants about possible exposure, through someone they have recently been in contact with who has subsequently been positively diagnosed. This specification complements the Bluetooth specification, which contains further information about the scheduling of broadcasts and the higher-level life cycle of the Bluetooth protocol.

External Functions

Concatenation

The symbol $||$ denotes concatenation.

HKDF

HKDF designates the **HKDF** function as defined by [IETF RFC 5869](#), using the SHA-256 hash function:

Output $\leftarrow HKDF(\text{Key}, \text{Salt}, \text{Info}, \text{OutputLength})$

AES

AES designates the encryption of a single AES-128 block:

Output $\leftarrow AES_{128}(\text{Key}, \text{Data})$

AES-CTR

AES-CTR designates the AES-128 block cipher in Counter Mode:

Ciphertext $\leftarrow AES_{128} - CTR(\text{Key}, \text{IV}, \text{Data})$

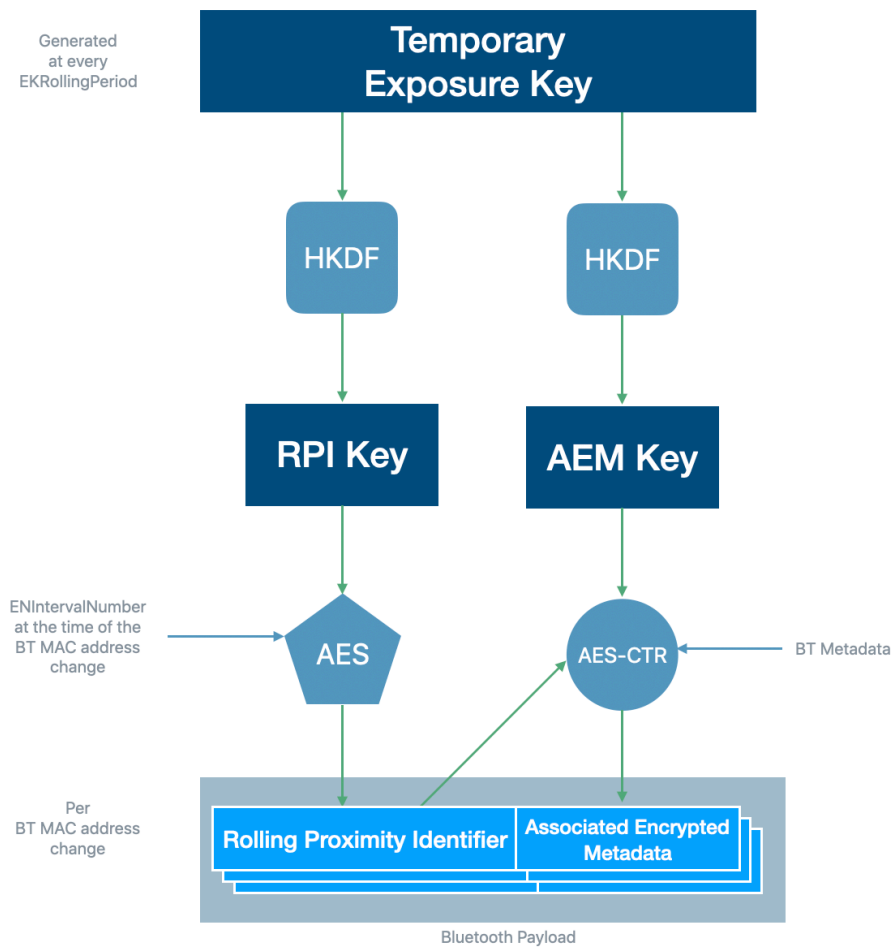
The cipher text output is of the same length as the input data, no padding is applied.

CRNG

The CRNG function designates a cryptographic random number generator:

Output $\leftarrow CRNG(\text{OutputLength})$

Key Schedule for Exposure Notification



To strengthen privacy, this protocol leverages a new concept — Bluetooth pseudorandom identifiers, referred to as *Rolling Proximity Identifiers*. Each Rolling Proximity Identifier is derived from a *Rolling Proximity Identifier Key*, which is in turn derived from a *Temporary Exposure Key* and a discretized representation of time. The Rolling Proximity Identifier changes at the same frequency as the Bluetooth randomized address, to prevent linkability and wireless tracking. Nonuser identifying *Associated Encrypted Metadata* is associated with Rolling Proximity Identifiers. The broadcast metadata from a user can only be decrypted later when the user tests positive.

In this protocol, the time is discretized in 10 minute intervals that are enumerated starting from Unix Epoch Time. `ENIntervalNumber` allows conversion of the current time to a number representing the interval it's in.

Temporary Exposure Keys roll at a frequent cadence called *TEKRollingPeriod*, which is set to 144, achieving a key validity of 24 hours. Each key is randomly and independently generated using a cryptographic random number generator. All devices sharing the same `TEKRollingPeriod` roll keys at the same time — at the beginning of an interval whose `ENIntervalNumber` is a multiple of `TEKRollingPeriod`.

ENIntervalNumber

This function provides a number for each 10 minute time window that's shared between all devices participating in the protocol. These time windows are derived from timestamps in Unix Epoch Time.

$$\text{ENIntervalNumber}(\text{Timestamp}) \leftarrow \frac{\text{Timestamp}}{60 \times 10}$$

ENIntervalNumber is encoded as a 32-bit (`uint32_t`) unsigned little-endian value.

TEKRollingPeriod

The TEKRollingPeriod is the duration for which a Temporary Exposure Key is valid (in multiples of 10 minutes). In our protocol, TEKRollingPeriod is defined as 144, achieving a key validity of 24 hours.

Temporary Exposure Key

When setting up the device for exposure detection, the first Temporary Exposure Key is generated on the device and associated with a ENIntervalNumber i , corresponding to the time from which the key is valid. That value is aligned with the TEKRollingPeriod and is derived as follows:

$$i \leftarrow \left\lfloor \frac{\text{ENIntervalNumber}(\text{Time at Key Generation})}{\text{EKRollingPeriod}} \right\rfloor \times \text{EKRollingPeriod}$$

The devices generate the 16-byte Temporary Exposure Key as follows:

$$tek_i \leftarrow \text{CRNG}(16)$$

The key is securely stored along with i . At the end of every TEKRollingPeriod, a new key is generated.

The total number of Temporary Exposure Keys stored on the device for a 14-day period is given by $\left\lceil \frac{144 \times 14}{\text{EKRollingPeriod}} \right\rceil$. Since Temporary Exposure Keys roll daily (that is, TEKRollingPeriod \leftarrow 144), the device stores 14 keys.

The use of 16-byte keys limits the server and device requirements for transferring and storing Diagnosis Keys while preserving low false-positive probabilities.

Rolling Proximity Identifier Key

The *Rolling Proximity Identifier Key* (RPIK) is derived from the Temporary Exposure Key and is used in order to derive the Rolling Proximity Identifiers.

$$\text{RPIK}_i \leftarrow \text{HKDF}(tek_i, \text{NULL}, \text{UTF8}(\text{"EN-RPIK"}), 16)$$

Rolling Proximity Identifier

Rolling Proximity Identifiers are privacy-preserving identifiers that are broadcast in Bluetooth payloads.

Each time the Bluetooth Low Energy MAC randomized address changes, we derive a new Rolling Proximity Identifier using the Rolling Proximity Identifier Key:

$$RPI_{i,j} \leftarrow AES_{128}(RPIK_i, PaddedData_j)$$

Where:

- j is the Unix Epoch Time at the moment the roll occurs
- $ENIN_j \leftarrow ENIntervalNumber(j)$
- PaddedData is the following sequence of 16 bytes:
 - $PaddedData_j[0...5] = UTF8("EN-RPI")$
 - $PaddedData_j[6...11] = 0x000000000000$
 - $PaddedData_j[12...15] = ENIN_j$

The use of 16-byte identifiers yields a low probability of collisions, and limits the risk of false-positive matches, while keeping device storage requirements low.

Associated Encrypted Metadata Key

The Associated Encrypted Metadata Keys are derived from the Temporary Exposure Keys in order to encrypt additional metadata.

$$AEMK_i \leftarrow HKDF(tek_i, NULL, UTF8("EN-AEMK"), 16)$$

Associated Encrypted Metadata

The Associated Encrypted Metadata is data encrypted along with the Rolling Proximity Identifier, and can only be decrypted later if the user broadcasting it tested positive and reveals their Temporary Exposure Key.

$$\text{Associated Encrypted Metadata}_{i,j} \leftarrow AES_{128} - CTR(AEMK_i, RPI_{i,j}, \text{Metadata})$$

The 16-byte Rolling Proximity Identifier and the appended encrypted metadata are broadcast over Bluetooth Low Energy wireless technology.

Positive Diagnosis

When a user tests positive, a limited set of Temporary Exposure Keys and their respective `ENIntervalNumber i` (describing when their validity started) are uploaded to the *Diagnosis Server*.

This set of Temporary Exposure Keys is limited to the time window in which the user could have been exposing other users (for example, the most recent 14 days). This subset of keys is referred to as *Diagnosis Keys*. If a user remains healthy and never tests positive, their Temporary Exposure Keys don't leave the device.

The Diagnosis Server aggregates the Diagnosis Keys from all users who have tested positive, and distributes them to all the user clients that are participating in exposure notification.

Value Matching of Positive Users

To identify any exposures, each client periodically fetches the list of new Diagnosis Keys from the Diagnosis Server. Because Diagnosis Keys are sets of Temporary Exposure Keys with their associated `ENIntervalNumber` i , each of the clients can again derive the sequence of Rolling Proximity Identifiers that were broadcast over Bluetooth from users who tested positive.

To do so, the clients use each of the Diagnosis Keys with the function defined, to derive the 144 Rolling Proximity Identifiers starting from `ENIntervalNumber` i . The clients match each of the derived identifiers against the sequence they found through Bluetooth scanning. A +/- two-hour tolerance window is allowed between when a Rolling Proximity Identifier derived from the Diagnosis Key was supposed to be broadcast, and the time at which it was scanned.

The Associated Encrypted Metadata does not have to be decrypted until a match occurs. Upon decryption, the data has to be appropriately sanitized and validated as the Associated Encrypted Metadata isn't authenticated.

Privacy Considerations

- The key schedule is fixed and defined by operating system components, preventing applications from including static or predictable information that could be used for tracking.
- A Temporary Exposure Key is required to correlate between a user's Rolling Proximity Identifiers. This reduces the risk of privacy loss from broadcasting the identifiers.
- Without the release of the Temporary Exposure Keys, it's computationally infeasible for an attacker to find a collision on a Rolling Proximity Identifier. This prevents a wide range of replay and impersonation attacks.
- When reporting Diagnosis Keys, the correlation of Rolling Proximity Identifiers by others is limited to 24 hour periods due to the use of Temporary Exposure Keys that change daily. The server must not retain metadata from clients uploading Diagnosis Keys after including those key in the aggregated list of Diagnosis Keys per day.

Test Vectors

Test vectors for interoperability testing between implementations of this specification are available upon request in a machine-readable format.

Revision History

v1.1 - April 23, 2020

- Renamed "Contact Tracing" to "Exposure Notification" throughout the document.
- Temporary Exposure Keys (previously known as Daily Tracing Keys) are now randomly generated and no longer derived.
- AES is now used instead of HMAC<SHA256> for improved performance.
- Encryption of associated metadata is now provided.
- Reformatted the title page and table of contents for consistency across documents.