

Finding Practically-exploitable Cryptographic Vulnerabilities in Matrix

6th Workshop on Attacks in Cryptography, 20th August 2023

Martin Albrecht (King's College London*)

Sofía Celi (Brave Software)

Benjamin Dowling (University of Sheffield)

Dan Jones (Royal Holloway, University of London)



me

martin.albrecht@kcl.ac.uk

cherenkov@riseup.net

b.dowling@sheffield.ac.uk

dan.jones@rhul.ac.uk

* Work completed whilst working for
Royal Holloway, University of London.



The image shows a browser window displaying the Matrix website. The browser's address bar shows 'matrix.org'. The navigation menu includes 'Discover', 'Develop', 'Foundation', 'Blog', 'FAQs', 'Matrix Live', 'Shop', and a 'Try Now' button. The main content area features the Matrix logo, a headline, and a 'Learn More' button. Below this, there are two columns of text: one on the left with the phrase 'Imagine a world...' and one on the right with the heading 'This is Matrix.'.

[matrix]

Discover Develop Foundation Blog FAQs Matrix Live Shop [Try Now](#)

[matrix]

An open network for secure, decentralized communication

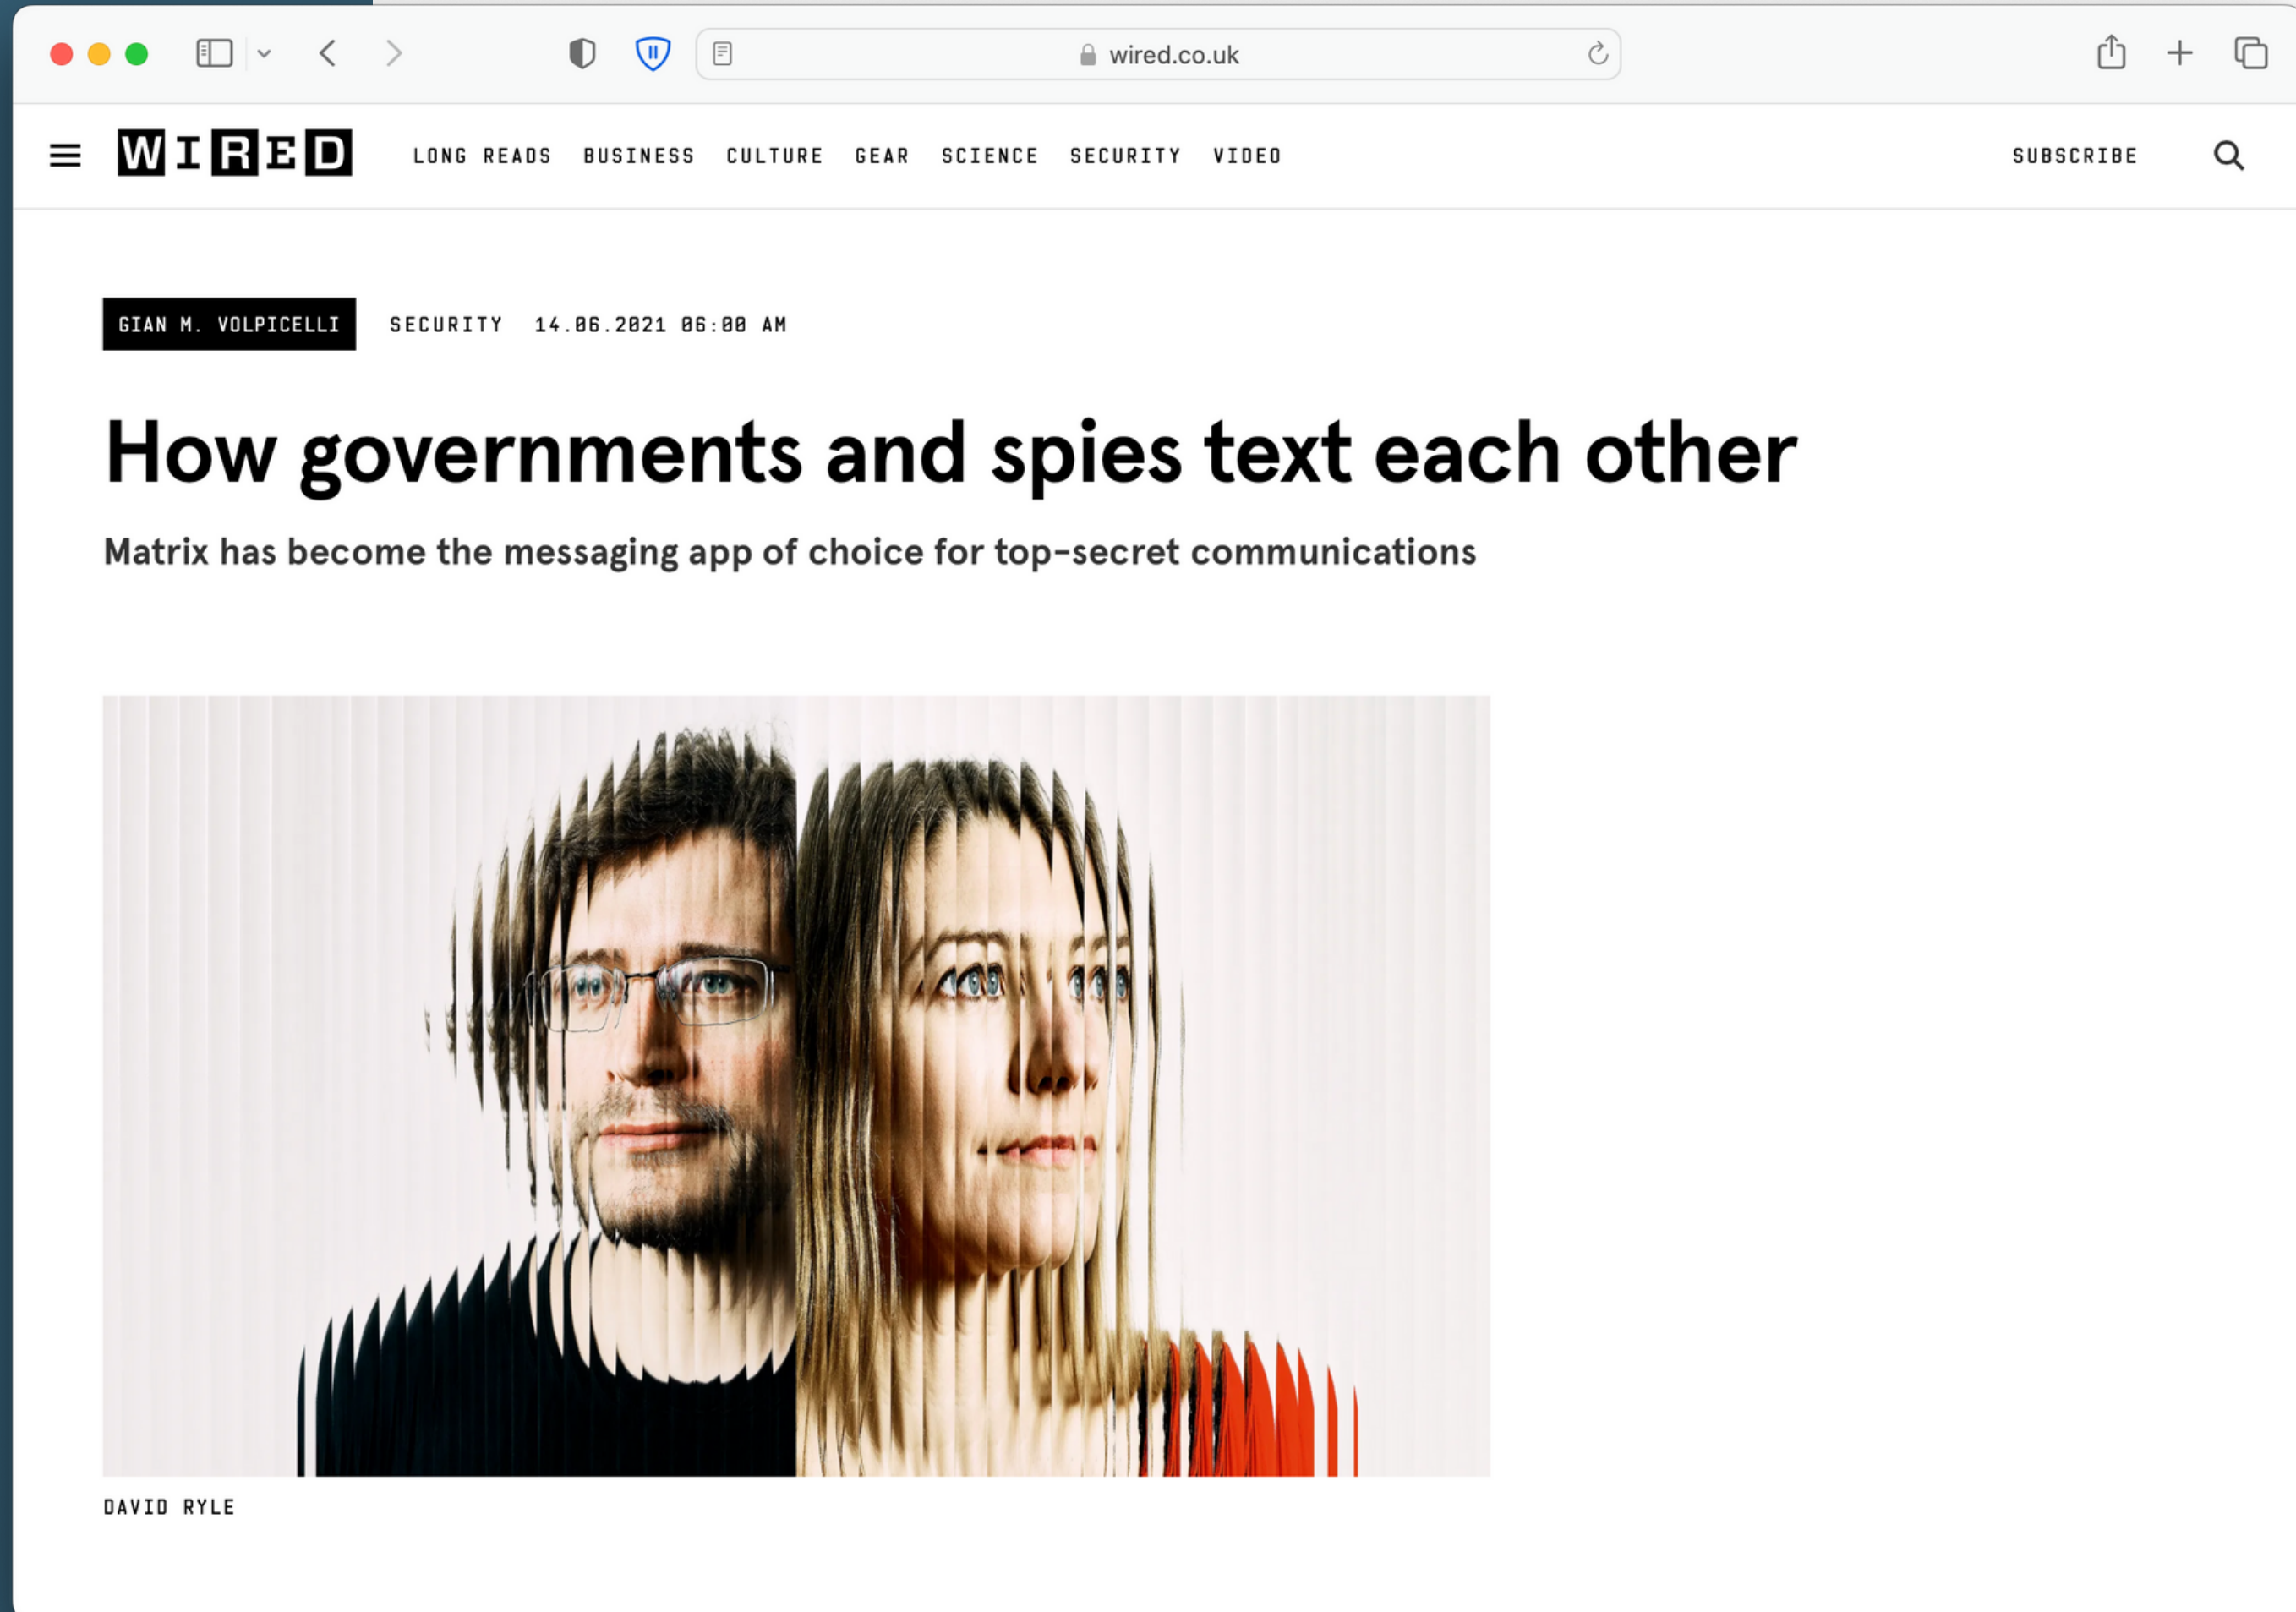
[Learn More](#)

Imagine a world...

...where it is as simple to message or call anyone as it is to send them an email.

This is Matrix.

Matrix is an open source project that publishes the [Matrix open standard](#) for secure, decentralized, real-time




wired.co.uk

W I R E D LONG READS BUSINESS CULTURE GEAR SCIENCE SECURITY VIDEO SUBSCRIBE

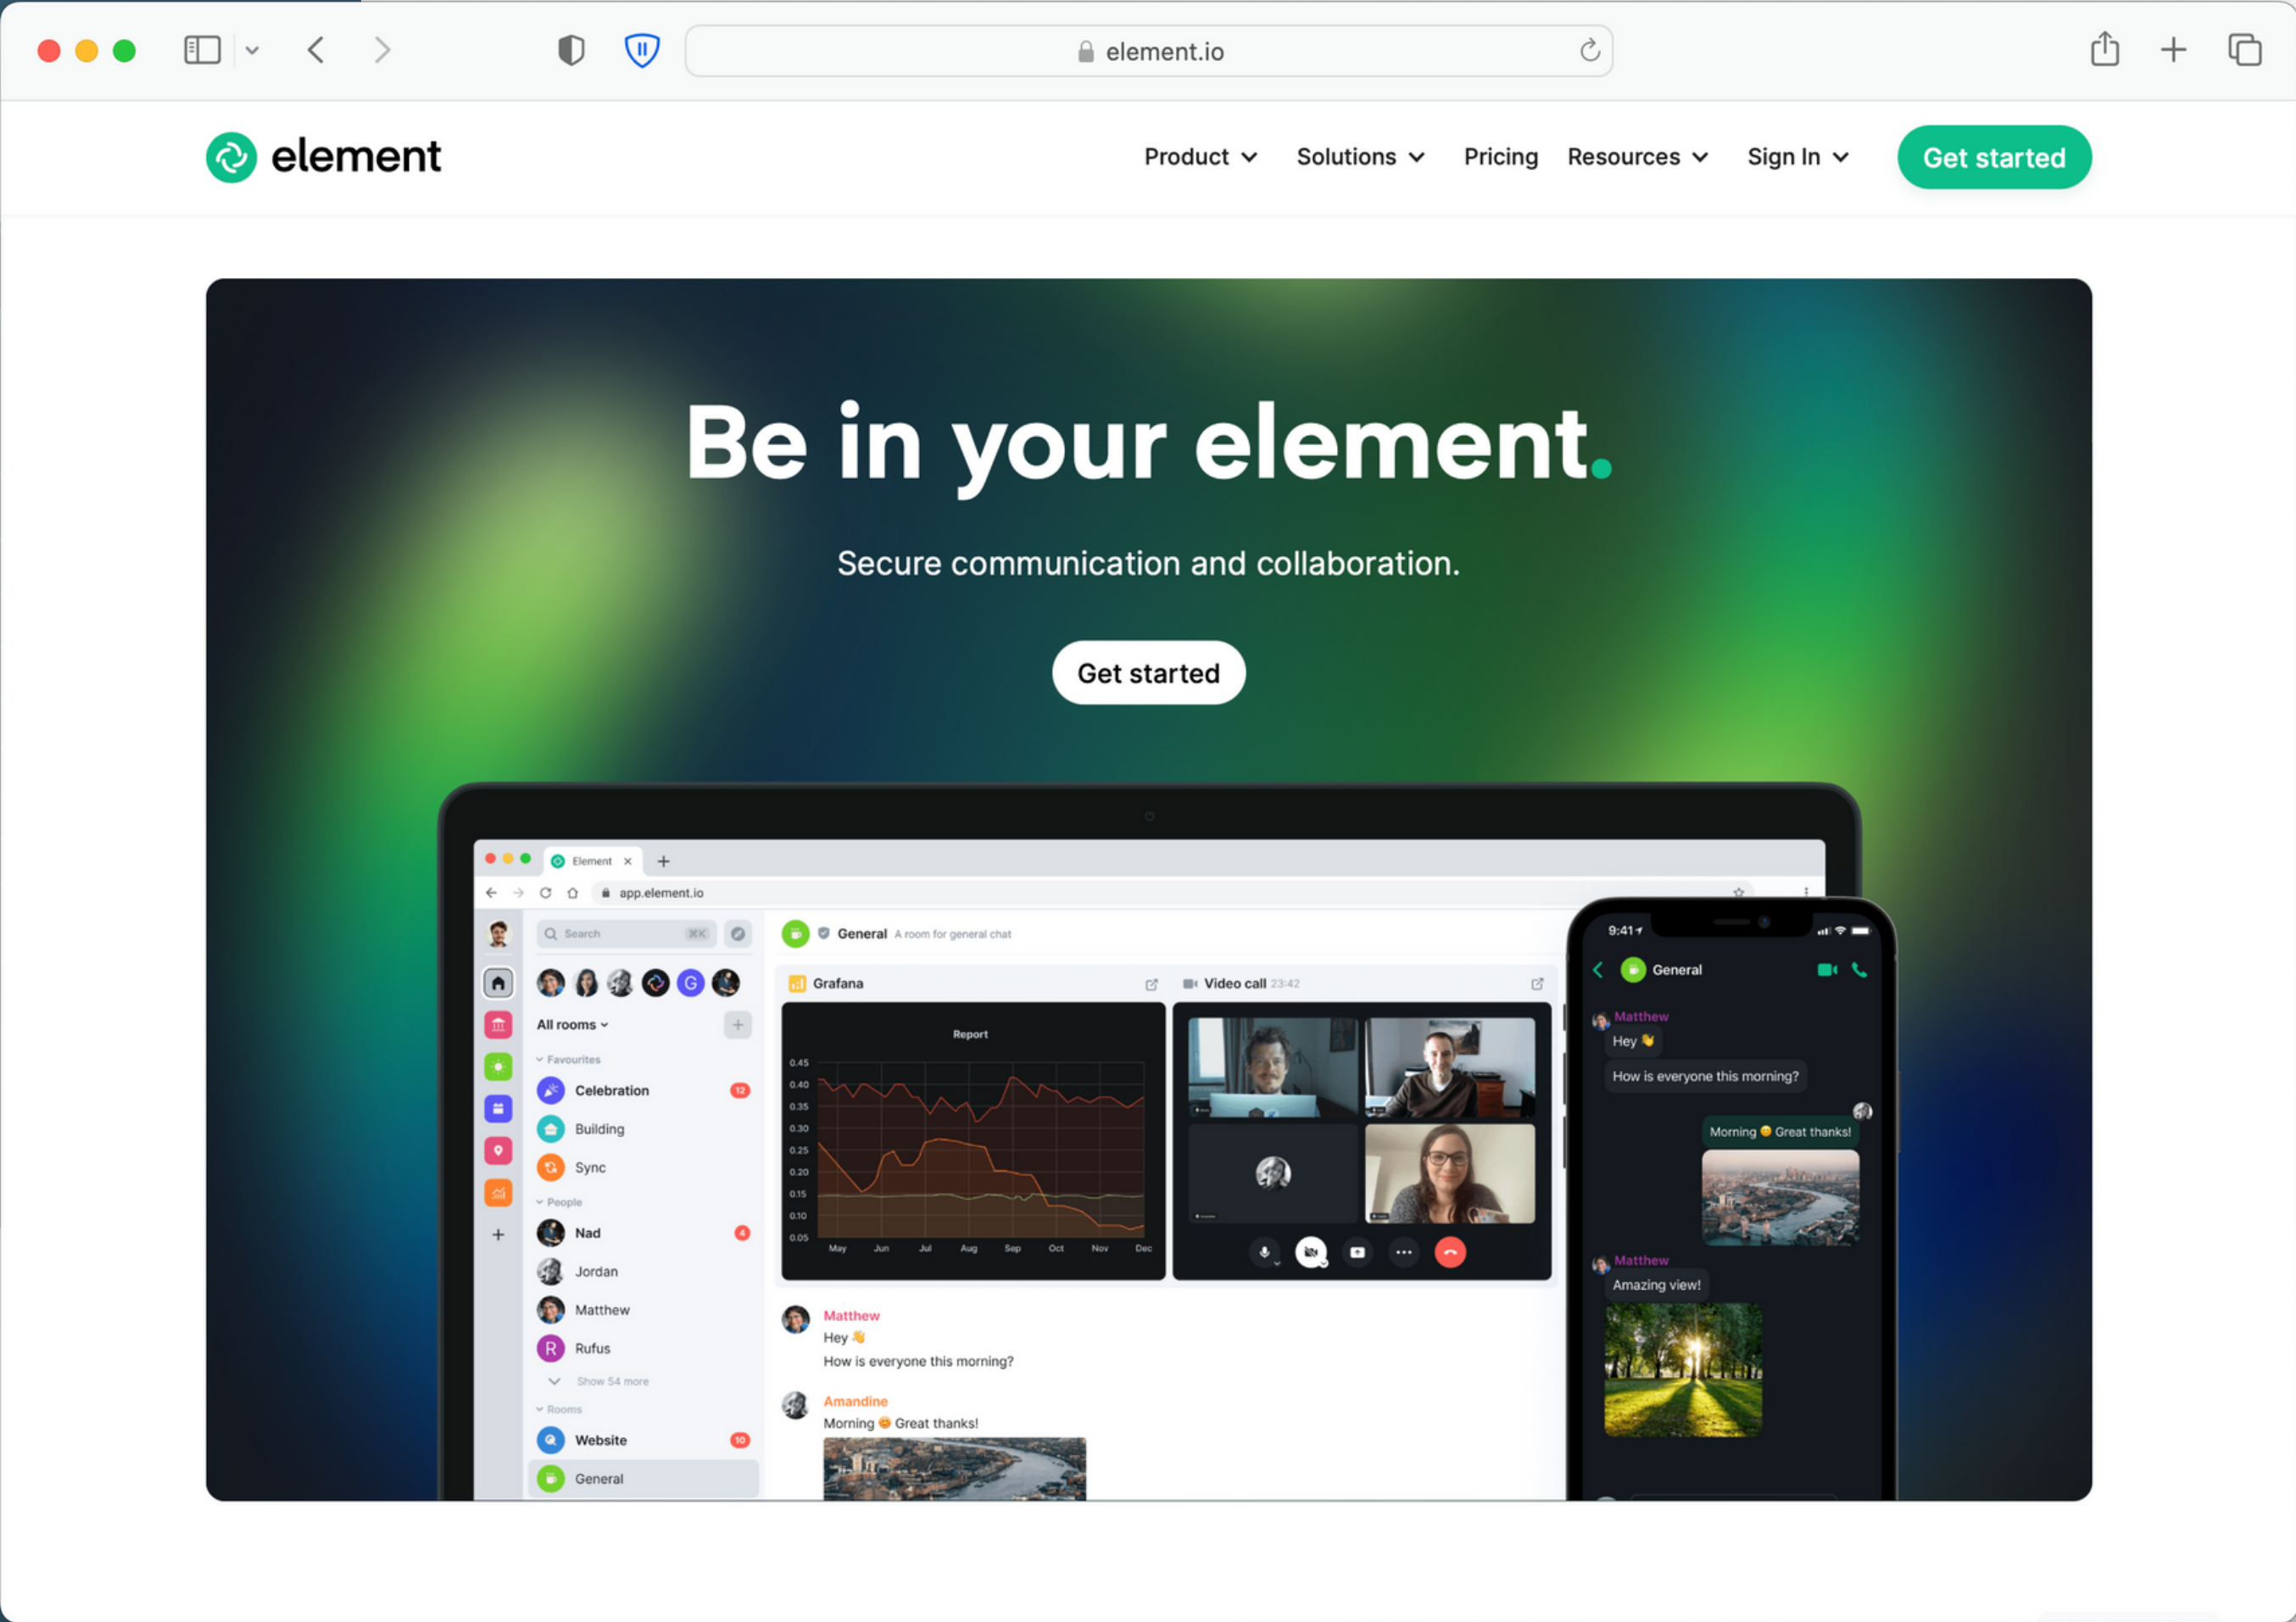
GIAN M. VOLPICELLI SECURITY 14.06.2021 06:00 AM

How governments and spies text each other

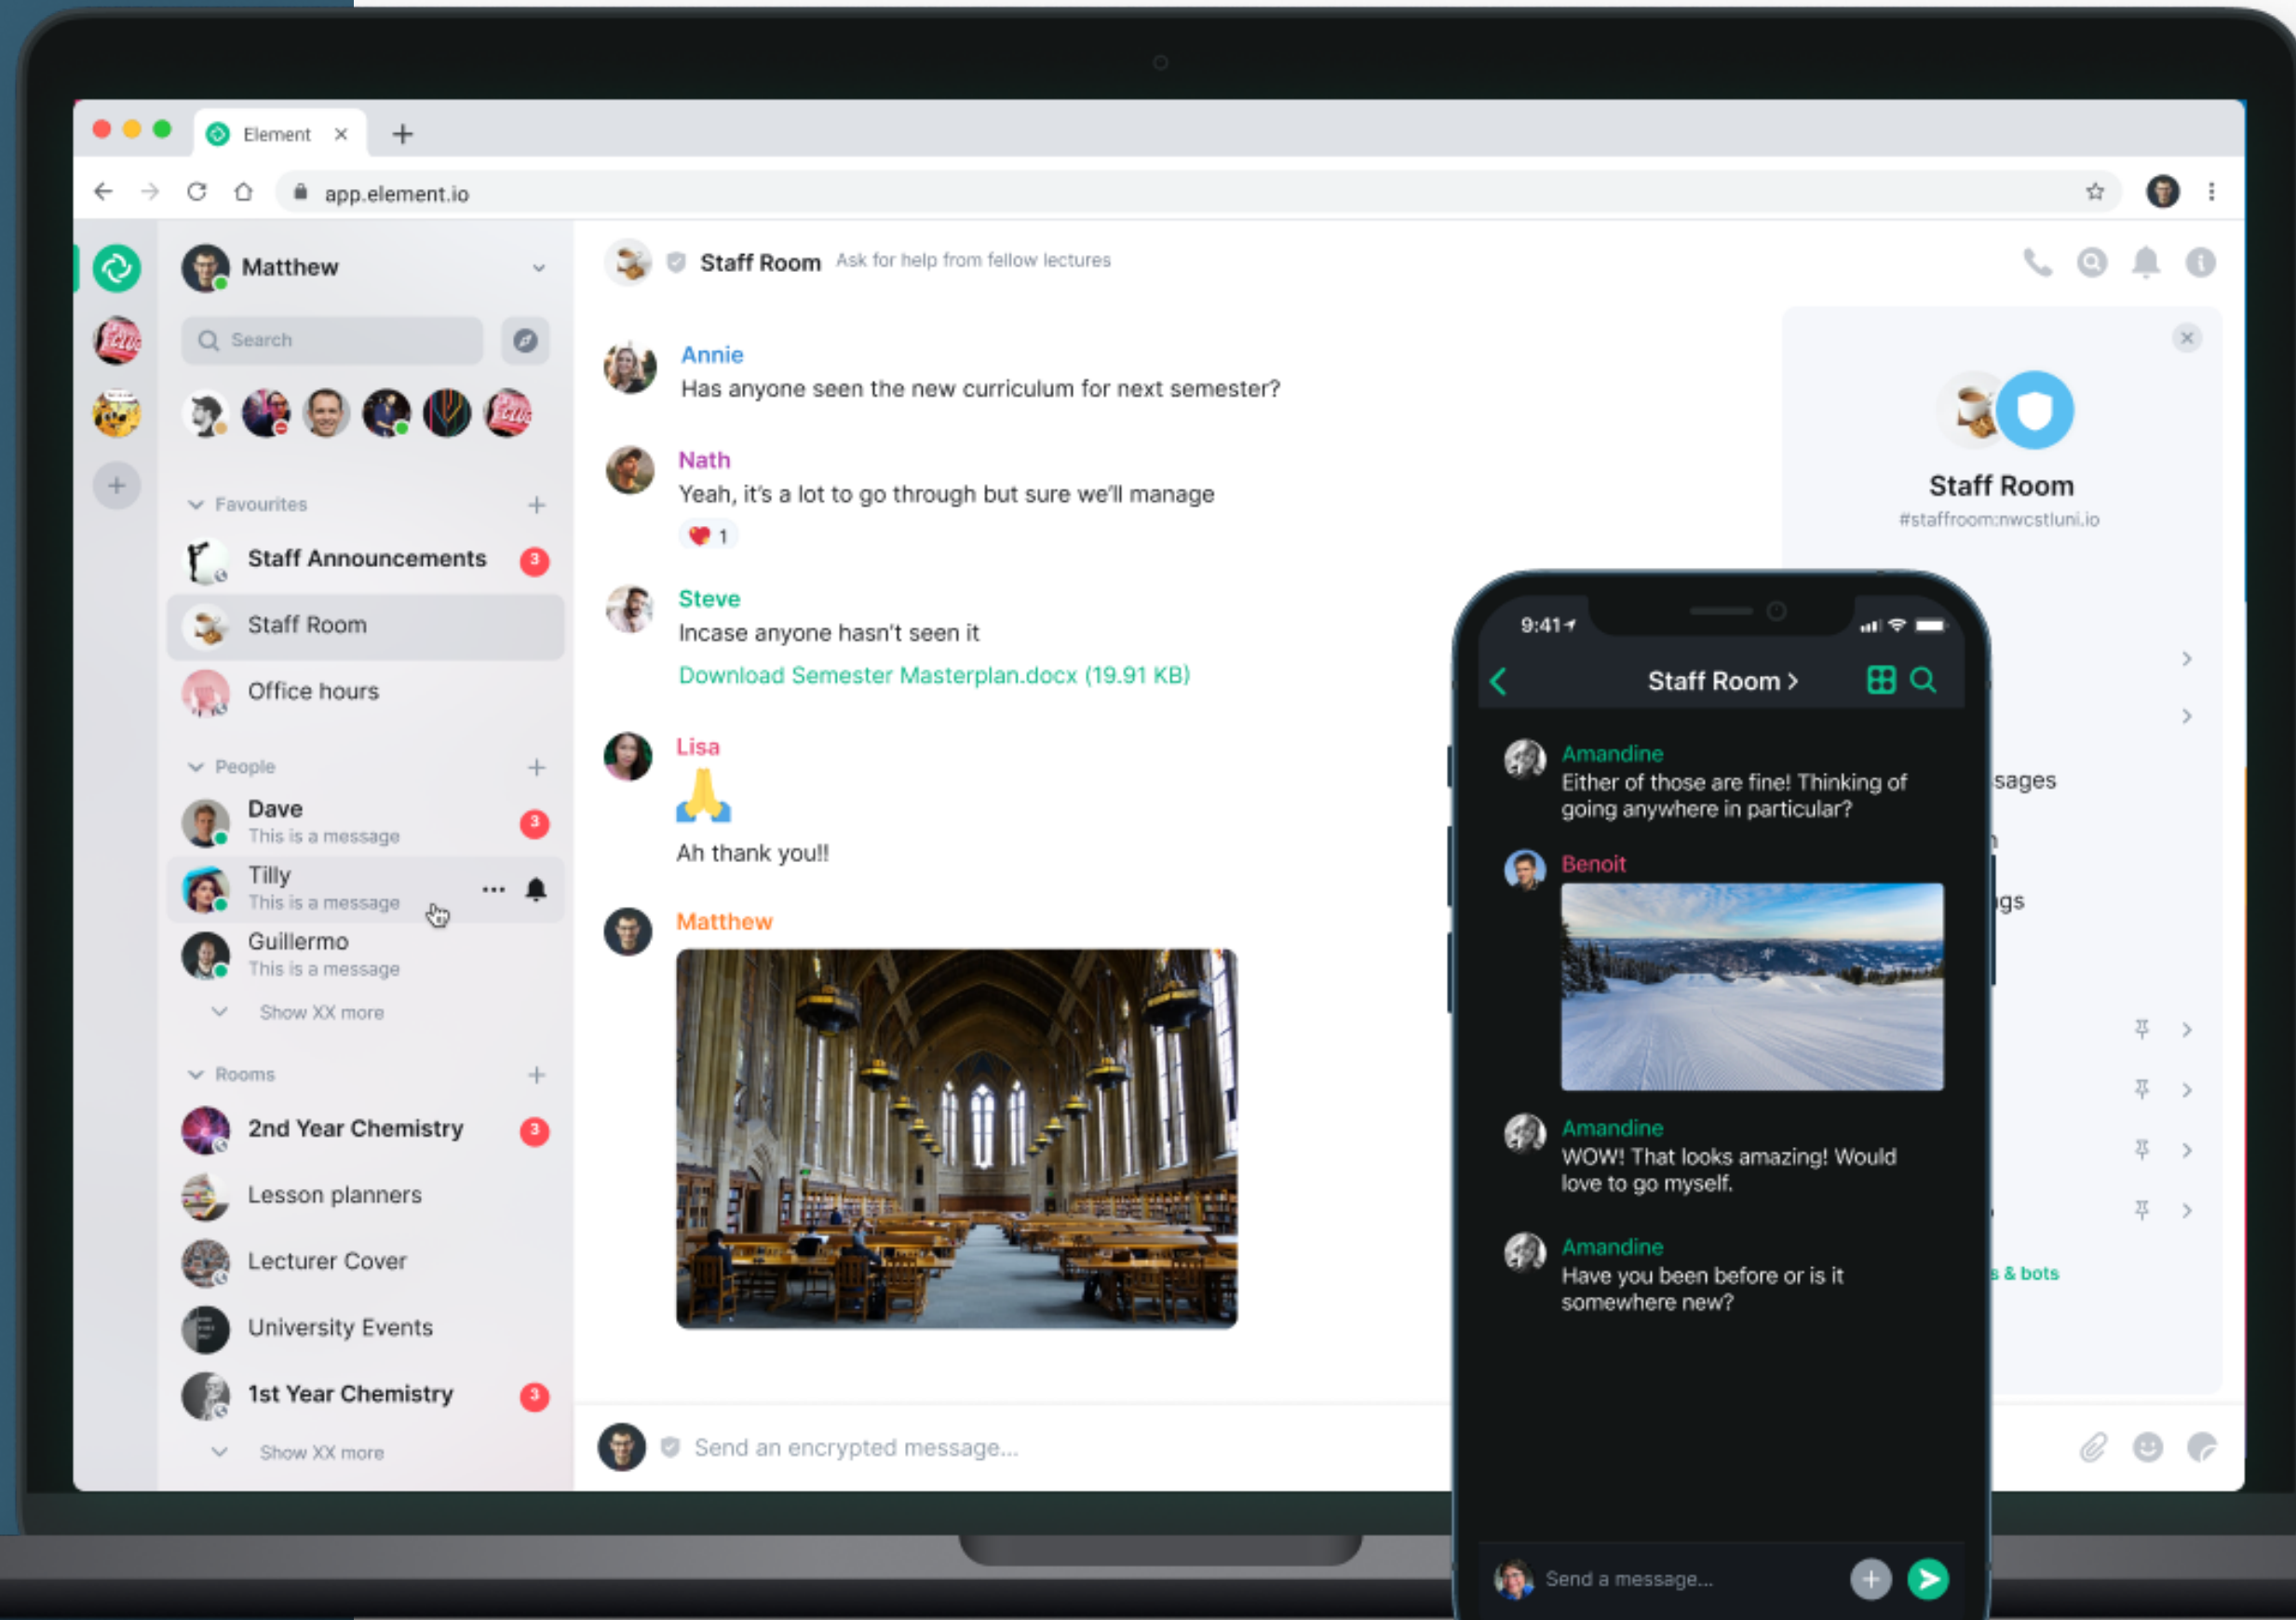
Matrix has become the messaging app of choice for top-secret communications




DAVID RYLE



Element





Practically-exploitable Cryptographic Vulnerabilities in Matrix

Martin R. Albrecht*, Sofía Celi†, Benjamin Dowling‡ and Daniel Jones§

* King's College London, martin.albrecht@kcl.ac.uk

† Brave Software, cherenkov@riseup.net

‡ Security of Advanced Systems Group, University of Sheffield, b.dowling@sheffield.ac.uk

§ Information Security Group, Royal Holloway, University of London, dan.jones@rhul.ac.uk

Abstract—We report several practically-exploitable cryptographic vulnerabilities in the Matrix standard for federated real-time communication and its flagship client and prototype implementation, Element. These, together, invalidate the confidentiality and authentication guarantees claimed by Matrix against a malicious server. This is despite Matrix' cryptographic routines being constructed from well-known and -studied cryptographic building blocks. The vulnerabilities we exploit differ in their nature (insecure by design, protocol confusion, lack of domain separation, implementation bugs) and are distributed broadly across the different subprotocols and libraries that make up the cryptographic core of Matrix and Element. Together, these vulnerabilities highlight the need for a systematic and formal analysis of the cryptography in the Matrix standard.

I. INTRODUCTION

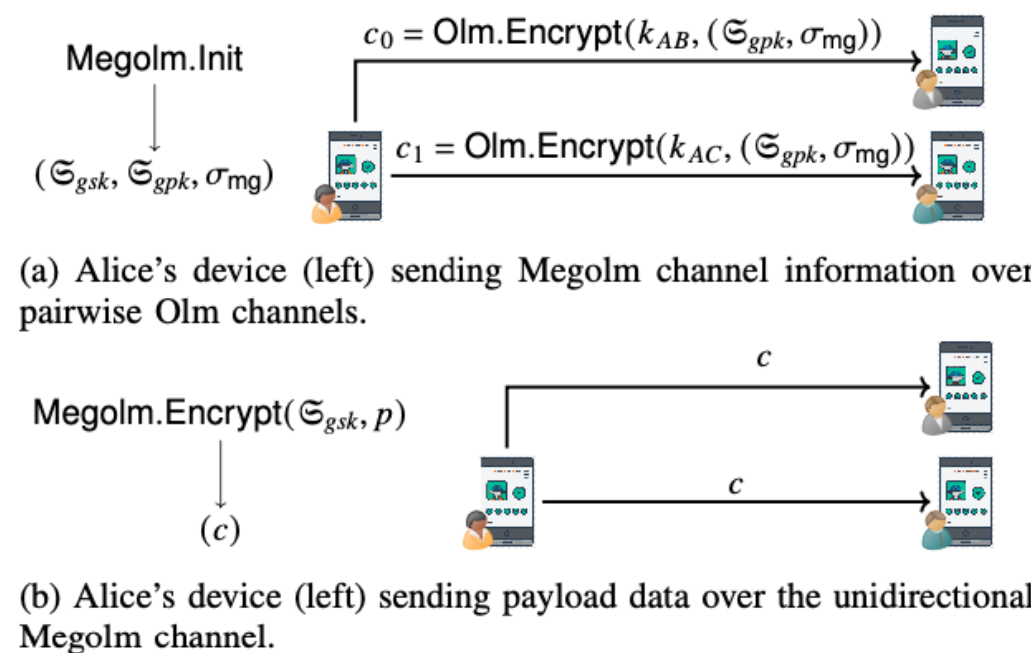



Fig. 1: Alice establishes a Megolm channel (a) and sends a


Matrix

A comprehensive secure messaging solution.




Matrix

A comprehensive secure messaging solution.




Matrix

A comprehensive secure messaging solution.




Matrix

A comprehensive secure messaging solution.



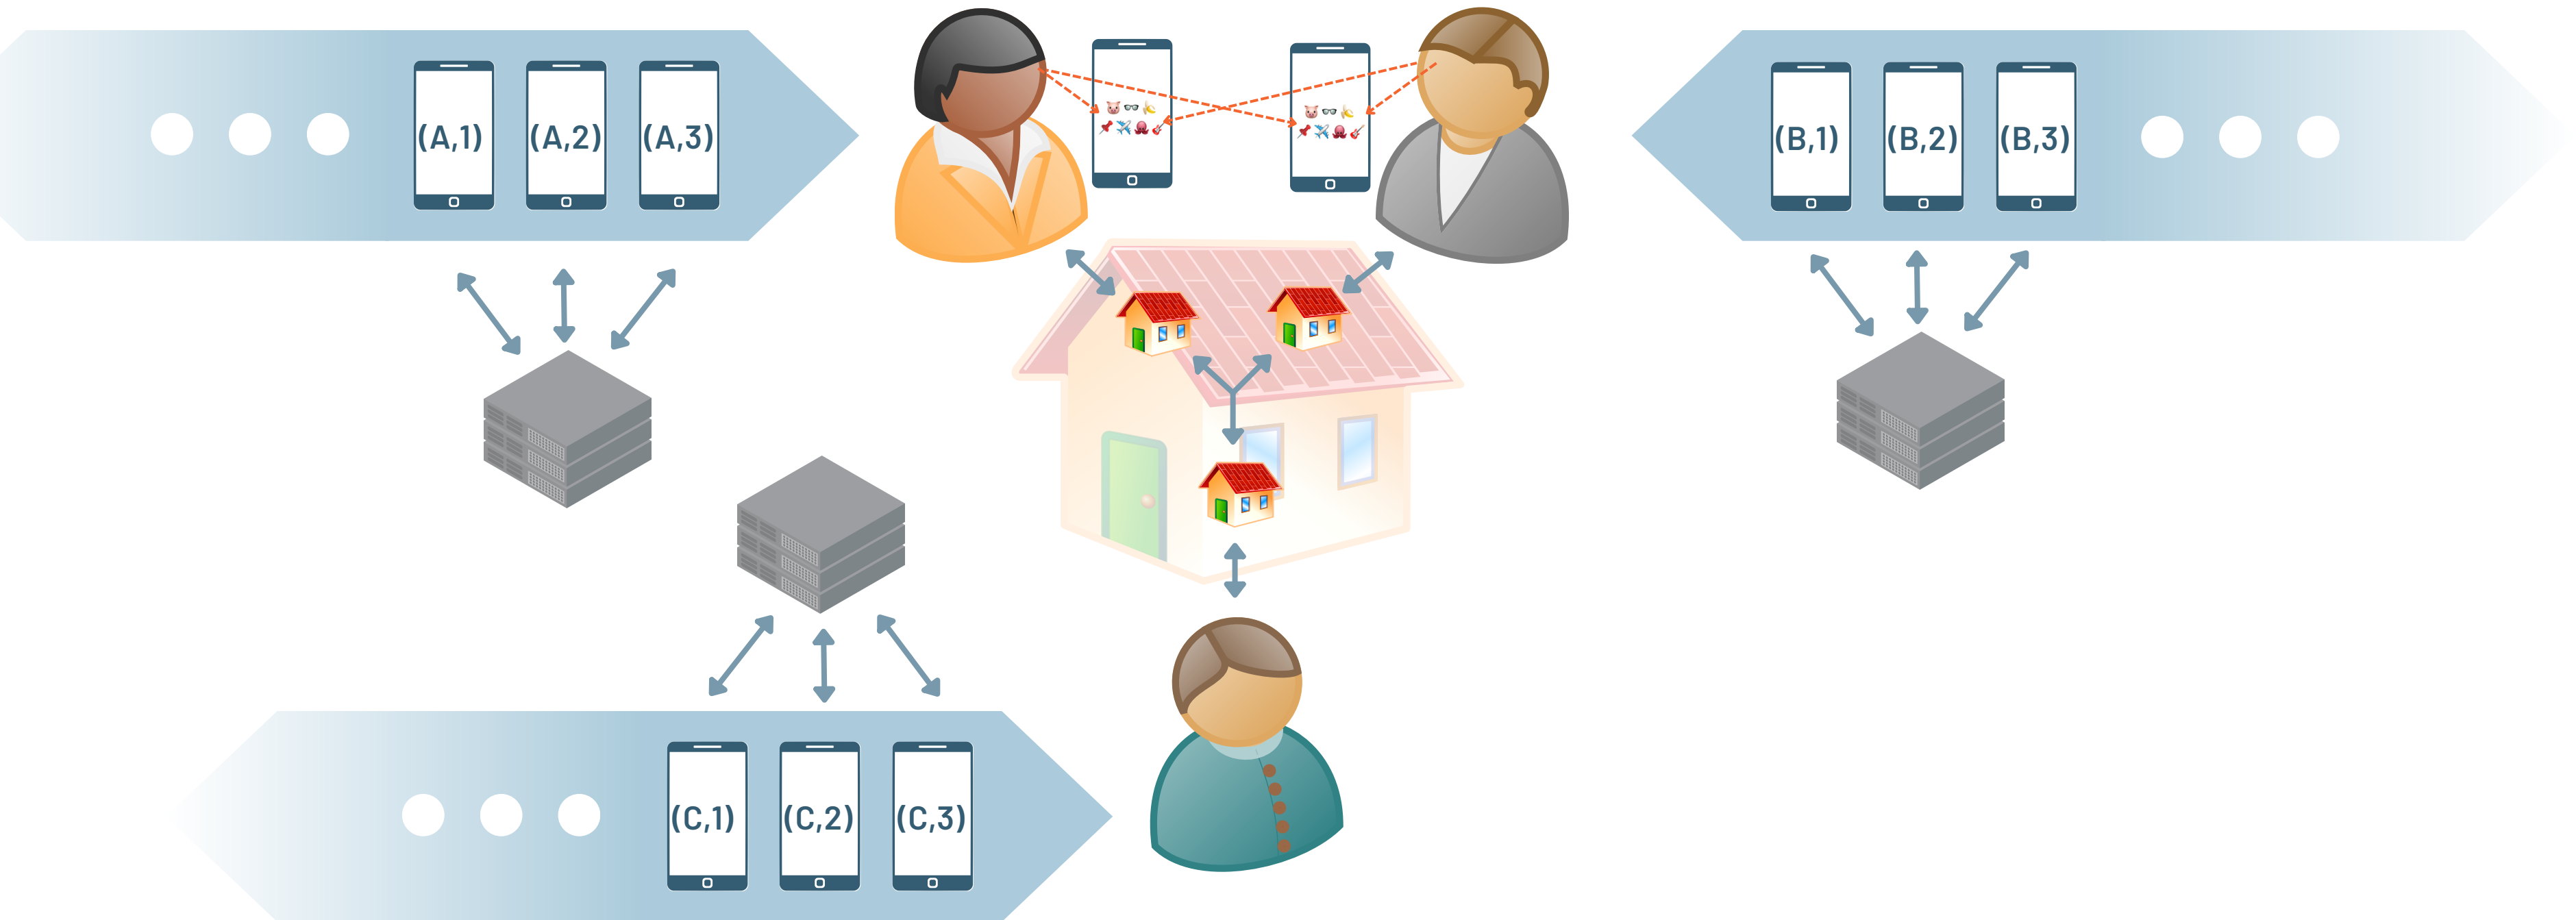
Matrix

A comprehensive secure messaging solution.




Matrix

A comprehensive secure messaging solution.



Matrix

A comprehensive secure messaging solution.



Research

The screenshot shows a web browser window with the URL `https://spec.matrix.org/unstable/client-server-api/#end-to-end-encryption`. The page title is "[matrix] specification — unstable version". The navigation menu includes "Foundation", "FAQs", and "Blog".

The left sidebar contains a table of contents for the "11.12 End-to-End Encryption" section, with "11.12.1 Key Distribution" highlighted. The main content area displays the "11.12. End-to-End Encryption" section, which includes a paragraph about optional end-to-end encryption and a subsection "11.12.1. Key Distribution".

The "11.12.1. Key Distribution" subsection contains a list item: "1. Bob publishes the public keys and supported algorithms for his device. This may include long-term identity keys, and/or one-time keys."

Below the list item is a diagram illustrating the key upload process. It shows two boxes representing "Bob's HS" (Home Server) and "Bob's Device". A double-headed arrow connects them, with the label `/keys/upload` positioned below the arrow.

```
graph TD
    HS[Bob's HS] <--> Device[Bob's Device]
    subgraph Label
        Label["/keys/upload"]
    end
```

Research

```
gitlab.matrix.org/matrix-org/olm X +
https://gitlab.matrix.org
# Megolm group ratchet
An AES-based cryptographic ratchet intended for group communications.
## Background
The Megolm ratchet is intended for encrypted messaging applications where there may be a large number of recipients of each message, thus precluding the use of peer-to-peer encryption systems such as [0lm][].
It also allows a recipient to decrypt received messages multiple times. For instance, in client/server applications, a copy of the ciphertext can be stored on the (untrusted) server, while the client need only store the session keys.
## Overview
Each participant in a conversation uses their own outbound session for encrypting messages. A session consists of a ratchet and an [Ed25519][] keypair.
Secrecy is provided by the ratchet, which can be wound forwards but not backwards, and is used to derive a distinct message key for each message.
Authenticity is provided via Ed25519 signatures.
The value of the ratchet, and the public part of the Ed25519 key, are shared with other participants in the conversation via secure peer-to-peer channels. Provided that peer-to-peer channel provides authenticity of the messages to the participants and deniability of the messages to third parties, the Megolm session will inherit those properties.
## The Megolm ratchet algorithm
The Megolm ratchet  $R_i$  consists of four parts,  $R_{i,j}$  for  $j \in \{0,1,2,3\}$ . The length of each part depends on the hash function in use (256 bits for this version of Megolm).
The ratchet is initialised with cryptographically-secure random data, and advanced as follows:
```

Client-Server API | Matrix Specif X +
https://spec.matrix.org/unstable/client-server-api/#end-to-end-encryption

[matrix] specification — unstable version Foundation FAQs Blog

- 11.11.1 Client behaviour
- 11.11.2 Security considerations
- 11.12 End-to-End Encryption**
 - 11.12.1 Key Distribution
 - 11.12.1.1 Overview
 - 11.12.1.2 Key algorithms
 - 11.12.1.3 Device keys
 - 11.12.1.4 Uploading keys
 - 11.12.1.5 Tracking the device list for a user
 - 11.12.1.6 Sending encrypted attachments
 - 11.12.1.6.1 Extensions to m.room.message msgtypes
 - 11.12.1.7 Claiming one-time keys
 - 11.12.2 Device verification
 - 11.12.2.1 Key verification framework
 - 11.12.2.2 Short Authentication String

11.12. End-to-End Encryption

Matrix optionally supports end-to-end encryption, allowing rooms to be created whose conversation contents are not decryptable or interceptable on any of the participating homeservers.

11.12.1. Key Distribution

Encryption and Authentication in Matrix is based around public-key cryptography. The Matrix protocol provides a basic mechanism for exchange of public keys, though an out-of-band channel is required to exchange fingerprints between users to build a web of trust.

11.12.1.1. Overview

1. Bob publishes the public keys and supported algorithms for his device. This may include long-term identity keys, and/or one-time keys.

```
+-----+ +-----+
| Bob's HS | | Bob's Device |
+-----+ +-----+
          |           |
          |<=====|
          /keys/upload
```

Research

```
gitlab.matrix.org/matrix-org/olm X
https://gitlab.matrix.org

# Megolm group ratchet
An AES-based cryptographic ratchet intended for group communications.

## Background
The Megolm ratchet is intended for encrypted messaging applications where there may be a large number of recipients of each message, thus precluding the use of peer-to-peer encryption systems such as [0lm][1].

It also allows a recipient to decrypt received messages, for example, in client/server applications, a copy of the key is stored on the (untrusted) server, while the client stores the key on the client.

## Overview
Each participant in a conversation uses their own key to encrypt messages. A session consists of a sequence of keys, each derived from the previous one. Secrecy is provided by the ratchet, which can be used to derive a distinct key for each message. Authenticity is provided via Ed25519 signatures. The value of the ratchet, and the public part of the key, is shared with other participants in the conversation via secure channels. Provided that peer-to-peer channels are secure, messages to the participants and deniability of the Megolm session will inherit those properties.

## The Megolm ratchet algorithm
The Megolm ratchet  $R_i$  consists of four parts  $R_{i,j}$  for  $j \in \{0,1,2,3\}$ . The length of each part is 256 bits for this version of Megolm.

The ratchet is initialised with cryptographic keys advanced as follows:
```

```
gitlab.matrix.org/matrix-org/olm X
https://gitlab.matrix.org

# Olm: A Cryptographic Ratchet
An implementation of the double cryptographic ratchet described by https://whispersystems.org/docs/specifications/doubleratchet/.

## Notation
This document uses  $\parallel$  to represent string concatenation. When  $\parallel$  appears on the right hand side of an  $=$  it means that the inputs are concatenated. When  $\parallel$  appears on the left hand side of an  $=$  it means that the output is split.

When this document uses  $\text{ECDH}(K_A, K_B)$  it means that each party computes a Diffie-Hellman agreement using their private key and the remote party's public key. So party  $A$  computes  $\text{ECDH}(K_B^{\text{public}}, K_A^{\text{private}})$  and party  $B$  computes  $\text{ECDH}(K_A^{\text{public}}, K_B^{\text{private}})$ .

Where this document uses  $\text{HKDF}(\text{salt}, \text{IKM}, \text{info}, L)$  it refers to the [HMAC-based key derivation function] with a salt value of  $\text{salt}$ , input key material of  $\text{IKM}$ , context string  $\text{info}$ , and output keying material length of  $L$  bytes.

## The Olm Algorithm
### Initial setup
The setup takes four [Curve25519] inputs: Identity keys for Alice and Bob,  $I_A$  and  $I_B$ , and one-time keys for Alice and Bob,  $E_A$  and  $E_B$ . A shared secret,  $S$ , is generated using [Triple Diffie-Hellman]. The initial 256 bit root key,  $R_0$ , and 256 bit chain key,  $C_{0,0}$ , are derived from the shared secret using an HMAC-based Key Derivation Function [SHA-256] as the hash function ([HKDF-SHA-256]) with default salt and "OLM_ROOT" as the info.


$$\begin{aligned} & \end{aligned}$$

```

Client-Server API | Matrix Specification

https://spec.matrix.org/unstable/client-server-api/#end-to-end-encryption

[matrix] specification — unstable version

Foundation FAQs Blog

- 11.11.1 Client behaviour
- 11.11.2 Security considerations
- 11.12 End-to-End Encryption
 - 11.12.1 Key Distribution
 - 11.12.1.1 Overview
 - 11.12.1.2 Key algorithms
 - 11.12.1.3 Device keys

11.12. End-to-End Encryption


Matrix optionally supports end-to-end encryption, allowing rooms to be created whose conversation contents are not decryptable or interceptable on any of the participating homeservers.

11.12.1. Key Distribution

Encryption and Authentication in Matrix is based around public-key cryptography. The Matrix protocol provides a basic mechanism for exchange of public keys, though an out-of-band channel is required to exchange fingerprints between users to build a web of trust.

11.12.1.1. Overview

1. Bob publishes the public keys and supported algorithms for his device. This may include long-term identity keys, and/or one-time keys.



Research

```
gitlab.matrix.org/matrix-org/olm X
https://gitlab.matrix.org

# Megolm group ratchet
An AES-based cryptographic ratchet intended for group communications.

## Background
The Megolm ratchet is intended for encrypted messaging applications where there may be a large number of recipients of each message, thus precluding the use of peer-to-peer encryption systems such as [0lm][1].

It also allows a recipient to decrypt received messages, for example, in client/server applications, a client can decrypt messages on the (untrusted) server, while the client can decrypt messages on the client.

## Overview
Each participant in a conversation uses their own key to encrypt messages. A session consists of a sequence of keys, each derived from the previous one. Secrecy is provided by the ratchet, which can be used to derive a distinct key for each message. Authenticity is provided via Ed25519 signatures. The value of the ratchet, and the public part of the keys, are shared with other participants in the conversation via secure channels. Provided that peer-to-peer channels are secure, messages to the participants and deniability of the Megolm session will inherit those properties.

## The Megolm ratchet algorithm
The Megolm ratchet  $R_i$  consists of four parts  $j \in \{0,1,2,3\}$ . The length of each part is 256 bits for this version of Megolm.

The ratchet is initialised with cryptographic keys advanced as follows:
```

```
gitlab.matrix.org/matrix-org/olm X
https://gitlab.matrix.org

# Olm: A Cryptographic Ratchet
An implementation of the double cryptographic ratchet described by https://whispersystems.org/docs/specifications/doubleratchet/.

## Notation
This document uses  $\parallel$  to represent string concatenation. When  $\parallel$  appears on the right hand side of an  $=$  it means that the inputs are concatenated. When  $\parallel$  appears on the left hand side of an  $=$  it means that the output is split.

When this document uses  $\operatorname{ECDH}(K_A, K_B)$  it means that each party computes a Diffie-Hellman agreement using their private key and the remote party's public key. So party  $A$  computes  $\operatorname{ECDH}(K_B^{\text{public}}, K_A^{\text{private}})$  and party  $B$  computes  $\operatorname{ECDH}(K_A^{\text{public}}, K_B^{\text{private}})$ .

Where this document uses  $\operatorname{HKDF}(\text{salt}, \text{IKM}, \text{info}, L)$  it refers to the [HMAC-based key derivation function] with a salt value of  $\text{salt}$ , input key material of  $\text{IKM}$ , context string  $\text{info}$ , and output keying material length of  $L$  bytes.

## The Olm Algorithm
### Initial setup
The setup takes four [Curve25519] inputs: Identity keys for Alice and Bob,  $I_A$  and  $I_B$ , and one-time keys for Alice and Bob,  $E_A$  and  $E_B$ . A shared secret,  $S$ , is generated using [Triple Diffie-Hellman]. The initial 256 bit root key,  $R_0$ , and 256 bit chain key,  $C_{0,0}$ , are derived from the shared secret using an HMAC-based Key Derivation Function [SHA-256] as the hash function ([HKDF-SHA-256]) with default salt and "OLM_ROOT" as the info.


$$\begin{aligned} \end{aligned}$$

```

Client-Server API | Matrix Specification

https://spec.matrix.org/unstable/client-server-api/#end-to-end-encryption

[matrix] specification — unstable version

Foundation FAQs Blog

- 11.11.1 Client behaviour
- 11.11.2 Security considerations
- 11.12 End-to-End Encryption
 - 11.12.1 Key Distribution
 - 11.12.1.1 Overview
 - 11.12.1.2 Key algorithms
 - 11.12.1.3 Device keys

11.12. End-to-End Encryption


Matrix optionally supports end-to-end encryption, allowing rooms to be created whose conversation contents are not decryptable or interceptable on any of the participating homeservers.

11.12.1. Key Distribution

Encryption and Authentication in Matrix is based on public key cryptography. The Matrix protocol provides a mechanism for exchange of public keys, the channel is required to exchange fingerprints to build a web of trust.

11.12.1.1. Overview

1. Bob publishes the public keys and supports them for his device. This may include long-term and/or one-time keys.



[m] Matrix.org - End-to-End Encryption

https://matrix.org/docs/leg...

Implementing End-to-End Encryption in Matrix clients

This guide is intended for authors of Matrix clients who wish to add support for end-to-end encryption. It is highly recommended that readers be familiar with the Matrix protocol and the use of access tokens before proceeding.

The libolm library

End-to-end encryption in Matrix is based on the Olm and Megolm cryptographic ratchets. The recommended starting point for any client authors is with the `libolm` library, which contains implementations of all of the cryptographic primitives required. The library itself is written in C/C++, but is architected in a way which makes it easy to write wrappers for higher-level languages.

Devices

We have a particular meaning for "device". As a user, I might have several devices (a desktop client, some web browsers, an Android device, an iPhone, etc). When I first use a client, it should register itself as a new device. If I log out and log in again as a different user, the client must register as a new device. Critically, the client must create a new set of keys (see below) for each "device".

The longevity of devices will depend on the client. In the web client, we create a new device every single time you log in. In a mobile client, it might be acceptable to reuse the device if a login session expires, **provided** the user is the same. **Never** share keys between different users.

Devices are identified by their `device_id` (which is unique within the scope of a given user).

Research

```
# Megolm group ratchet
An AES-based cryptographic ratchet intended for group communication.

## Background
The Megolm ratchet is intended for encrypted messaging applications. It may be a large number of recipients of each message, thus precluding peer-to-peer encryption systems such as [Olm][1].

It also allows a recipient to decrypt received messages. For example, in an instance, in client/server applications, a copy of the key is stored on the (untrusted) server, while the client ratchets the key.

## Overview
Each participant in a conversation uses their own key to encrypt and decrypt messages. A session consists of a sequence of keys. The key for a message is derived from the previous key. Secrecy is provided by the ratchet, which can be used to derive a distinct key for each message. Authenticity is provided via Ed25519 signatures. The value of the ratchet, and the public part of the key, is shared with other participants in the conversation via direct channels. Provided that peer-to-peer channels are secure, messages to the participants and deniability are maintained. The Megolm session will inherit those properties.

## The Megolm ratchet algorithm
The Megolm ratchet  $R_i$  consists of four parts  $j \in \{0,1,2,3\}$ . The length of each part is 256 bits for this version of Megolm.

The ratchet is initialised with cryptographic primitives advanced as follows:
```

```
# Olm: A Cryptographic Ratchet
An implementation of the Olm ratchet is available at https://whispersystem.org/olm/

## Notation
This document uses the notation  $A \parallel B$  to denote the concatenation of the inputs A and B. The notation  $A \oplus B$  denotes the XOR of A and B. When this document refers to a key, it means the key for the current message. So party A's key is  $K_A$  and party B's key is  $K_B$ .
```

```
## The Olm Algorithm
### Initial setup
The setup takes four inputs: two keys  $I_A$  and  $I_B$ , and one-time keys for Alice and Bob,  $E_A$  and  $E_B$ . A shared secret,  $S$ , is generated using [Triple Diffie-Hellman][1]. The initial 256 bit root key,  $R_0$ , and 256 bit chain key,  $C_{0,0}$ , are derived from the shared secret using an HMAC-based Key Derivation Function [SHA-256][2] as the hash function ([HKDF-SHA-256][3]) with default salt and "OLM_ROOT" as the info.


$$\begin{aligned} & \text{...} \\ & \text{...} \end{aligned}$$

```

```
Client-Server API | Matrix Specification
https://spec.matrix.org/unstable/client-server-api/#end-to-end-encryption

SAS.ts - matrix-org/matrix-js-sdk
https://sourcegraph.com/github.com/matrix-org/matrix-js-sdk@706b4d6

context:global
matrix-org/matrix-js-sdk @ 706b4d6 > / src / crypto / verification / SAS.ts

Files Symbols
Search symbols...

EVENTS
EmojiMapping
EventHandlerMap
HASHES_LIST
HASHES_SET
IGeneratedSas
decimal
emoji
ISasEvent
cancel
confirm
mismatch
sas
KEY_AGREEMENT_LIST
KEY_AGREEMENT_SET
MAC_LIST
MAC_SET
SAS
NAME
deVerification

1 /*
2 Copyright 2018 - 2021 The Matrix.org Foundation C.I.C.
3
4 Licensed under the Apache License, Version 2.0 (the "License");
5 you may not use this file except in compliance with the License.
6 You may obtain a copy of the License at
7
8 http://www.apache.org/licenses/LICENSE-2.0
9
10 Unless required by applicable law or agreed to in writing, software
11 distributed under the License is distributed on an "AS IS" BASIS,
12 WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
13 See the License for the specific language governing permissions and
14 limitations under the License.
15 */
16
17 /**
18 * Short Authentication String (SAS) verification.
19 * @module crypto/verification/SAS
20 */
21
22 import anotherjson from 'another-json';
23 import { Utility, SAS as OlmSAS } from '@matrix-org/olm';
24
25 import { VerificationBase as Base, SwitchStartEventError, VerificationEvent } from './Verification';
26 import {
27   errorFactory,
28   newInvalidMessageError,
29   newKeyMismatchError,
30   newUnknownMethodError,
31   newUserCancelledError,
32 } from './Error';
33 import { logger } from '../logger';
34 import { IContent, MatrixEvent } from '../models/event';
35
36 const START_TYPE = "m.key.verification.start";
37
38 const EVENTS = [
39   "m.key.verification.accept",
```

Foundation FAQs Blog

Implementing End-to-End Encryption in Matrix clients

This guide is intended for authors of Matrix clients who wish to add support for end-to-end encryption. It is highly recommended that readers be familiar with the Matrix protocol and the use of access tokens before proceeding.

The libolm library

End-to-end encryption in Matrix is based on the Olm and Megolm cryptographic ratchets. The recommended starting point for any client authors is with the `libolm` library, which contains implementations of all of the cryptographic primitives required. The library itself is written in C/C++, but is architected in a way which makes it easy to write wrappers for higher-level languages.

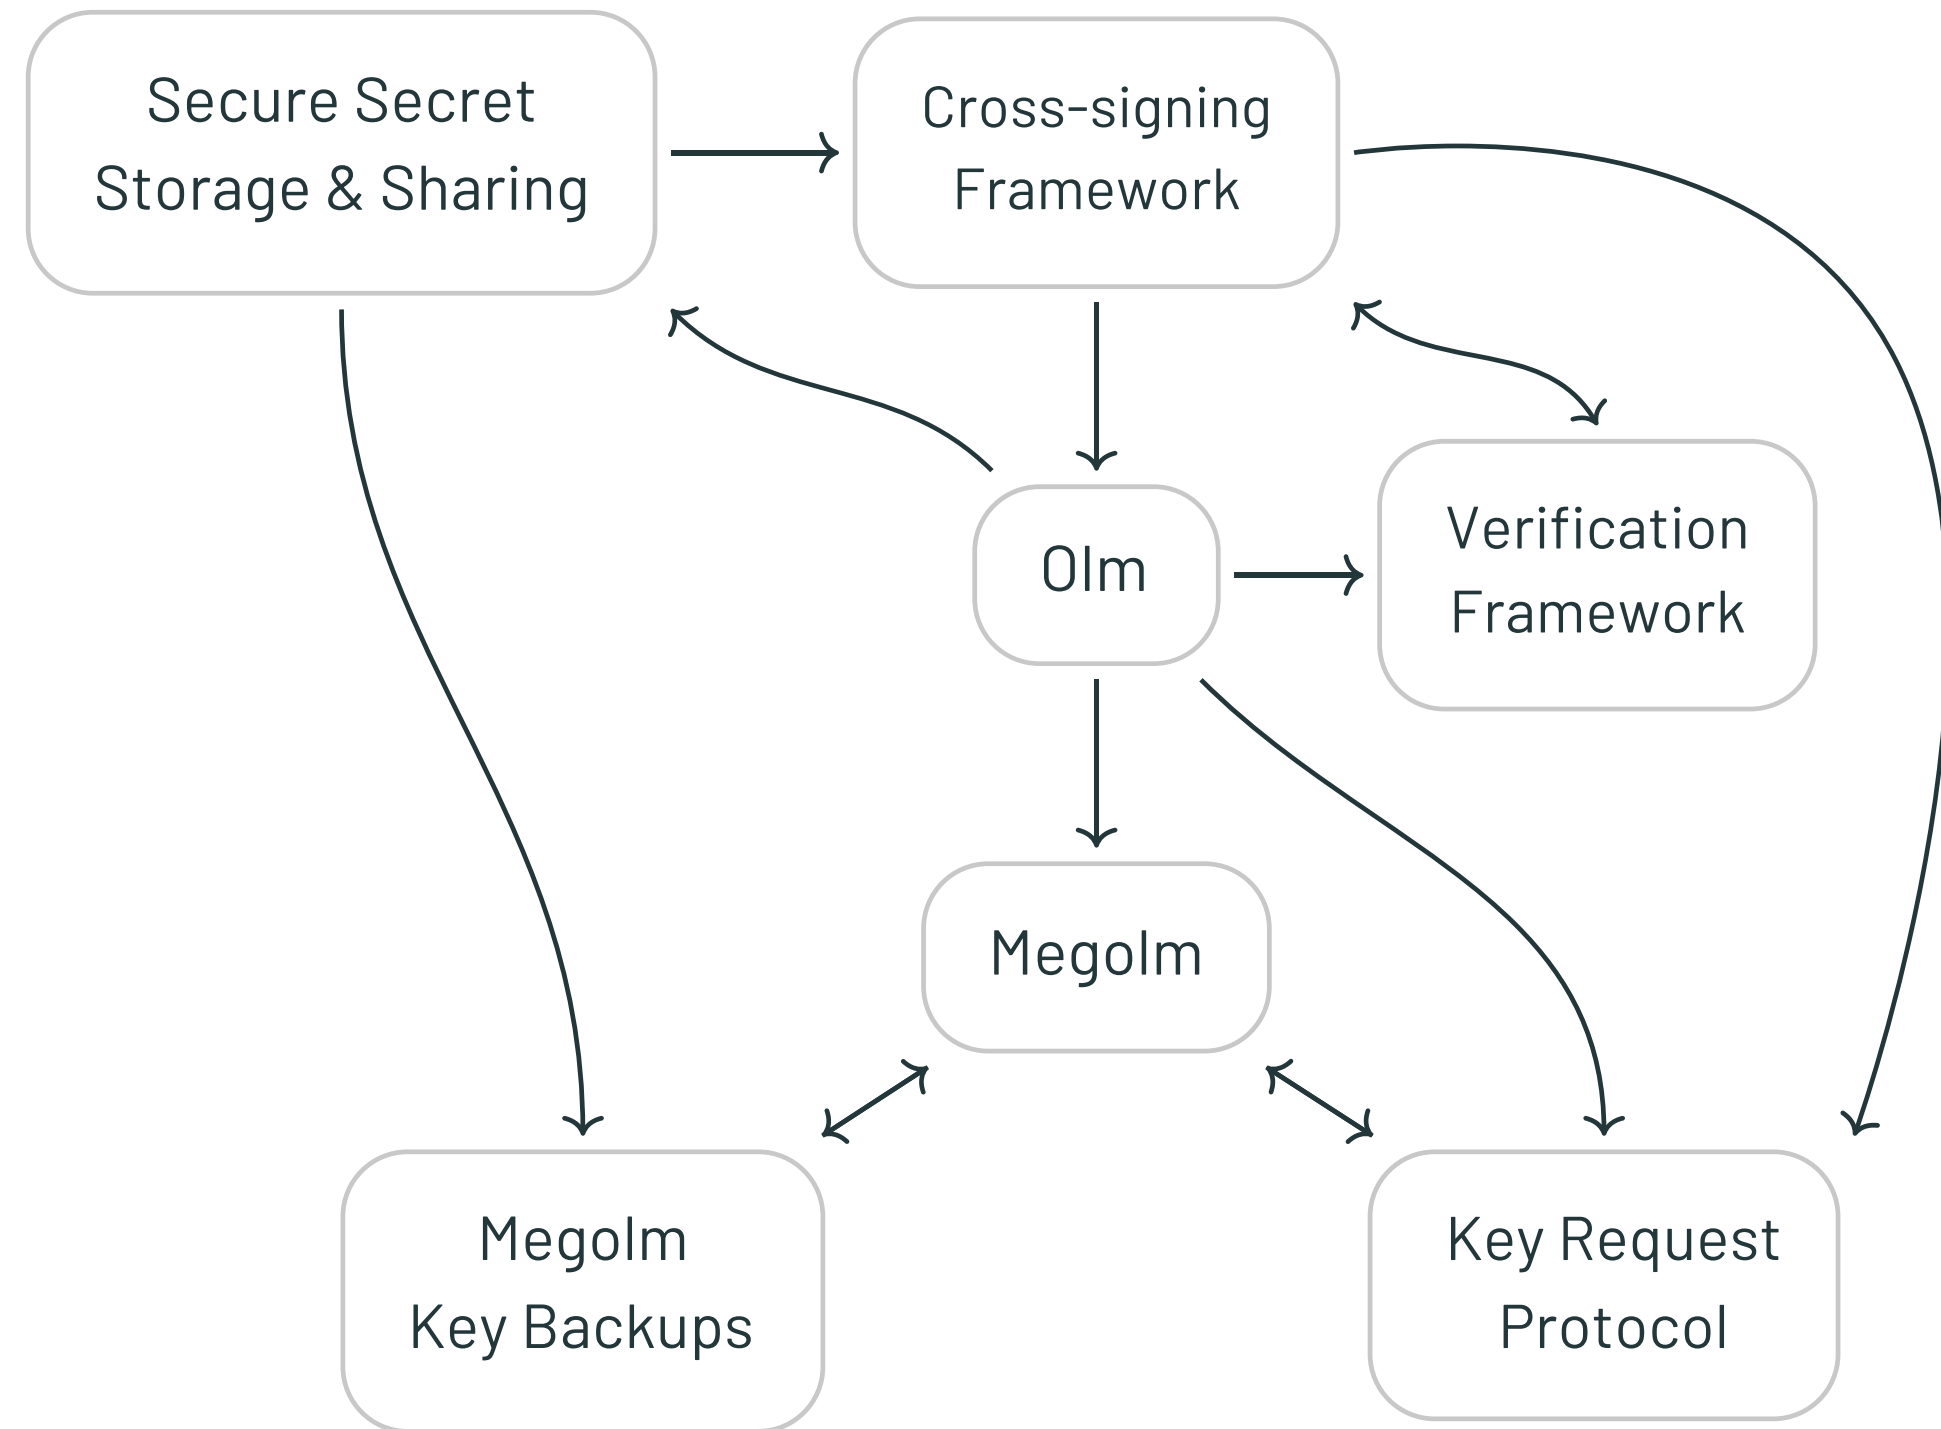
Devices

We have a particular meaning for "device". As a user, I might have several devices (a desktop client, some web browsers, an Android device, an iPhone, etc). When I first use a client, it should register itself as a new device. If I log out and log in again as a different user, the client must register as a new device. Critically, the client must create a new set of keys (see below) for each "device".

The longevity of devices will depend on the client. In the web client, we create a new device every single time you log in. In a mobile client, it might be acceptable to reuse the device if a login session expires, **provided** the user is the same. **Never** share keys between different users.

Devices are identified by their `device_id` (which is unique within the scope of a given user).

Matrix Overview

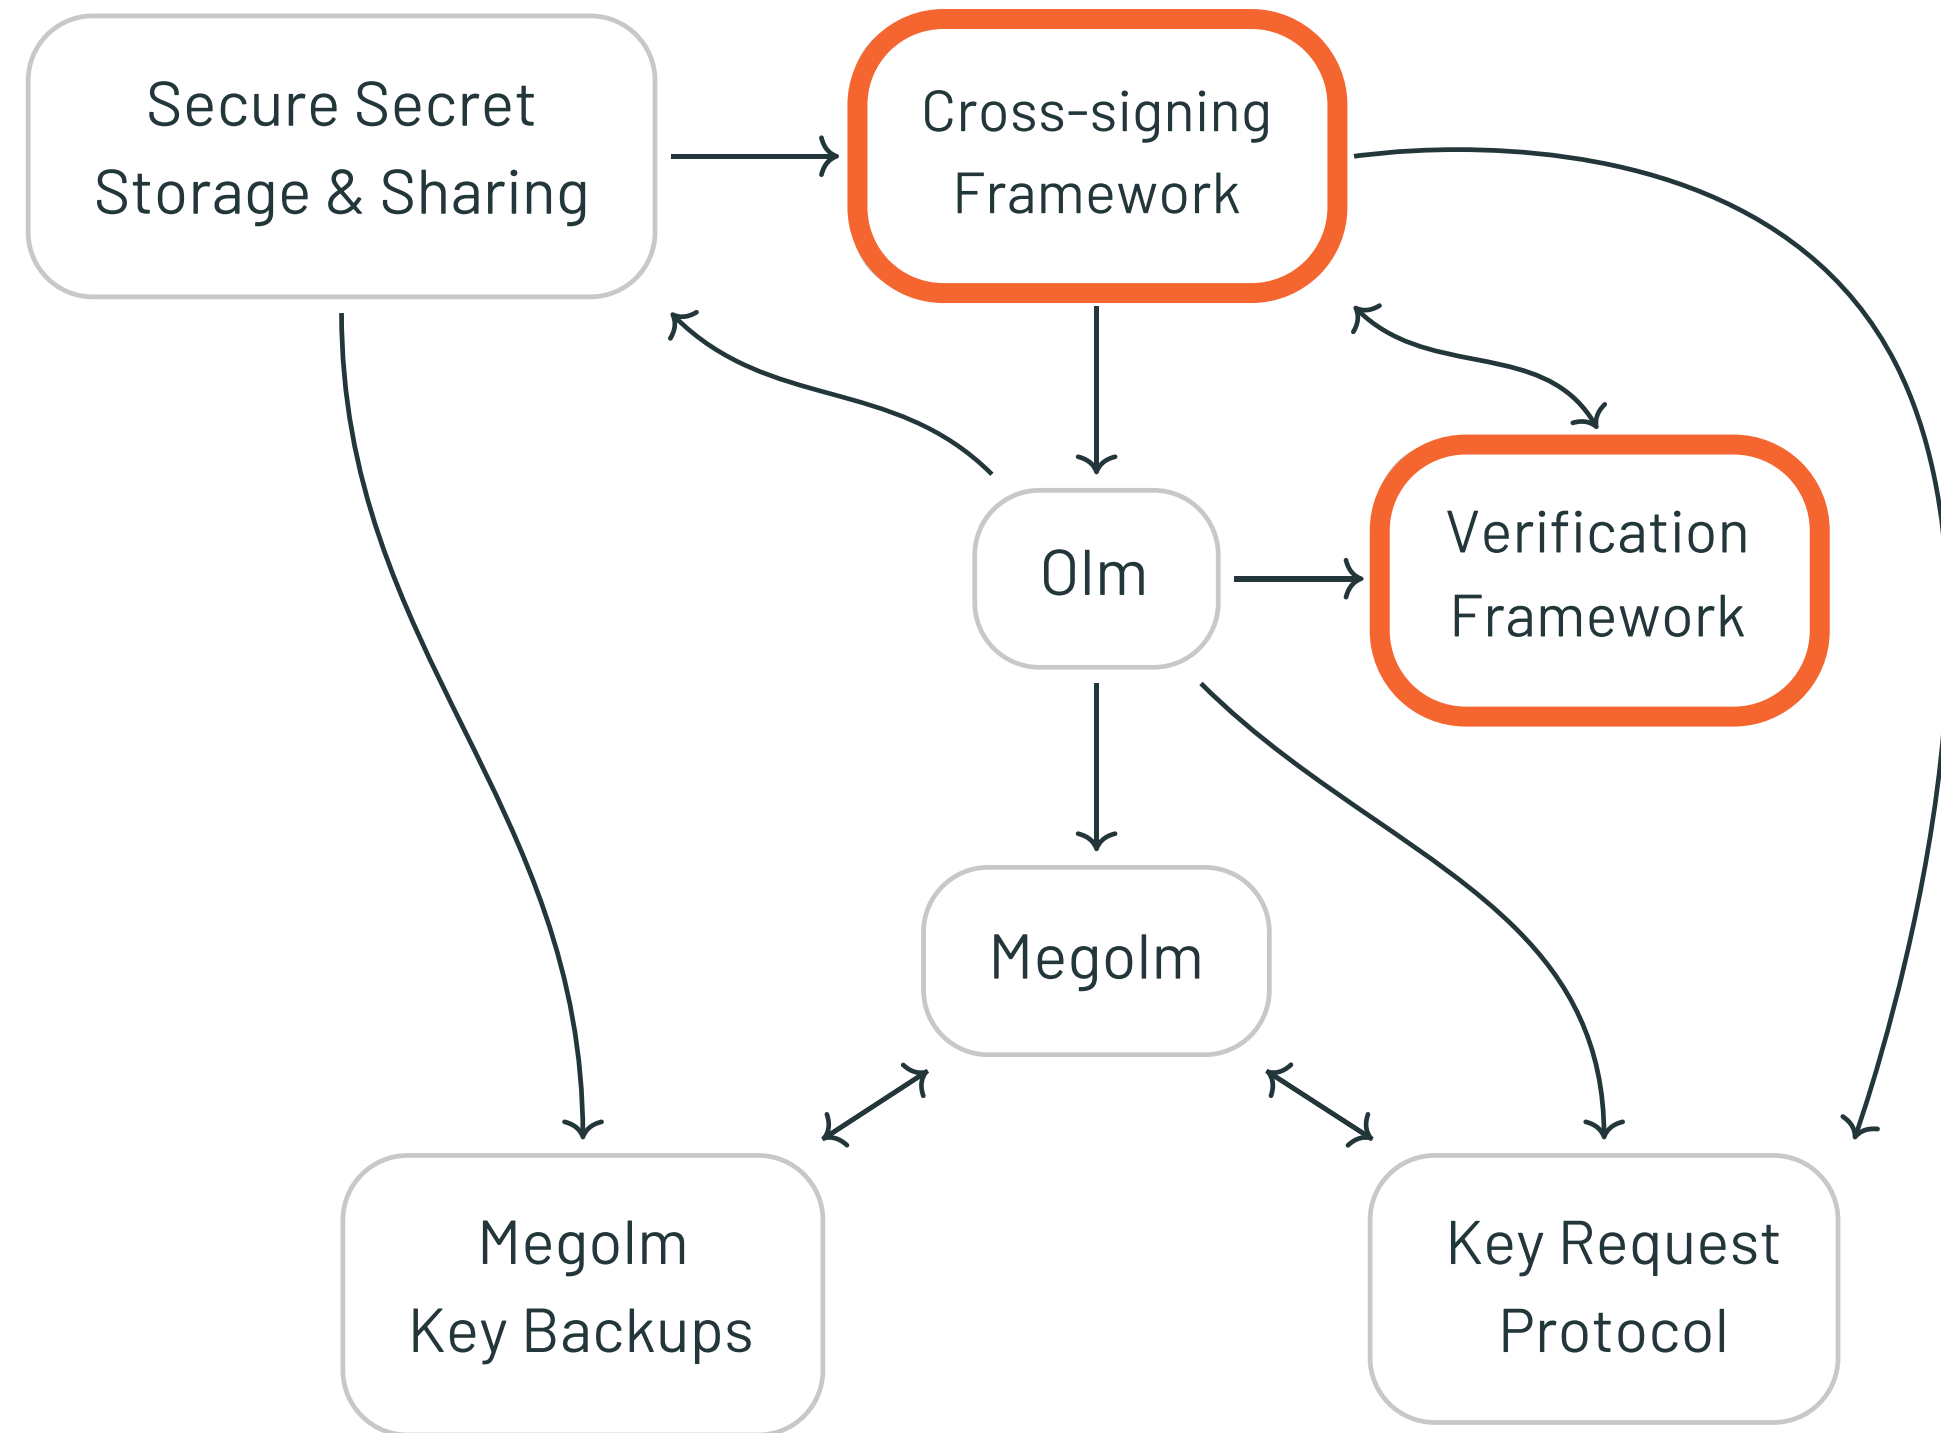


Cross-signing

- Cryptographic identities for users and their devices.

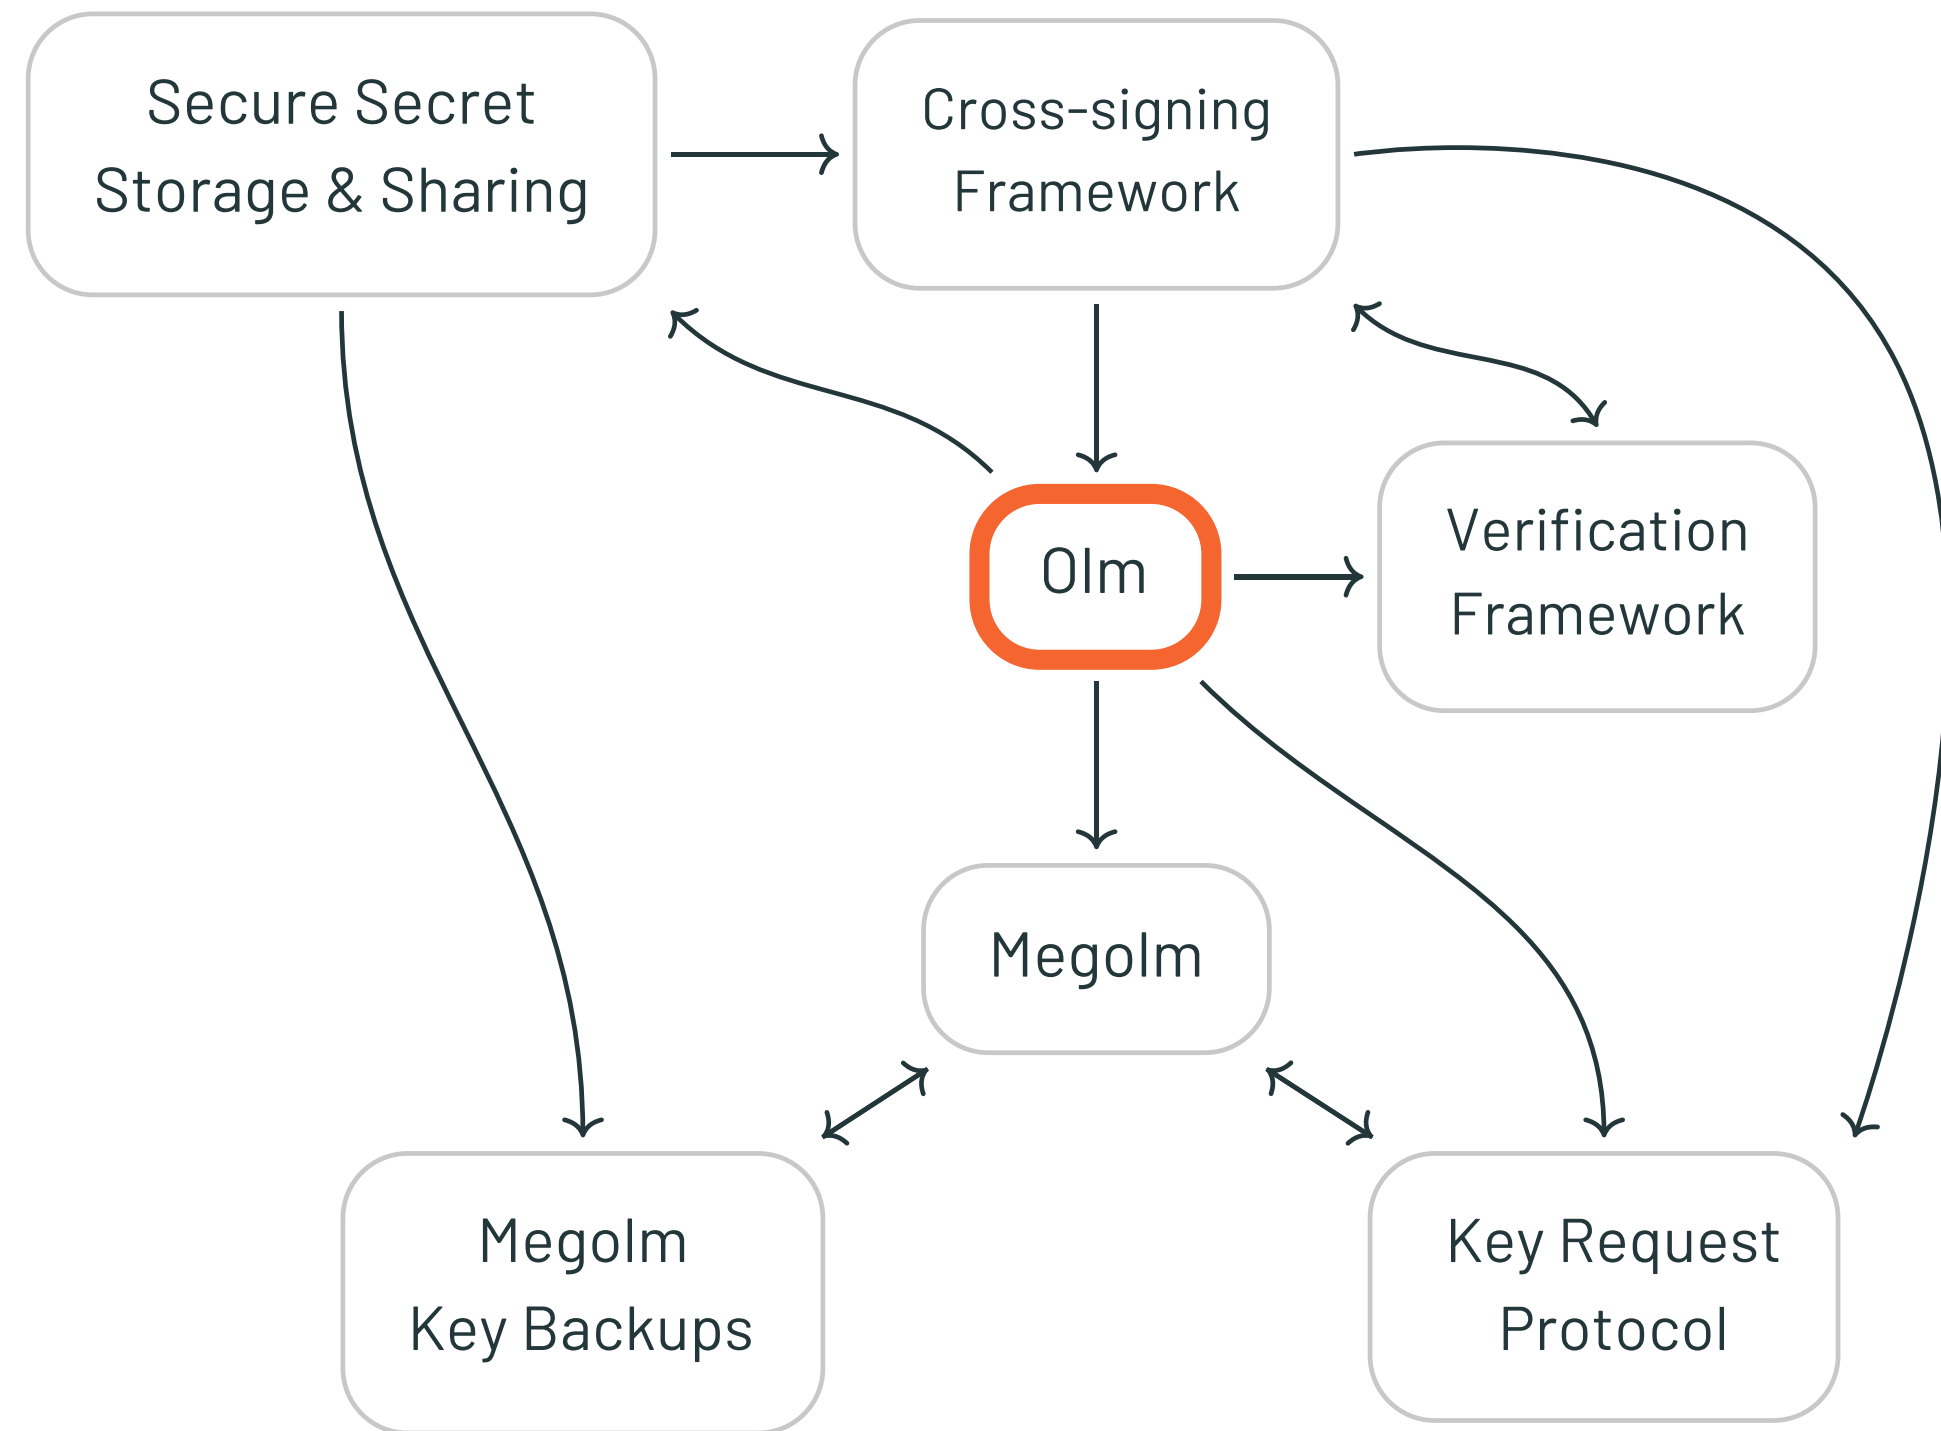
& Verification

- *Self-verification*
Users sign their own devices to indicate trust.
- *Cross-signing*
Users sign each other's identities.



Olm

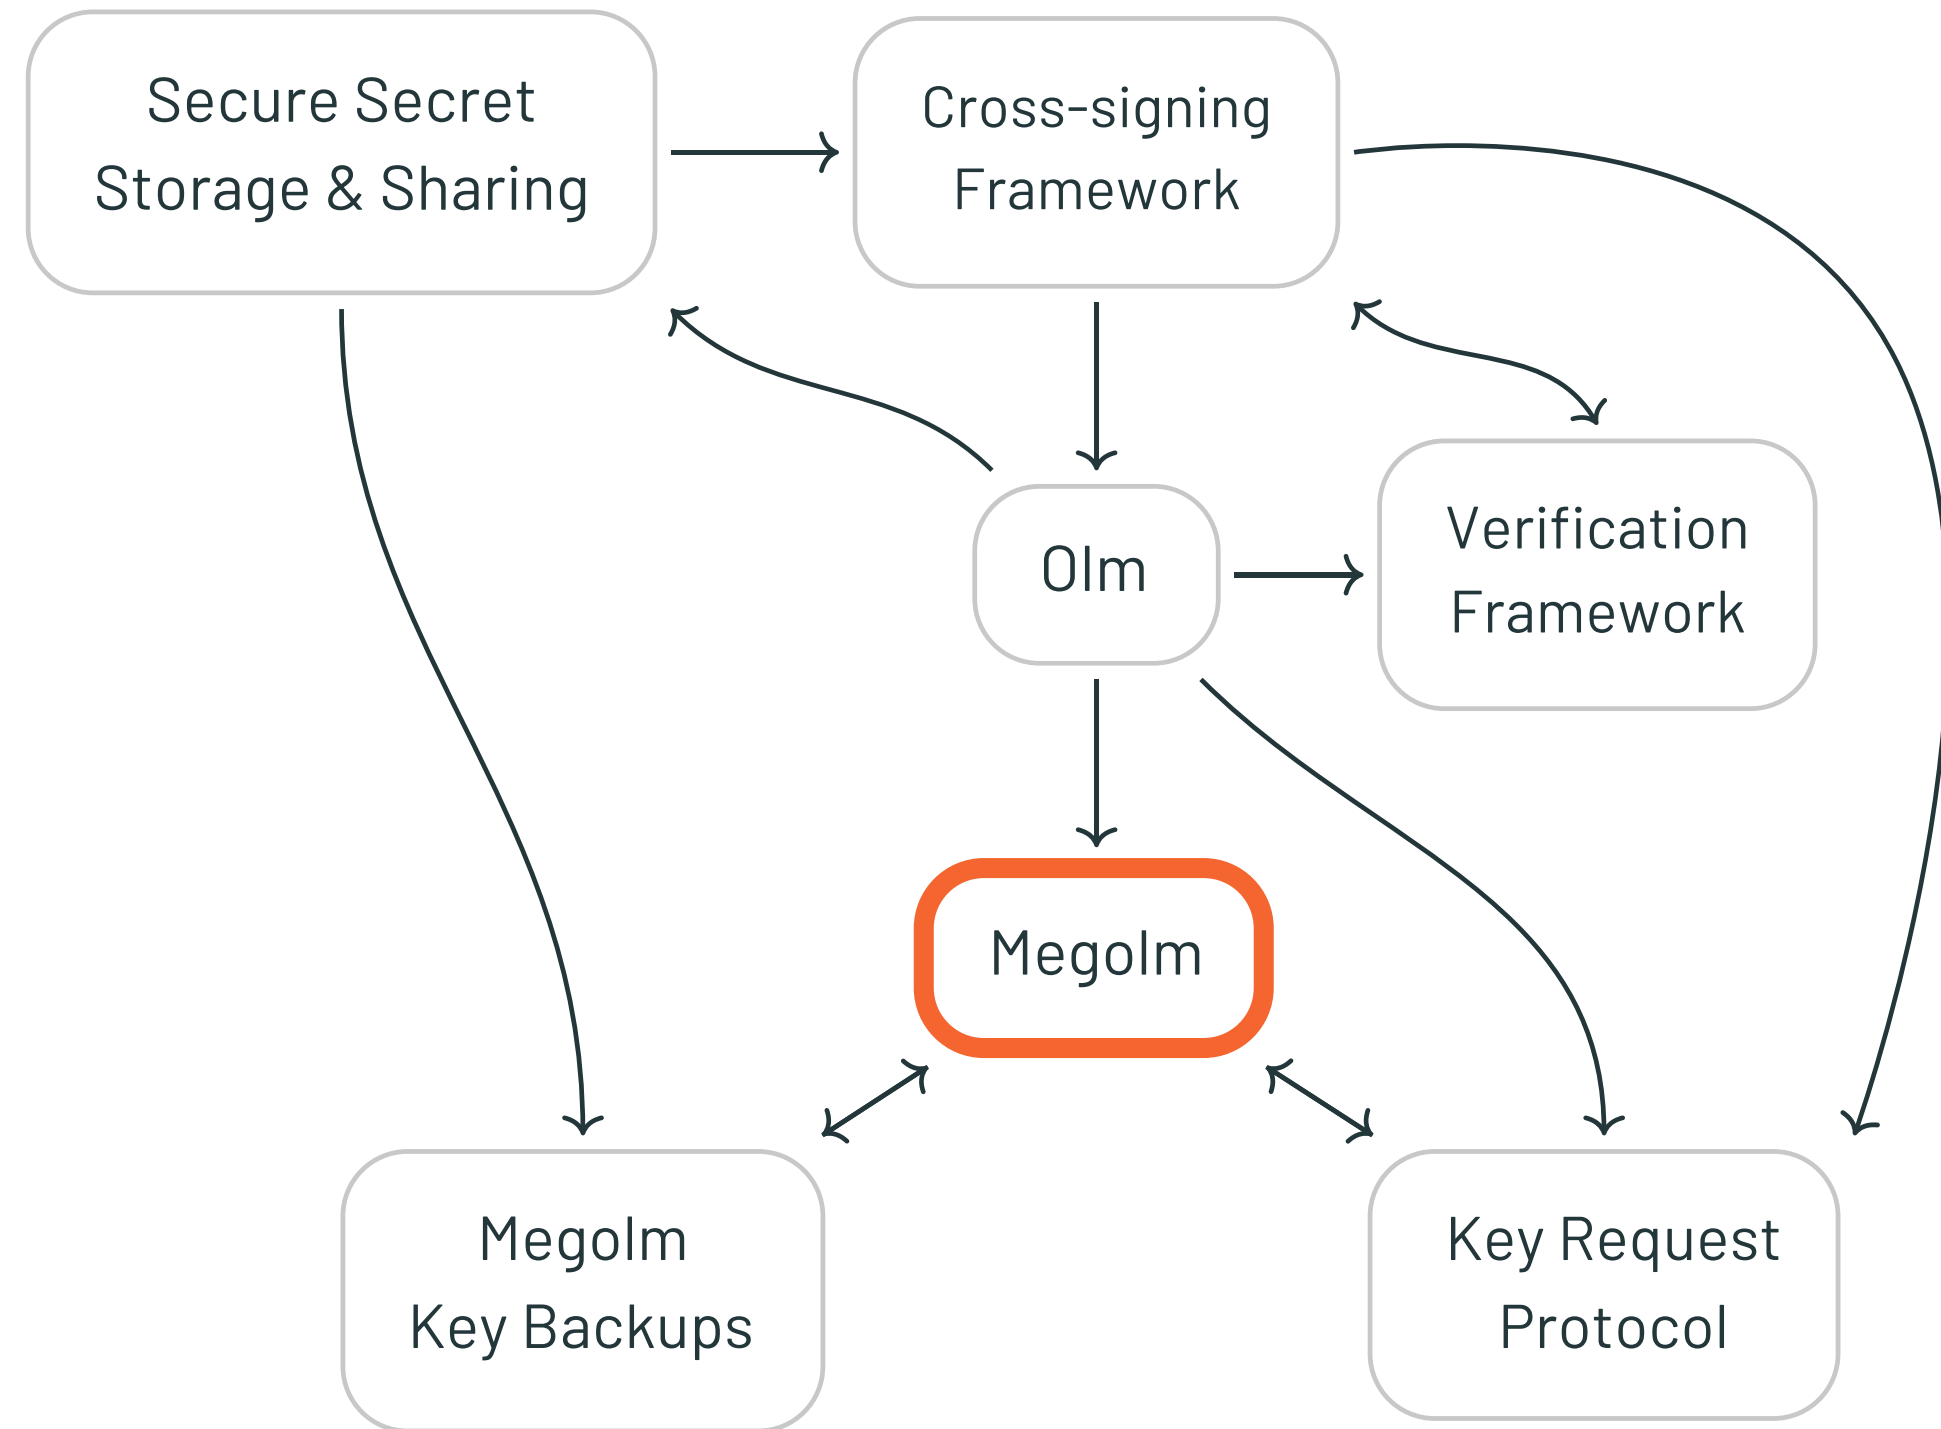
- Secure pairwise channels between devices.
- Initial key exchange via *Triple Diffie-Hellman*.
- Continuous key exchange via *Double Ratchet*.
- Signalling layer for other sub-protocols in Matrix.



Megolm

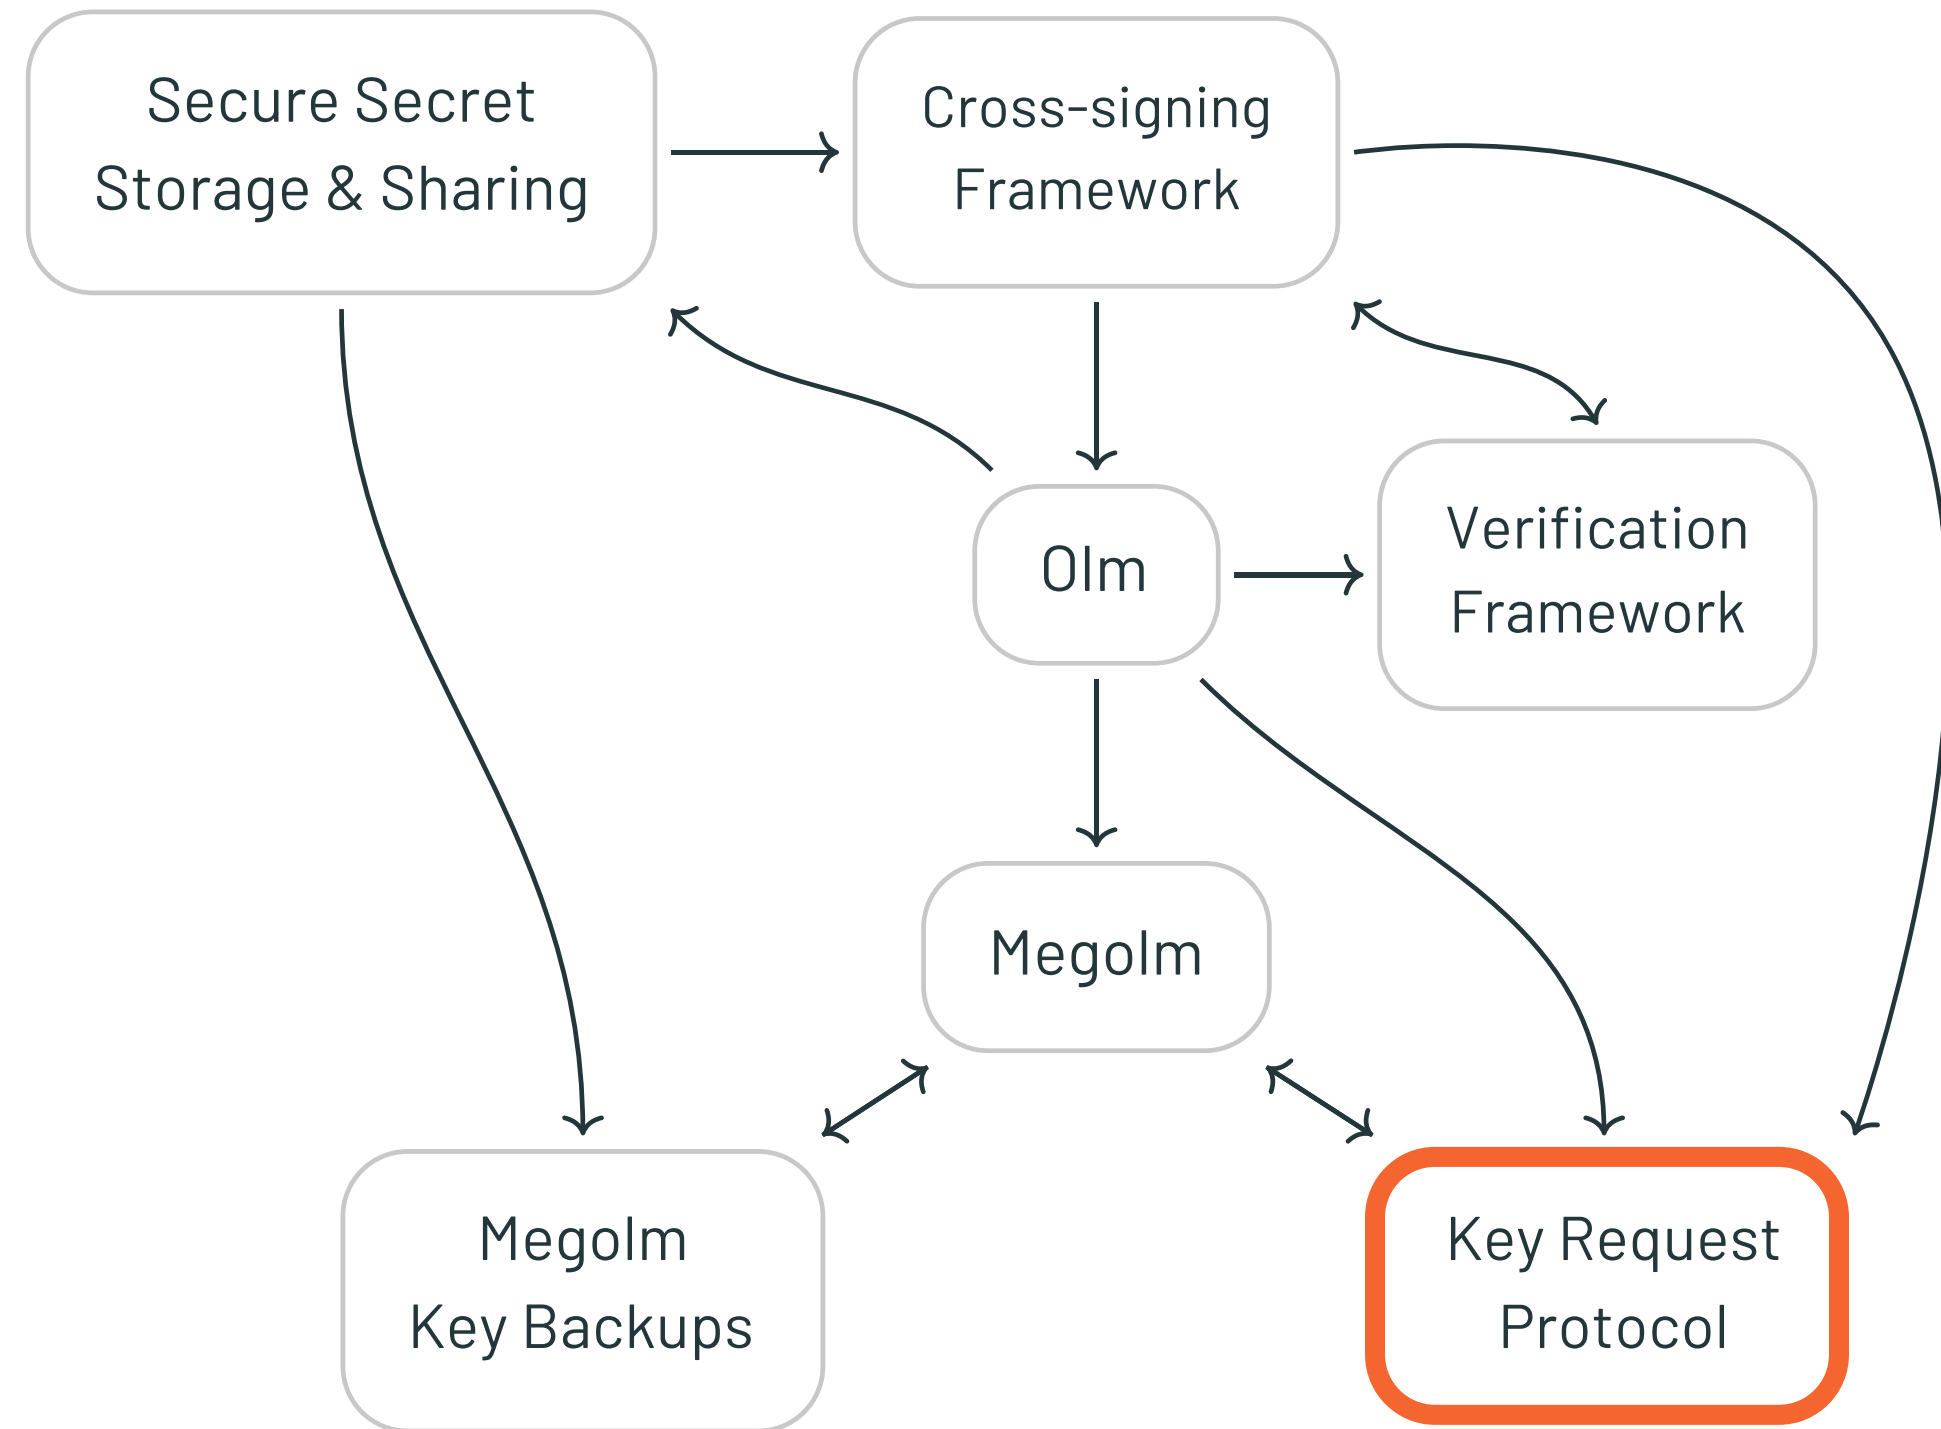
- Secure one-to-many channel.
- Symmetric ratchet for forward security.
- Session keys are distributed over Olm.
- Each sender maintains their own Megolm session.

↪ Compose together to form group chat.



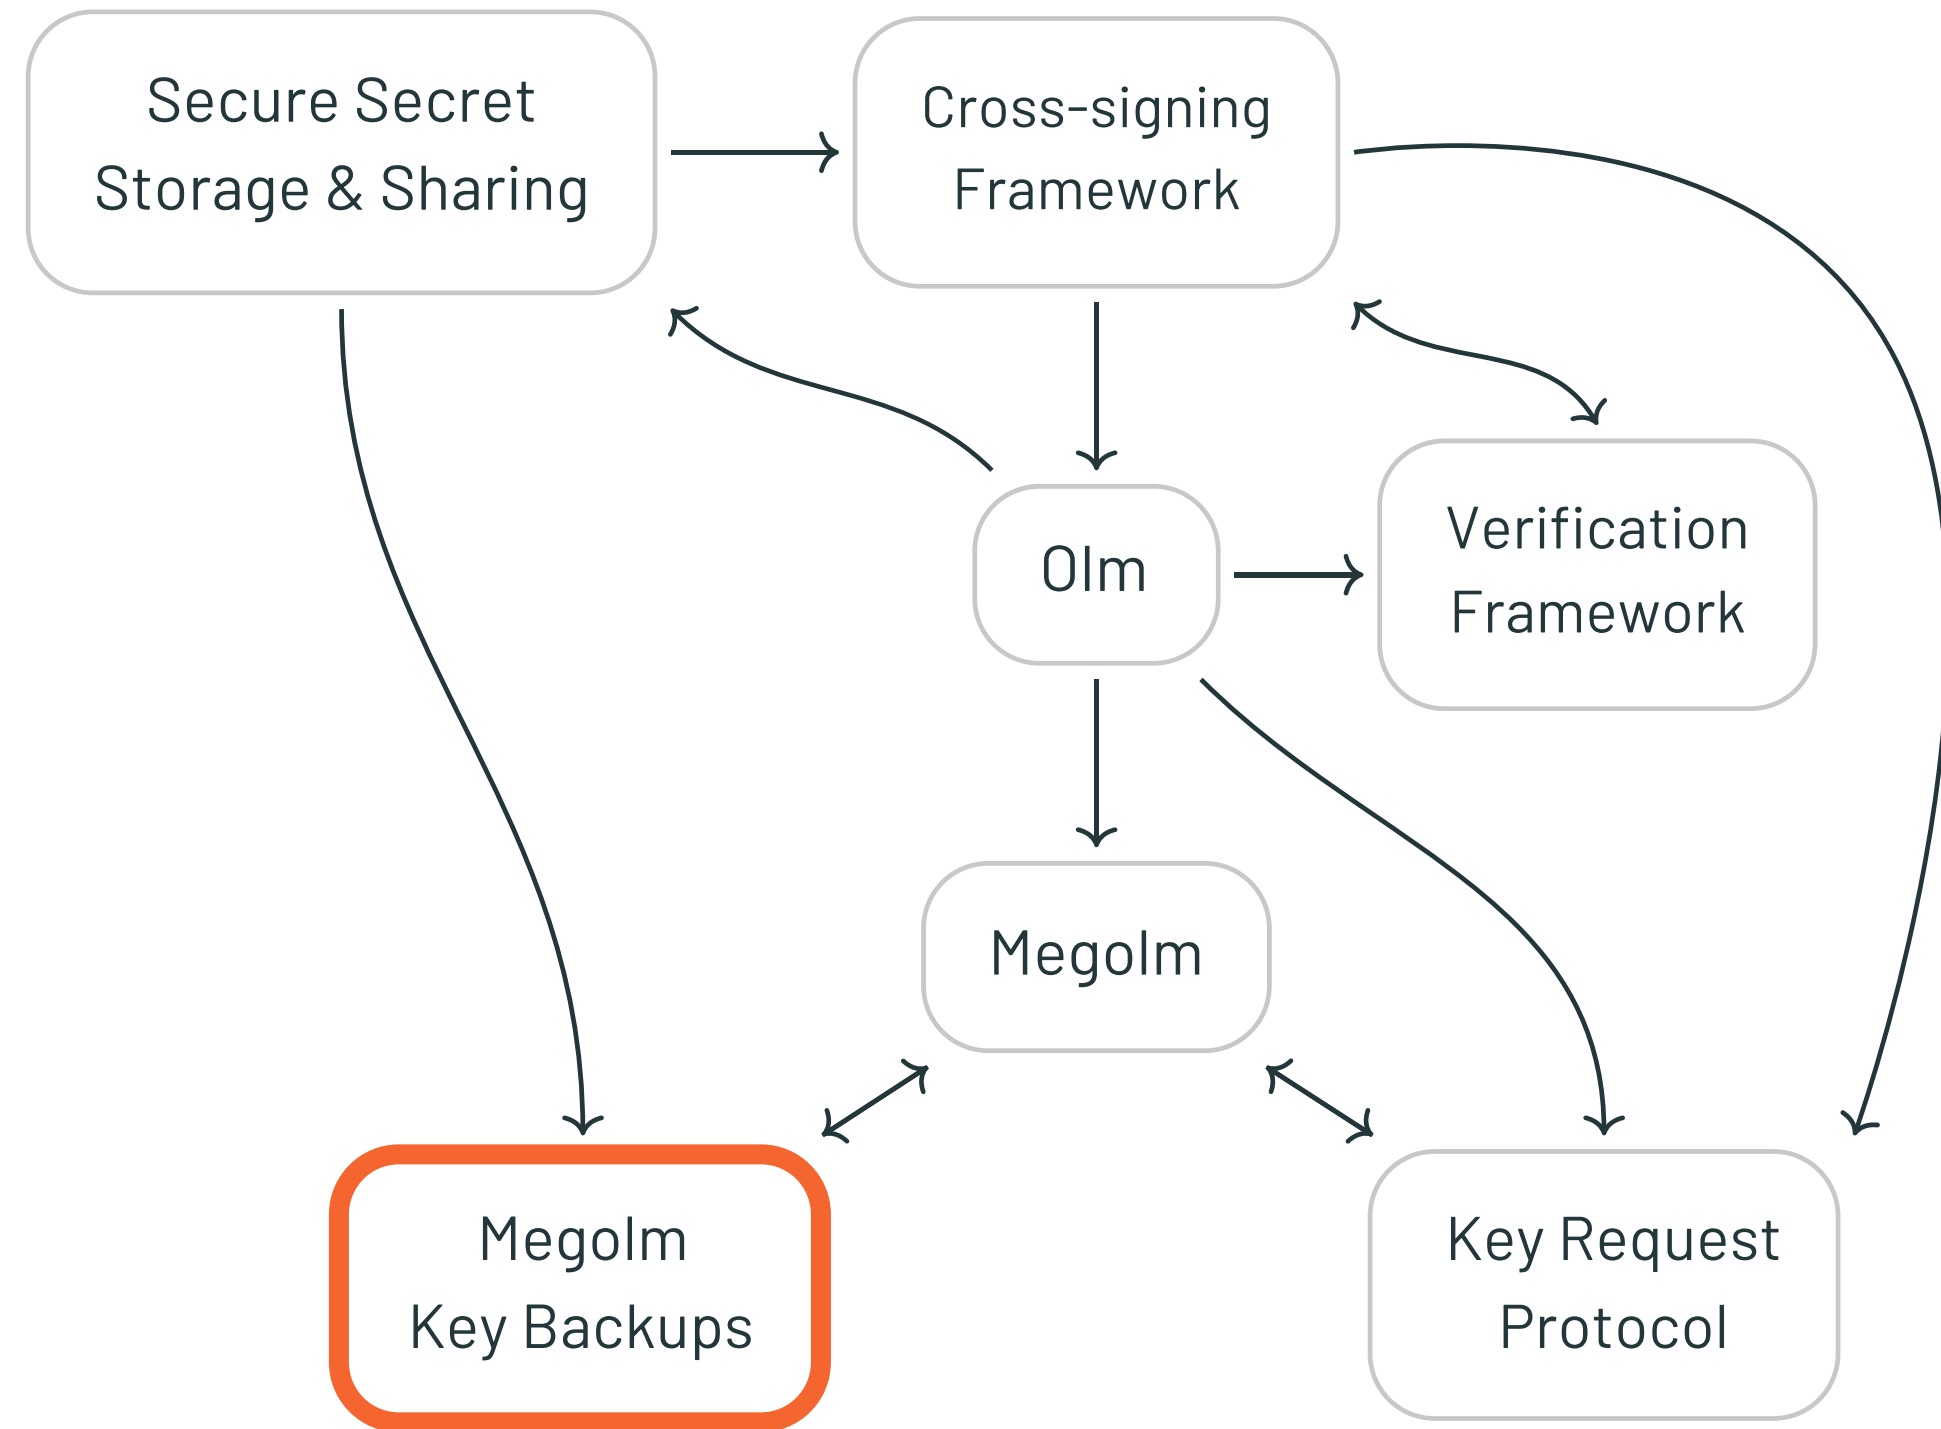
Key Request Protocol

- Request-response protocol.
- Allows devices to request and share keys between each other.
- Responding device must ensure requesting device is entitled to the keys.
- Keys are shared over Olm.



Megolm Key Backups

- Asynchronous alternative to Key Request protocol.
- Inbound Megolm sessions are encrypted and saved to the homeserver.
- Encrypt using a secret key shared between a user's devices.



Secure Secret Storage & Sharing

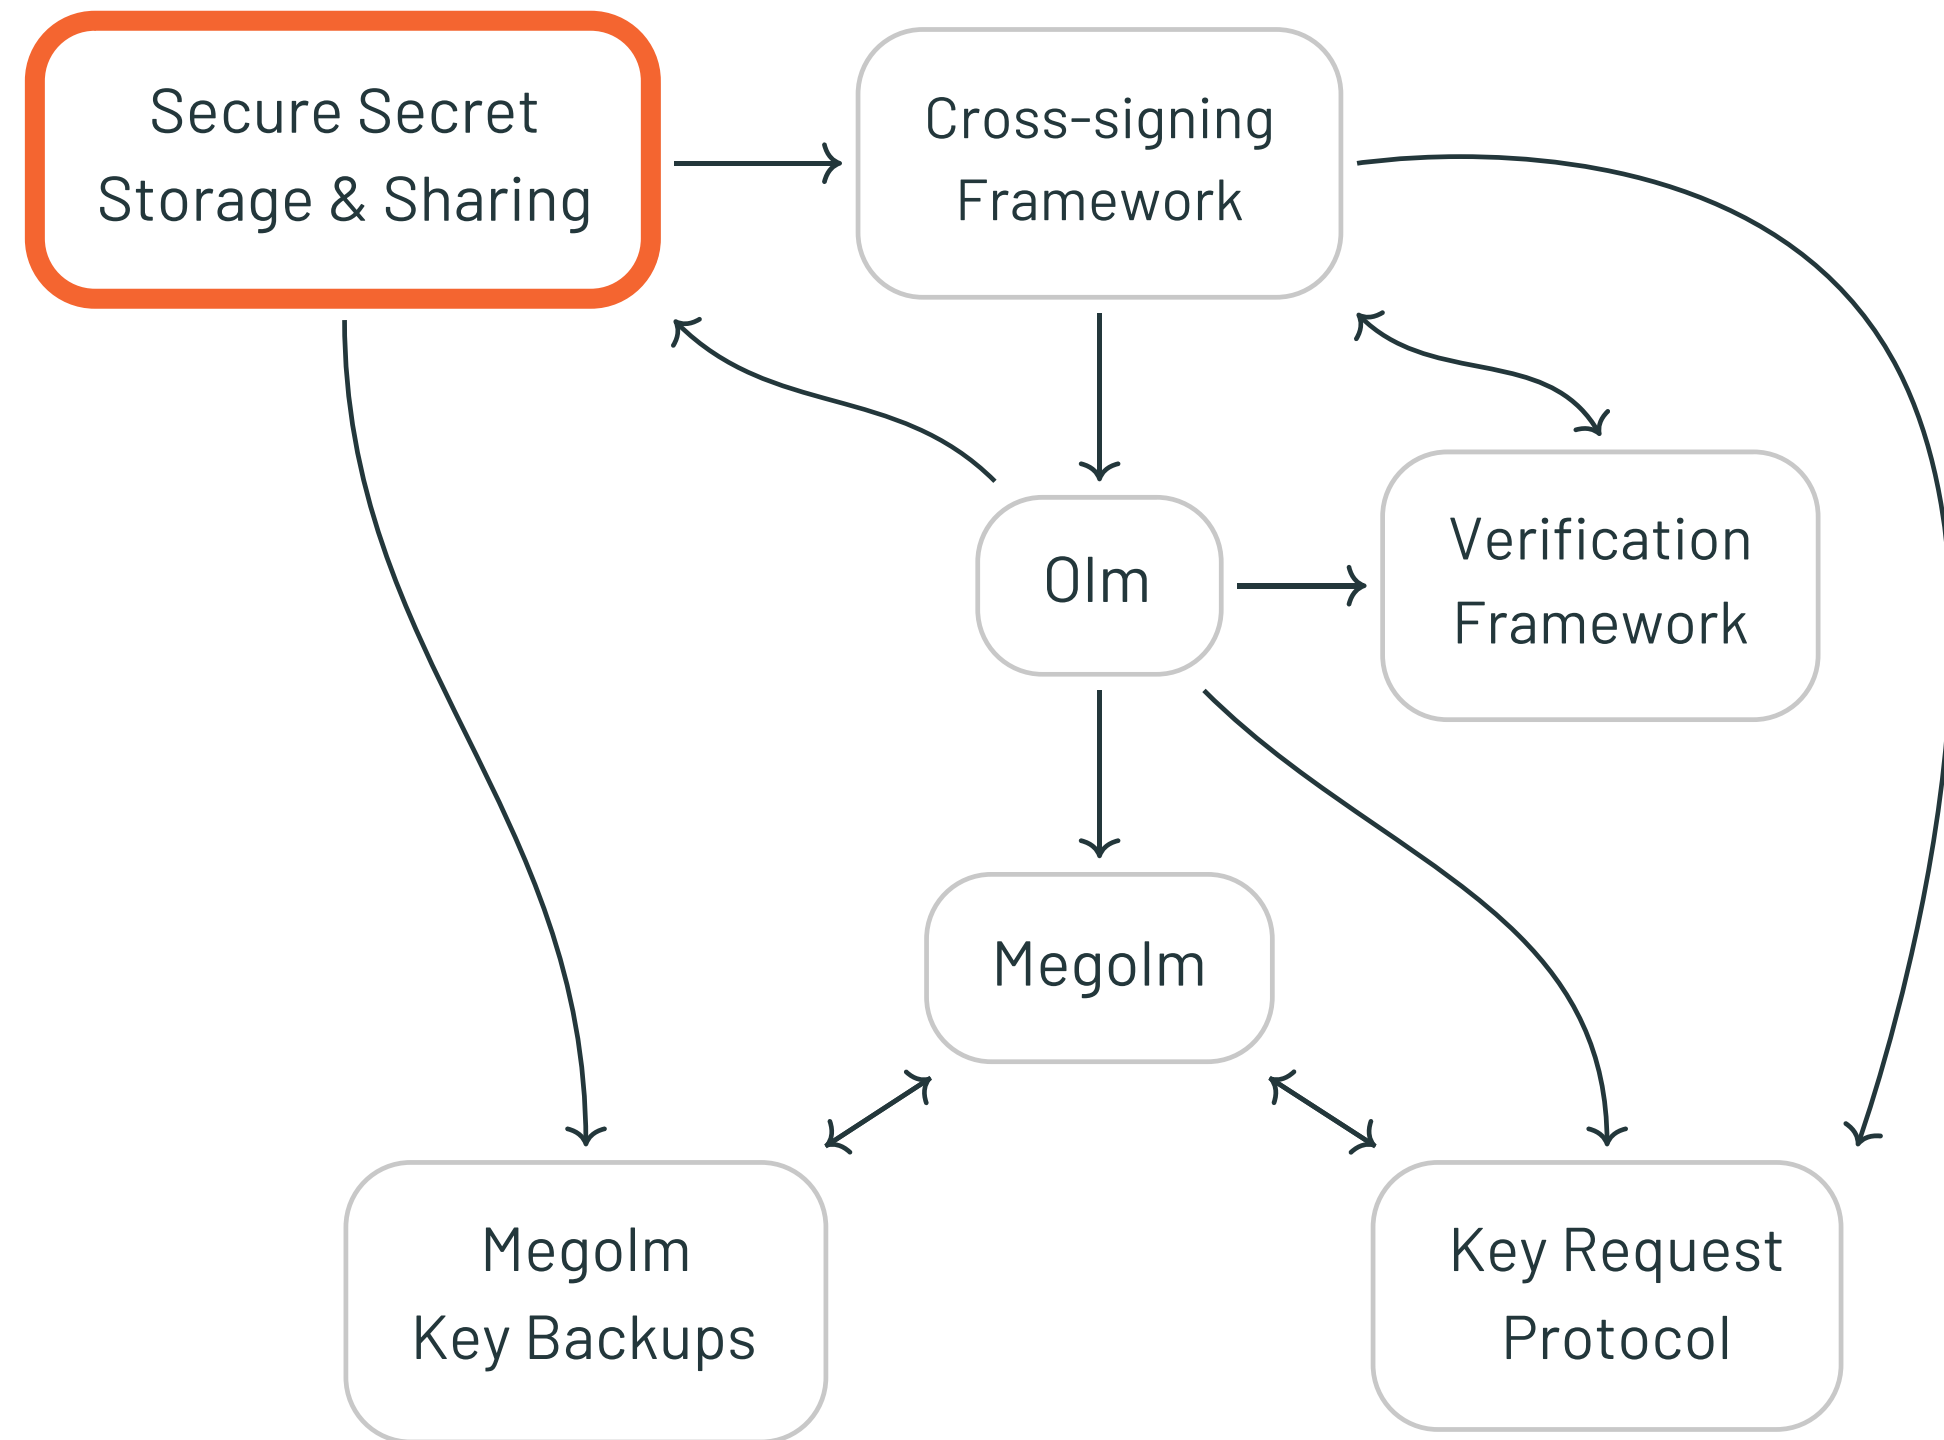
Backup, recover and share user-level secrets.

Secret Storage:

- Encrypt secrets and store on homeserver.
- Shared symmetric key (may be password-derived).

Secret Sharing:

- Use Olm to share secrets to newly verified devices.



Secure Secret Storage & Sharing

Backup, recover and share *user-level secrets*.

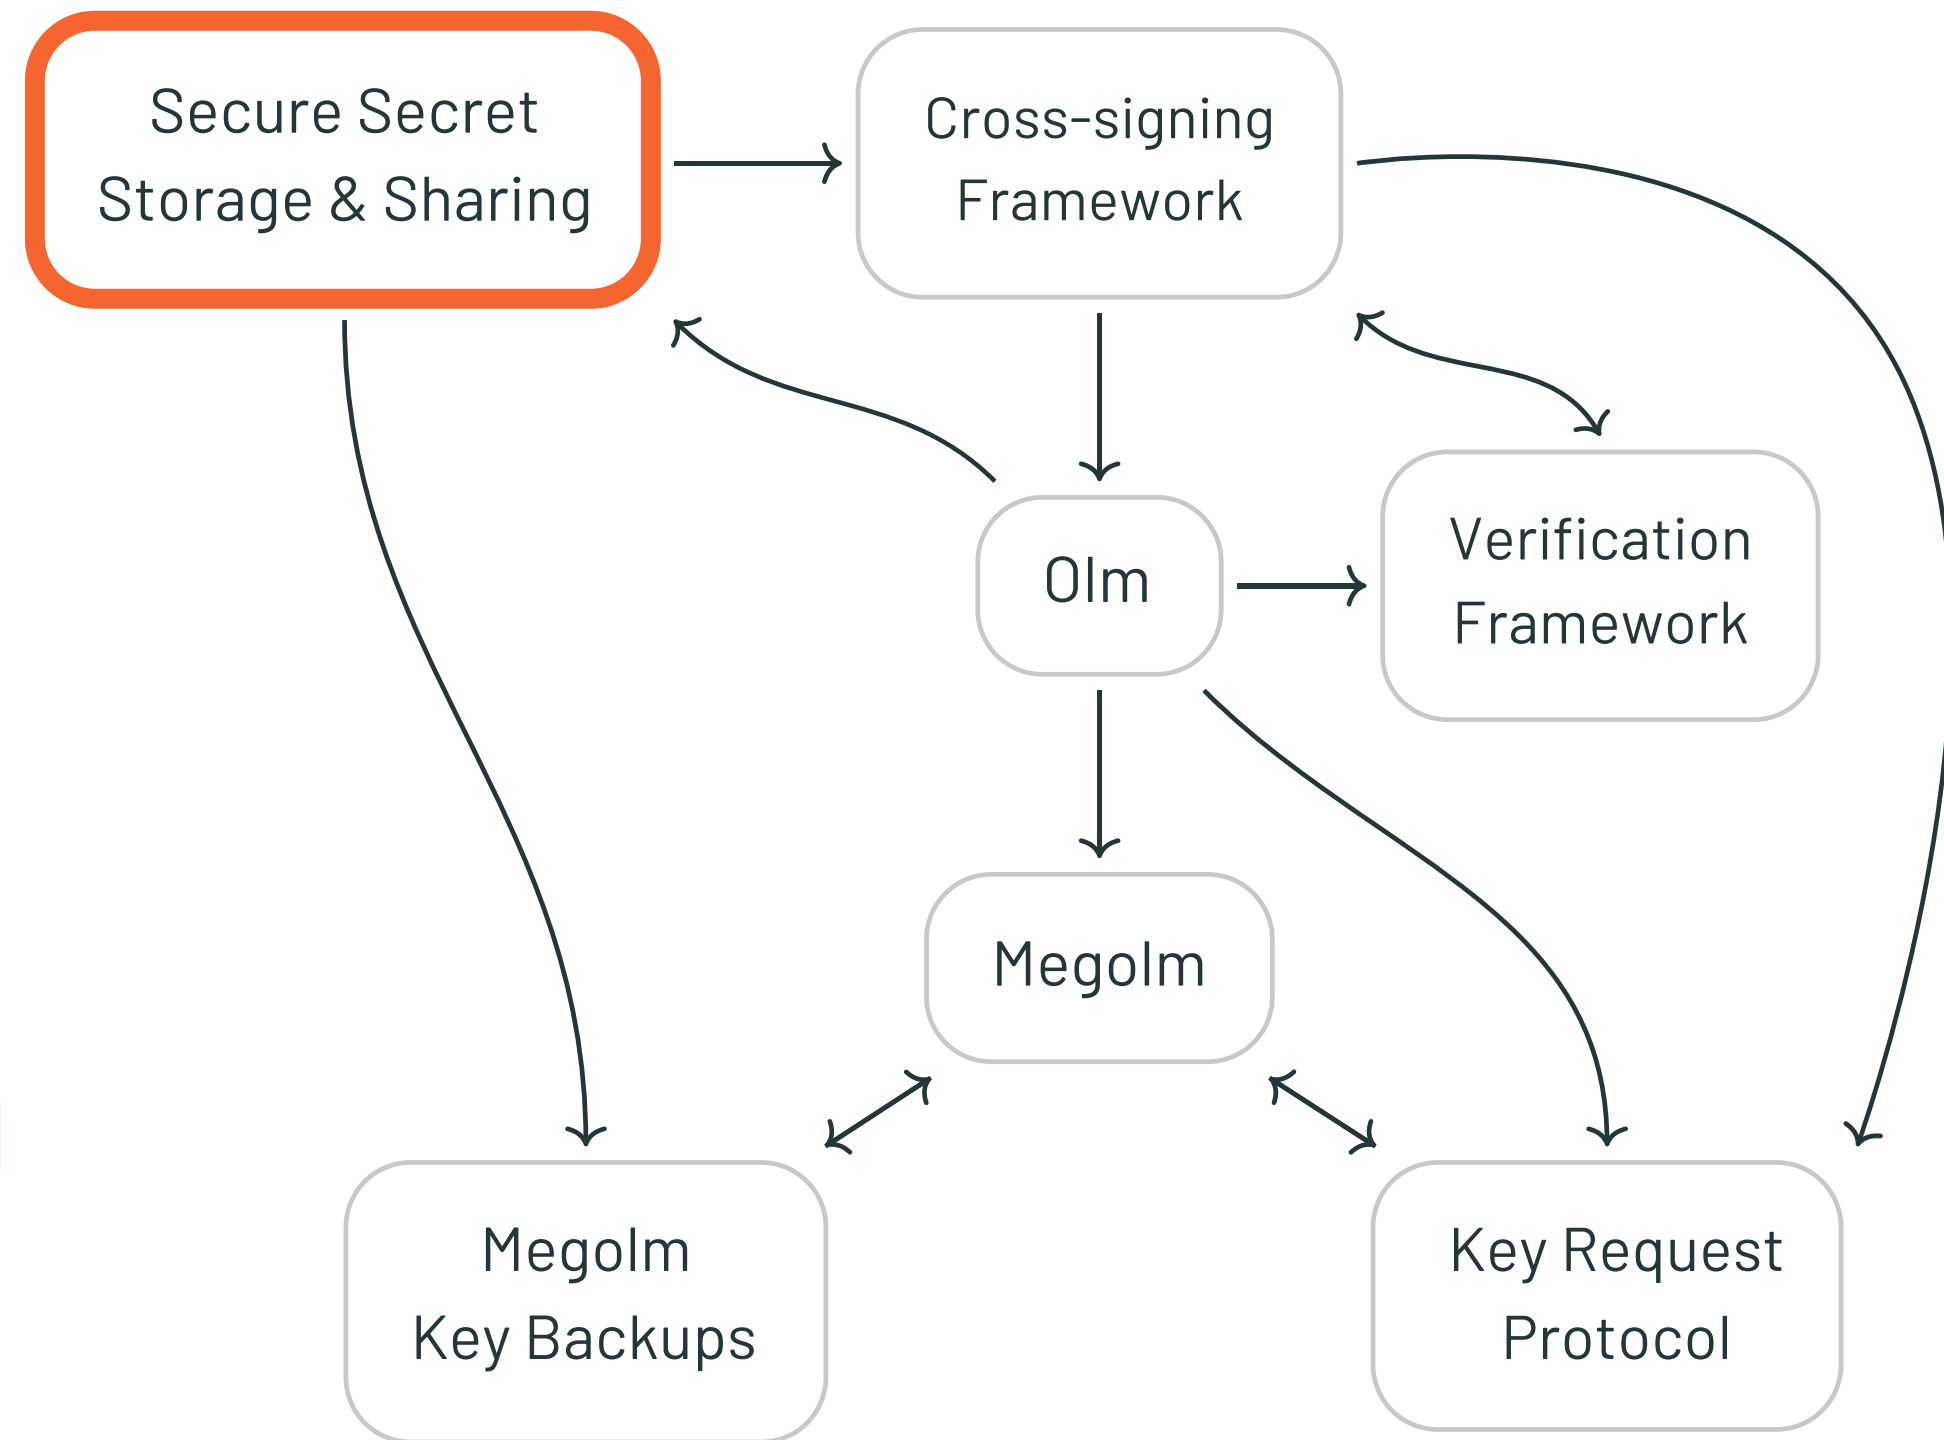
Secret Storage:

- Encrypt secrets and store on homeserver.
- Shared symmetric key (may be password-derived).

Secret Sharing:

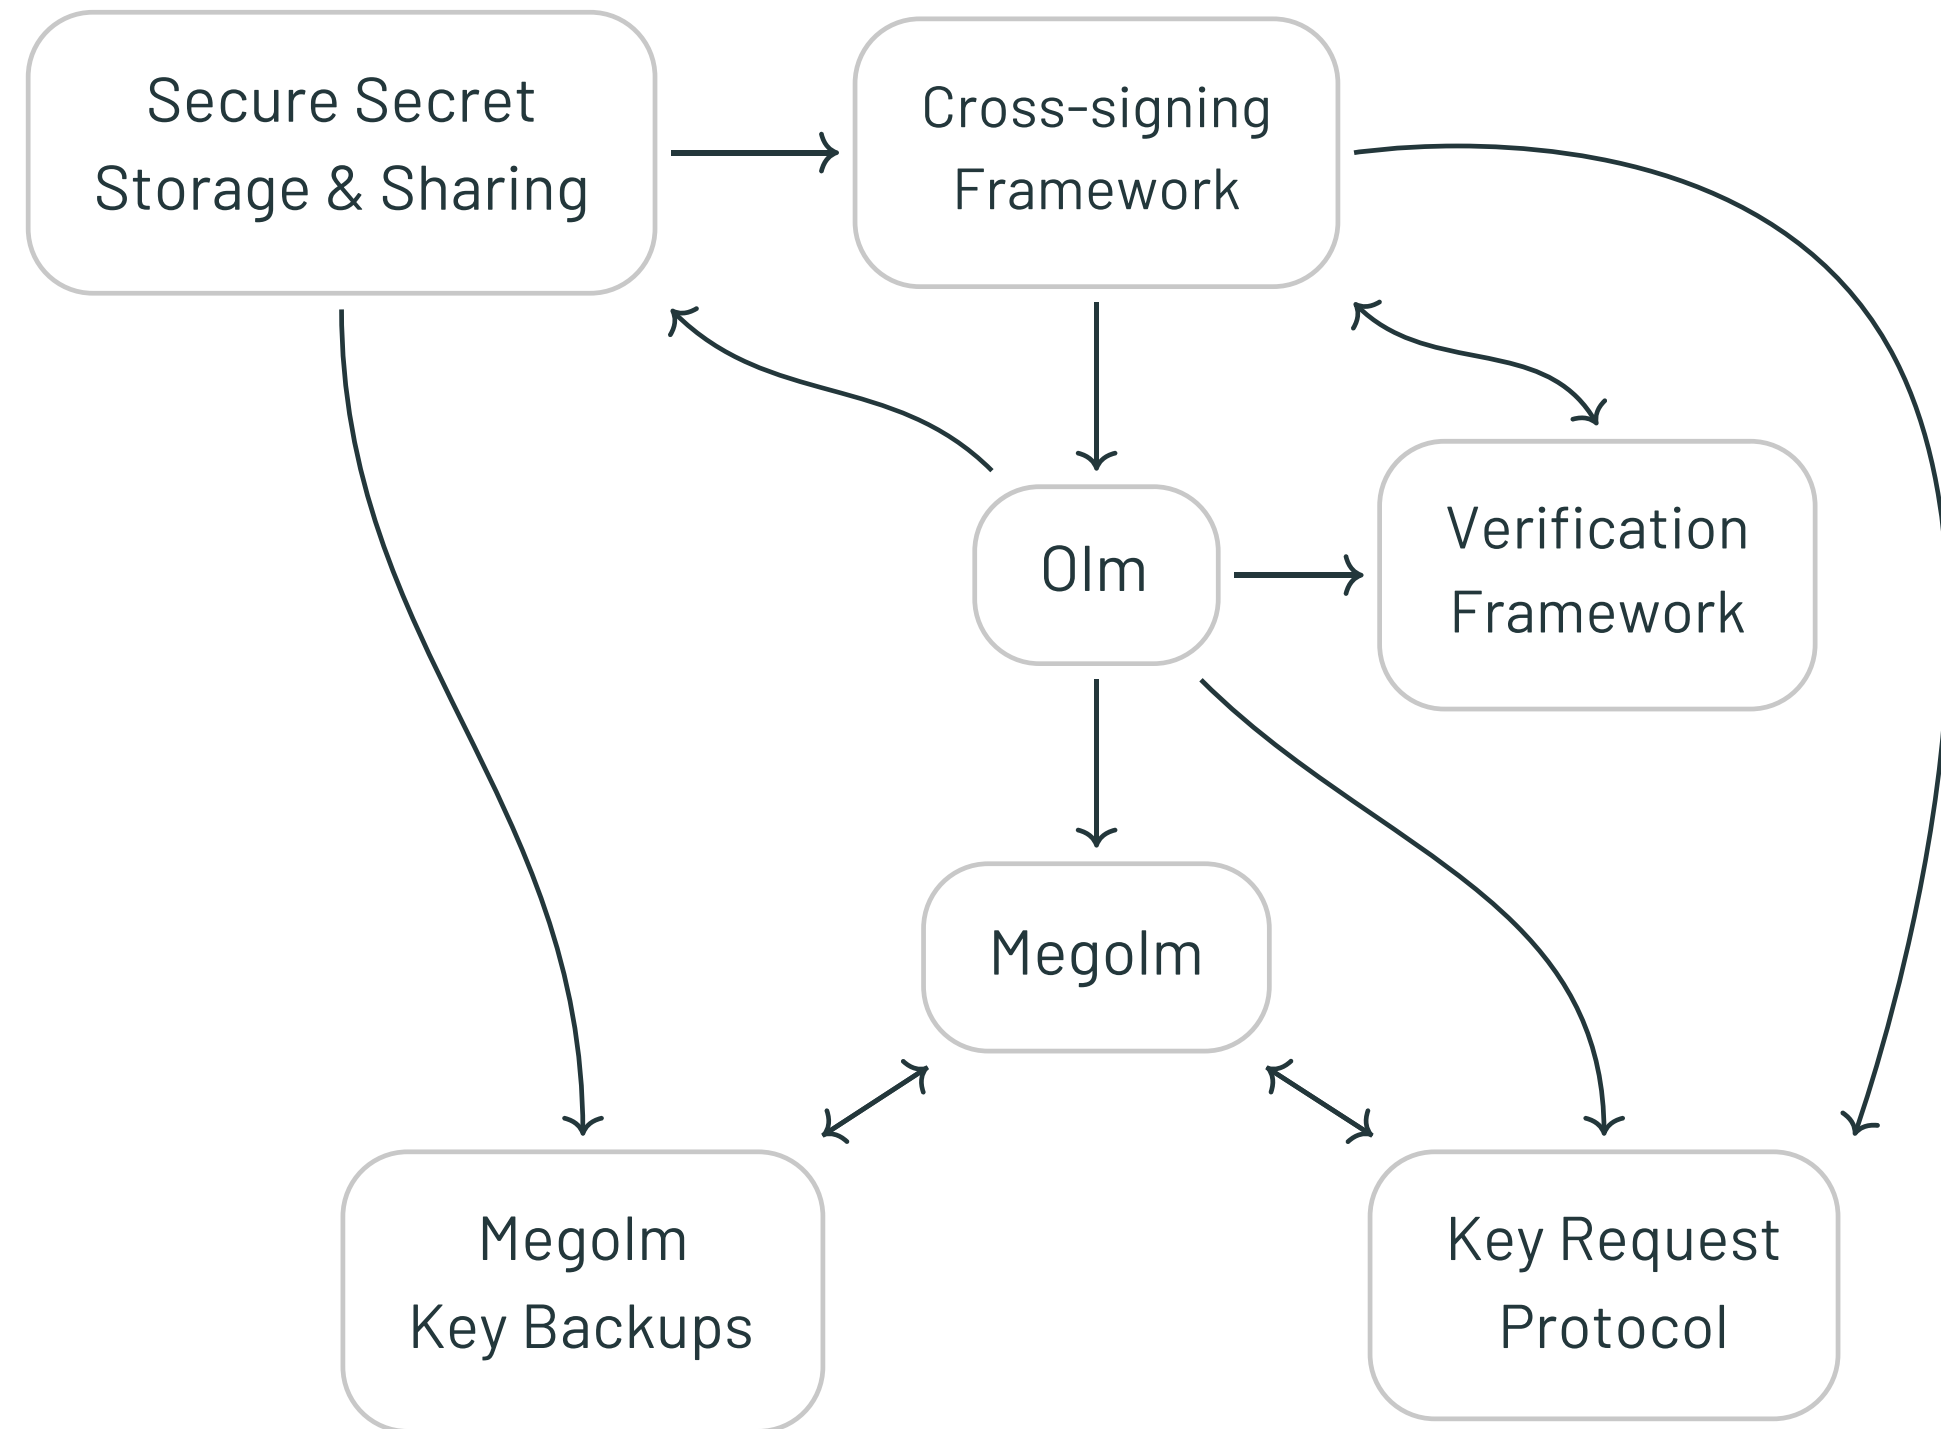
- Use Olm to share secrets to newly verified devices.

Root secret for a user's account



Modelling Matrix & Finding Attacks

We found many of these attacks while formalising each sub-protocol as part of the modelling work.

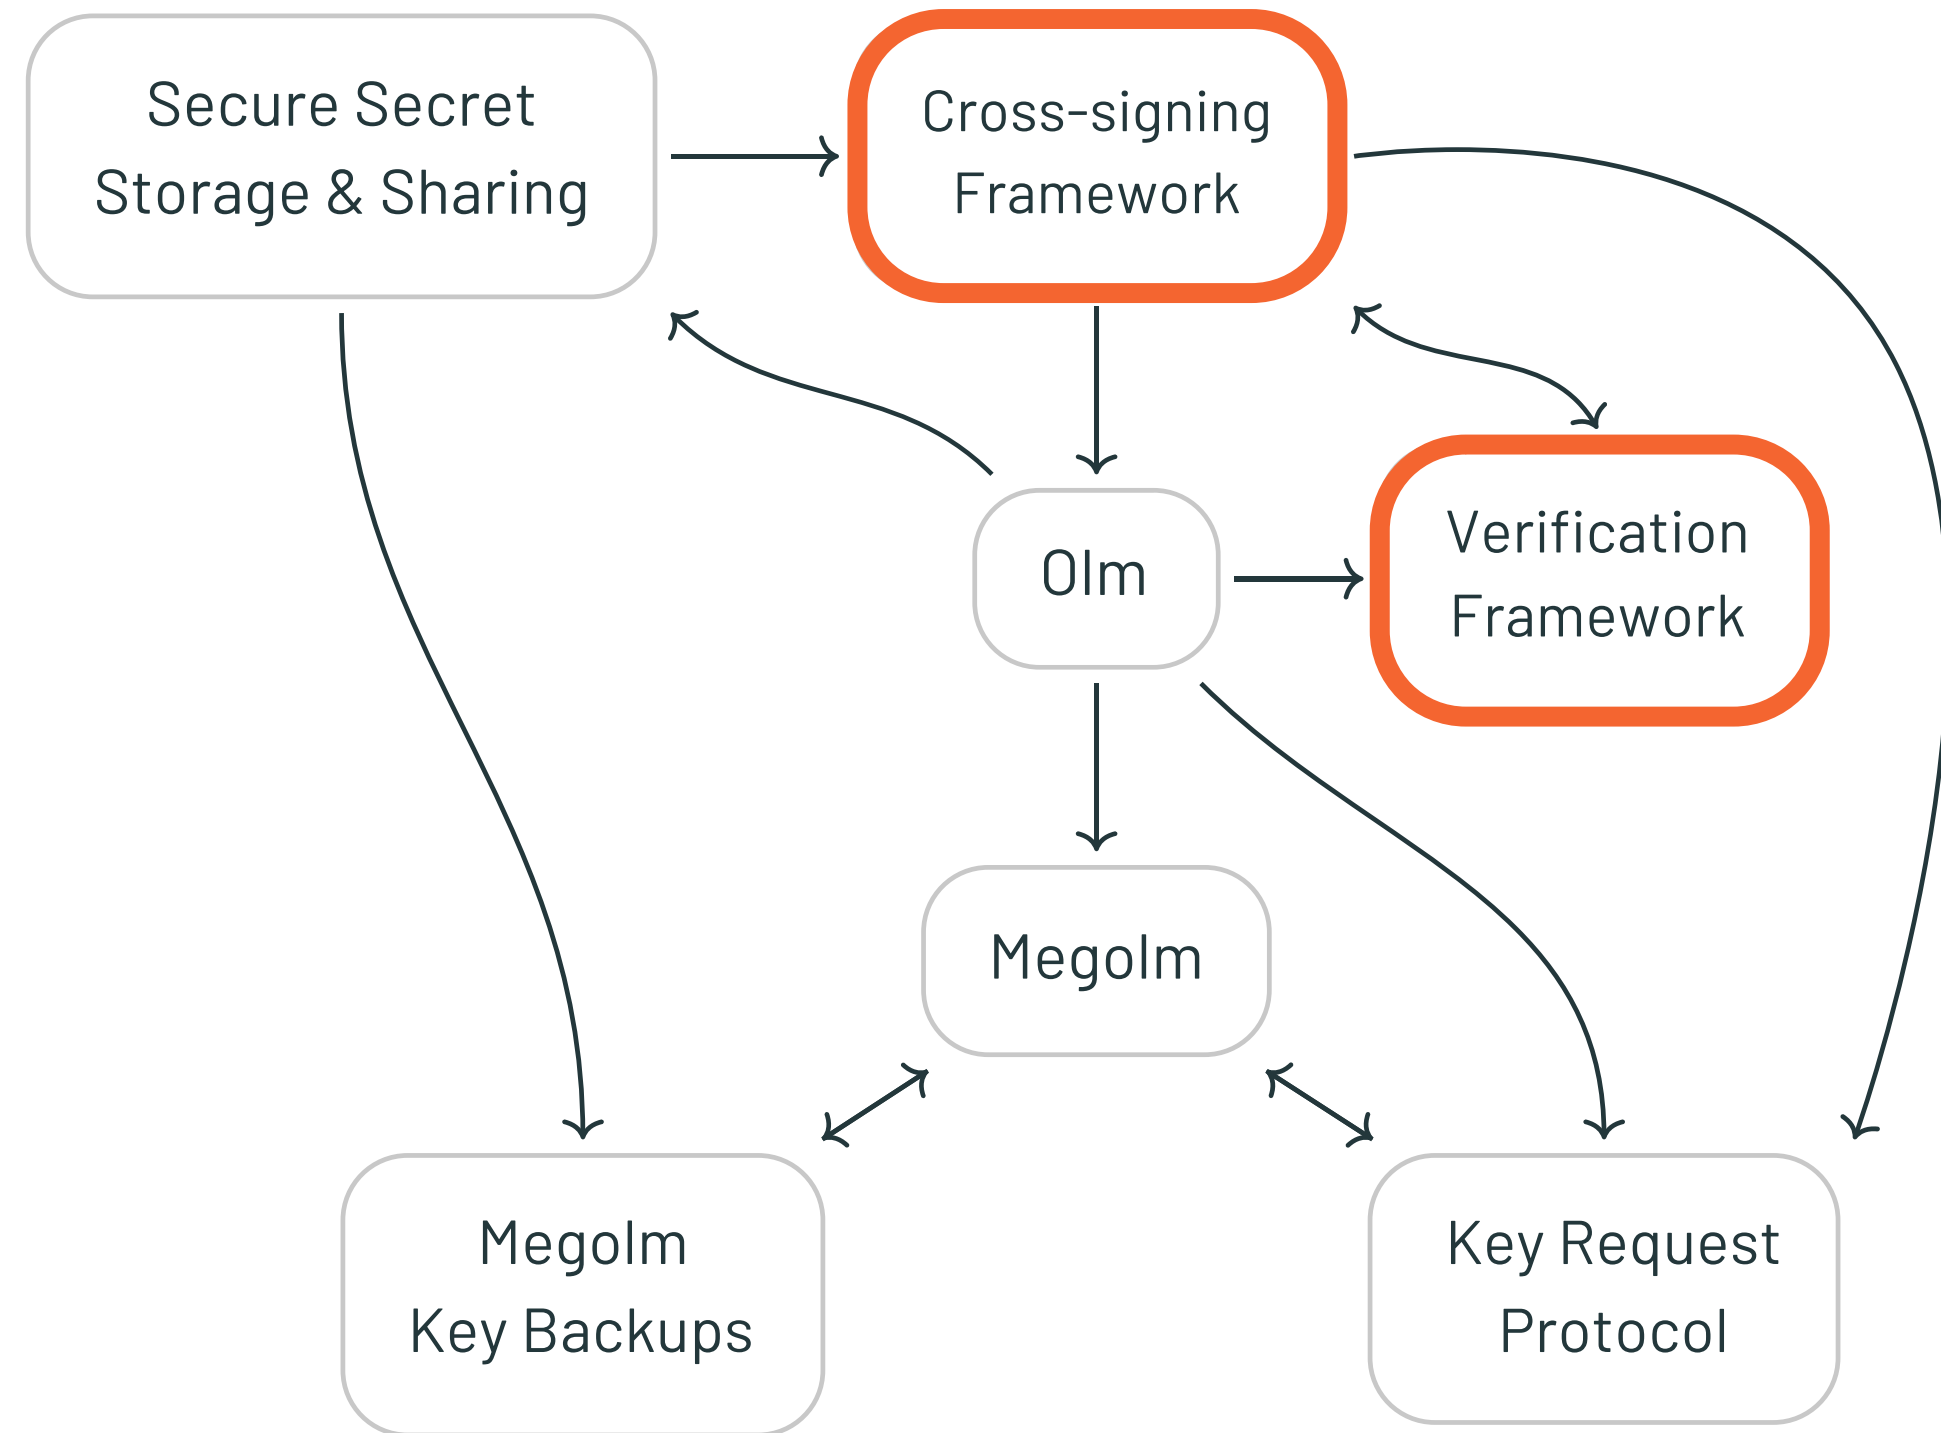


Cross-signing

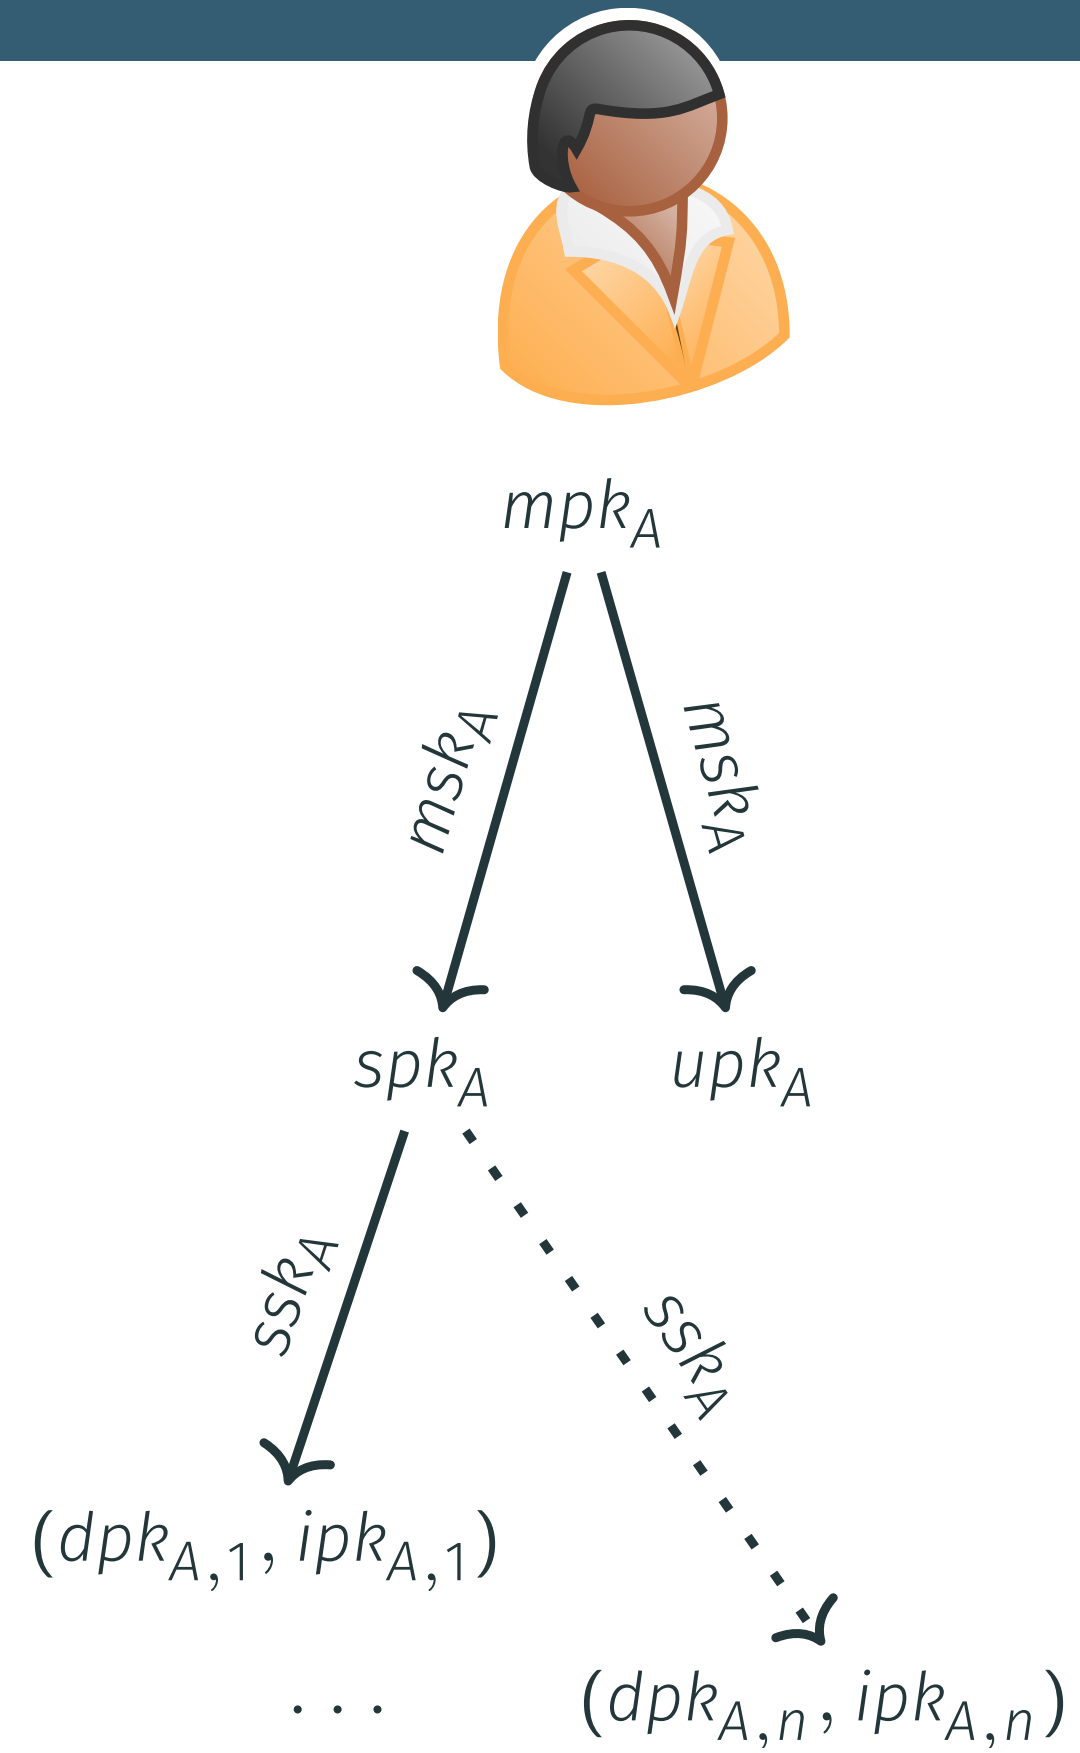
- Cryptographic identities for users and their devices.

& Verification

- *Self-verification*
Users sign their own devices to indicate trust.
- *Cross-signing*
Users sign each other's identities.

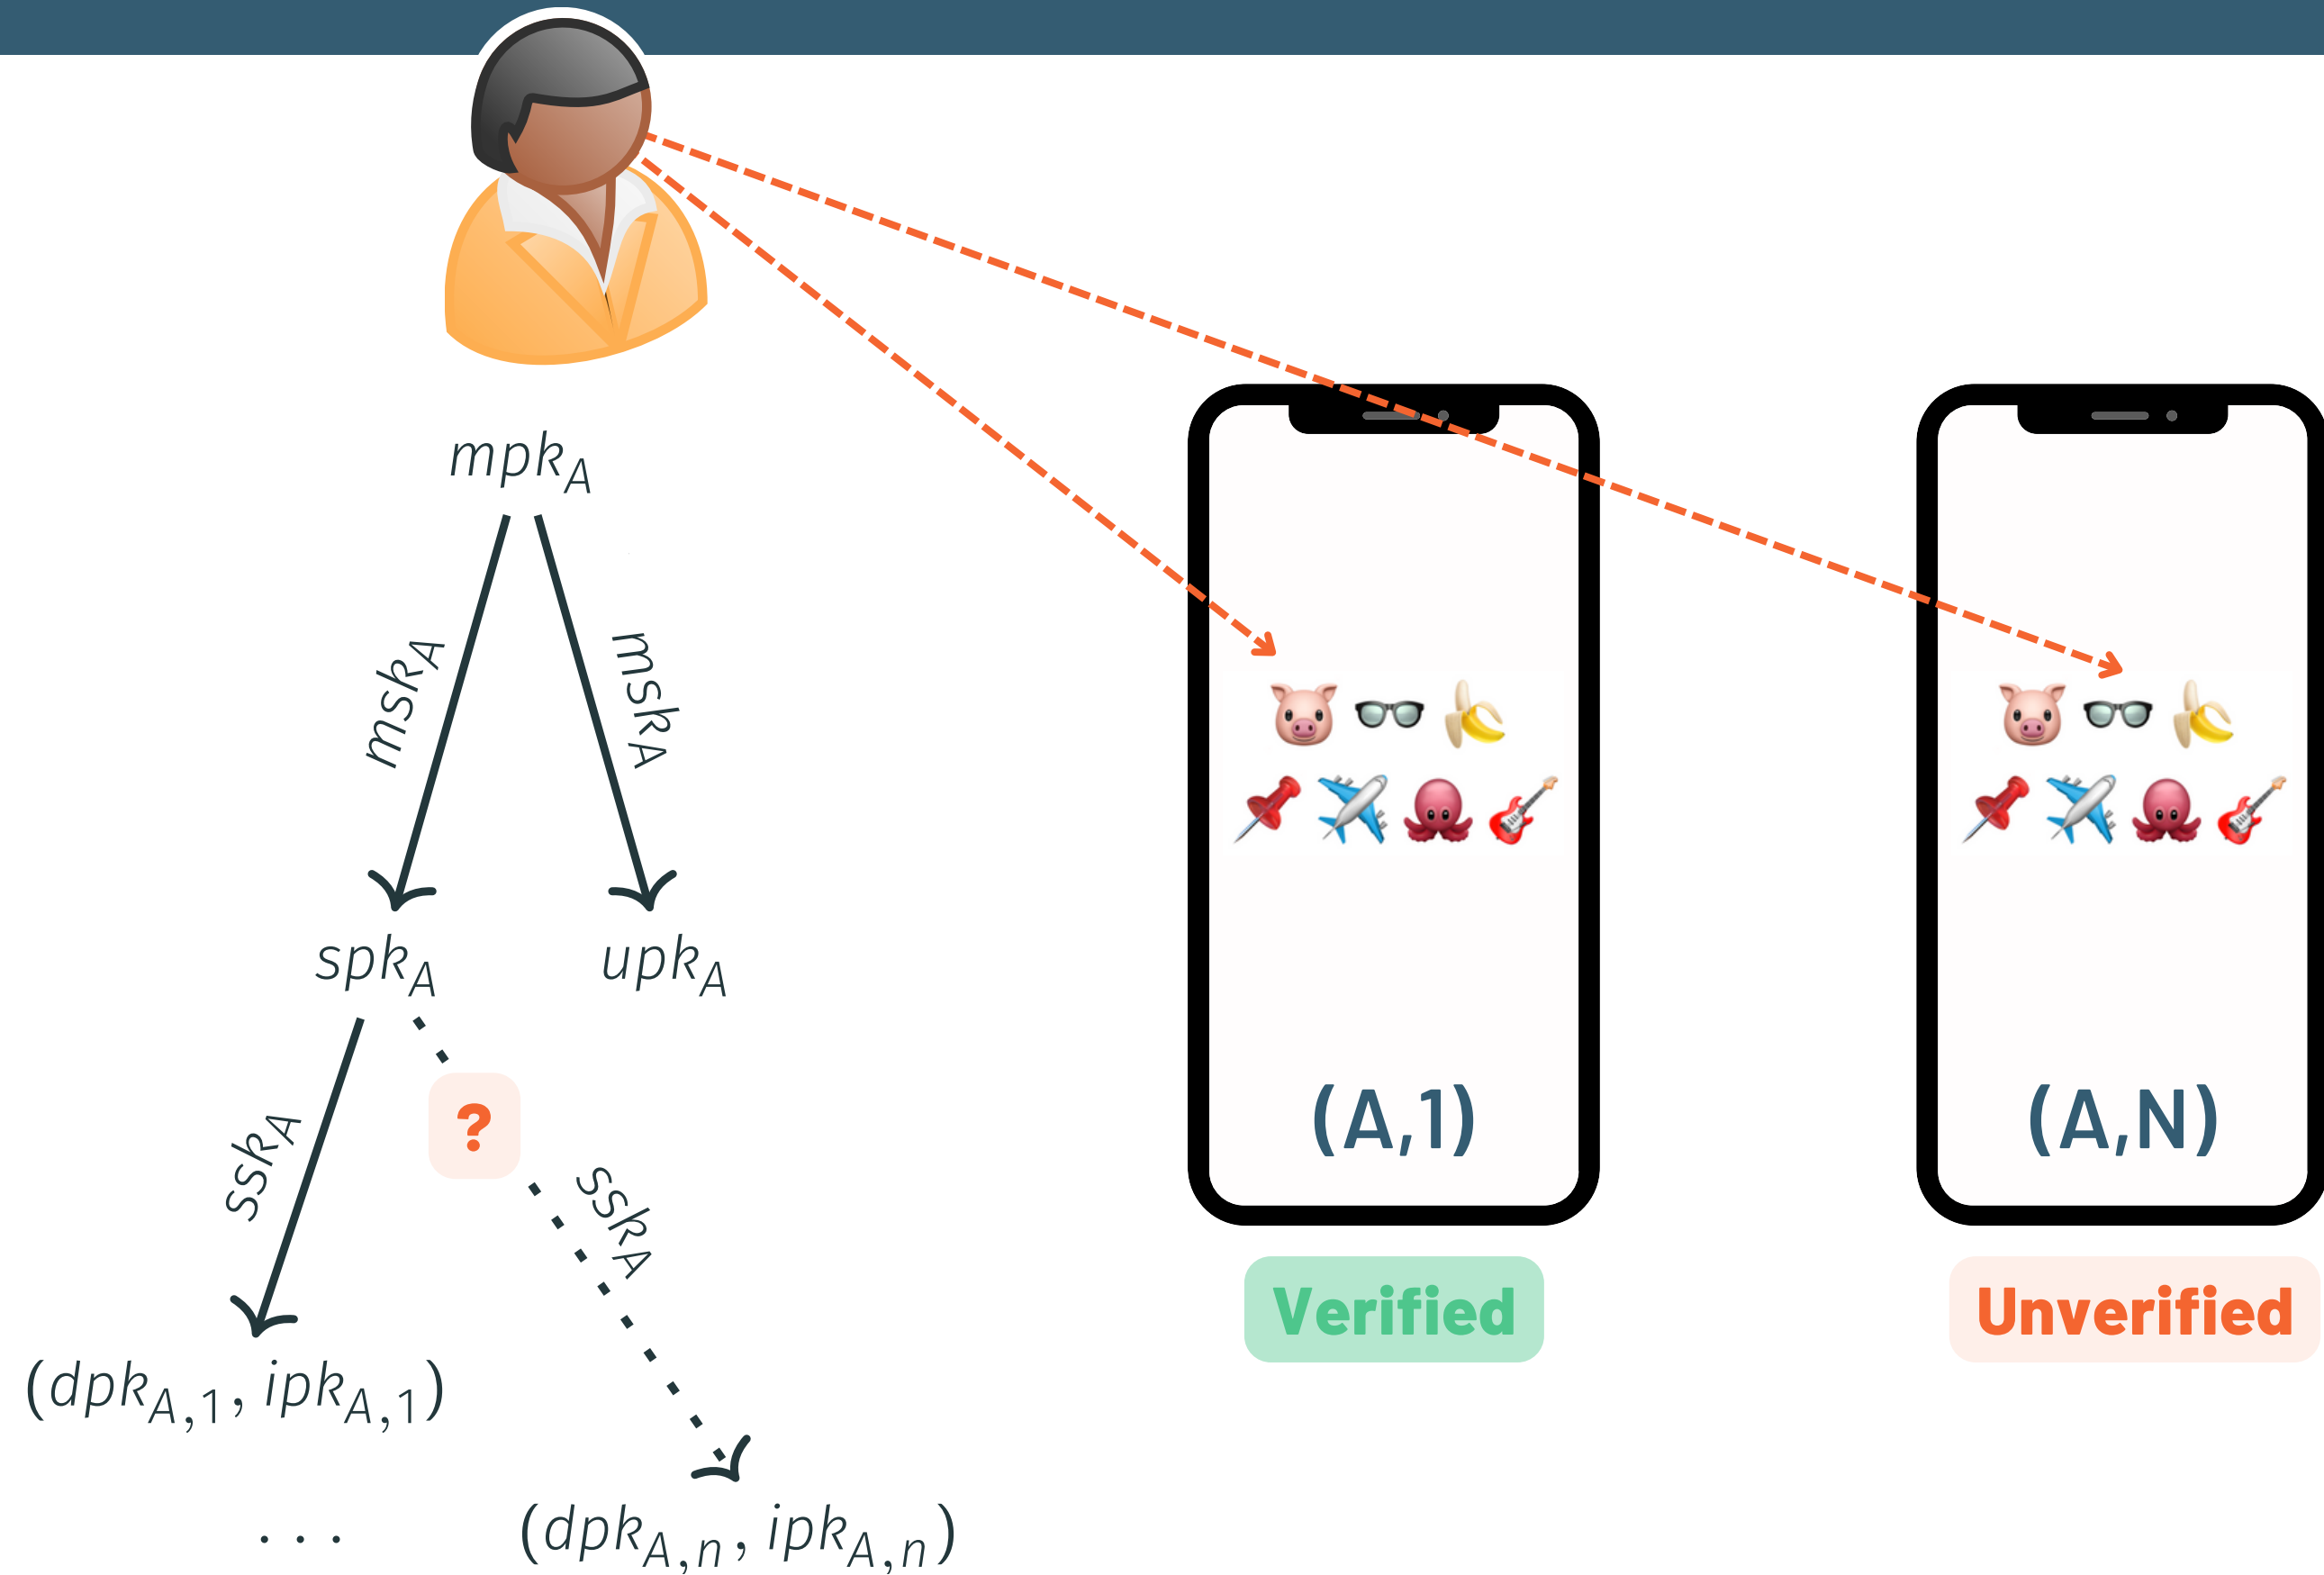


Cross-signing and Verification





Cross-signing and Verification

Device-to-device verification




Cross-signing and Verification

Device-to-device verification


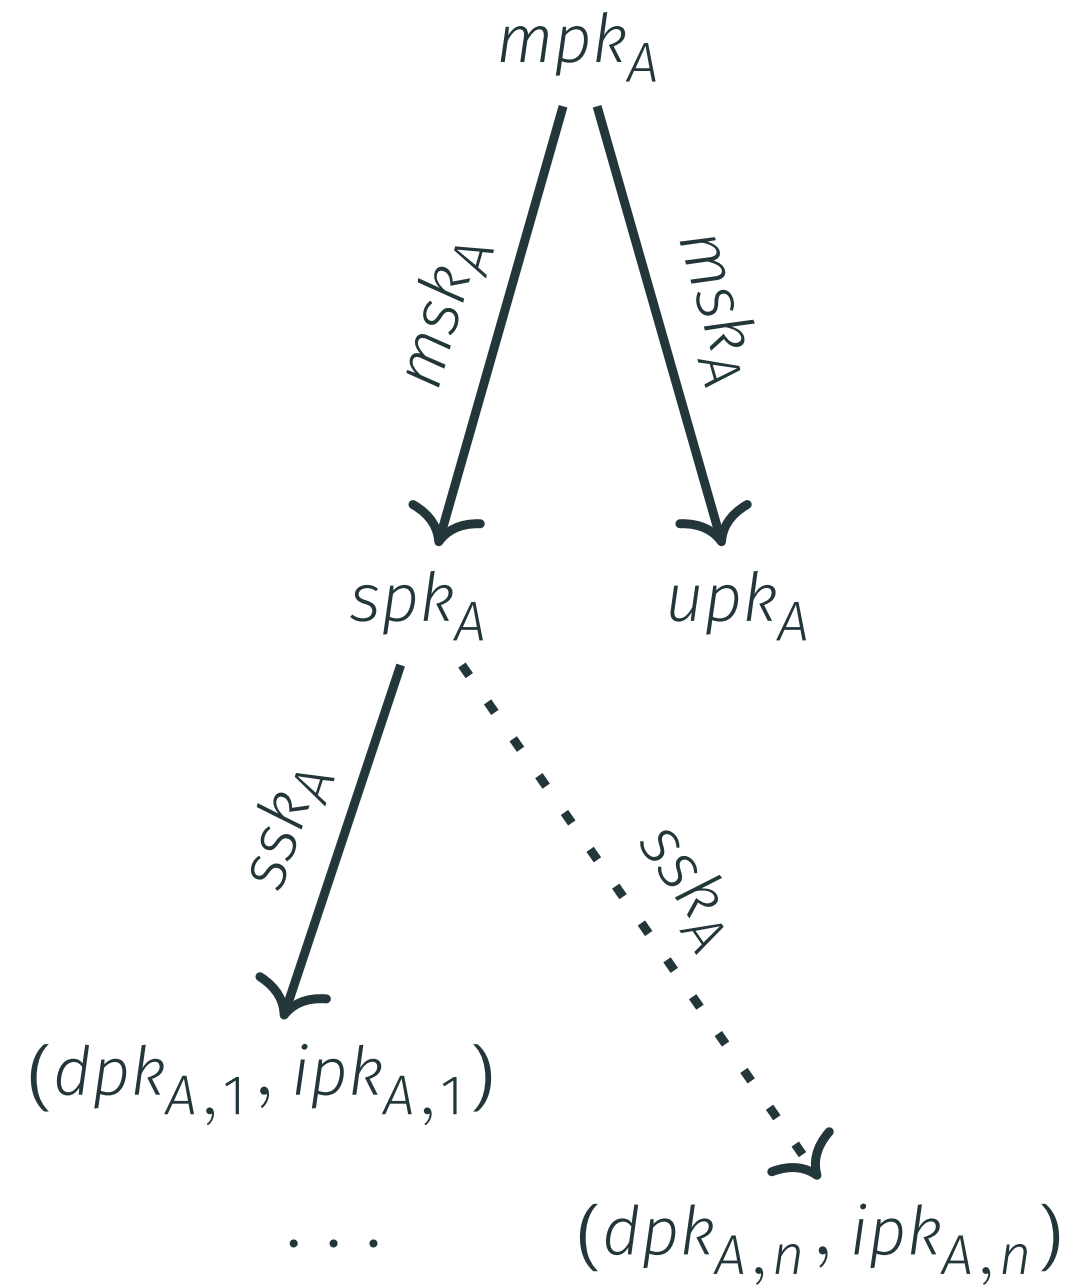


Verified



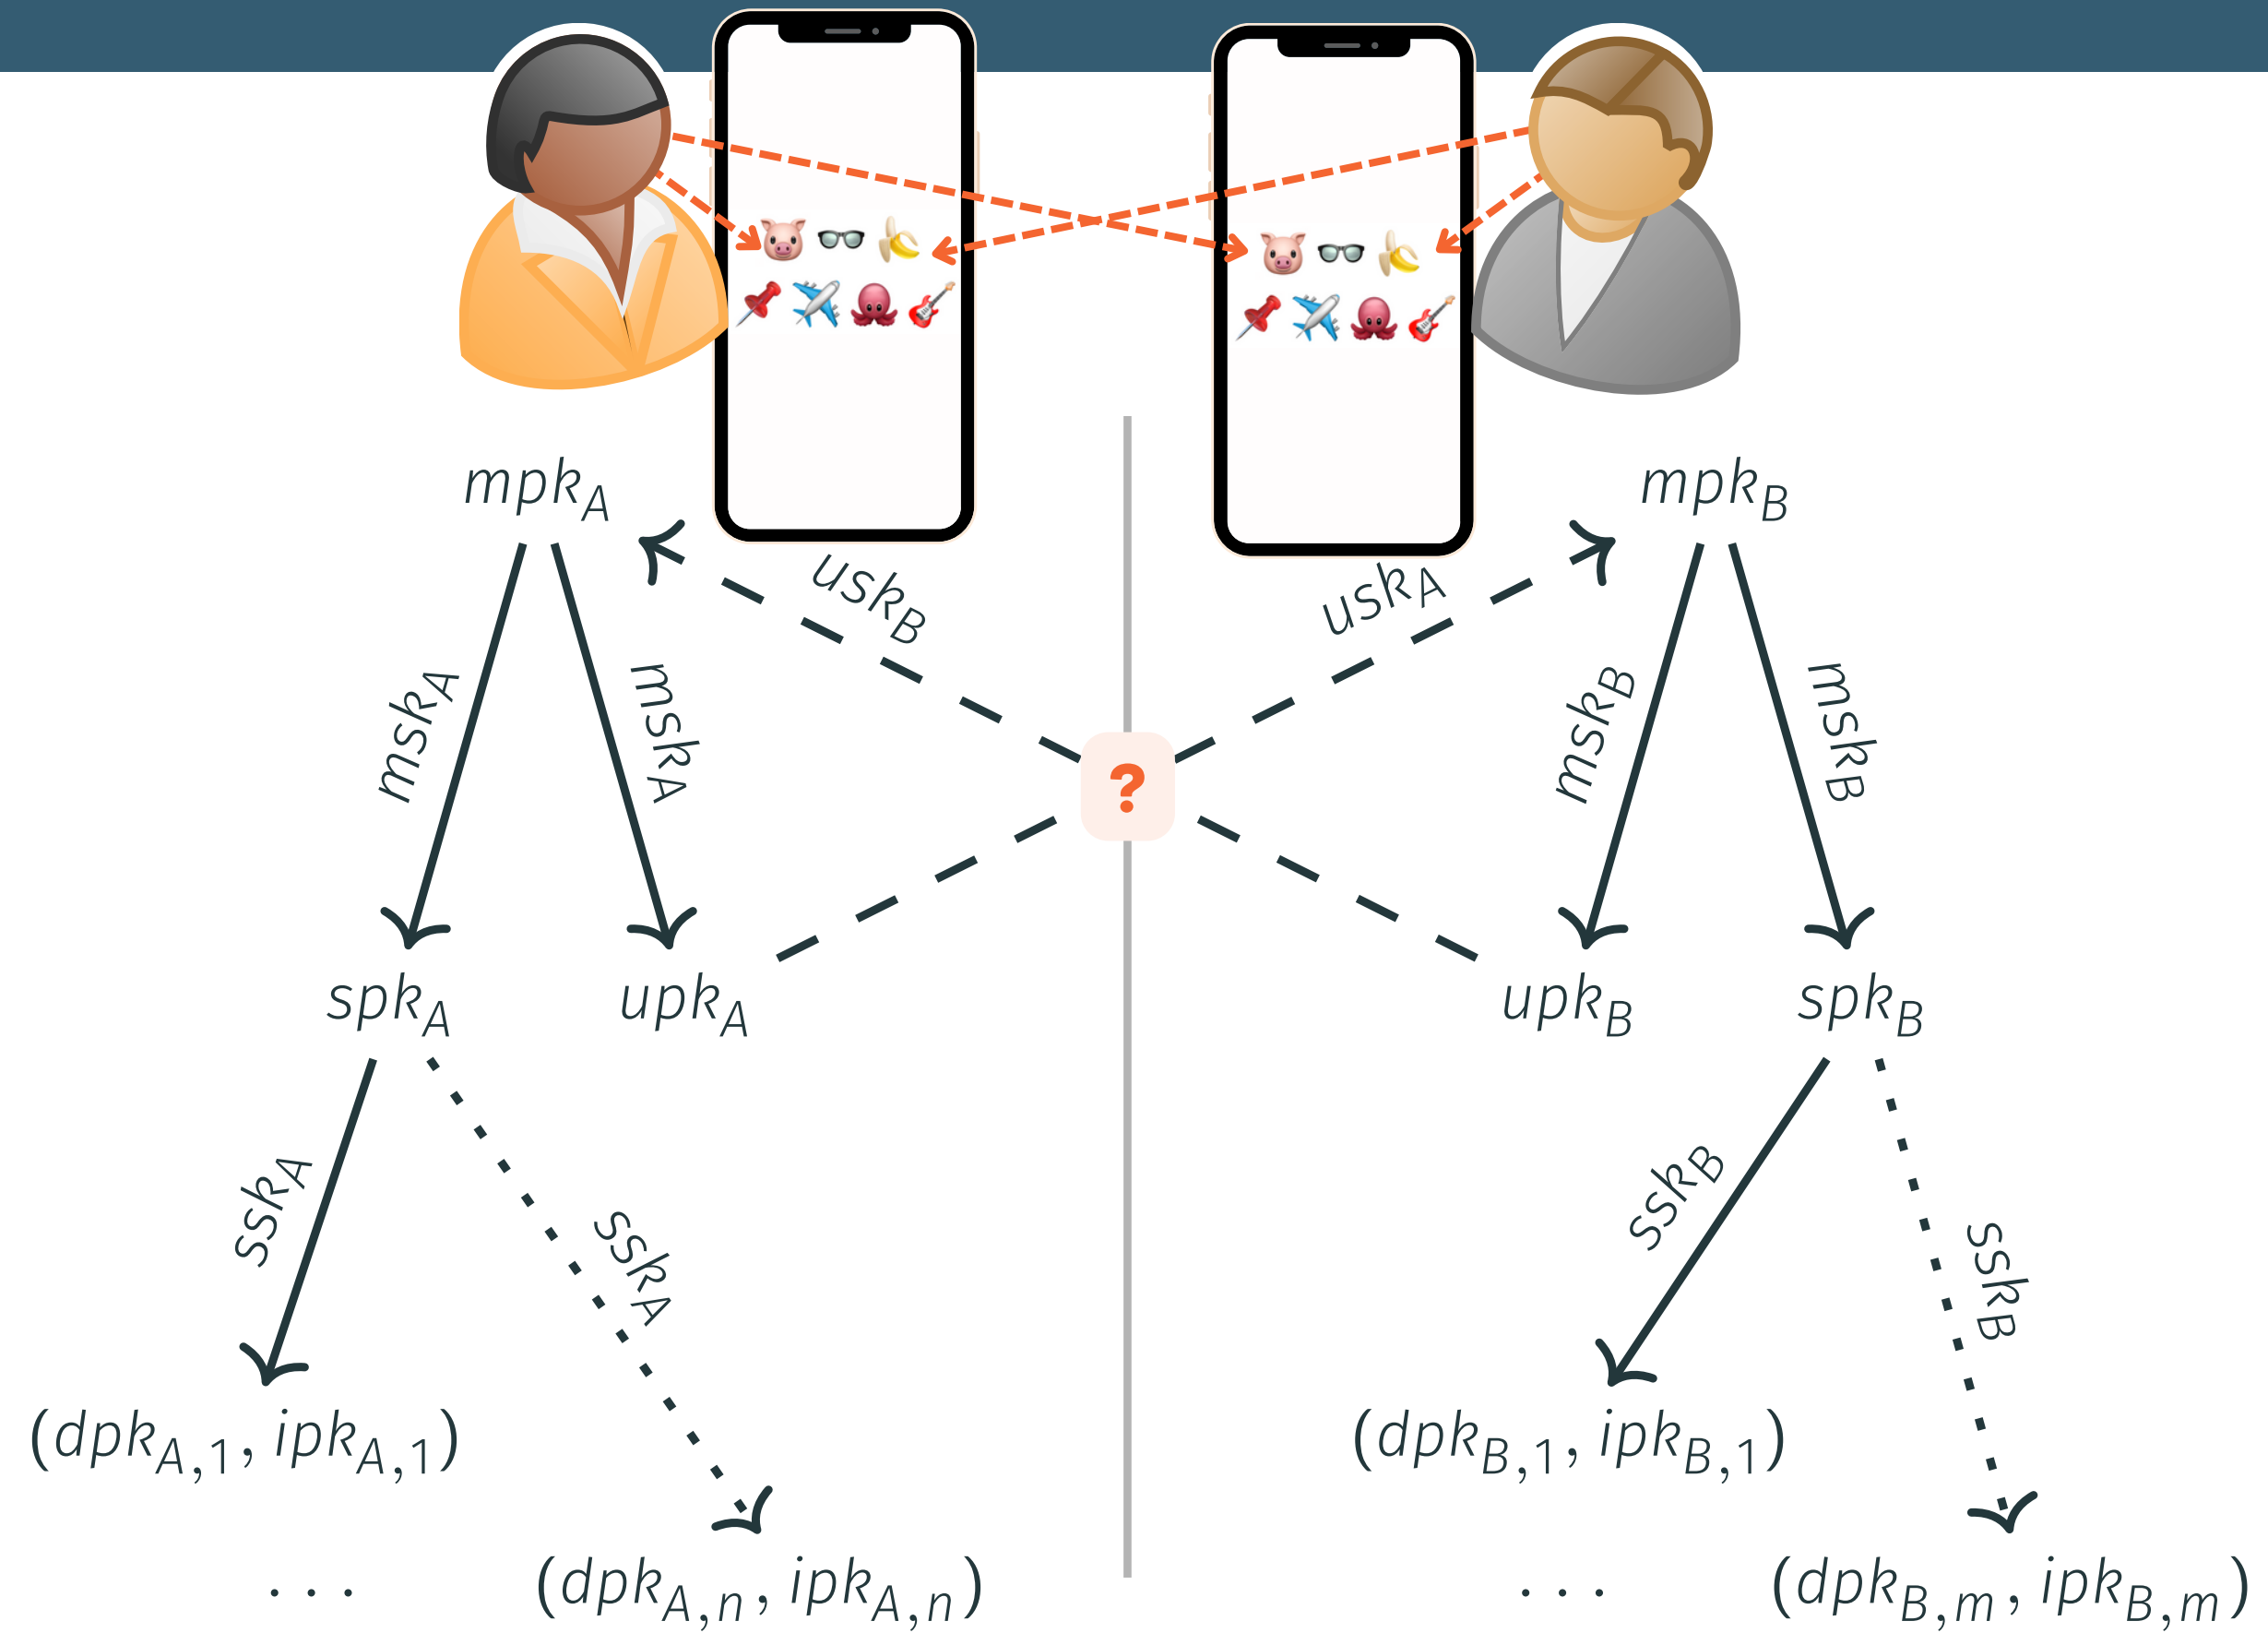
Verified

Cross-signing and Verification



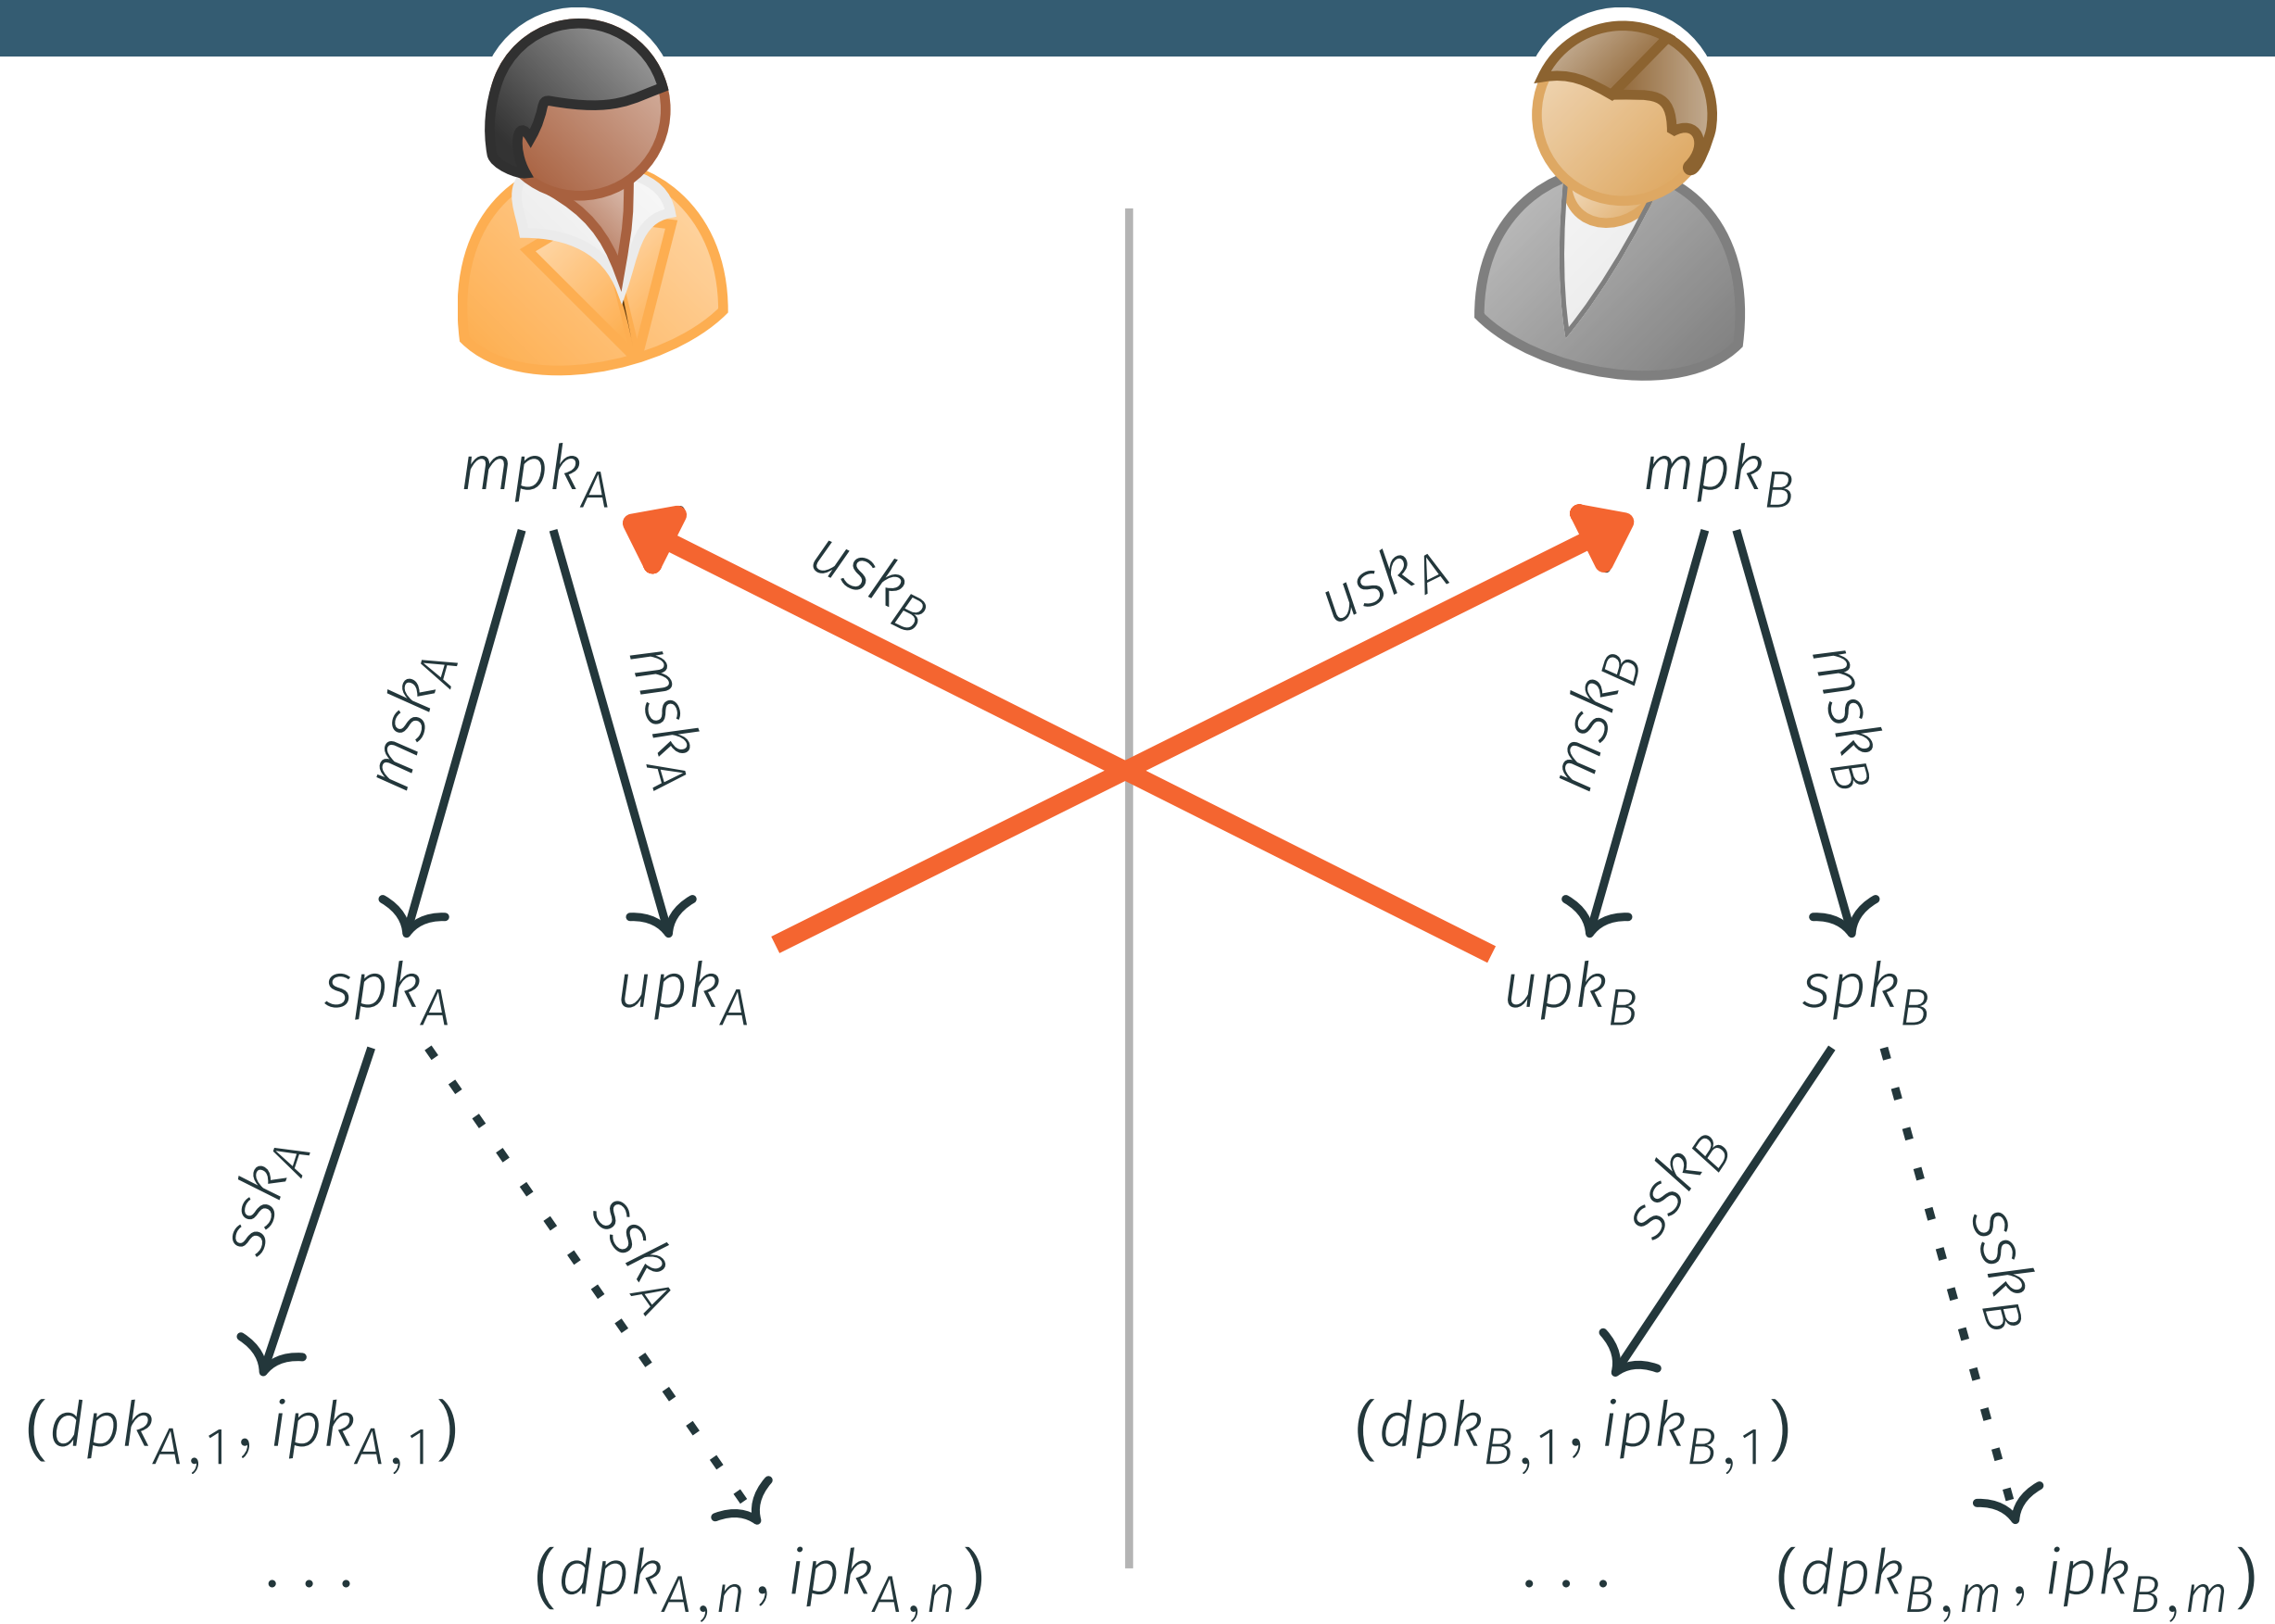
Cross-signing and Verification

User-to-user verification




Cross-signing and Verification

User-to-user verification




Cross-signing and Verification

Verifying device identities



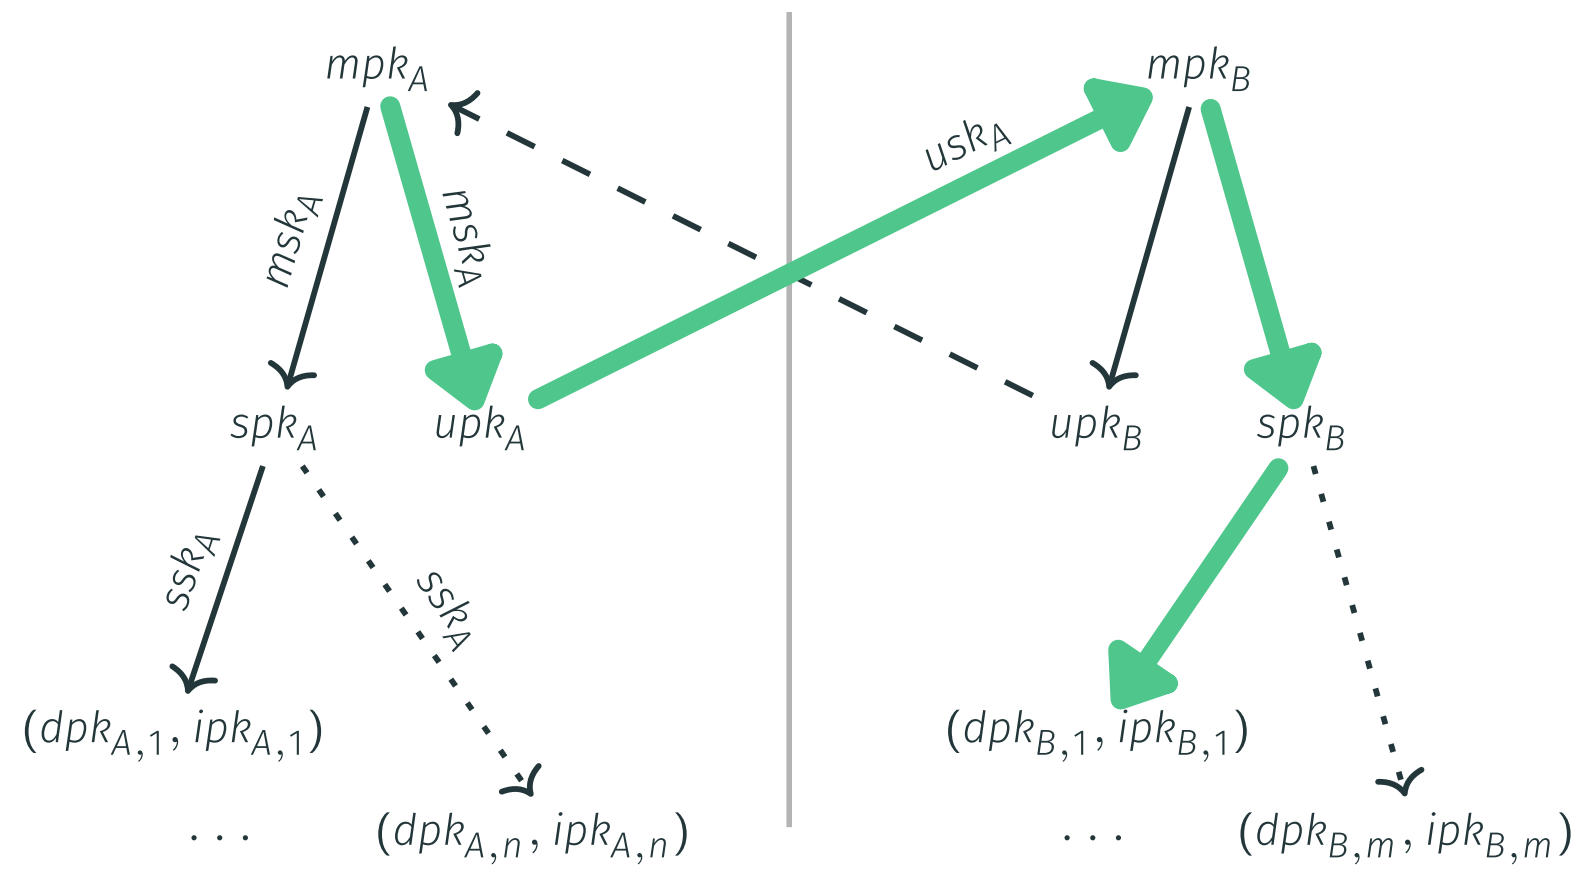

Cross-signing and Verification

Verifying device identities



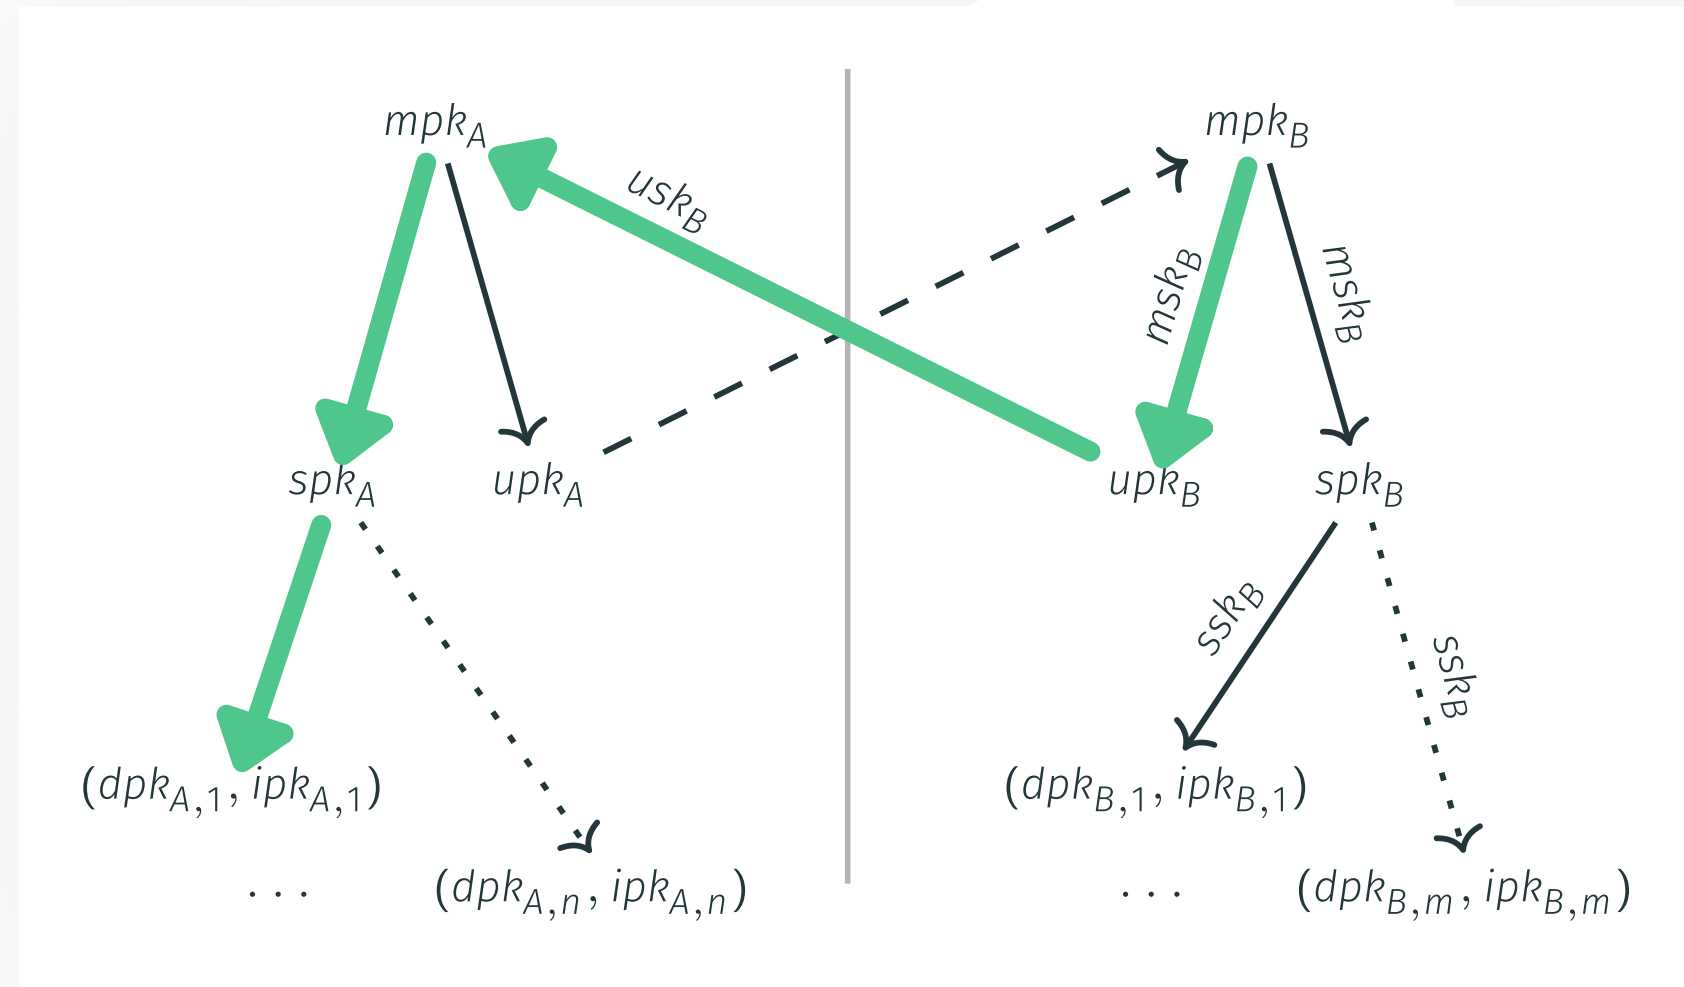
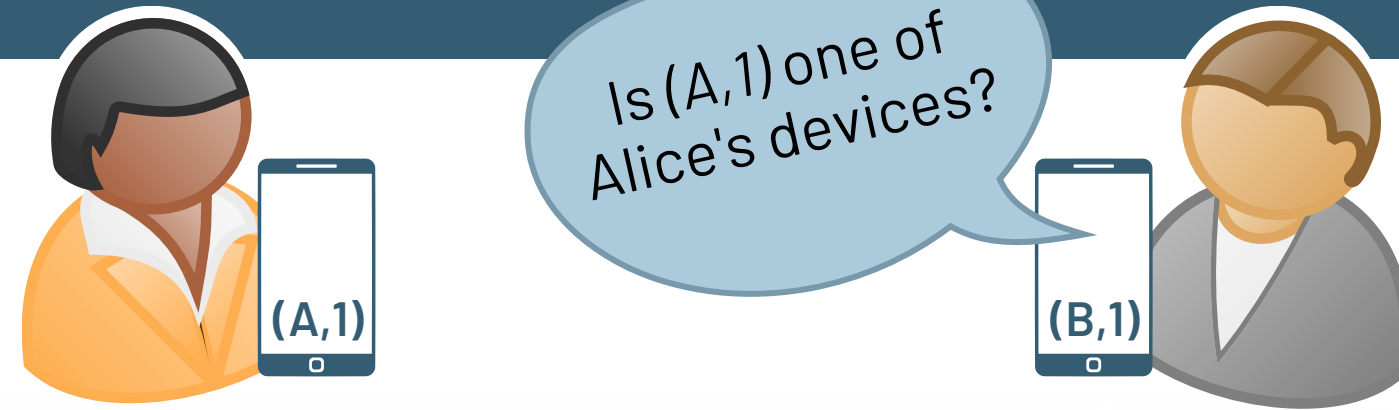
Cross-signing and Verification

Verifying device identities




Cross-signing and Verification

Verifying device identities




Cross-signing and Verification

Short Authentication String Protocol



Cross-signing and Verification


Short Authentication String Protocol



**1. Setup tamperproof channel
with known identities**

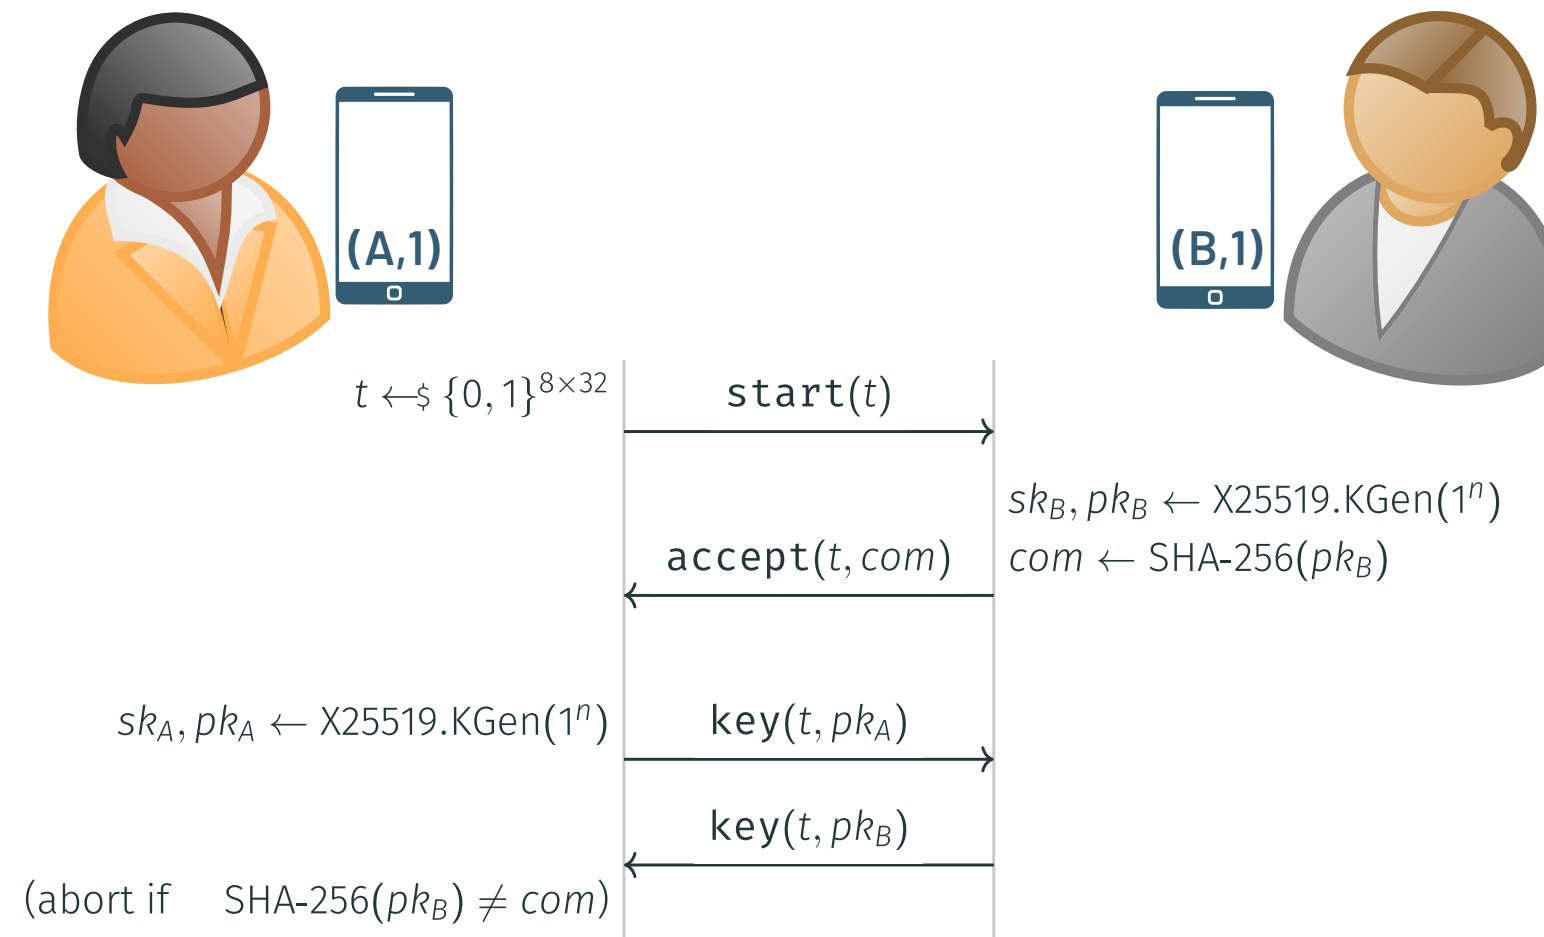
Cross-signing and Verification

Short Authentication String Protocol



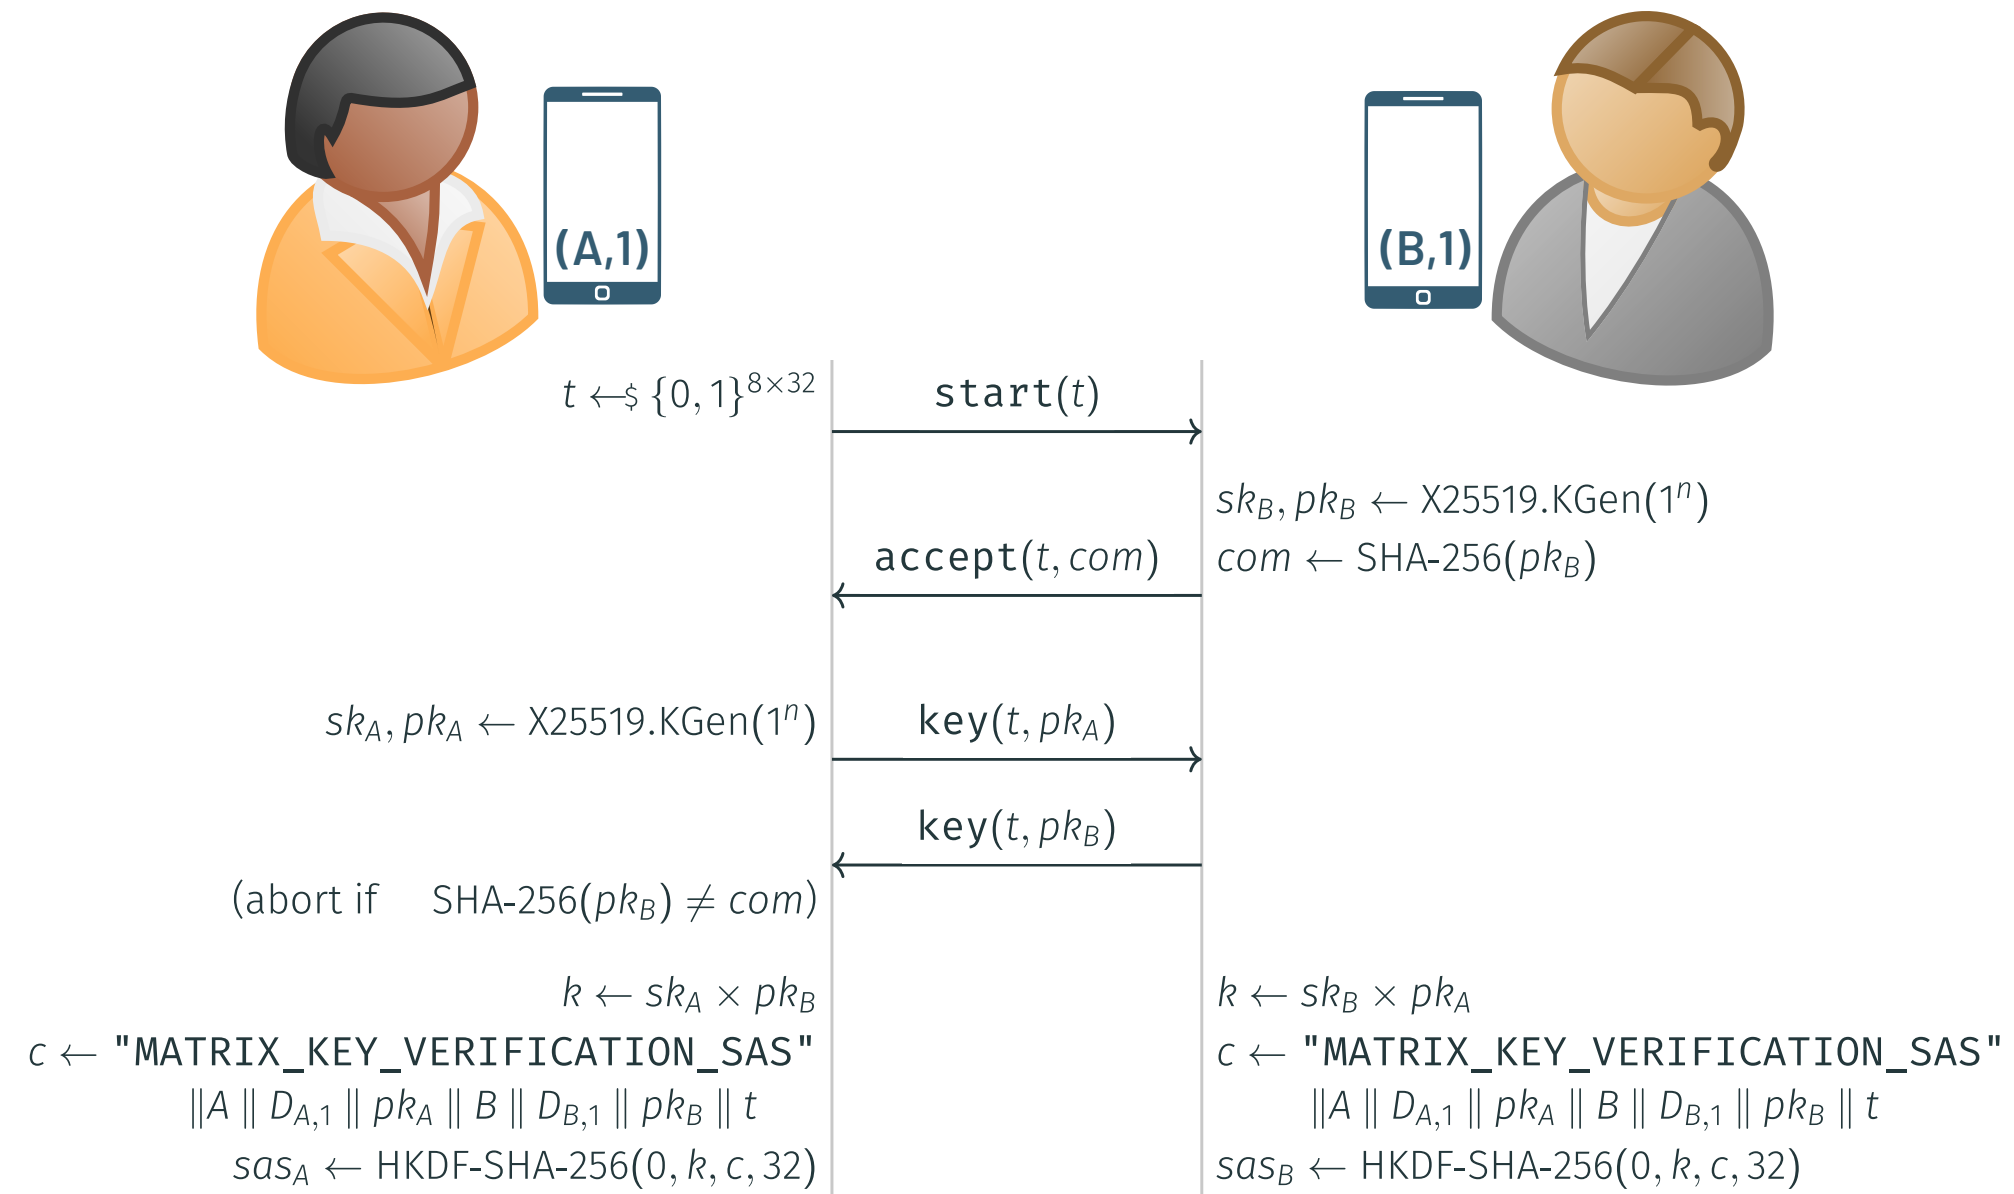
Cross-signing and Verification

Short Authentication String Protocol



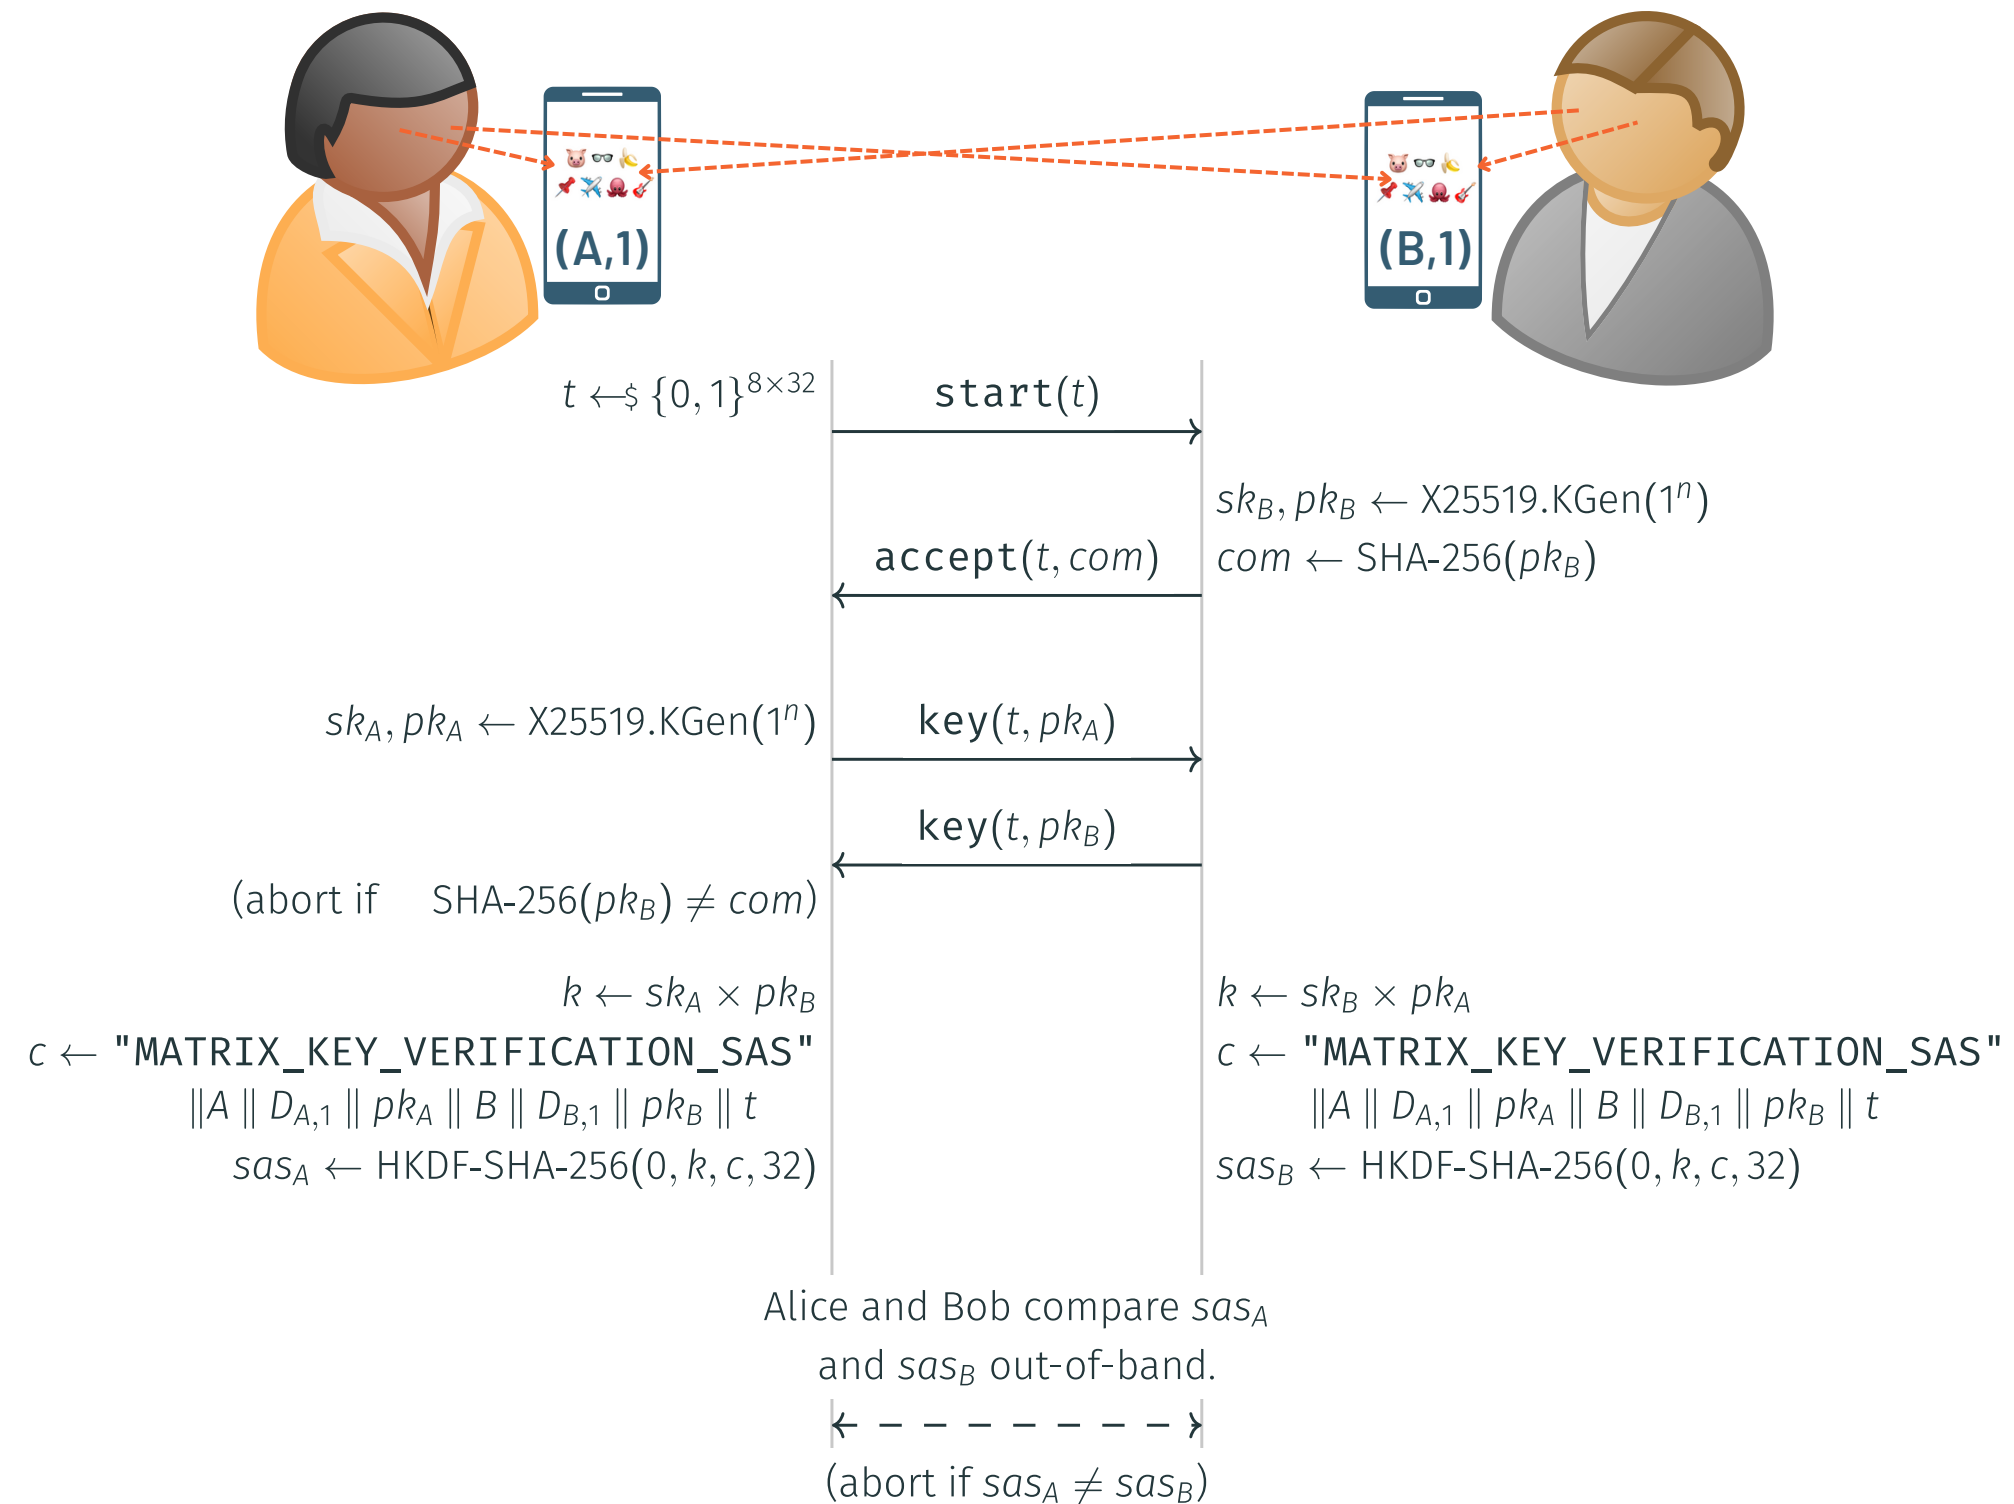
Cross-signing and Verification

Short Authentication String Protocol



Cross-signing and Verification

Short Authentication String Protocol



Cross-signing and Verification

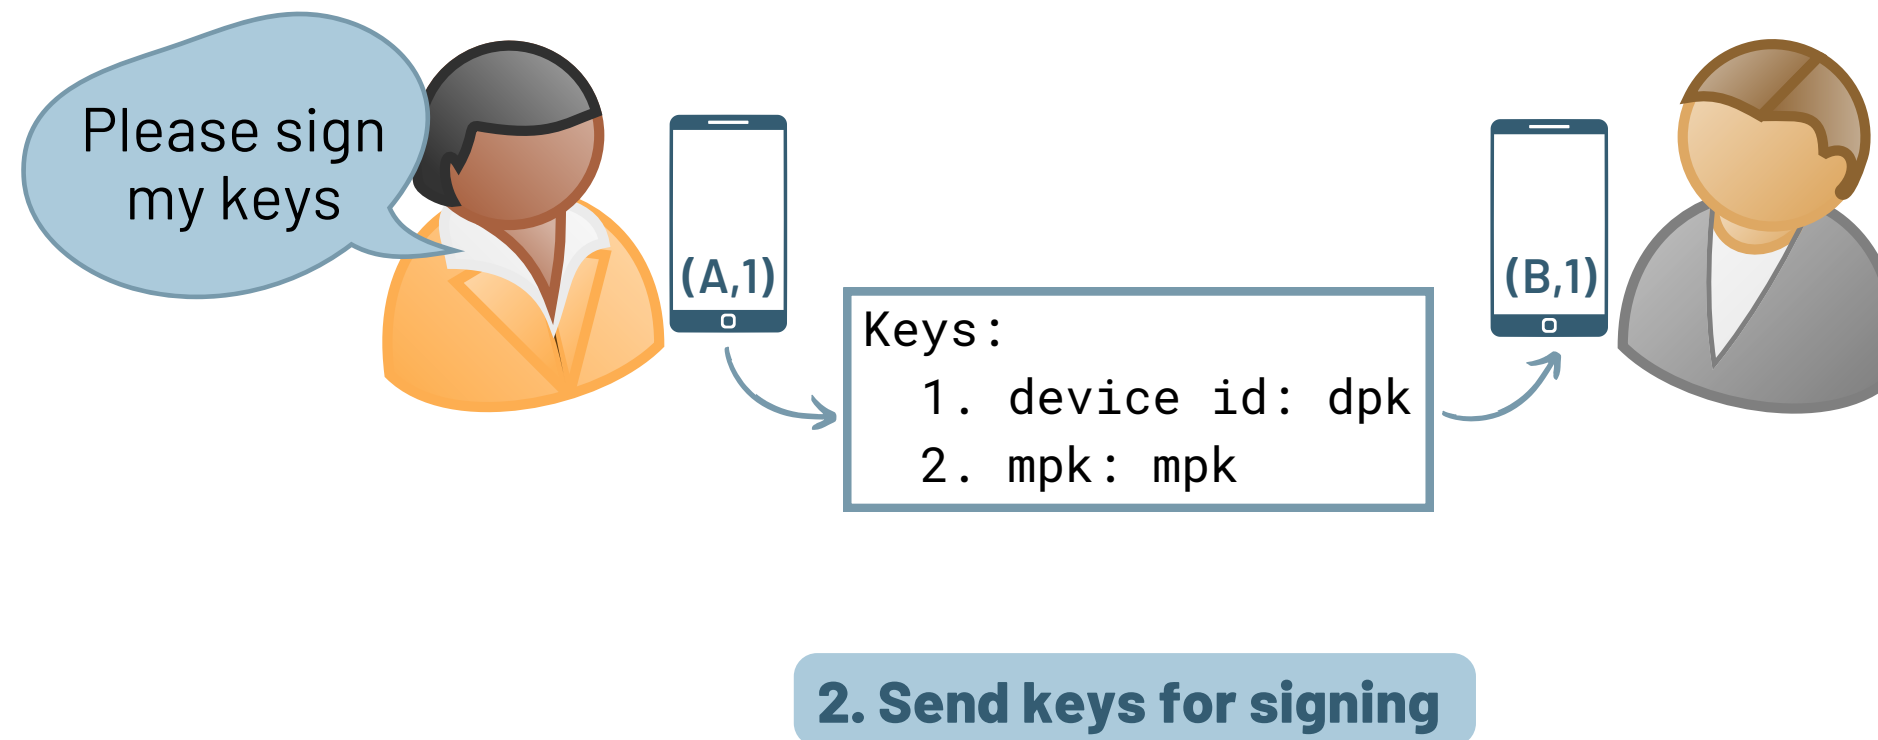
Short Authentication String Protocol



Alice and Bob may use k to communicate securely through the homeserver

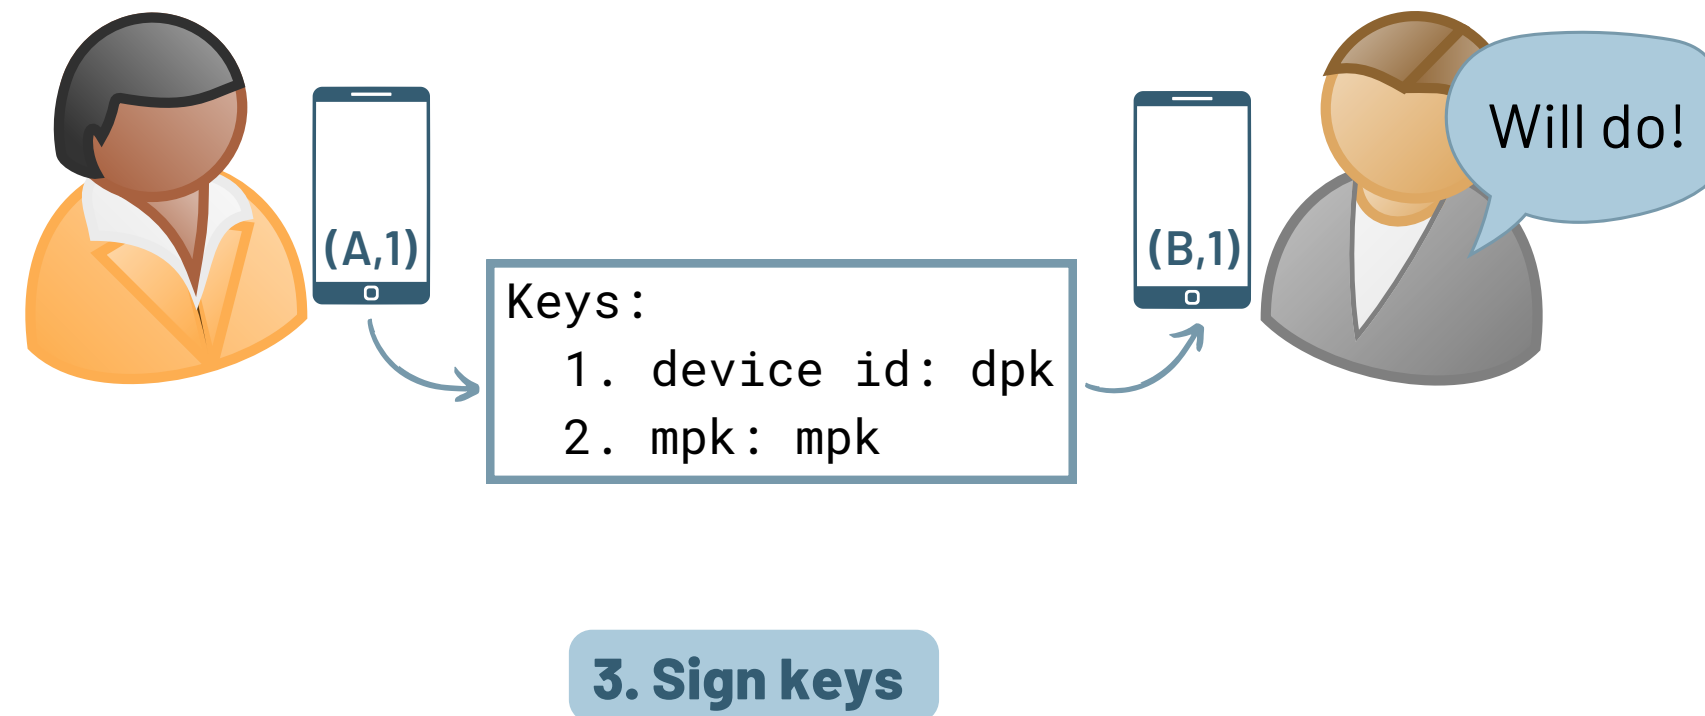
Cross-signing and Verification

Short Authentication String Protocol



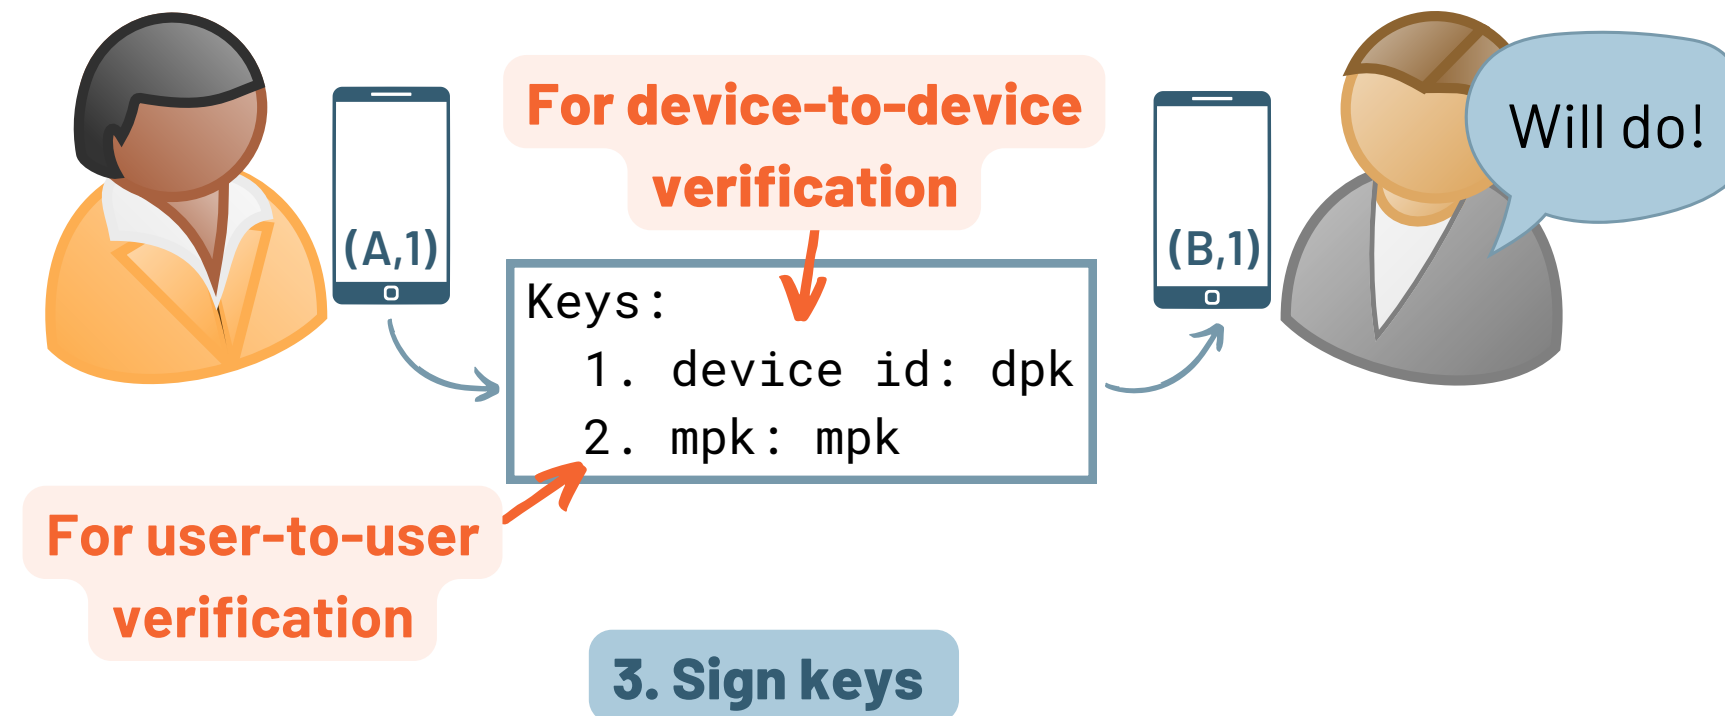
Cross-signing and Verification

Short Authentication String Protocol



Cross-signing and Verification

Short Authentication String Protocol



Cross-signing and Verification

Short Authentication String Protocol

The screenshot shows a web browser window displaying the Matrix specification page for cross-signing. The browser's address bar shows the URL `https://spec.matrix.org/unstable/client-server-api/#cross-signing`. The page title is "[matrix] specification — unstable version". The navigation menu includes "Foundation", "FAQs", and "Blog".

The left sidebar contains a table of contents with the following items:

- 11.12.2.2.6 SAS method: emoji
- 11.12.2.3 Cross-signing
 - 11.12.2.3.1 Key and signature security**
- 11.12.2.4 QR codes
 - 11.12.2.4.1 QR code format
 - 11.12.2.4.2 Verification messages specific to QR codes
- 11.12.3 Sharing keys between devices
 - 11.12.3.1 Key requests
 - 11.12.3.2 Server-side key backups
 - 11.12.3.2.1 Recovery key
 - 11.12.3.2.2 Backup algorithm:
`m.megolm_backup.v1.curve25519-sha2-aes-sha2`
 - 11.12.3.3 Key exports
 - 11.12.3.3.1 Key export format
- 11.12.4 Messaging Algorithms

The main content area features a diagram at the top showing two sets of data blocks connected by dashed lines, representing cross-signing between devices. Below the diagram, the text explains that **Verification methods** can be used to verify a user's master key by using the master public key, encoded using unpadded base64, as the device ID, and treating it as a normal device. For example, if Alice and Bob verify each other using SAS, Alice's `m.key.verification.mac` message to Bob may include `"ed25519:alices+master+public+key": "alices+master+public+key"` in the `mac` property. Servers therefore must ensure that device IDs will not collide with cross-signing public keys.

The text further states that cross-signing private keys can be stored on the server or shared with other devices using the **Secrets** module. When doing so, the master, user-signing, and self-signing keys are identified using the names `m.cross_signing.master`, `m.cross_signing.user_signing`, and `m.cross_signing.self_signing`, respectively, and the keys are base64-encoded before being encrypted.

11.12.2.3.1. Key and signature security

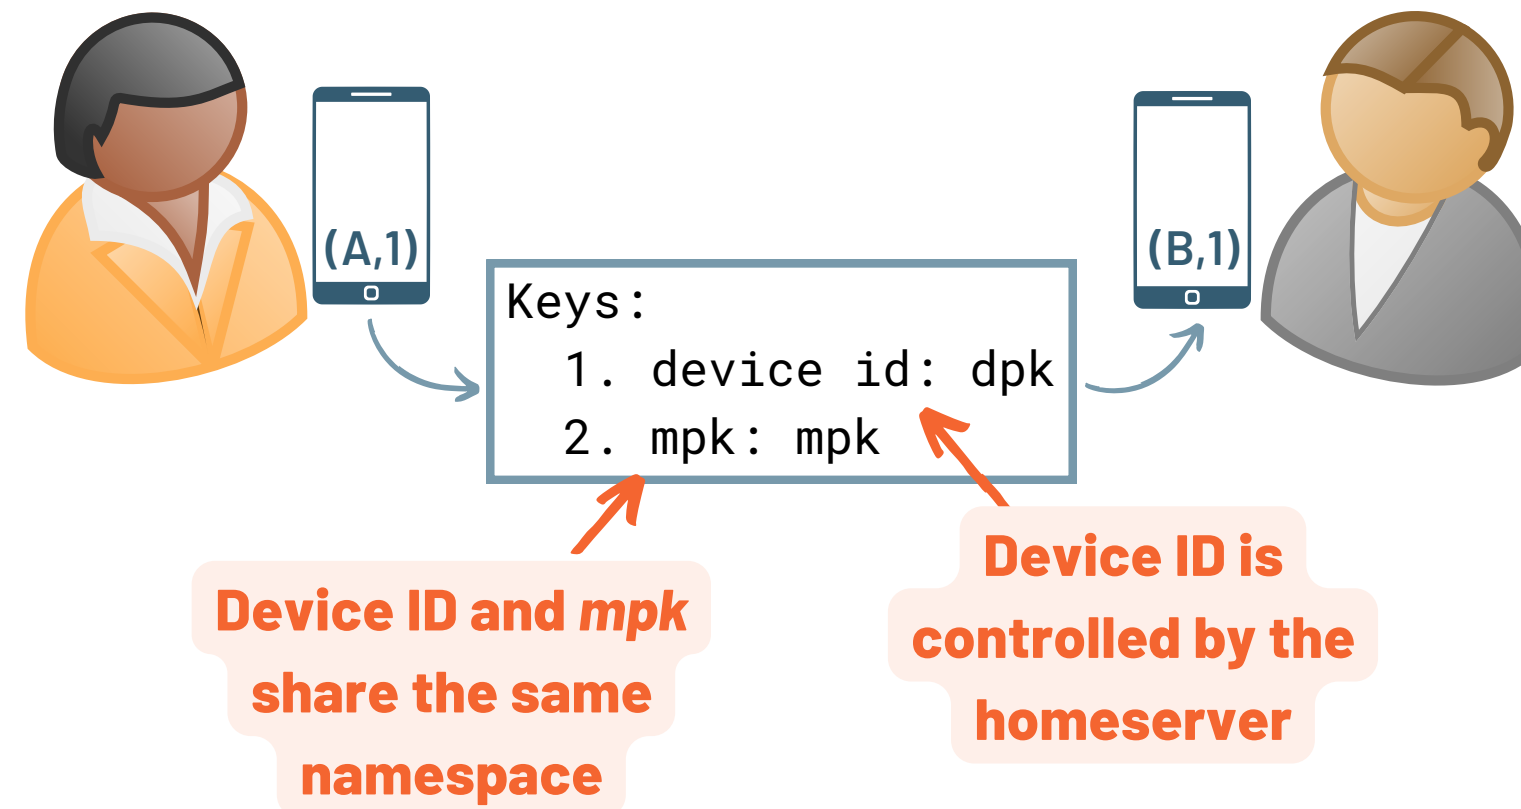
A user's master key could allow an attacker to impersonate that user to other users, or other users to that user. Thus clients must ensure that the private part of the master key is treated securely. If clients do not have a secure means of storing the master key (such as a secret storage system provided by the operating system), then clients must not store the private part.

If a user's client sees that any other user has changed their master key, that client must notify the user about the change before allowing communication between the users to continue.

Since device key IDs (`ed25519:DEVICE_ID`) and cross-signing key IDs (`ed25519:PUBLIC_KEY`) occupy the same namespace, clients must ensure that they

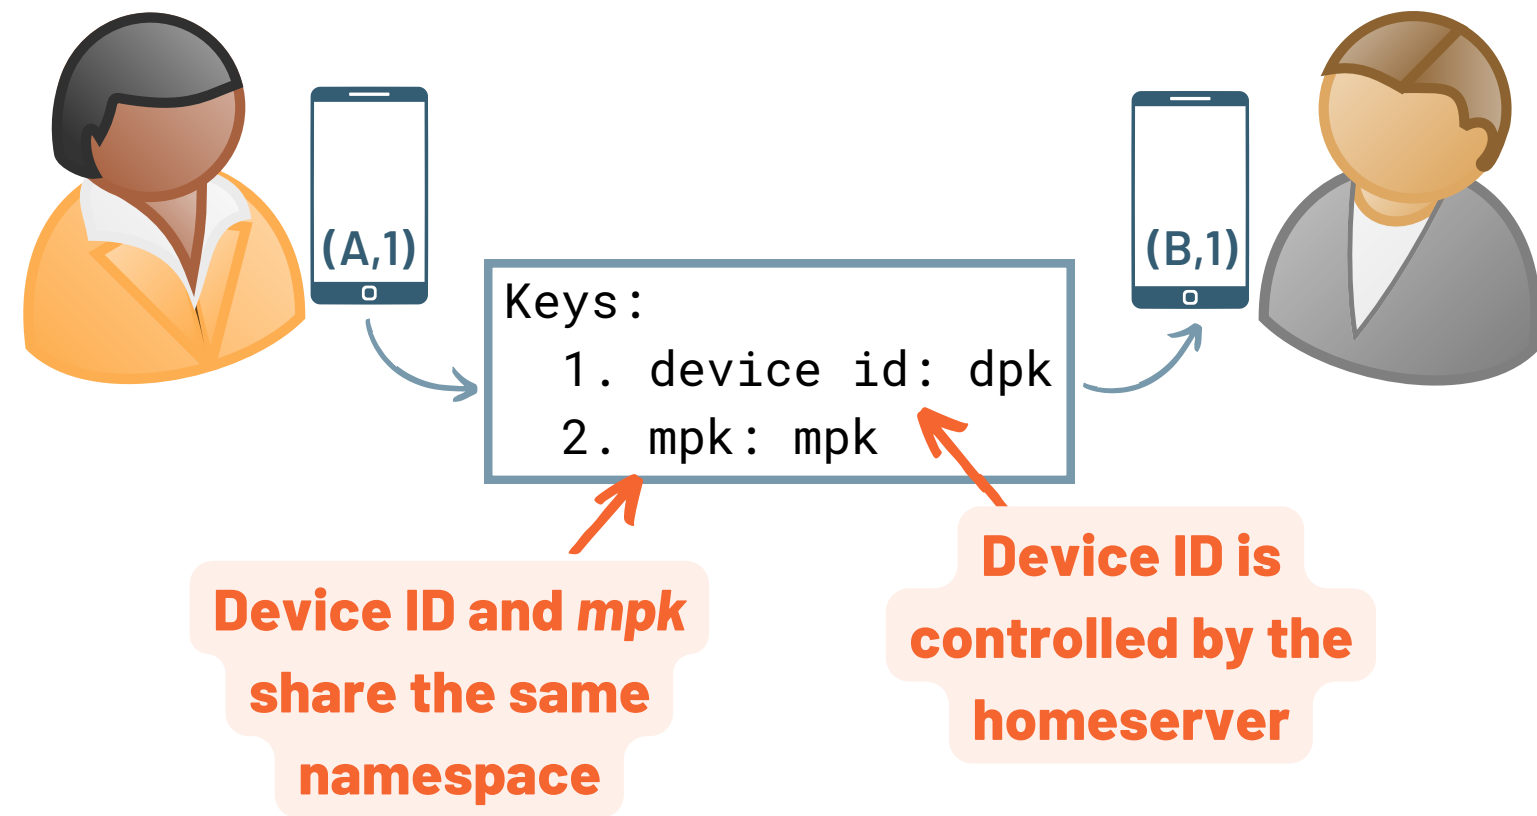
Cross-signing and Verification

Short Authentication String Protocol



User/Device Confusion in Out-of-Band Verification

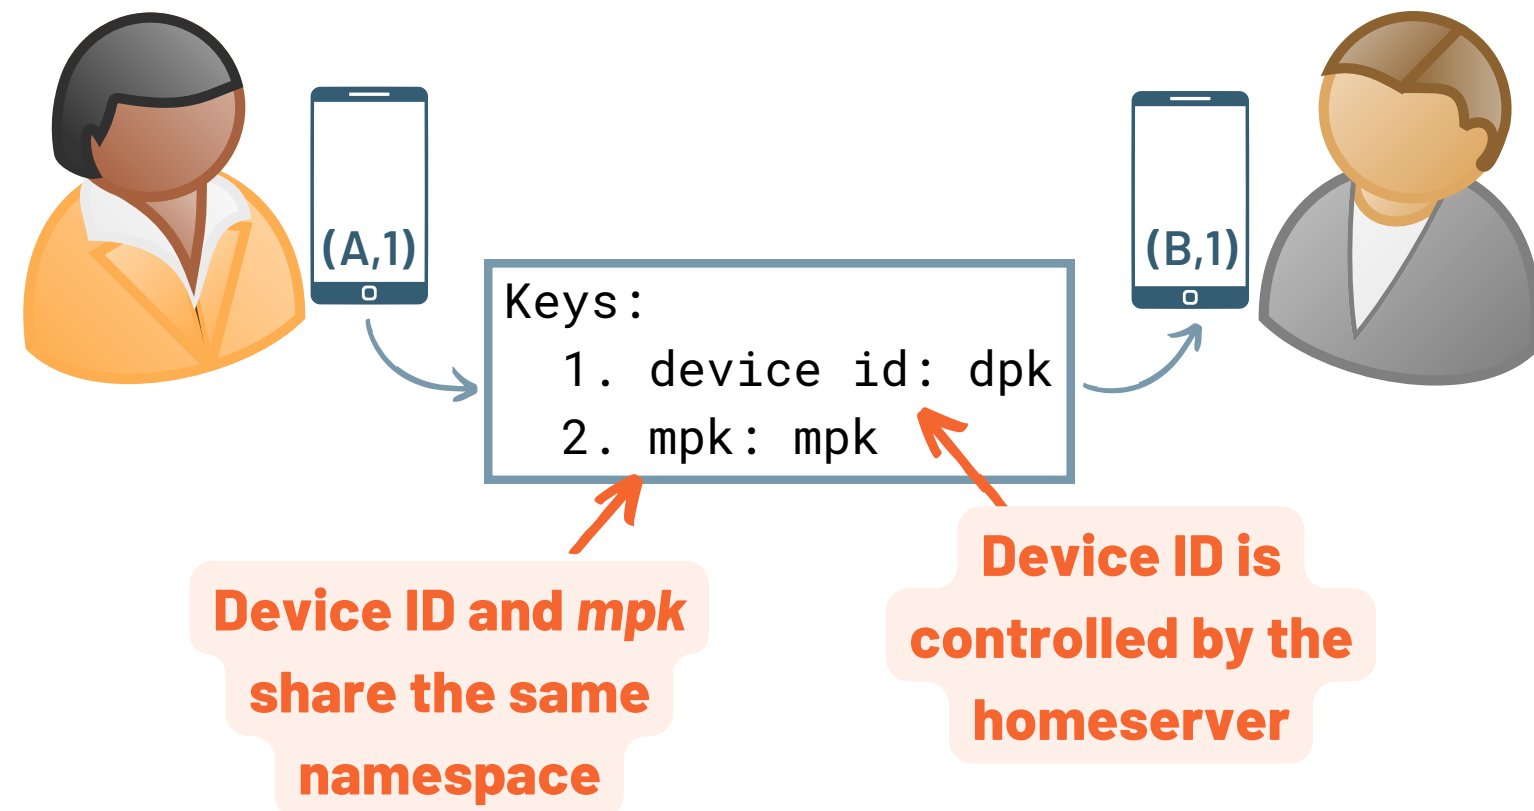
Aim: Trick clients into signing an attacker controlled identity.



User/Device Confusion in Out-of-Band Verification

Aim: Trick clients into signing an attacker controlled identity.

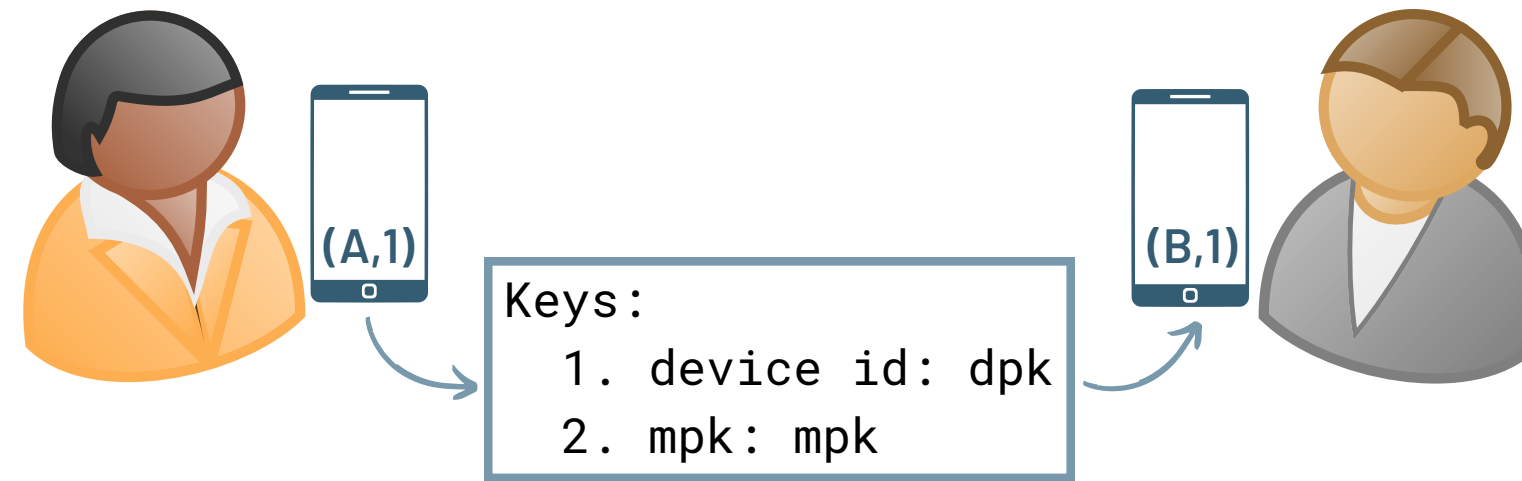
If we set Alice's device identifier to a valid cross-signing key, **can we trick Bob into signing it as if it were Alice's?**



User/Device Confusion in Out-of-Band Verification

Aim: Trick clients into signing an attacker controlled identity.

*If we set Alice's device identifier to a valid cross-signing key, can we trick Bob into signing it as if it were Alice's? **Yes***

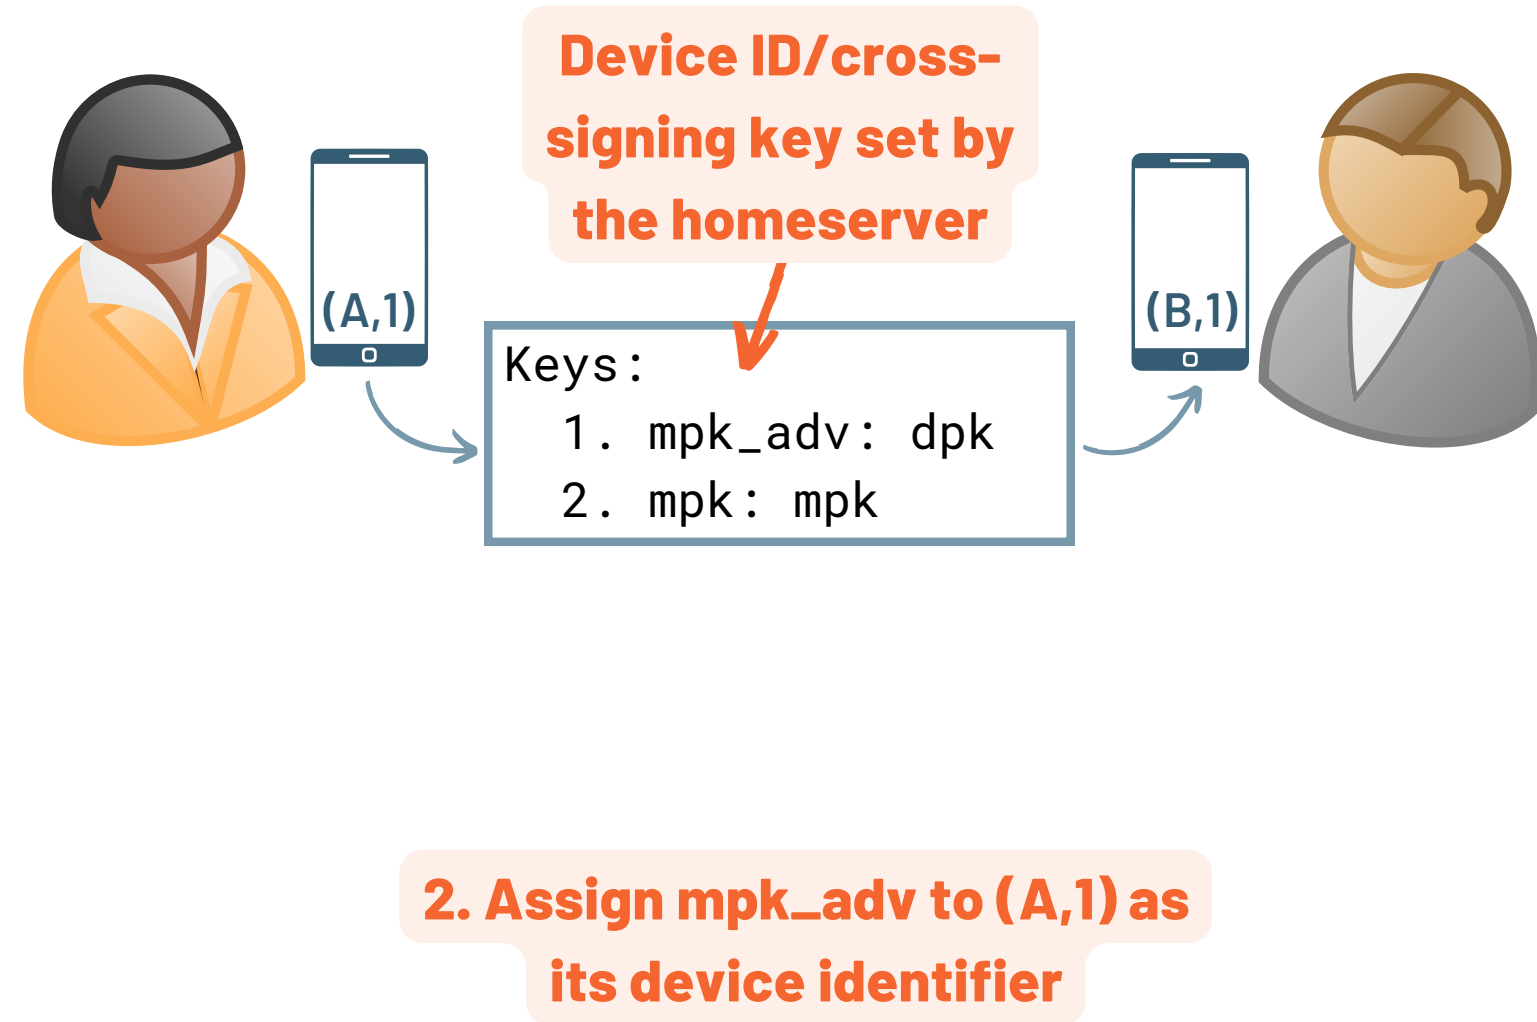


1. Construct a string mpk_adv as a valid Device ID and user cross-signing key (mpk)

User/Device Confusion in Out-of-Band Verification

Aim: Trick clients into signing an attacker controlled identity.

*If we set Alice's device identifier to a valid cross-signing key, can we trick Bob into signing it as if it were Alice's? **Yes***

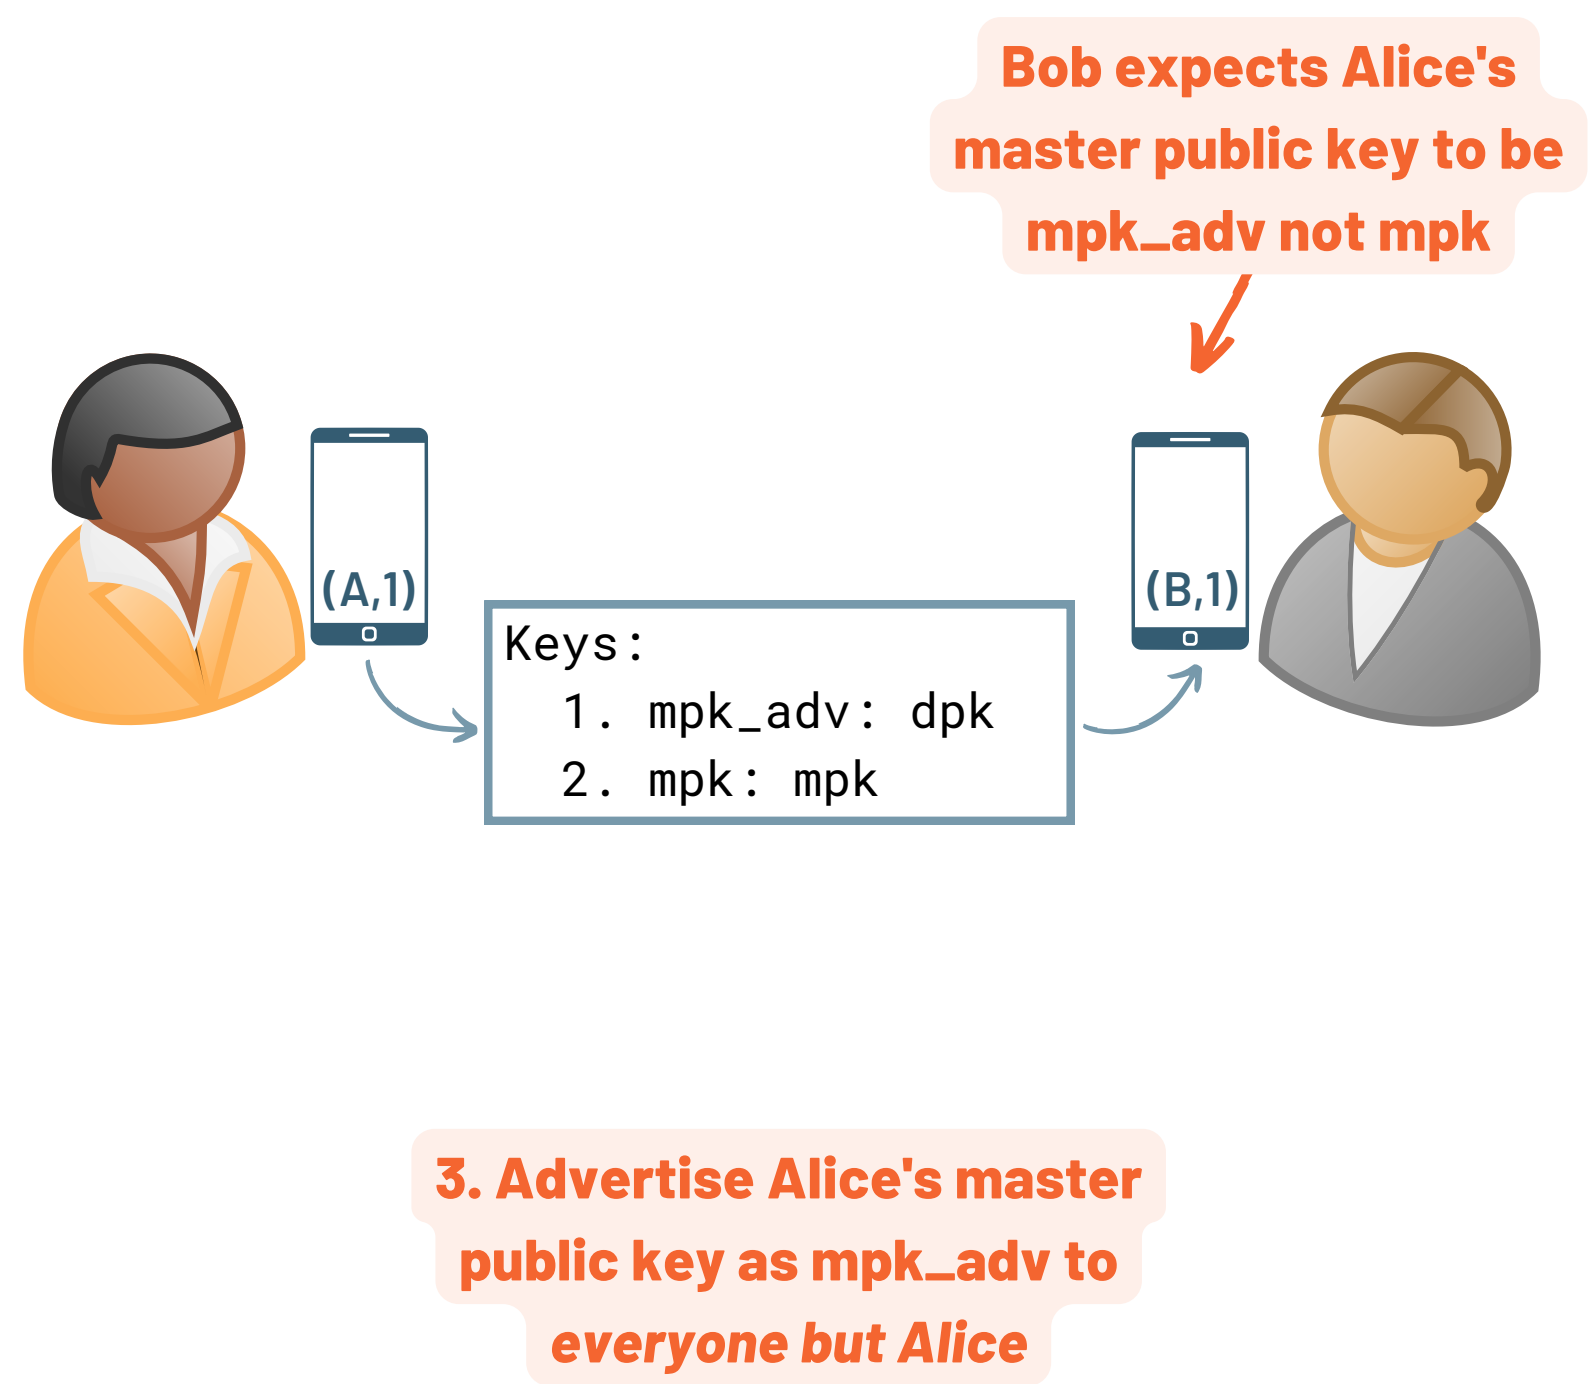


User/Device Confusion in Out-of-Band Verification

Aim: Trick clients into signing an attacker controlled identity.

*If we set Alice's device identifier to a valid cross-signing key, can we trick Bob into signing it as if it were Alice's? **Yes***

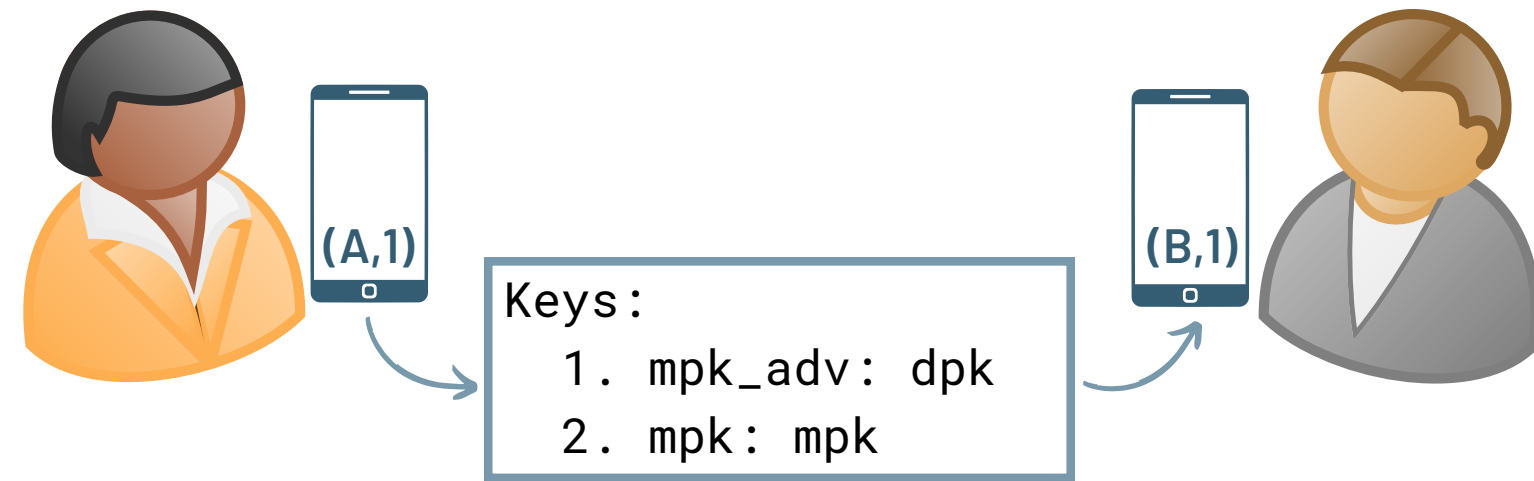
ATTACK



User/Device Confusion in Out-of-Band Verification

Aim: Trick clients into signing an attacker controlled identity.

*If we set Alice's device identifier to a valid cross-signing key, can we trick Bob into signing it as if it were Alice's? **Yes***



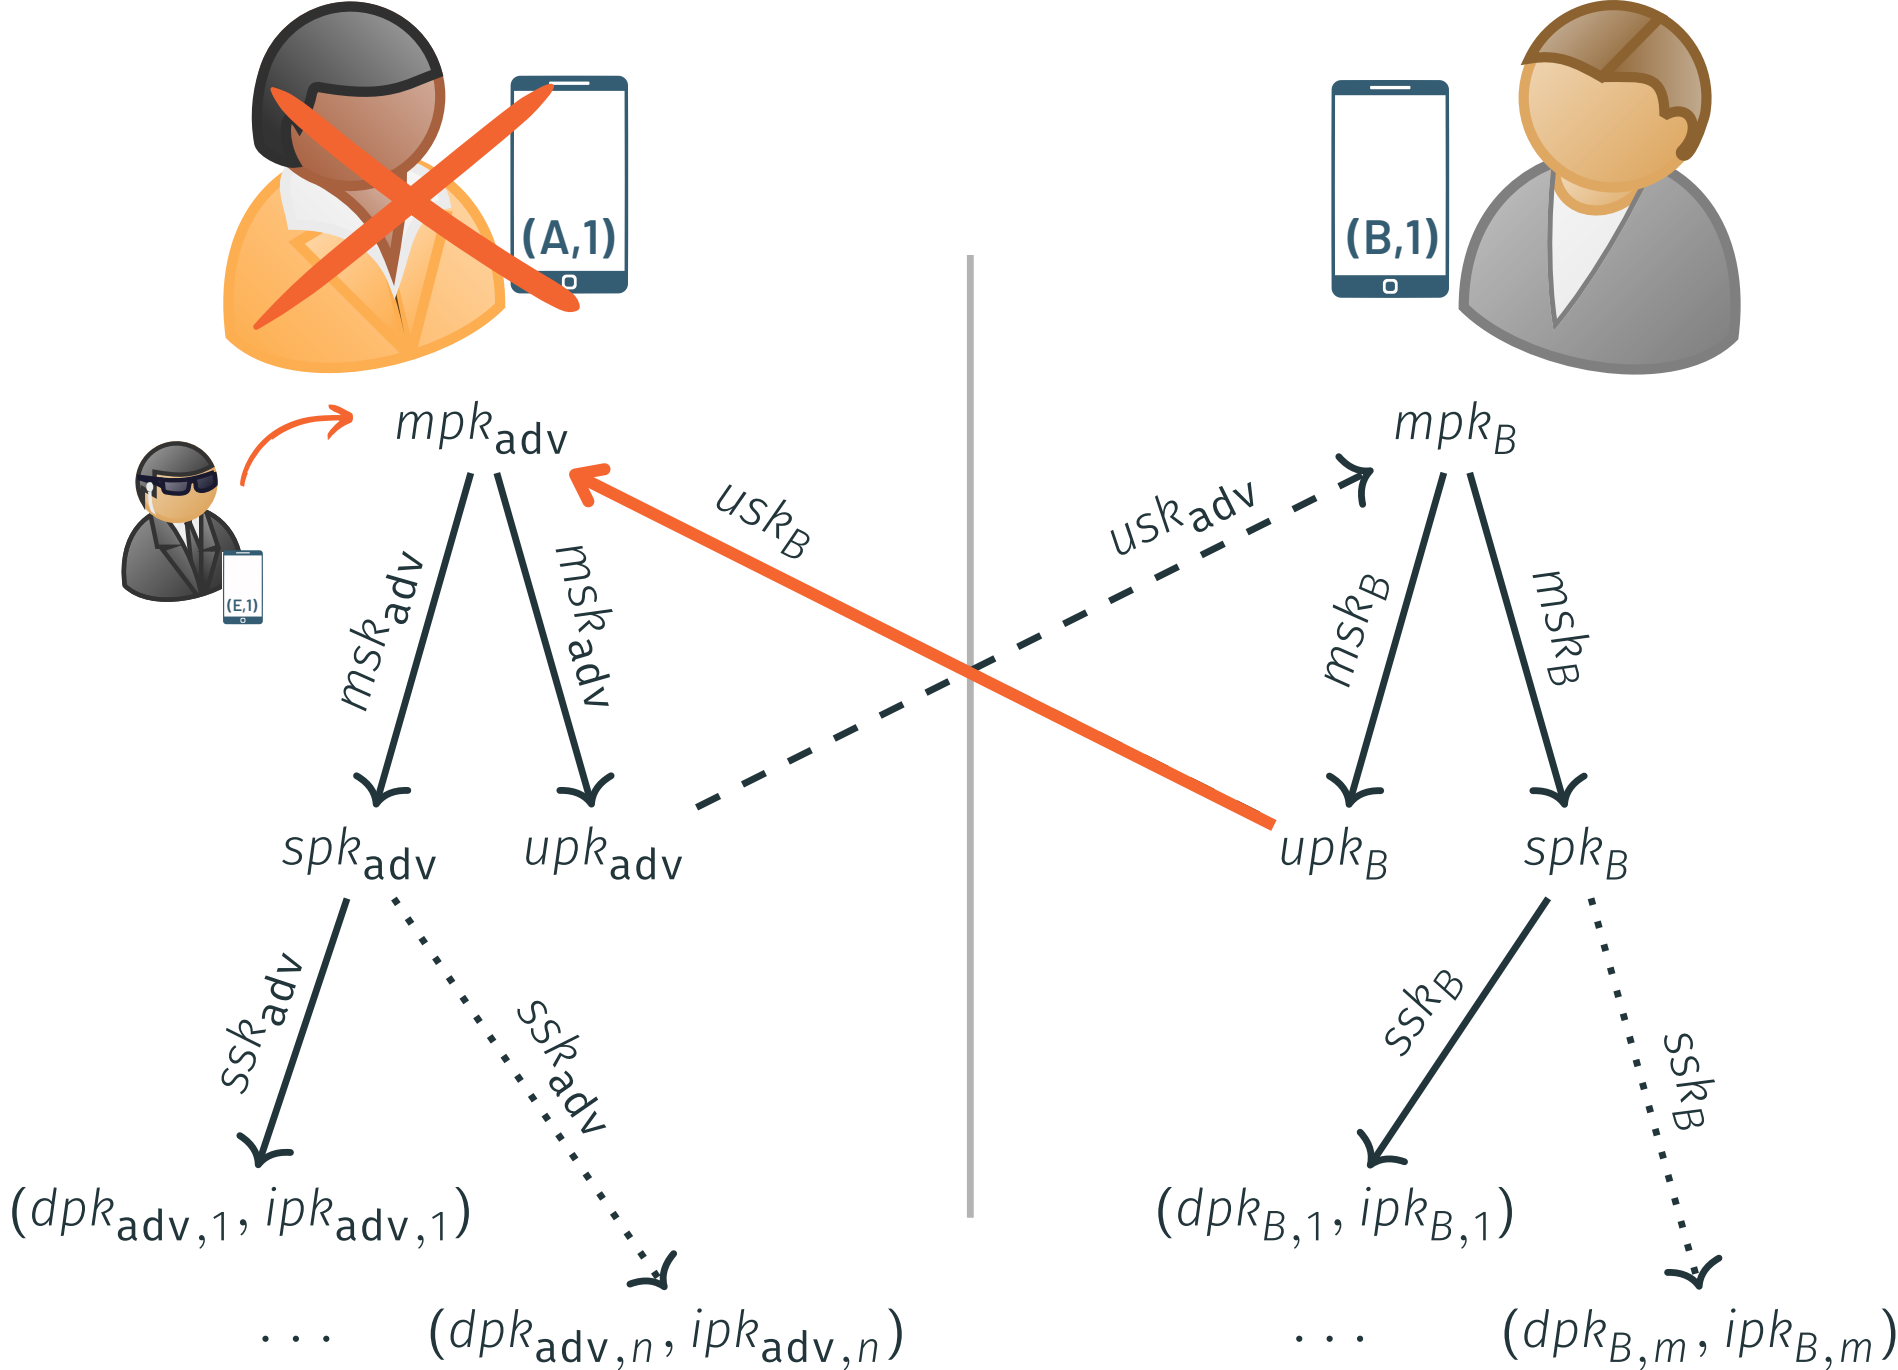
4. Alice's client asks Bob to sign mpk_adv!

User/Device Confusion in Out-of-Band Verification

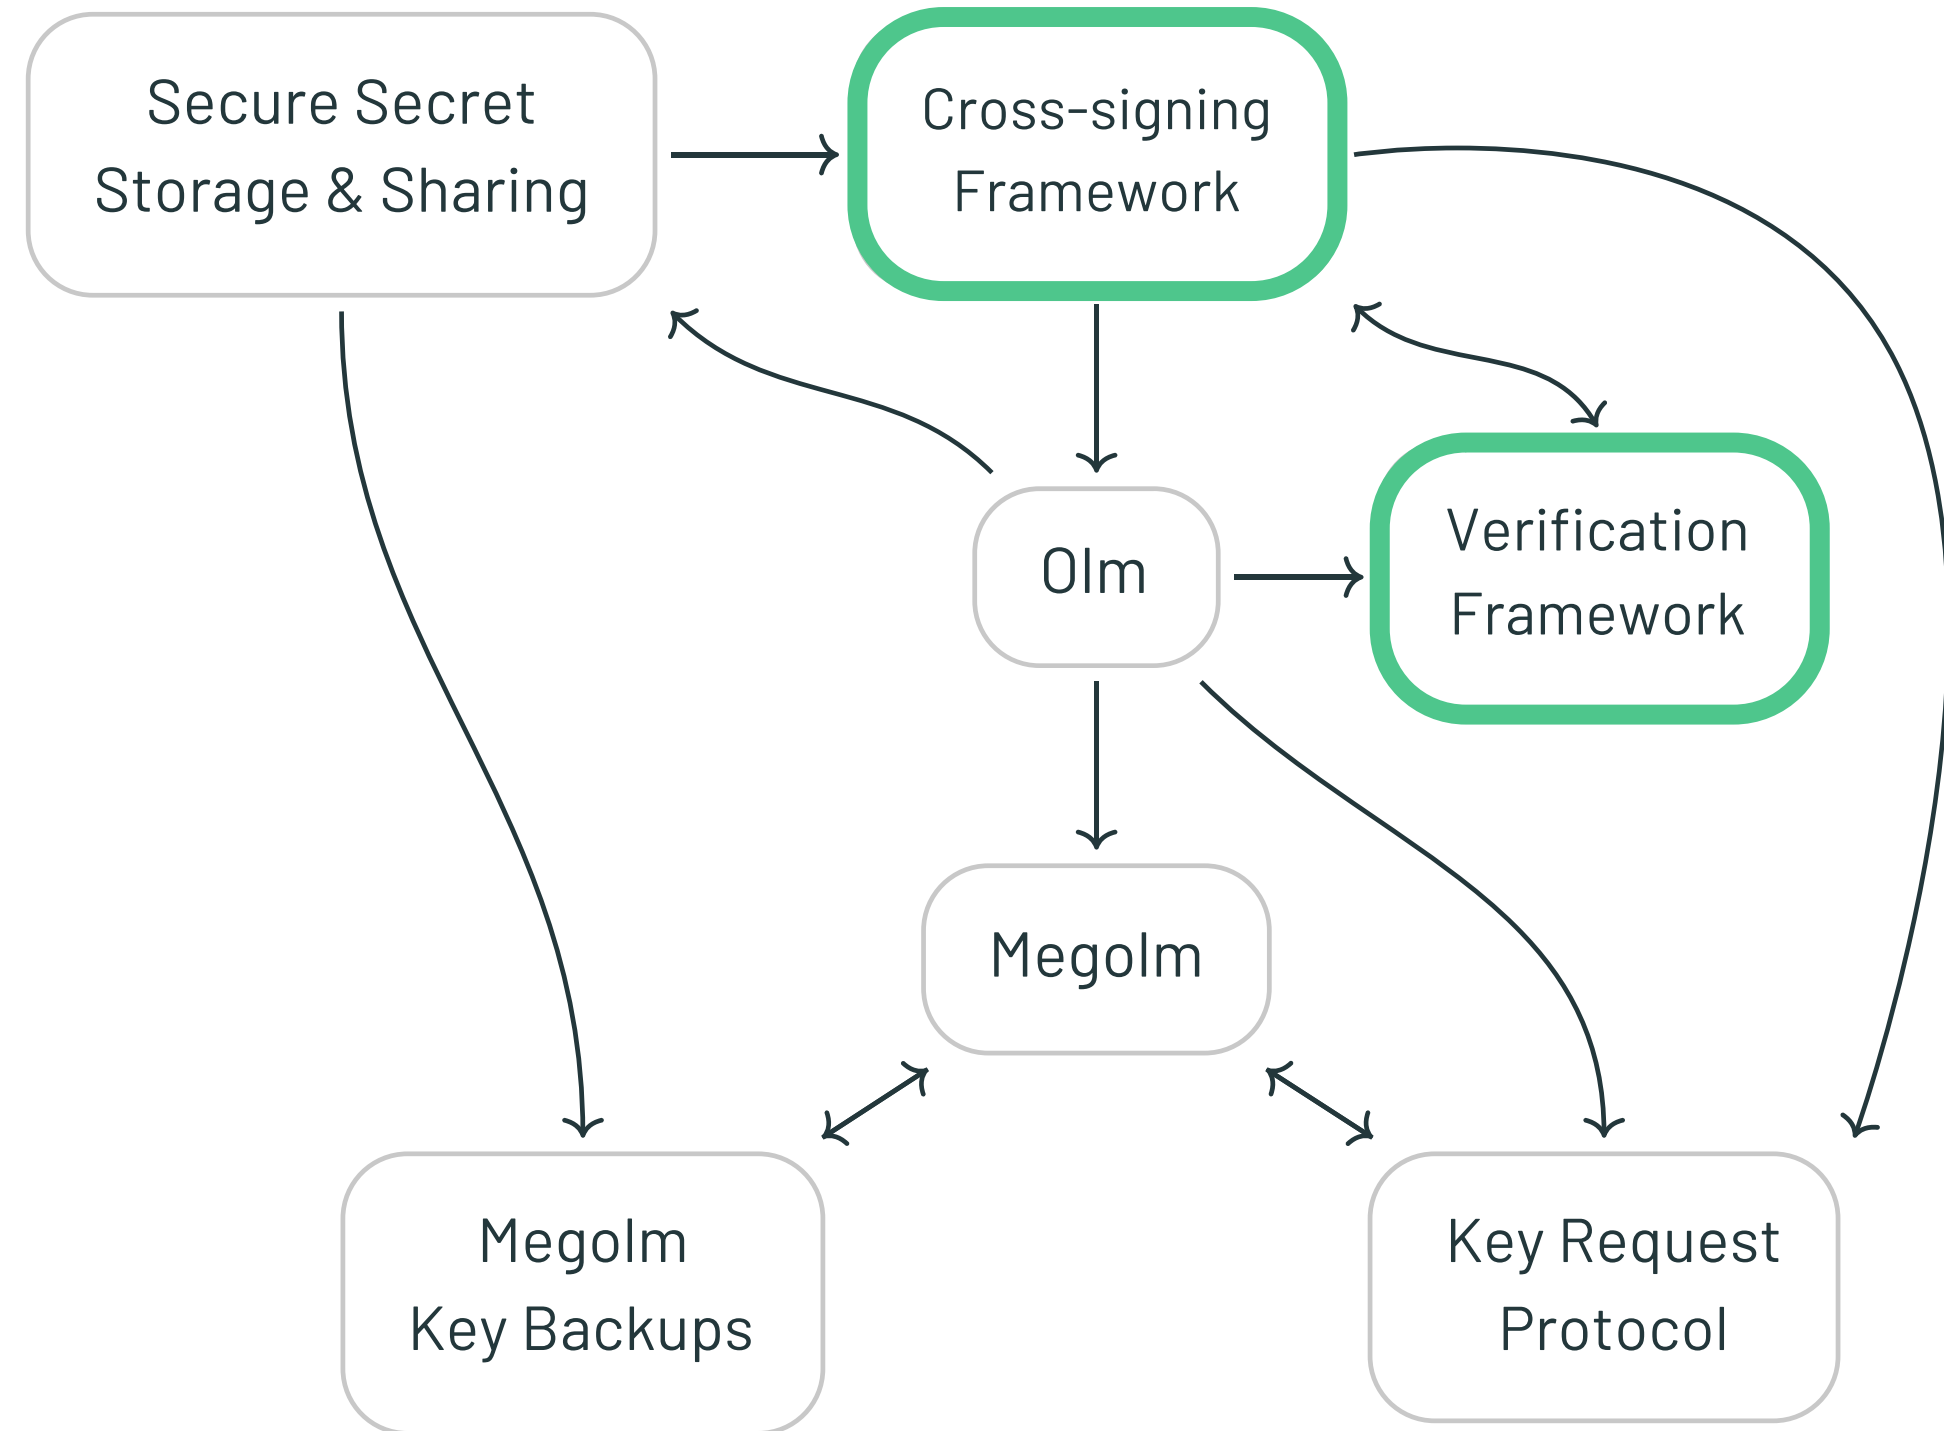
Aim: Trick clients into signing an attacker controlled identity.

If we set Alice's device identifier to a valid cross-signing key, can we trick Bob into signing it as if it were Alice's? **Yes**

Active MITM




Modelling Matrix & Finding Attacks



Megolm

1. Alice initialises and distributes session.

$$(\mathfrak{S}_{gsk}^A, \mathfrak{S}_{gpk}^A, \sigma_{mg}^A) \leftarrow \text{Megolm.Init}()$$


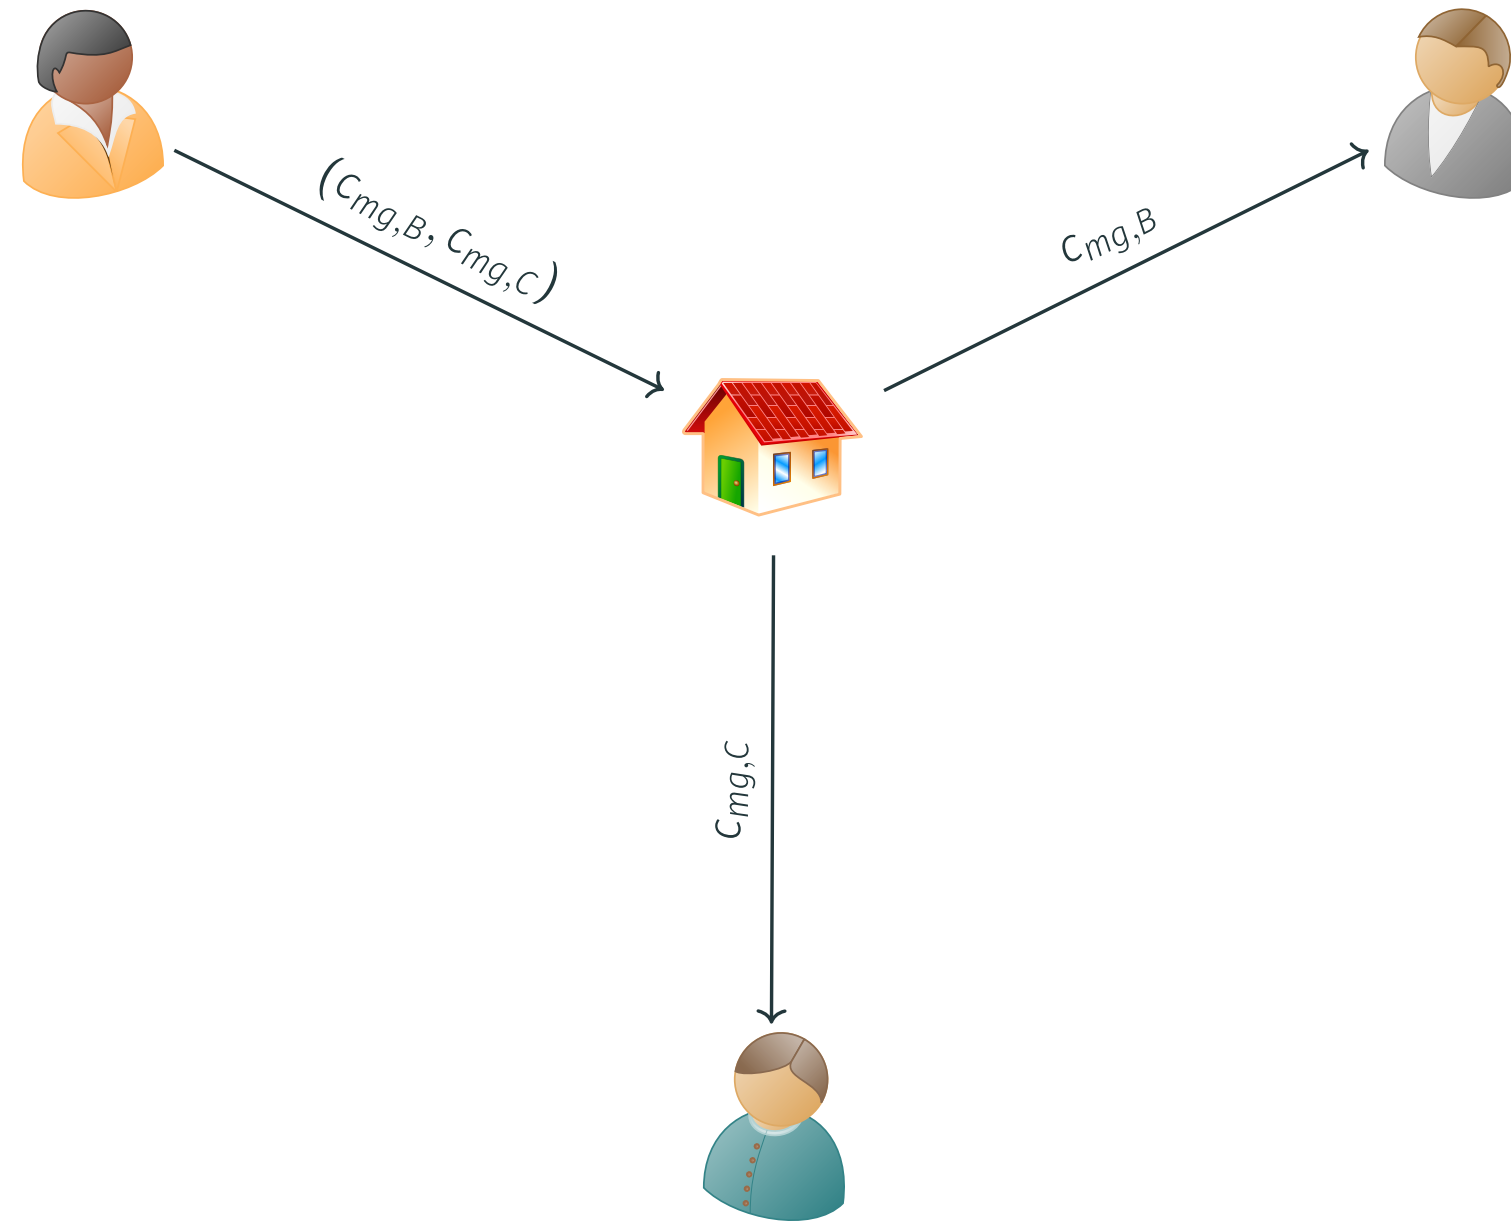
Megolm.Init()

 $i \leftarrow 0$ $R \leftarrow \{0, 1\}^{1024}$ $(gsk, gpk) \leftarrow \text{Ed25519.KGen}(1^n)$ $\text{ver} \leftarrow \mathbf{0x03}$ $\sigma_{mg} \leftarrow \text{Ed25519.Sign}(gsk, (\text{ver}, i, R, gpk))$ $\mathfrak{S}_{gsk} \leftarrow (\text{ver}, i, R, gsk, gpk)$ $\mathfrak{S}_{gpk} \leftarrow (\text{ver}, i, R, gpk)$ $\text{return } \mathfrak{S}_{gsk}, \mathfrak{S}_{gpk}, \sigma_{mg}$

Megolm

1. Alice initialises and distributes session.

$(\mathfrak{S}_{gsk}^A, \mathfrak{S}_{gpk}^A, \sigma_{mg}^A) \leftarrow \text{Megolm.Init}()$
 $(\pi_{olm}^B, c_{mg,B}) \leftarrow \text{Olm.Enc}(\pi_{olm}^B, \mathfrak{S}_{gpk}^A \parallel \sigma_{mg}^A)$
 $(\pi_{olm}^C, c_{mg,C}) \leftarrow \text{Olm.Enc}(\pi_{olm}^C, \mathfrak{S}_{gpk}^A \parallel \sigma_{mg}^A)$



Megolm

2. Bob and Claire receive and verify Alice's session.

$$\begin{aligned} (\mathfrak{S}_{gsk}^A, \mathfrak{S}_{gpk}^A, \sigma_{mg}^A) &\leftarrow \text{Megolm.Init}() \\ (\pi_{olm}^B, c_{mg,B}) &\leftarrow \text{Olm.Enc}(\pi_{olm}^B, \mathfrak{S}_{gpk}^A \parallel \sigma_{mg}^A) \\ (\pi_{olm}^C, c_{mg,C}) &\leftarrow \text{Olm.Enc}(\pi_{olm}^C, \mathfrak{S}_{gpk}^A \parallel \sigma_{mg}^A) \end{aligned}$$



$(c_{mg,B}, c_{mg,C})$



$c_{mg,B}$



$$(\pi_{olm}^A, \mathfrak{S}_{gpk}^A \parallel \sigma_{mg}^A) \leftarrow \text{Olm.Dec}(\pi_{olm}^A, c_{mg,B})$$

$c_{mg,C}$



$$(\pi_{olm}^A, \mathfrak{S}_{gpk}^A \parallel \sigma_{mg}^A) \leftarrow \text{Olm.Dec}(\pi_{olm}^A, c_{mg,C})$$

Megolm

3. Bob and Claire generate and distribute their own sessions.

Megolm Sessions

$(\mathcal{G}_{gsk}^A, \mathcal{G}_{gpk}^A, \sigma_{mg}^A)$
 $(B, \mathcal{G}_{gpk}^B, \sigma_{mg}^B)$
 $(C, \mathcal{G}_{gpk}^C, \sigma_{mg}^C)$



Megolm Sessions

$(\mathcal{G}_{gsk}^B, \mathcal{G}_{gpk}^B, \sigma_{mg}^B)$
 $(A, \mathcal{G}_{gpk}^A, \sigma_{mg}^A)$
 $(C, \mathcal{G}_{gpk}^C, \sigma_{mg}^C)$



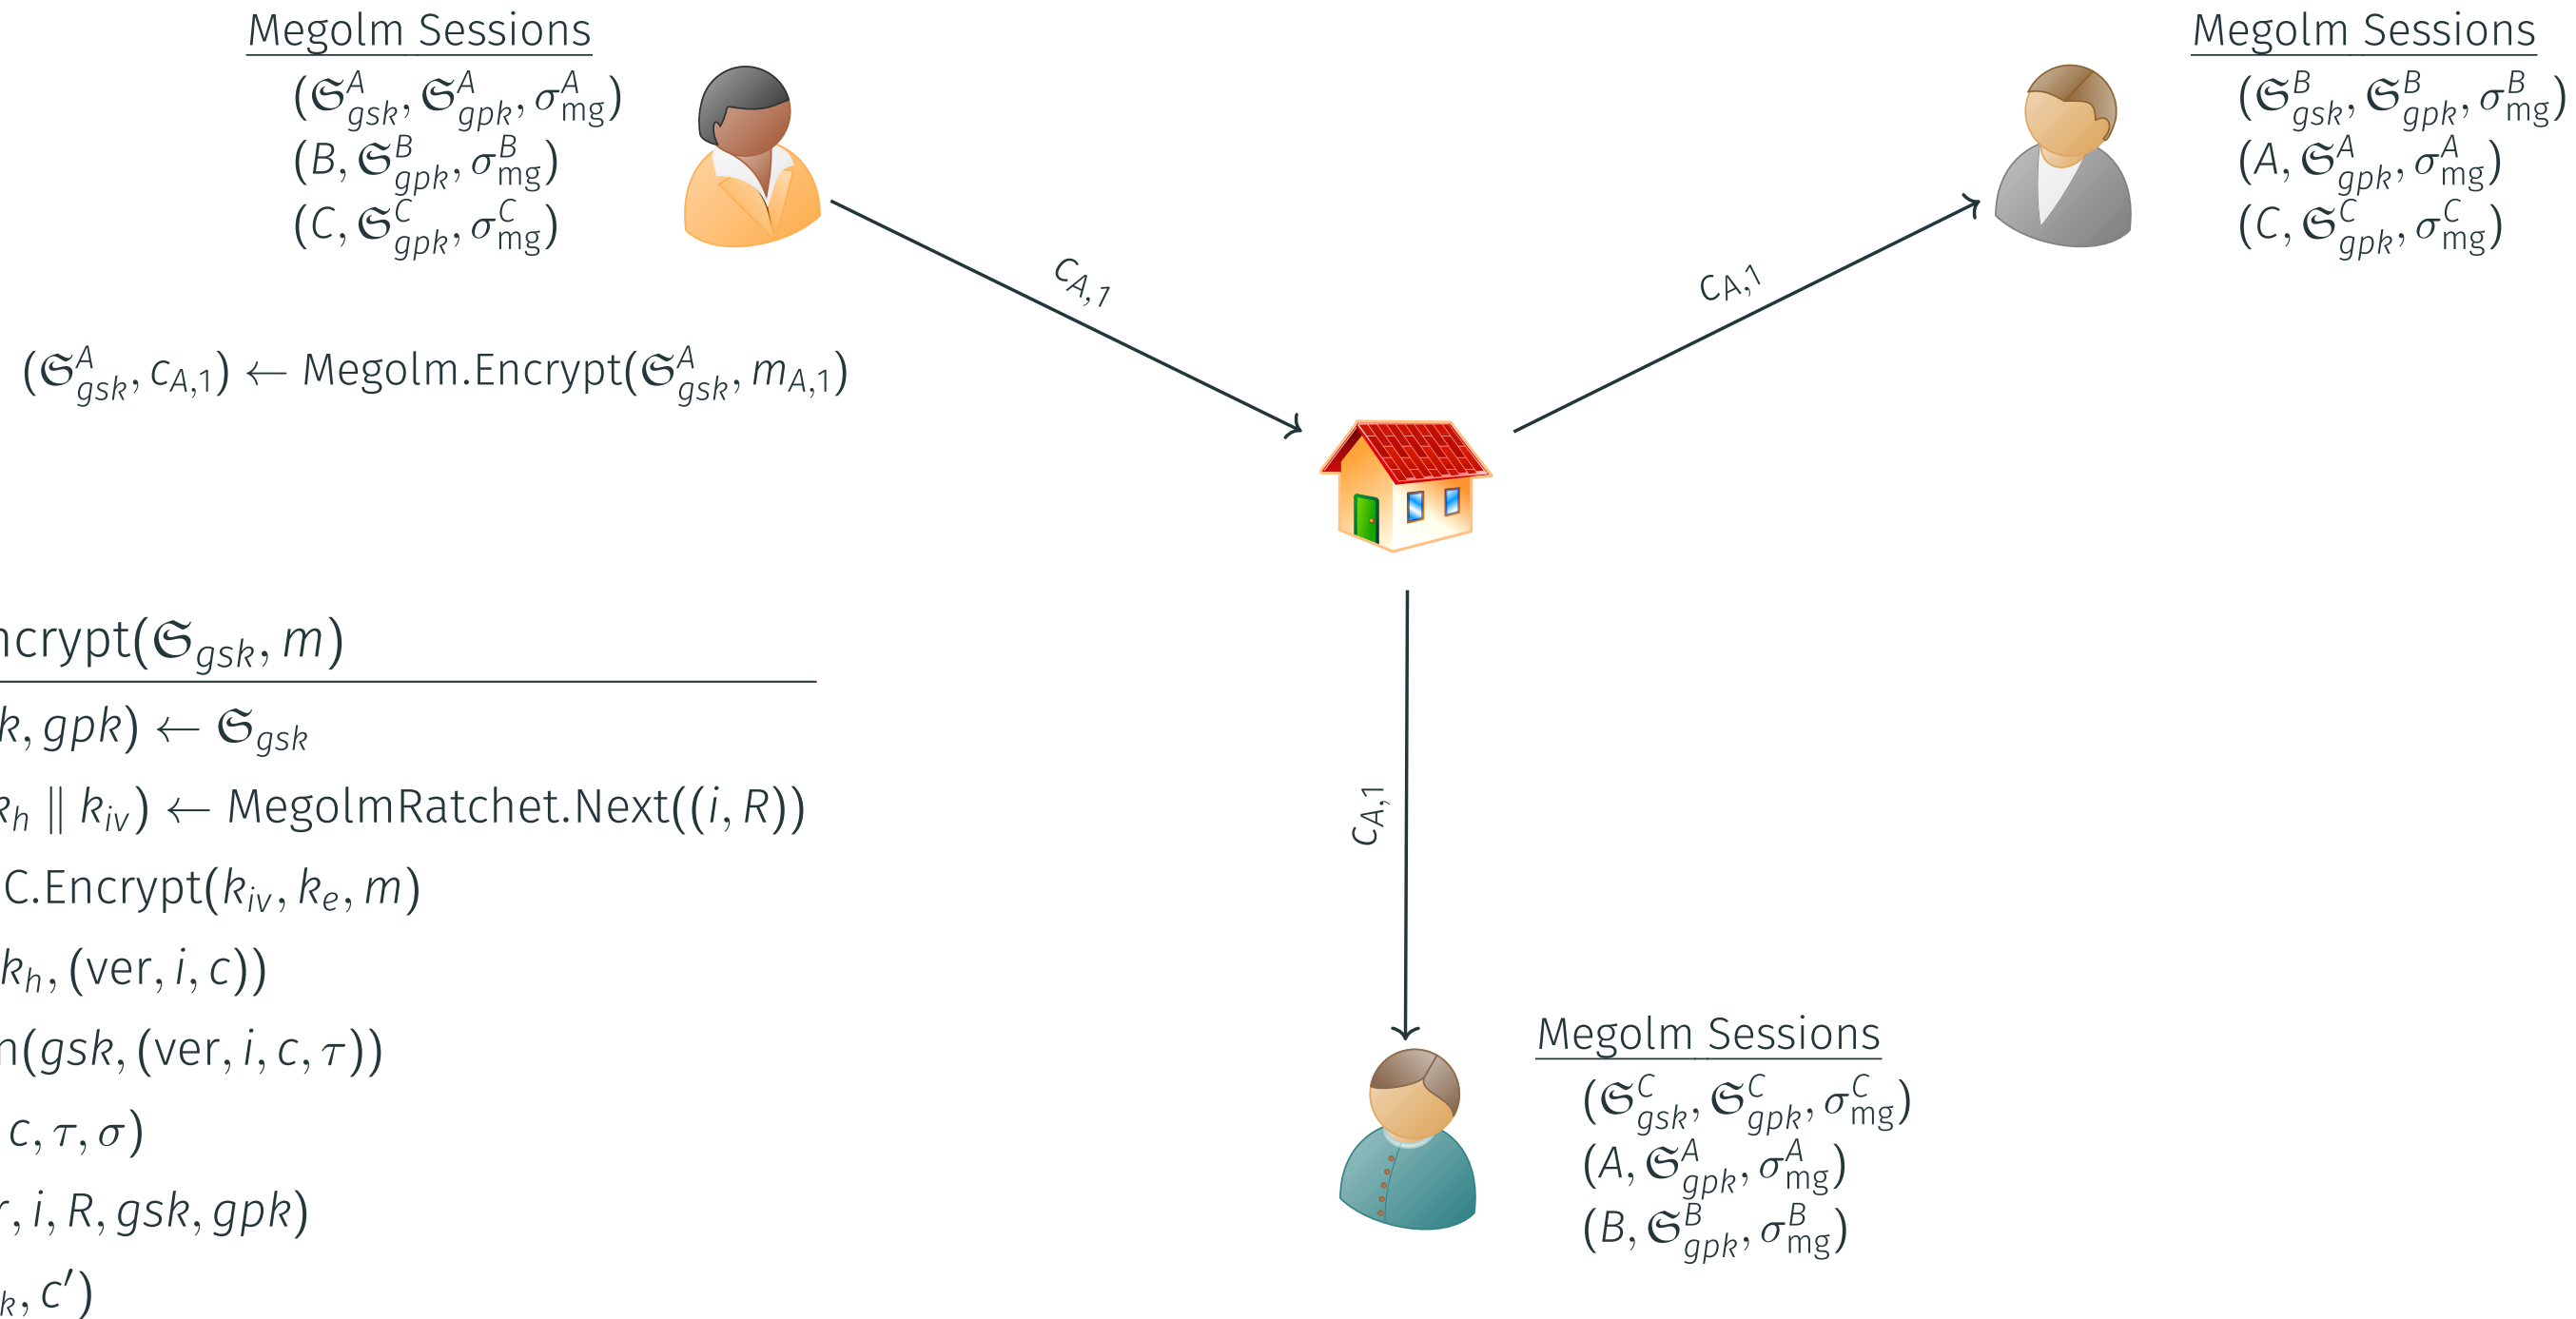
Megolm Sessions

$(\mathcal{G}_{gsk}^C, \mathcal{G}_{gpk}^C, \sigma_{mg}^C)$
 $(A, \mathcal{G}_{gpk}^A, \sigma_{mg}^A)$
 $(B, \mathcal{G}_{gpk}^B, \sigma_{mg}^B)$



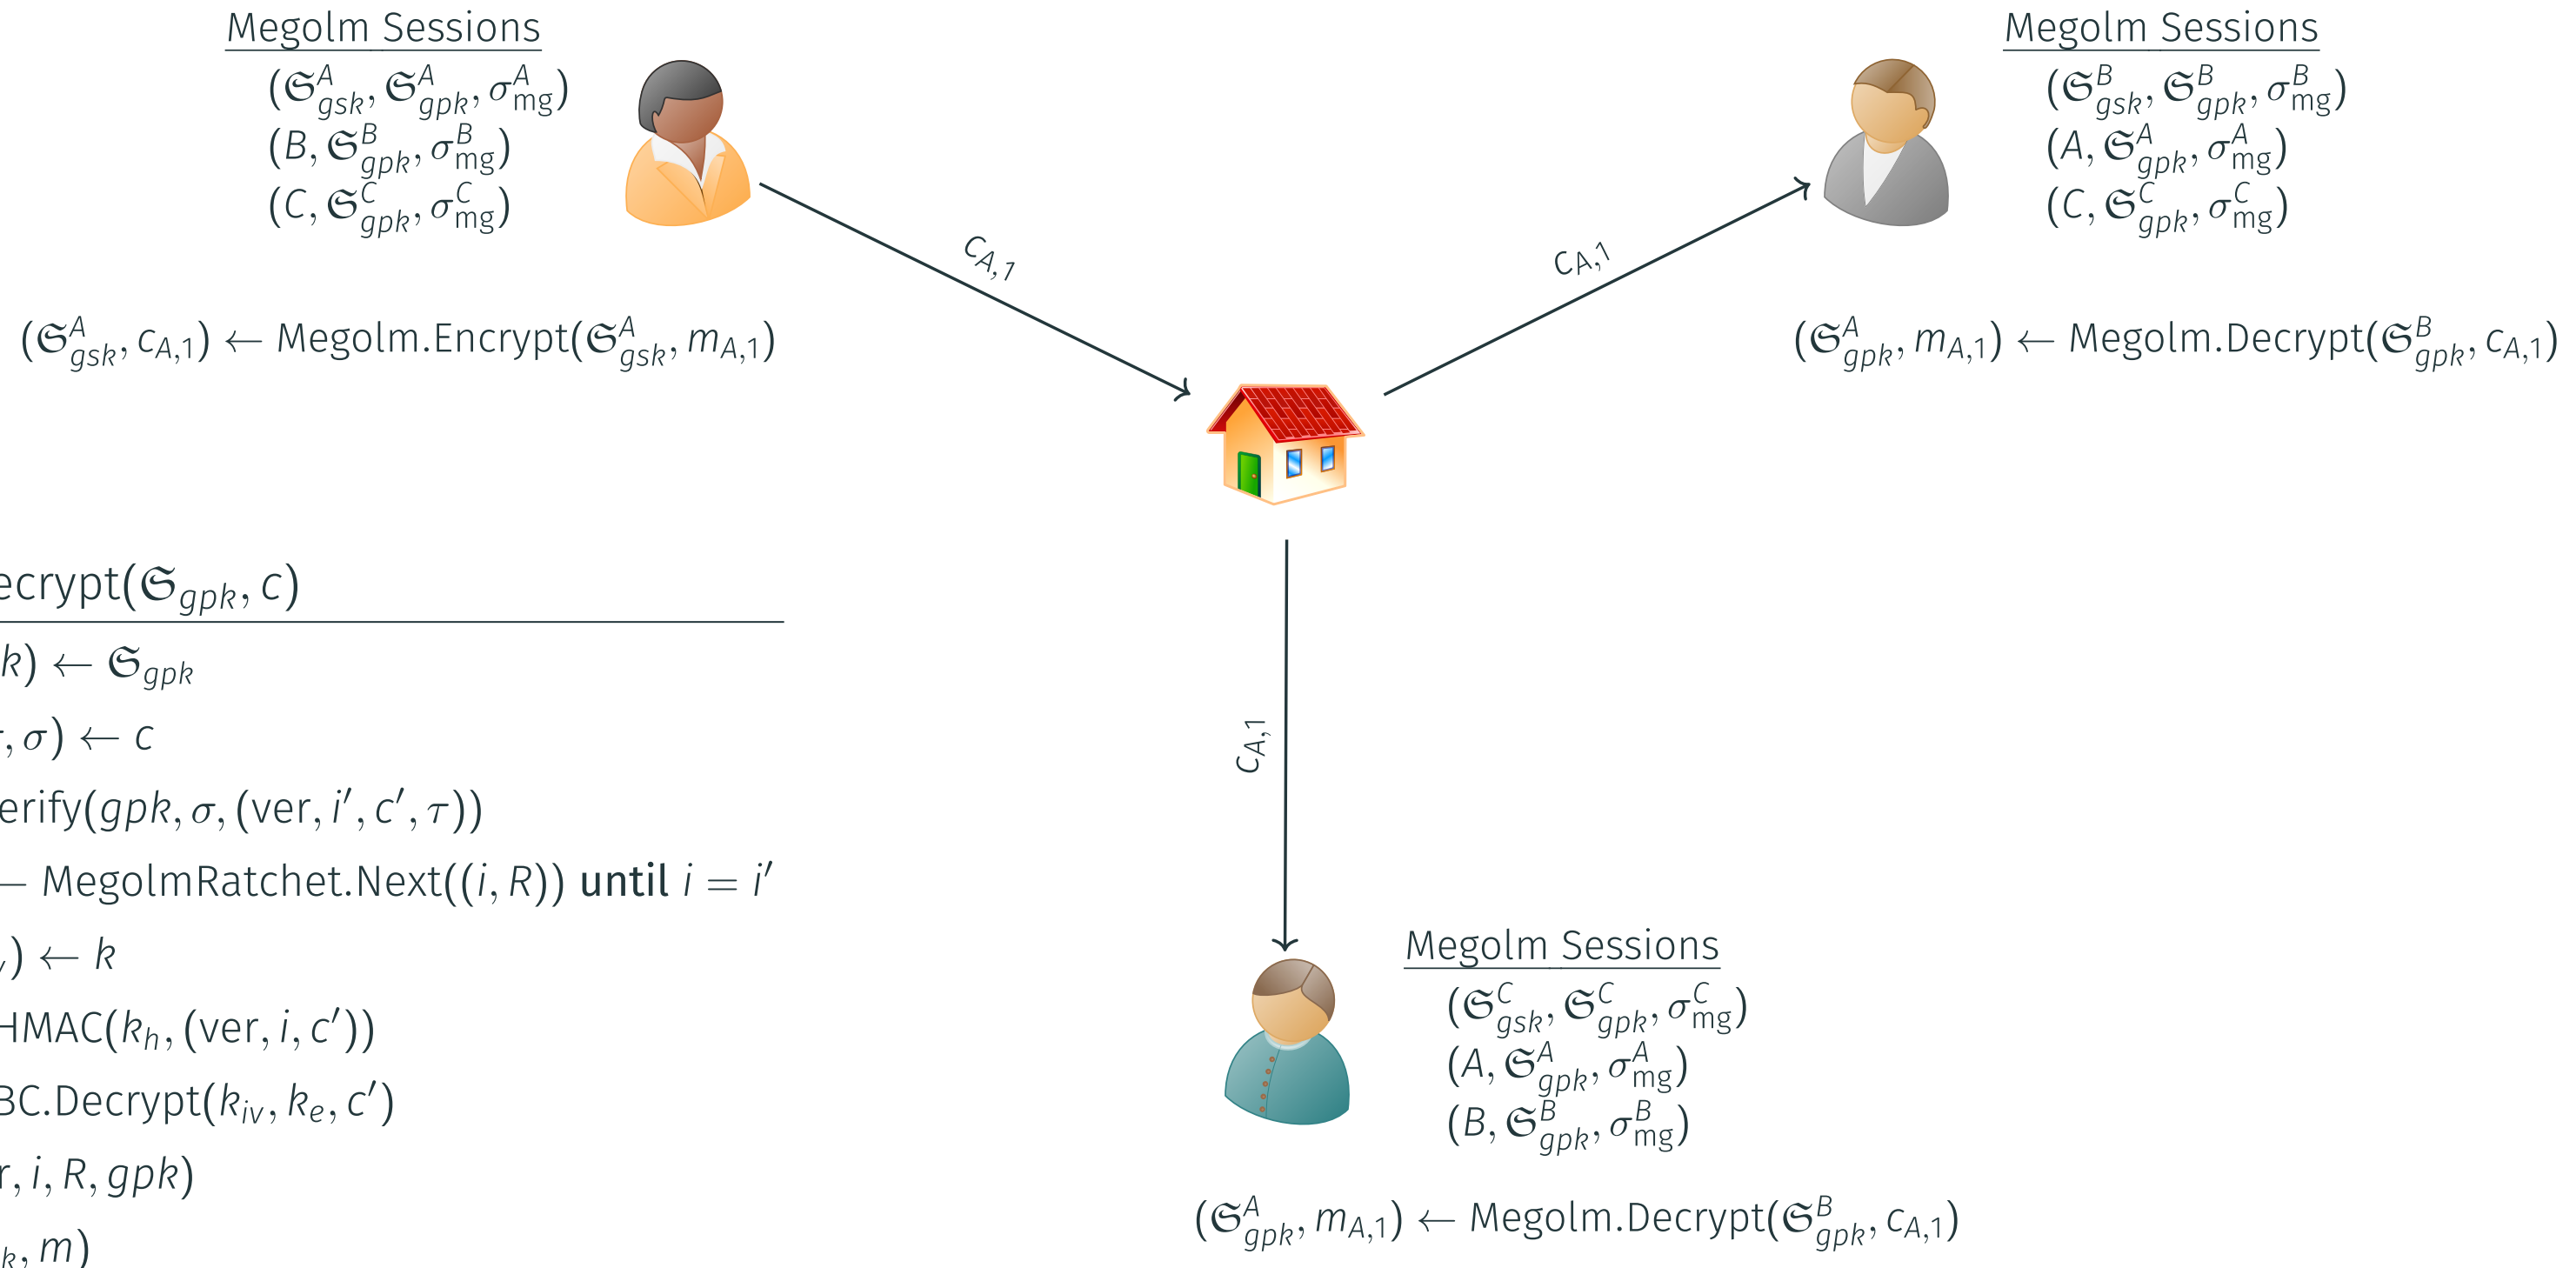
Megolm

4. Alice sends a Megolm message to Bob and Claire.

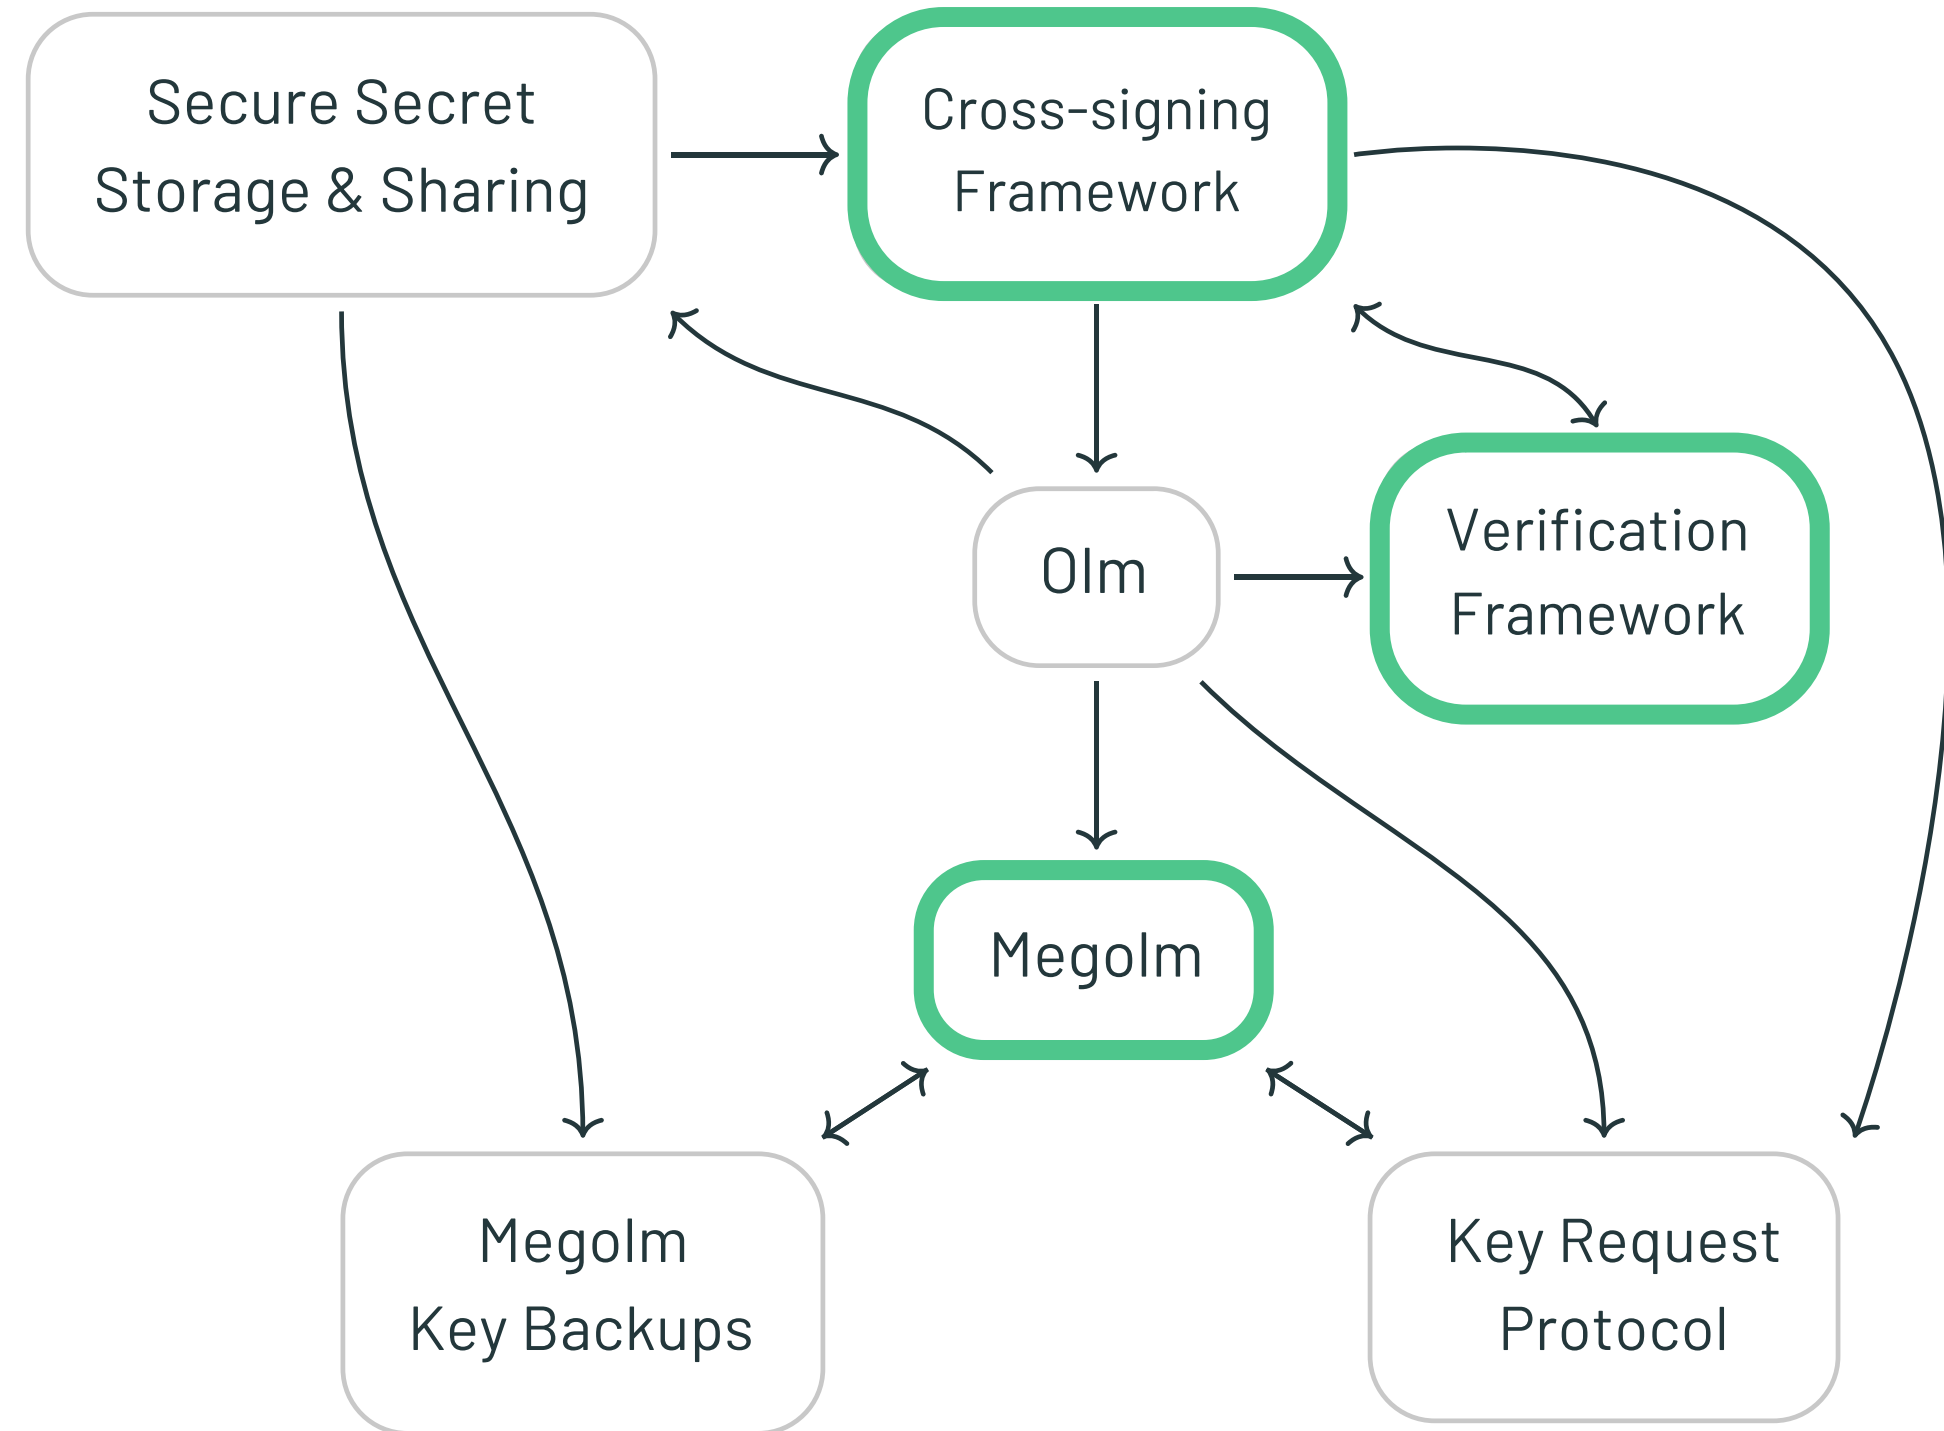


Megolm

5. Bob and Claire decrypt and verify Alice's message.

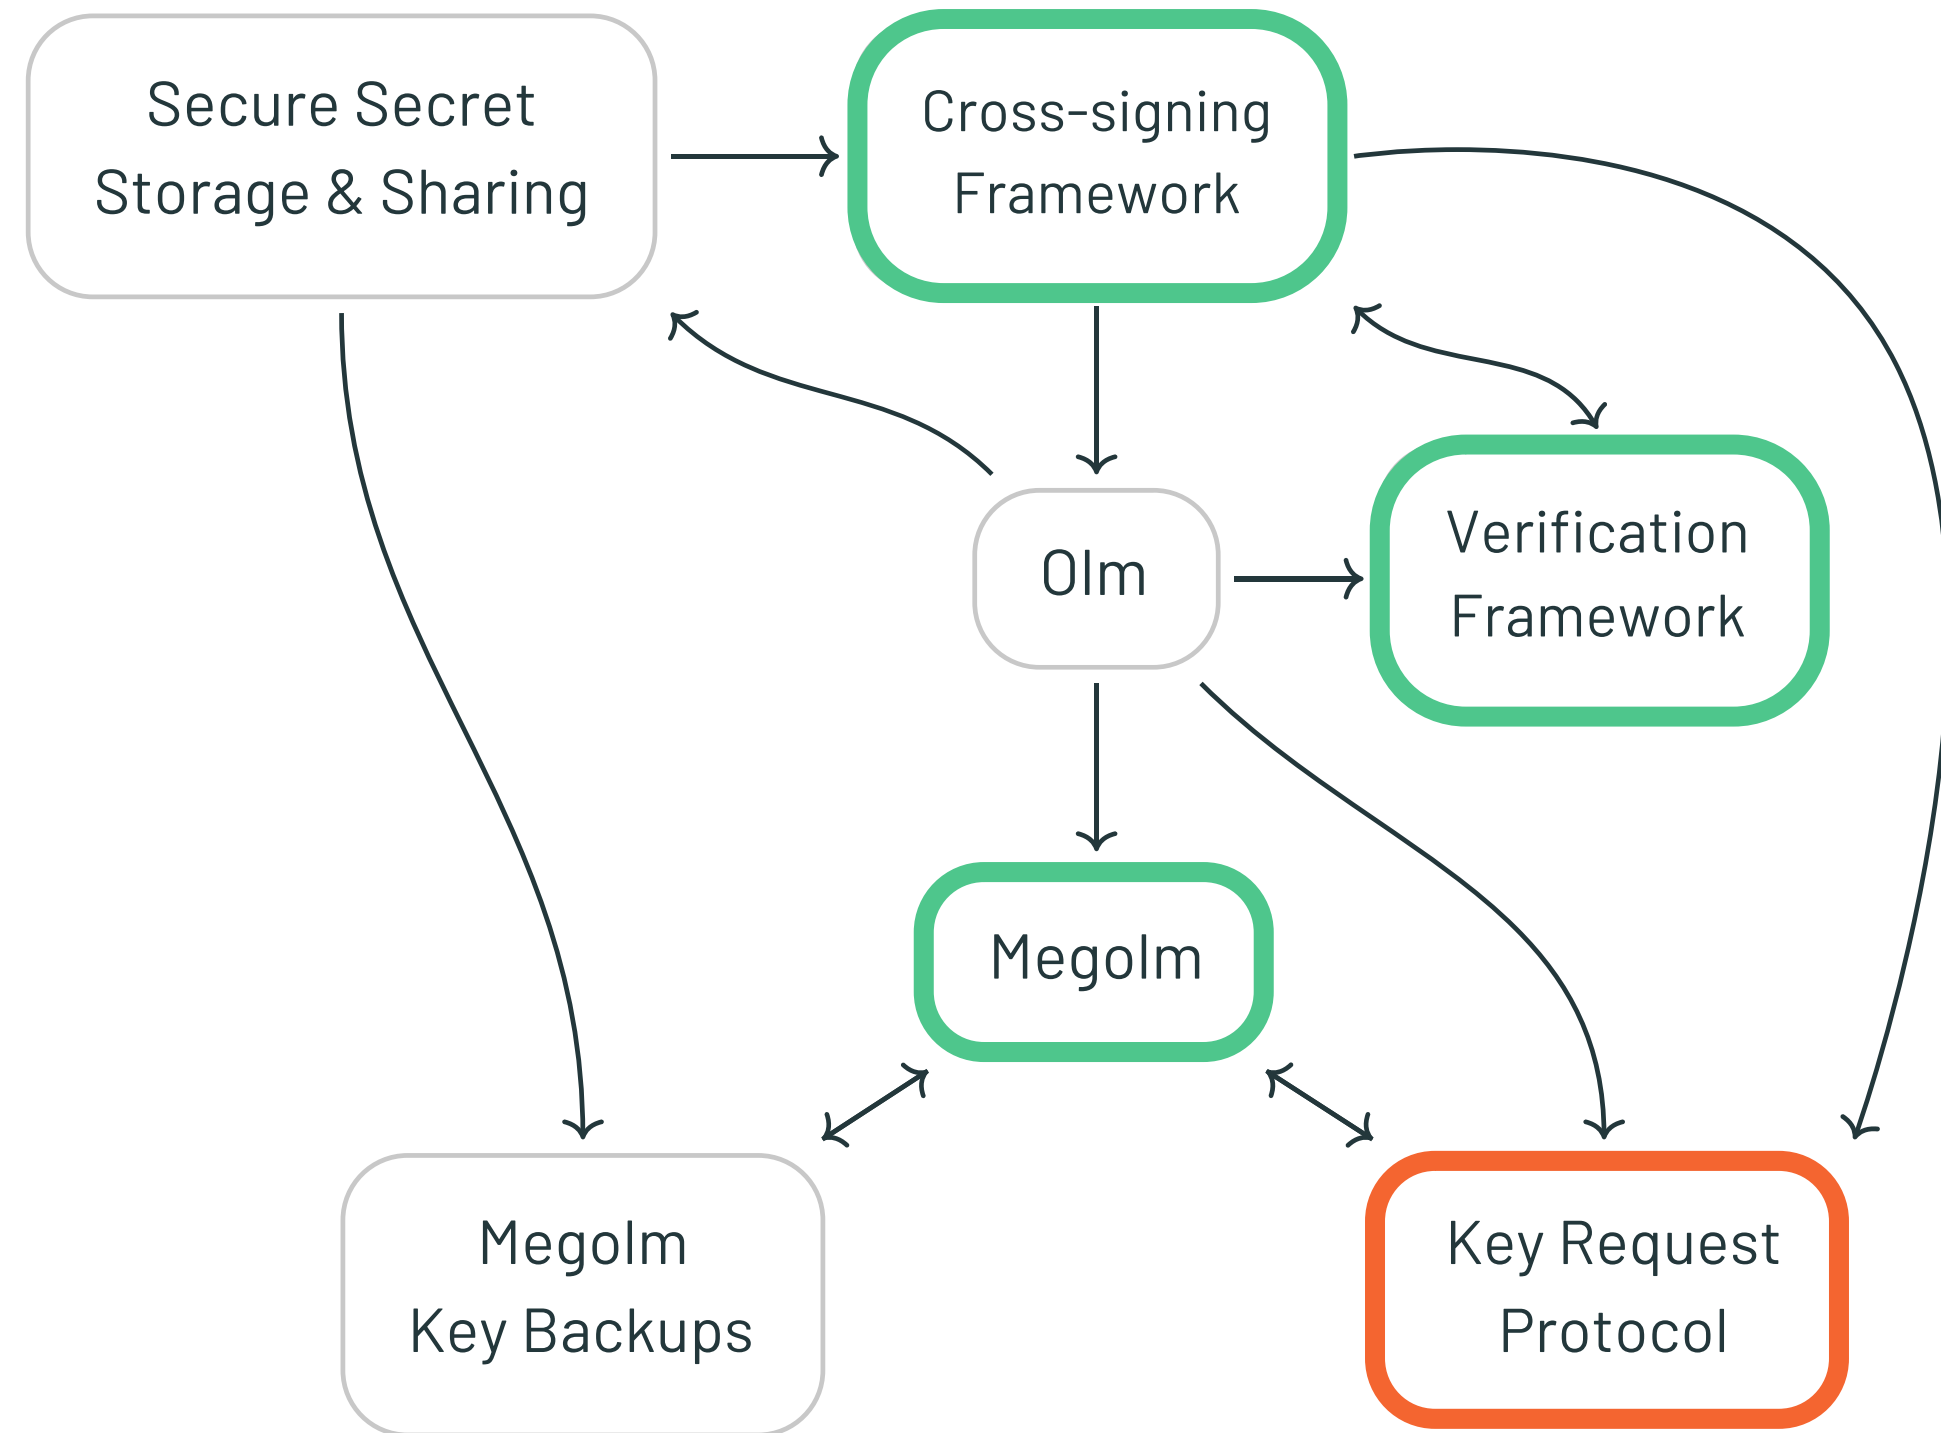


Modelling Matrix & Finding Attacks



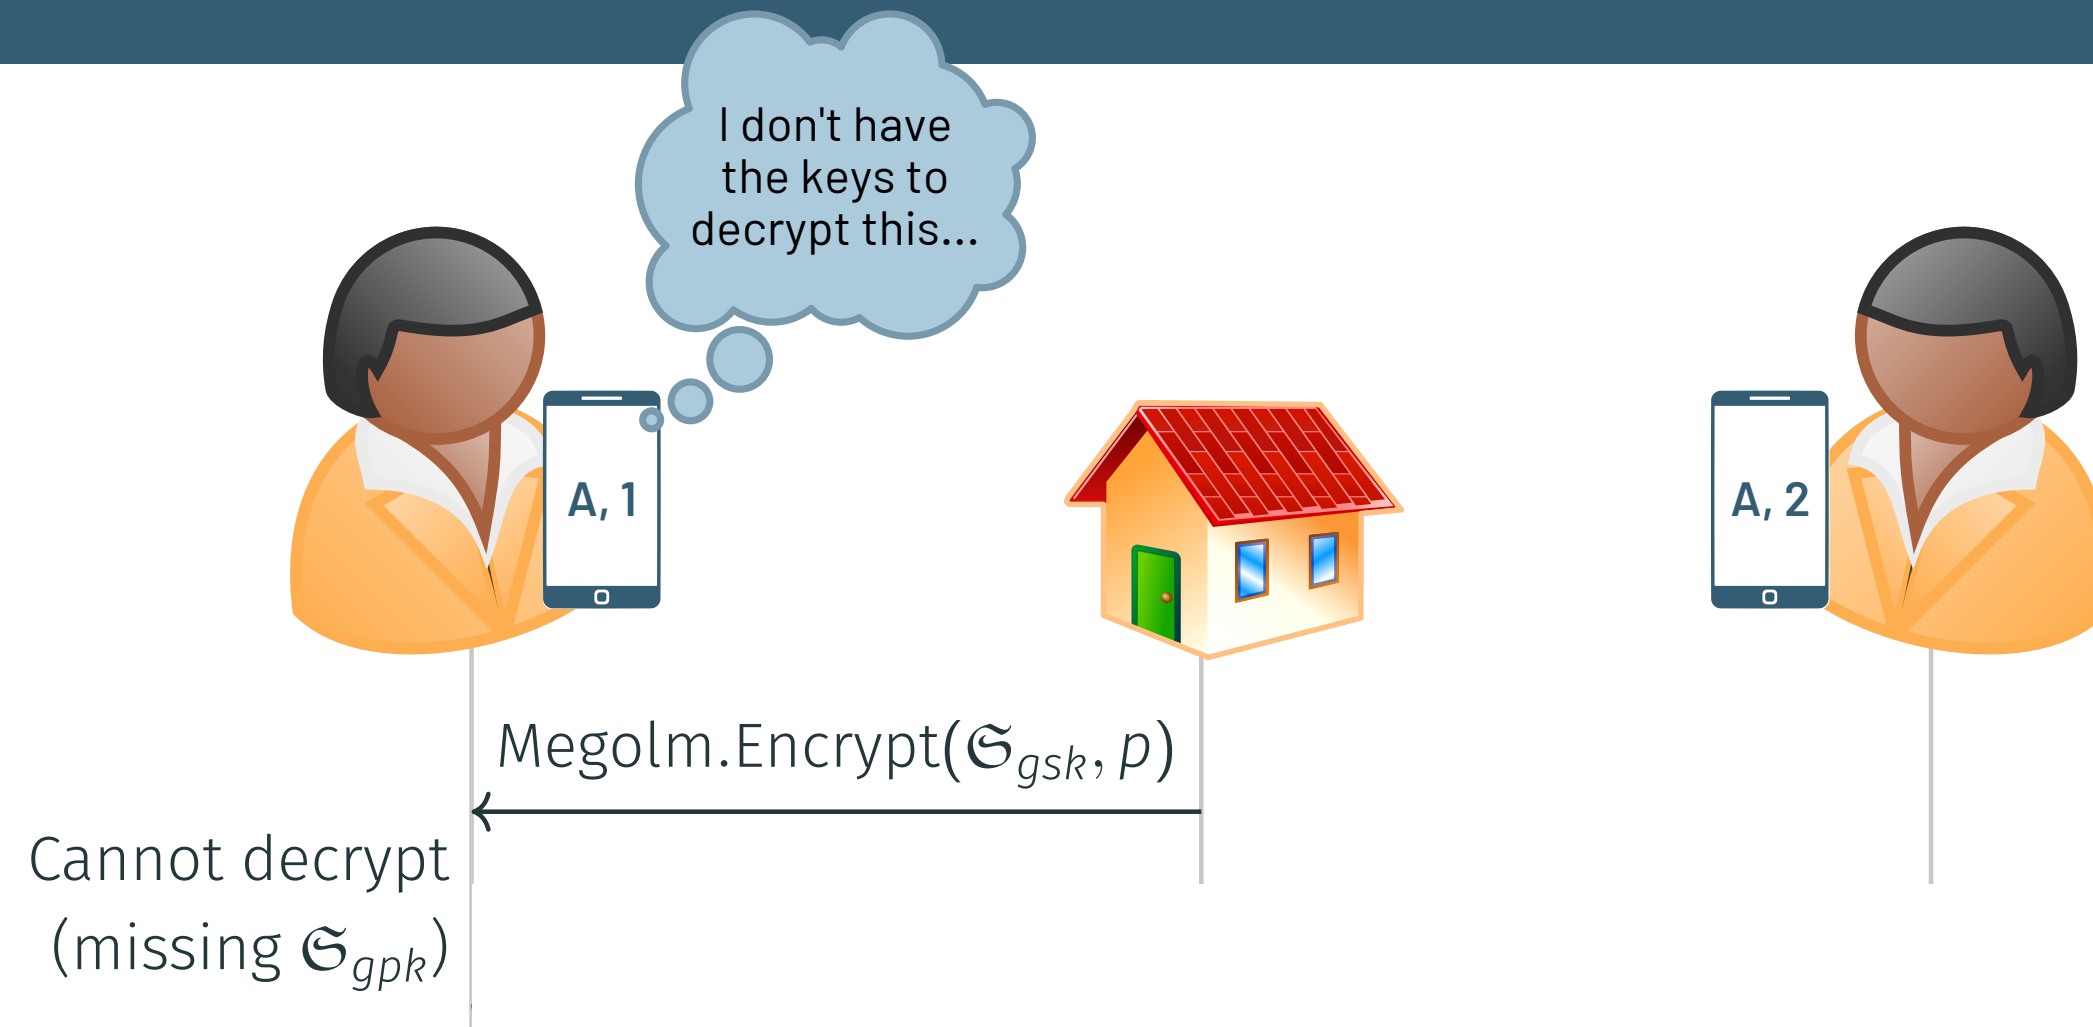
Key Request Protocol

- Request-response protocol.
- Allows devices to request and share keys between each other.
- Responding device must ensure requesting device is entitled to the keys.
- Keys are shared over Olm.



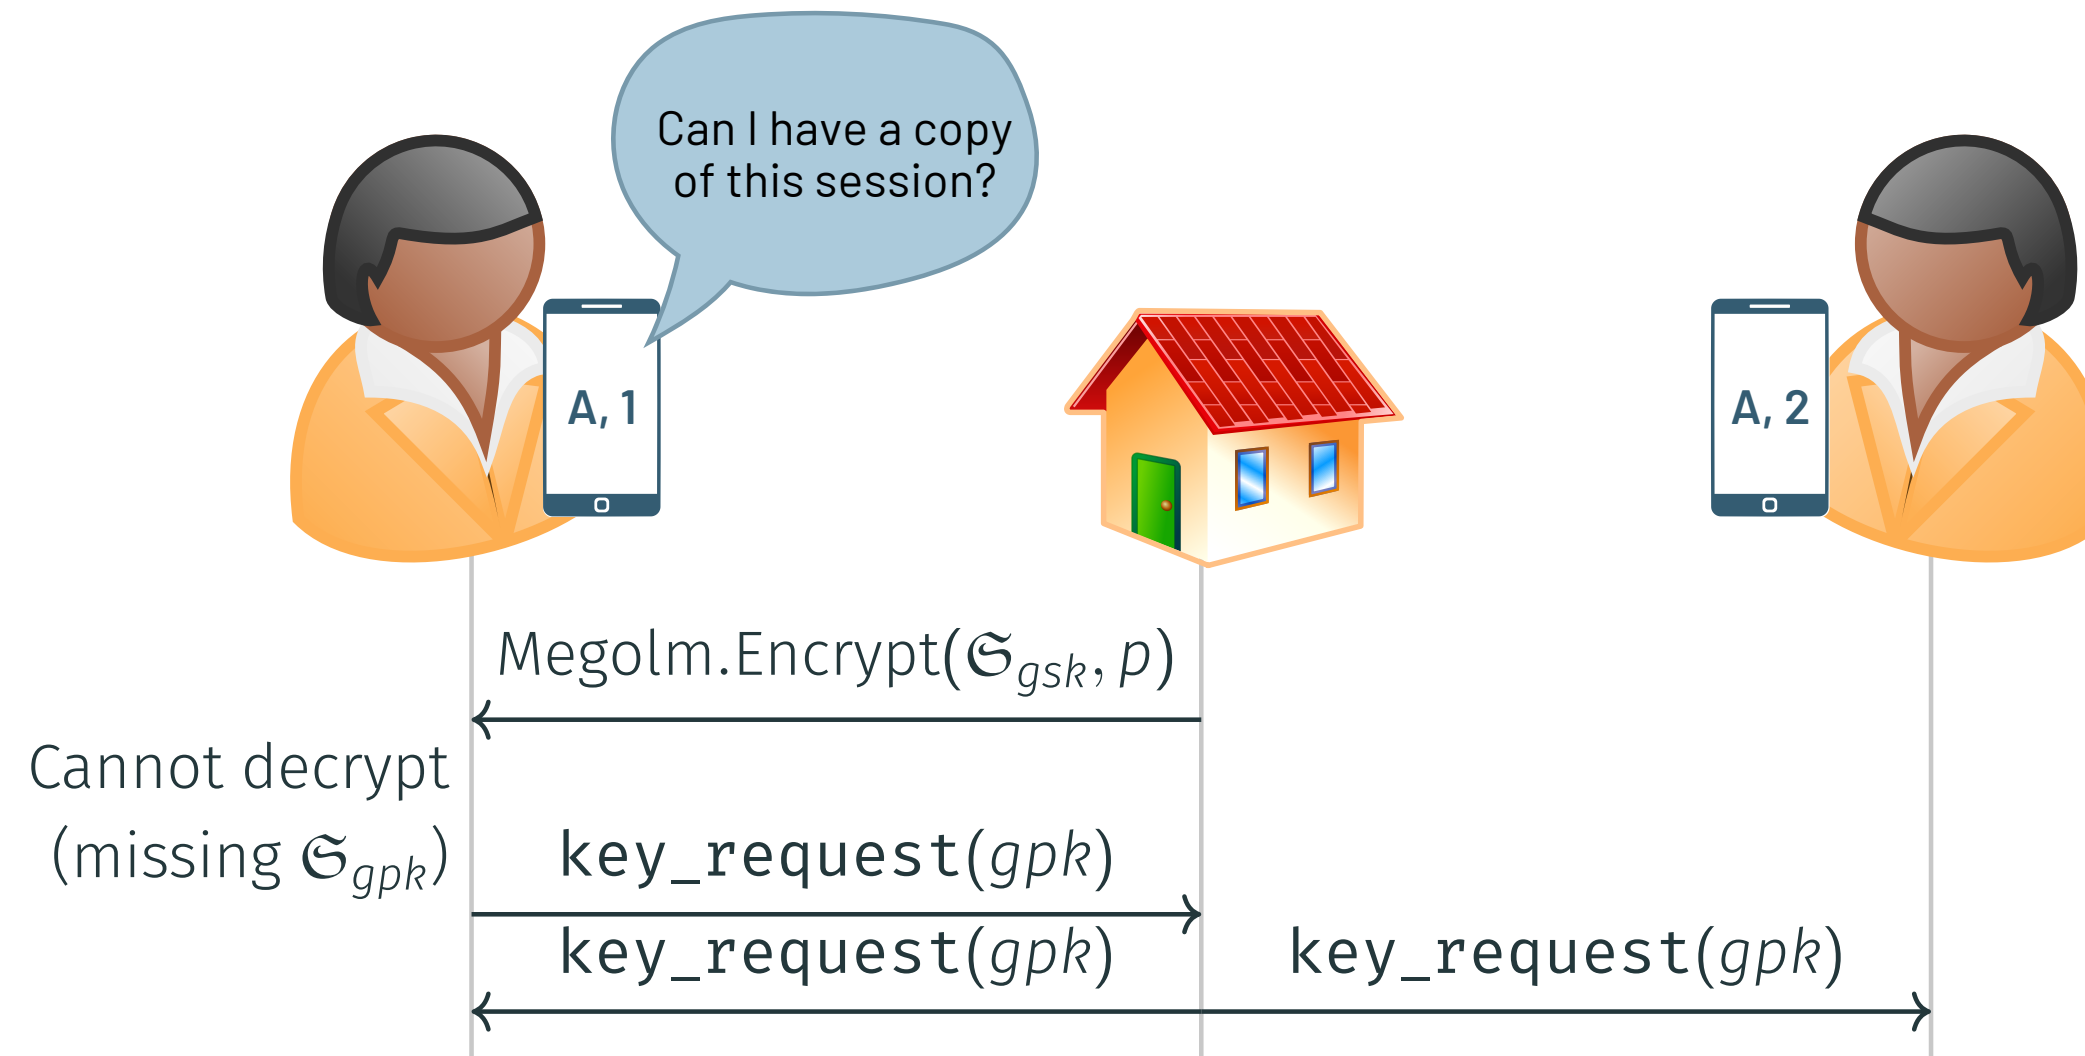
Key Request Protocol

1. Alice's device receives a ciphertext it can't decrypt.



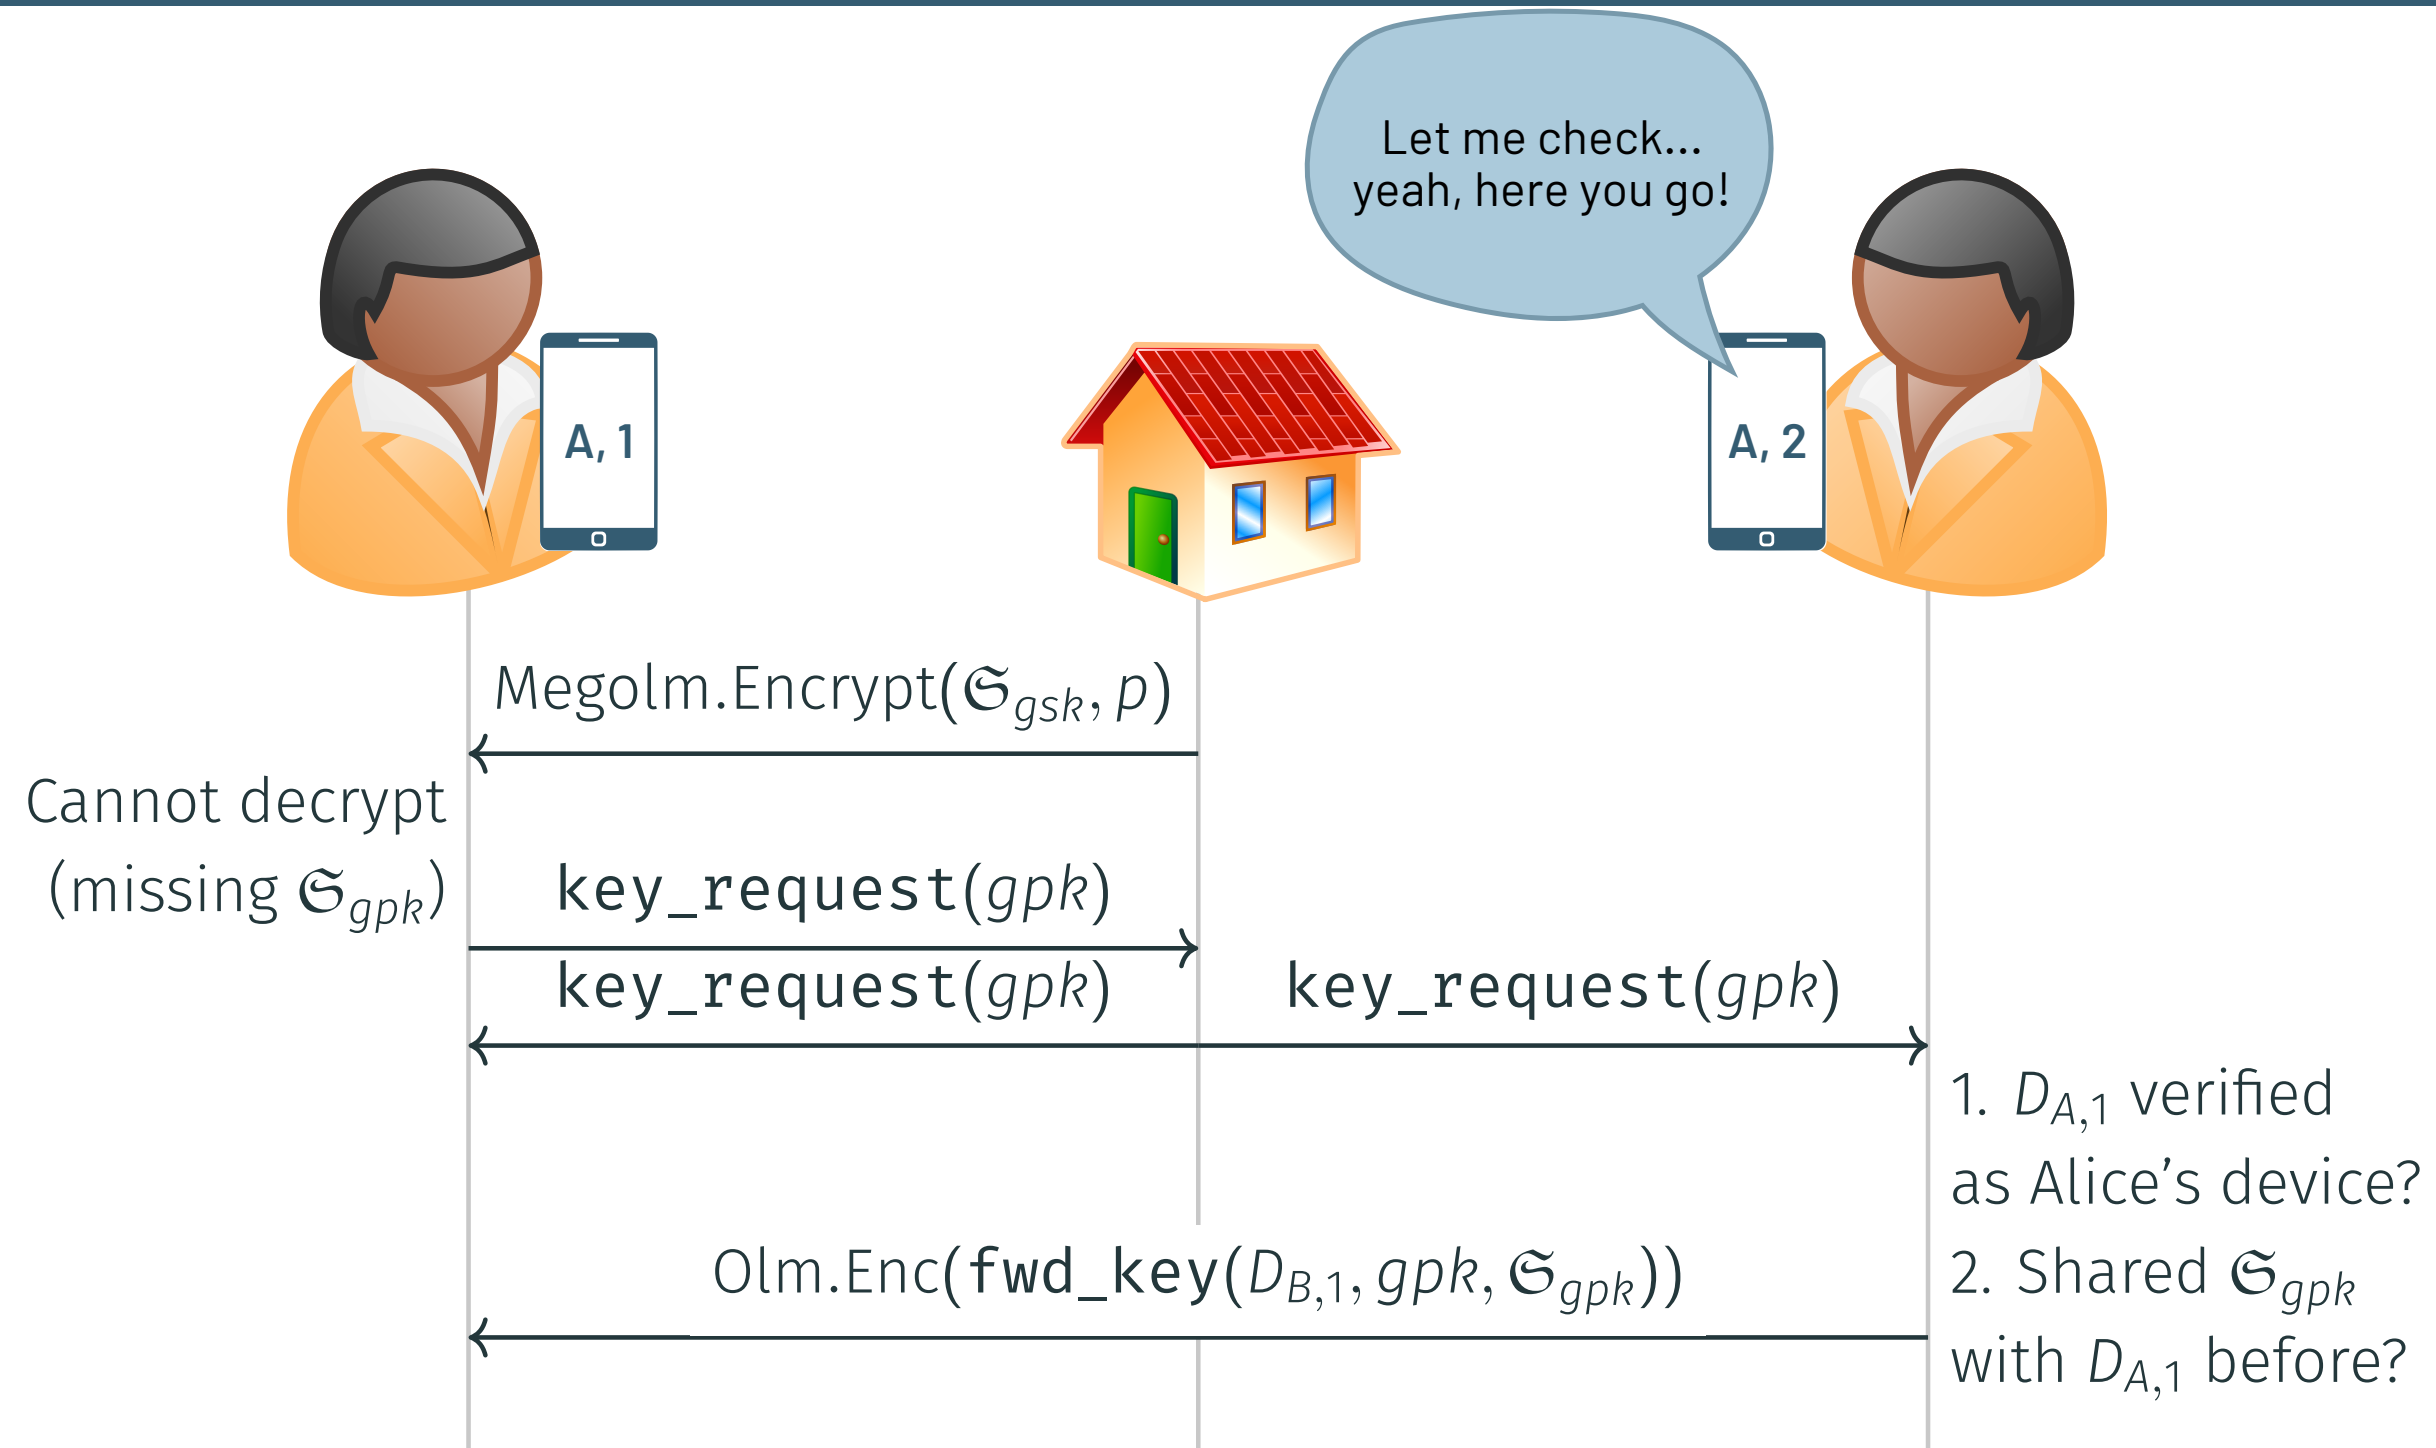
Key Request Protocol

2. Alice's client requests a copy of the decryption keys.



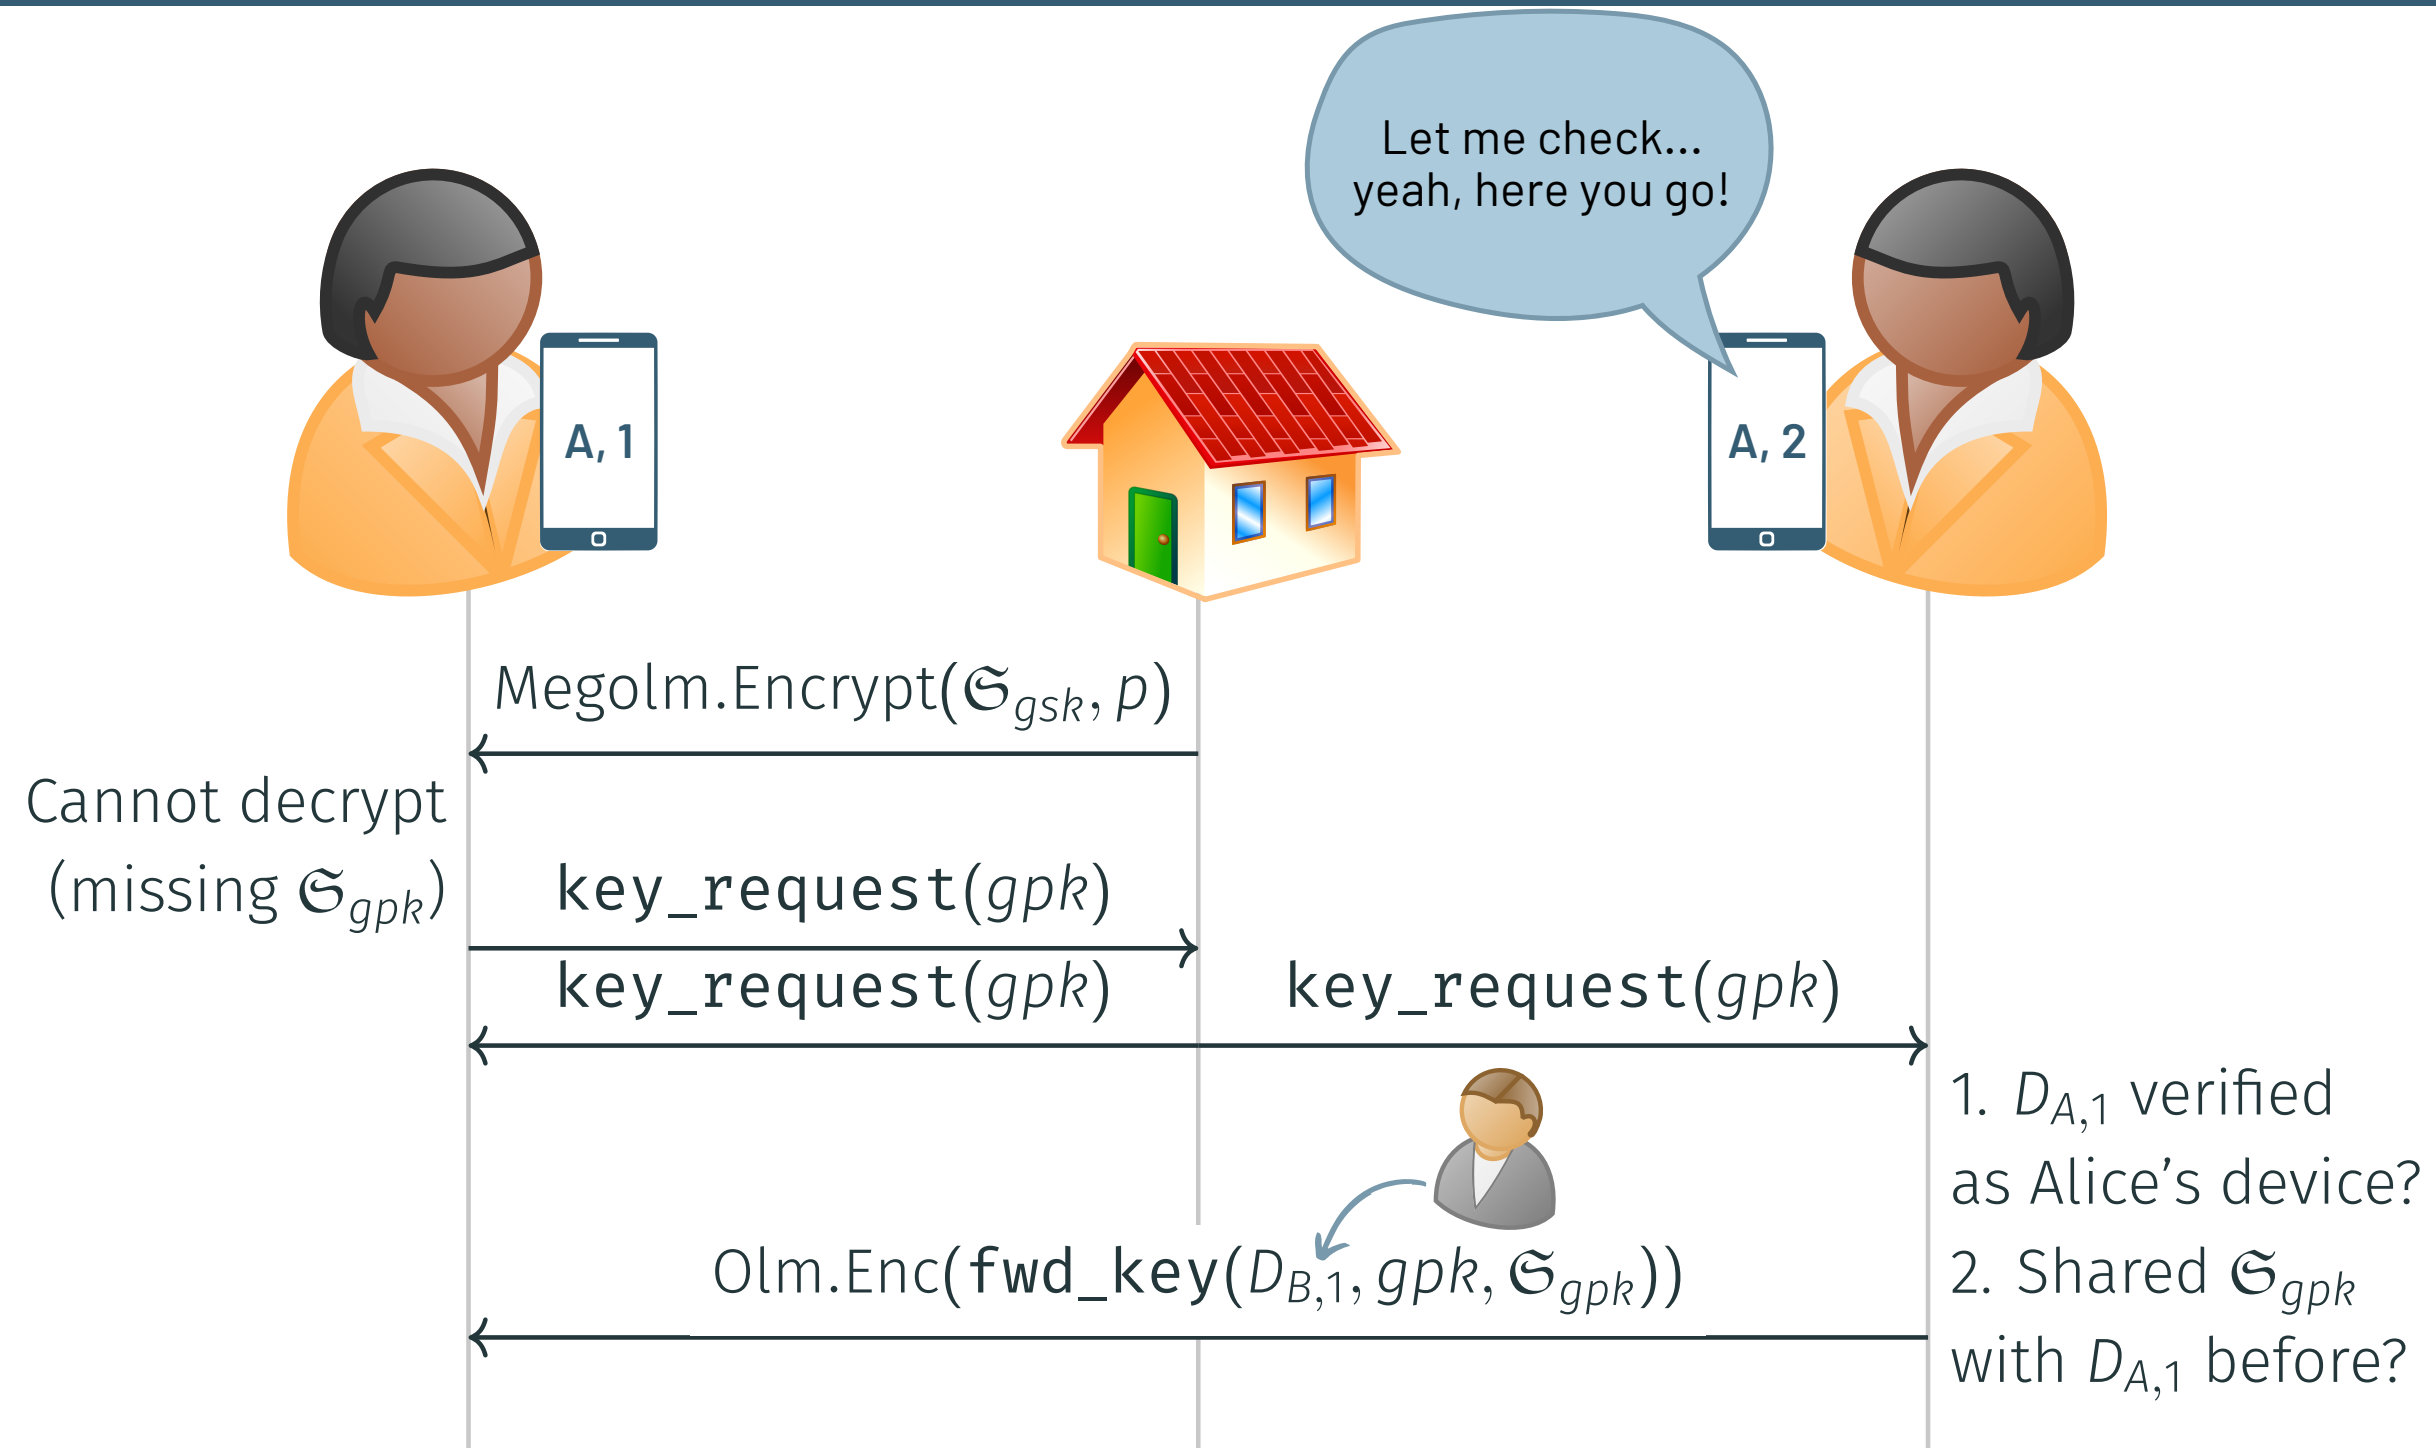
Key Request Protocol

3. Another device (possibly) shares their copy of the keys.



Key Request Protocol

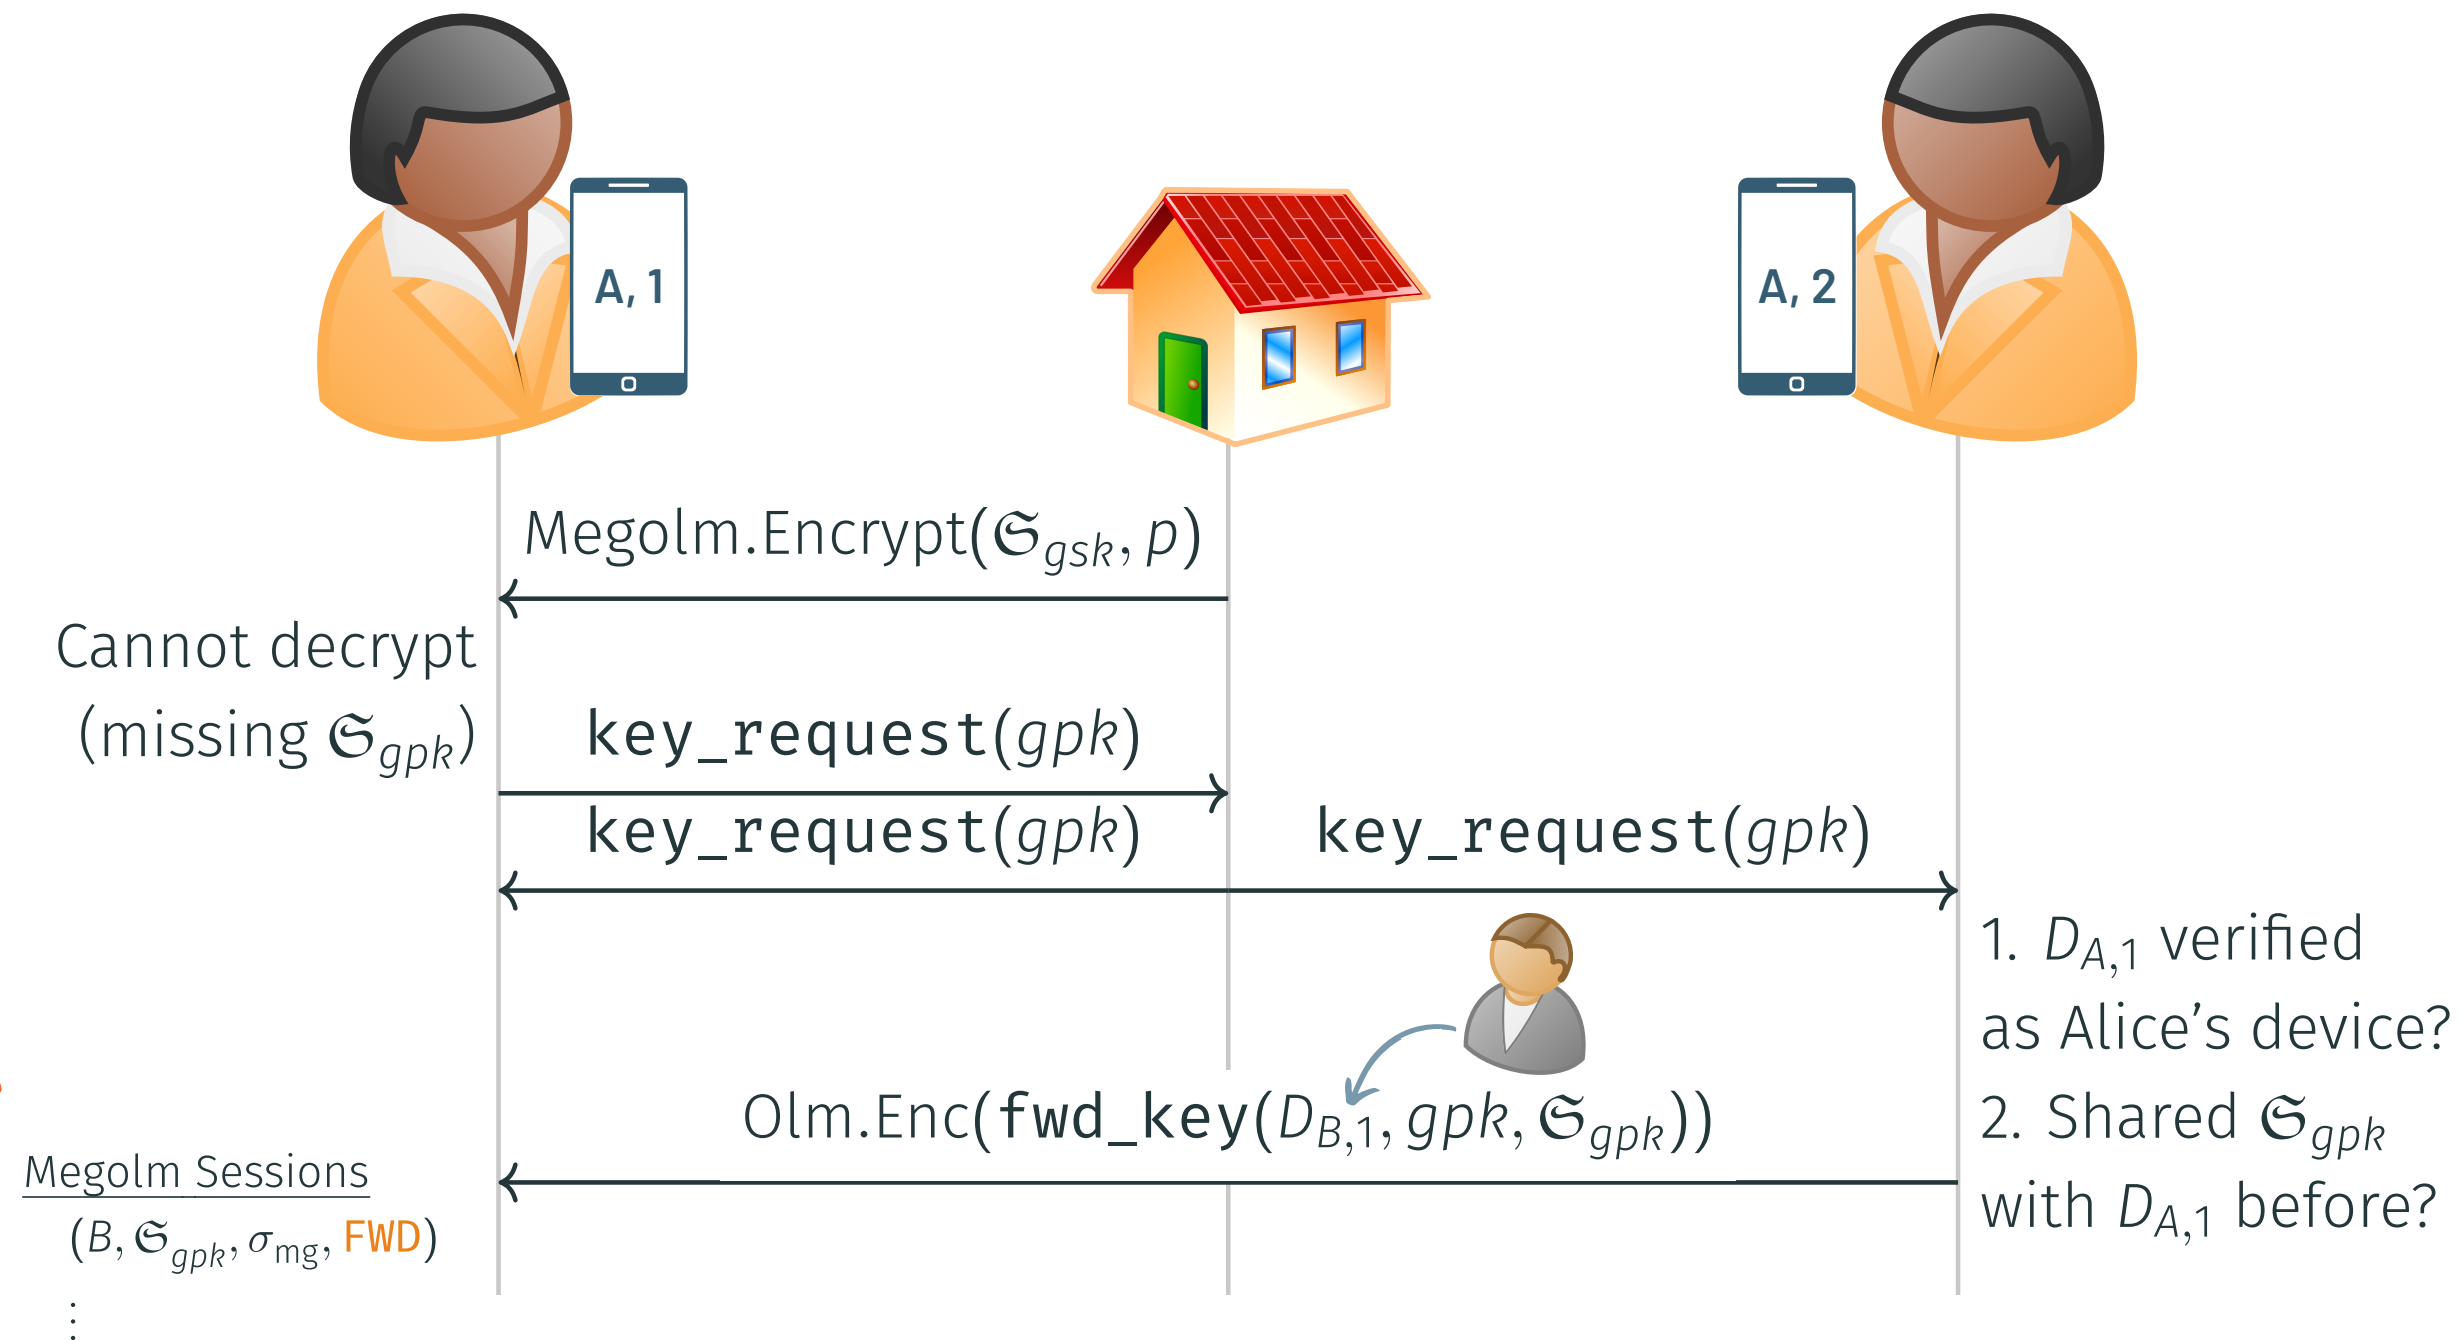
3. Another device (possibly) shares their copy of the keys.



Key Request Protocol

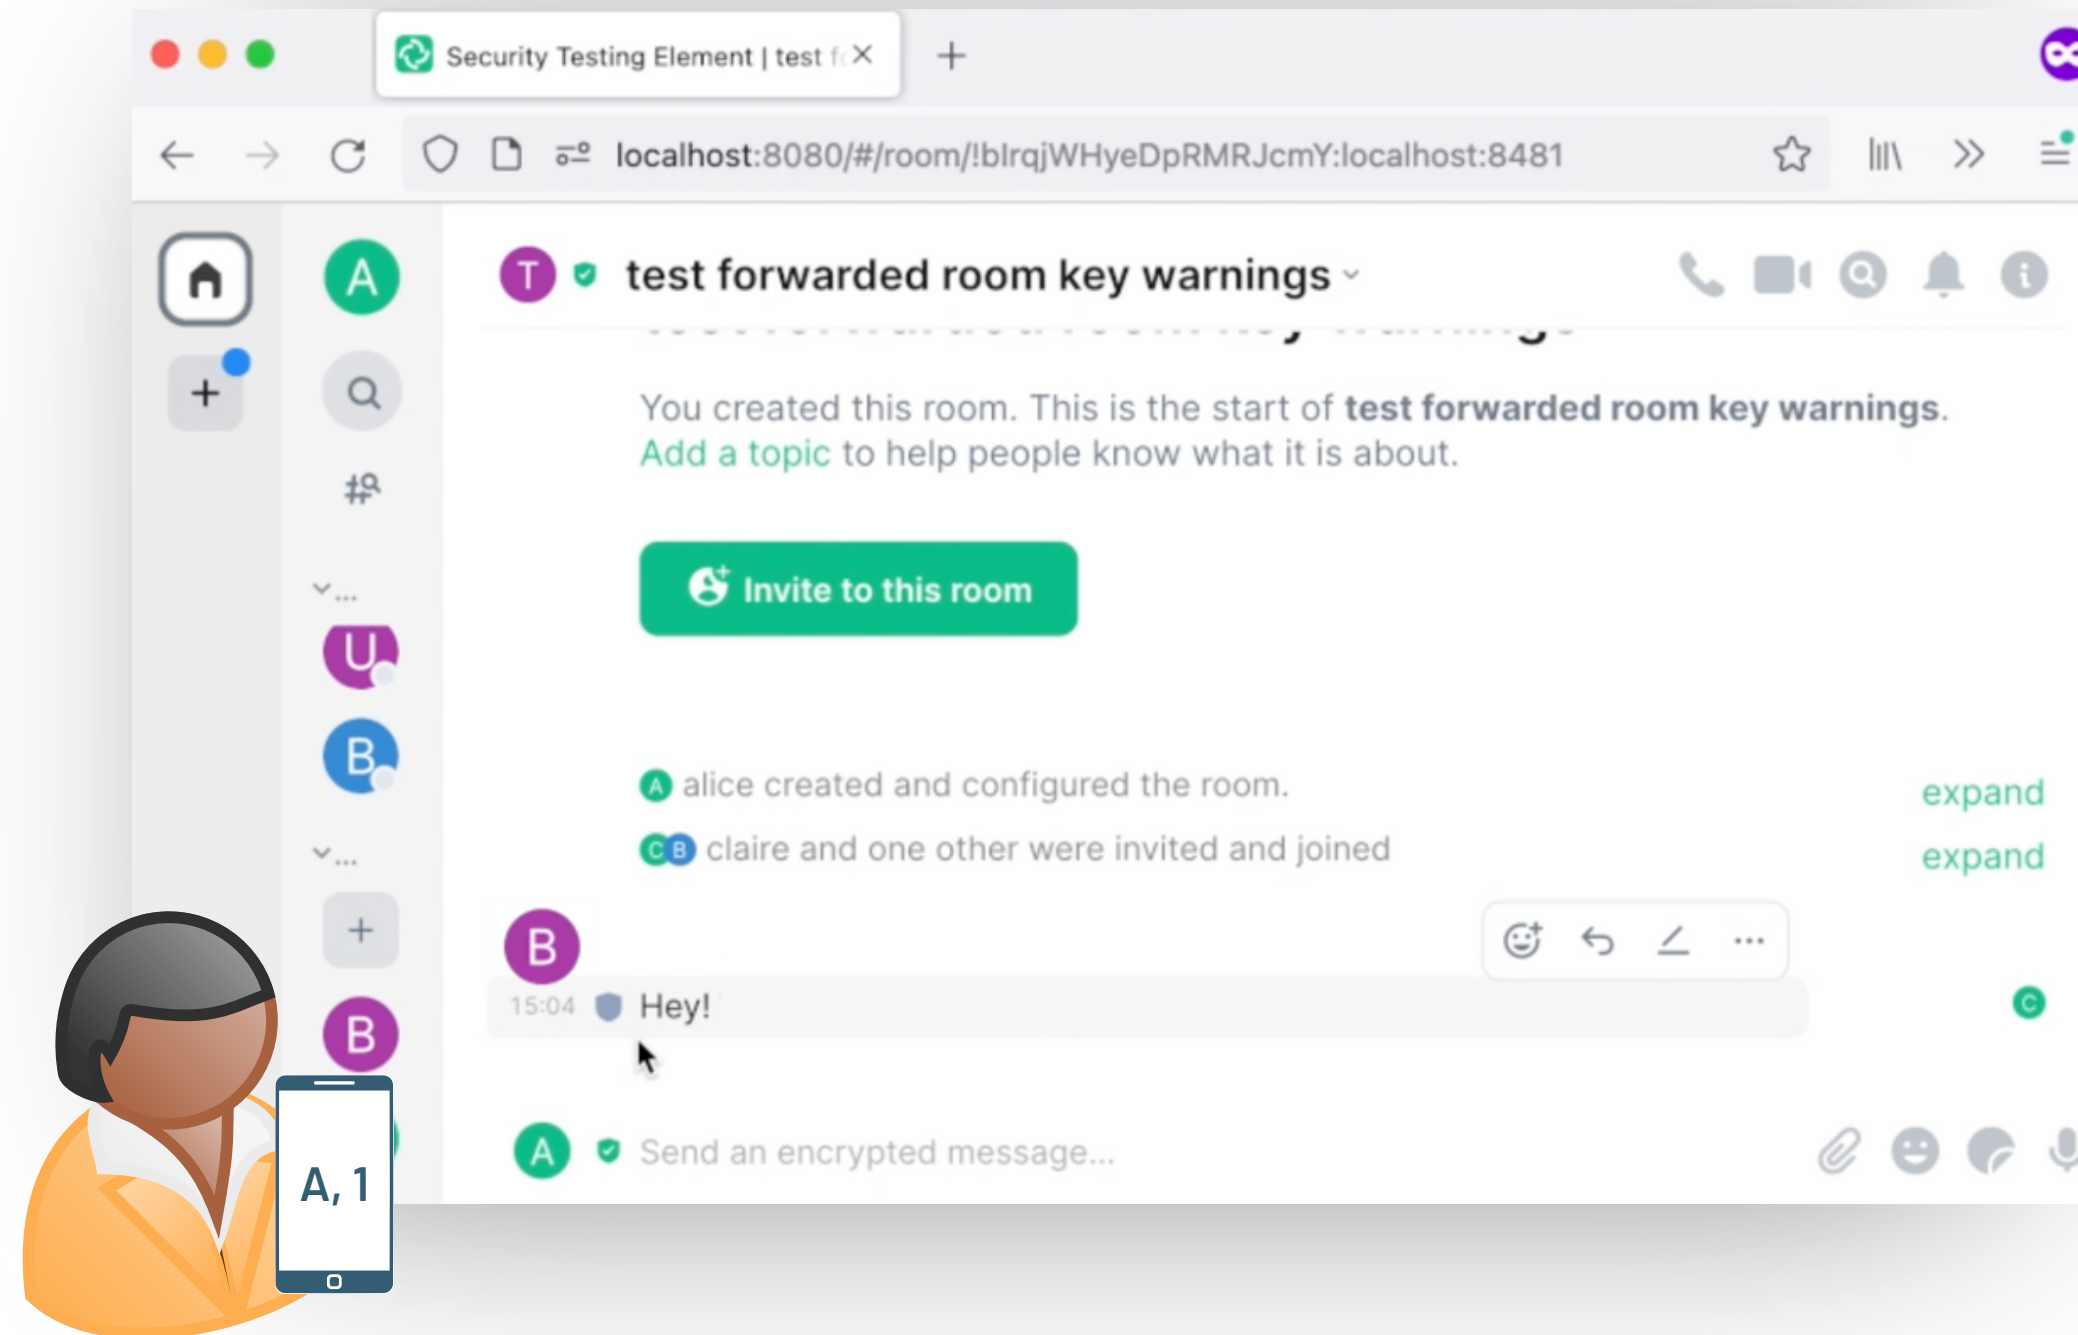
4. (Possibly) accept shared keys then decrypt message.

(A,1) is trusting (A,2) to set the inbound session for other senders (e.g. Bob).



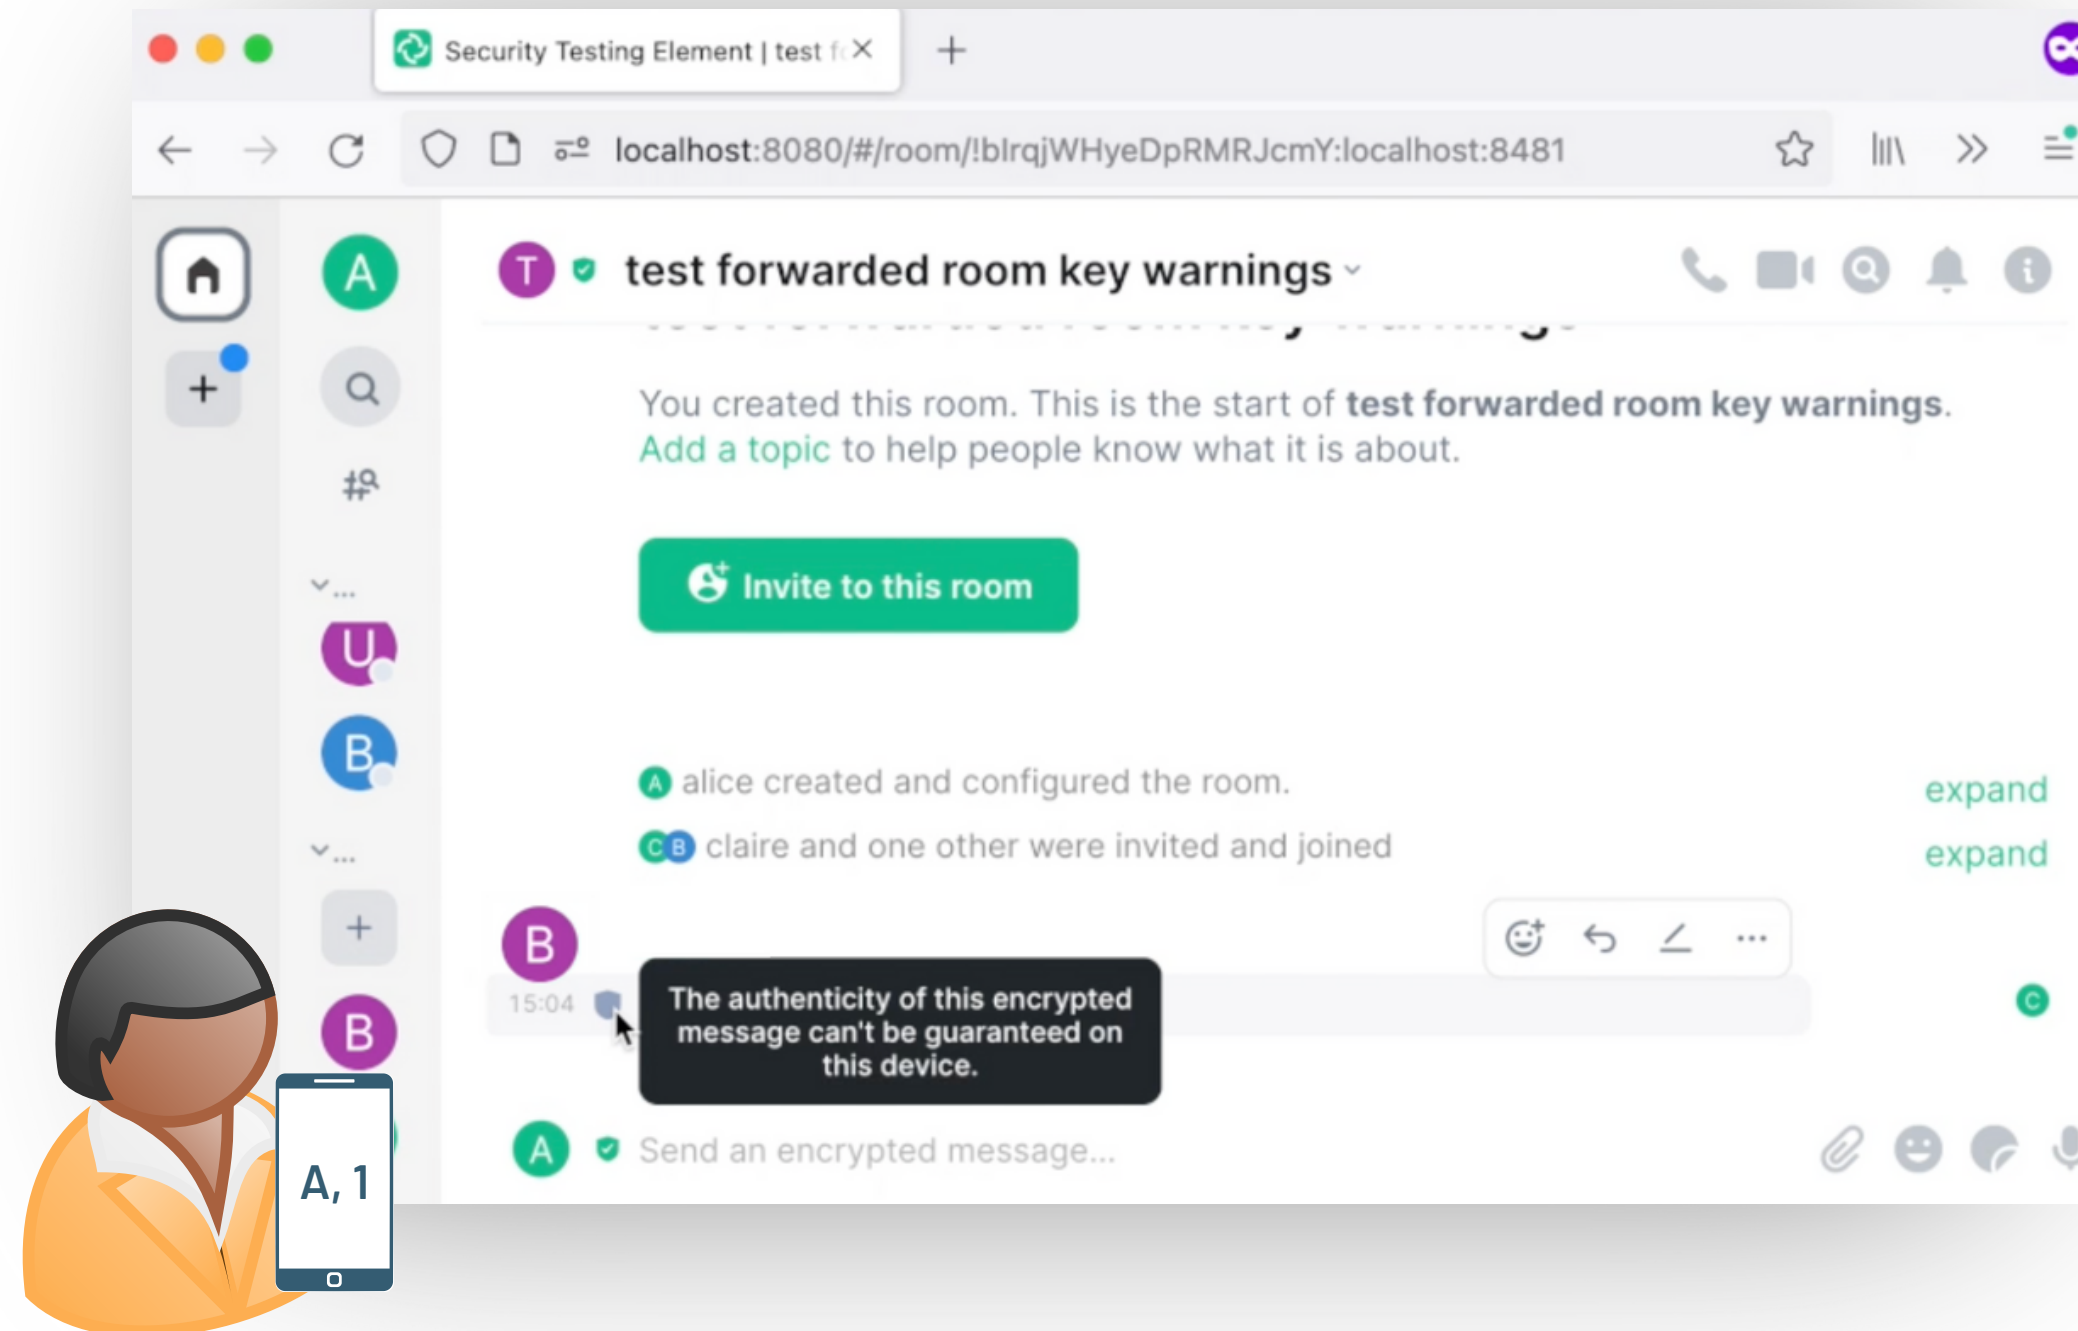
Key Request Protocol

4. (Possibly) accept shared keys then decrypt message.



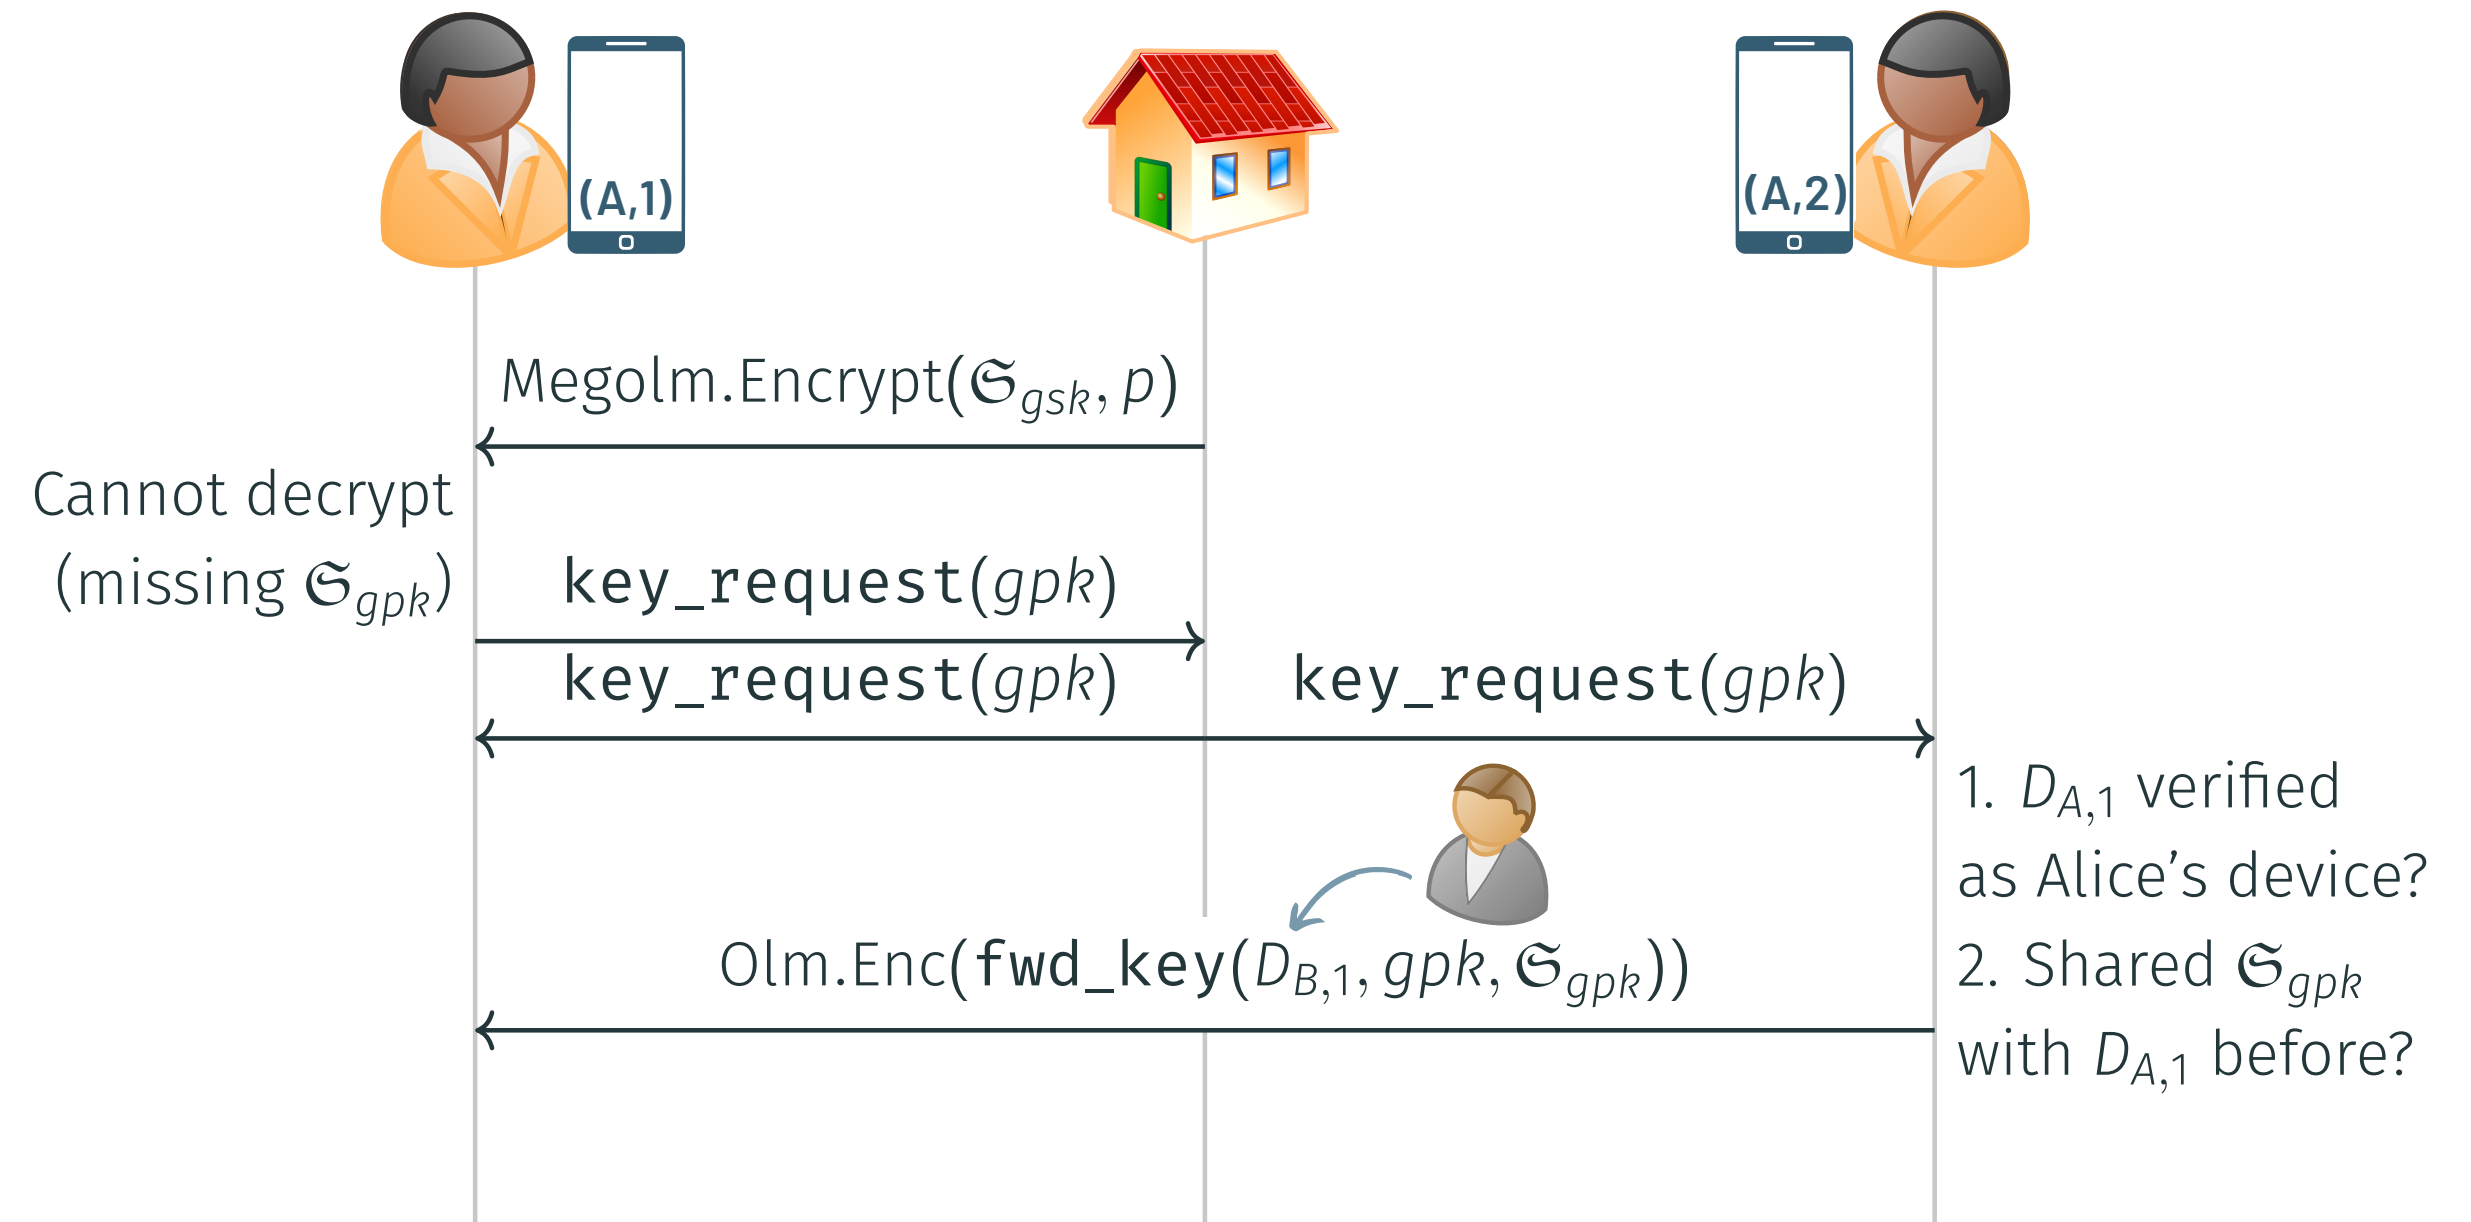
Key Request Protocol

4. (Possibly) accept shared keys then decrypt message.



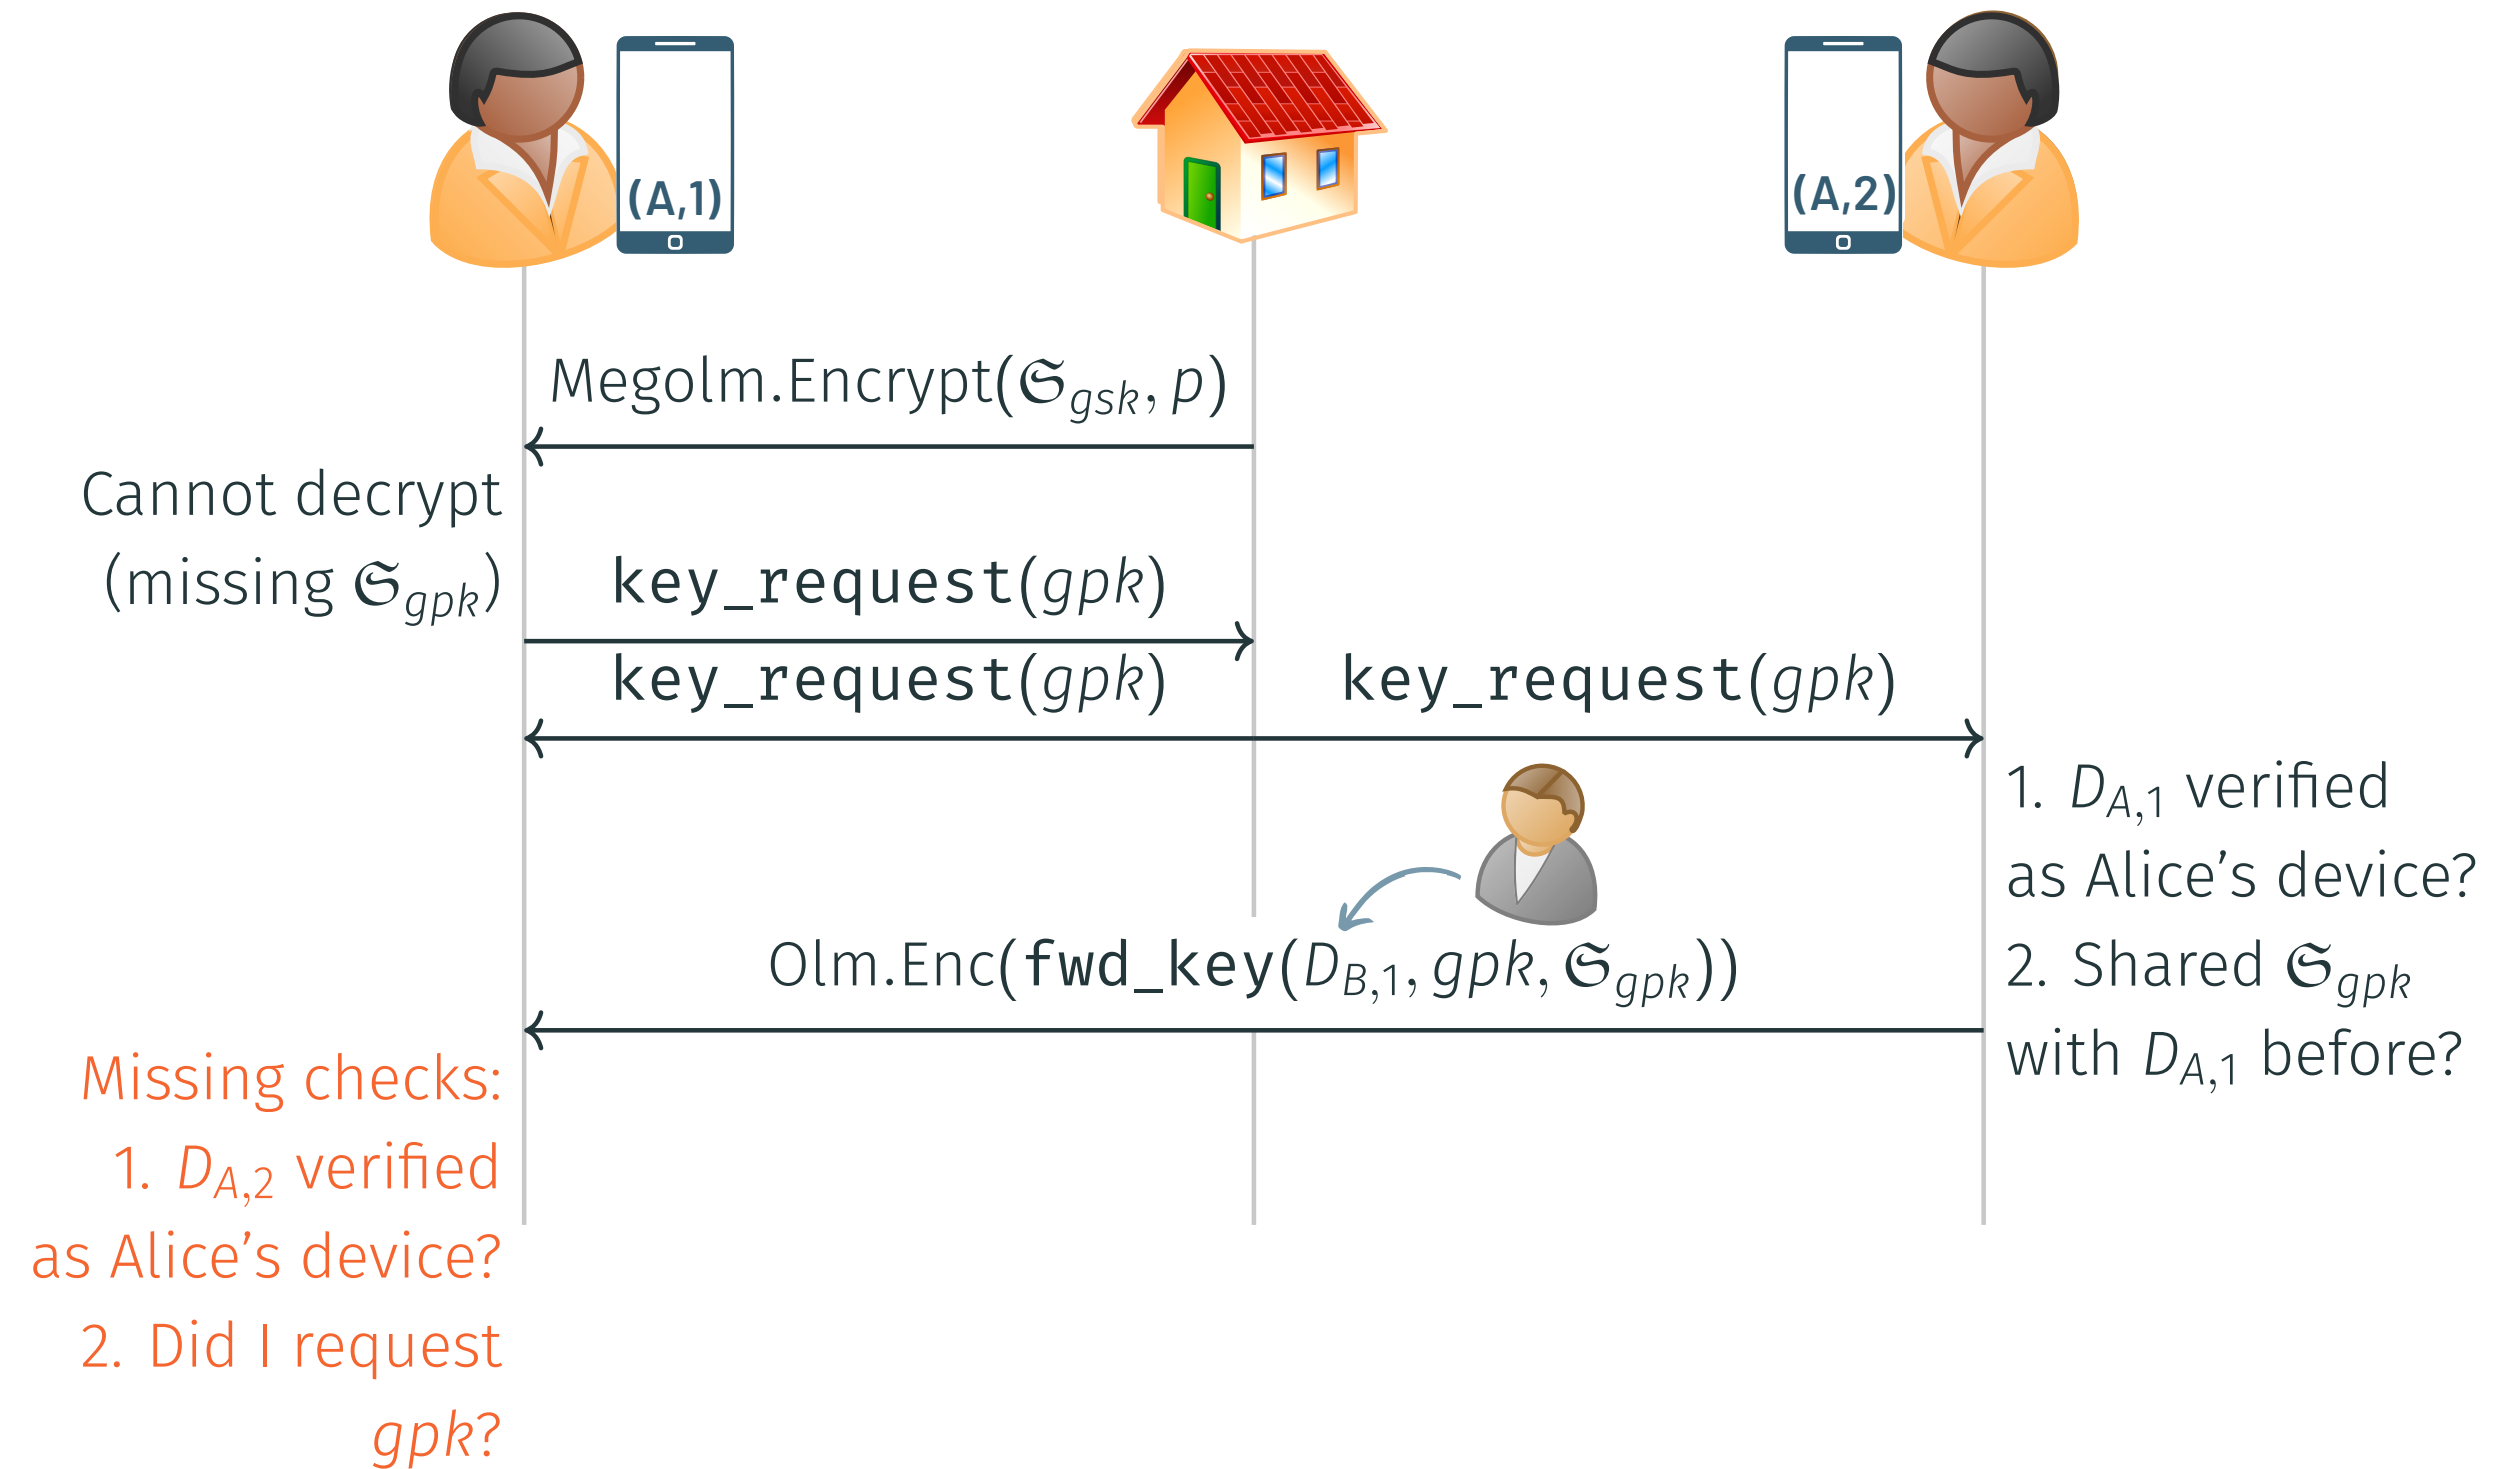
Impersonation through key sharing

ATTACK



Impersonation through key sharing

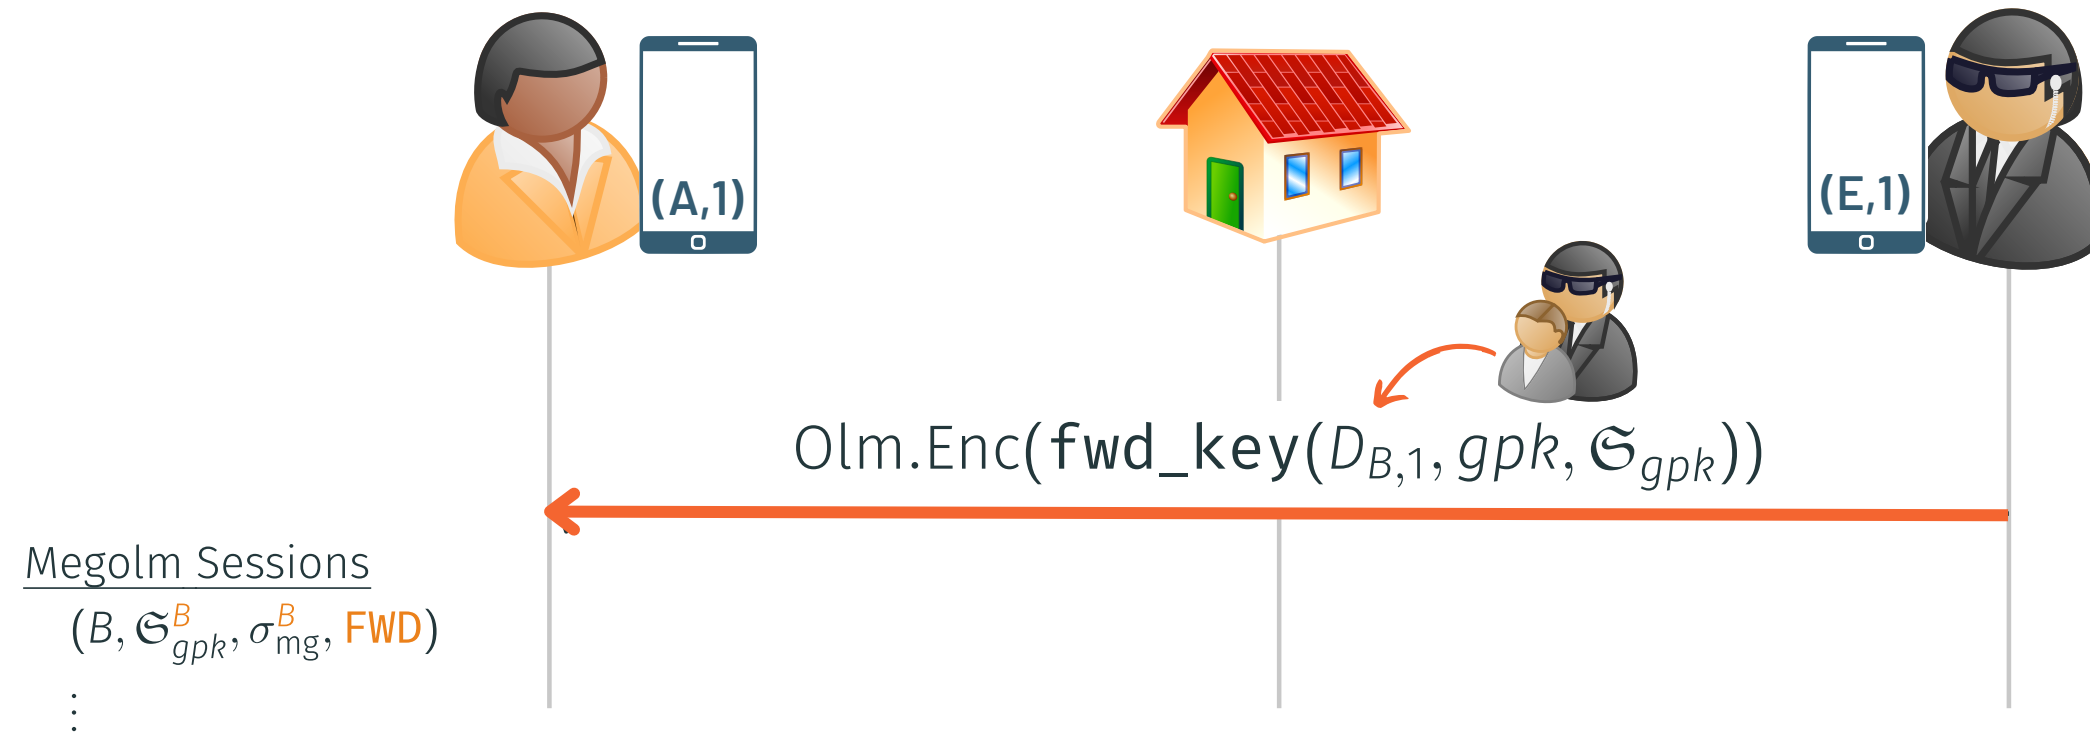
- Missing checks on the receiving side.



Impersonation through key sharing

- Missing checks on the receiving side.
- Allows attackers to inject an inbound Megolm session as if it were someone else's.

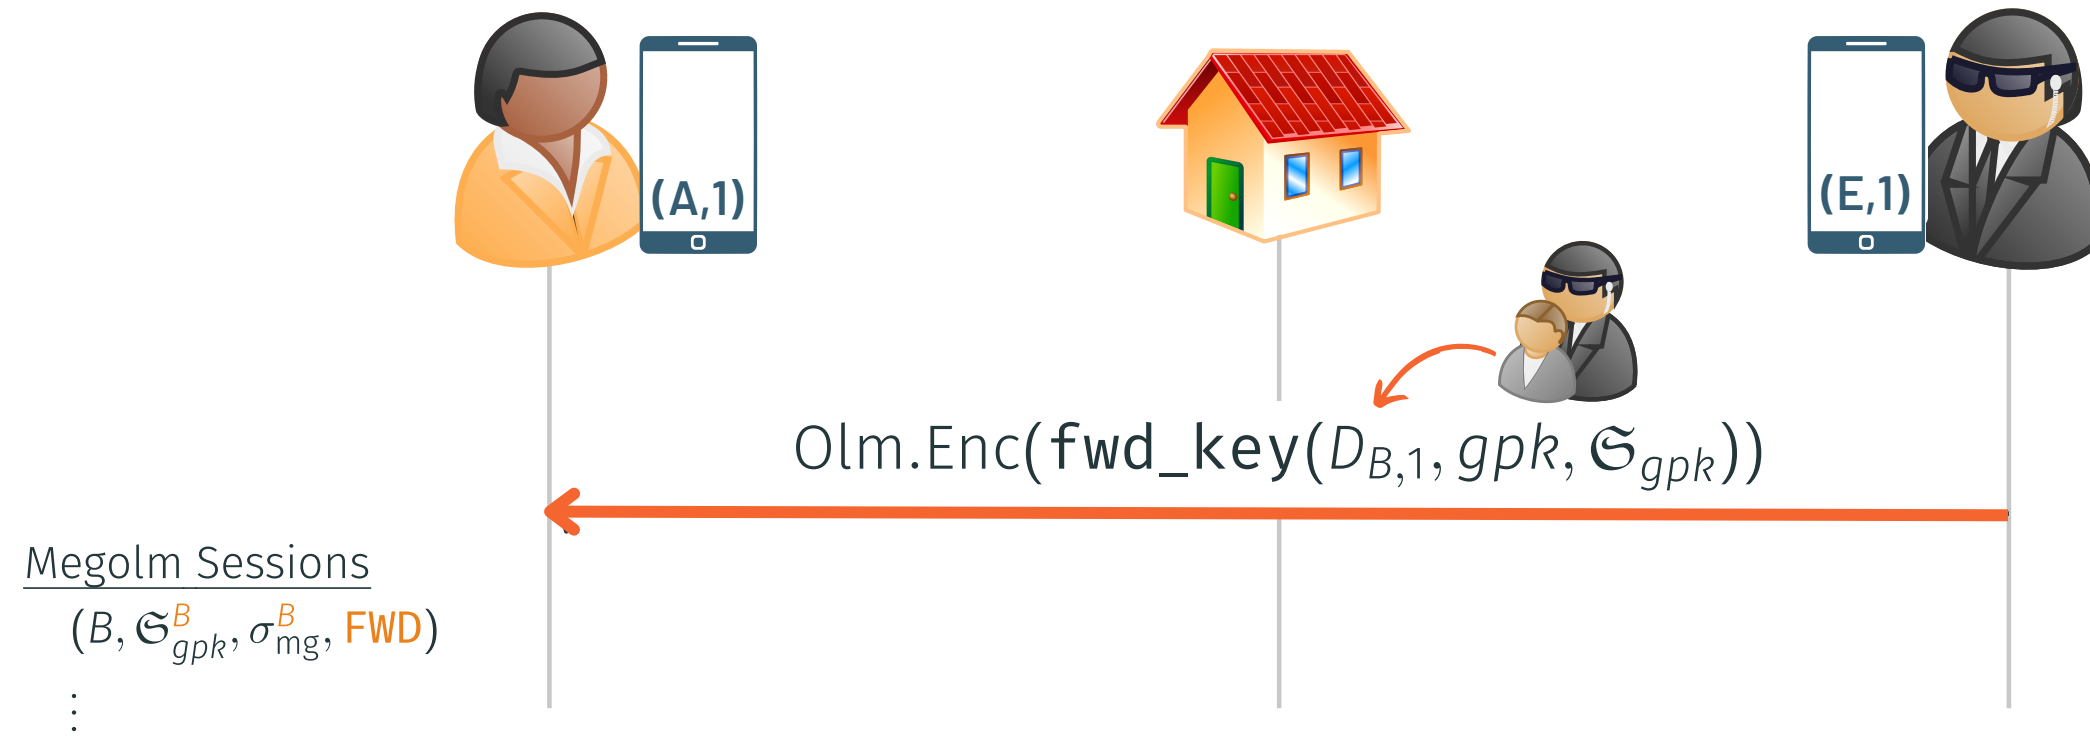
↪ Impersonation *



Impersonation through key sharing

- Missing checks on the receiving side.
- Allows attackers to inject an inbound Megolm session as if it were someone else's.

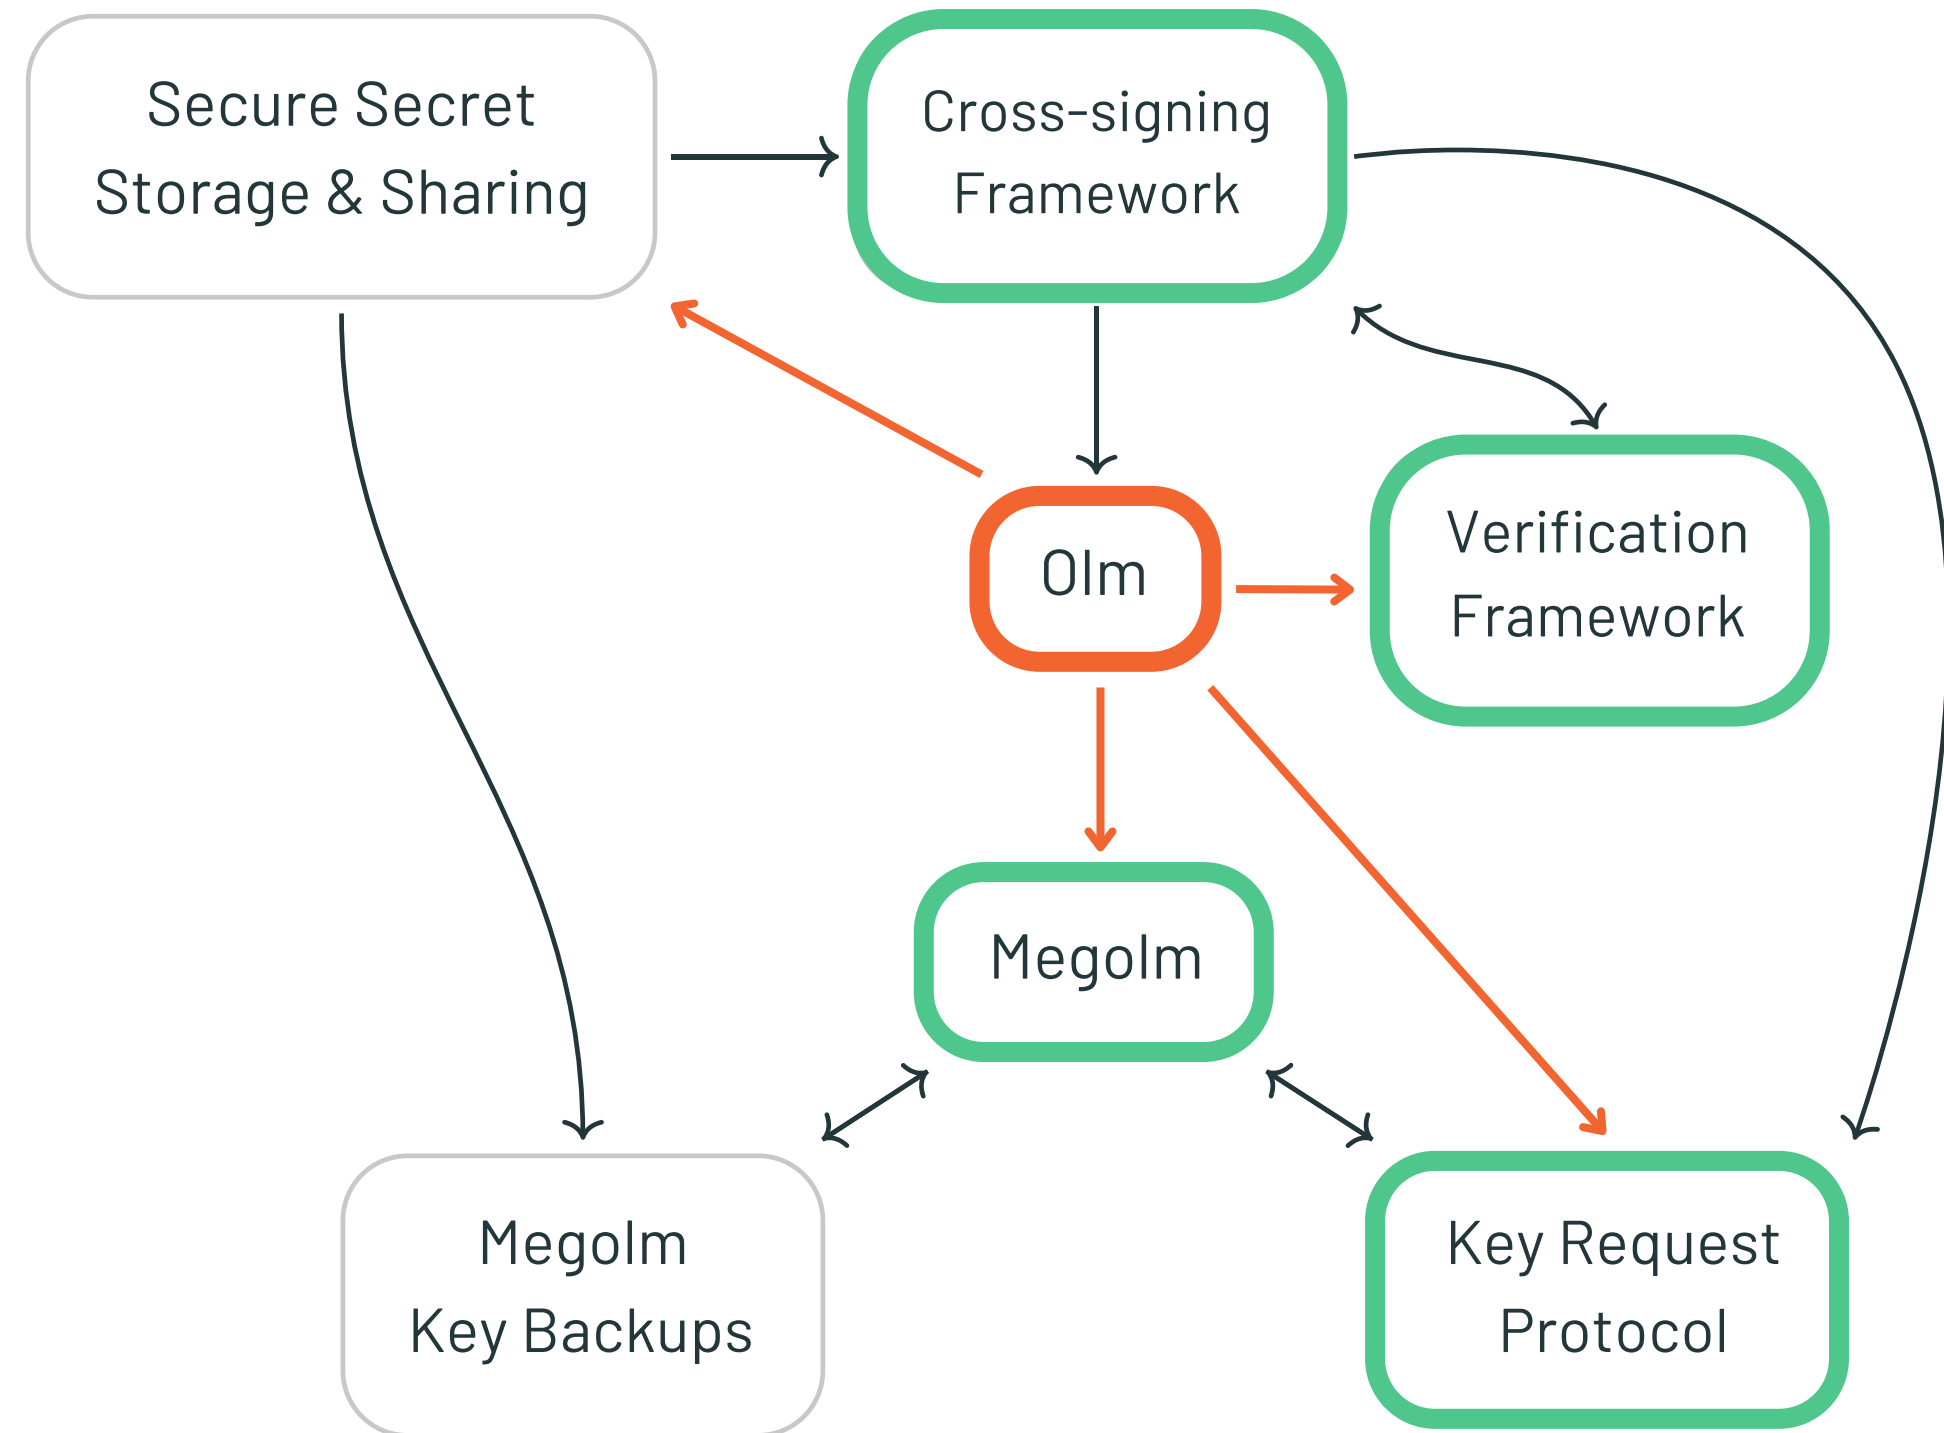
↪ Impersonation *



* The authenticity of this encrypted message can't be guaranteed on this device.

Modelling Olm's Composition

How to differentiate messages from different protocols and purposes?



Modelling Olm's Composition

Olm and Megolm are *generally* interchangeable in the specification.

The screenshot shows a web browser window displaying the Matrix specification page for the 'm.room.encrypted' event type. The browser's address bar shows the URL 'https://spec.matrix.org/unstable/client-server-api/#events-8'. The page header includes the Matrix logo, the text 'specification — unstable version', and navigation links for 'Foundation', 'FAQs', and 'Blog'. A sidebar on the left contains a table of contents with the following items: 'aes-sha2', '11.12.3.3 Key exports', '11.12.3.3.1 Key export format', '11.12.4 Messaging Algorithms', '11.12.4.1 Messaging Algorithm Names', '11.12.4.2 m.olm.v1.curve25519 aes-sha2', '11.12.4.2.1 Recovering from undecryptable messages', '11.12.4.3 m.megolm.v1.aes-sha2', '11.12.5 Protocol definitions', '11.12.5.1 Events', '11.12.5.2 Key management API', and '11.12.5.3 Extensions'. The main content area is titled 'm.room.encrypted' and contains a description: 'This event type is used when sending encrypted events. It can be used either within a room (in which case it will have all of the normal properties in Room events), or as a to-device event.' Below the description is a box labeled 'Event type:' with the value 'Message event'. A section titled 'Content' contains a table with the following data:

Name	Type	Description
algorithm	enum	Required: The encryption algorithm used to encrypt this event. The value of this field determines which other properties will be present. One of: [m.olm.v1.curve25519-aes-sha2, m.megolm.v1.aes-sha2].
ciphertext	string {string: CiphertextInfo}	Required: The encrypted content of the event. Either the encrypted

Modelling Olm's Composition

Olm and Megolm are *generally* interchangeable in the specification.

With some special casing.

The screenshot shows a web browser window displaying the Matrix specification page for the `m.forwarded_room_key` event type. The page title is "Client-Server API | Matrix Specification" and the URL is `https://spec.matrix.org/unstable/client-server-api/#mforwarded_room_key`. The page header includes the Matrix logo, "specification — unstable version", and links for "Foundation", "FAQs", and "Blog".

The left sidebar shows a navigation menu with the following items:

- 11.12 Protocol definitions
 - 11.12.5.1 Events**
 - 11.12.5.2 Key management API
 - 11.12.5.3 Extensions to /sync
- 11.12.6 Reporting that decryption keys are withheld
- 11.13 Secrets
 - 11.13.1 Storage
 - 11.13.1.1 Key storage
 - 11.13.1.2 Secret storage
 - 11.13.1.2.1 `m.secret_storage.hmac-sha2`
 - 11.13.1.2.2 Key representation

▼ `m.forwarded_room_key`

This event type is used to forward keys for end-to-end encryption. It is encrypted as an `m.room.encrypted` event using [Olm](#), then sent as a [to-device](#) event.

Event type: Message event

Content

Name	Type	Description
<code>algorithm</code>	string	Required: The encryption algorithm the key in this event is to be used with.
<code>forwarding_curve25519_key_chain</code>	[string]	Required: Chain of Curve25519 keys. It starts out empty, but each time the key is forwarded to another

Modelling Olm's Composition

How to differentiate messages from different protocols and purposes?

```
Matrix.Decrypt( $st_{mt}$ ,  $type$ ,  $alg$ ,  $c$ )
```

```
// Decrypt ciphertext
```

```
if ( $type = m.room.encrypted$ )  $\wedge$  ( $alg = mego\l m$ ) then
```

```
    // Select correct Megolm session for decryption
```

```
     $\mathfrak{S}'_{gpk}, ipk_{snd}, type, m \leftarrow Megolm.Decrypt(\mathfrak{S}_{gpk}, c)$ 
```

```
elseif ( $type = m.room.encrypted$ )  $\wedge$  ( $alg = olm$ ) then
```

```
    // Select correct Olm session for decryption
```

```
     $st_{olm}, ipk_{snd}, type, m \leftarrow Olm.Dec(st_{olm}, c)$ 
```

```
// Handle plaintext
```

```
if ( $type = m.room.message$ ) then
```

```
    // Handle room message
```

```
elseif ( $type = m.room\_key$ ) then
```

```
    // Handle initial Megolm session distribution
```

```
elseif (...) then
```

```
    // Handle other message types
```

Modelling Olm's Composition

How to differentiate messages from different protocols and purposes?

```
Matrix.Decrypt( $st_{mt}$ ,  $type$ ,  $alg$ ,  $c$ )

---

// Decrypt ciphertext  
if ( $type = m.room.encrypted$ )  $\wedge$  ( $alg = mego\text{lm}$ ) then  
    // Select correct Megolm session for decryption  
     $\mathfrak{S}'_{gpk}, ipk_{snd}, type, m \leftarrow \text{Megolm.Decrypt}(\mathfrak{S}_{gpk}, c)$   
elseif ( $type = m.room.encrypted$ )  $\wedge$  ( $alg = olm$ ) then  
    // Select correct Olm session for decryption  
     $st_{olm}, ipk_{snd}, type, m \leftarrow \text{Olm.Dec}(st_{olm}, c)$   
// Handle plaintext  
if ( $type = m.room.message$ ) then  
    // Handle room message  
elseif ( $type = m.room\_key$ ) then  
    // Handle initial Megolm session distribution  
elseif (...) then  
    // Handle other message types
```

Modelling Olm's Composition

How to differentiate messages from different protocols and purposes?

```
Matrix.Decrypt( $st_{mt}$ ,  $type$ ,  $alg$ ,  $c$ )
```

```
// Decrypt ciphertext
```

```
if ( $type = m.room.encrypted$ )  $\wedge$  ( $alg = mego\!lm$ ) then
```

```
    // Select correct Megolm session for decryption
```

```
     $\mathfrak{S}'_{gpk}, ipk_{snd}, type, m \leftarrow Megolm.Decrypt(\mathfrak{S}_{gpk}, c)$ 
```

```
elseif ( $type = m.room.encrypted$ )  $\wedge$  ( $alg = olm$ ) then
```

```
    // Select correct Olm session for decryption
```

```
     $st_{olm}, ipk_{snd}, type, m \leftarrow Olm.Dec(st_{olm}, c)$ 
```

```
// Handle plaintext
```

```
if ( $type = m.room.message$ ) then
```

```
    // Handle room message
```

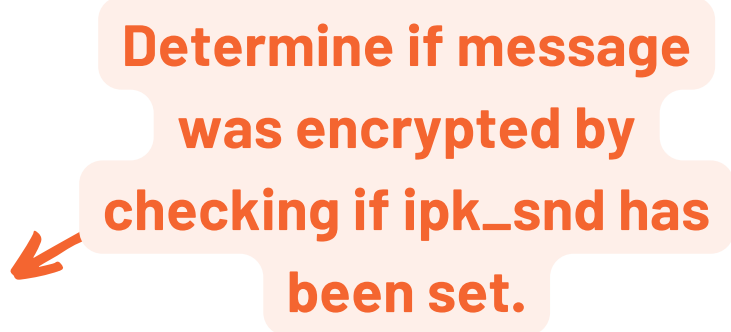
```
elseif ( $type = m.room_key$ ) then
```

```
    // Handle initial Megolm session distribution
```

```
elseif (...) then
```

```
    // Handle other message types
```

Determine if message was encrypted by checking if `ipk_snd` has been set.



Modelling Olm's Composition

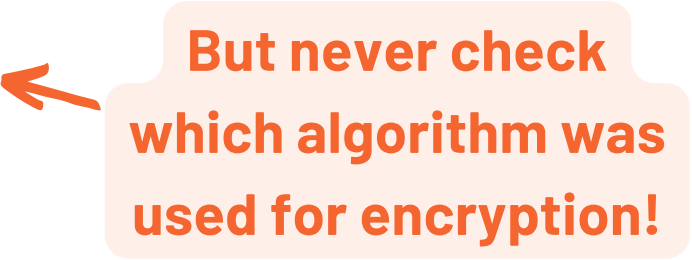
How to differentiate messages from different protocols and purposes?

```
Matrix.Decrypt( $st_{mt}$ , type, alg, c)

---

// Decrypt ciphertext  
if (type = m.room.encrypted)  $\wedge$  (alg = mego1m) then  
    // Select correct Megolm session for decryption  
     $\mathcal{G}'_{gpk}, ipk_{snd}, type, m \leftarrow \text{Megolm.Decrypt}(\mathcal{G}_{gpk}, c)$   
elseif (type = m.room.encrypted)  $\wedge$  (alg = olm) then  
    // Select correct Olm session for decryption  
     $st_{olm}, ipk_{snd}, type, m \leftarrow \text{Olm.Dec}(st_{olm}, c)$   
// Handle plaintext  
if (type = m.room.message) then  
    // Handle room message  
elseif (type = m.room_key) then  
    // Handle initial Megolm session distribution  
elseif (...) then  
    // Handle other message types
```

**But never check
which algorithm was
used for encryption!**



Olm/Megolm Protocol Confusion

Keys sent as part of Megolm's initial key share are fully trusted.
But they are sent over Olm.

Can we share a session key over Megolm?

ATTACK

Megolm Sessions
($B, \mathcal{G}_{gpk}^B, \sigma_{mg}^B, \text{FWD}$)
⋮



Megolm Sessions
($\mathcal{G}_{gsk}^B, \mathcal{G}_{gpk}^B, \sigma_{mg}^B$)
⋮



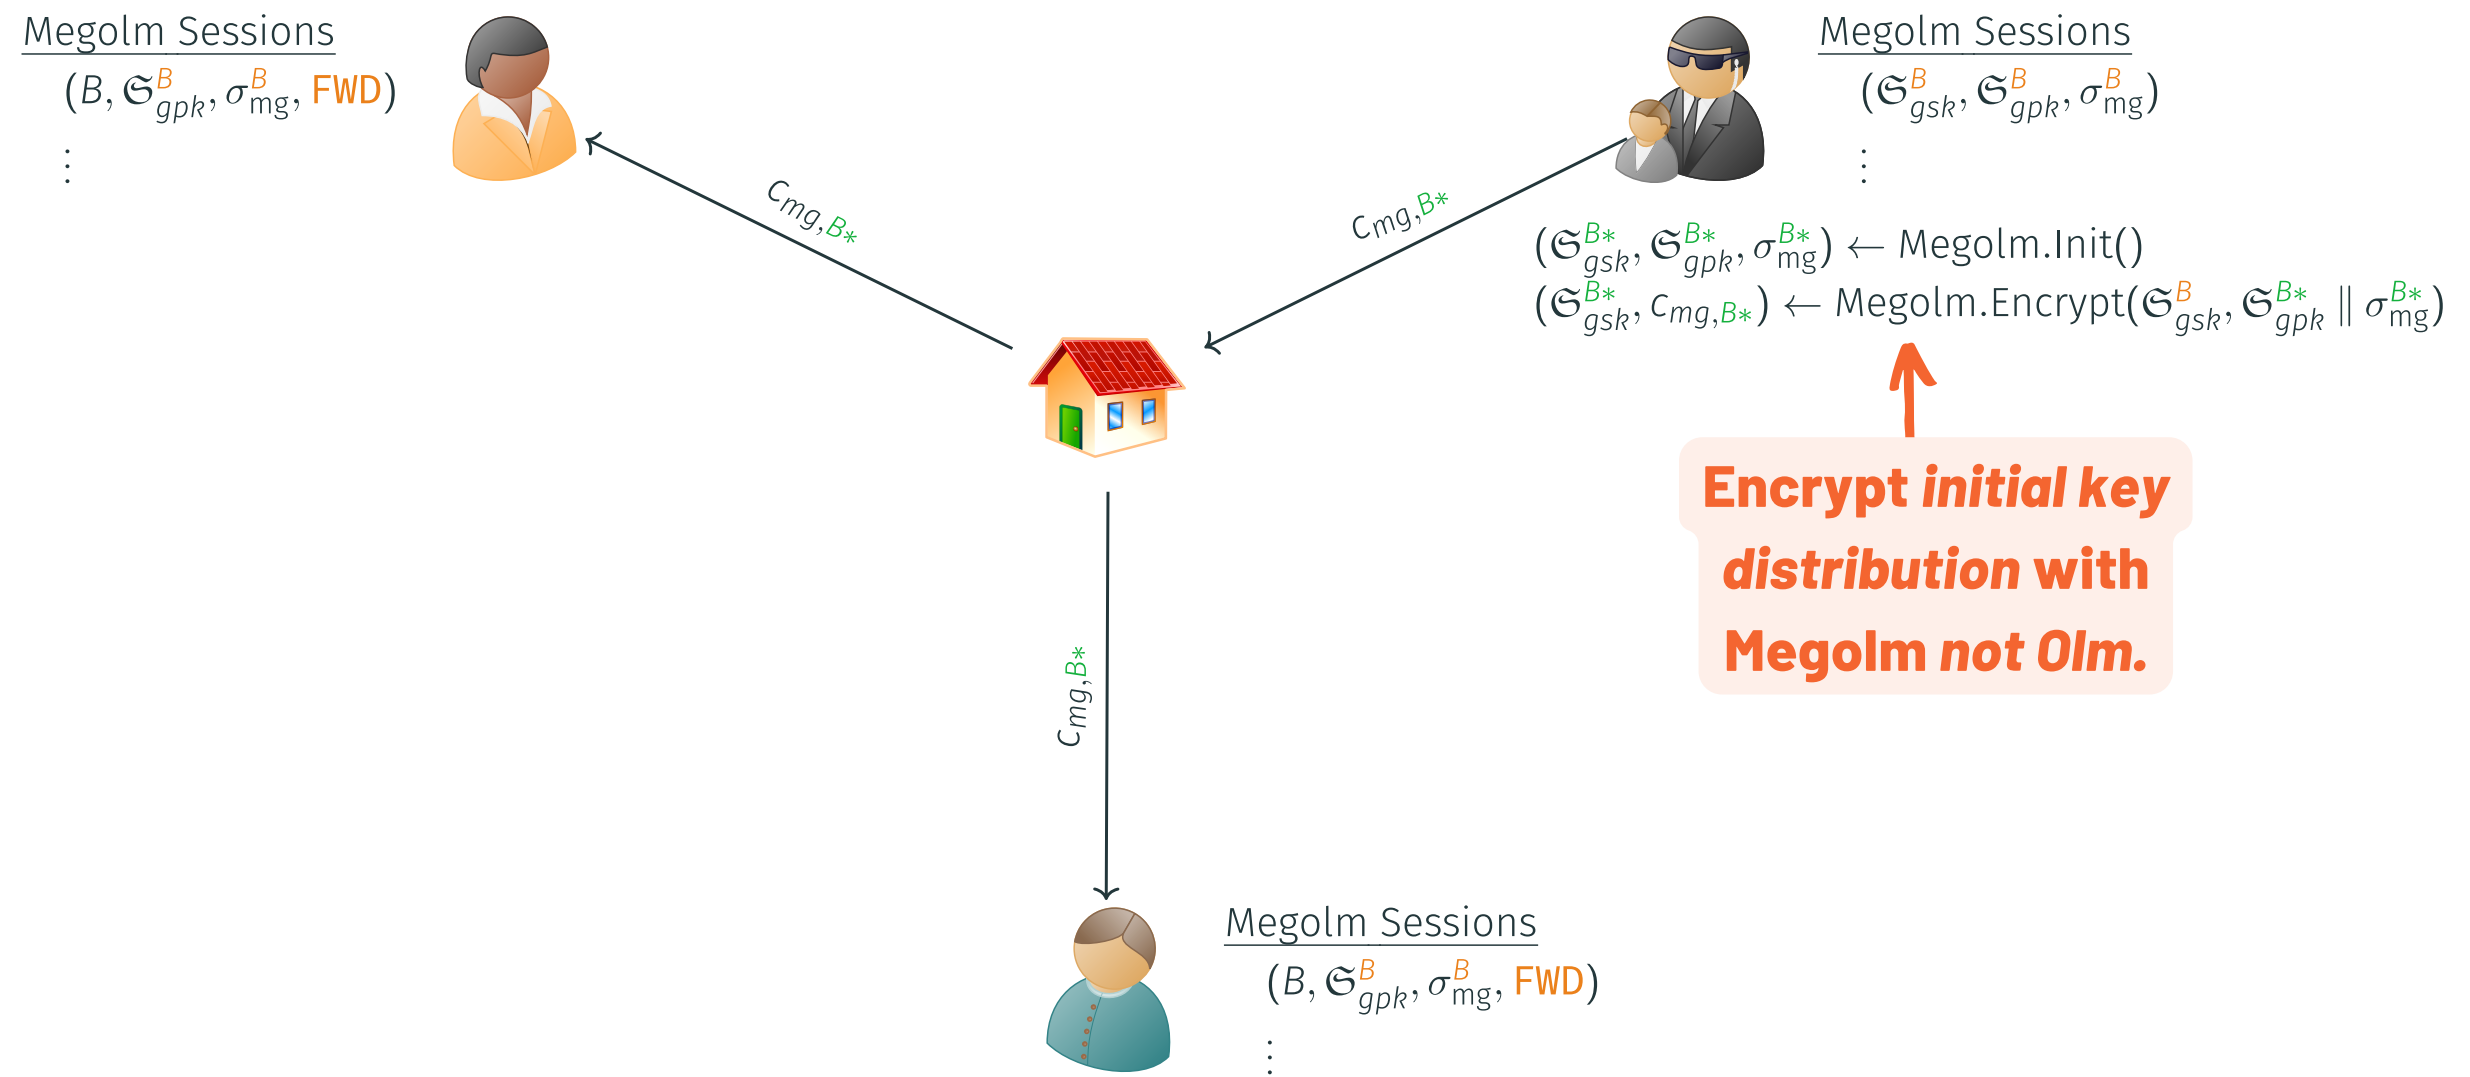
Megolm Sessions
($B, \mathcal{G}_{gpk}^B, \sigma_{mg}^B, \text{FWD}$)
⋮

Olm/Megolm Protocol Confusion

Keys sent as part of Megolm's *initial key share* are *fully trusted*.
But they are sent over Olm.

Can we share a session key over Megolm? **Yes**

ATTACK

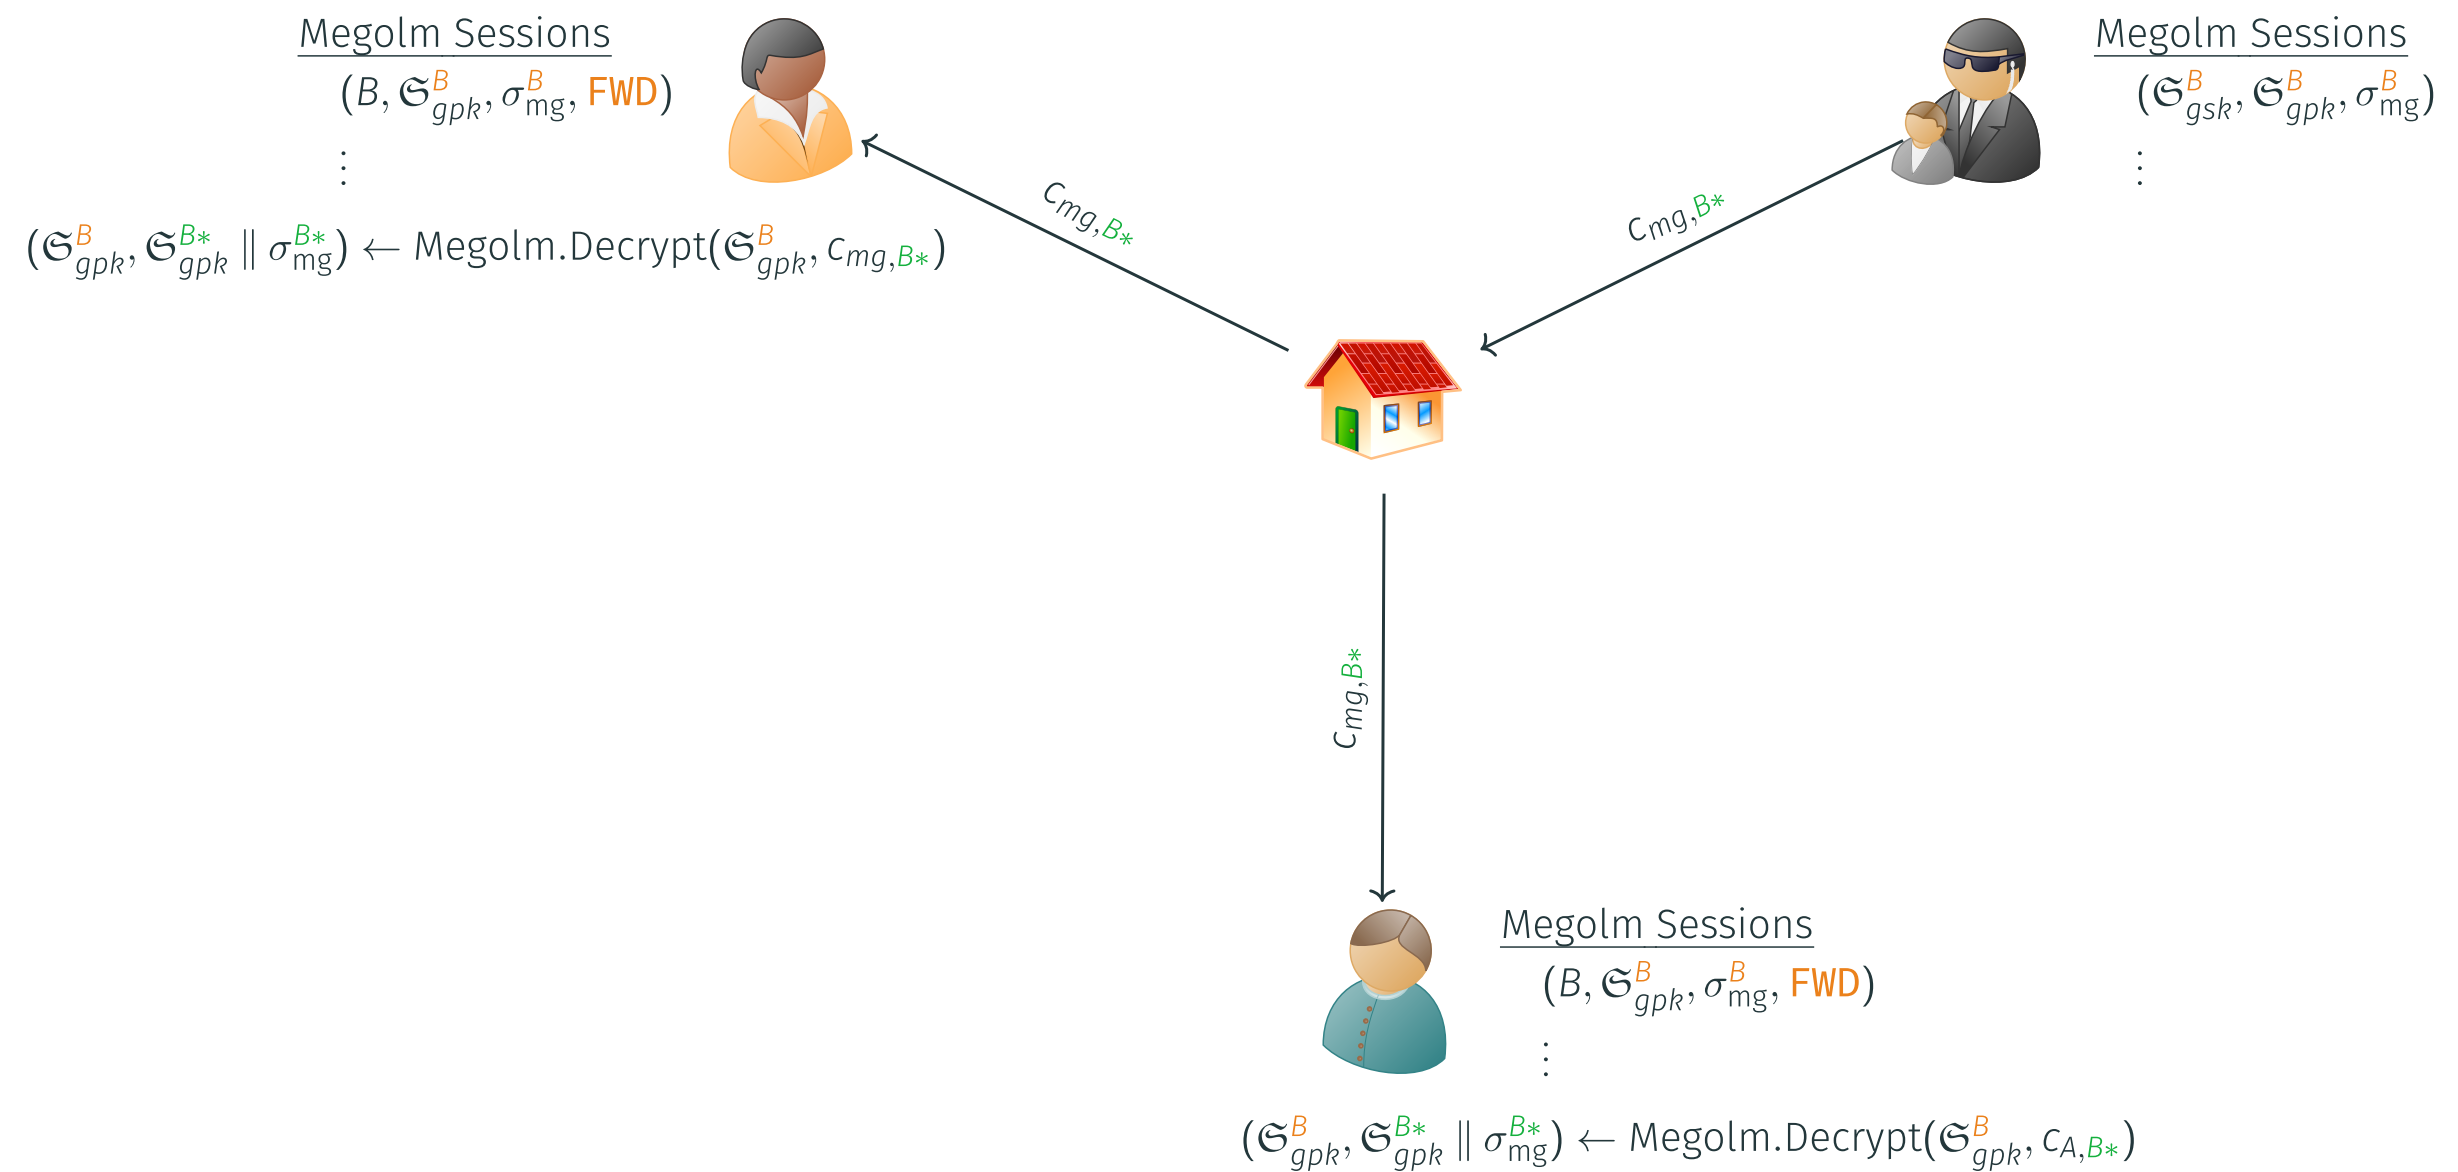


Olm/Megolm Protocol Confusion

Keys sent as part of Megolm's *initial key share* are *fully trusted*.
But they are sent over Olm.

Can we share a session key over Megolm? **Yes**

ATTACK



Olm/Megolm Protocol Confusion

Keys sent as part of Megolm's *initial key share* are *fully trusted*.
But they are sent over Olm.

Can we share a session key over Megolm? **Yes**

Megolm Sessions
(B, $\mathfrak{G}_{gpk}^B, \sigma_{mg}^B$, FWD)
(B, $\mathfrak{G}_{gpk}^{B^*}, \sigma_{mg}^{B^*}$)
⋮



Megolm Sessions
($\mathfrak{G}_{gsk}^B, \mathfrak{G}_{gpk}^B, \sigma_{mg}^B$)
($\mathfrak{G}_{gsk}^{B^*}, \mathfrak{G}_{gpk}^{B^*}, \sigma_{mg}^{B^*}$)
⋮

Messages from B* are not marked as FWD.

Megolm Sessions
(B, $\mathfrak{G}_{gpk}^B, \sigma_{mg}^B$, FWD)
(B, $\mathfrak{G}_{gpk}^{B^*}, \sigma_{mg}^{B^*}$)
⋮



Olm/Megolm Protocol Confusion

Keys sent as part of Megolm's *initial key share* are *fully trusted*. But they are sent over Olm.

Can we share a session key over Megolm? **Yes**

↪ Impersonation *(no caveat)*

~~The authenticity of this encrypted message can't be guaranteed on this device.~~

ATTACK

Megolm Sessions
(B, $\mathcal{G}_{gpk}^B, \sigma_{mg}^B$, FWD)
(B, $\mathcal{G}_{gpk}^{B^*}, \sigma_{mg}^{B^*}$)
⋮



Megolm Sessions
($\mathcal{G}_{gsk}^B, \mathcal{G}_{gpk}^B, \sigma_{mg}^B$)
($\mathcal{G}_{gsk}^{B^*}, \mathcal{G}_{gpk}^{B^*}, \sigma_{mg}^{B^*}$)
⋮

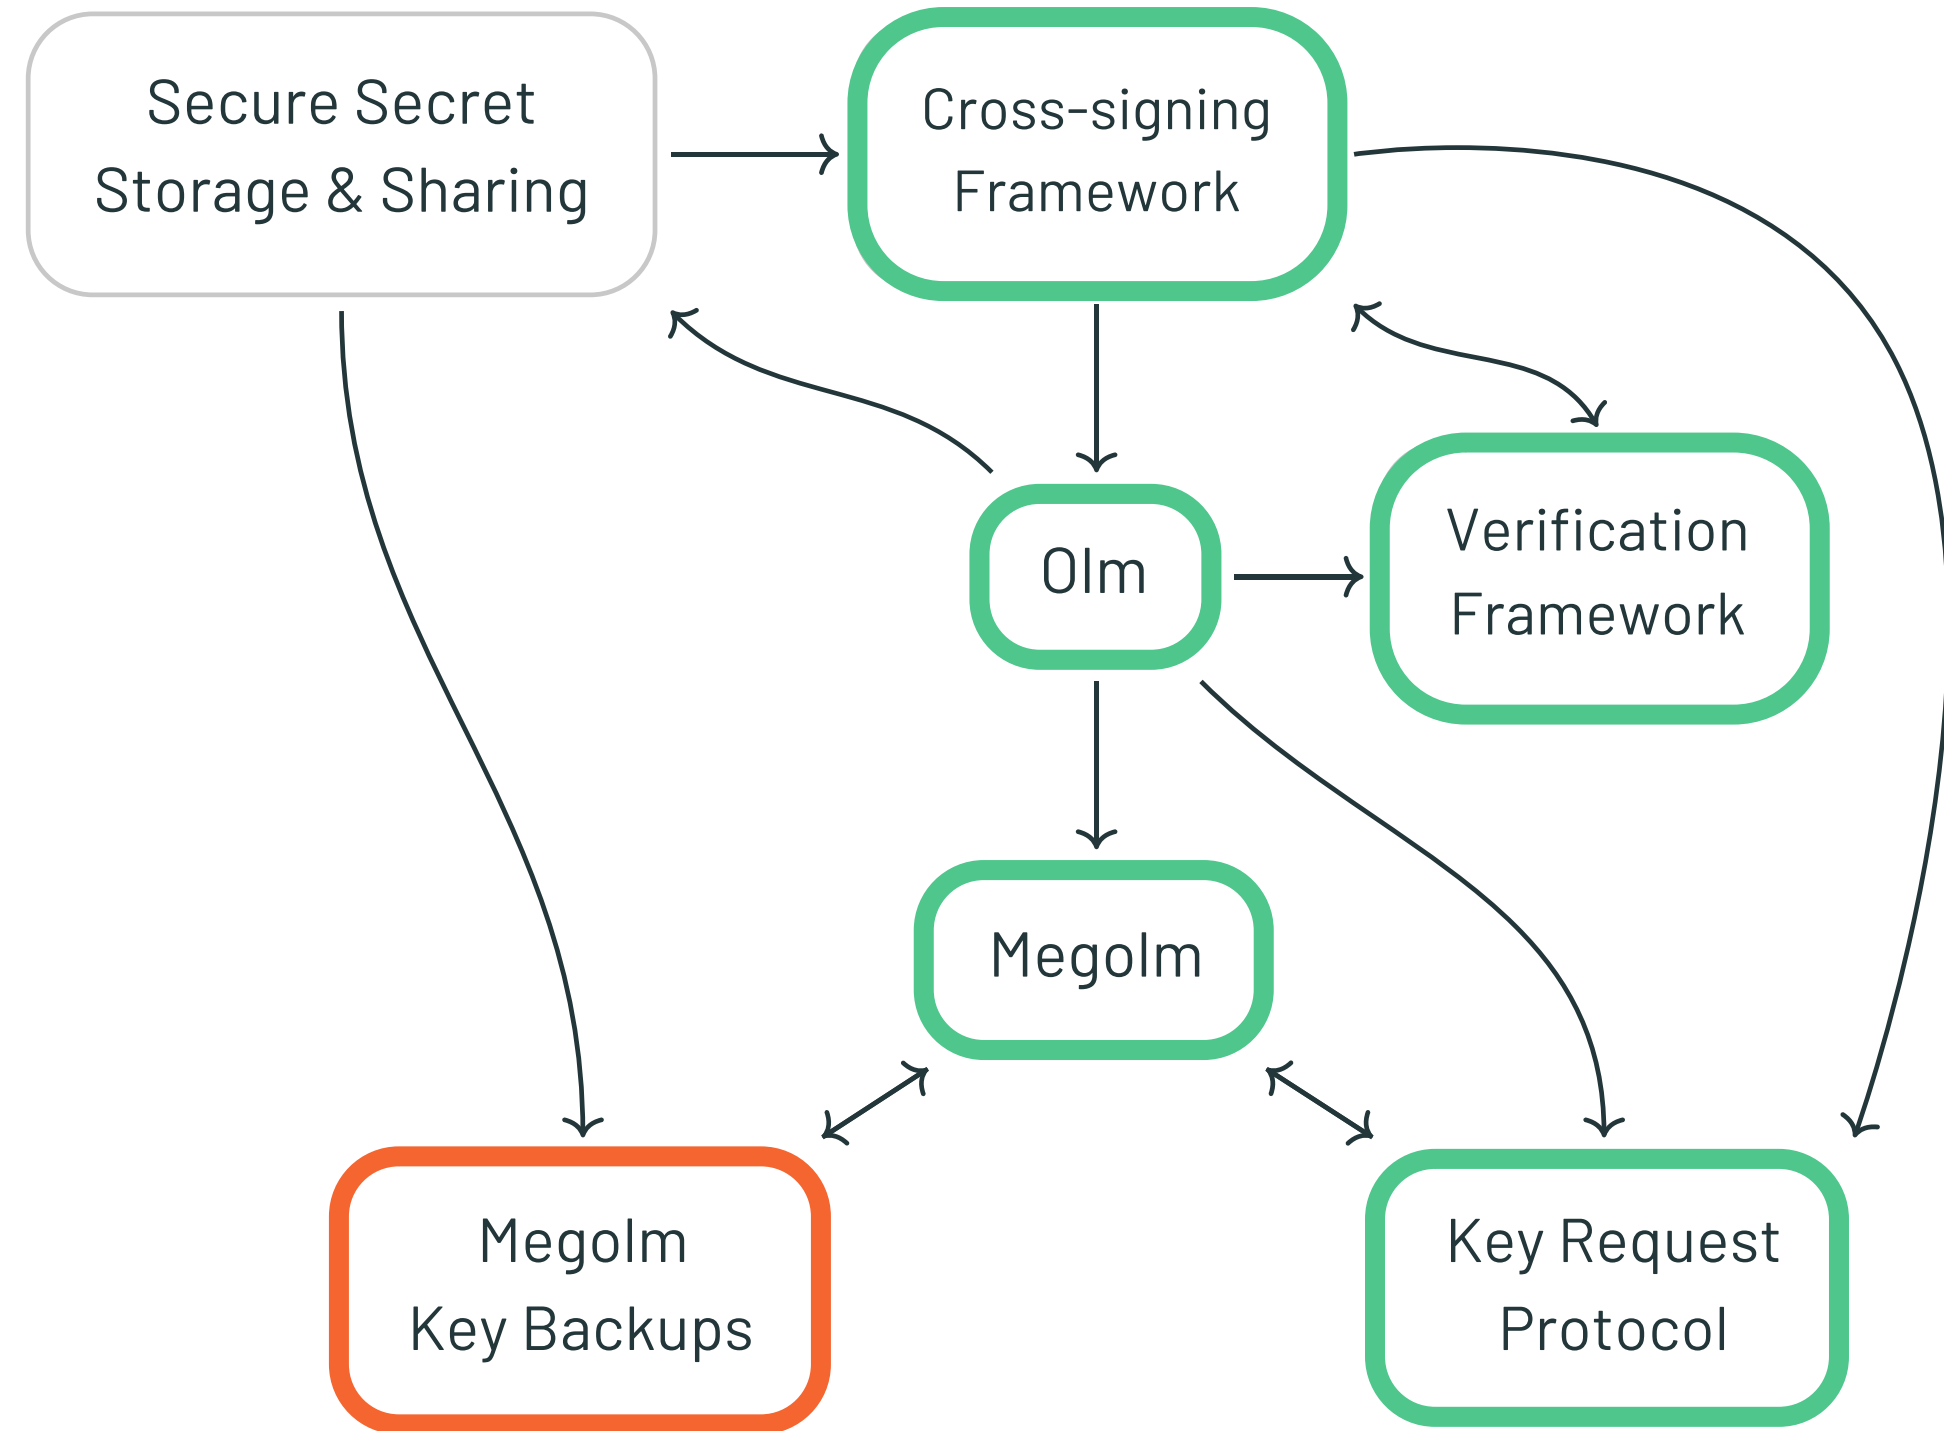
Messages from B* are not marked as FWD.

Megolm Sessions
(B, $\mathcal{G}_{gpk}^B, \sigma_{mg}^B$, FWD)
(B, $\mathcal{G}_{gpk}^{B^*}, \sigma_{mg}^{B^*}$)
⋮



Megolm Key Backups

- Asynchronous alternative to Key Request protocol.
- Inbound Megolm sessions are encrypted and saved to the homeserver.
- Encrypt using a secret key shared between a user's devices.



Megolm Key Backups

1. Setup



$bsk, bpk \leftarrow X25519.KGen(1^n)$

$auth_data(bpk)$

$auth_data(bpk)$

$auth_data(bpk)$

$D^{A,1}$ marks bpk as trusted since it has bsk

$D^{A,2}$ marks bpk as untrusted w/out bsk

Megolm Key Backups

1. Setup



$bsk, bpk \leftarrow X25519.KGen(1^n)$

$auth_data(bpk)$

$auth_data(bpk)$

$auth_data(bpk)$

$D^{A,1}$ marks bpk as trusted since it has bsk

$D^{A,2}$ marks bpk as untrusted w/out bsk

$D^{A,1}$ and $D^{A,2}$ share bsk using SSSS (or out-of-band)

$D^{A,2}$ marks bpk as trusted since it has bsk

...

...

...

Megolm Key Backups

2. Backup




1. Generate ephemeral key for ECDH with bpk
2. Derive encryption keys with HKDF
3. Encode session
4. Encrypt-then-MAC



Megolm Key Backups

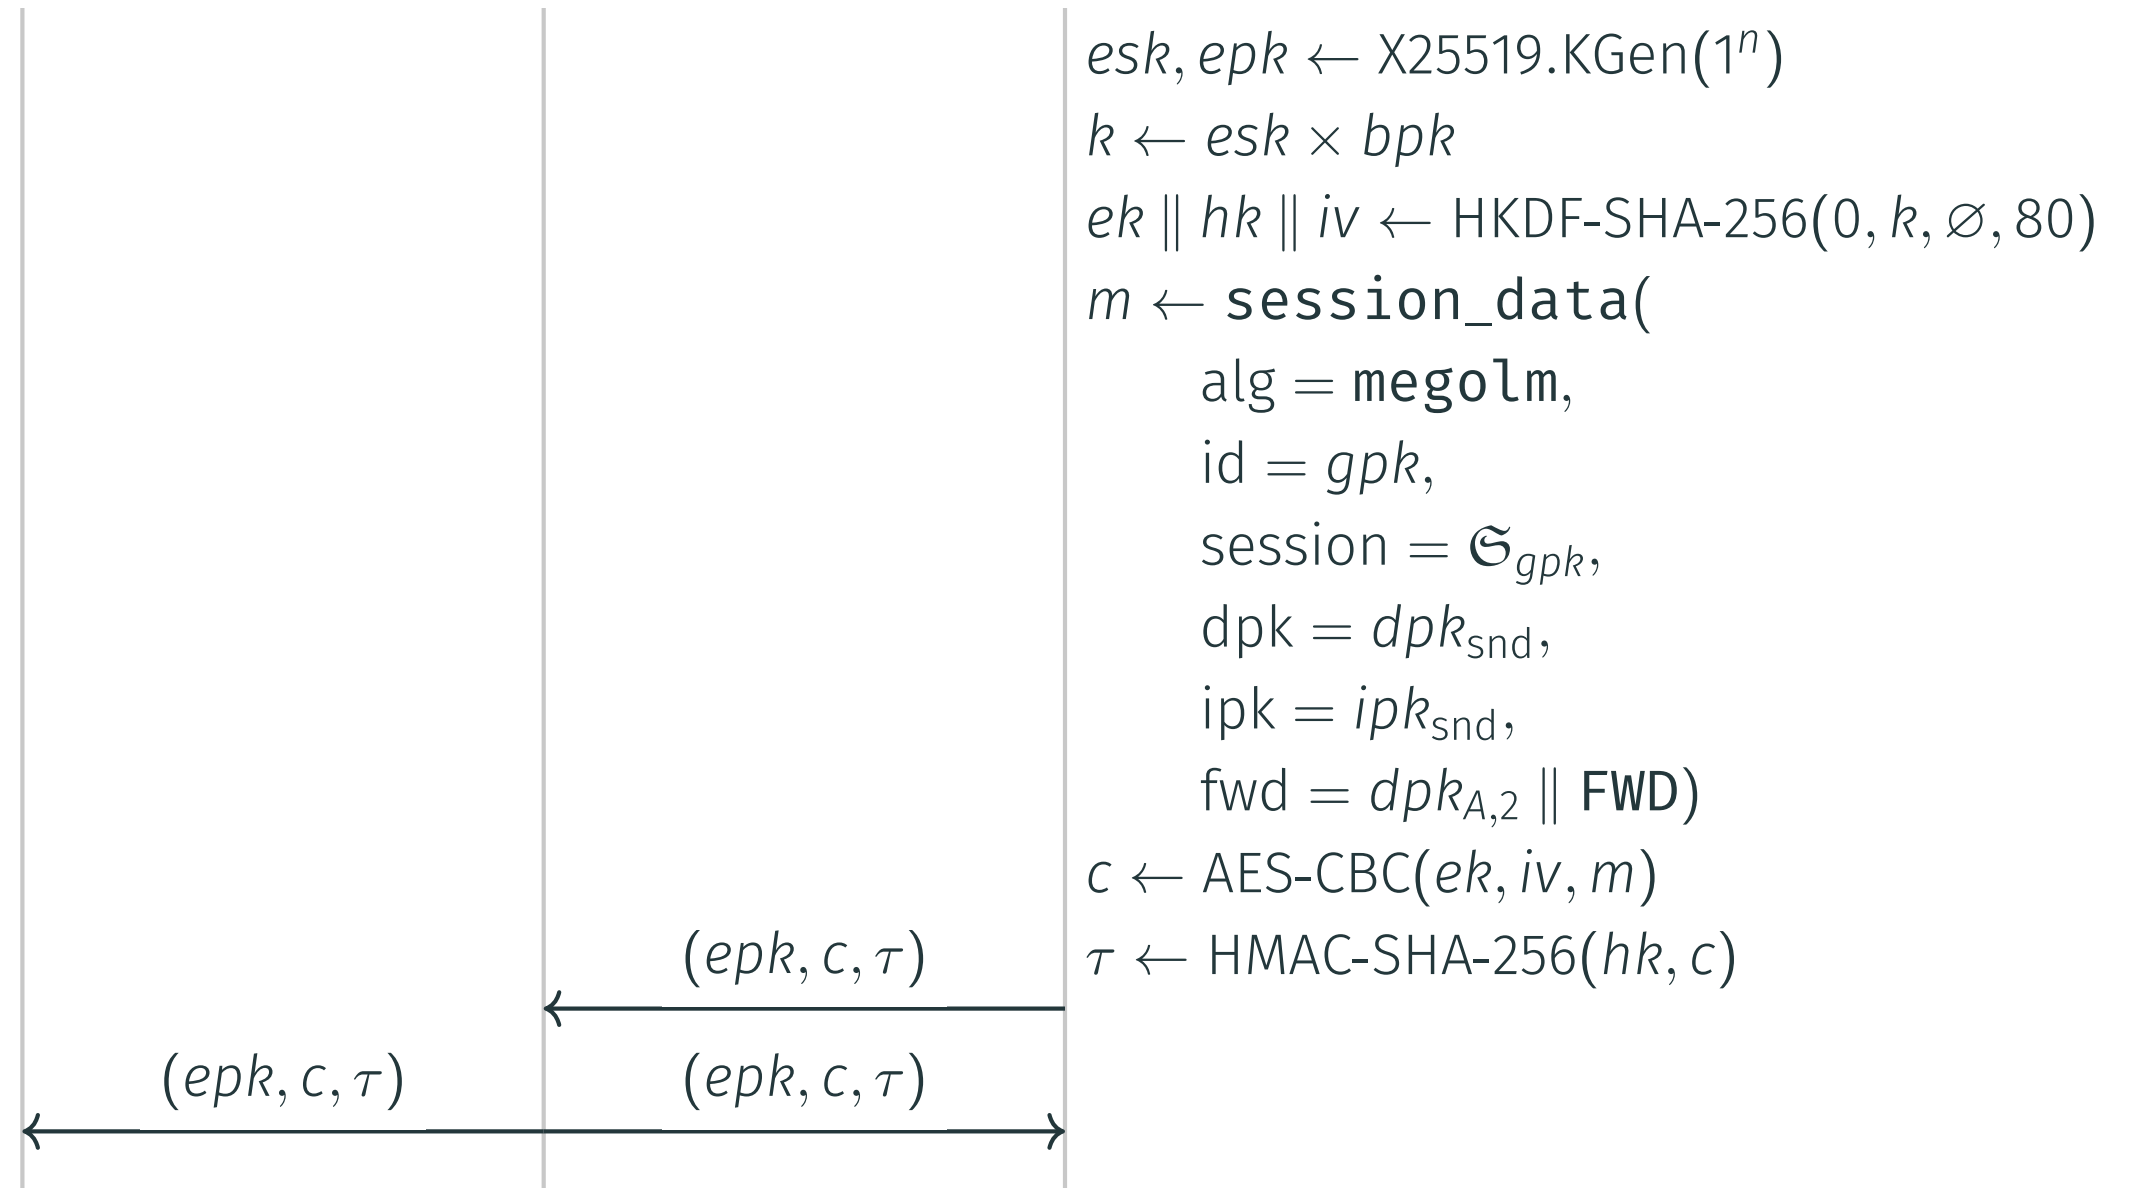
2. Backup



$esk, epk \leftarrow X25519.KGen(1^n)$
 $k \leftarrow esk \times bpk$
 $ek \parallel hk \parallel iv \leftarrow HKDF-SHA-256(0, k, \emptyset, 80)$
3. Encode session
4. Encrypt-then-MAC

Megolm Key Backups

2. Backup



Secure Secret Storage & Sharing

Backup, recover and share *user-level secrets*.

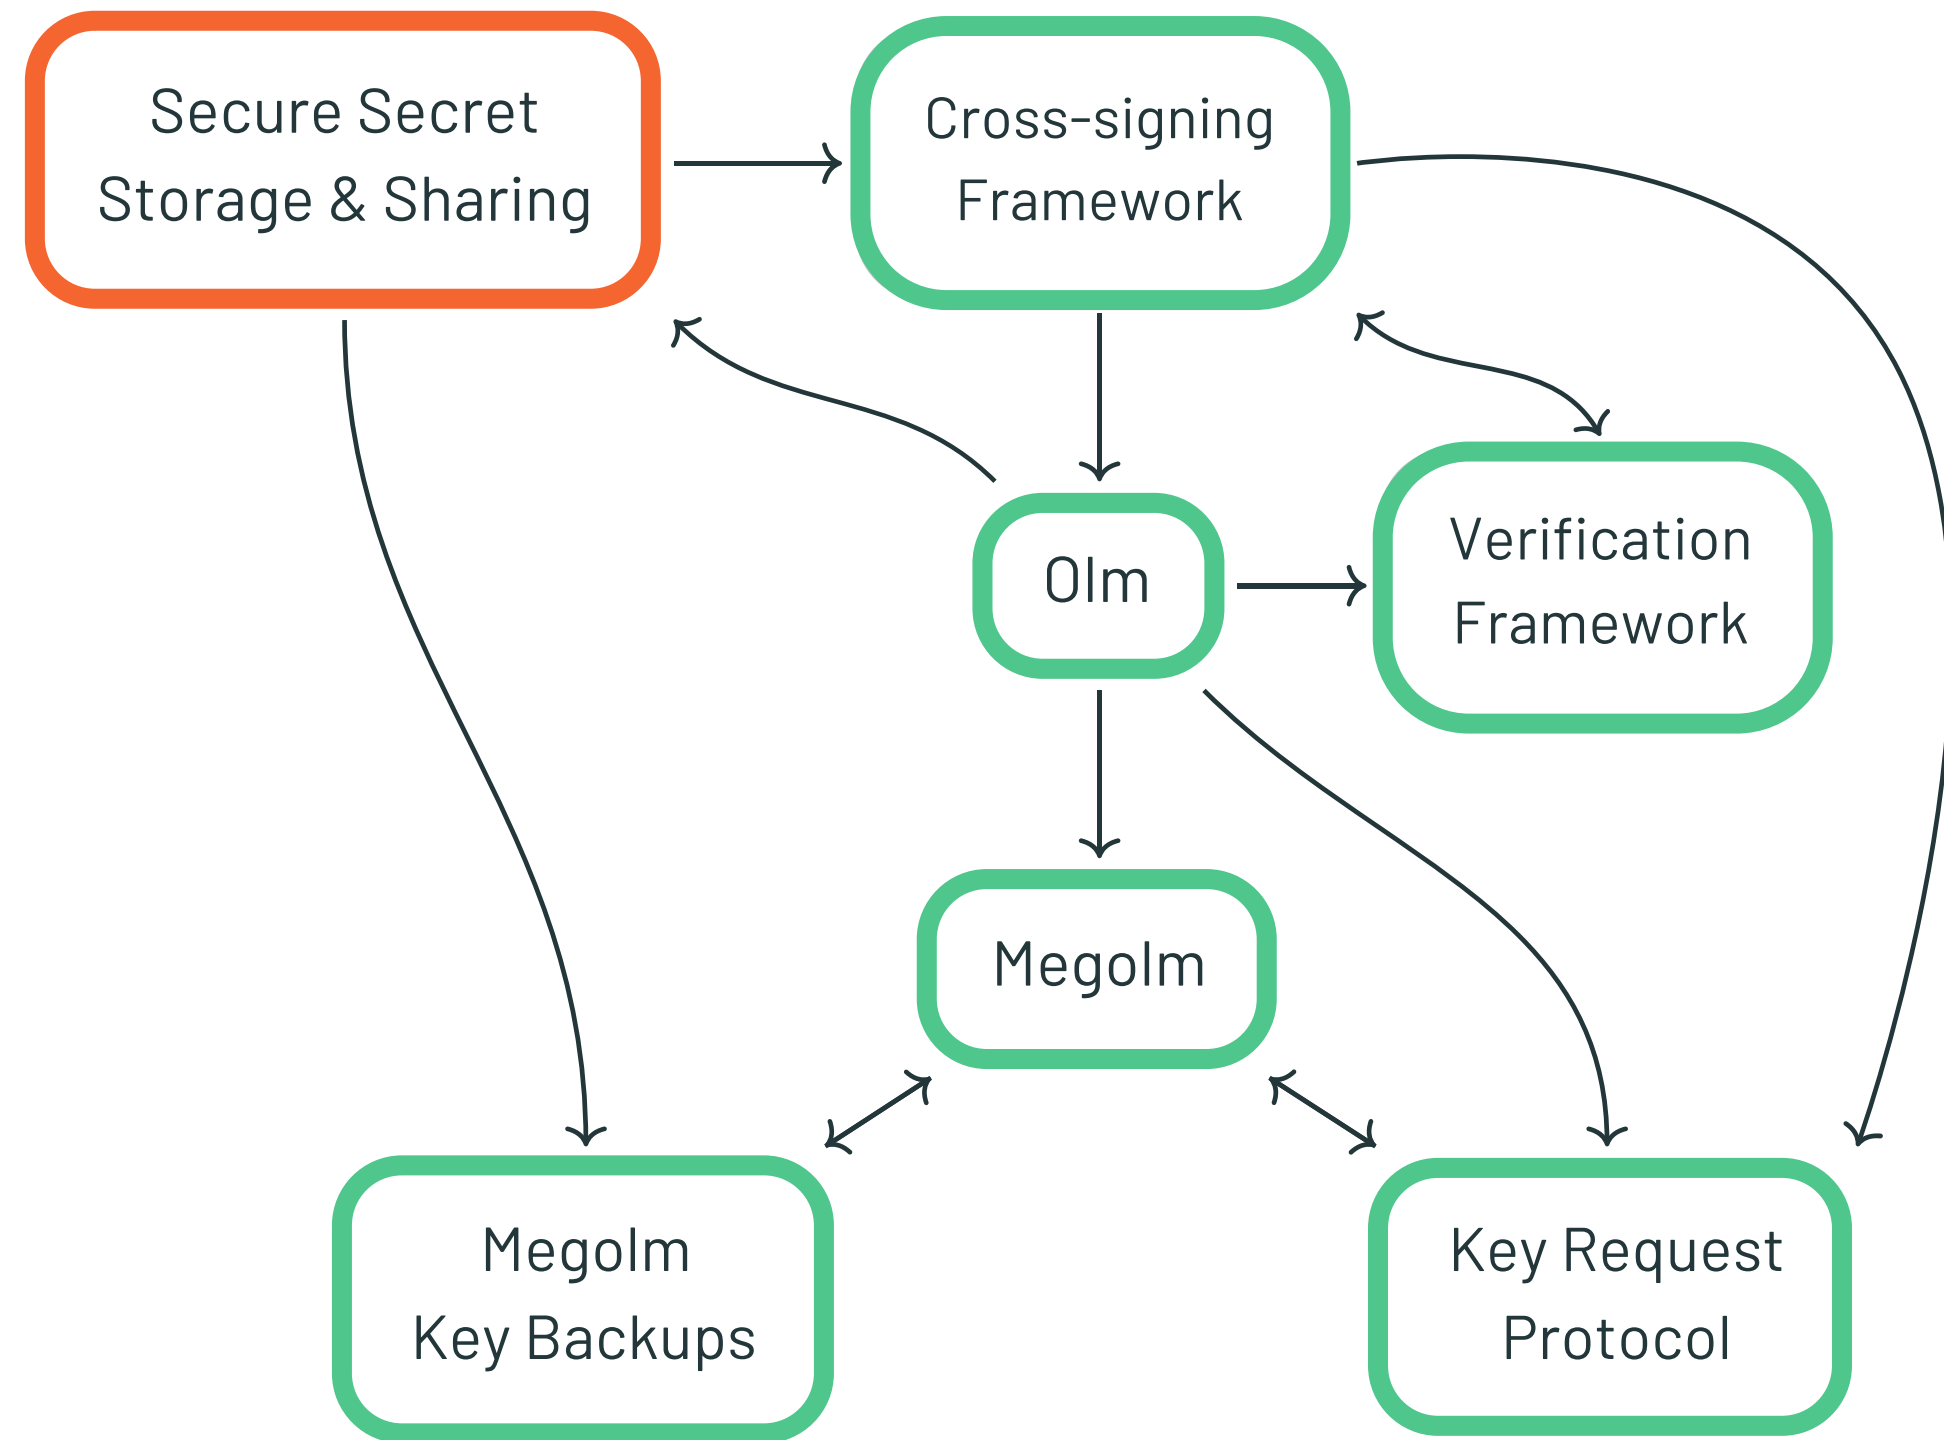
E.g. Alice's (*msk*, *usk*, *ssk*).

Secret Storage:

- Encrypt secrets and store on homeserver.
- Shared symmetric key (may be password-derived).

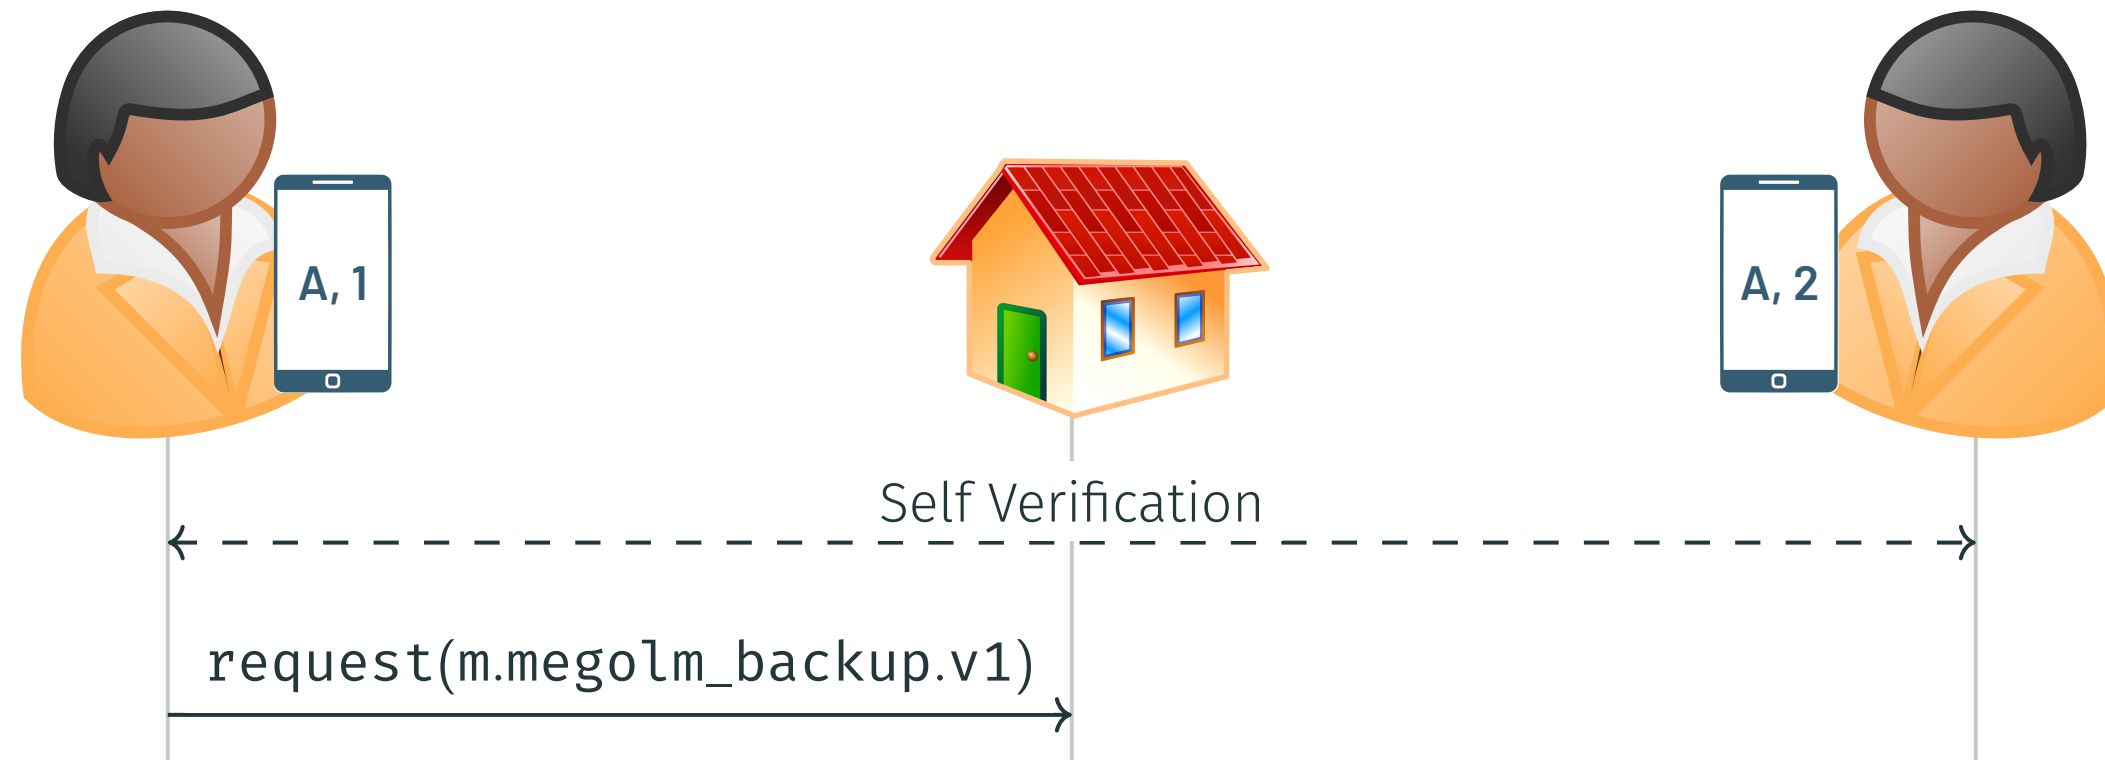
Secret Sharing:

- Use Olm to share secrets to newly verified devices.



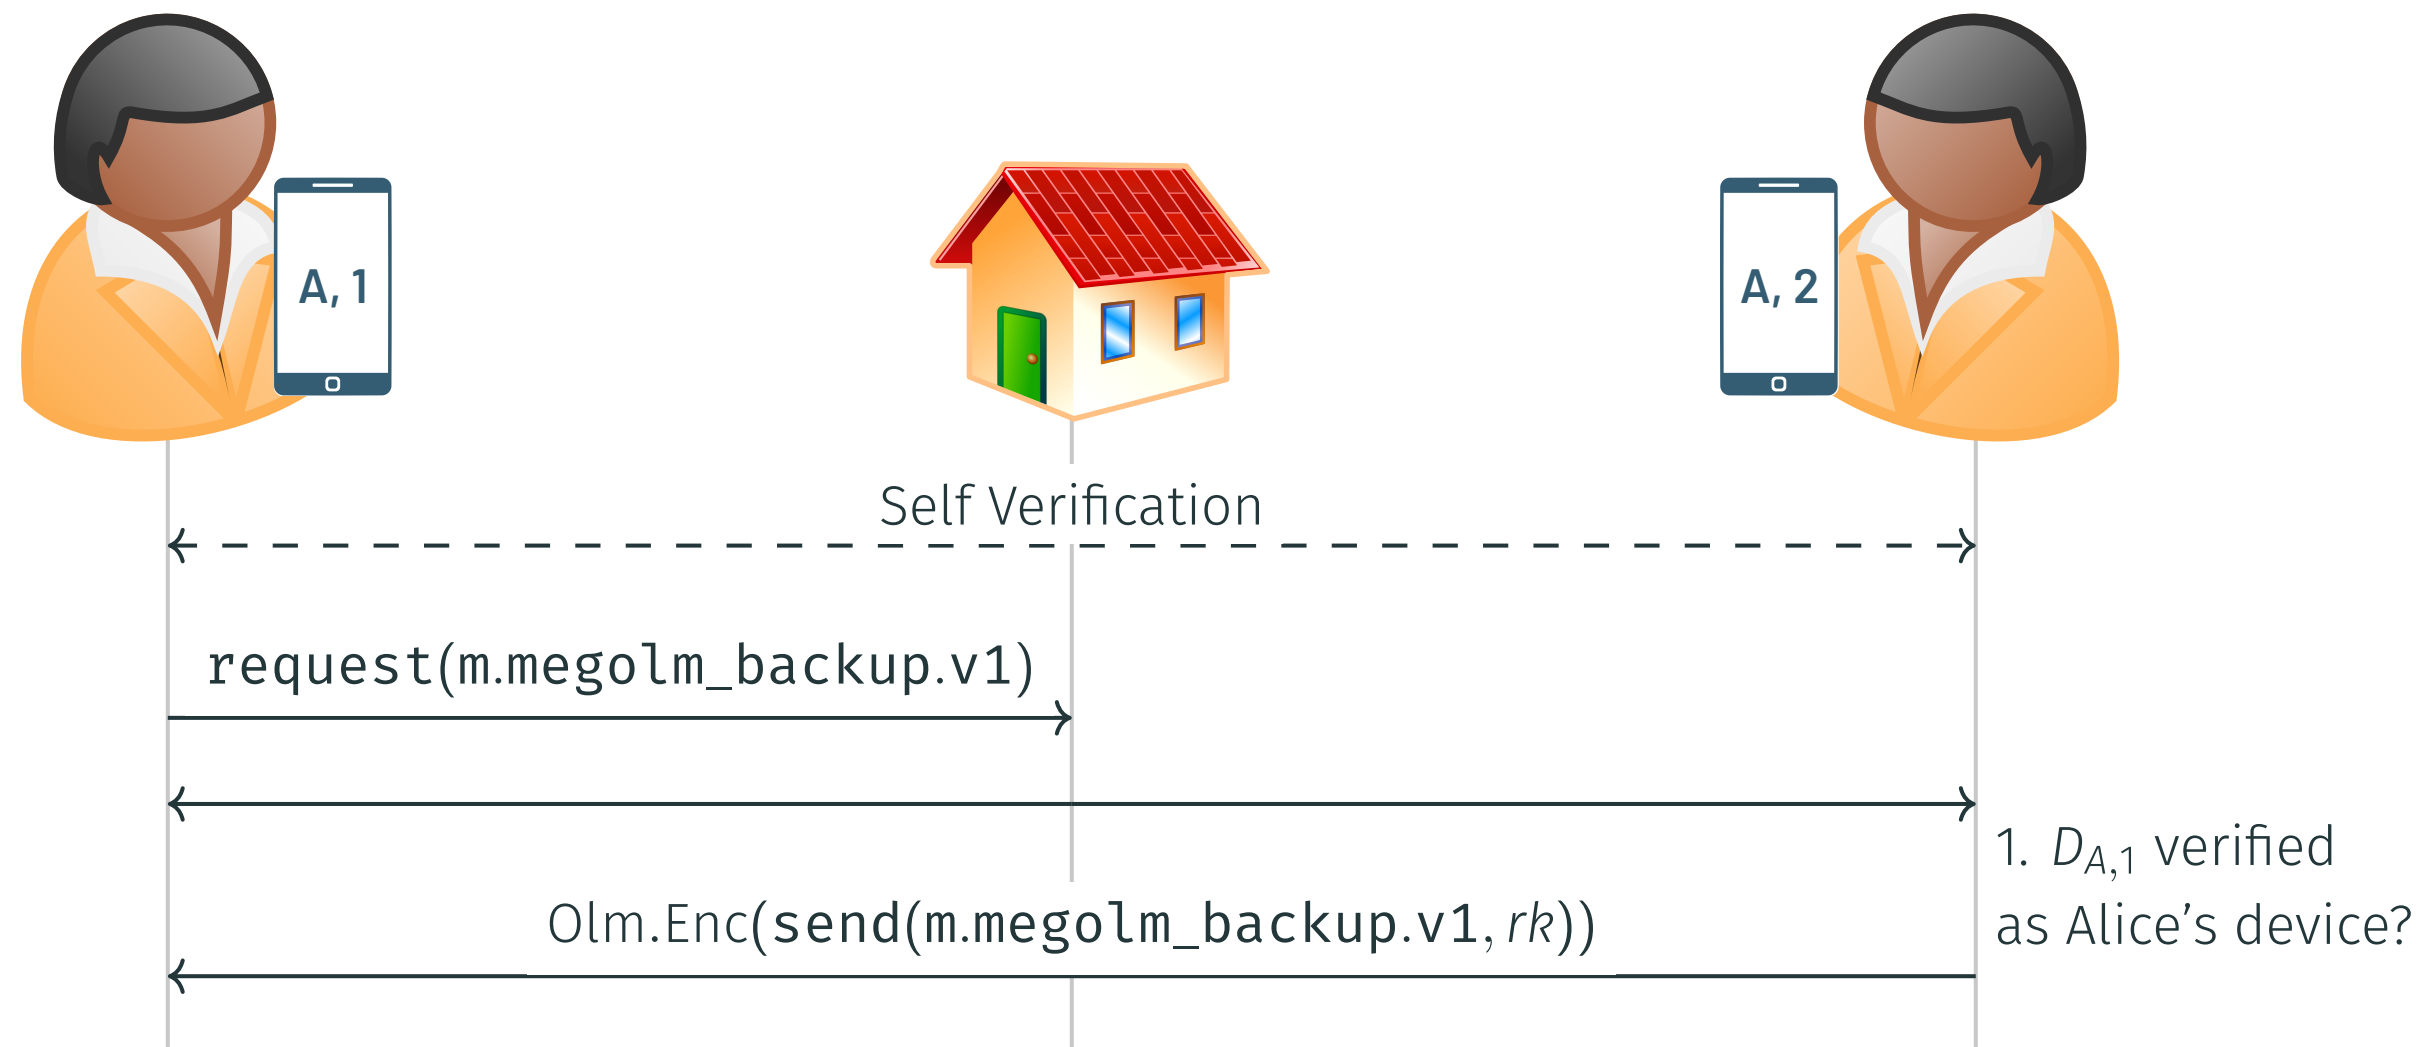
Secure Secret Sharing

Example: Sharing secret for Megolm Key Backups



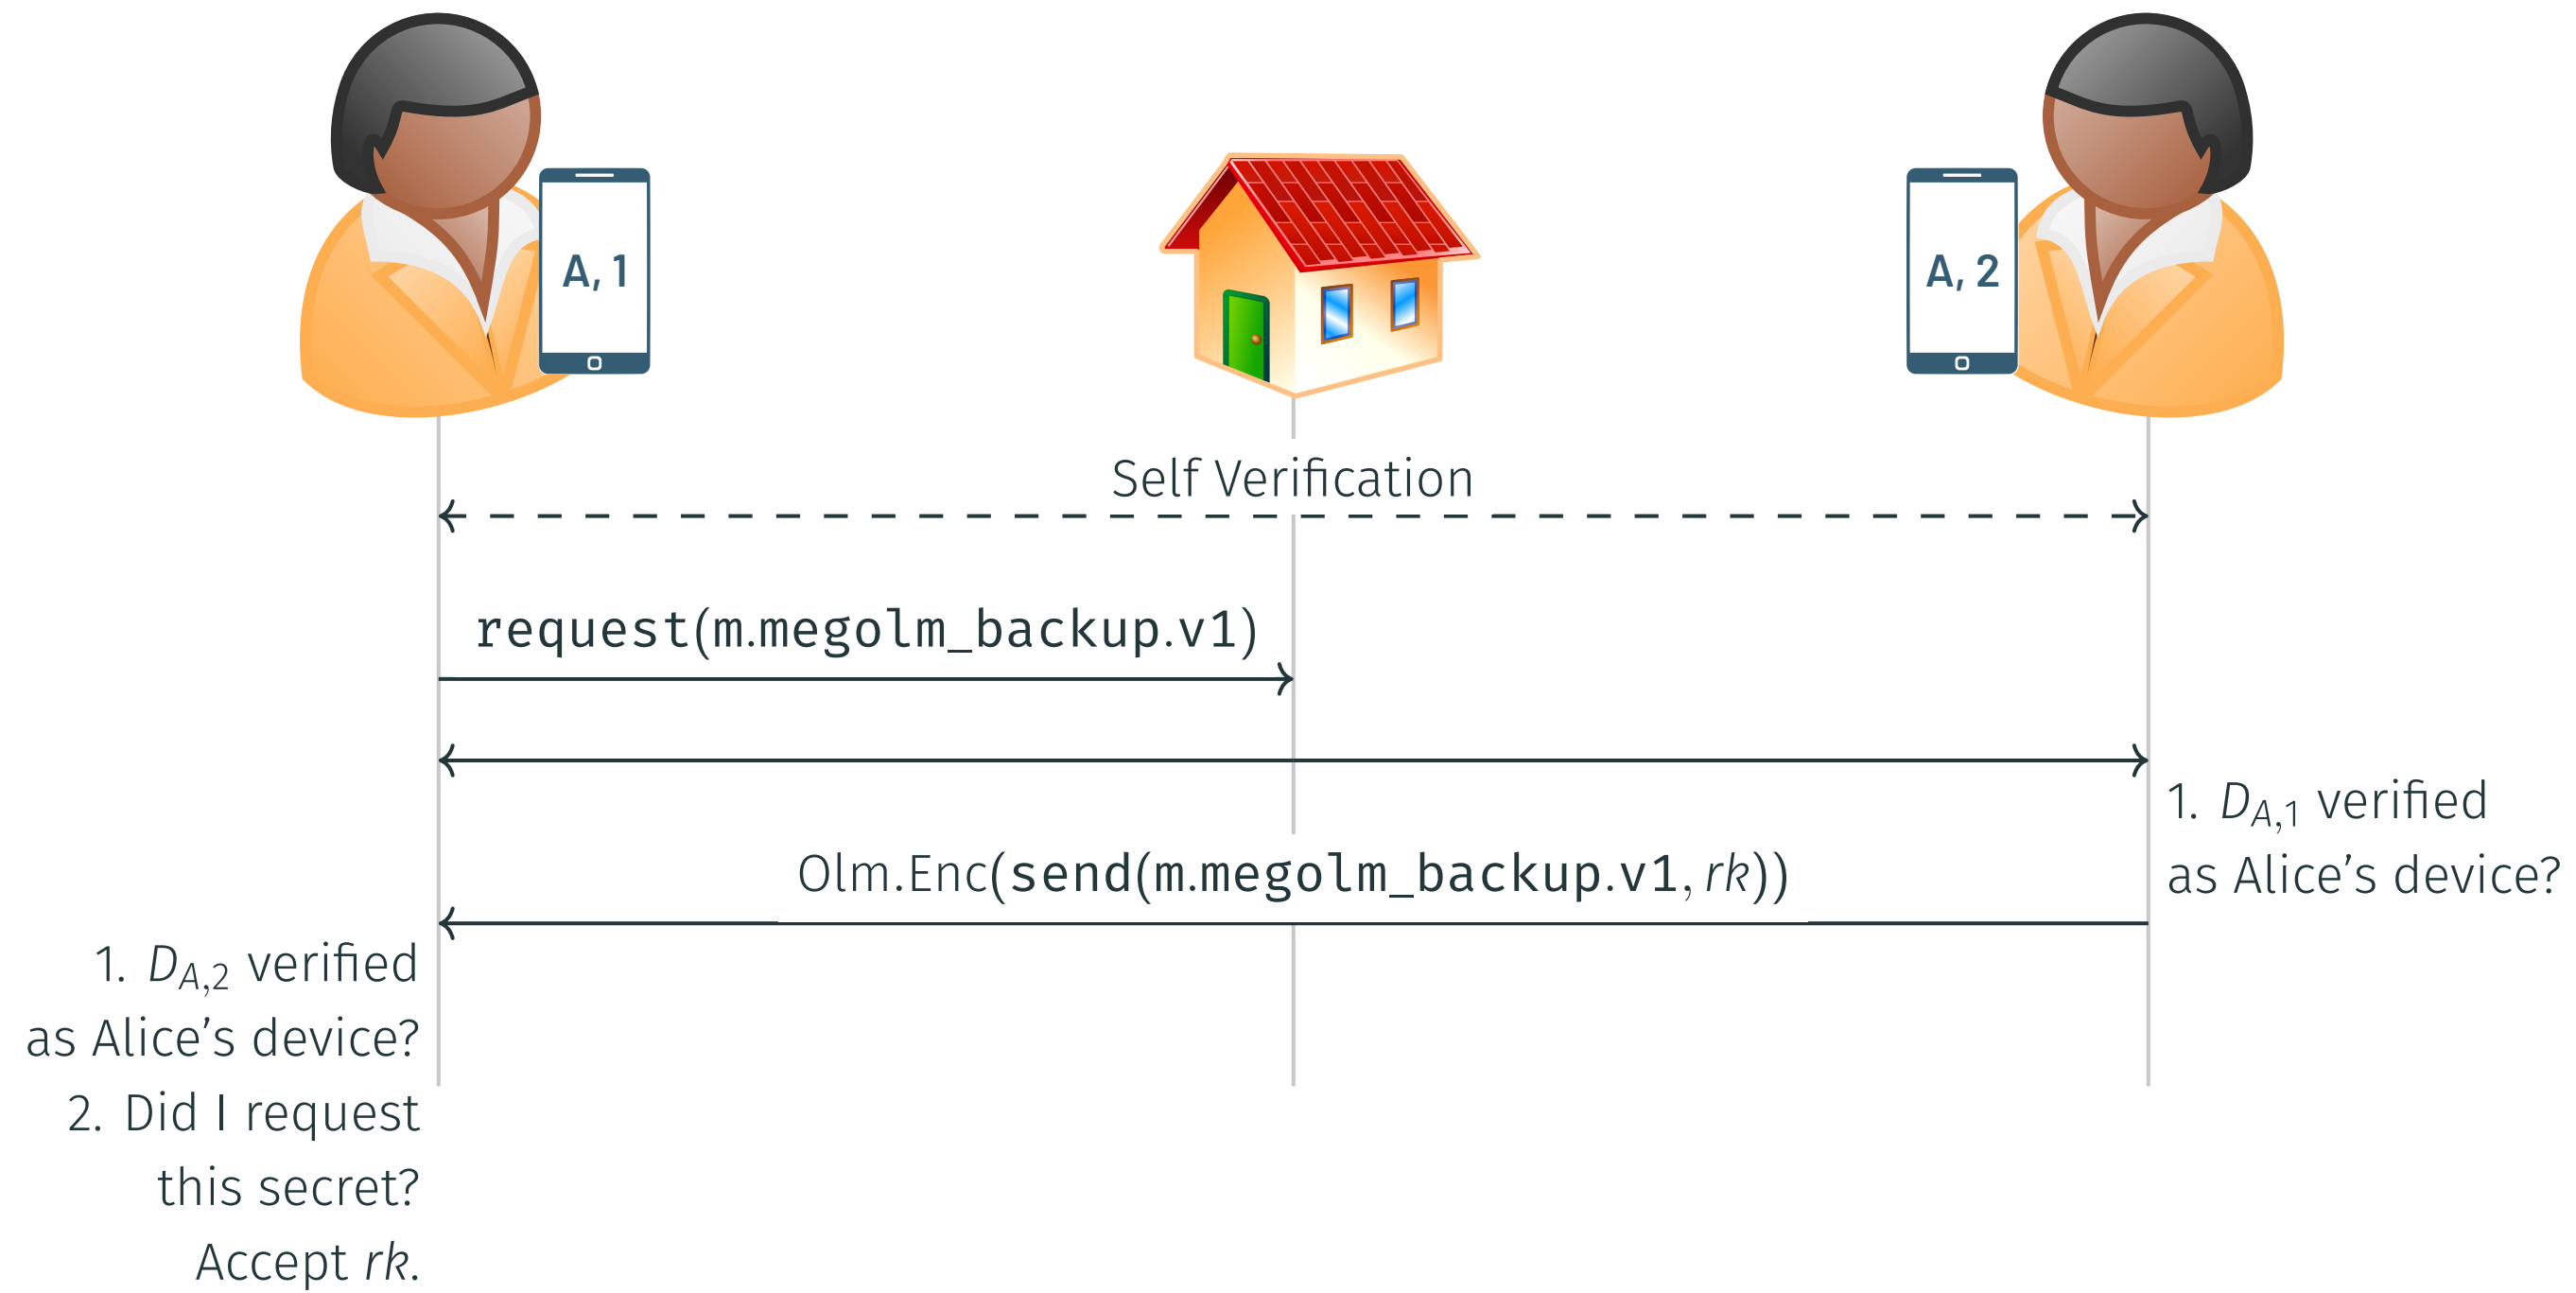
Secure Secret Sharing

Example: Sharing secret for Megolm Key Backups



Secure Secret Sharing

Example: Sharing secret for Megolm Key Backups



Adversary controlled backup keys

Since Secure Secret Sharing uses Olm to share secrets between devices (incl. Megolm backup keys)

Can we pull off the same trick off again?



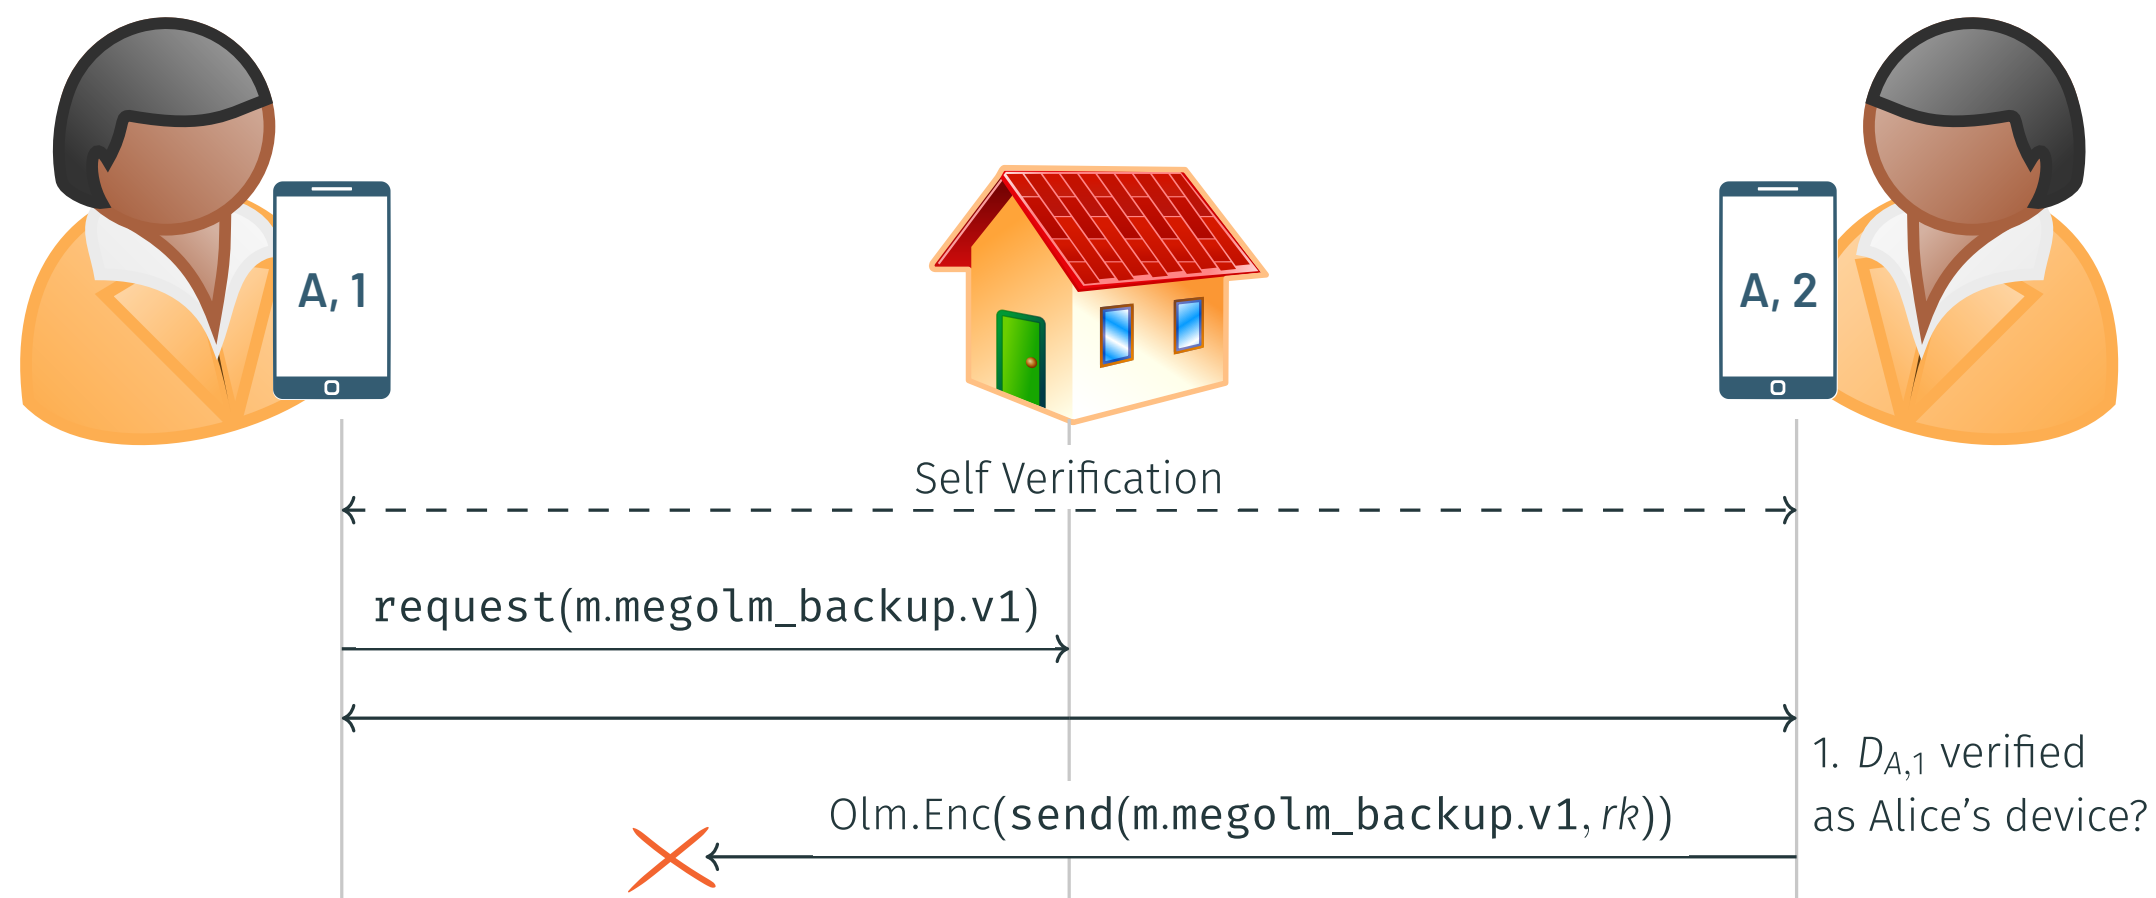
ATTACK

Adversary controlled backup keys

Since Secure Secret Sharing uses Olm to share secrets between devices (incl. Megolm backup keys)

Can we pull off the same trick off again?

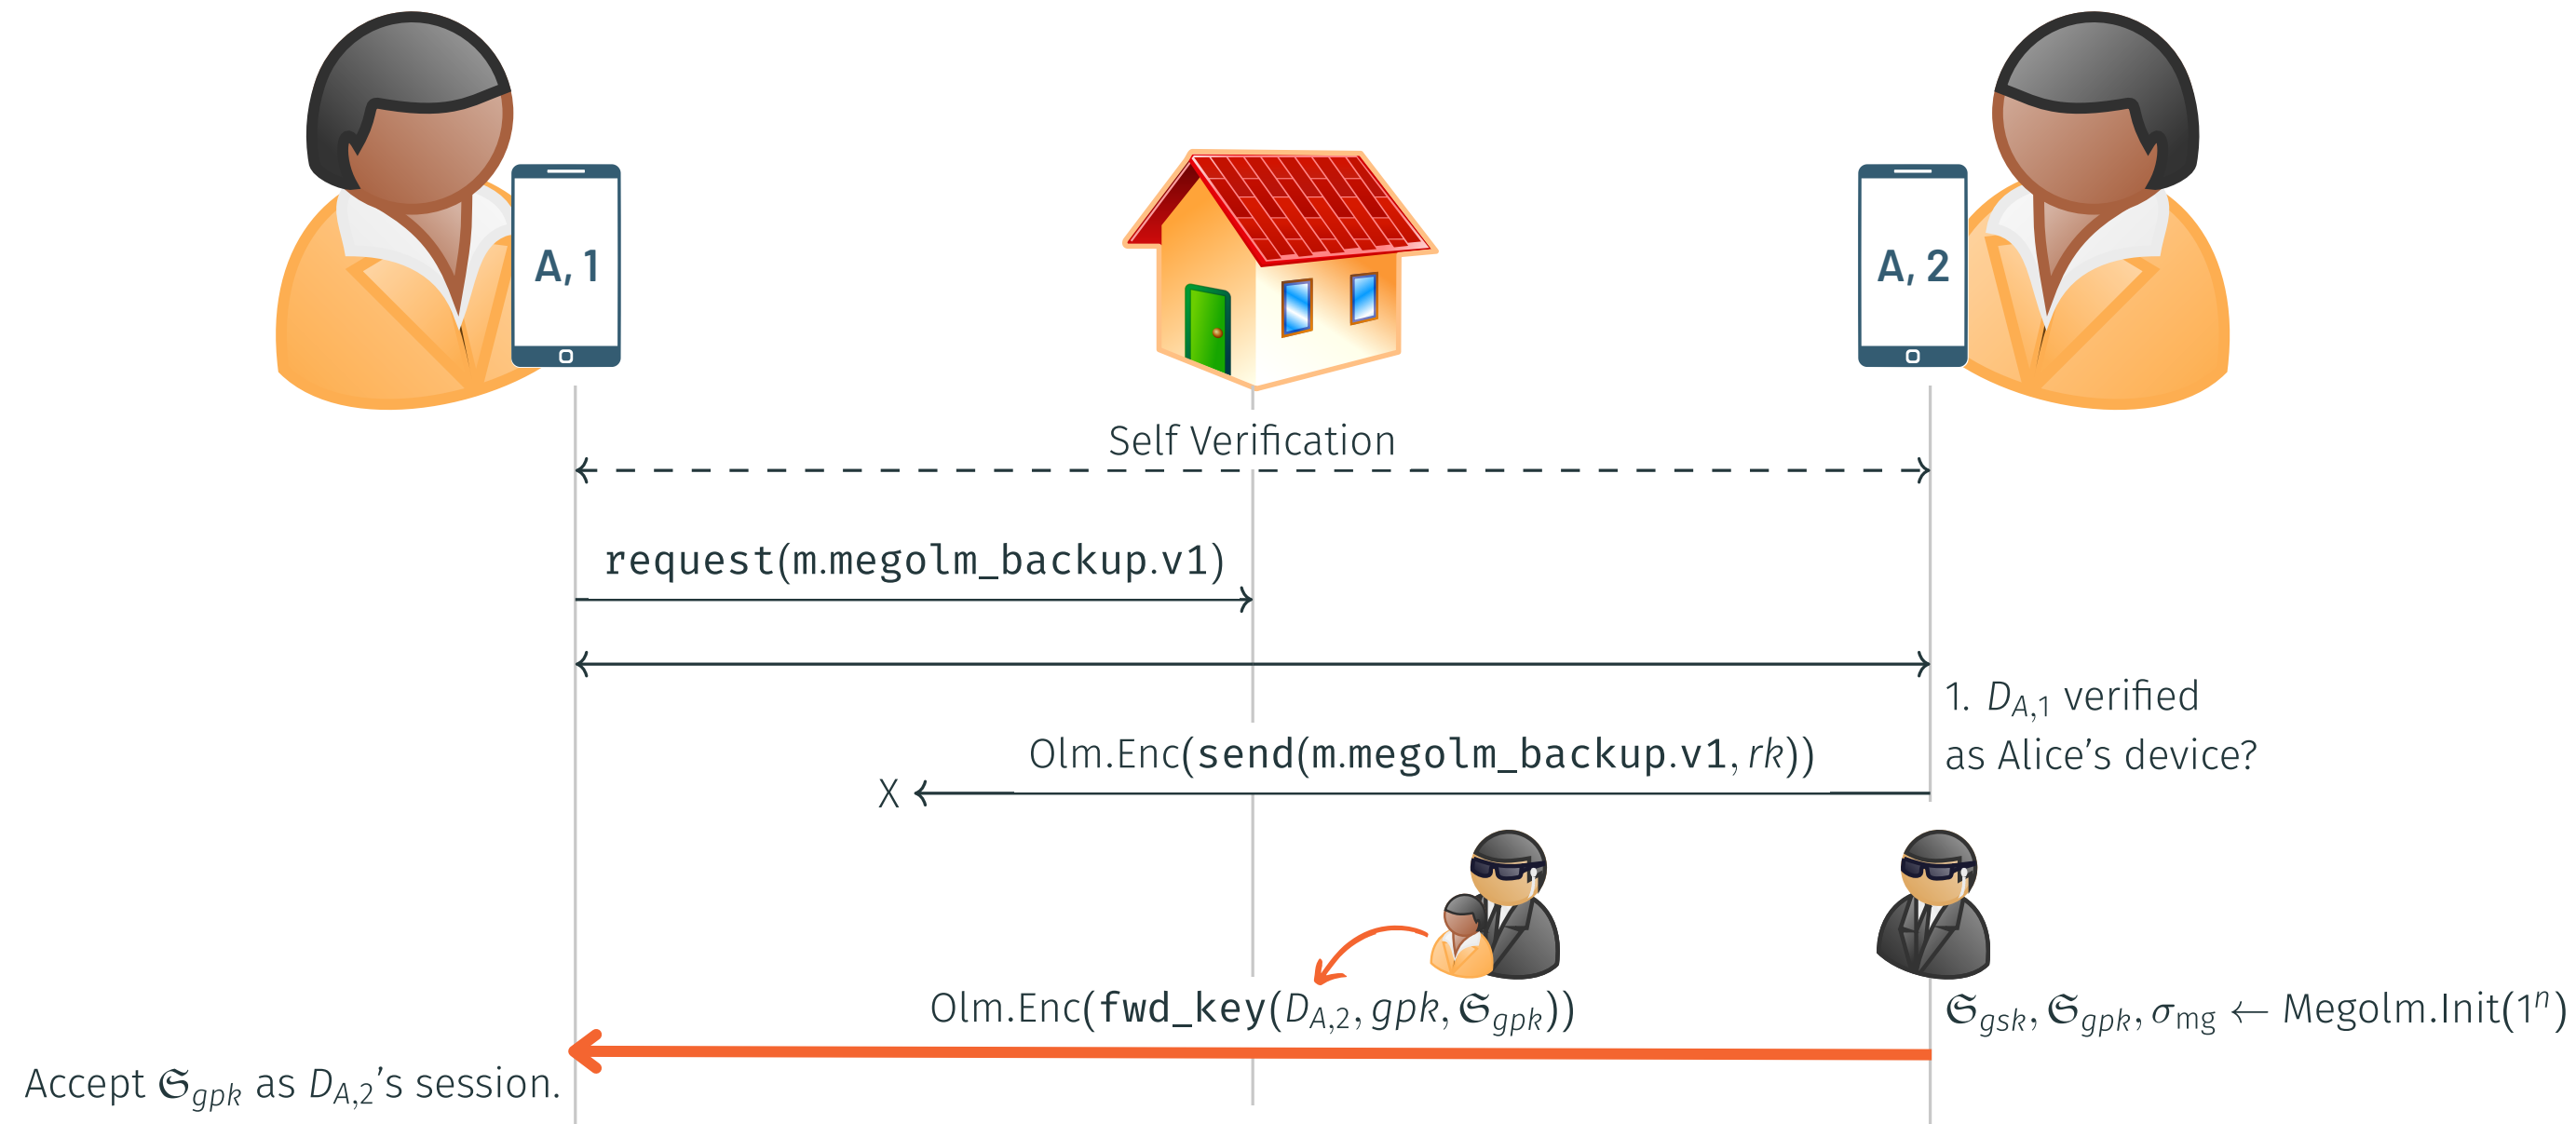
ATTACK



Adversary controlled backup keys

Since Secure Secret Sharing uses Olm to share secrets between devices (incl. Megolm backup keys)

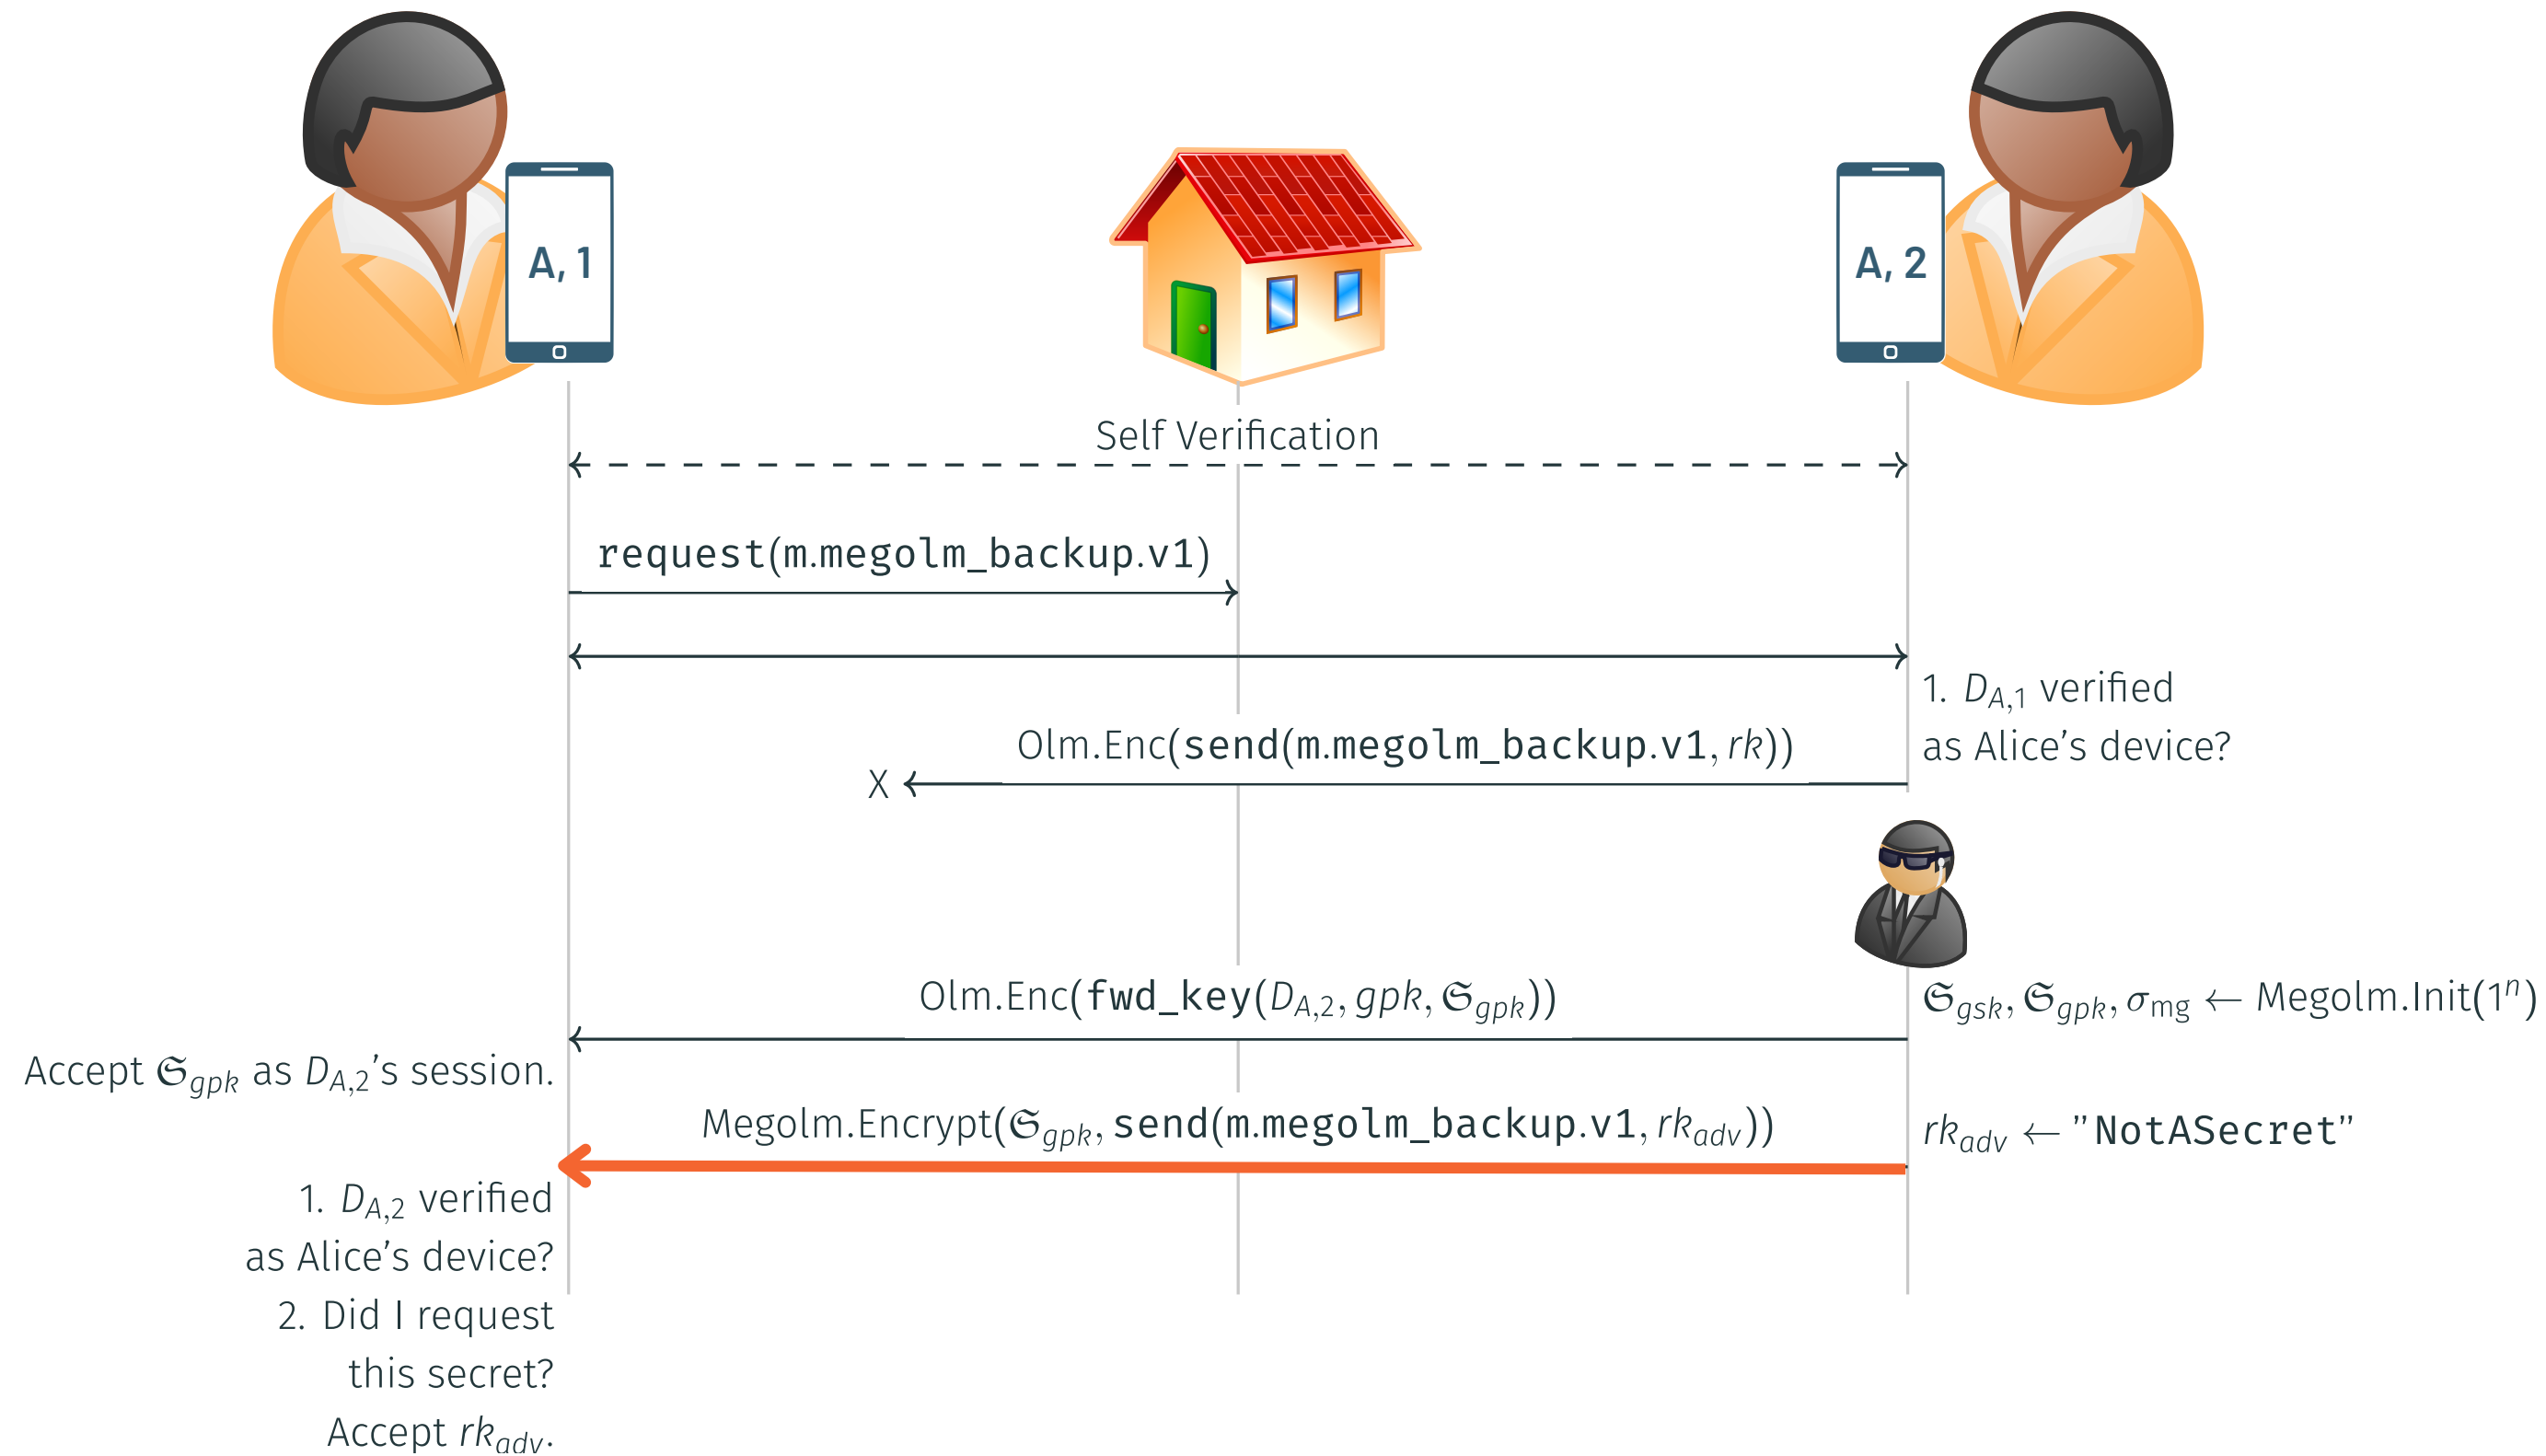
Can we pull off the same trick off again?



Adversary controlled backup keys

Since Secure Secret Sharing uses Olm to share secrets between devices (incl. Megolm backup keys)

Can we pull off the same trick off again?

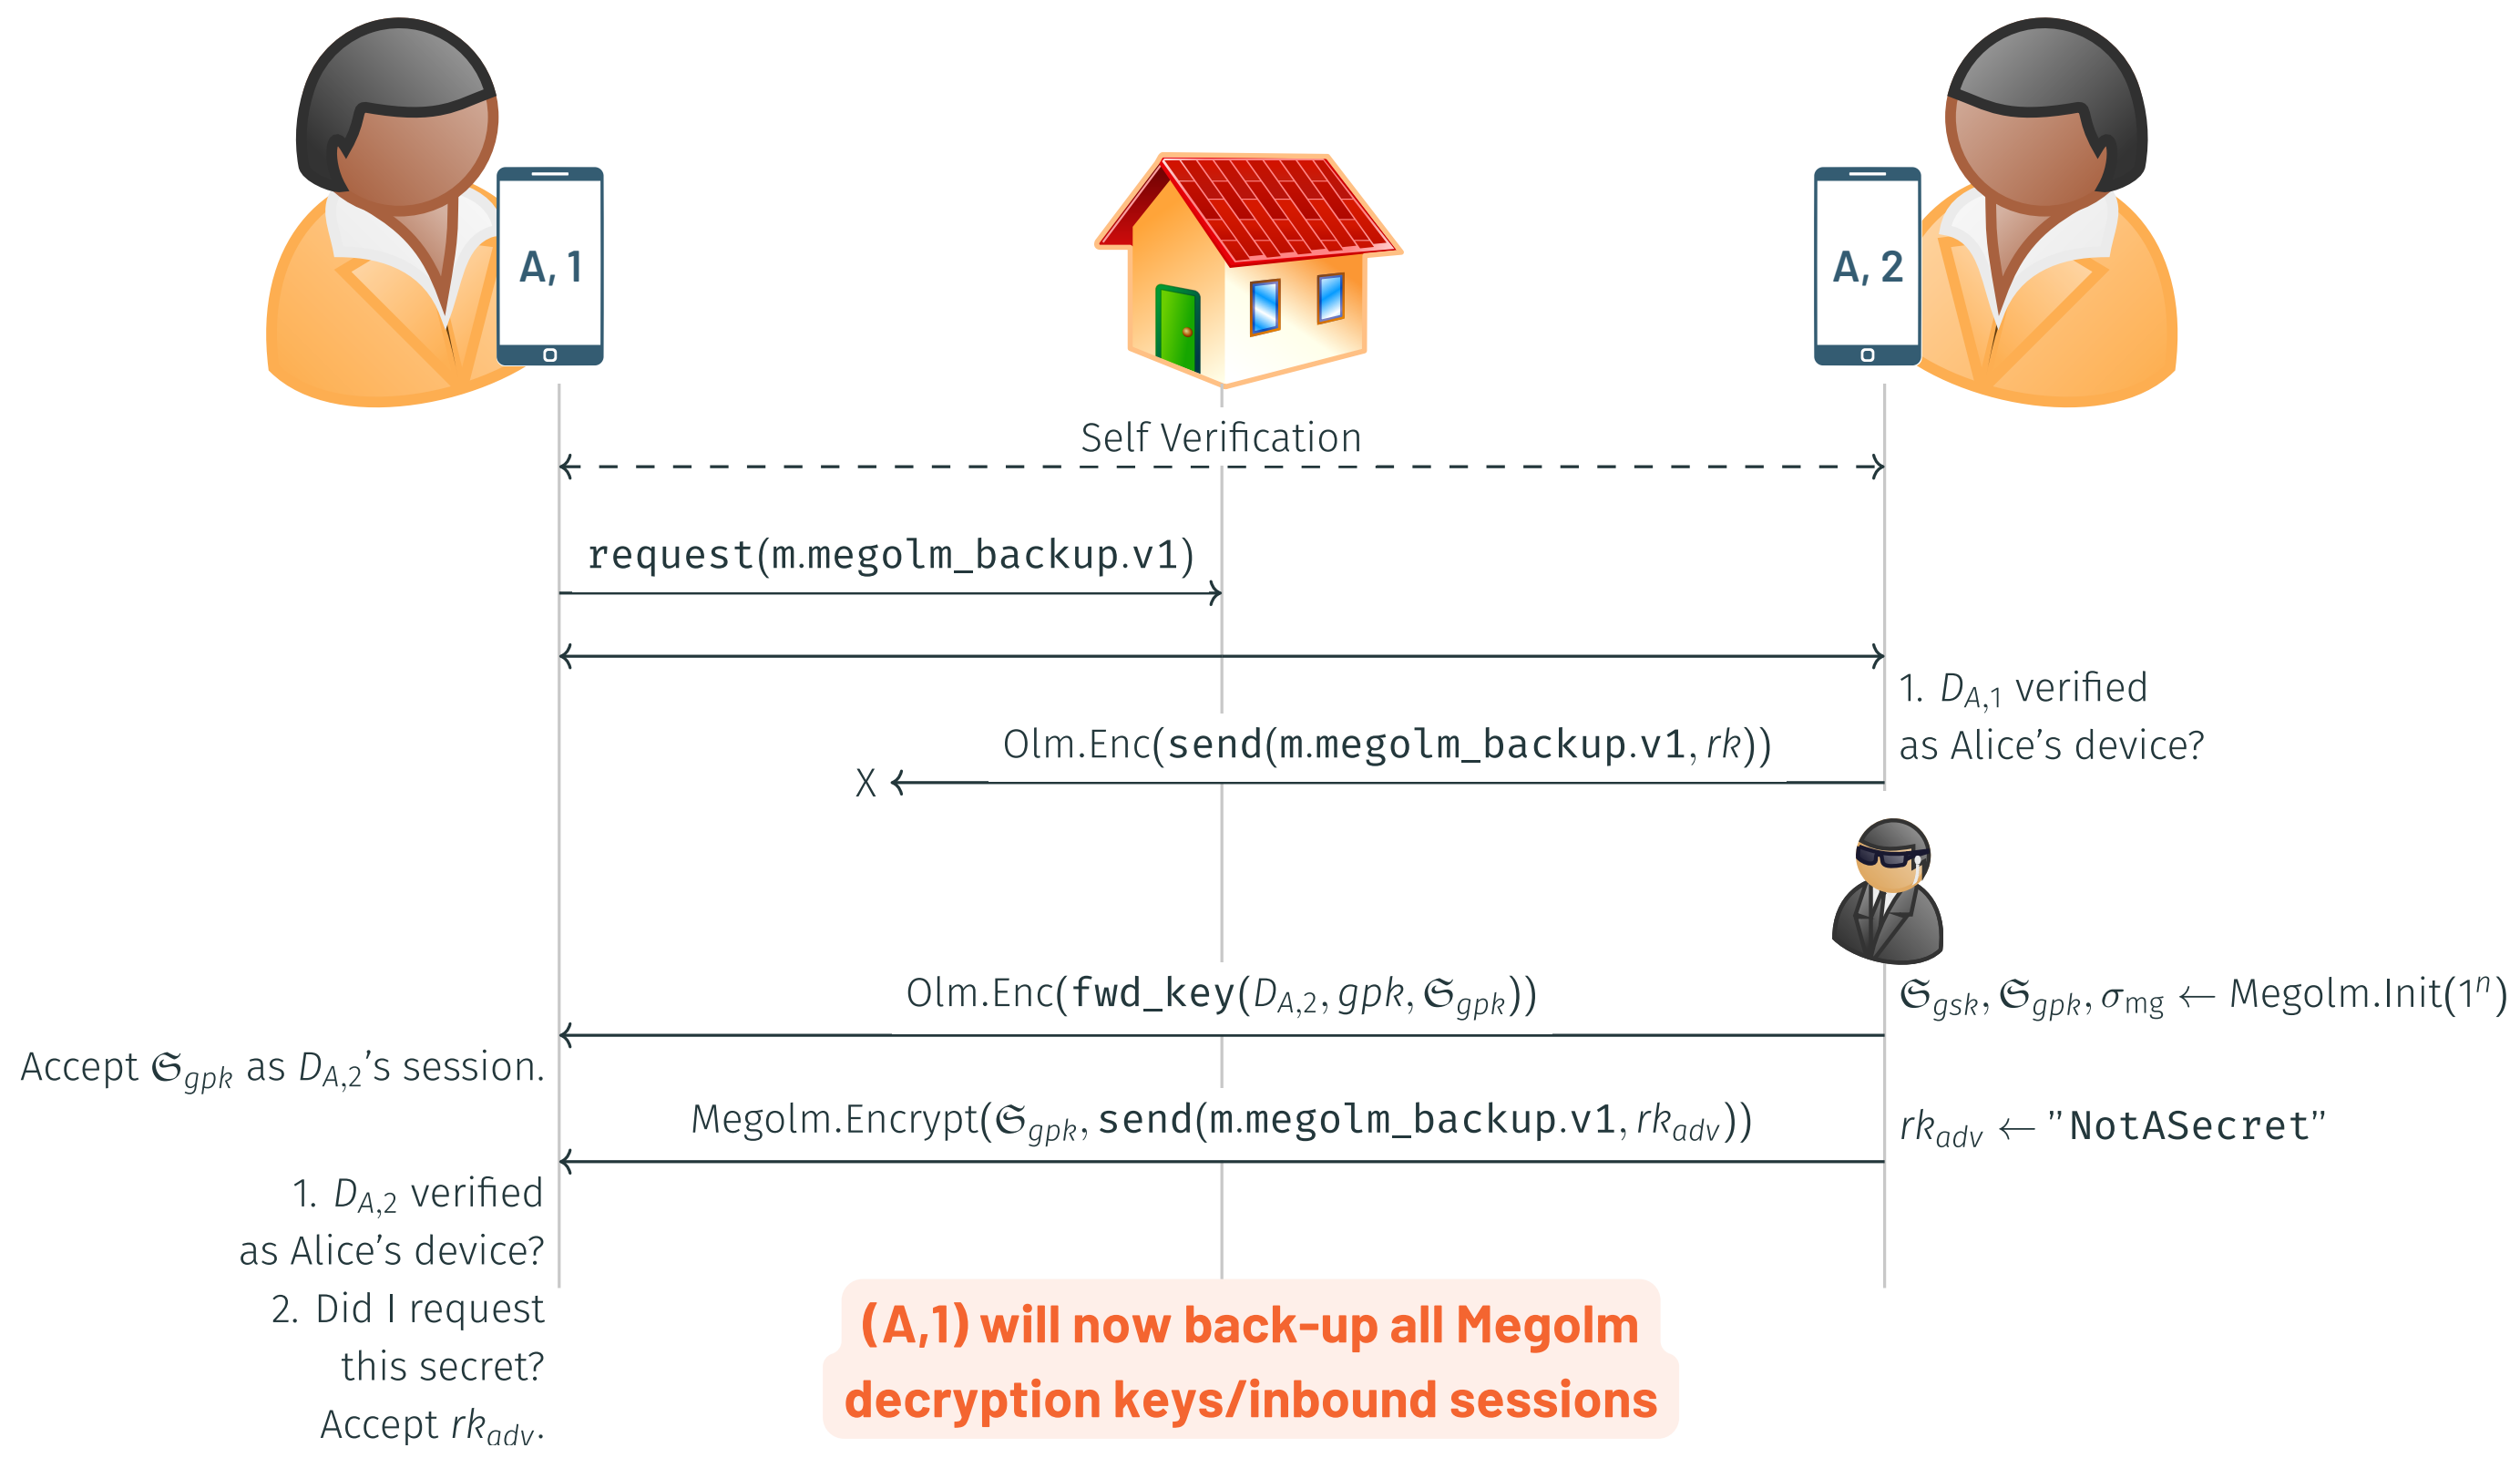


Adversary controlled backup keys

Since Secure Secret Sharing uses Olm to share secrets between devices (incl. Megolm backup keys)

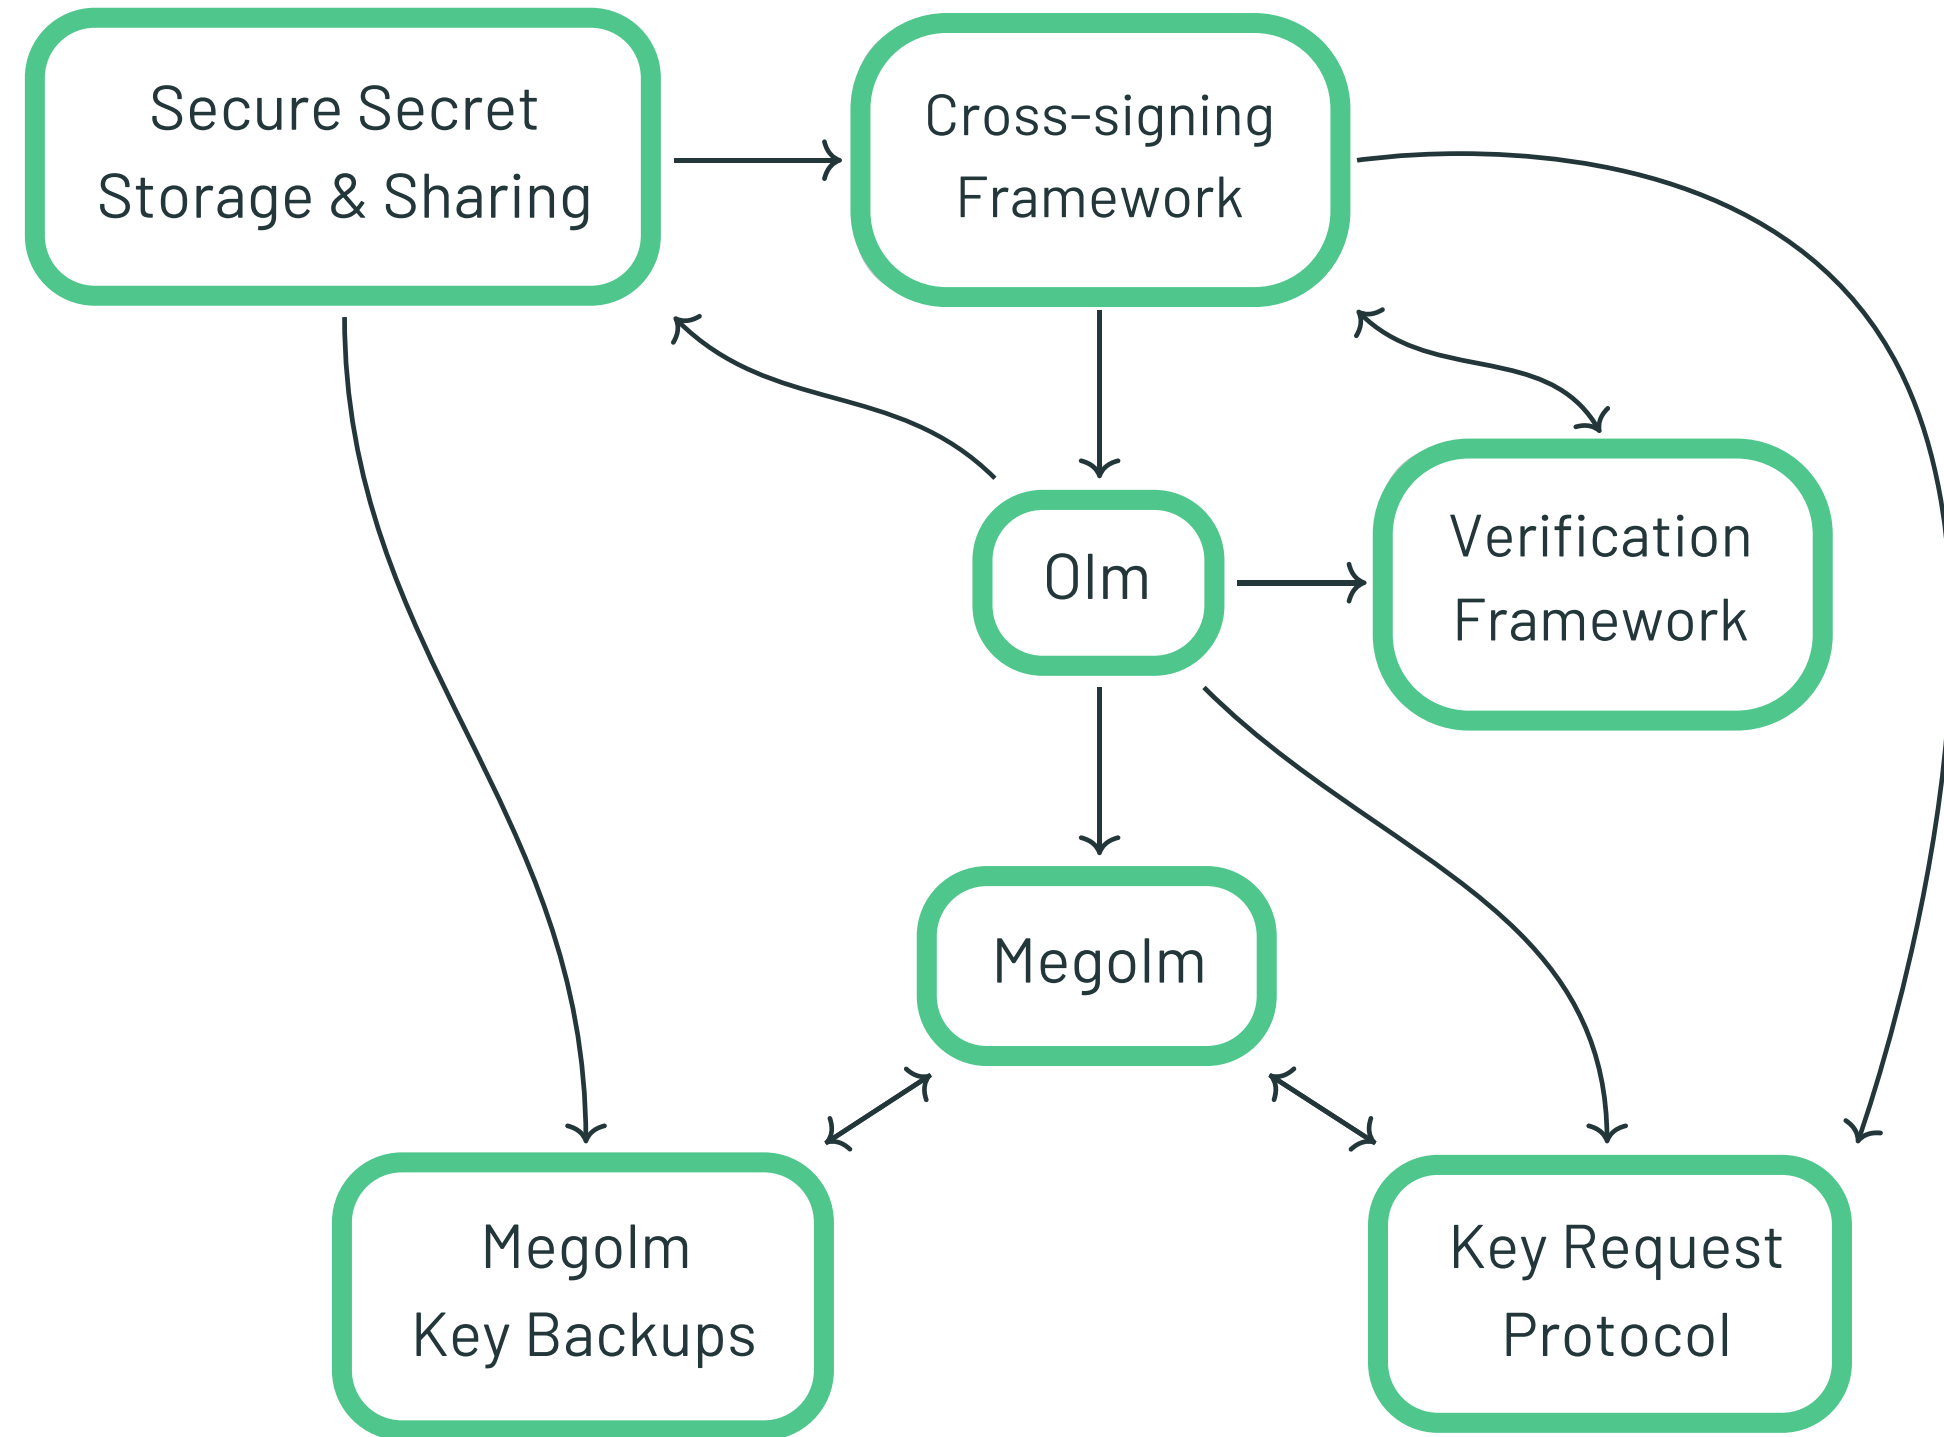
Can we pull off the same trick off again? **Yes**

Confidentiality break

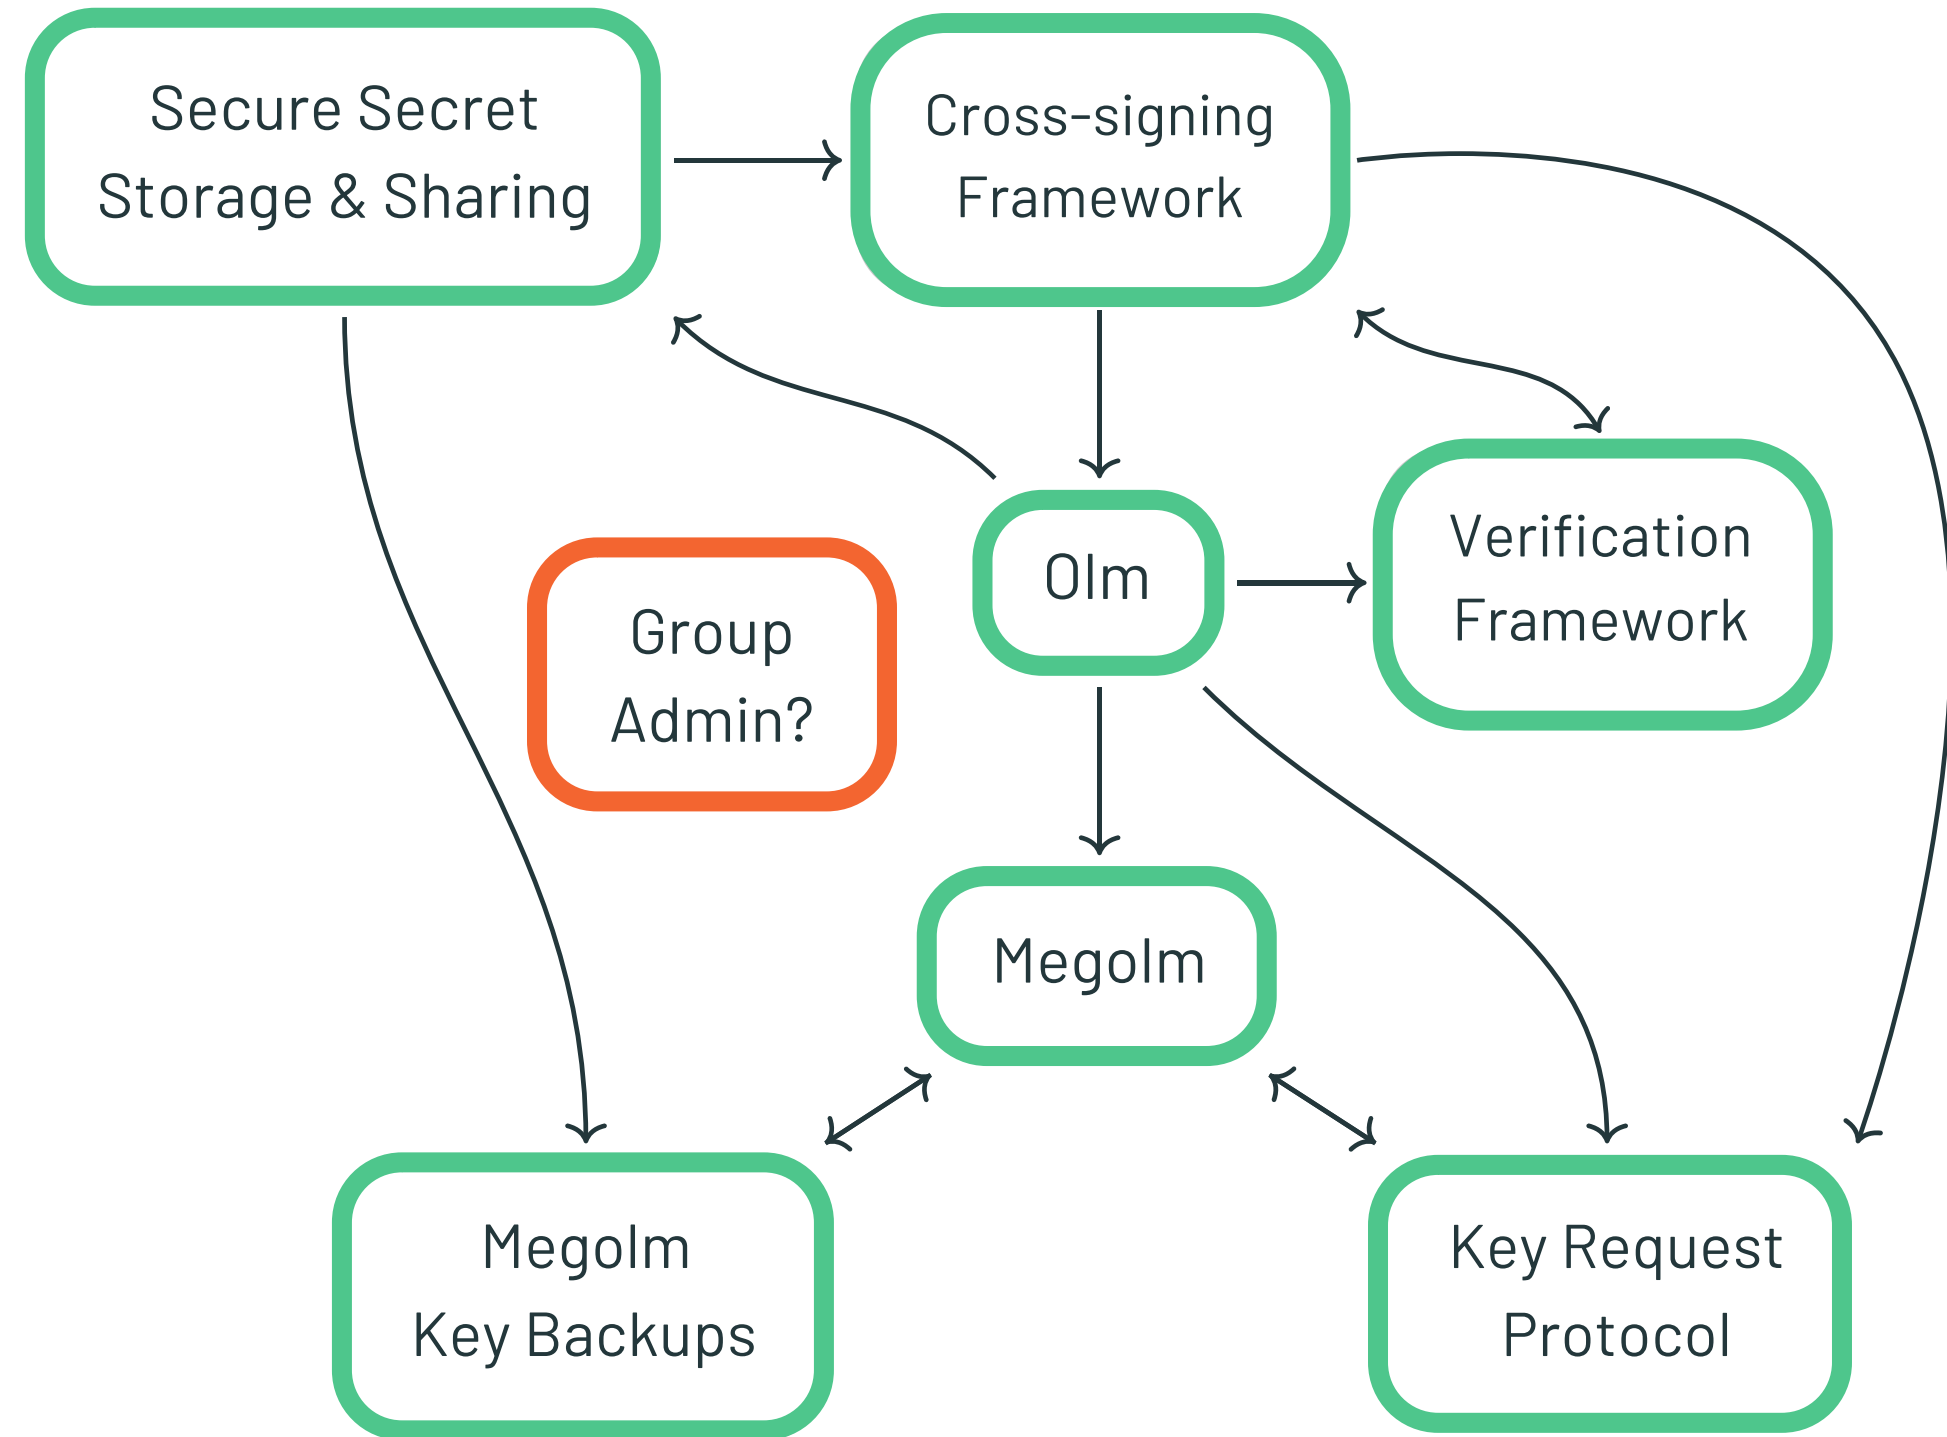


(A,1) will now back-up all Megolm decryption keys/inbound sessions

Modelling Matrix & Finding Attacks

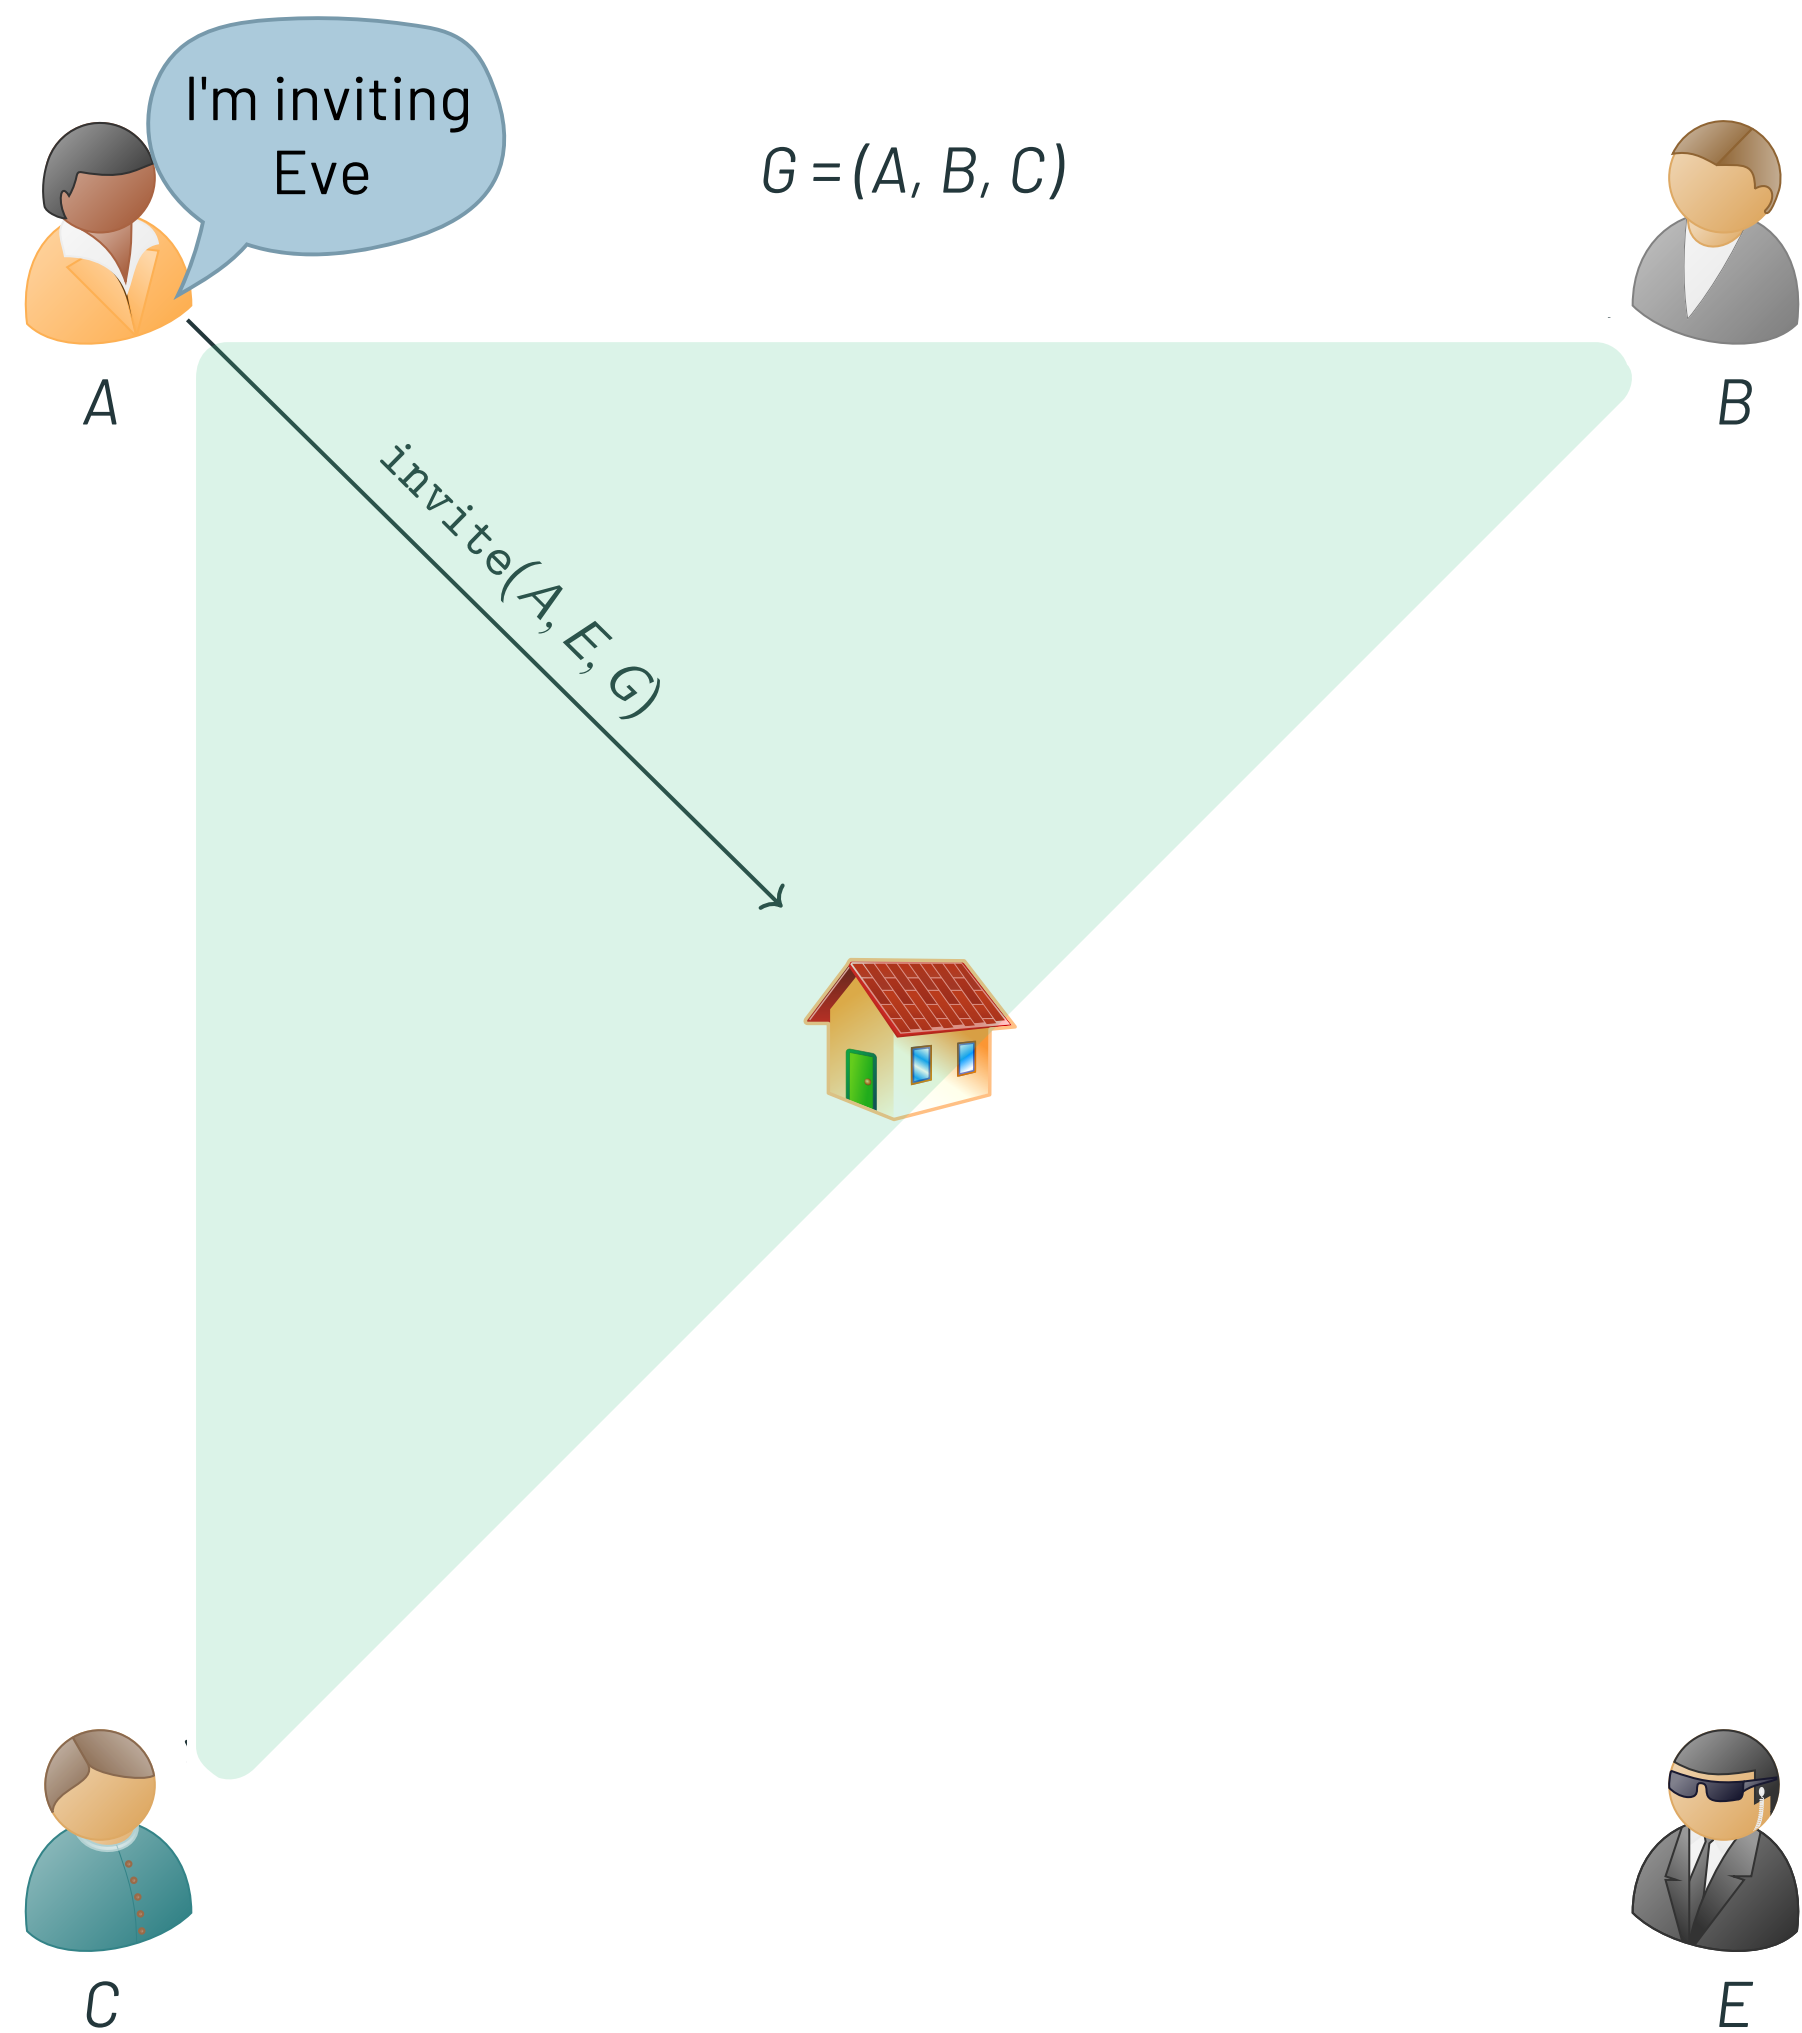


Group Administration




Group Administration

- Group membership is managed through events.



Group Administration

- Group membership is managed through events.
- Like messages in the room.




Group Administration

- Group membership is managed through events.
- Like messages in the room.





A

$G = (A, B, C, E)$



B



C



E

Group Administration

- Group membership is managed through events.
- Like messages in the room.
- These events are *unauthenticated*.



A

$G = (A, B, C, E)$



B



C



E

Server Control of Group Membership

Aim: add a server-controlled user to the group attack.

Can the server forge group invites?



A

$G = (A, B, C)$



B

ATTACK



C

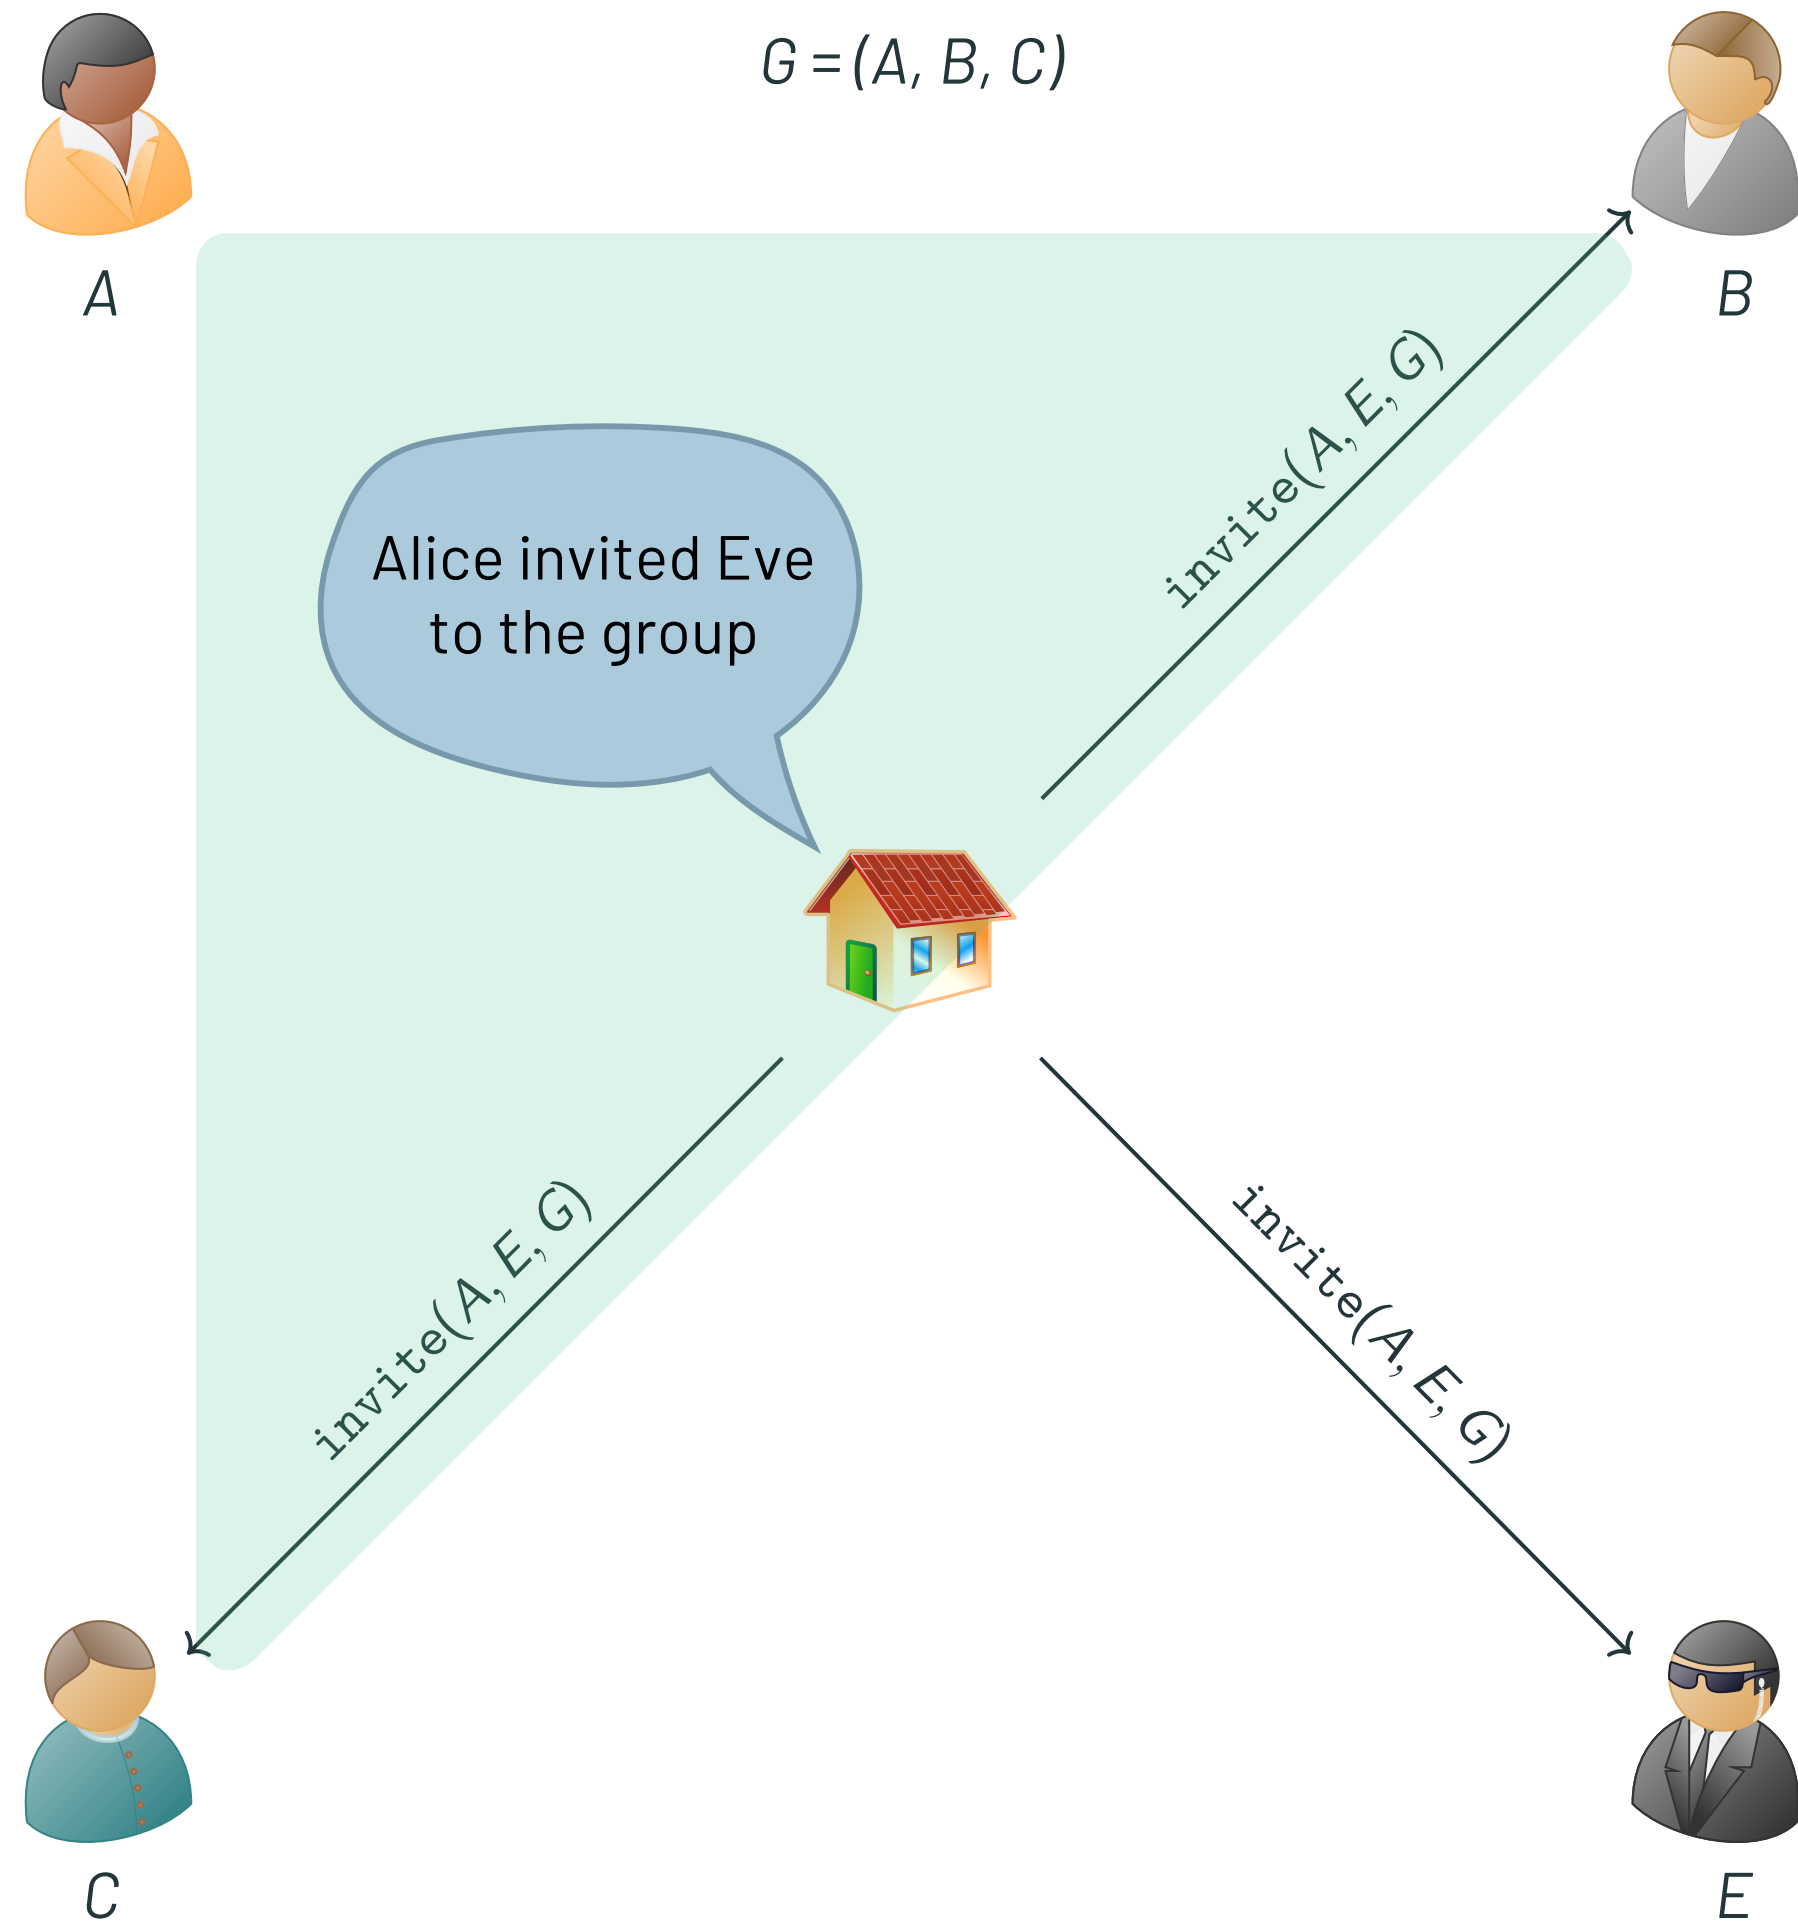


E

Server Control of Group Membership

Aim: add a server-controlled user to the group attack.

Can the server forge group invites?



ATTACK

Server Control of Group Membership

Aim: add a server-controlled user to the group attack.

Can the server forge group invites? **Yes**

↪ Confidentiality break *

* detectable in the interface

* for future messages only



A

$G = (A, B, C, E)$



B



C



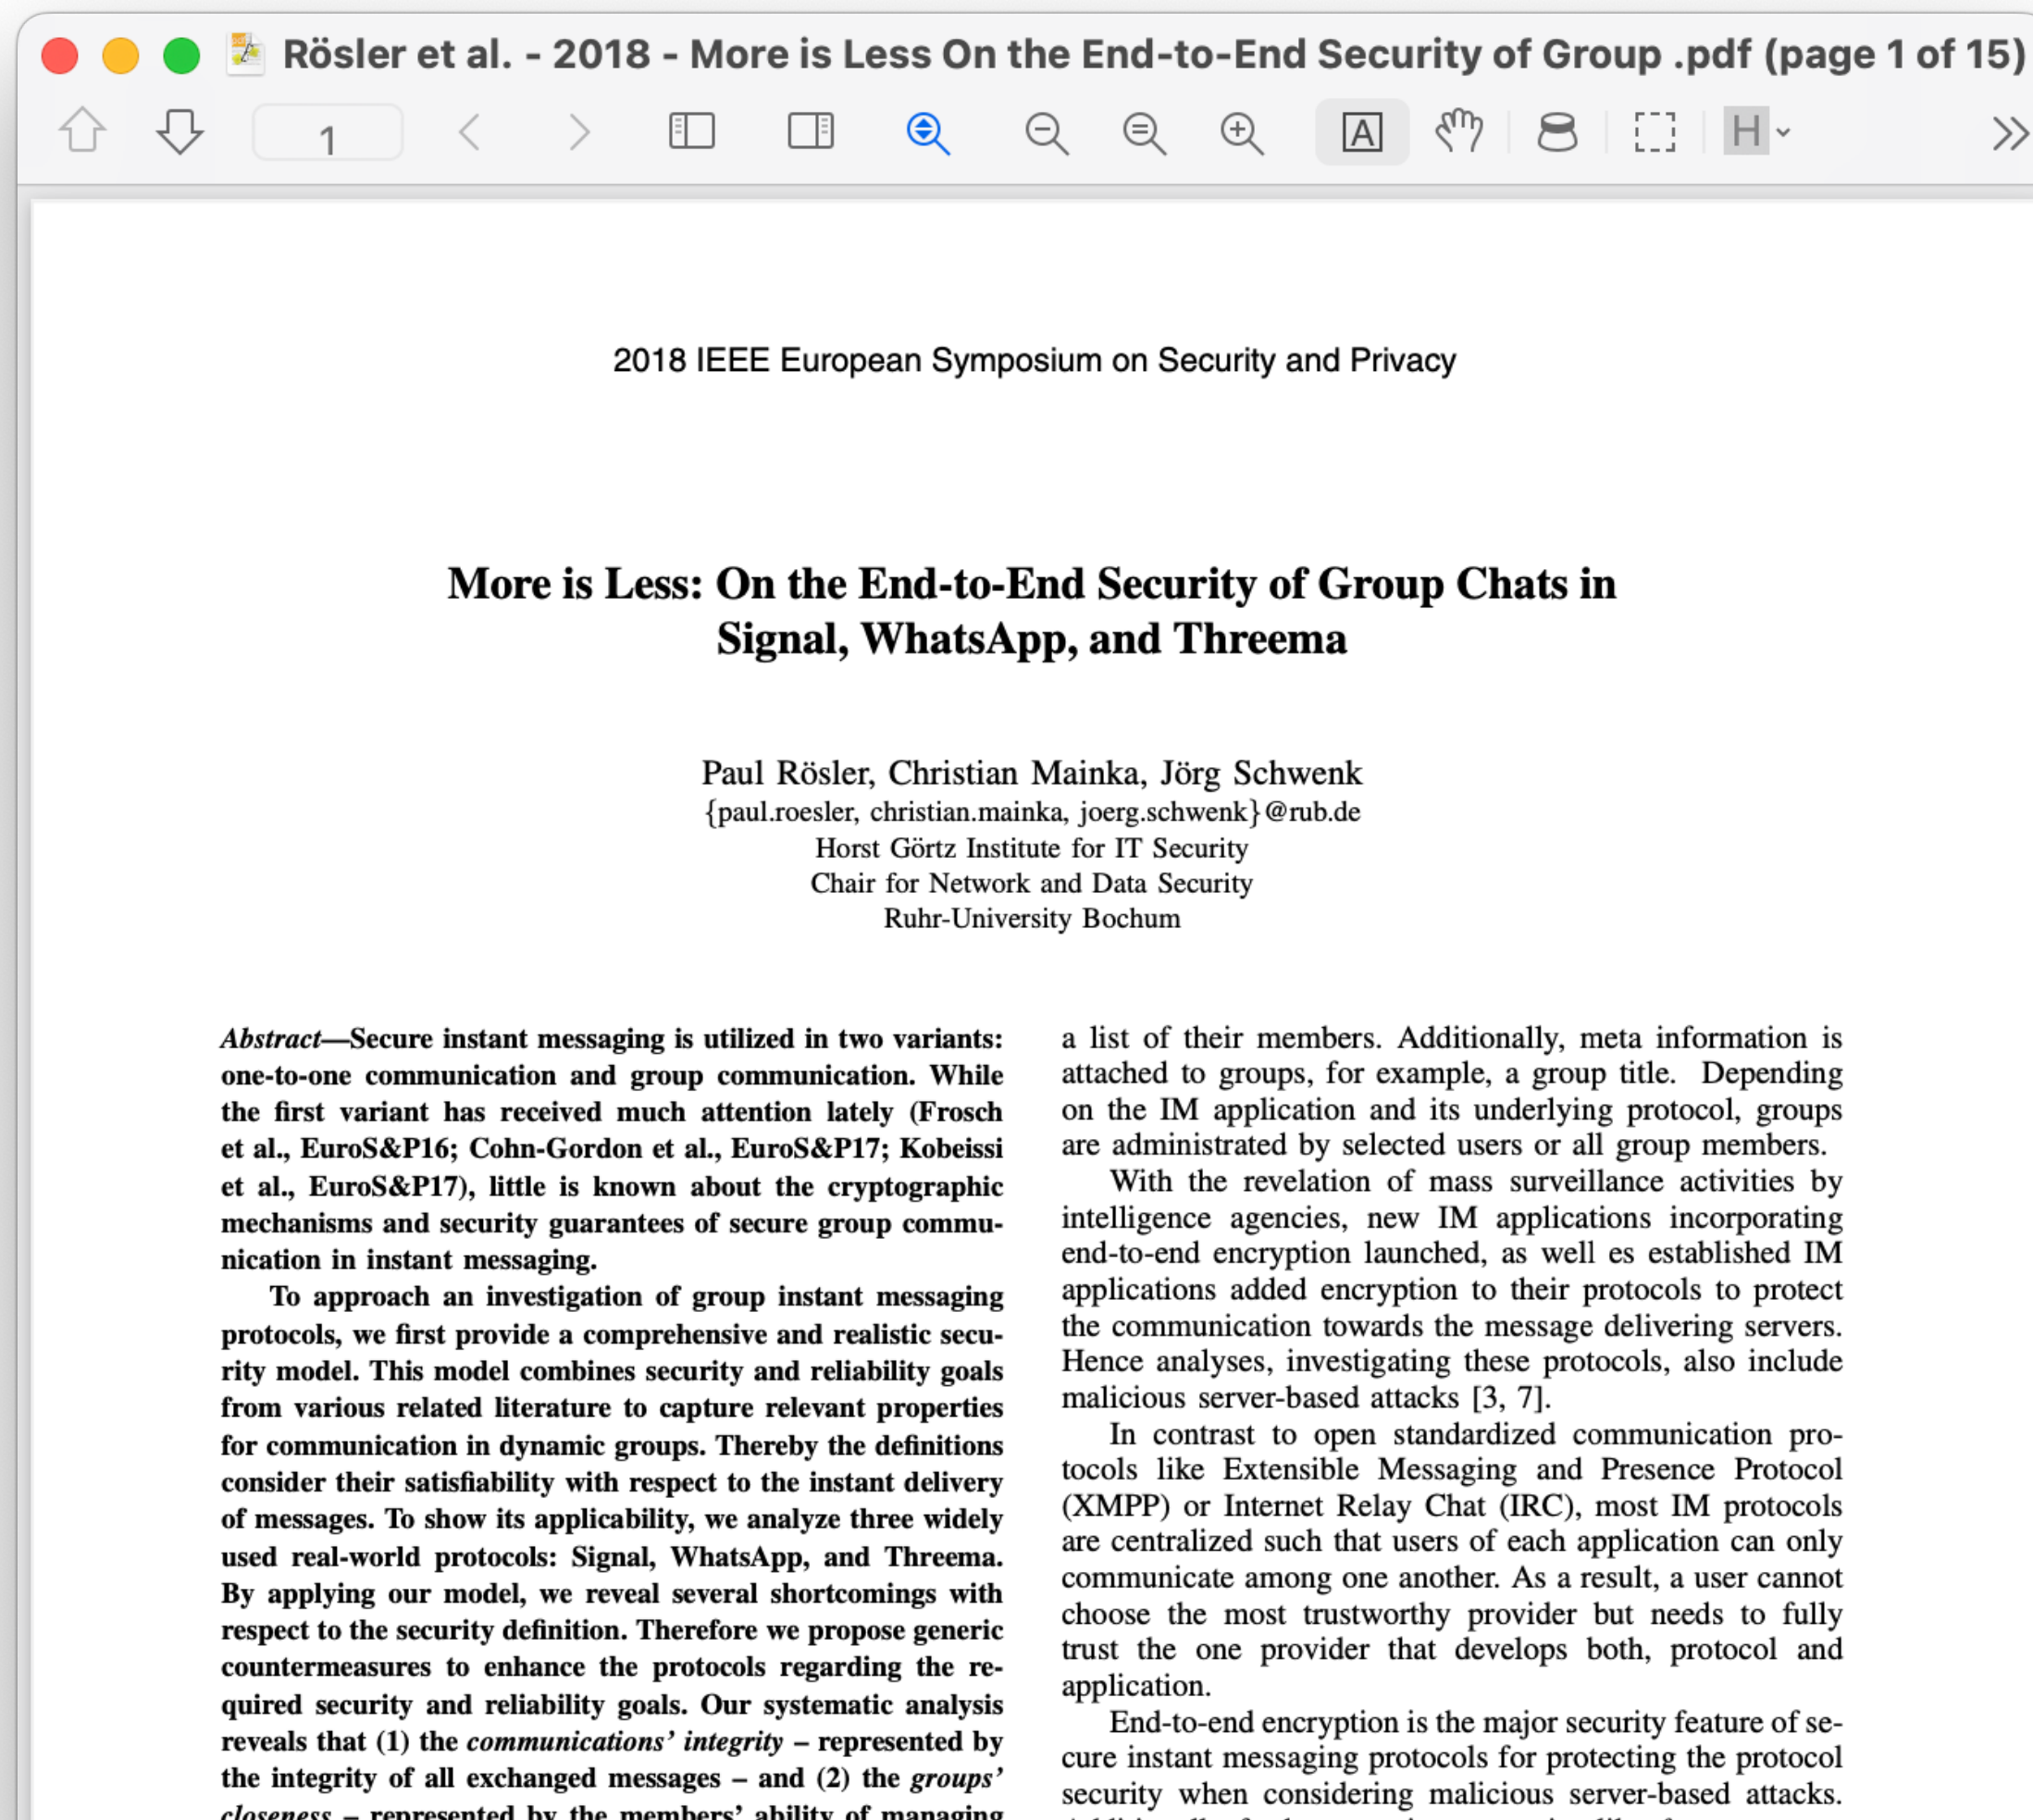
E

ATTACK

Server Control of Group Membership

Common issue among real-world group messaging protocols.

More is Less: both Signal and WhatsApp were vulnerable to a *burgle-into-the-group* attack.



Server Control of Group Membership

Common issue among real-world group messaging protocols.



The screenshot shows a web browser displaying the GitHub repository for `matrix-org/matrix-spec-proposals`. The page is for a specific proposal, `3917-cryptographic-membership.md`, which was recently updated by user `duxovni`. The commit message is "Fix room ID in example creation event". The page title is "MSC3917: Cryptographically Constrained Room Membership". The main content of the proposal is visible, starting with a paragraph: "In the current Matrix protocol, room membership events are not cryptographically signed, except by homeservers during federation. This means that a malicious homeserver can easily insert additional members into an end-to-end encrypted room. The falsified members will not receive keys for past messages, since those are only shared by existing members when they invite new members, but the falsified members will still be provided with keys for all new messages. Although the new member joining the room will be visible to all of the existing members, making it more difficult to perform such an attack undetected, it would still be preferable to have a means for clients to independently verify that a member actually belongs in a room." A second paragraph follows: "This proposal provides a method for clients to sign room membership events such that the room memberships form a tree of signatures rooted in the creation of the room, ensuring that every member belongs to a chain of invitations that ultimately leads back to the room's creator. This establishes a cryptographically verifiable bounding set of possible members of a room, significantly raising the barrier for homeservers to inject unauthorized members into the room."

Attacks Summary

1

COMPLETE BREAK
User/Device Confusion in
Out-of-Band Verification

2

AUTHENTICATION BREAK *
Impersonation
through Key Sharing

6

(THEORETICAL)
**IND-CCA
BREAK**

3

CONFIDENTIALITY BREAK
Adversary Controlled
Megolm Backup Key

4

AUTHENTICATION BREAK
Protocol Confusion

5

TRIVIAL CONFIDENTIALITY BREAK *
Server Control of
Group Membership

* Detectable in the user interface.

Attacks Summary

1

COMPLETE BREAK
User/Device Confusion in
Out-of-Band Verification

2

AUTHENTICATION BREAK*
Impersonation
through Key Sharing

6

(THEORETICAL)
IND-CCA
BREAK

3

CONFIDENTIALITY BREAK
Adversary Controlled
Megolm Backup Key

4

AUTHENTICATION BREAK
Protocol Confusion

5

TRIVIAL CONFIDENTIALITY BREAK*
Server Control of
Group Membership

* Detectable in the user interface.

Attacks Summary


1

COMPLETE BREAK
User/Device Confusion in
Out-of-Band Verification



2

AUTHENTICATION BREAK*
Impersonation
through Key Sharing



6
(THEORETICAL)
IND-CCA
BREAK



3

CONFIDENTIALITY BREAK
Adversary Controlled
Megolm Backup Key



4

AUTHENTICATION BREAK
Protocol Confusion



5

TRIVIAL CONFIDENTIALITY BREAK*
Server Control of
Group Membership

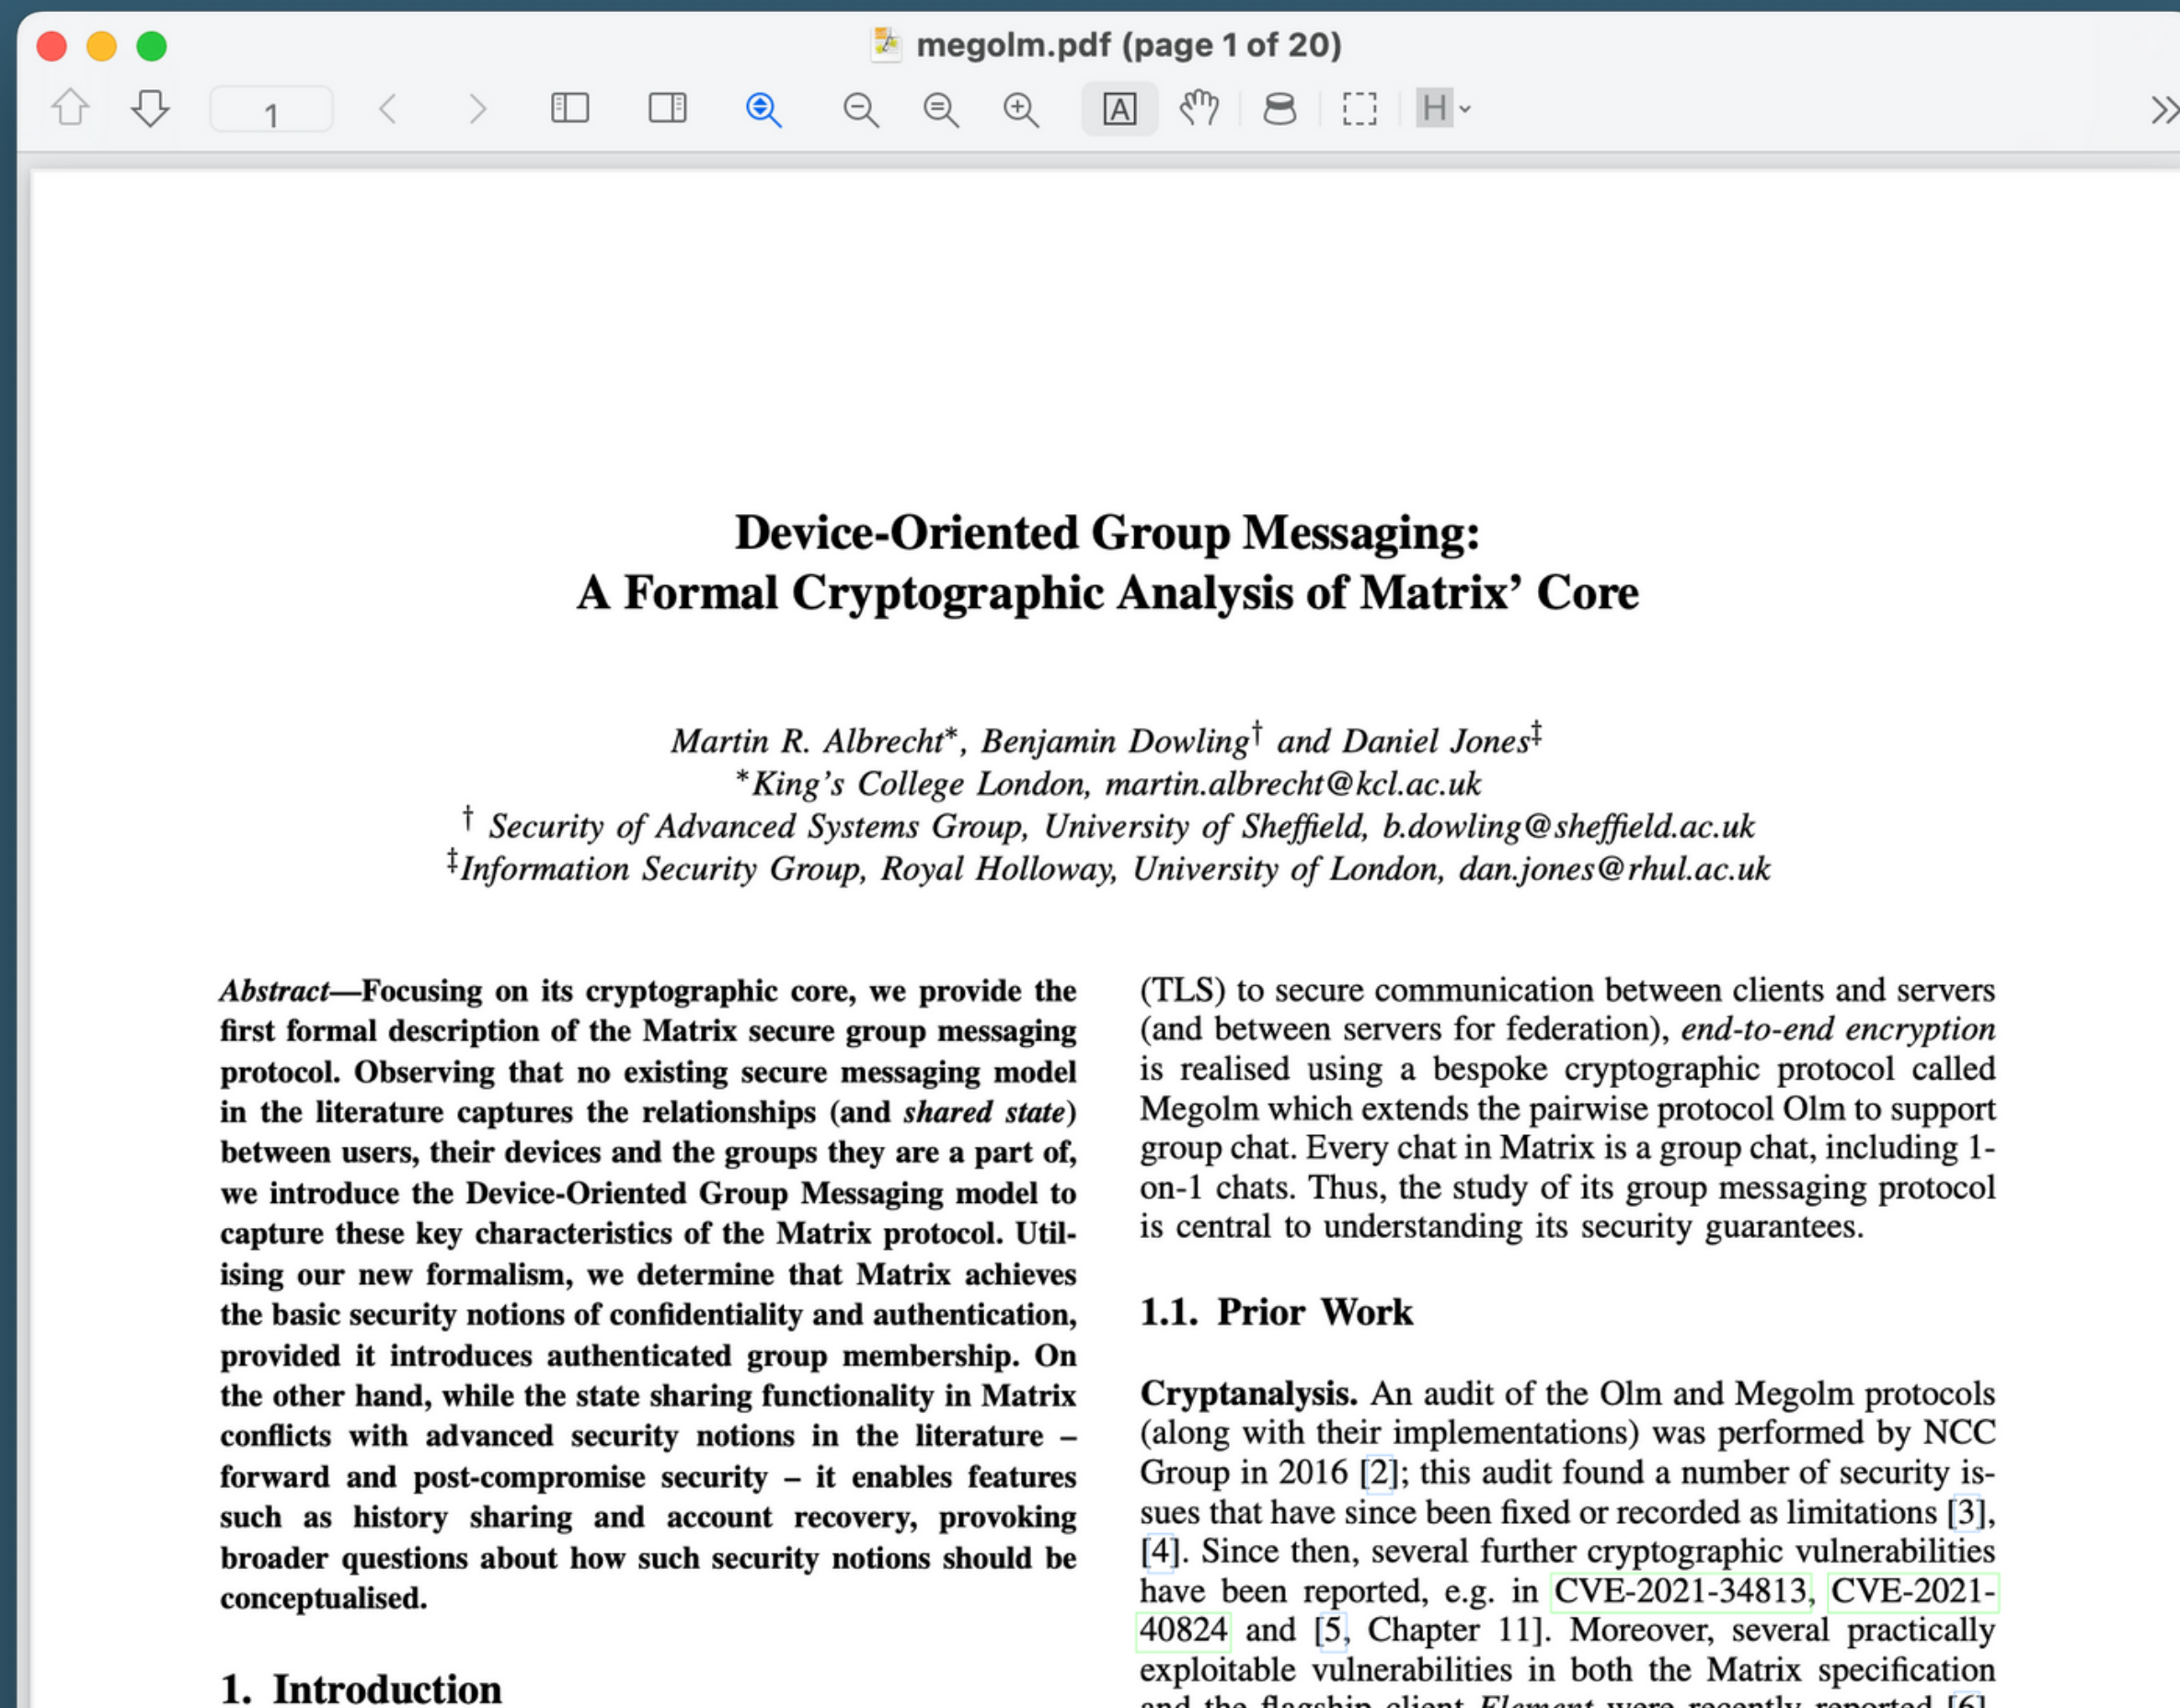


* Detectable in the user interface.

Modelling Matrix

Completed formalisation of Matrix' cryptographic core.

Security analysis and proof focuses on subset aiming to capture how the security of messages is affected by device-to-device interactions and state sharing.



Interested?

Find our paper at
<https://ia.cr/2023/485>.


Keep an eye out for the
follow-up modelling paper!

Image Credits

- Home from Nuvola Icon set by David Vignoni (LGPL)



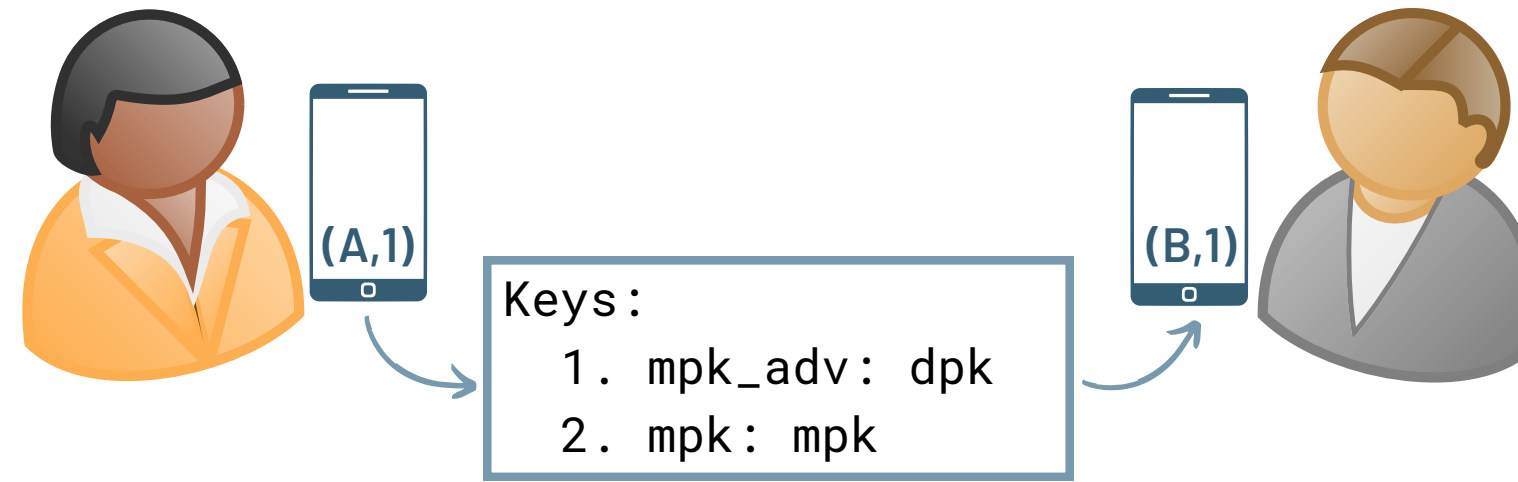
- tikzpeople package by Nils Fleischhacker



- Matrix screenshots from Matrix.org

- Element screenshots from element.io

SAS Attack



```

1. SAS.SendMAC(A, DA,i, B, DB,j, mpk, dpk, k, t)
-----
c ← "MATRIX_KEY_VERIFICATION_MAC" || A || DA,i || B || DB,j || t
iddev ← "ed25519:" || DA,i
macdev ← SAS.CalcMAC(k, dpk, c || iddev)
idcs ← "ed25519:" || mpk
maccs ← SAS.CalcMAC(k, mpk, c || idcs)
ms ← ((iddev, macdev), (idcs, maccs))
ks ← SAS.CalcMAC(k, sort(iddev, idcs), c || "KEY_IDS")
return (ms, ks)

```

```

2. SAS.VerifyMAC(A, DA,i, B, DB,j, mac, k, t)
-----
(iddev, macdev), (idcs, maccs), ks ← mac
c ← "MATRIX_KEY_VERIFICATION_MAC" || B || DB,j || A || DA,i || t
ks' ← SAS.CalcMAC(k, sort(iddev, idcs), c || "KEY_IDS")
assert ks' = ks
v ← ∅
for (id, mac) in ((iddev, macdev), (idcs, maccs))
    "ed25519:" || DB,j ← id
    // Check if this is a device verification request
    dpk ← HS.QueryKey("dpk", B, DB,j)
    if dpk ≠ ⊥ then
        D ← x
        if mac = SAS.CalcMAC(k, dpk, c || id) then
            v ← v ∪ {(B, D)}
    // Check if this is a cross-signing verification request
    elseif (x = HS.QueryKey("mpk", B)
            ∩ mac = SAS.CalcMAC(k, x, c || id)) then
        mpk ← x
        v ← v ∪ {(B, mpk)}
return v

```

```

3. SAS.SignDevice(A, DA,i)
-----
// Check whether DA,i is a cross-signing identity
mpk ← HS.QueryKey("mpk", A)
if DA,i = mpk then
    return UserVerified(A, mpk)
// Otherwise, DA,i refers to a device
else
    return DeviceVerified(A, DA,i)

```