



Auditing Key Transparency

Advancing WhatsApp tampering protection for millions

Mari Galicer, Thibault Meunier
Cloudflare

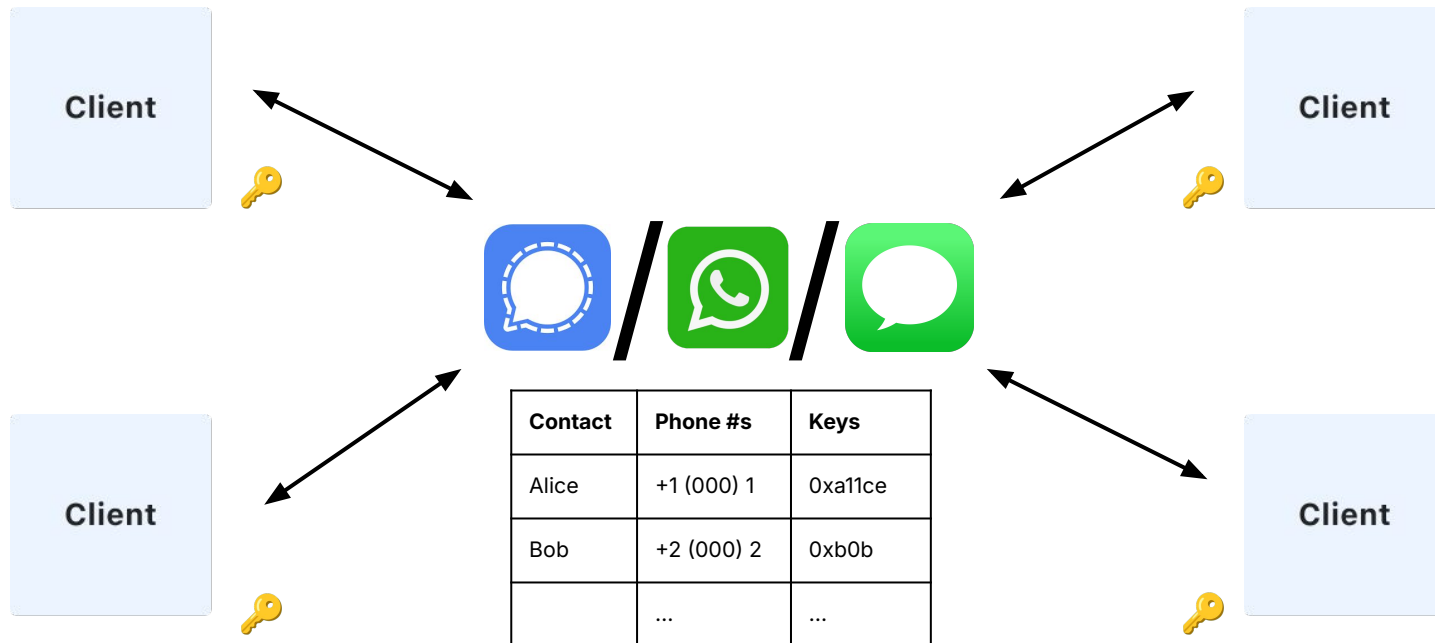
Kevin Lewi
WhatsApp / Meta

Agenda

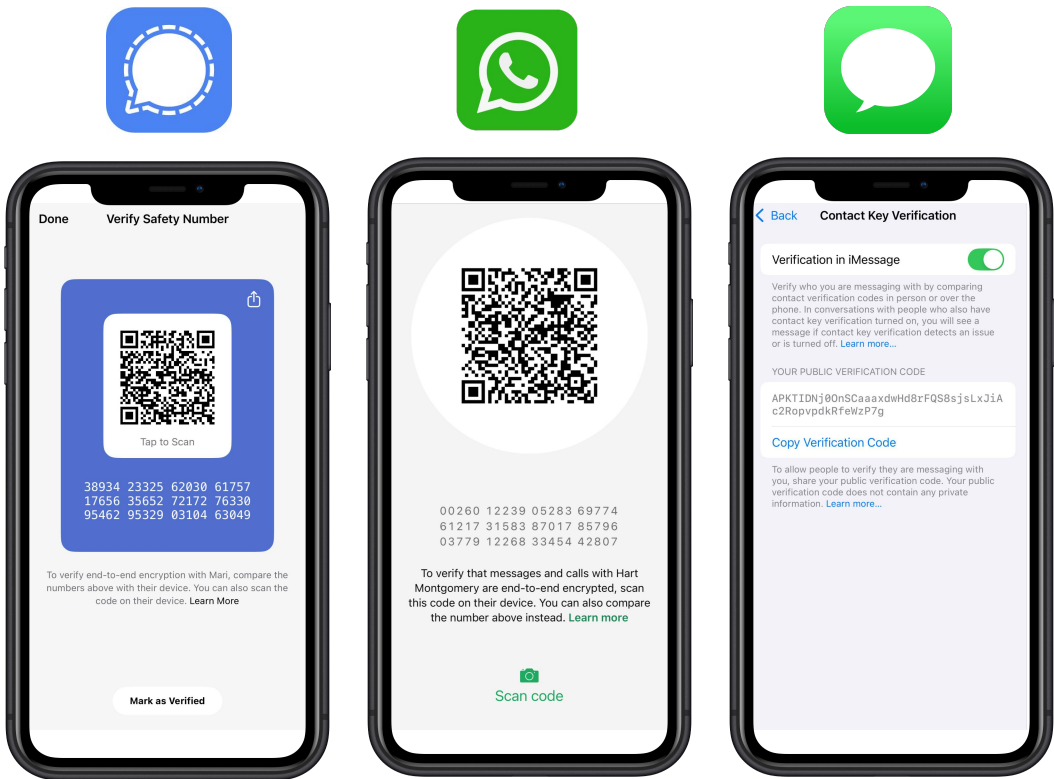
- 1 A refresher on Key Transparency
- 2 Auditor system design
- 3 Real-world deployment
- 4 What's next

A refresher on Key Transparency

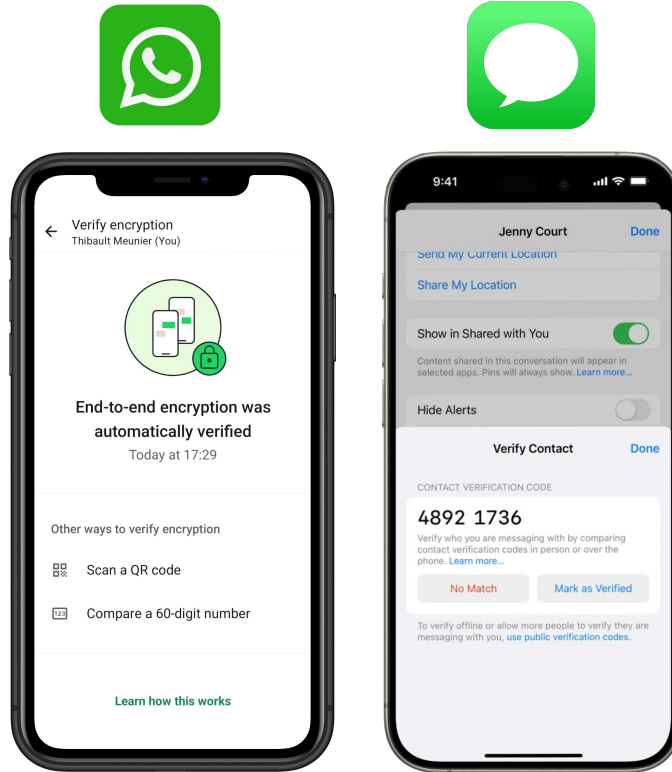
E2EE Messaging




QR Codes



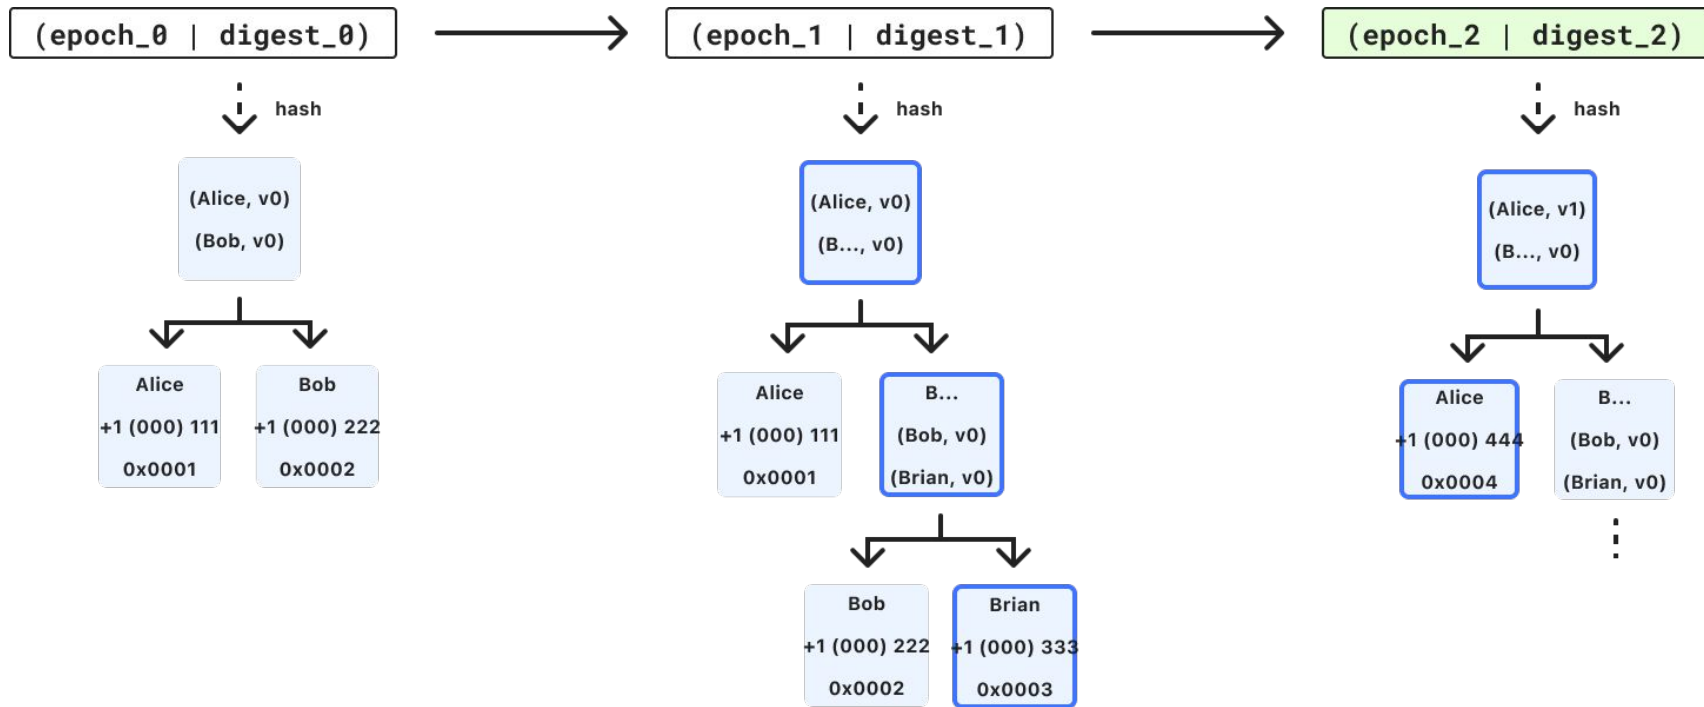
Key Transparency



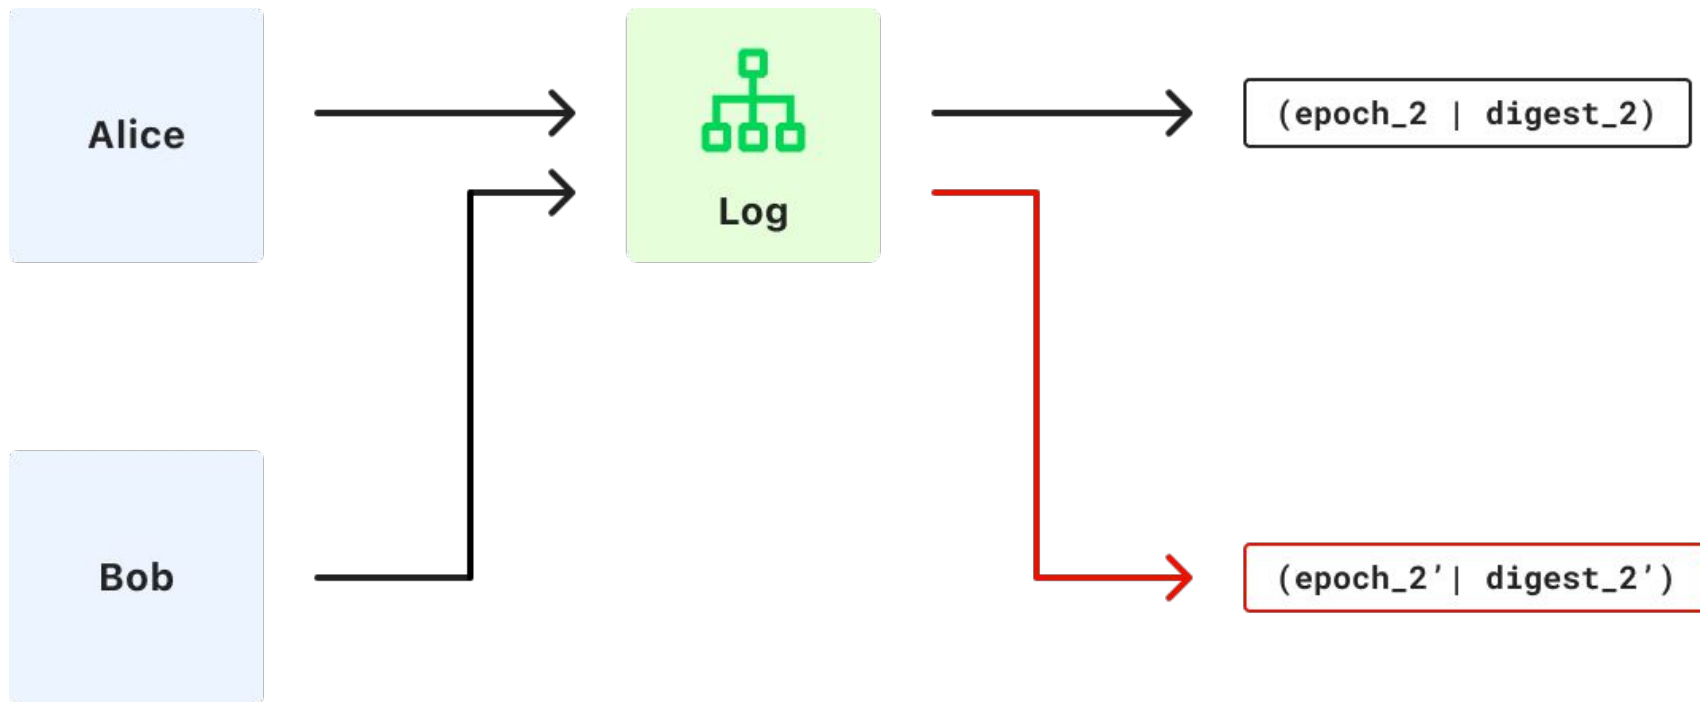
Key Transparency – Client and AKD



Key Transparency – Auditable Key Directory (AKD)

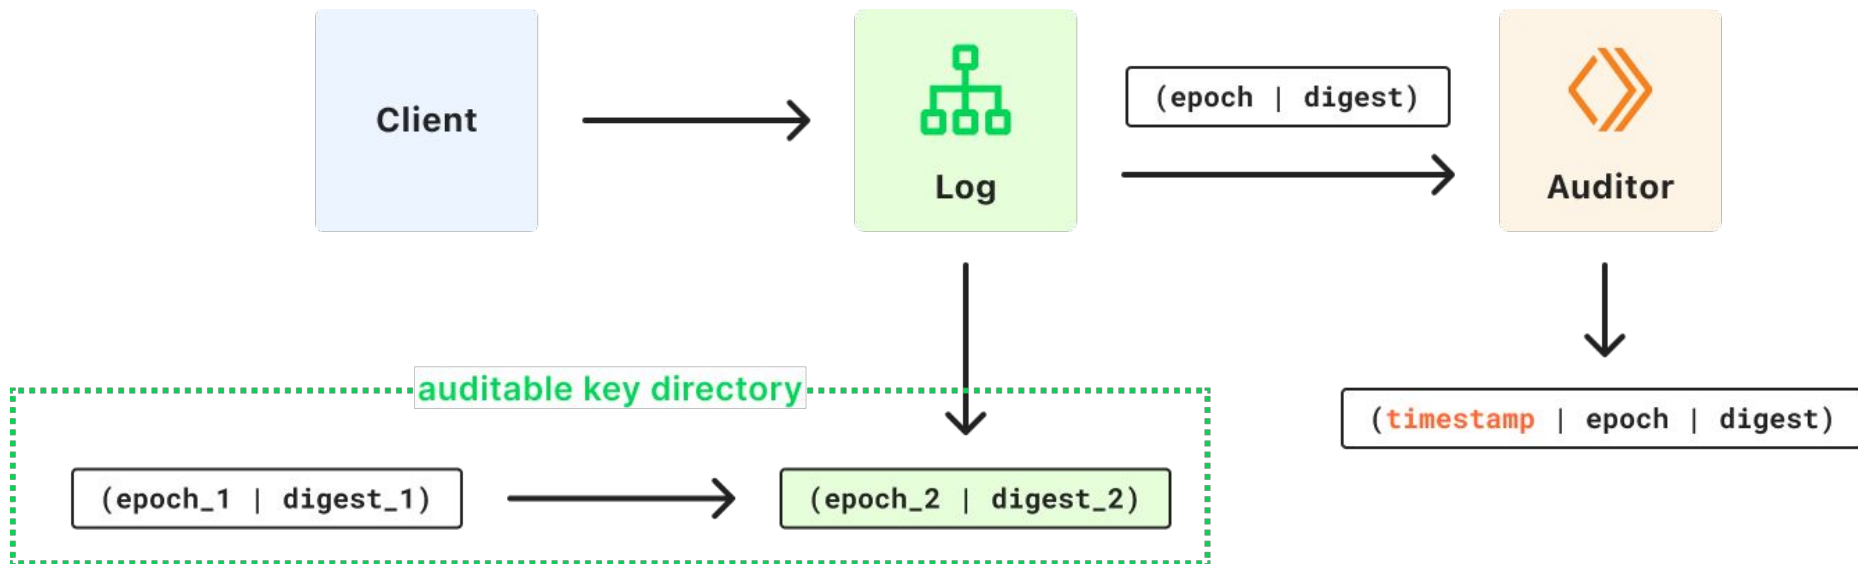


Key Transparency – split view attack



Auditor system design

Introducing the Auditor



What is the Auditor responsible for?



Witnessing

Ensures epochs are unique and in sequential order.



Monitoring

Making sure that the AKD is correctly constructed and that all epochs transitions are valids.




Privacy preserving


Does not see users' private information: no name, no phone number, no public key. It is a trusted third party.

How does that relate to Certificate Transparency


Certificate Transparency



Key Transparency



Validating epoch uniqueness and transitions



Real world deployment

A command line interface

```
> # Hello RWC!  
> # Let's do a demo
```

Recorded with
[charmbracelet/vhs](https://github.com/charmbracelet/vhs)

Code on
[cloudflare/plexi](https://github.com/cloudflare/plexi)

A command line interface

```
> # Install Plexi
> cargo install plexi
    Updating crates.io index
    Ignored package `plexi v0.1.2` is already installed, use --force
to override
> █
```

Recorded with
[charmbracelet/vhs](https://github.com/charmbracelet/vhs)

Code on
[cloudflare/plexi](https://github.com/cloudflare/plexi)

A command line interface

```
> # List audited logs
> plexi ls --remote-url 'https://plexi.key-transparency.cloudflare.com'
,
test.11092024
test.nl.cloudflare.plexi.example.com
test.whatsapp.key-transparency.v1
whatsapp.key-transparency.v1
>
```

Recorded with
[charmbracelet/vhs](https://github.com/charmbracelet/vhs)

Code on
[cloudflare/plexi](https://github.com/cloudflare/plexi)

A command line interface

```
> # Audit the latest epoch
> plexi audit \
  --remote-url 'https://akd-auditor.cloudflare.com' \
  --namespace 'whatsapp.key-transparency.v1' \
  --long
Audit proof verification enabled. It can take a few seconds
.....
```

Recorded with
[charmbracelet/vhs](https://github.com/charmbracelet/vhs)

Code on
[cloudflare/plexi](https://github.com/cloudflare/plexi)

A command line interface

```
Audit proof verification enabled. It can take a few seconds
```

```
.....
```

```
Namespace
```

```
  Name          : whatsapp.key-transparency.v1
```

```
  Ciphersuite    : ed25519(protobuf)
```

```
Signature (2025-03-21T16:57:41Z)
```

```
  Epoch height   : 1001282
```

```
  Epoch digest   : 5ebc1ef0b528acab3f6aa47fa7b728f8318dd751c87e
```

```
3eb18939546805e07475
```

```
  Signature      : 54595ddc6c20c04e2183cc6001268f692f6c52fc6e8e
```

```
acb05b5db08b142ae390f1d7c92c8bb55cebfb42097afda8e2037d1e8da78737dc9ea
```

```
968da13d6e5903
```

```
  Signature verification: success
```

```
  Proof verification   : success
```

```
> █
```

Recorded with
[charmbracelet/vhs](https://charmbracelet.com/vhs/)

Code on
[cloudflare/plexi](https://cloudflare.com/plexi)

Real world deployment

... in reality

Incidents

Privacy Eng Pages App Fri 11:41 PM




Incident [1466211](#) triggered by [Privacy Eng](#): 1 alert for Plexi Worker Production Epoch Not Increasing

38 replies Mon 11:04 AM




During an incident, epochs fail to progress

Latest Epoch Over Time



When epochs get bigger, verification latency increases

Verification Latency



More incidents



Publishes every 30s

Theory

The Log publishes new heads.

Practice

When that failed, the backlog of updates grew, increasing the proof size 6x, going beyond our initial provisioning threshold.



Global ordering

Theory

Global ordering scales easily.

Practice

In practice, the auditor signs a timestamp. This means that it's hard to replay signatures.



Integrity

Theory

No party gets corrupted.

Practice

This is true! So far, both the Log and the Auditor managed to remain in a non-corrupted state, despite hiccups.

What's next

Transparency – a timeline

Certificate Transparency

Google [launches](#) their first Certificate Transparency Log.

[RFC 6962](#) published at the IETF.

Key transparency paper

[CONIKS](#) introduces Key Transparency.

One more paper

[SEEMLess](#) formalises some of CONIKS designs and improves performance with a new data structure.

Public auditing

Cloudflare releases [Plexi Auditor](#) in collaboration with WhatsApp, the talk you are listening to.

2013

2014

2015

2017

2019

2023

2024

2025+

Keybase

End-to-end encrypted messaging app relying on [signature chains](#).

Google Key Transparency

[Inspired](#) by CONIKS and Certificate Transparency.

Key Transparency Logs

[Parakeet](#) paper makes SEEMLess practical at scale.
[iMessage](#), [Proton](#), [WhatsApp](#) launch their Log.
IETF forms the [keytrans](#) working group.

More?

RFC, auditing network, adoption

Why there aren't more auditors... yet!



Technical expertise

Key Transparency is new

There needs to be more expertise and understanding about the guarantees it provides with and without auditing.



Implementations

There are a lot

CONIKS, CT-based, AKD, tlog-based.
[IETF keytrans](#) is developing a standard.




Performance at scale

Real world is big

Scale makes the system more expensive to audit, or less performant than on a off-the-shelf device.

Where can I see / use it?



Last updated: 2025-03-20T12:34:12Z

Key Transparency aims to secure the distribution of public keys for end-to-end encrypted (E2EE) messaging systems, such as Whatsapp. It achieves this by building a verifiable append-only data structure called a Log, similar to [Certificate Transparency](#).

Cloudflare verifies Key Transparency Logs to ensure the transparency of end-to-end encrypted messaging public keys. This component is called an Auditor. Cloudflare provides an [API](#) for anyone to monitor the work of the Auditor, and verify the state of its associated Logs locally. This local validation can be done with [cloudflare/plexi](#) cli for instance.

Log status

Name	Status	Updated
WhatsApp	Online ●	2025-03-20T12:34:10Z

Log list

WhatsApp

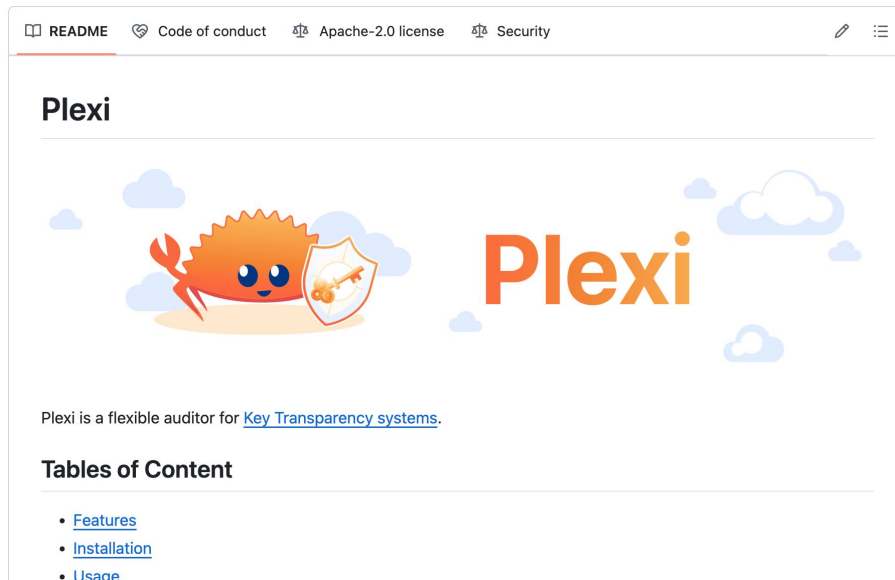
Status: Online ●

Last update: 2025-03-20T12:34:10Z


Latest epoch: [997875](#) [↗](#)

Root: [458298/3ae9497069cc722dc9e00f8251da87071646a57dae2fc7882f1d8214961d80bd](#) [↗](#)

Log name: whatsapp.key-transparency.v1




Towards a transparency ecosystem

Key Transparency Dashboard

GitHub

Log list

WhatsApp

Status: Online 

Last update: 2025-03-20T10:45:20Z

Latest epoch: [997875](#)


Root: [458298/3ae9497069cc722dc9e00f8251da87071646a57dae2fc7882f1d8214961d80bd](#)

Log name: whatsapp.key-transparency.v1



Your app here



Status: Online 

Last update: 2025-03-20T12:34:10Z

Latest epoch: [997875](#)

Root: [458298/3ae9497069cc722dc9e00f8251da87071646a57dae2fc7882f1d8214961d80bd](#)

Log name: my-e2ee-messenger.key-transparency.v1

Thank you

 blog.cloudflare.com

 engineering.fb.com



Backup slides

Key Transparency Tradeoffs



Active vs Passive

Active: Users report to a 3rd party/gossip

Passive: 3rd party signature



Sync vs Async

Sync: 3rd party is on-path for publishing

Async: Delayed detection



Human vs Automated

WhatsApp can prompt for confirmation

Automated systems have no direct interventions



Two signatures

The witness provides a sync signature.

The monitor does an async signature.

They have the same format but not the same public key.



Auditor endpoint

No endpoint: lightweight auditor

Endpoint: more accountability and trust



Frequency

The faster you publish epochs, the smaller the epochs, but the more availability you need.

Multiple papers and deployment

CONIKS: Bringing Key Transparency to End Users SEEM/less: Secure End-to-End Encrypted Messaging with *less* Trust

Marcela S. Melara and Aaron Blankstein, *Princeton University*; Joseph Bonneau, *Stanford University* and *The Electronic Frontier Foundation*; Edward W. Felten and Michael J. Freedman, *Princeton University*

<https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/melara>

Parakeet: Practical Key Transparency for End-to-End Encrypted Messaging

Harjasleen Malvai^{*†}, Lefteris Kokoris-Kogias^{‡§}, Alberto Sonnino^{‡¶}, Esha Ghosh^{||}, Ercan Oztürk^{**}, Kevin Lewi^{**}, and Sean Lawlor^{**}

^{*}UIUC, [†]IC3, [‡]Mysten Labs, [§]IST Austria, [¶]University College London (UCL), ^{||}Microsoft Research, ^{**}Meta

OPTIKS: An Optimized Key Transparency System

Julia Len, *Cornell Tech*; Melissa Chase, Esha Ghosh, Kim Laine, and Radames Cruz Moreno, *Microsoft Research*

<https://www.usenix.org/conference/usenixsecurity24/presentation/len>

ELEKTRA: Efficient Lightweight multi-dEvice Key TRAnsparency*

Julia Len[†]
Cornell Tech
New York, USA
jl原因@cs.cornell.edu

Melissa Chase
Microsoft Research
Redmond, USA
melissac@microsoft.com


Esha Ghosh
Microsoft Research
Redmond, USA
esha.ghosh@microsoft.com

Daniel Jost
New York University
New York, USA
daniel.jost@cs.nyu.edu

Balachandar Kesavan
Zoom Video Communications
New York, USA
surya.heronhay@zoom.us

Antonio Marcedone
Zoom Video Communications
New York, USA
antonio.marcedone@zoom.us

Gossip



The Log is corrupted, now what

1. Users update are stale
2. Restore from backups
3. Start from a last known good state, or from scratch
4. Discuss with the auditor, will need to provision a new namespace
5. Communicate about it
6. Cut a new release of WhatsApp