

Certificate Transparency with Privacy


Saba Eskandarian, Eran Messeri, Joe Bonneau, Dan Boneh
Stanford Google NYU Stanford

Certificate Authorities



Public Key

Certificate Authorities



apo-CA-lypse

An update on attempted man-in-the-middle attacks

August 29, 2011

**FINAL REPORT ON DIGINOTAR HACK SHOWS TOTAL
COMPROMISE OF CA SERVERS**



DigiNotar
Internet Trust Services



apo-CA-lypse

An update on attempted man-in-the-middle attacks

August 29, 2011

**FINAL REPORT ON DIGINOTAR HACK SHOWS TOTAL
COMPROMISE OF CA SERVERS**



DigiNotar
Internet Trust Services

Distrusting WoSign and StartCom Certificates


October 31, 2016

Outline

- **Certificate Transparency**
- Redaction of private subdomains
- Privacy-preserving proof of misbehavior

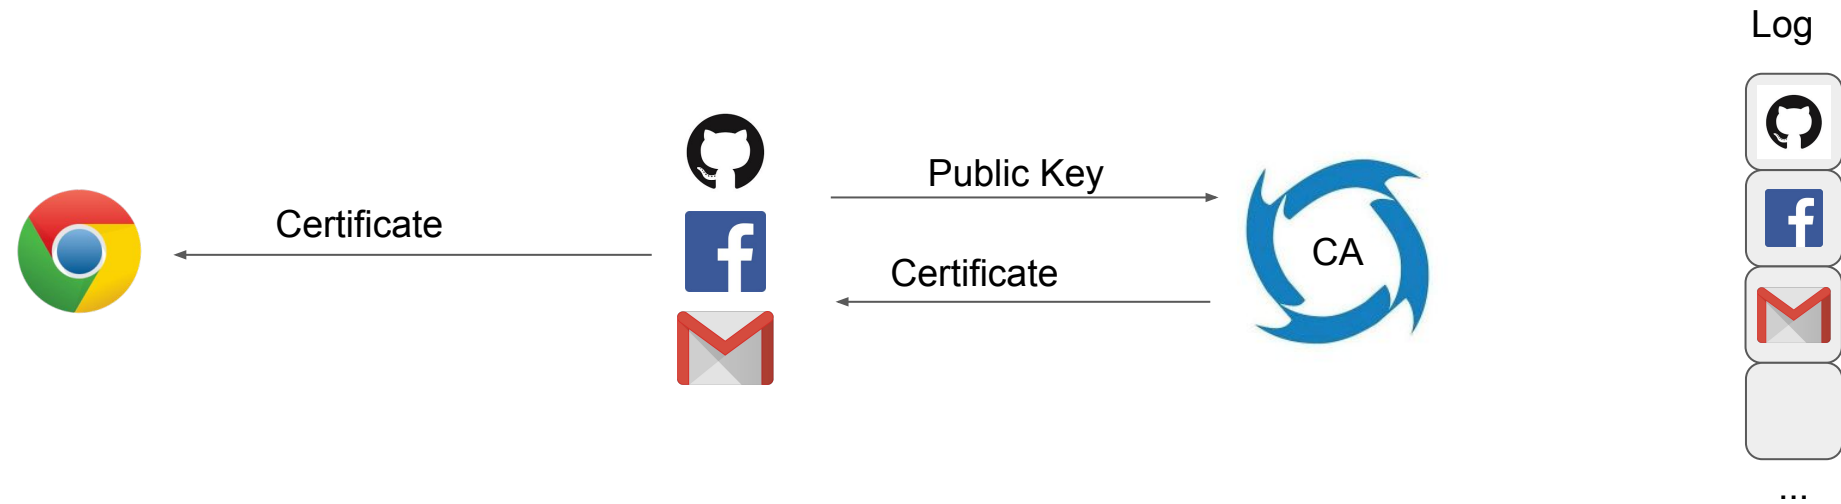
Certificate Transparency (CT)

Idea: public, verifiable log of all certificates



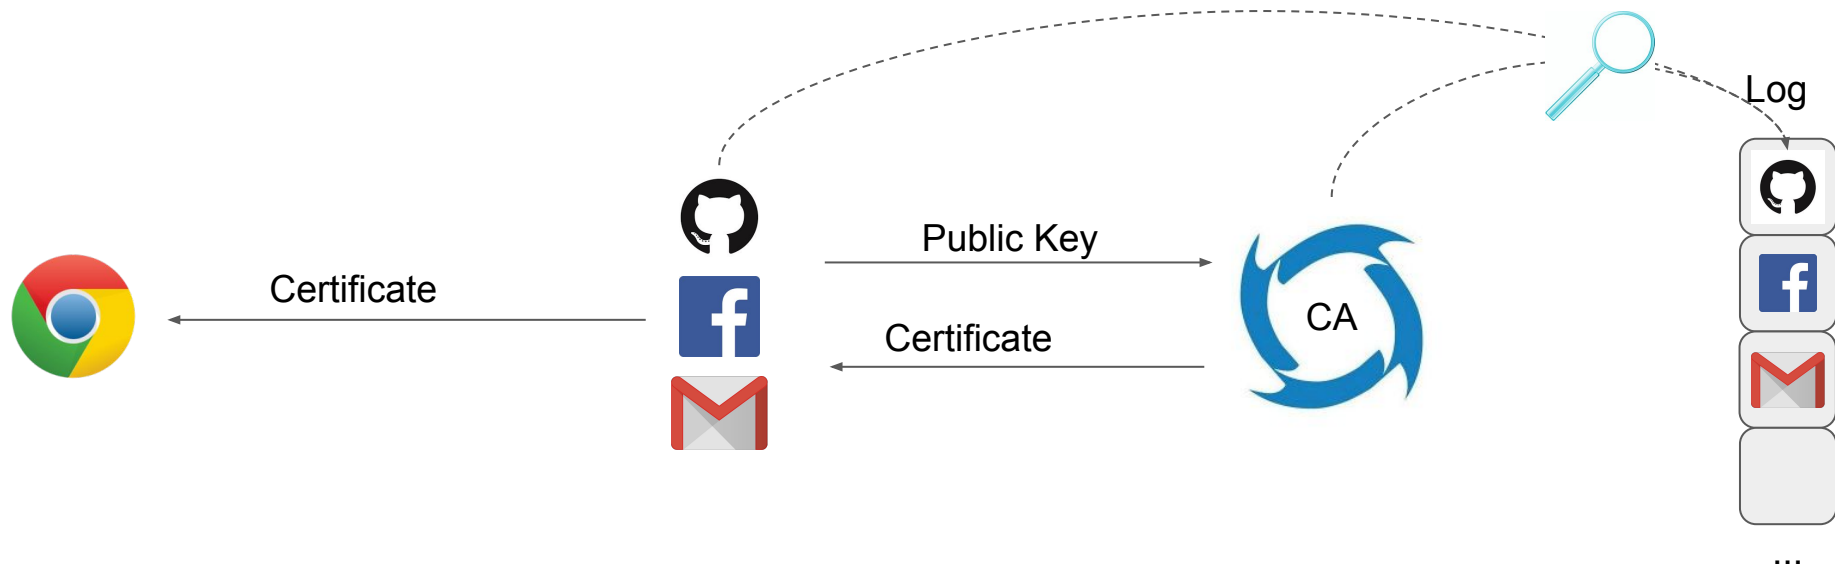
Certificate Transparency (CT)

Idea: public, verifiable log of all certificates




Certificate Transparency (CT)

Idea: public, verifiable log of all certificates



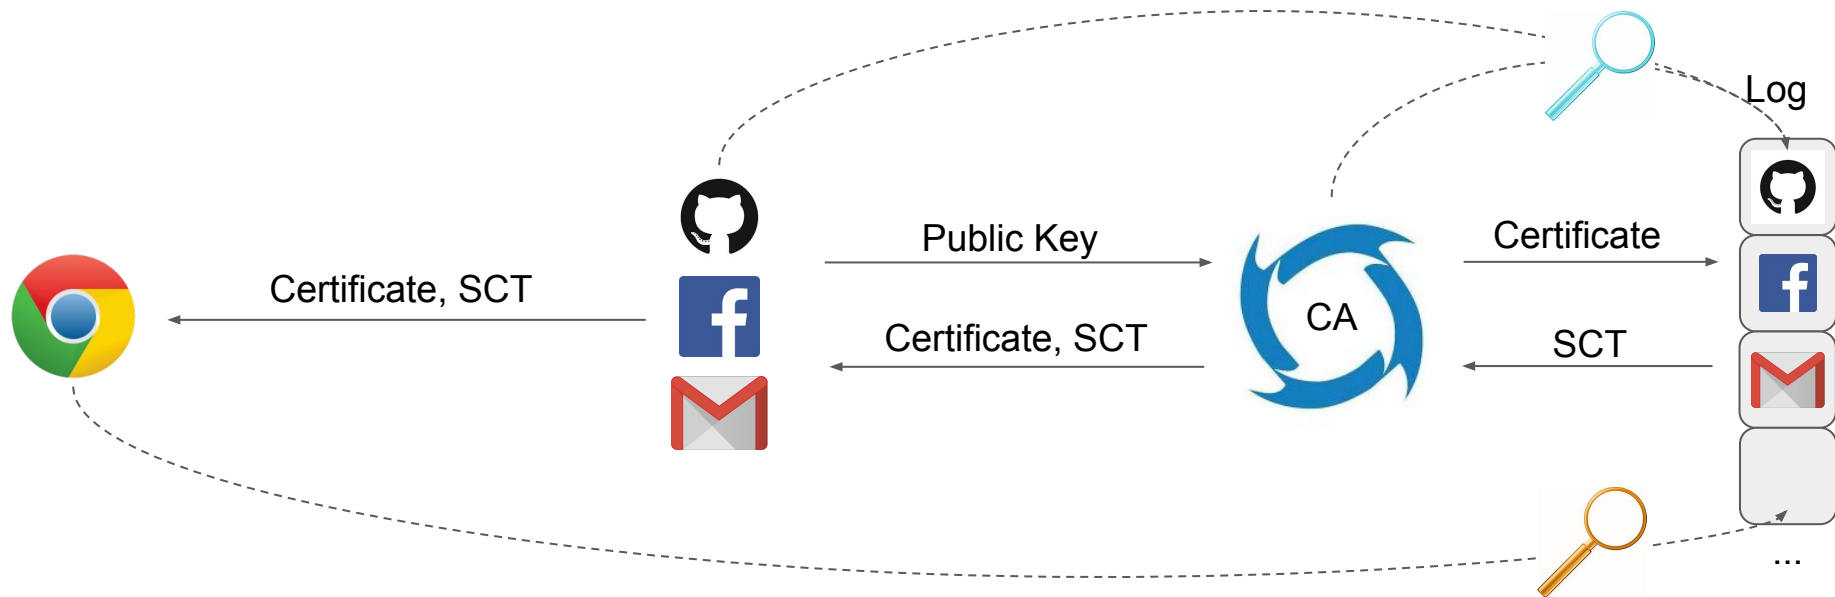
Certificate Transparency (CT)

Idea: public, verifiable log of all certificates




Certificate Transparency (CT)

Idea: public, verifiable log of all certificates



Certificate Transparency (CT)

Idea: public, verifiable log of all certificates



CT logging required by chrome for all sites starting October 2017!

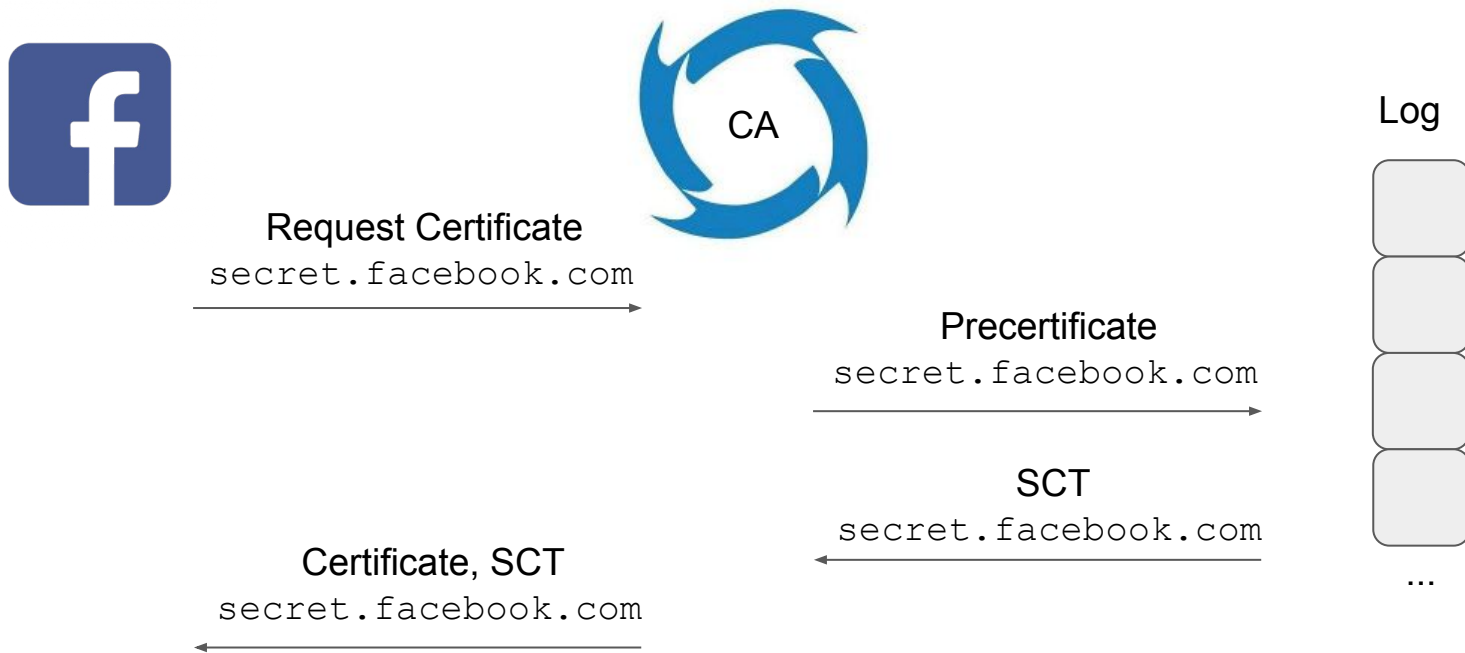
Transparency and Privacy?



Outline

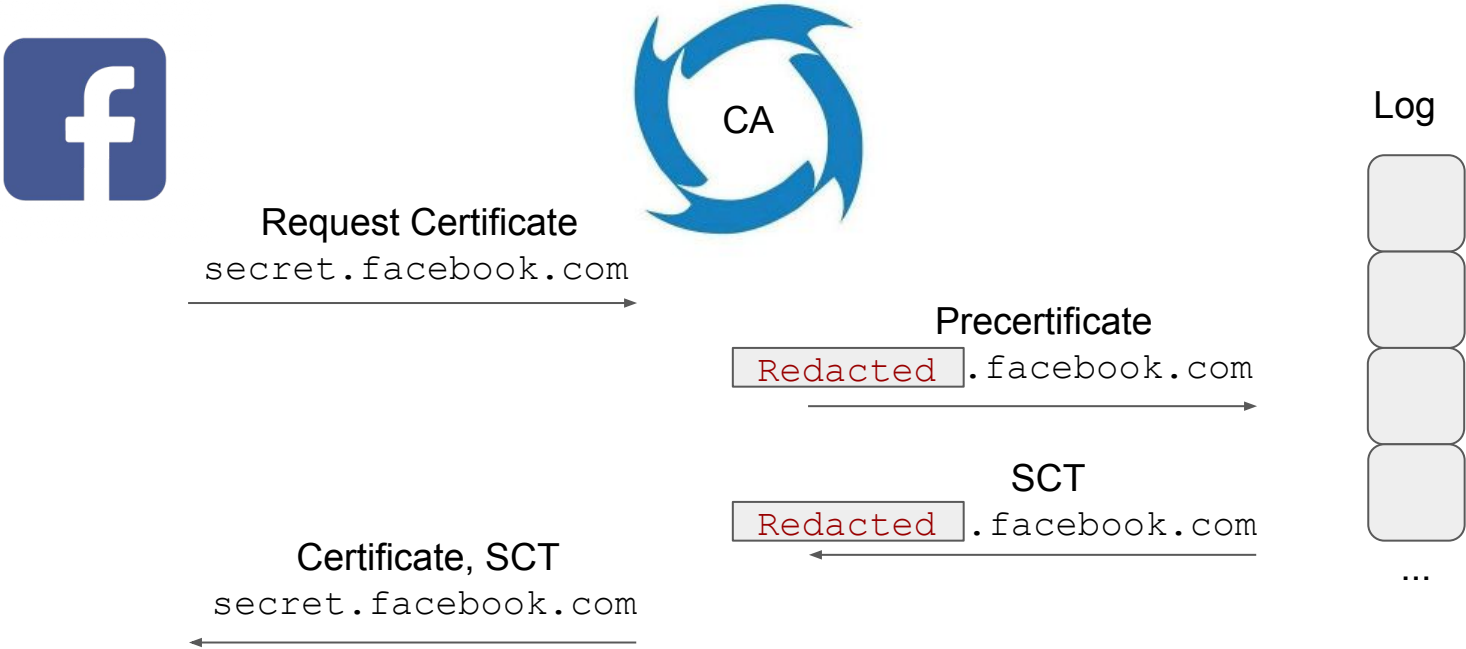
- Certificate Transparency
- **Redaction of private subdomains**
- Privacy-preserving proof of misbehavior

Redaction: keeping secrets on a public log



Problem: `secret.facebook.com` is publicly visible on the log!

Redaction: keeping secrets on a public log




Problem: `secret.facebook.com` is publicly visible on the log!

Tools: Commitments

Usage:

$c \leftarrow \text{Commit}(m, r)$

$\text{Verify}(c, m, r)$



Security Properties:

Hiding: given commitment $\text{Commit}(m, r)$, can't find m


Binding: given commitment $\text{Commit}(m, r)$, can't decommit to $m' \neq m$

Tools: Commitments

Usage:

$c \leftarrow \text{Commit}(m, r)$

$\text{Verify}(c, m, r)$



Security Properties:

Hiding: given commitment $\text{Commit}(m, r)$, can't find m



Binding: given commitment $\text{Commit}(m, r)$, can't decommit to $m' \neq m$

Tools: Commitments

Usage:

$c \leftarrow \text{Commit}(m, r)$

$\text{Verify}(c, m, r)$



$\text{Verify}(\text{chest}, \text{val}, r)$

Security Properties:

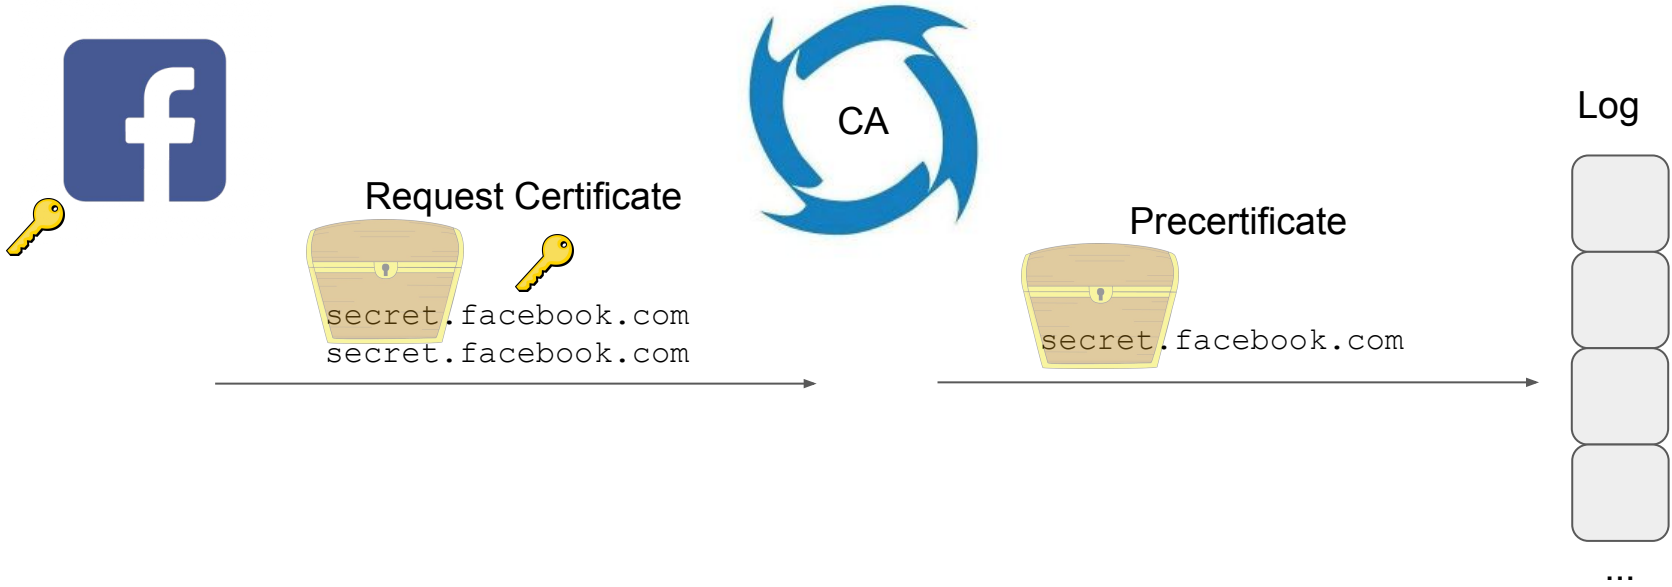
Hiding: given commitment $\text{Commit}(m, r)$, can't find m

Binding: given commitment $\text{Commit}(m, r)$, can't decommit to $m' \neq m$

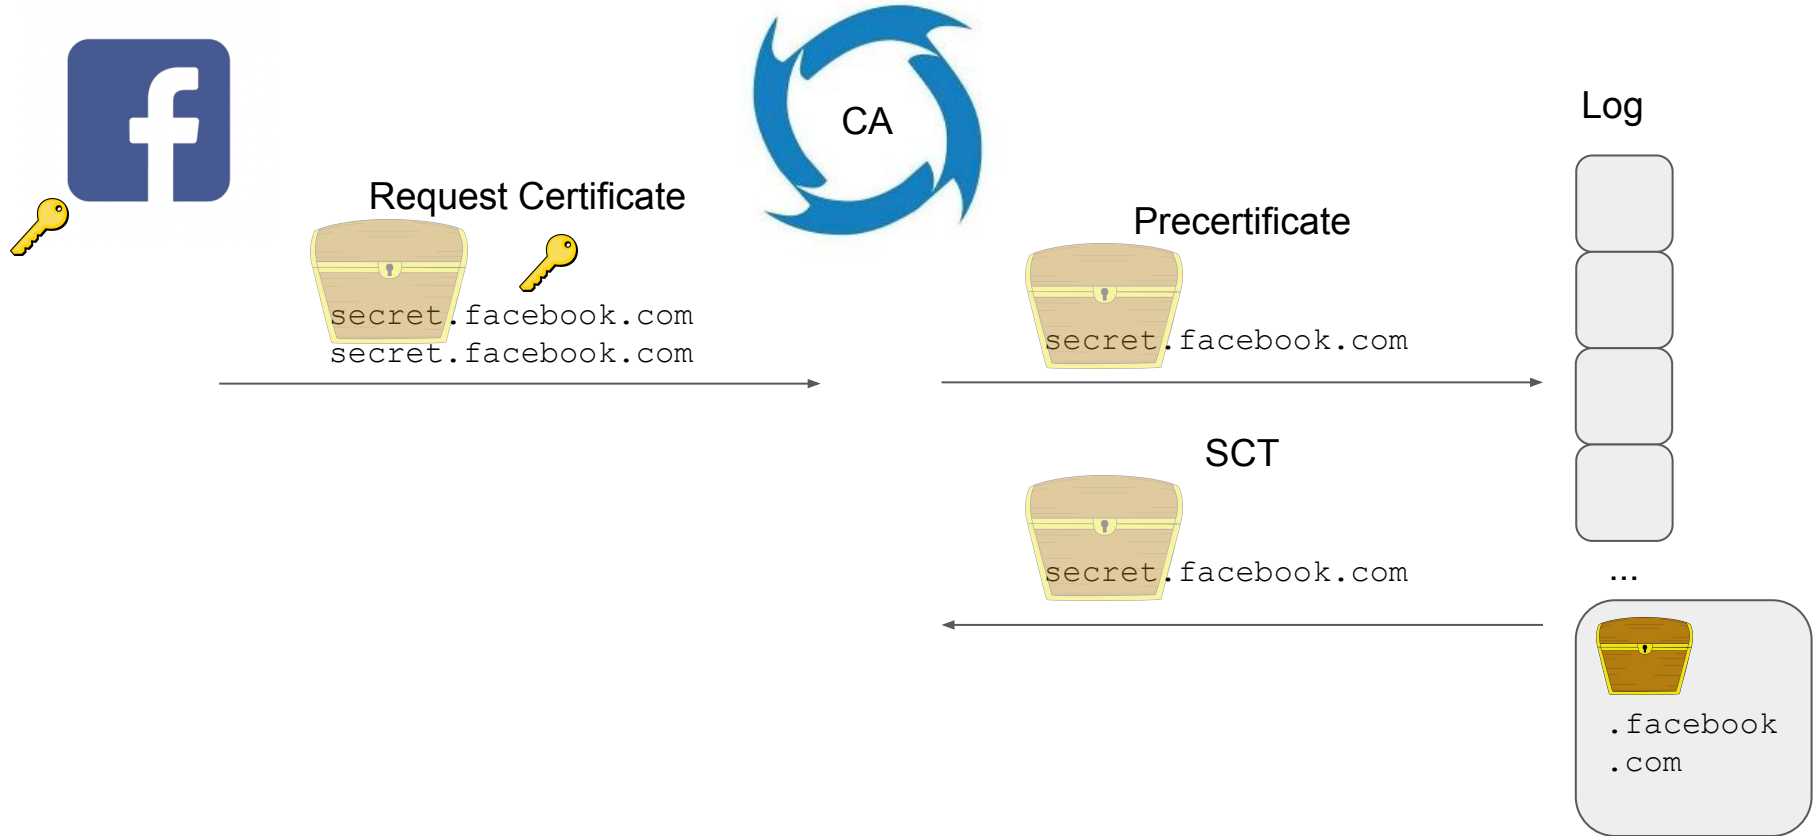
Subdomain Redaction via Commitments




Subdomain Redaction via Commitments



Subdomain Redaction via Commitments



Subdomain Redaction via Commitments



Subdomain Redaction via Commitments




Page Request: `secret.facebook.com`

Subdomain Redaction via Commitments



Page Request: `secret.facebook.com`

A horizontal arrow pointing from left to right, indicating the direction of the page request.

Subdomain Redaction via Commitments



Page Request: `secret.facebook.com`



Verify(, `secret`, )

Security

How can a monitor still check the log?

Knowledge of number of entries per domain owner reveals extra certificates


Why can't a malicious site or CA reuse an existing redacted SCT?

Binding property of commitment

Outline

- Certificate Transparency
- Redaction of private subdomains
- **Privacy-preserving proof of misbehavior**

Privacy-Compromising Proof of Exclusion




Excluded
SCT




`secret.facebook.com`

Privacy-Compromising Proof of Exclusion



Excluded
SCT



secret.facebook.com



Goals

- Auditor proves to vendor that an SCT is missing from log
- Auditor does not reveal domain name, vendor only learns that log is misbehaving

Goals

- Auditor proves to vendor that an SCT is missing from log
- Auditor does not reveal domain name, vendor only learns that log is misbehaving

Then:

- Vendor can investigate log
- Vendor can **blindly** revoke missing certificate (by pushing a revocation value to all browsers)

Goals


- Auditor proves to vendor that an SCT is missing from log
- Auditor does not reveal domain name, vendor only learns that log is misbehaving

Then:

- Vendor can investigate log
- Vendor can **blindly** revoke missing certificate (by pushing a revocation value to all browsers)

Assumption: timestamps in order


What Does Auditor Prove?



Excluded
SCT



What Does Auditor Prove?

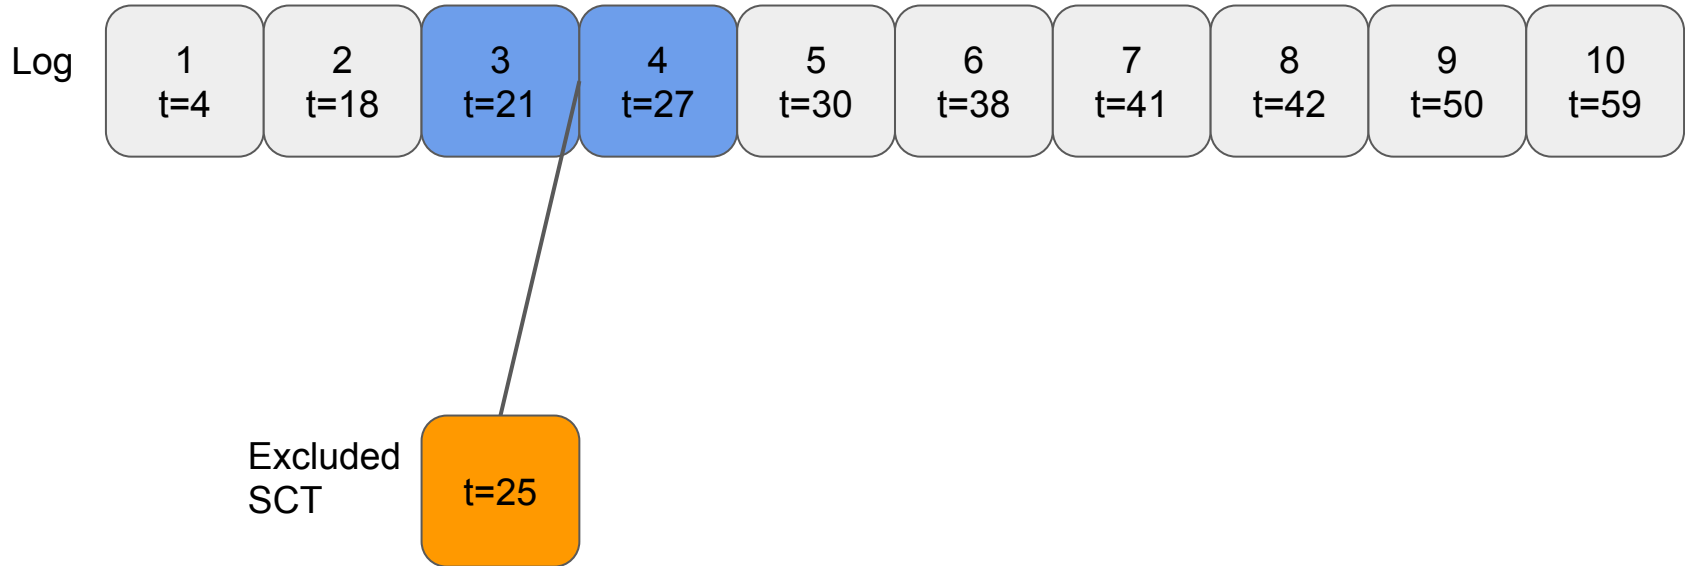


Excluded
SCT




Assumption: timestamps in order

What Does Auditor Prove?




Assumption: timestamps in order

What Does Auditor Prove?





What Does Auditor Prove?




What about privacy?!




Tools: Additively Homomorphic Commitments




Tools: Additively Homomorphic Commitments




Tools: Additively Homomorphic Commitments





Tools: Zero-Knowledge Proofs






Tools: Zero-Knowledge Proofs







Tools: Zero-Knowledge Proofs





Tools: Zero-Knowledge Proofs




Tools: Zero-Knowledge Proofs



Tools: Zero-Knowledge Proofs




Proof of Exclusion




What about
privacy?!



Proof of Exclusion



X Y Z



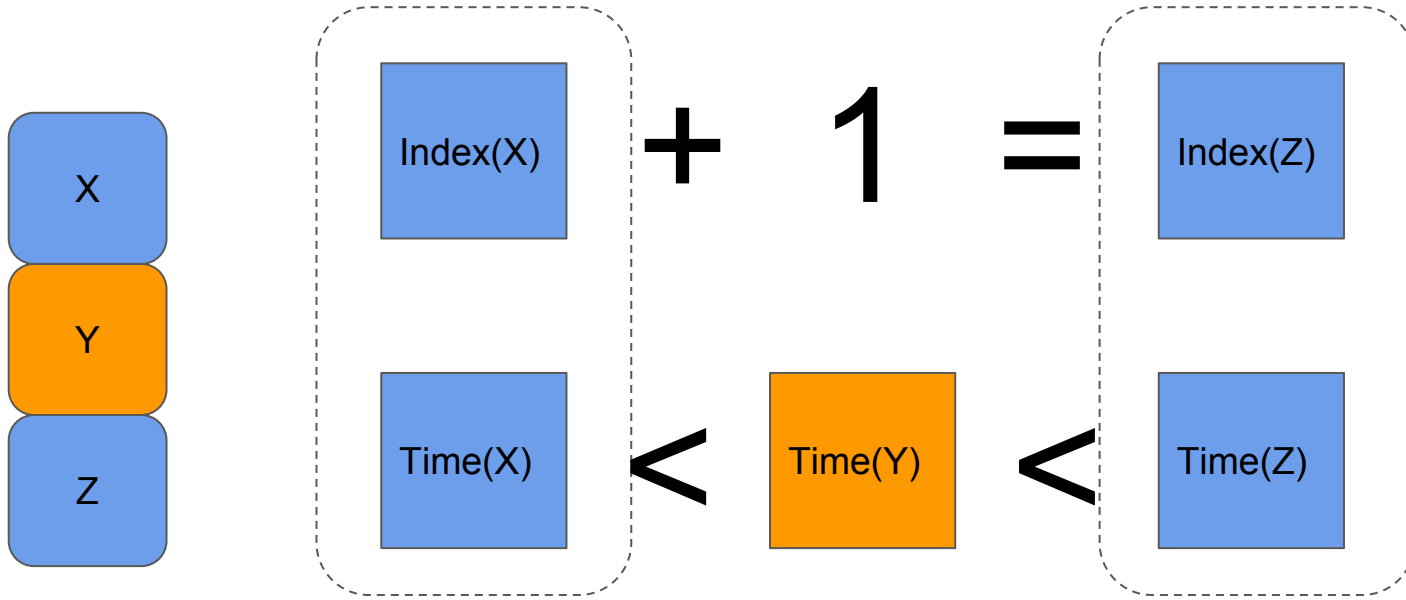
What about privacy?!





Proof of Exclusion



Proof of Exclusion




Proof of Exclusion




Proof of Exclusion

- X
- Y
- Z



Proof of Exclusion

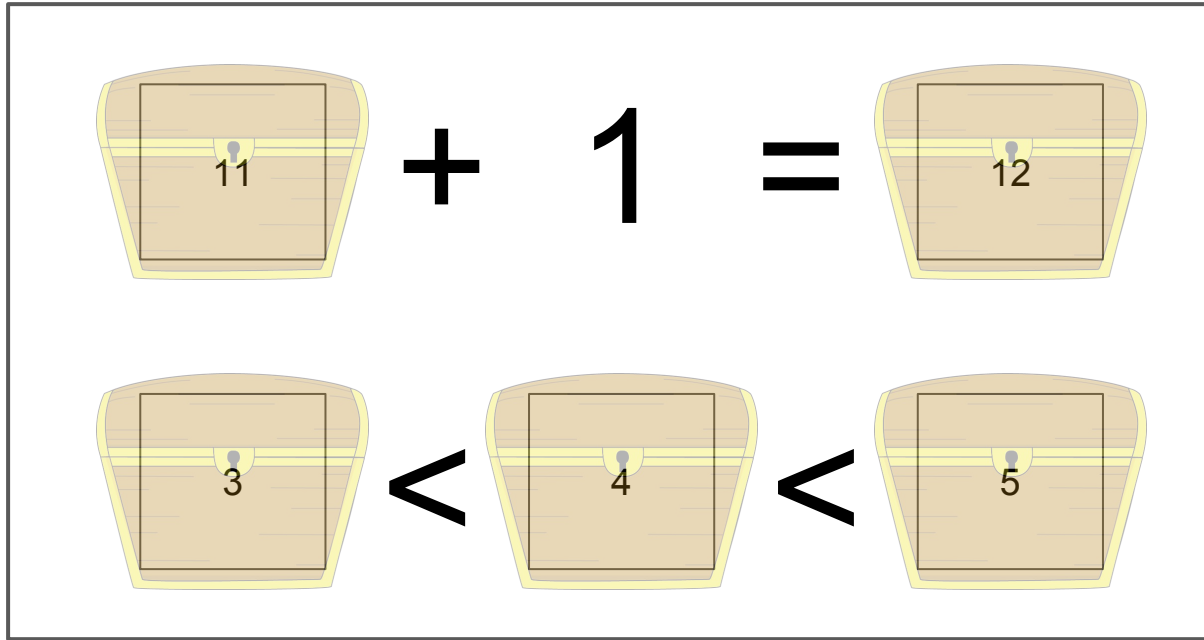

- X
- Y
- Z



Are these numbers *really* from the log?




Proof of Exclusion



hehehe...



Proof of Exclusion




Needed for
proof




Index(X)

Time(X)


Proof of Exclusion




New
signatures
from log







Needed for
proof





Proof of Exclusion




New signatures from log






Needed for proof



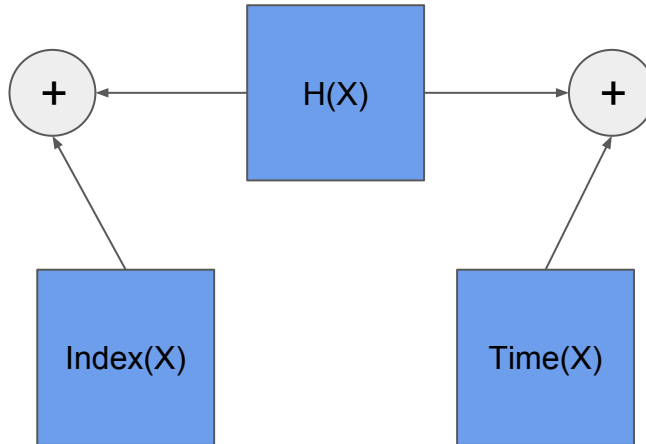
Proof of Exclusion




New signatures from log







Needed for proof



Proof of Exclusion




New signatures from log








Needed for proof


Proof of Exclusion



New signatures from log



Needed for proof



Performance Numbers

Online Costs

Proof Size: 333 kB

Time to generate: 5.0 seconds

Time to verify: 2.3 seconds

Offline Costs (storage)

Growth of log entry: 480 bytes

Growth of SCT: 160 bytes

Revocation notice size: 32 bytes

Summary

- CT is an exciting new feature of our web infrastructure
- Transparency raises new privacy concerns
- Work on privacy-preserving solutions to two issues:
 - Compatibility between CT and need for private domain names
 - Reporting CT log misbehavior without revealing private information

See paper for details and security proofs: <https://arxiv.org/pdf/1703.02209.pdf>