Private Access Tokens

draft-private-access-tokens-01

Tommy Pauly, Chris Wood, Jana Iyengar,
Steven Valdez, Scott Hendrickson

SECDISPATCH

IETF 112, November 2021, Virtual

Agenda

Motivation

Protocol architecture

Deployment considerations

Dispatching

What problem are we trying to solve?

Servers often use client IP addresses as an identification mechanism

Servers can recognize these addresses over time. They can use them to rate-limit access to their server.

The Metropolitan Times

Providing your daily dose of news

SUBSCRIBE Log In

Tampa Bay Buccaneers win the 2021 Super Bowl

Jackson Terry 5 min read 02/08/2021

Raymond James Stadium Decorated for the Super Bowl – Image Credit: US Customs and Border Protection

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. Scelerisque fermentum dui faucibus in ornare. Vestibulum lorem sed risus ultricies tristique nulla aliquet enim. Morbi tristique senectus et netus et. Feugiat in ante metus dictum at tempor commodo ullamcorper. Laoreet sit amet cursus sit. A diam sollicitudin tempor id eu nisl. Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua. *Feugiat scelerisque varius morbi enim nunc faucibus a pellentesque.*

https://en.wikipedia.org/wiki/Paywall#/media/File:Metered_Paywall_Example.svg

IP addresses are also bad at identifying correctly in many cases, like behind large NATs

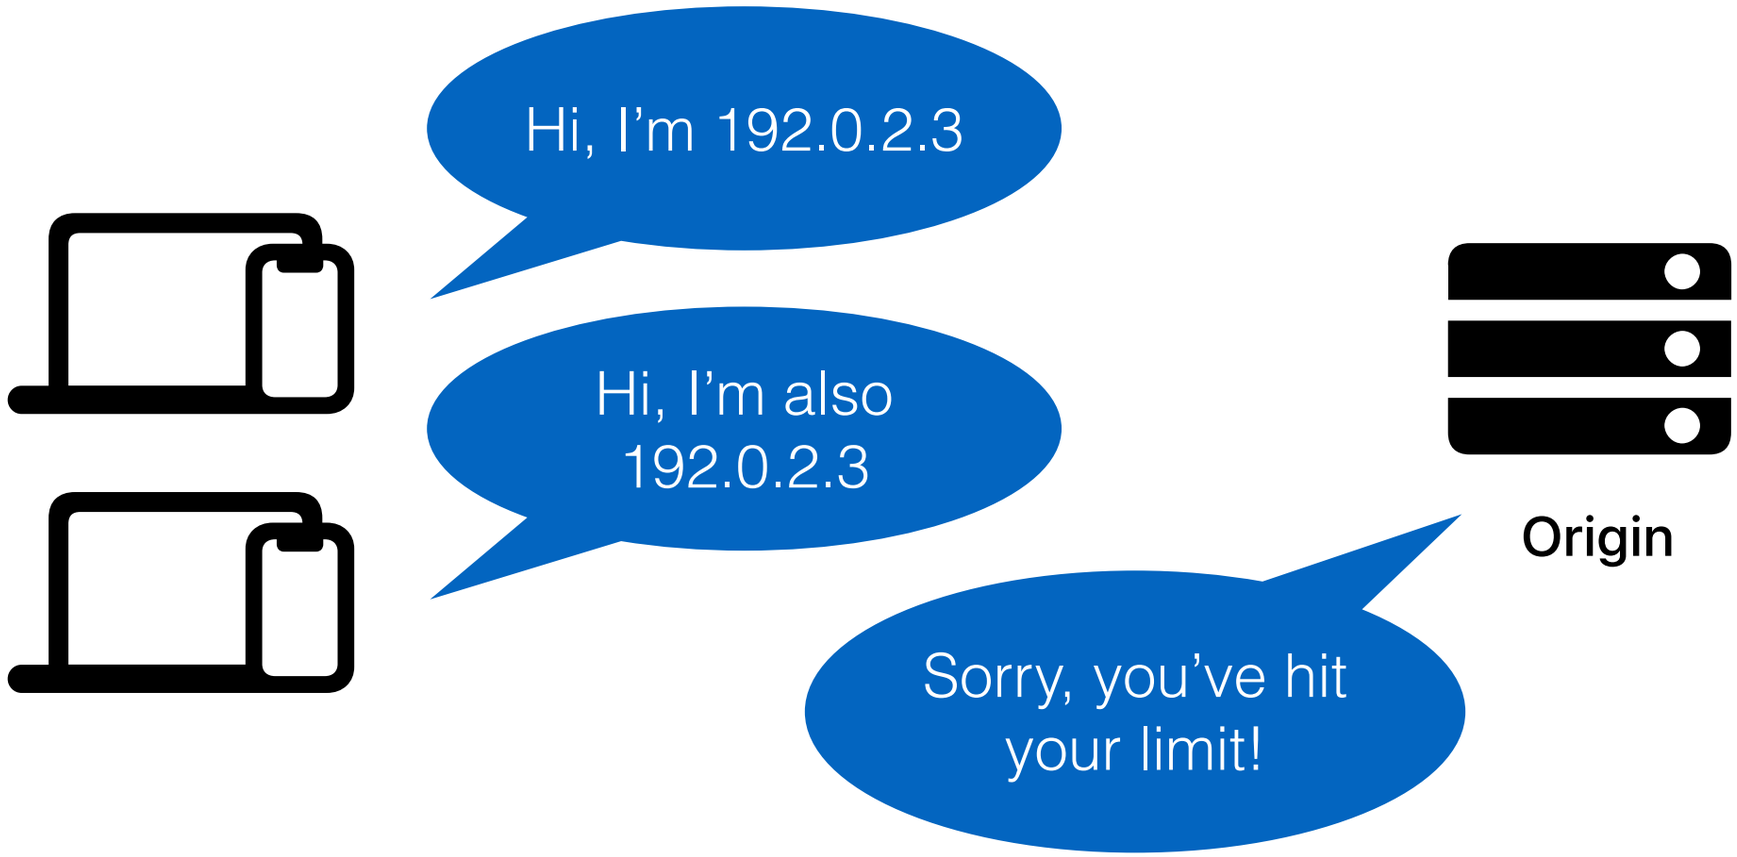

Proxies (Private Relay), VPNs, and Tor all improve IP address privacy

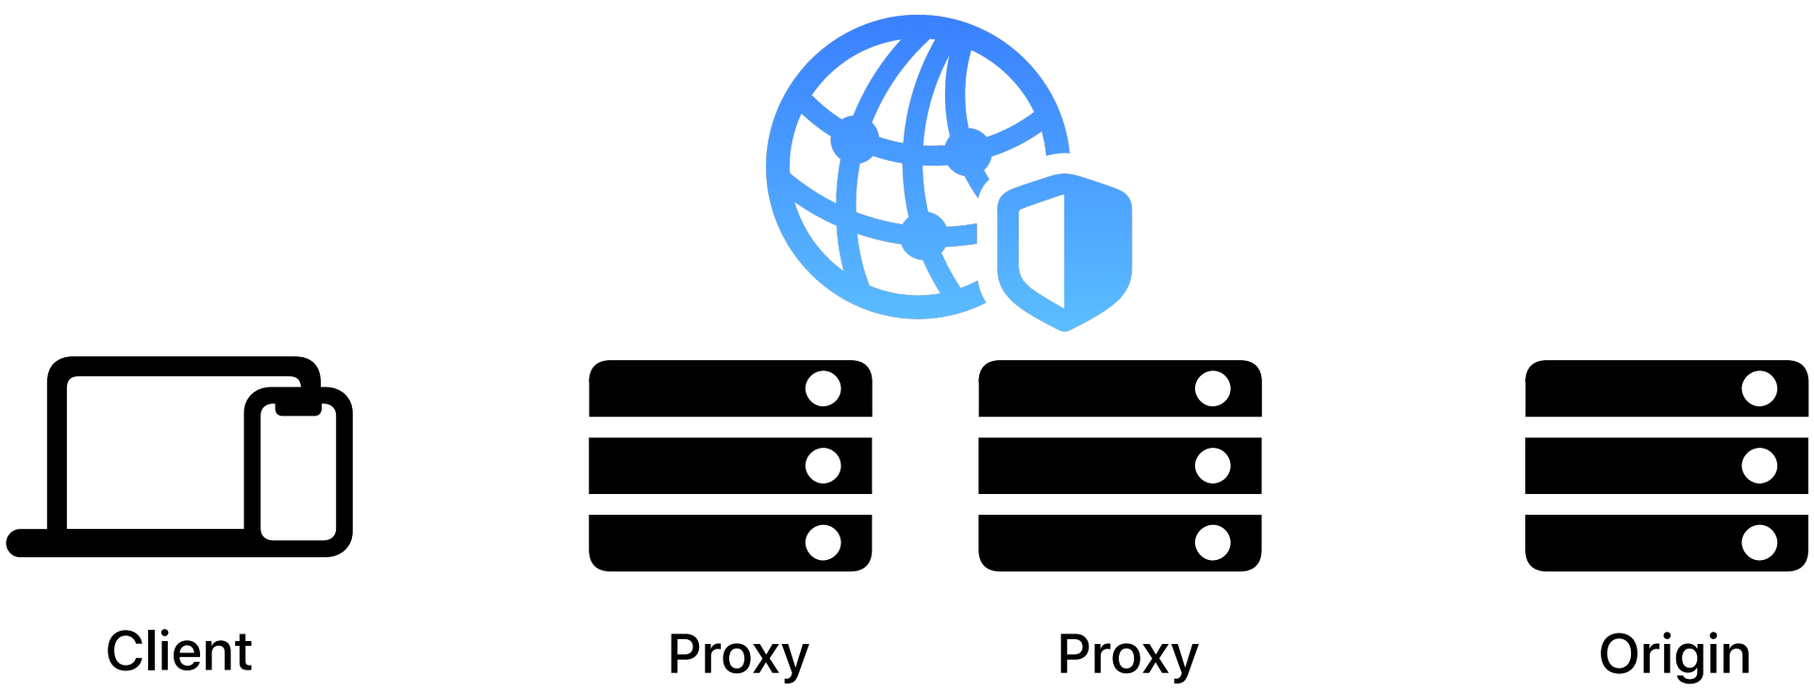

However, this makes rate-limiting harder

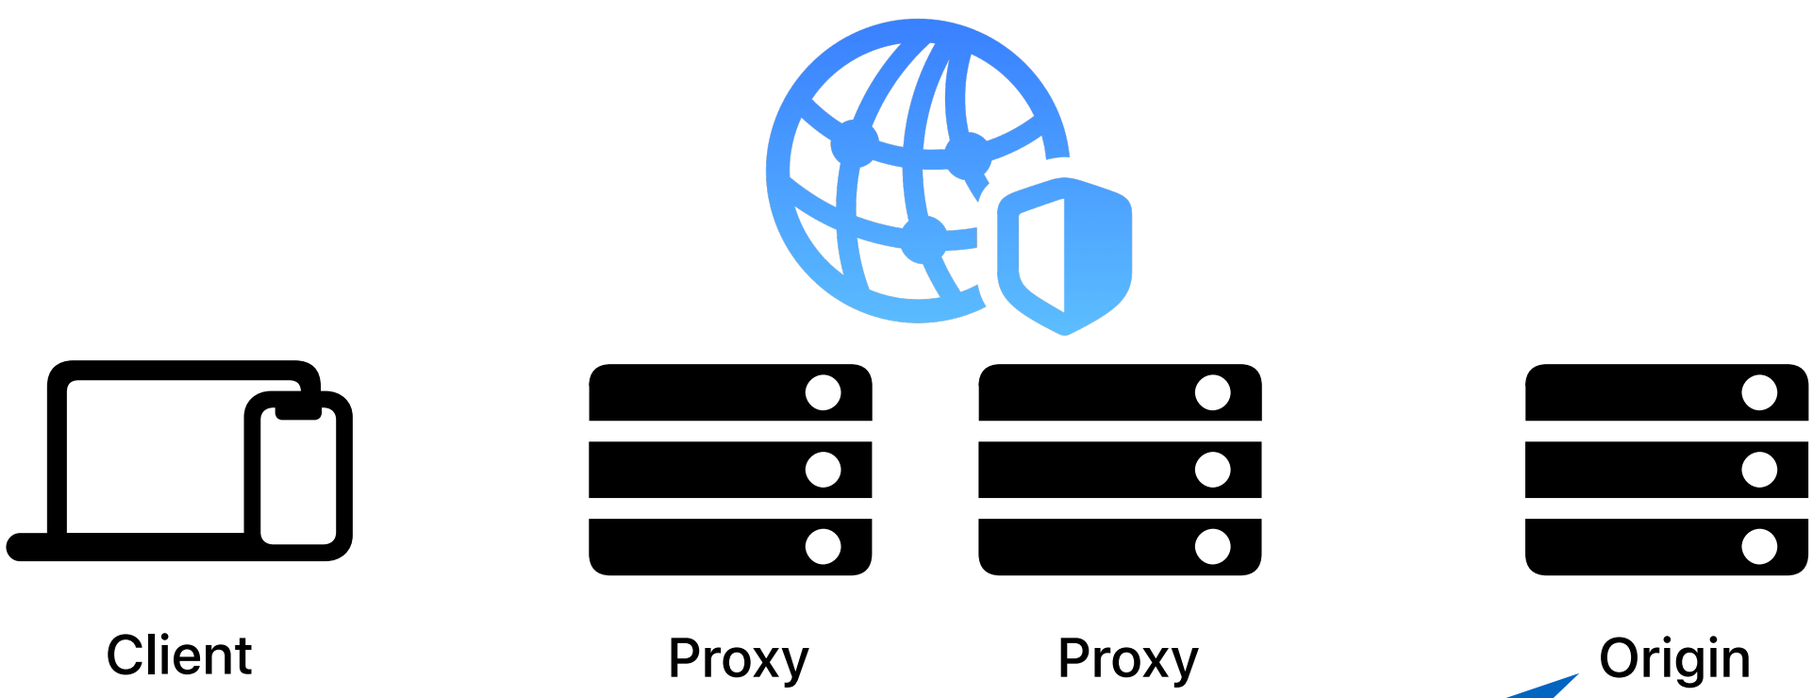

Allow rate-limits to work, regardless of IP address

Don’t introduce a new stable identifier

Where is this useful?

Anonymous access based on limited client state, like per-origin rate-limiting

Not for cases where you log in, since that is a stronger identity

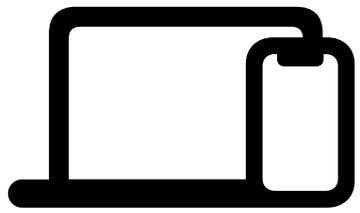

Anonymous access, no rate limit (stateless)

Read Wikipedia

Use a search engine

Anonymous access, rate limited (stateful)

Account log-in / creation

Read newspaper article

Authenticated access

Upload to a social media account

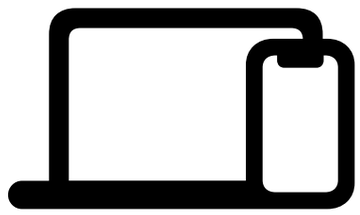

Anonymous access, rate limited (stateful)

Account log-in / creation

Read newspaper article

Private Access Tokens solve this use case

Client can prove to the Origin that it has performed fewer than N accesses in a time window

No entity can correlate user identity with browsing history

How do Private Access Tokens work?

Token Challenge and Request

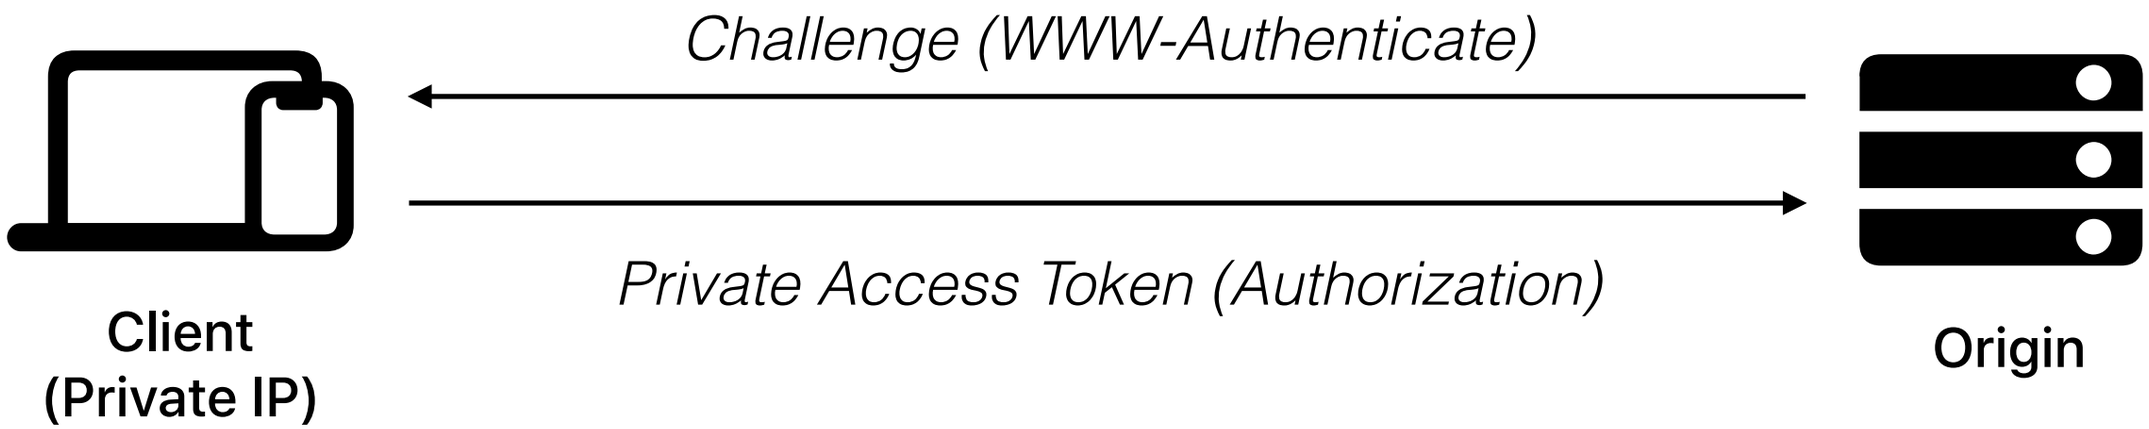

Clients access Origin, potentially using a private/shared IP address

Origins can challenge clients on sensitive operations (creating an account, reading an article without logging in)

Clients fetch an unlinkable token for the origin, and present it

Token Issuance

Who can issue per-origin tokens?

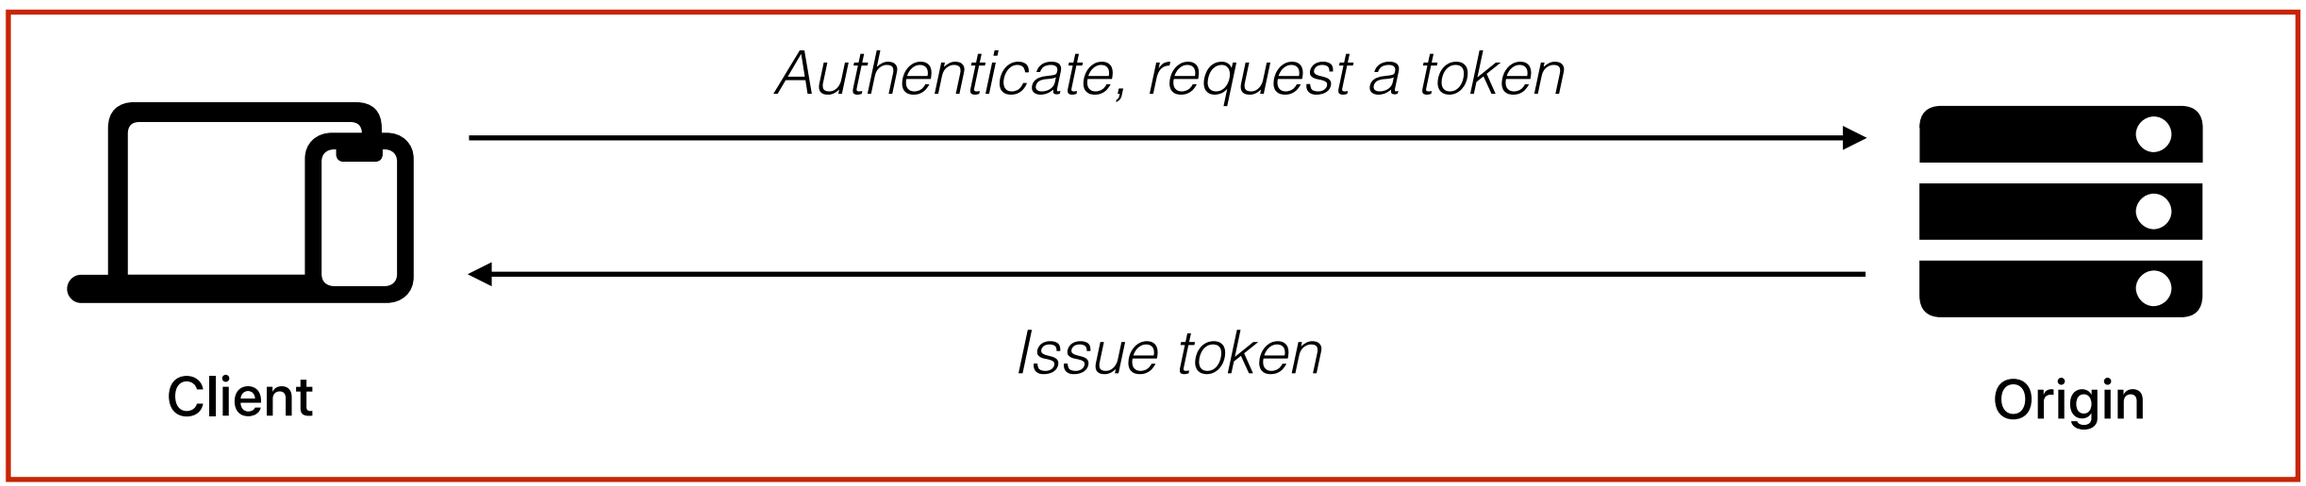

Origin? No!

Client doesn’t want to, or can’t, authenticate to the origin

Token Issuance

Who can issue per-origin tokens?

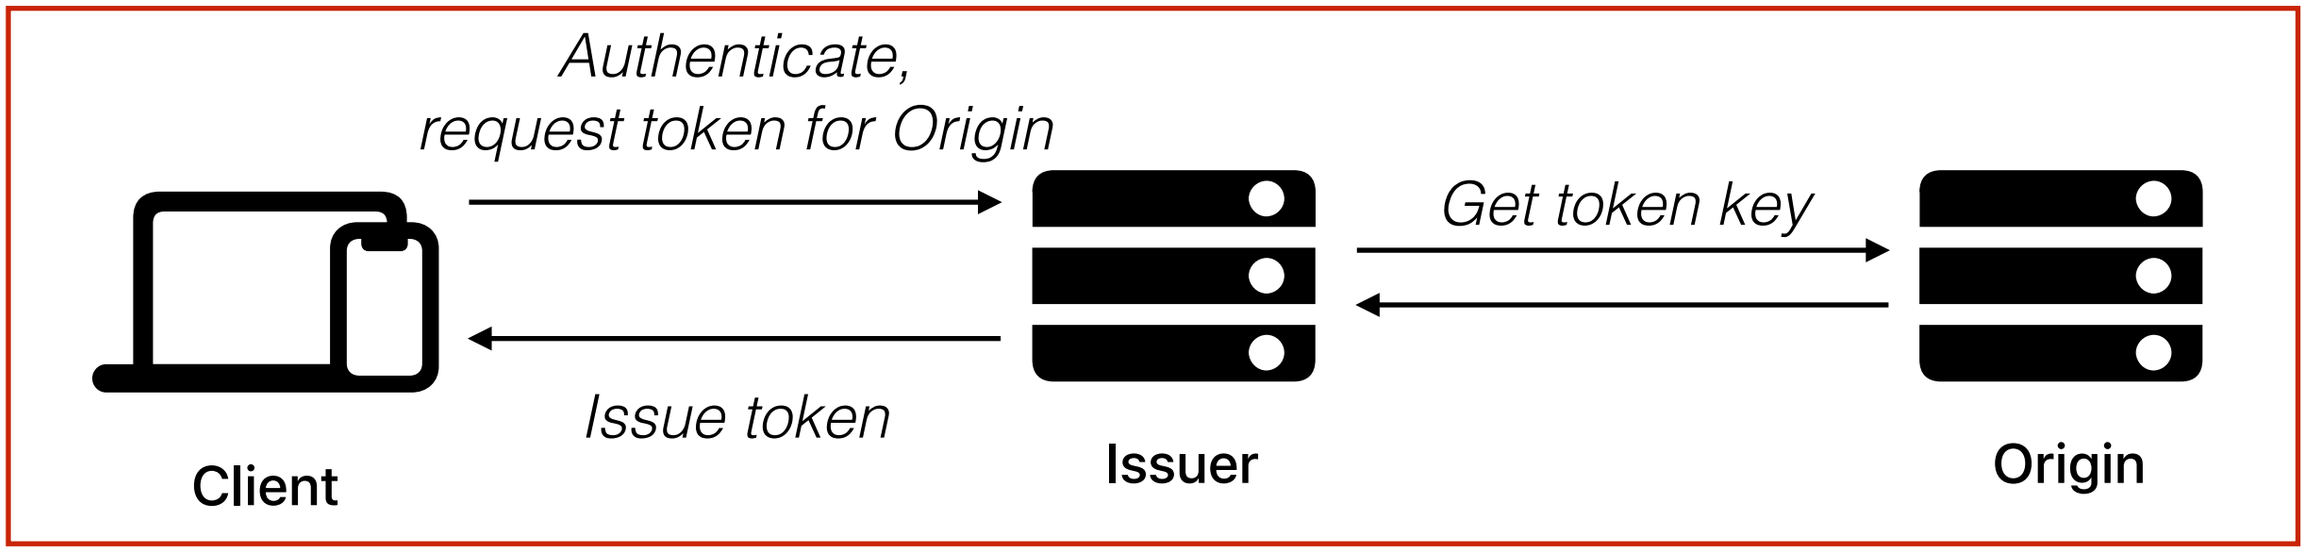

Trusted issuer? No!

Issuer would learn client browsing history

Token Issuance

Who can issue per-origin tokens?

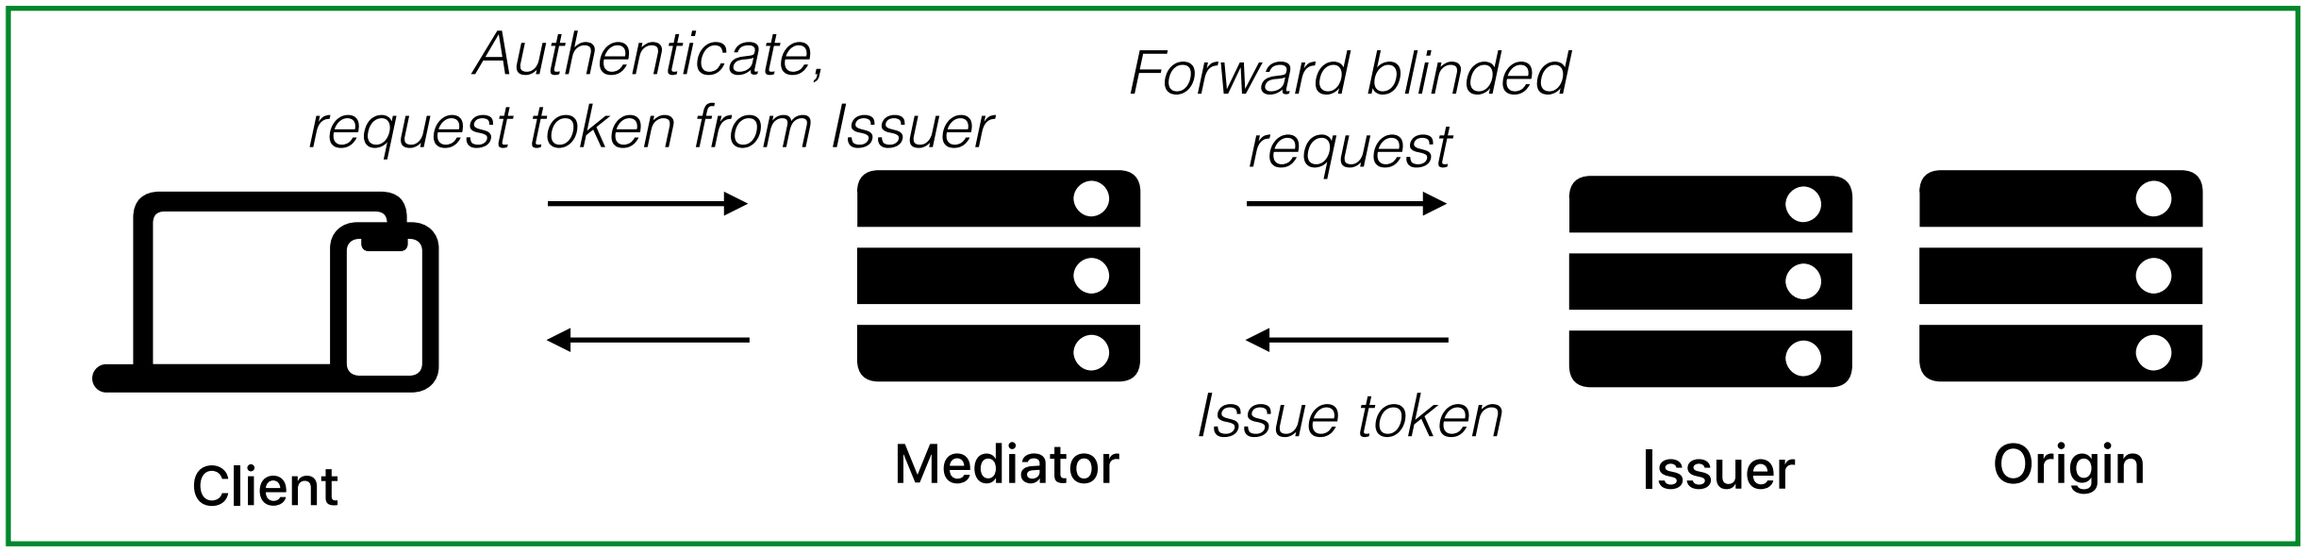

Combination of client-trusted Mediator and origin-trusted Issuer

Mediator checks, then hides, client identity. Mediator only sees Issuer name, not Origin

Issuer enforces policy on behalf of the Origin

Full Protocol

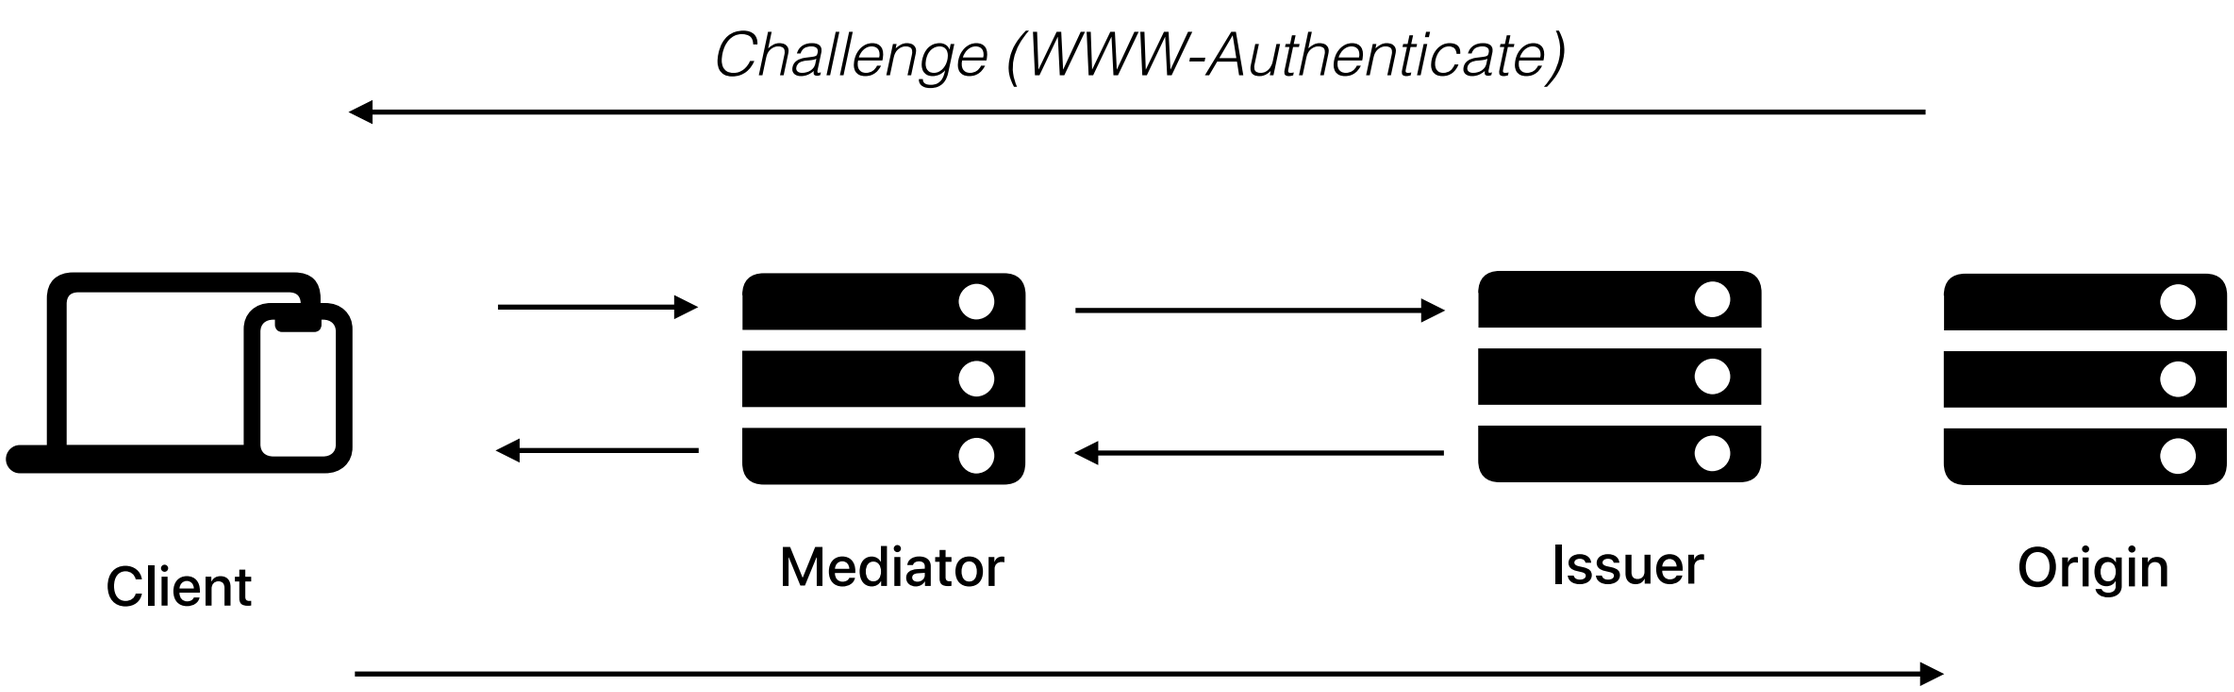

Configuration and state

Clients only tell Mediators about an “ANON_ORIGIN” *; Actual origin name is encrypted to Issuer

Mediators keep a count of tokens issued for each client per “ANON_ORIGIN”

Issuers define a “policy window”, which defines when the count on the Mediator rolls over

* Mediators can detect if clients lie about ANON_ORIGIN -> Origin mappings

Cryptographic Dependencies

Challenge and Redemption (Origin)

RSA Blind Signatures

Issuance (Client, Mediator, Issuer)

RSA Blind Signatures Client

HPKE

Blinded DH with Schnorr Proof-of-Knowledge (see CFRG presentation for details)

How is this deployed?

Deployment Expectations

Clients choose trusted Mediators

- Based on device certs, verified account logins, etc

Origins choose trusted Issuers

- Each Issuer should serve many Origins
- Existing CDN, hosting, or security service relationships

Mediators and Issuers need mutual trust

- Should be different entities for best privacy properties

Architecture

Each Mediator serves many Clients, each Issuer serves many Origins

This protects Client and Origin identities

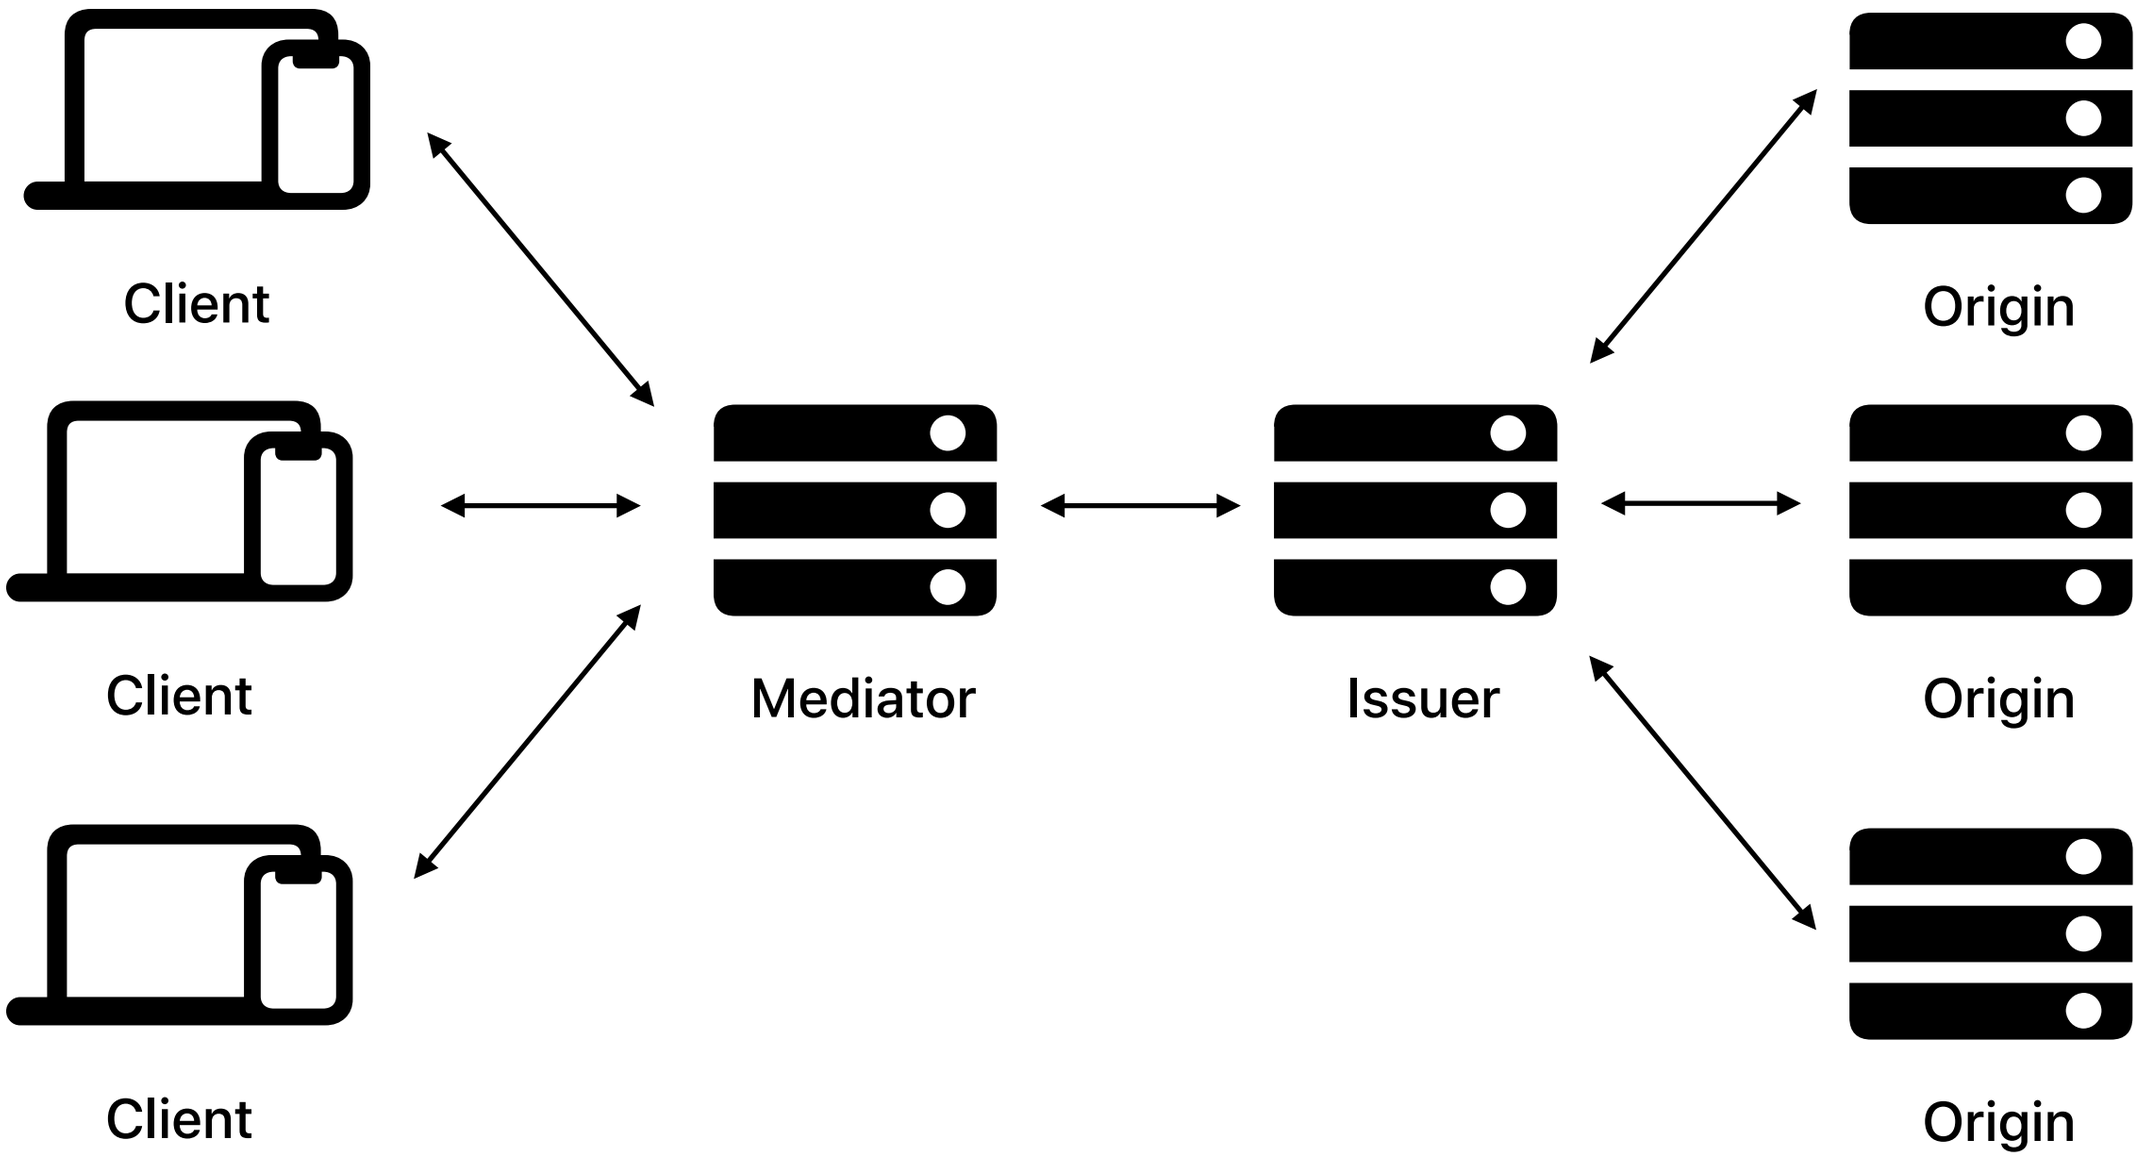

Client Identity

Mediators are responsible for determining what a "client" is

- PATs don't require one mechanism for this
  - Needs to be something that the ecosystem agrees is hard to forge

Users can have (few) multiple identities

- Different devices and accounts
- Limited in ability to amplify

Avoiding centralization

Mediators and Issuers are entities that help represent many clients and origins

- We should avoid letting this become an ecosystem that consolidates down to a few entities

- It needs to be easy for new Mediator and Issuer services to enter

- Avoid situations where Issuers (on behalf of Origins) only allow a handful of Mediators

PATs may actually be able to have less centralization than other alternatives

Origins can already prefer to use a fast sign-in to prefer known partners, who may be sharing data

Without an alternative for clients to use, pressure to avoid captchas can move towards signing in with major services more

Privacy Pass allows a client to present a token from some other origin

Redeeming origins can choose to discriminate based on where tokens came from, to prefer major services

With PATs, Origins don’t see Mediators — they can’t discriminate based on how the client authenticated

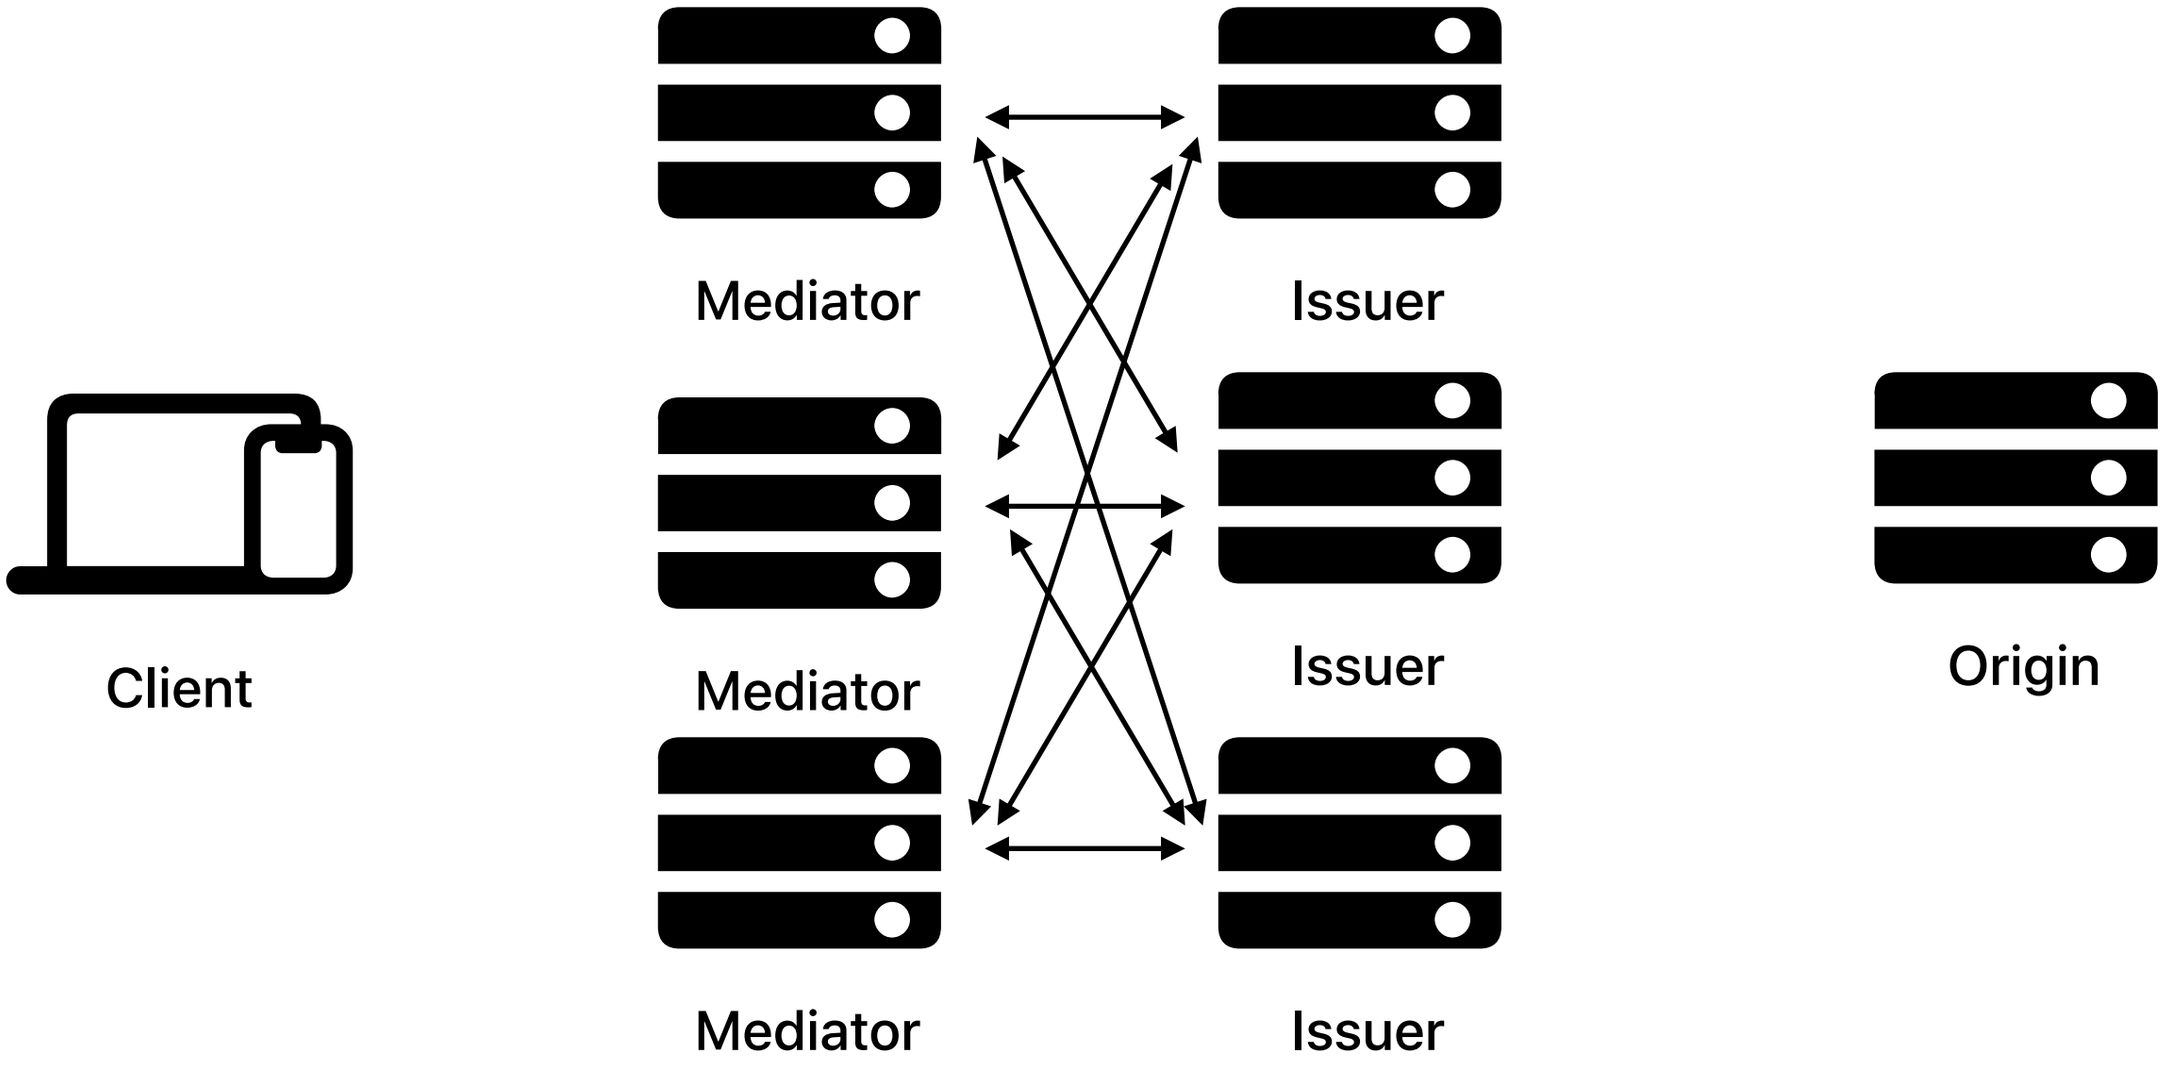

If Issuers start rejecting new Mediators, it could be publicly reported and audited

Similar to entities deciding to reject Certificate Authorities

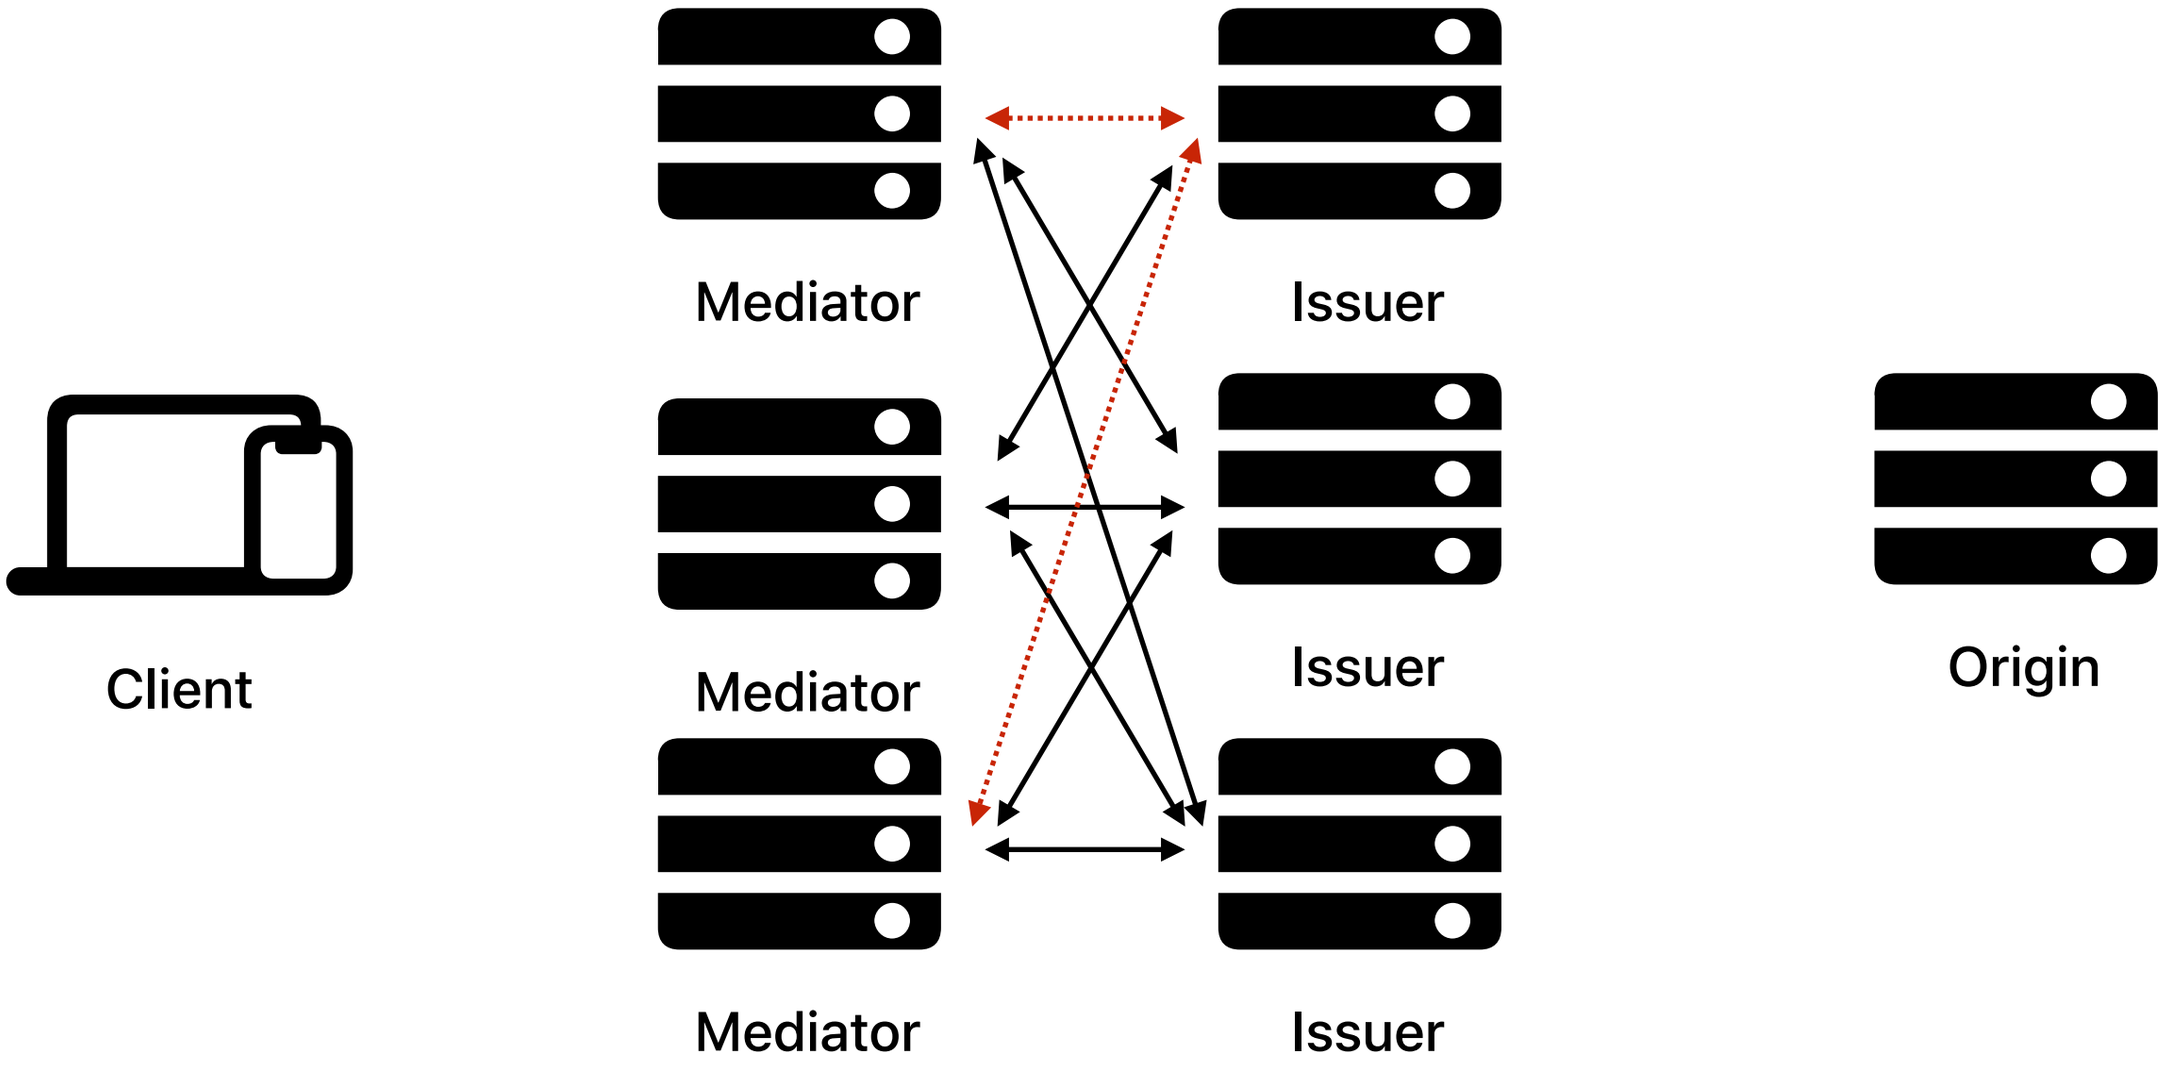

Where should this work be done?

Related Work

Private Access Tokens (PATs) differ from *Privacy Pass* in four key ways:

- Per-client per-origin state (not unlimited access)
- Per-origin tokens (no cross-origin spending)
- Online challenge-based (limiting token hoarding)
- Publicly verifiable (offline verification)

Is this a more generic form of Privacy Pass?

Dispatch

Where should this work be done?

- Privacy Pass Working Group
- Short-lived Working Group (like OHAI)