

Workgroup: Network Working Group

Internet-Draft:

draft-ietf-privacypass-architecture-16

Published: 25 September 2023

Intended Status: Informational

Expires: 28 March 2024

Authors: A. Davidson J. Iyengar C. A. Wood

 LIP Fastly Cloudflare

The Privacy Pass Architecture

Abstract

This document specifies the Privacy Pass architecture and requirements for its constituent protocols used for authorization based on privacy-preserving authentication mechanisms. It describes the conceptual model of Privacy Pass and its protocols, its security and privacy goals, practical deployment models, and recommendations for each deployment model that helps ensure the desired security and privacy goals are fulfilled.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 March 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the

Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
- [3. Architecture](#)
 - [3.1. Overview](#)
 - [3.2. Use Cases](#)
 - [3.3. Privacy Goals and Threat Model](#)
 - [3.4. Redemption Protocol](#)
 - [3.5. Issuance Protocol](#)
 - [3.5.1. Attester Role](#)
 - [3.5.2. Issuer Role](#)
 - [3.5.3. Issuance Metadata](#)
 - [3.5.4. Future Issuance Protocol Requirements](#)
 - [3.6. Information Flow](#)
 - [3.6.1. Token Challenge Flow](#)
 - [3.6.2. Attestation Flow](#)
 - [3.6.3. Issuance Flow](#)
 - [3.6.4. Token Redemption Flow](#)
- [4. Deployment Models](#)
 - [4.1. Shared Origin, Attester, Issuer](#)
 - [4.2. Joint Attester and Issuer](#)
 - [4.3. Joint Origin and Issuer](#)
 - [4.4. Split Origin, Attester, Issuer](#)
- [5. Deployment Considerations](#)
 - [5.1. Discriminatory Treatment](#)
 - [5.2. Centralization](#)
- [6. Privacy Considerations](#)
 - [6.1. Partitioning by Issuance Metadata](#)
 - [6.2. Partitioning by Issuance Consistency](#)
 - [6.3. Partitioning by Side-Channels](#)
- [7. Security Considerations](#)
 - [7.1. Token Caching](#)
- [8. IANA Considerations](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Appendix A. Acknowledgements](#)
- [Authors' Addresses](#)

1. Introduction

Privacy Pass is an architecture for authorization based on privacy-preserving authentication mechanisms. In other words, relying parties authenticate clients in a privacy-preserving way, i.e., without learning any unique, per-client information through the

authentication protocol, and then make authorization decisions on the basis of that authentication succeeding or failing. Possible authorization decisions might be to provide clients with read access to a particular resource or write access to a particular resource.

Typical approaches for authorizing clients, such as through the use of long-term state (cookies), are not privacy-friendly since they allow servers to track clients across sessions and interactions. Privacy Pass takes a different approach: instead of presenting linkable state-carrying information to servers, e.g., a cookie indicating whether or not the client is an authorized user or has completed some prior challenge, clients present unlinkable proofs that attest to this information. These proofs, or tokens, are private in the sense that a given token cannot be linked to the protocol interaction where that token was initially issued.

At a high level, the Privacy Pass architecture consists of two protocols: redemption and issuance. The redemption protocol, described in [[AUTHSCHEME](#)], runs between Clients and Origins (servers). It allows Origins to challenge Clients to present tokens for consumption. Origins verify the token to authenticate the Client -- without learning any specific information about the Client -- and then make an authorization decision on the basis of the token verifying successfully or not. Depending on the type of token, e.g., whether or not it can be cached, the Client either presents a previously obtained token or invokes an issuance protocol, such as [[ISSUANCE](#)], to acquire a token to present as authorization.

This document describes requirements for both redemption and issuance protocols and how they interact. It also provides recommendations on how the architecture should be deployed to ensure the privacy of clients and the security of all participating entities.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

The following terms are used throughout this document:

Client: An entity that seeks authorization to an Origin. Using [[RFC9334](#)] terminology, Clients implement the RATS Attester role.

Token: A cryptographic authentication message used for authorization decisions, produced as output from an issuance mechanism or protocol.

Origin:

An entity that consumes tokens presented by Clients and uses them to make authorization decisions.

Token challenge: The mechanism by which Origins request tokens from Clients, often denoted TokenChallenge.

Token request: A message used by Clients to request a token from an Issuer, often denoted TokenRequest.

Token response: A message used by Issuers to send tokens to a Client, often denoted TokenResponse.

Redemption: The mechanism by which Clients present tokens to Origins for the purposes of authorization.

Issuer: An entity that issues tokens to Clients for properties attested to by the Attester.

Issuance: The mechanism by which an Issuer produces tokens for Clients.

Attester: An entity that attests to properties of Clients for the purposes of token issuance. Using [[RFC9334](#)] terminology, Attesters implement the RATS Verifier role.

Attestation procedure: The procedure by which an Attester determines whether or not a Client has the specific set of properties that are necessary for token issuance.

The trust relationships between each of the entities in this list is further elaborated upon in [Section 3.3](#).

3. Architecture

The Privacy Pass architecture consists of four logical entities -- Client, Origin, Issuer, and Attester -- that work in concert for token redemption and issuance. This section presents an overview of Privacy Pass, a high-level description of the threat model and privacy goals of the architecture, and the goals and requirements of the redemption and issuance protocols. Deployment variations for the Origin, Issuer, and Attester in this architecture, including considerations for implementing these entities, are further discussed in [Section 4](#).

3.1. Overview

This section describes the typical interaction flow for Privacy Pass.

1. A Client interacts with an Origin by sending a request or otherwise interacting with the Origin in a way that triggers a response containing a token challenge. The token challenge indicates a specific Issuer to use.
2. If the Client already has a token available that satisfies the token challenge, e.g., because the Client has a cache of previously issued tokens, it can skip to [step 6](#) and redeem its token; see [Section 7.1](#) for security considerations of cached tokens.
3. If the Client does not have a token available and decides it wants to obtain one (or more) bound to the token challenge, it then invokes the issuance protocol. As a prerequisite to the issuance protocol, the Client runs the deployment-specific attestation process that is required for the designated Issuer. Client attestation can be done via proof of solving a CAPTCHA, checking device or hardware attestation validity, etc.; see [Section 3.5.1](#) for more details.
4. If the attestation process completes successfully, the client creates a token request to send to the designated Issuer (generally via the Attester, though it is not required to be sent through the Attester). The Attester and Issuer might be functions on the same server, depending on the deployment model (see [Section 4](#)). Depending on the attestation process, it is possible for attestation to run alongside the issuance protocol, e.g., where Clients send necessary attestation information to the Attester along with their token request. If the attestation process fails, the Client receives an error and issuance aborts without a token.
5. The Issuer generates a token response based on the token request, which is returned to the Client (generally via the Attester). Upon receiving the token response, the Client computes a token from the token challenge and token response. This token can be validated by anyone with the per-Issuer key, but cannot be linked to the content of the token request or token response.
6. If the Client has a token, it includes it in a subsequent request to the Origin, as authorization. This token is sent only once in reaction to a challenge; Clients do not send tokens more than once, even if they receive duplicate or redundant challenges. The Origin validates that the token was generated by

the expected Issuer and has not already been redeemed for the corresponding token challenge. If the Client does not have a token, perhaps because issuance failed, the client does not reply to the Origin's challenge with a new request.

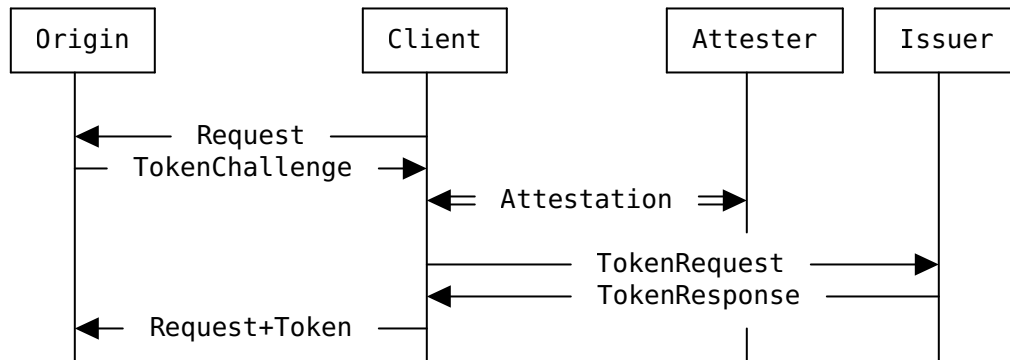


Figure 1: Privacy pass redemption and issuance protocol interaction

3.2. Use Cases

Use cases for Privacy Pass are broad and depend greatly on the deployment model as discussed in [Section 4](#). The initial motivating use case for Privacy Pass [[PrivacyPassCloudflare](#)] was to help rate-limit malicious or otherwise abusive traffic from services such as Tor [[DMS2004](#)]. The generalized and evolved architecture described in this document also work for this use case. However, for added clarity, some more possible use cases are described below.

*Low-quality, anti-fraud signal for open Internet services. Tokens can attest that the Client behind the user agent is likely not a bot attempting to perform some form of automated attack such as credential stuffing. Example attestation procedures for this use case might be solving some form of CAPTCHA or presenting evidence of a valid, unlocked device in good standing.

*Privacy-preserving authentication for proprietary services. Tokens can attest that the Client is a valid subscriber for a proprietary service, such as a deployment of Oblivious HTTP [[OHTTP](#)].

3.3. Privacy Goals and Threat Model

The end-to-end flow for Privacy Pass described in [Section 3.1](#) involves three different types of contexts:

Redemption context: The interactions and set of information shared between the Client and Origin, i.e., the information that is provided or otherwise available to the Origin during redemption

that might be used to identify a Client and construct a token challenge. This context includes all information associated with redemption, such as the timestamp of the event, Client-visible information (including the IP address), and the Origin name.

Issuance context: The interactions and set of information shared between the Client, Attester, and Issuer, i.e., the information that is provided or otherwise available to Attester and Issuer during issuance that might be used to identify a Client. This context includes all information associated with issuance, such as the timestamp of the event, any Client-visible information (including the IP address), and the Origin name (if revealed during issuance). This does not include the token challenge in its entirety, as that is kept secret from the Issuer during the issuance protocol.

Attestation context: The interactions and set of information shared between the Client and Attester only, for the purposes of attesting the validity of the Client, that is provided or otherwise available during attestation that might be used to identify the Client. This context includes all information associated with attestation, such as the timestamp of the event and any Client-visible information, including information needed for the attestation procedure to complete.

The privacy goals of Privacy Pass assume a threat model in which Origins trust specific Issuers to produce tokens, and Issuers in turn trust one or more Attesters to correctly run the attestation procedure with Clients. This arrangement ensures that tokens which validate for a given Issuer were only issued to a Client that successfully completed attestation with an Attester that the Issuer trusts. Moreover, this arrangement means that if an Origin accepts tokens issued by an Issuer that trusts multiple Attesters, then a Client can use any one of these Attesters to issue and redeem tokens for the Origin. Whether or not these different entities in the architecture collude for learning redemption, issuance, or attestation contexts, as well as the necessary preconditions for context unlinkability, depends on the deployment model; see [Section 4](#) for more details.

The mechanisms for establishing trust between each entity in this arrangement are deployment-specific. For example, in settings where Clients interact with Issuers through an Attester, Attesters and Issuers might use mutually authenticated TLS to authenticate one another. In settings where Clients do not communicate with Issuers through an Attester, the Attesters might convey this trust via a digital signature that Issuers can verify.

Clients explicitly trust Attesters to perform attestation correctly and in a way that does not violate their privacy. In particular, this means that Attesters which may be privy to private information about Clients are trusted to not disclose this information to non-colluding parties. Colluding parties are assumed to have access to the same information; see [Section 4](#) for more about different deployment models and non-collusion assumptions. However, Clients assume Issuers and Origins are malicious.

Given this threat model, the privacy goals of Privacy Pass are oriented around unlinkability based on redemption, issuance, and attestation contexts, as described below.

1. Origin-Client unlinkability. This means that given two redemption contexts, the Origin cannot determine if both redemption contexts correspond to the same Client or two different Clients. Informally, this means that a Client in a redemption context is indistinguishable from any other Client that might use the same redemption context. The set of Clients that share the same redemption context is referred to as a redemption anonymity set.
2. Issuer-Client unlinkability. This is similar to Origin-Client unlinkability in that a Client in an issuance context is indistinguishable from any other Client that might use the same issuance context. The set of Clients that share the same issuance context is referred to as an issuance anonymity set.
3. Attester-Origin unlinkability. This is similar to Origin-Client and Issuer-Client unlinkability. It means that given two attestation contexts, the Attester cannot determine if both contexts correspond to the same Origin or two different Origins. The set of Origins that share the same attestation context is referred to as an attestation anonymity set.
4. Redemption context unlinkability. Given two redemption contexts, the Origin cannot determine which issuance and attestation contexts each redemption corresponds to, even with the collaboration of the Issuer and Attester (should these be different parties). This means that any information that may be learned about the Client during the issuance and attestation flows cannot be used by the Origin to compromise Client privacy.

These unlinkability properties ensure that only the Client is able to correlate information that might be used to identify them with activity on the Origin. The Attester, Issuer, and Origin only receive the information necessary to perform their respective functions.

The manner in which these unlinkability properties are achieved depends on the deployment model, type of attestation, and issuance protocol details. For example, as discussed in [Section 4](#), in some cases it is necessary to use an anonymization service such as Tor [[DMS2004](#)] which hides Clients IP addresses. In general, anonymization services ensures that all Clients which use the service are indistinguishable from one another, though in practice there may be small distinguishing features (TLS fingerprints, HTTP headers, etc). Moreover, Clients generally trust these services to not disclose private Client information (such as IP addresses) to untrusted parties. Failure to use an anonymization service when interacting with Attesters, Issuers, or Origins can allow the set of possible Clients to be partitioned by the Client's IP address, and can therefore lead to unlinkability violations. Similarly, malicious Origins may attempt to link two redemption contexts together by using Client-specific Issuer public keys. See [Section 5](#) and [Section 6](#) for more information.

The remainder of this section describes the functional properties and security requirements of the redemption and issuance protocols in more detail. [Section 3.6](#) describes how information flows between Issuer, Origin, Client, and Attester through these protocols.

3.4. Redemption Protocol

The Privacy Pass redemption protocol, described in [[AUTHSCHEME](#)], is an authorization protocol wherein Clients present tokens to Origins for authorization. Normally, redemption is preceded by a challenge, wherein the Origin challenges Clients for a token with a TokenChallenge ([\[AUTHSCHEME\]](#), [Section 2.1](#)) and, if possible, Clients present a valid Token ([\[AUTHSCHEME\]](#), [Section 2.2](#)) in reaction to the challenge. This interaction is shown below.




Figure 2: Challenge and redemption protocol interaction

Alternatively, when configured to do so, Clients may opportunistically present Token values to Origins without a corresponding TokenChallenge.

The structure and semantics of the TokenChallenge and Token messages depend on the issuance protocol and token type being used; see [\[AUTHSCHEME\]](#) for more information.

The challenge provides the client with the information necessary to obtain tokens that the server might subsequently accept in the redemption context. There are a number of ways in which the token may vary based on this challenge, including:

- *Issuance protocol. The challenge identifies the type of issuance protocol required for producing the token. Different issuance protocols have different security properties, e.g., some issuance protocols may produce tokens that are publicly verifiable, whereas others may not have this property.

- *Issuer identity. Token challenges identify which Issuers are trusted for a given issuance protocol. As described in [Section 3.3](#), the choice of Issuer determines the type of Attesters and attestation procedures possible for a token from the specified Issuer, but the Origin does not learn exactly which Attester was used for a given token issuance event.

- *Redemption context. Challenges can be bound to a given redemption context, which influences a client's ability to pre-fetch and cache tokens. For example, an empty redemption context always allows tokens to be issued and redeemed non-interactively, whereas a fresh and random redemption context means that the redeemed token must be issued only after the client receives the challenge. See [Section 2.1.1](#) of [\[AUTHSCHEME\]](#) for more details.

- *Per-Origin or cross-Origin. Challenges can be constrained to the Origin for which the challenge originated (referred to as per-Origin tokens), or can be used across multiple Origins (referred to as cross-Origin tokens). The set of Origins for which a cross-Origin token is applicable is referred to as the cross-Origin set. Opting into this set is done by explicitly agreeing on the contents of the set. Every Origin in a cross-Origin set, by opting in, agrees to admit tokens for any other Origin in the set. See [Section 2.1.1](#) of [\[AUTHSCHEME\]](#) for more information on how this set is created.

Origins that admit cross-Origin tokens bear some risk of allowing tokens issued for one Origin to be spent in an interaction with another Origin. In particular, cross-Origin tokens issued based on a challenge for one Origin can be redeemed at another Origin in the cross-Origin set, which can make it difficult to regulate token consumption. Depending on the use case, Origins may need to maintain state to track redeemed tokens. For example, Origins that accept cross-Origin tokens across shared redemption contexts SHOULD track

which tokens have been redeemed already in those redemption contexts, since these tokens can be issued and then spent multiple times for any such challenge. Note that Clients which redeem the same token to multiple Origins do risk those Origins being able to link Client activity together, which can disincentivize this behavior. See [Section 2.1.1](#) of [\[AUTHSCHEME\]](#) for discussion.

How Clients respond to token challenges can have privacy implications. For example, if an Origin allows the Client to choose an Issuer, then the choice of Issuer can reveal information about the Client used to partition anonymity sets; see [Section 6.2](#) for more information about these privacy considerations.

3.5. Issuance Protocol

The Privacy Pass issuance protocol, described in [\[ISSUANCE\]](#), is a two-message protocol that takes as input a TokenChallenge from the redemption protocol ([\[AUTHSCHEME\]](#), [Section 2.1](#)) and produces a Token ([\[AUTHSCHEME\]](#), [Section 2.2](#)), as shown in [Figure 1](#).

The structure and semantics of the TokenRequest and TokenResponse messages depend on the issuance protocol and token type being used; see [\[ISSUANCE\]](#) for more information.

Clients interact with the Attester and Issuer to produce a token for a challenge. The context in which an Attester vouches for a Client during issuance is referred to as the attestation context. This context includes all information associated with the issuance event, such as the timestamp of the event and Client-visible information, including the IP address or other information specific to the type of attestation done.

Each issuance protocol may be different, e.g., in the number and types of participants, underlying cryptographic constructions used when issuing tokens, and even privacy properties.

Clients initiate the issuance protocol using the token challenge, a randomly generated nonce, and public key for the Issuer, all of which are the Client's private input to the protocol and ultimately bound to an output Token; see [Section 2.2](#) of [\[AUTHSCHEME\]](#) for details. Future specifications may change or extend the Client's input to the issuance protocol to produce Tokens with a different structure.

Token properties vary based on the issuance protocol. Important properties supported in this architecture are described below.

1. Public verifiability. This means the Token can be verified using the Issuer public key. A Token that does not have public verifiability can only be verified using the Issuer secret key.

2. Public metadata. This means that the Token can be cryptographically bound to arbitrary public information. See [Section 6.1](#) for privacy considerations of public metadata.
3. Private metadata. This means that the Token can be cryptographically bound to arbitrary private information, i.e., information that the Client does not observe during Token issuance or redemption. See [Section 6.1](#) for privacy considerations of private metadata.

The issuance protocol itself can be any interactive protocol between Client, Issuer, or other parties that produces a valid token bound to the Client's private input, subject to the following security requirements.

1. Unconditional input secrecy. The issuance protocol MUST NOT reveal anything about the Client's private input, including the challenge and nonce, to the Attester or Issuer, regardless of the hardness assumptions of the underlying cryptographic protocol(s). This property is sometimes also referred to as blindness.
2. One-more forgery security. The issuance protocol MUST NOT allow malicious Clients or Attesters (acting as Clients) to forge tokens offline or otherwise without interacting with the Issuer directly.
3. Concurrent security. The issuance protocol MUST be safe to run concurrently with arbitrarily many Clients, Attesters and Issuers.

See [Section 3.5.4](#) for requirements on new issuance protocol variants and related extensions.

In the sections below, we describe the Attester and Issuer roles in more detail.

3.5.1. Attester Role

In Privacy Pass, attestation is the process by which an Attester bears witness to, confirms, or authenticates a Client so as to verify properties about the Client that are required for Issuance. Issuers trust Attesters to perform attestation correctly, i.e., to implement attestation procedures in a way that are not subverted or bypassed by malicious Clients.

[[RFC9334](#)] describes an architecture for attestation procedures. Using that architecture as a conceptual basis, Clients are RATS attesters that produce attestation evidence, and Attesters are RATS verifiers that appraise the validity of attestation evidence.

The type of attestation procedure is a deployment-specific option and outside the scope of the issuance protocol. Example attestation procedures are below.

- *Solving a CAPTCHA. Clients that solve CAPTCHA challenges can be attested to have this capability for the purpose of being ruled out as a bot or otherwise automated Client.

- *Presenting evidence of Client device validity. Some Clients run on trusted hardware that are capable of producing device-level attestation evidence.

- *Proving properties about Client state. Clients can be associated with state and the Attester can verify this state. Examples of state include the Client's geographic region and whether the Client has a valid application-layer account.

Attesters may support different types of attestation procedures.

Each attestation procedure has different security properties. For example, attesting to having a valid account is different from attesting to running on trusted hardware. Supporting multiple attestation procedures is an important step towards ensuring equitable access for Clients; see [Section 5.1](#).

The role of the Attester in the issuance protocol and its impact on privacy depends on the type of attestation procedure, issuance protocol, and deployment model. For instance, increasing the number of required attestation procedures could decrease the overall anonymity set size. As an example, the number of Clients that have solved a CAPTCHA in the past day, that have a valid account, and that are running on a trusted device is less than the number of Clients that have solved a CAPTCHA in the past day. See [Section 6.2](#) for more discussion of how the variety of attestation procedures can negatively impact Client privacy.

Depending on the issuance protocol, the Issuer might learn information about the Origin. To ensure Issuer-Client unlinkability, the Issuer should be unable to link that information to a specific Client. For such issuance protocols where the Attester has access to Client-specific information, such as is the case for attestation procedures that involve Client-specific information (such as application-layer account information) or for deployment models where the Attester learns Client-specific information (such as Client IP addresses), Clients trust the Attester to not share any Client-specific information with the Issuer. In deployments where the Attester does not learn Client-specific information, or where the Attester and Issuer are operated by the same entity (such as in the

Joint Attester and Issuer model described in [Section 4.2](#)), the Client does not need to explicitly trust the Attester in this regard.

Issuers trust Attesters to correctly and reliably perform attestation. However, certain types of attestation can vary in value over time, e.g., if the attestation procedure is compromised. Broken attestation procedures are considered exceptional events and require configuration changes to address the underlying cause. For example, if an attestation procedure is compromised or subverted because of a zero-day exploit on devices that implement the attestation procedure, then the corresponding attestation procedure should be untrusted until the exploit is patched. Addressing changes in attestation quality is therefore a deployment-specific task. In Split Attester and Issuer deployments (see [Section 4.4](#)), Issuers can choose to remove compromised Attesters from their trusted set until the compromise is patched.

From the perspective of an Origin, tokens produced by an Issuer with at least one compromised Attester cannot be trusted assuming the Origin does not know which attestation procedure was used for issuance. This is because the Origin cannot distinguish between tokens that were issued via compromised Attesters and tokens that were issued via uncompromised Attesters absent some distinguishing information in the tokens themselves or from the Issuer. As a result, until the attestation procedure is fixed, the Issuer cannot be trusted by Origins. Moreover, as a consequence, any tokens issued by an Issuer with a compromised attester may no longer be trusted by Origins, even if those tokens were issued to Clients interacting with an uncompromised Attester.

3.5.2. Issuer Role

In Privacy Pass, the Issuer is responsible for completing the issuance protocol for Clients that complete attestation through a trusted Attester. As described in [Section 3.5.1](#), Issuers explicitly trust Attesters to correctly and reliably perform attestation. Origins explicitly trust Issuers to only issue tokens from trusted Attesters. Clients do not explicitly trust Issuers.

Depending on the deployment model case, issuance may require some form of Client anonymization service, similar to an IP-hiding proxy, so that Issuers cannot learn information about Clients. This can be provided by an explicit participant in the issuance protocol, or it can be provided via external means, such as through the use of an IP-hiding proxy service like Tor [[DMS2004](#)]. In general, Clients SHOULD minimize or remove identifying information where possible when invoking the issuance protocol.

Issuers are uniquely identifiable by all Clients with a consistent identifier. In a web context, this identifier might be the Issuer host name. Issuers maintain one or more configurations, including issuance key pairs, for use in the issuance protocol. Each configuration is assumed to have a unique and canonical identifier, sometimes referred to as a key identifier or key ID. Issuers can rotate these configurations as needed to mitigate risk of compromise; see [Section 6.2](#) for more considerations around configuration rotation. The Issuer public key for each active configuration is made available to Origins and Clients for use in the issuance and redemption protocols.

3.5.3. Issuance Metadata

Certain instantiations of the issuance protocol may permit public or private metadata to be cryptographically bound to a token. As an example, one trivial way to include public metadata is to assign a unique Issuer public key for each value of metadata, such that N keys yields $\log_2(N)$ bits of metadata. Metadata may be public or private.

Public metadata is that which clients can observe as part of the token issuance flow. Public metadata can either be transparent or opaque. For example, transparent public metadata is a value that the client either generates itself, or the Issuer provides during the issuance flow and the client can check for correctness. Opaque public metadata is metadata the client can see but cannot check for correctness. As an example, the opaque public metadata might be a "fraud detection signal", computed on behalf of the Issuer, during token issuance. Generally speaking, Clients cannot determine if this value is generated in a way that permits tracking.

Private metadata is that which Clients cannot observe as part of the token issuance flow. Such instantiations can be built on the Private Metadata Bit construction from Kreuter et al. [[KLOR20](#)] or the attribute-based VOPRF from Huang et al. [[HIJK21](#)].

Metadata can be arbitrarily long or bounded in length. The amount of permitted metadata may be determined by application or by the underlying cryptographic protocol. The total amount of metadata bits included in a token is the sum of public and private metadata bits. Every bit of metadata can be used to partition the Client issuance or redemption anonymity sets; see [Section 6.1](#) for more information.

3.5.4. Future Issuance Protocol Requirements

The Privacy Pass architecture and ecosystem are both intended to be receptive to extensions that expand the current set of

functionalities through new issuance protocols. Each new issuance protocol and extension MUST adhere to the following requirements:

1. Include a detailed analysis of the privacy impacts of the extension, why these impacts are justified, and guidelines on how to use the protocol to mitigate or minimize negative deployment or privacy consequences discussed in [Section 5](#) and [Section 6](#), respectively.
2. Adhere to the guidelines specified in [Section 3.5.2](#) for managing Issuer public key data.
3. Clearly specify how to interpret and validate TokenChallenge and Token messages that are exchanged during the redemption protocol.

3.6. Information Flow

The end-to-end process of redemption and issuance protocols involves information flowing between Issuer, Origin, Client, and Attester. That information can have implications on the privacy goals that Privacy Pass aims to provide as outlined in [Section 3.3](#). In this section, we describe the flow of information between each party. How this information affects the privacy goals in particular deployment models is further discussed in [Section 4](#).

3.6.1. Token Challenge Flow

To use Privacy Pass, Origins choose an Issuer from which they are willing to accept tokens. Origins then construct a token challenge using this specified Issuer and information from the redemption context it shares with the Client. This token challenge is then delivered to a Client. The token challenge conveys information about the Issuer and the redemption context, such as whether the Origin desires a per-Origin or cross-Origin token. Any entity that sees the token challenge might learn things about the Client as known to the Origin. This is why input secrecy is a requirement for issuance protocols, as it ensures that the challenge is not directly available to the Issuer.

3.6.2. Attestation Flow

Clients interact with the Attester to prove that they meet some required set of properties. In doing so, Clients contribute information to the attestation context, which might include sensitive information such as application-layer identities, IP addresses, and so on. Clients can choose whether or not to contribute this information based on local policy or preference.

3.6.3. Issuance Flow

Clients use the issuance protocol to produce a token bound to a token challenge. In doing so, there are several ways in which the issuance protocol contributes information to the attestation or issuance contexts. For example, a token request may contribute information to the attestation or issuance contexts as described below.

*Issuance protocol. The type of issuance protocol can contribute information about the Issuer's capabilities to the attestation or issuance contexts, as well as the capabilities of a given Client. For example, if a Client is presented with multiple issuance protocol options, then the choice of which issuance protocol to use can contribute information about the Client's capabilities.

*Issuer configuration. Information about the Issuer configuration, such as its identity or the public key used to validate tokens it creates, can be revealed during issuance and contribute to the attestation or issuance contexts.

*Attestation information. The issuance protocol can contribute information to the attestation or issuance contexts based on what attestation procedure the Issuer uses to trust a token request. In particular, a token request that is validated by a given Attester means that the Client which generated the token request must be capable of the completing the designated attestation procedure.

*Origin information. The issuance protocol can contribute information about the Origin that challenged the Client in [Section 3.6.1](#). In particular, a token request designated for a specific Issuer might imply that the resulting token is for an Origin that trusts the specified Issuer. However, this is not always true, as some token requests can correspond to cross-Origin tokens, i.e., they are tokens that would be accepted at any Origin that accepts the cross-Origin token.

Moreover, a token may contribute information to the issuance attestation or contexts as described below.

*Origin information. The issuance protocol can contribute information about the Origin in how it responds to a token request. For example, if an Issuer learns the Origin during issuance and is also configured to respond in some way on the basis of that information, and the Client interacts with the Issuer transitively through the Attester, that response can reveal information to the Attester.

*Token. The token produced by the issuance protocol can contain information from the issuance context. In particular, depending on the issuance protocol, tokens can contain public or private

metadata, and Issuers can choose that metadata on the basis of information in the issuance context.

Exceptional cases in the issuance protocol, such as when either the Attester or Issuer aborts the protocol, can contribute information to the attestation or issuance contexts. The extent to which information in this context harms the Issuer-Client or Attester-Origin unlinkability goals in [Section 3.3](#) depends on deployment model; see [Section 4](#). Clients can choose whether or not to contribute information to these contexts based on local policy or preference.

3.6.4. Token Redemption Flow

Clients send tokens to Origins during the redemption protocol. Any information that is added to the token during issuance can therefore be sent to the Origin. Information can either be explicitly passed in a token, or it can be implicit in the way the Client responds to a token challenge. For example, if a Client fails to complete issuance, and consequently fails to redeem a token for a token challenge, this can reveal information to the Origin that it might not otherwise have access to. However, an Origin cannot necessarily distinguish between a Client that fails to complete issuance and one that ignores the token challenge altogether.

4. Deployment Models

The Origin, Attester, and Issuer portrayed in [Figure 1](#) can be instantiated and deployed in a number of ways. The deployment model directly influences the manner in which attestation, issuance, and redemption contexts are separated to achieve Origin-Client, Issuer-Client, and Attester-Origin unlinkability.

This section covers some expected deployment models and their corresponding security and privacy considerations. Each deployment model is described in terms of the trust relationships and communication patterns between Client, Attester, Issuer, and Origin. Entities drawn within the same bounding box are assumed to be operated by the same party and are therefore able to collude. Collusion can enable linking of attestation, issuance, and redemption contexts. Entities not drawn within the same bounding box are assumed to not collude, meaning that entities operated by separate parties that do not collude. Mechanisms for enforcing non-collusion are out of scope for this architecture.

4.1. Shared Origin, Attester, Issuer

In this model, the Origin, Attester, and Issuer are all operated by the same entity, as shown in [Figure 3](#).

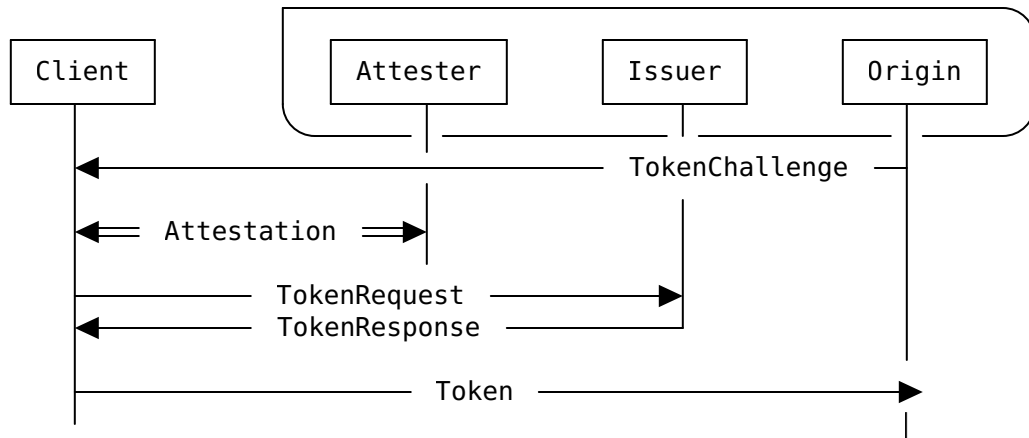


Figure 3: Shared Deployment Model

This model represents the initial deployment of Privacy Pass, as described in [[PrivacyPassCloudflare](#)]. In this model, the Attester, Issuer, and Origin share the attestation, issuance, and redemption contexts. As a result, attestation mechanisms that can uniquely identify a Client, e.g., requiring that Clients authenticate with some type of application-layer account, are not appropriate, as they could lead to unlinkability violations.

Origin-Client, Issuer-Client, and Attester-Origin unlinkability requires that issuance and redemption events be separated over time, such as through the use of tokens that correspond to token challenges with an empty redemption context (see [Section 3.4](#)), or be separated over space, such as through the use of an anonymizing service when connecting to the Origin.

4.2. Joint Attester and Issuer

In this model, the Attester and Issuer are operated by the same entity that is separate from the Origin. The Origin trusts the joint Attester and Issuer to perform attestation and issue Tokens. Clients interact with the joint Attester and Issuer for attestation and issuance. This arrangement is shown in [Figure 4](#).

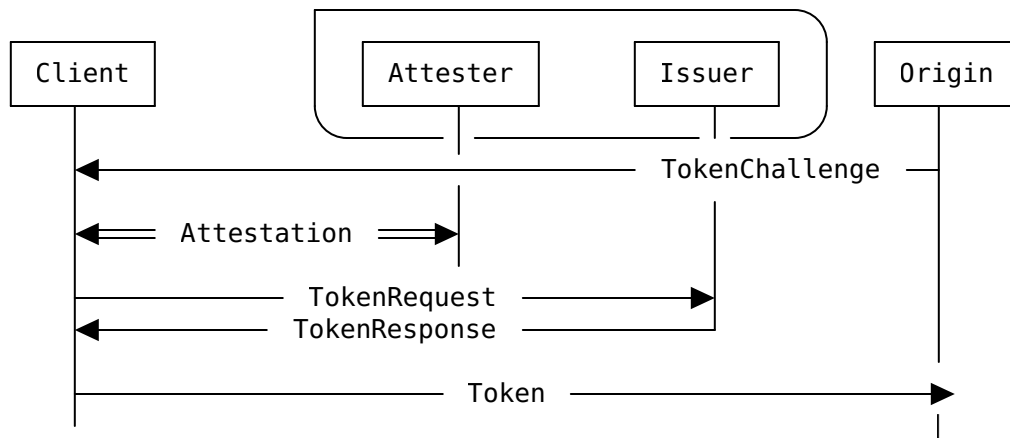


Figure 4: Joint Attester and Issuer Deployment Model

This model is useful if an Origin wants to offload attestation and issuance to a trusted entity. In this model, the Attester and Issuer share an attestation and issuance context for the Client, which is separate from the Origin's redemption context.

Similar to the Shared Origin, Attester, Issuer model, Issuer-Client and Origin-Client unlinkability in this model requires issuance and redemption events, respectively, be separated over time or space. Attester-Origin unlinkability requires that the Attester and Issuer do not learn the Origin for a particular token request. For this reason, issuance protocols that require the Issuer to learn information about the Origin, such as that which is described in [RATE-LIMITED], are not appropriate since they could lead to Attester-Origin unlinkability violations through the Origin name.

4.3. Joint Origin and Issuer

In this model, the Origin and Issuer are operated by the same entity, separate from the Attester, as shown in the figure below. The Issuer accepts token requests that come from trusted Attesters. Since the Attester and Issuer are separate entities, this model requires some mechanism by which Issuers establish trust in the Attester (as described in [Section 3.3](#)). For example, in settings where the Attester is a Client-trusted service that directly communicates with the Issuer, one way to establish this trust is via mutually-authenticated TLS. However, alternative authentication mechanisms are possible. This arrangement is shown in [Figure 5](#).

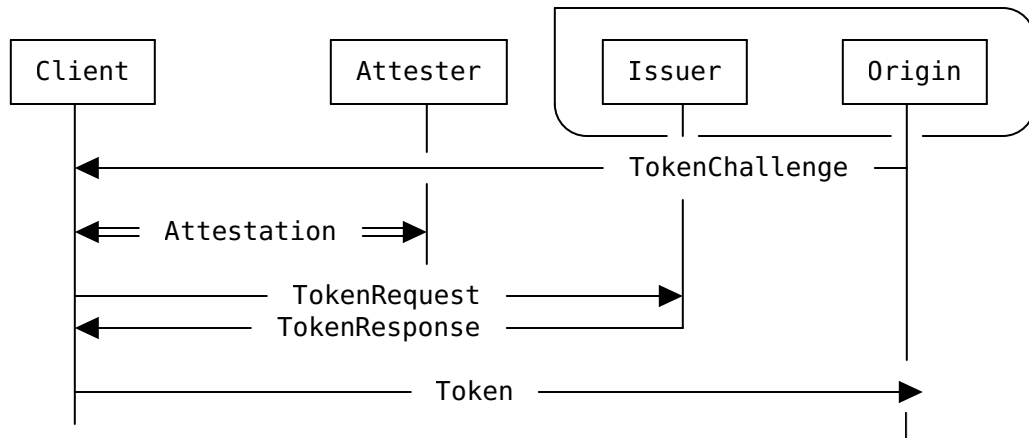


Figure 5: Joint Origin and Issuer Deployment Model

This model is useful for Origins that require Client-identifying attestation, e.g., through the use of application-layer account information, but do not otherwise want to learn information about individual Clients beyond what is observed during the token redemption, such as Client IP addresses.

In this model, attestation contexts are separate from issuer and redemption contexts. As a result, any type of attestation is suitable in this model.

Moreover, any type of token challenge is suitable assuming there is more than one Origin involved, since no single party will have access to the identifying Client information and unique Origin information. Issuers that produce tokens for a single Origin are not suitable in this model since an Attester can infer the Origin from a token request, as described in [Section 3.6.3](#). However, since the issuance protocol provides input secrecy, the Attester does not learn details about the corresponding token challenge, such as whether the token challenge is per-Origin or cross-Origin.

4.4. Split Origin, Attester, Issuer

In this model, the Origin, Attester, and Issuer are all operated by different entities. As with the joint Origin and Issuer model, the Issuer accepts token requests that come from trusted Attesters, and the details of that trust establishment depend on the issuance protocol and relationship between Attester and Issuer; see [Section 3.3](#). This arrangement is shown in [Figure 1](#).

This is the most general deployment model, and is necessary for some types of issuance protocols where the Attester plays a role in token issuance; see [\[RATE-LIMITED\]](#) for one such type of issuance protocol.

In this model, the Attester, Issuer, and Origin have a separate view of the Client: the Attester sees potentially sensitive Client identifying information, such as account identifiers or IP addresses, the Issuer sees only the information necessary for issuance, and the Origin sees token challenges, corresponding tokens, and Client source information, such as their IP address. As a result, attestation, issuance, and redemption contexts are separate, and therefore any type of token challenge is suitable in this model as long as there is more than a single Origin.

As in the Joint Origin and Issuer model in [Section 4.3](#), and as described in [Section 3.6.3](#), if the Issuer produces tokens for a single Origin, then per-Origin tokens are not appropriate since the Attester can infer the Origin from a token request.

5. Deployment Considerations

[Section 4](#) discusses deployment models that are possible in practice. Beyond possible implications on security and privacy properties of the resulting system, Privacy Pass deployments can impact the overall ecosystem in two important ways: (1) discriminatory treatment of Clients and the gated access to otherwise open services, and (2) centralization. This section describes considerations relevant to these topics.

5.1. Discriminatory Treatment

Origins can use tokens as a signal for distinguishing between Clients that are capable of completing attestation with one Attester trusted by the Origin's chosen Issuer, and Clients that are not capable of doing the same. A consequence of this is that Privacy Pass could enable discriminatory treatment of Clients based on Attestation support. For example, an Origin could only authorize Clients that successfully authenticate with a token, prohibiting access to all other Clients.

The type of attestation procedures supported for a particular deployment depends greatly on the use case. For example, consider a proprietary deployment of Privacy Pass that authorizes clients to access a resource such as an anonymization service. In this context, it is reasonable to support specific types of attestation procedures that demonstrate Clients can access the resource, such as with an account or specific type of device. However, in open deployments of Privacy Pass that are used to safeguard access to otherwise open or publicly accessible resources, diversity in attestation procedures is critically important so as to not discriminate against Clients that choose certain software, hardware, or identity providers.

In principle, Issuers should strive to mitigate discriminatory behavior by providing equitable access to all Clients. This can be done by working with a set of Attesters that are suitable for all Clients. In practice, this may require tradeoffs in what type of attestation Issuers are willing to trust so as to enable more widespread support. In other words, trusting a variety of Attesters with a diverse set of attestation procedures would presumably increase support among Clients, though most likely at the expense of decreasing the overall value of tokens issued by the Issuer.

For example, to disallow discriminatory behavior between Clients with and without device attestation support, an Issuer may wish to support Attesters that support CAPTCHA-based attestation. This means that the overall attestation value of a Privacy Pass token is bound by the difficulty in spoofing or bypassing either one of these attestation procedures.

5.2. Centralization

A consequence of limiting the number of participants (Attesters or Issuers) in Privacy Pass deployments for meaningful privacy is that it forces concentrated centralization amongst those participants. [\[CENTRALIZATION\]](#) discusses several ways in which this might be mitigated. For example, a multi-stakeholder governance model could be established to determine what candidate participants are fit to operate as participants in a Privacy Pass deployment. This is precisely the system used to control the Web's trust model.

Alternatively, Privacy Pass deployments might mitigate this problem through implementation. For example, rather than centralize the role of attestation in one or few entities, attestation could be a distributed function performed by a quorum of many parties, provided that neither Issuers nor Origins learn which Attester implementations were chosen. As a result, Clients could have more opportunities to switch between attestation participants.

6. Privacy Considerations

The previous section discusses the impact of deployment details on Origin-Client, Issuer-Client, and Attester-Origin unlinkability. The value these properties affords to end users depends on the size of anonymity sets in which Clients or Origins are unlinkable. For example, consider two different deployments, one wherein there exists a redemption anonymity set of size two and another wherein there is a redemption anonymity set of size 2^{32} . Although Origin-Client unlinkability guarantees that the Origin cannot link any two requests to the same Client based on these contexts, respectively, the probability of determining the "true" Client is higher the smaller these sets become.

In practice, there are a number of ways in which the size of anonymity sets may be reduced or partitioned, though they all center around the concept of consistency. In particular, by definition, all Clients in an anonymity set share a consistent view of information needed to run the issuance and redemption protocols. An example type of information needed to run these protocols is the Issuer public key. When two Clients have inconsistent information, these Clients effectively have different redemption contexts and therefore belong in different anonymity sets.

The following sections discuss issues that can influence anonymity set size. For each issue, we discuss mitigations or safeguards to protect against the underlying problem.

6.1. Partitioning by Issuance Metadata

Any public or private metadata bits of information can be used to further segment the size of the Client's anonymity set. Any Issuer that wanted to track a single Client could add a single metadata bit to Client tokens. For the tracked Client it would set the bit to 1, and 0 otherwise. Adding additional bits provides an exponential increase in tracking granularity similarly to introducing more Issuers (though with more potential targeting).

For this reason, deployments should take the amount of metadata used by an Issuer in creating redemption tokens must into account -- together with the bits of information that Issuers may learn about Clients otherwise. Since this metadata may be useful for practical deployments of Privacy Pass, Issuers must balance this against the reduction in Client privacy.

The number of permitted metadata values often depends on deployment-specific details. In general, limiting the amount of metadata permitted helps limit the extent to which metadata can uniquely identify individual Clients. Failure to bound the number of possible metadata values can therefore lead to reduction in Client privacy. Most token types do not admit any metadata, so this bound is implicitly enforced.

6.2. Partitioning by Issuance Consistency

Anonymity sets can be partitioned by information used for the issuance protocol, including: metadata, Issuer configuration (keys), and Issuer selection.

Any issuance metadata bits of information can be used to partition the Client anonymity set. For example, any Issuer that wanted to track a single Client could add a single metadata bit to Client tokens. For the tracked Client it would set the bit to 1, and 0 otherwise. Adding additional bits provides an exponential increase in

tracking granularity similarly to introducing more Issuers (though with more potential targeting).

The number of active Issuer configurations also contributes to anonymity set partitioning. In particular, when an Issuer updates their configuration and the corresponding key pair, any Client that invokes the issuance protocol with this configuration becomes part of a set of Clients which also ran the issuance protocol using the same configuration. Issuer configuration updates, e.g., due to key rotation, are an important part of hedging against long-term private key compromise. In general, key rotations represent a trade-off between Client privacy and Issuer security. Therefore, it is important that key rotations occur on a regular cycle to reduce the harm of an Issuer key compromise.

Lastly, if Clients are willing to issue and redeem tokens from a large number of Issuers for a specific Origin, and that Origin accepts tokens from all Issuers, partitioning can occur. As an example, if a Client obtains tokens from many Issuers and an Origin later challenges that Client for a token from each Issuer, the Origin can learn information about the Client. This arrangement might happen if Clients request tokens from different Issuers, each of which have different Attester preferences. Each per-Issuer token that a Client holds essentially corresponds to a bit of information about the Client that Origin learns. Therefore, there is an exponential loss in privacy relative to the number of Issuers.

The fundamental problem here is that the number of possible issuance configurations, including the keys in use and the Issuer identities themselves, can partition the Client anonymity set. To mitigate this problem, Clients SHOULD bound the number of active issuance configurations per Origin as well as across Origins. Moreover, Clients SHOULD employ some form of consistency mechanism to ensure that they receive the same configuration information and are not being actively partitioned into smaller anonymity sets. See [\[CONSISTENCY\]](#) for possible consistency mechanisms. Depending on the deployment, the Attester might assist the Client in applying these consistency checks across clients. Failure to apply a consistency check can allow Client-specific keys to violate Origin-Client unlinkability.

6.3. Partitioning by Side-Channels

Side-channel attacks, such as those based on timing correlation, could be used to reduce anonymity set size. In particular, for interactive tokens that are bound to a Client-specific redemption context, the anonymity set of Clients during the issuance protocol consists of those Clients that started issuance between the time of the Origin's challenge and the corresponding token redemption.

Depending on the number of Clients using a particular Issuer during that time window, the set can be small. Applications should take such side channels into consideration before choosing a particular deployment model and type of token challenge and redemption context.

7. Security Considerations

This document describes security and privacy requirements for the Privacy Pass redemption and issuance protocols. It also describes deployment models built around non-collusion assumptions and privacy considerations for using Privacy Pass within those models. Ensuring Client privacy -- separation of attestation and redemption contexts -- requires active work on behalf of the Client, especially in the presence of malicious Issuers and Origins. Implementing mitigations discussed in [Section 4](#) and [Section 6](#) is therefore necessary to ensure that Privacy Pass offers meaningful privacy improvements to end-users.

7.1. Token Caching

Depending on the Origin's token challenge, Clients can request and cache more than one token using an issuance protocol. Cached tokens help improve privacy by separating the time of token issuance from the time of token redemption, and also allow Clients to reduce the overhead of receiving new tokens via the issuance protocol.

As a consequence, Origins that send token challenges which are compatible with cached tokens need to take precautions to ensure that tokens are not replayed. This is typically done via keeping track of tokens that are redeemed for the period of time in which cached tokens would be accepted for particular challenges.

Moreover, since tokens are not intrinsically bound to Clients, it is possible for malicious Clients to collude and share tokens in a so-called "hoarding attack." As an example of this attack, many distributed Clients could obtain cacheable tokens and then share them with a single Client to redeem in a way that would violate an Origin's attempt to limit tokens to any one particular Client. In general, mechanisms for mitigating hoarding attacks depend on the deployment model and use case. For example, in the Joint Origin and Issuer model, comparing the issuance and redemption contexts can help detect hoarding attacks, i.e., if the distribution of redemption contexts is noticeably different from the distribution of issuance contexts. Rate limiting issuance, either at the Client, Attester, or Issuer, can also help mitigate these attacks.

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Normative References

[AUTHSCHEME] Pauly, T., Valdez, S., and C. A. Wood, "The Privacy Pass HTTP Authentication Scheme", Work in Progress, Internet-Draft, draft-ietf-privacypass-auth-scheme-13, 12 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-auth-scheme-13>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/rfc/rfc2119>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

9.2. Informative References

[CENTRALIZATION]

Nottingham, M., "Centralization, Decentralization, and Internet Standards", Work in Progress, Internet-Draft, draft-nottingham-avoiding-internet-centralization-14, 14 September 2023, <<https://datatracker.ietf.org/doc/html/draft-nottingham-avoiding-internet-centralization-14>>.

[CONSISTENCY] Davidson, A., Finkel, M., Thomson, M., and C. A. Wood, "Key Consistency and Discovery", Work in Progress, Internet-Draft, draft-ietf-privacypass-key-consistency-01, 10 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-key-consistency-01>>.

[DMS2004] Dingledine, R., Mathewson, N., and P. Syverson, "Tor: The Second-Generation Onion Router", August 2004, <<https://svn.torproject.org/svn/projects/design-paper/tor-design.html>>.

[HIJK21]

Huang, S., Iyengar, S., Jeyaraman, S., Kushwah, S., Lee, C. K., Luo, Z., Mohassel, P., Raghunathan, A., Shaikh, S., Sung, Y. C., and A. Zhang, "PrivateStats: De-Identified Authenticated Logging at Scale", January 2021, <<https://research.fb.com/privatestats>>.

[ISSUANCE] Celi, S., Davidson, A., Valdez, S., and C. A. Wood, "Privacy Pass Issuance Protocol", Work in Progress, Internet-Draft, draft-ietf-privacypass-protocol-14, 14

September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-protocol-14>>.

- [KLOR20] Kreuter, B., Lepoint, T., Orrù, M., Raykova, M., and Springer International Publishing, "Anonymous Tokens with Private Metadata Bit", Advances in Cryptology – CRYPTO 2020, pp. 308-336, DOI 10.1007/978-3-030-56784-2_11, 2020, <https://doi.org/10.1007/978-3-030-56784-2_11>.
- [OHTTP] Thomson, M. and C. A. Wood, "Oblivious HTTP", Work in Progress, Internet-Draft, draft-ietf-ohai-ohttp-10, 25 August 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-ohai-ohttp-10>>.
- [PrivacyPassCloudflare] Sullivan, N., "Cloudflare Supports Privacy Pass", n.d., <<https://blog.cloudflare.com/cloudflare-supports-privacy-pass/>>.
- [PrivacyPassExtension] "Privacy Pass Browser Extension", n.d., <<https://github.com/privacypass/challenge-bypass-extension>>.
- [RATE-LIMITED] Hendrickson, S., Iyengar, J., Pauly, T., Valdez, S., and C. A. Wood, "Rate-Limited Token Issuance Protocol", Work in Progress, Internet-Draft, draft-privacypass-rate-limit-tokens-03, 6 July 2022, <<https://datatracker.ietf.org/doc/html/draft-privacypass-rate-limit-tokens-03>>.
- [RFC9334] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedures (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <<https://www.rfc-editor.org/rfc/rfc9334>>.

Appendix A. Acknowledgements

The authors would like to thank Eric Kinnear, Scott Hendrickson, Tommy Pauly, Christopher Patton, Benjamin Schwartz, Martin Thomson, Steven Valdez and other contributors of the Privacy Pass Working Group for many helpful contributions to this document.

Authors' Addresses

Alex Davidson
LIP
Lisbon
Portugal

Email: alex.davidson92@gmail.com

Jana Iyengar
Fastly

Email: jri@fastly.com

Christopher A. Wood
Cloudflare
101 Townsend St
San Francisco,
United States of America

Email: caw@heapingbits.net