

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 23 August 2021

A. Davidson
LIP
C.A. Wood
Cloudflare
19 February 2021

Privacy Pass Architectural Framework
draft-ietf-privacypass-architecture-01

Abstract

This document specifies the architectural framework for constructing secure and anonymity-preserving instantiations of the Privacy Pass protocol. It provides recommendations on how the protocol ecosystem should be constructed to ensure the privacy of clients, and the security of all participating entities.

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 23 August 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
3. Ecosystem participants	4
3.1. Servers	5
3.2. Clients	6
3.2.1. Client identifying information	6
4. Key management and discovery	6
4.1. Key rotation	7
4.2. Ciphersuites	8
5. Server running modes	8
5.1. Single-Verifier	8
5.2. Delegated-Verifier	8
5.3. Asynchronous-Verifier	9
5.4. Public-Verifier	9
5.5. Bounded number of servers	10
6. Metadata	10
6.1. Client privacy implications	11
7. Client-Server trust relationship	11
8. Privacy considerations	12
8.1. Server key rotation	12
8.2. Large numbers of servers	13
8.2.1. Allowing larger number of servers	14
8.3. Partitioning of server key material	14
8.4. Tracking and identity leakage	15
8.5. Client incentives for anonymity reduction	15
9. Security considerations	15
9.1. Double-spend protection	16
9.2. Token exhaustion	16
9.3. Avoiding server centralization	16
10. Protocol parametrization	16
10.1. Justification	17
10.2. Example parameterization	18
10.3. Allowing more servers	19
11. Extension integration policy	19
12. Existing applications	19
12.1. Cloudflare challenge pages	19
12.2. Trust Token API	20
12.3. Zero-knowledge Access Passes	20
12.4. Basic Attention Tokens	20
12.5. Token Based Services	20
13. References	20
13.1. Normative References	20
13.2. Informative References	21
Appendix A. Contributors	22
Authors' Addresses	22

1. Introduction

The Privacy Pass protocol provides an anonymity-preserving mechanism for authorization of clients with servers. The protocol is detailed in [I-D.ietf-privacypass-protocol] and is intended for use in the application-layer.

The way that the ecosystem around the protocol is set up can have significant impacts on the stated privacy and security guarantees of the protocol. For instance, the number of servers issuing Privacy Pass tokens, along with the number of registered clients, determines the anonymity set of each individual client. Moreover, this can be influenced by other factors, such as: the key rotation policy used by each server; and, the number of supported ciphersuites. There are also client behavior patterns that can reduce the effective security of the server.

In this document, we will provide a structural framework for building the ecosystem around the Privacy Pass protocol. The core of the document also includes policies for the following considerations.

- * How server key material should be managed and accessed.
- * Compatible server issuance and redemption running modes and associated expectations.
- * How clients should evaluate server trust relationships.
- * Security and privacy properties of the protocol.
- * A concrete assessment and parametrization of the privacy budget associated with different settings of the above policies.
- * The incorporation of potential extensions into the wider ecosystem.

Finally, we will discuss existing applications that make use of the Privacy Pass protocol, and highlight how these may fit with the proposed framework.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The following terms are used throughout this document.

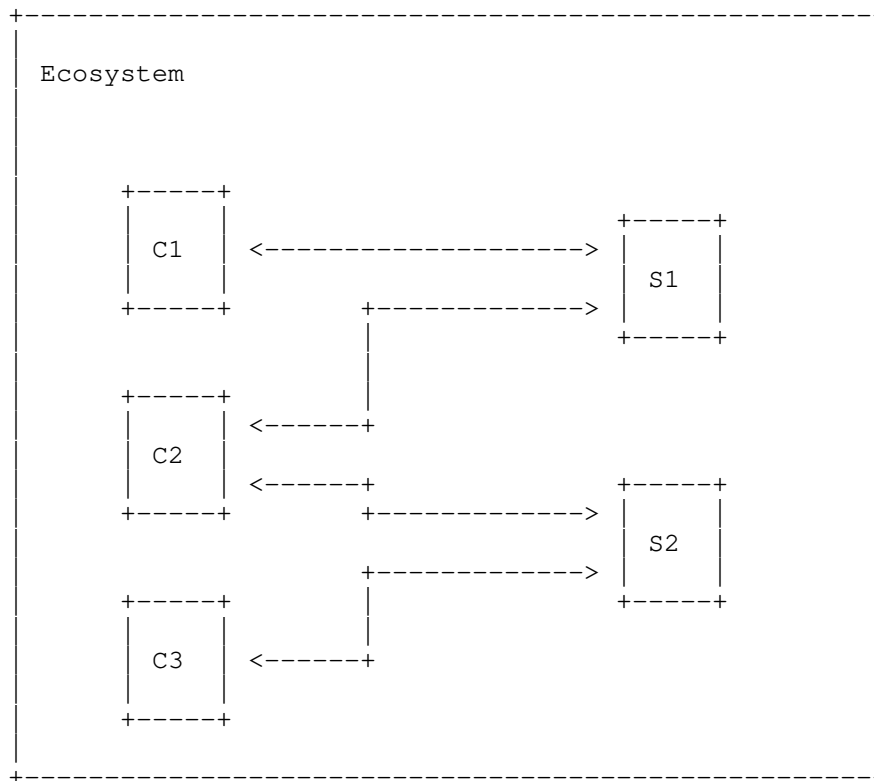
- * Server: An entity that issues anonymous tokens to clients. In symmetric verification cases, the server must also verify tokens. Also referred to as the server.
- * Client: An entity that seeks authorization from a server.

We assume that all protocol messages are encoded into raw byte format before being sent. We use the TLS presentation language [RFC8446] to describe the structure of the data that is communicated and stored.

3. Ecosystem participants

The Privacy Pass ecosystem refers to the global framework in which multiple instances of the Privacy Pass protocol operate. This refers to all servers that support the protocol, or any extension of it, along with all of the clients that may interact with these servers.

The ecosystem itself, and the way it is constructed, is critical for evaluating the privacy of each individual client. We assume that a client's privacy refers to fraction of users that it represents in the anonymity set that it belongs to. We discuss this more in Section 8.



In the above diagram, the arrows indicate the open channels between a client and a server. An open channel indicates that a client accepts Privacy Pass tokens from this server.

If no channel exists, this means that the client chooses not to accept tokens from (or redeem tokens with) that particular server. We discuss the roles of servers and clients further in Section 3.1 and Section 3.2, respectively.

3.1. Servers

Generally, servers in the Privacy Pass ecosystem are entities whose primary function is to undertake the role of the "server" in [I-D.ietf-privacypass-protocol]. To facilitate this, the server **MUST** hold a Privacy Pass protocol keypair at any given time. The server public key **MUST** be made available to all clients in such a way that key rotations and other updates are publicly visible. The server **MAY** also require additional state for ensuring this. We provide a wider discussion in Section 4.

Note that, in the core protocol instantiation from [I-D.ietf-privacypass-protocol], the redemption phase is a symmetric protocol. This means that the server is the same server that ultimately processes token redemptions from clients. However, plausible extensions to the protocol specification may allow public verification of tokens by entities which do not hold the secret Privacy Pass keying material. We highlight possible client and server configurations in Section 5.

The server must be uniquely identifiable by all clients with a consistent identifier.

3.2. Clients

Clients in the Privacy Pass ecosystem are entities whose primary function is to undertake the role of the "Client" in [I-D.ietf-privacypass-protocol]. Clients are assumed to only store data related to the tokens that it has been issued by the server. This storage is used for constructing redemption requests.

Clients MAY choose not to accept tokens from servers that they do not trust. See Section 7 for a wider discussion.

3.2.1. Client identifying information

Privacy properties of this protocol do not take into account other possibly identifying information available in an implementation, such as a client's IP address. Servers which monitor IP addresses may use this to track client redemption patterns over time. Clients cannot check whether servers monitor such identifying information. Thus, clients SHOULD minimize or remove identifying information where possible, e.g., by using anonymity-preserving tools such as Tor to interact with servers.

4. Key management and discovery

The key material and protocol configuration that a server uses to issue tokens corresponds to a number of different pieces of information.

- * The ciphersuite that the server is using.
- * The public keys that are active for the server.

The way that the server publishes and maintains this information impacts the effective privacy of the clients; see Section 8 for more details. The fundamental requirement for key management and discovery is that servers must be unable to target specific clients with unique keys without detection. There are a number of ways in which this might be implemented:

- * Servers use a verifiable, tamper-free registry from which clients discover keys. Similar to related mechanisms and protocols such as Certificate Transparency [RFC6962], this may require external auditors or additional client behavior to ensure the registry state is consistent for all clients.
- * Clients use an anonymity-preserving tool such as Tor to discover keys from multiple network vantage points. This is done to ensure consistent keys to seemingly different clients.
- * Clients embed server keys into software.

Specific mechanisms for key management and discovery are out of scope for this document.

4.1. Key rotation

Token issuance associates all issued tokens with a particular choice of key. If a server issues tokens with many keys, then this may harm the anonymity of the Client. For example, they would be able to map the Client's access patterns by inspecting which key each token they possess has been issued under.

To prevent against this, servers **MUST** only use one private key for issuing tokens at any given time. Servers **MAY** use one or more keys for redemption to allow servers for seamless key rotation.

Servers may rotate keys as a means of revoking tokens issued under old or otherwise expired keys. Alternatively, servers may include expiration information as metadata alongside the token; See Section 6 for more discussion about metadata constraints. Both techniques are equivalent since they cryptographically bind expiration to individual tokens.

Key rotations should be limited in frequency for similar reasons. See Section 10 for guidelines on what frequency of key rotations are permitted.

4.2. Ciphersuites

Since a server is only permitted to have a single active issuing key, this implies that only a single ciphersuite is allowed per issuance period. If a server wishes to change their ciphersuite, they **MUST** do so during a key rotation.

5. Server running modes

We provide an overview of some of the possible frameworks for configuring the way that servers run in the Privacy Pass ecosystem. In short, servers may be configured to provide symmetric issuance and redemption with clients. While some servers may be configured as proxies that accept Privacy Pass data and send it to another server that actually processes issuance or redemption data. Finally, we also consider instances of the protocol that may permit public verification.

The intention with providing each of these running modes is to cover the different applications that utilize variants of the Privacy Pass protocol. We **RECOMMEND** that any Privacy Pass server implementation adheres to one of these frameworks.

5.1. Single-Verifier

The simplest way of considering the Privacy Pass protocol is in a setting where the same server plays the role of server and verifier, we call this "Single-Verifier" (SV).

Let S be the server, and C be the client. When S wants to issue tokens to C , they invoke the issuance protocol where C generates their own inputs, and S uses their secret key sk_S . In this setting, C can only perform token redemption with S . When a token redemption is required, C and S invoke the redemption phase of the protocol, where C uses an issued token from a previous exchange, and S uses sk_S to validate the redemption.

5.2. Delegated-Verifier

In this setting, each client C obtains issued tokens from a server S via the issuance phase of the protocol. The difference is that C can prove that they hold a valid authorization with any verifier V . We still only consider S to hold their own secret key. We name this mode "Delegated-Verifier" (DV).

When C interacts with V , V can ask C to provide proof of authorization to the separate server S . The first stage of the redemption phase of the protocol is invoked between C and V , which

sees C send an unused redemption token to V. This message is then used in a redemption exchange between V and S, where V plays the role of the Client. Then S sends the result of the redemption verification to V, and V uses this result to determine whether C has a valid token.

5.3. Asynchronous-Verifier

This setting is inspired by recently proposed APIs such as [TrustTokenAPI]. It is similar to the DV configuration, except that the verifiers V no longer interact with the server S. Only C interacts with S, and this is done asynchronously to the authorization request from V. Hence "Asynchronous-Verifier" (AV).

When V invokes a redemption for C, C then invokes a redemption exchange with S in a separate session. If verification is carried out successfully by S, S instead returns a Signed Redemption Record (SRR) that contains the following information:

```
"result": {  
  "timestamp": "2019-10-09-11:06:11",  
  "verifier": "V",  
},  
"signature": sig,
```

The "signature" field carries a signature evaluated over the contents of "result" using a long-term signing key for the server S, of which the corresponding public key is well-known to C and V. This would need to be published alongside other public key data for S. Then C can prove that they hold a valid authorization from S to V by sending the SRR to V. The SRR can be verified by V by verifying the signature, using the well-known public key for S.

Such records can be cached to display again in the future. The server can also add an expiry date to the record to determine when the client must refresh the record.

5.4. Public-Verifier

We consider the case where client redemptions can be verified publicly using the server public key. This allows for defining extensions of Privacy Pass that use public-key cryptography to allow public verification.

In this case, the client C obtains a redemption token from S. The redemption token is publicly verifiable in the sense that any entity that knows the public key for S can verify the token. This running mode is known as "Public-Verifier" (PV).

5.5. Bounded number of servers

Each of the configurations above can be generalized to settings where a bounded number of servers are allowed, and verifiers can invoke authorization verification for any of the available servers.

As we will discuss later in Section 8, configuring a large number of servers can lead to privacy concerns for the clients in the ecosystem. Therefore, we are careful to ensure that the number of servers is kept strictly bounded. The actual servers can be replaced with different servers as long as the total never exceeds this bound. Moreover, server replacements also have an effect on client anonymity that is similar to when a key rotation occurs.

See Section 8 for more details about maintaining privacy with multiple servers.

6. Metadata

Certain instantiations of Privacy Pass 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. The total amount of metadata bits included in a token is the sum of public and private metadata bits. See Section 10 for discussion about metadata limits.

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 server 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 server, during token issuance. In normal circumstances, clients cannot determine if this value is correct or otherwise a tracking vector.

Private metadata is that which clients cannot observe as part of the token issuance flow. In [I-D.ietf-privacypass-protocol], it is possible to include private metadata to redemption tokens. The core protocol instantiation that is described does not include additional metadata. However, future instantiations may use this functionality to provide redemption verifiers with additional information about the user. Such instantiations may be built on the Private Metadata Bit construction from Kreuter et al. [KLOR20] or the attribute-based VOPRF from Huang et al. [HIJK21].

Metadata may also be arbitrarily long or bounded in length. The amount of permitted metadata may be determined by application or by the underlying cryptographic protocol.

6.1. Client privacy implications

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

For this reason, the amount of metadata used by a server in creating redemption tokens must be taken into account - together with the bits of information that server's may learn about clients otherwise. Since this metadata may be useful for practical deployments of Privacy Pass, servers must balance this against the reduction in client privacy. In general, servers should permit no more than 32 bits of metadata, as this can uniquely identify each possible user. We discuss this more in Section 10.

7. Client-Server trust relationship

It is important, based on the architecture above, that any client can determine whether it would like to interact with a given server in the ecosystem. Note that this decision must be taken before a client issues a valid redemption to the server, since redemptions reveal the anonymity set that the client belongs to.

This judgement can be based on a multitude of factors, associated with the way that a server presents itself in the ecosystem. A non-exhaustive list of server characteristics that a client MAY want to check are the following.

- * Which key registry a server posts their key updates to.
- * How frequent key updates are issued, and which ciphersuite they use.
- * The reason given to initiate the redemption.

To aid client trust decisions, a server can publish a "Privacy Pass policy" that documents the procedures that the server uses to ensure that client privacy is respected. If a server does not publish such a document then the client may choose to use its own judgement, or to reject the server altogether.

It should be noted that the client trust decision can be made apriori by specifying an allow-list of all servers that it accepts tokens from. This means that these checks do not have to be performed online.

8. Privacy considerations

In the Privacy Pass protocol [I-D.ietf-privacypass-protocol], redemption tokens intentionally encode very little information beyond which key was used to sign them. The protocol intentionally uses components that provide cryptographic guarantees of this fact. However, even with these guarantees, the way that the ecosystem is constructed can be used to identify clients based on this limited information.

The goal of the Privacy Pass ecosystem is to construct an environment that can easily measure (and maximize) the relative anonymity of any client that is part of it. An inherent feature of being part of this ecosystem is that any client can only remain private relative to the entire space of users using the protocol. Moreover, by owning tokens for a given set of keys, the client's anonymity set shrinks to the total number of clients controlling tokens for the same keys.

In the following, we consider the possible ways that servers and servers can leverage their position to try and reduce the anonymity sets that clients belong to (or, user segregation). For each case, we provide mitigations that the Privacy Pass ecosystem must implement to prevent these actions.

8.1. Server key rotation

Techniques to introduce client "segregation" can be used to reduce client anonymity. Such techniques are closely linked to the type of key schedule that is used by the server. When a server rotates their key, any client that invokes the issuance protocol in this key cycle will be part of a group of possible clients owning valid tokens for this key. To mechanize this attack strategy, a server could introduce a key rotation policy that forces clients into small key cycles. Thus, reducing the size of the anonymity set for these clients.

We RECOMMEND that servers should only invoke key rotation for fairly large periods of time such as between 1 and 12 weeks. Key rotations represent a trade-off between client privacy and continued server security. Therefore, it is still important that key rotations occur on a fairly regular cycle to reduce the harmfulness of a server key compromise.

With an active user-base, a week gives a fairly large window for clients to participate in the Privacy Pass protocol and thus enjoy the anonymity guarantees of being part of a larger group. The low ceiling of 12 weeks prevents a key compromise from being too destructive. If a server realizes that a key compromise has occurred then the server should sample a new key, and upload the public key to the key registry; invoking any revocation procedures that may apply for the old key.

8.2. Large numbers of servers

Similarly to the server rotation dynamic that is raised above, if there are a large number of servers then segregation can occur. In the FV, AV and PV running modes (Section 5), a verifier OV can trigger redemptions for any of the available servers. Each redemption token that a client holds essentially corresponds to a bit of information about the client that OV can learn. Therefore, there is an exponential loss in anonymity relative to the number of servers that there are.

For example, if there are 32 servers, then OV learns 32 bits of information about the client. If the distribution of server trust is anything close to a uniform distribution, then this is likely to uniquely identify any client amongst all other Internet users. Assuming a uniform distribution is clearly the worst-case scenario, and unlikely to be accurate, but it provides a stark warning against allowing too many servers at any one time.

In cases where clients can hold tokens for all servers at any given time, a strict bound SHOULD be applied to the active number of servers in the ecosystem. We propose that allowing no more than 4 servers at any one time is highly preferable (leading to a maximum of 64 possible user segregations). However, as highlighted in Section 10, having a very large user base (> 5 million users), could potentially allow for larger values. server replacements should only occur with the same frequency as config rotations as they can lead to similar losses in anonymity if clients still hold redemption tokens for previously active servers.

In addition, we RECOMMEND that trusted registries indicate at all times which servers are deemed to be active. If a client is asked to invoke any Privacy Pass exchange for an server that is not declared active, then the client SHOULD refuse to retrieve the server configuration during the protocol.

8.2.1. Allowing larger number of servers

The bounds on the numbers of servers that we proposed above are very restrictive. This is due to the fact that we considered a situation where a client could be issued (and forced to redeem) tokens for any issuing key.

An alternative system is to ensure a robust strategy for ensuring that clients only possess redemption tokens for a similarly small number of servers at any one time. This prevents a malicious verifier from being able to invoke redemptions for many servers since the client would only be holding redemption tokens for a small set of servers. When a client is issued tokens from a new server and already has tokens from the maximum number of servers, it simply deletes the oldest set of redemption tokens in storage and then stores the newly acquired tokens.

For example, if clients ensure that they only hold redemption tokens for 4 servers, then this increases the potential size of the anonymity sets that the client belongs to. However, this doesn't protect clients completely as it would if only 4 servers were permitted across the whole system. For example, these 4 servers could be different for each client. Therefore, the selection of servers they possess tokens for is still revealing. Understanding this trade-off is important in deciding the effective anonymity of each client in the system.

8.3. Partitioning of server key material

If there are multiple key registries, or if a key registry colludes with an server, then it is possible to provide a split-view of an server's key material to different clients. This would involve posting different key material in different locations, or actively modifying the key material at a given location.

Key registries should operate independently of server's in the ecosystem, and within the guidelines stated in Section 4. Any client should follow the recommendations in Section 7 for determining whether an server and its key material should be trusted.

8.4. Tracking and identity leakage

Privacy losses may be encountered if too many redemptions are allowed in a short burst. For instance, in the Internet setting, this may allow delegated or asynchronous verifiers to learn more information from the metadata that the client may hold (such as first-party cookies for other domains). Mitigations for this issue are similar to those proposed in Section 8.2 for tackling the problem of having large number of servers.

In AV, cached SRRs and their associated server public keys have a similar tracking potential to first party cookies in the browser setting. These considerations will be covered in a separate document, detailing Privacy Pass protocol integration into the wider web architecture [I-D.ietf-privacypass-http-api].

8.5. Client incentives for anonymity reduction

Clients may see an incentive in accepting all tokens that are issued by a server, even if the tokens fail later verification checks. This is because tokens effectively represent a form of currency that they can later redeem for some sort of benefit. The verification checks that are put in place are there to ensure that the client does not sacrifice their anonymity. However, a client may judge the "monetary" benefit of owning tokens to be greater than their own privacy.

Firstly, a client behaving in this way would not be compliant with the protocol, as laid out in [I-D.ietf-privacypass-protocol].

Secondly, acting in this way only affects the privacy of the immediate client. There is an exception if a large number of clients colluded to accept bad data, then any client that didn't accept would be part of a smaller anonymity set. However, such a situation would be identical to the situation where the total number of clients in the ecosystem is small. Therefore, the reduction in the size of the anonymity set would be equivalent; see Section 8.2 for more details.

9. Security considerations

We present a number of security considerations that prevent malicious clients from abusing the protocol.

9.1. Double-spend protection

All issuing server should implement a robust storage-query mechanism for checking that tokens sent by clients have not been spent before. Such tokens only need to be checked for each server individually. But all servers must perform global double-spend checks to avoid clients from exploiting the possibility of spending tokens more than once against distributed token checking systems. For the same reason, the global data storage must have quick update times. While an update is occurring it may be possible for a malicious client to spend a token more than once.

9.2. Token exhaustion

When a client holds tokens for an server, it is possible for any verifier to invoke that client to redeem tokens for that server. This can lead to an attack where a malicious verifier can force a client to spend all of their tokens for a given server. To prevent this from happening, methods should be put into place to prevent many tokens from being redeemed at once.

For example, it may be possible to cache a redemption for the entity that is invoking a token redemption. If the verifier requests more tokens then the client simply returns the cached token that it returned previously. This could also be handled by simply not redeeming any tokens for verification if a redemption had already occurred in a given time window.

In AV, the client instead caches the SRR that it received in the asynchronous redemption exchange with the server. If the same verifier attempts another redemption request, then the client simply returns the cached SRR. The SRRs can be revoked by the server, if need be, by providing an expiry date or by signaling that records from a particular window need to be refreshed.

9.3. Avoiding server centralization

[[OPEN ISSUE: explain potential and mitigations for server centralization]]

10. Protocol parametrization

We provide a summary of the parameters that we use in the Privacy Pass protocol ecosystem. These parameters are informed by both privacy and security considerations that are highlighted in Section 8 and Section 9, respectively. These parameters are intended as a single reference point for those implementing the protocol.

Firstly, let U be the total number of users, I be the total number of servers. We let M be the total number of metadata bits that are allowed to be added by any given server. Assuming that each user accept tokens from a uniform sampling of all the possible servers, as a worst-case analysis, this segregates users into a total of 2^I buckets. As such, we see an exponential reduction in the size of the anonymity set for any given user. This allows us to specify the privacy constraints of the protocol below, relative to the setting of A .

parameter	value
Minimum anonymity set size (A)	5000
Recommended key lifetime (L)	2 - 24 weeks
Recommended key rotation frequency (F)	$L/2$
Maximum additional metadata bits (M)	1
Maximum allowed servers (I)	$(\log_2(U/A) - 1) / 2$
Maximum active issuance keys	1
Maximum active redemption keys	2
Minimum cryptographic security parameter	128 bits

Table 1

10.1. Justification

We make the following assumptions in these parameter choices.

- * Inferring the identity of a user in a 5000-strong anonymity set is difficult.
- * After 2 weeks, all clients in a system will have rotated to the new key.

In terms of additional metadata, the only concrete applications of Privacy Pass that use additional metadata require just a single bit. Therefore, we set the ceiling of permitted metadata to 1 bit for now, this may be revisited in future revisions.

The maximum choice of I is based on the equation $1/2 * U/2^{(2I)} = A$. This is derived from the fact that permitting I servers lead to 2^I segregations of the total user-base U . Moreover, if we permit $M = 1$, then this effectively halves the anonymity set for each server, and thus we incur a factor of $2I$ in the exponent. By reducing I , we limit the possibility of performing the attacks mentioned in Section 8.

We must also account for each user holding issued data for more than one possible active keys. While this may also be a vector for monitoring the access patterns of clients, it is likely to unavoidable that clients hold valid issuance data for the previous key epoch. This also means that the server can continue to verify redemption data for a previously used key. This makes the rotation period much smoother for clients.

For privacy reasons, it is recommended that key epochs are chosen that limit clients to holding issuance data for a maximum of two keys. By choosing $F = L/2$ then the minimum value of F is a week, since the minimum recommended value of L is 2 weeks. Therefore, by the initial assumption, then all users should only have access to only two keys at any given time. This reduces the anonymity set by another half at most.

Finally, the minimum security parameter size is related to the cryptographic security offered by the protocol that is run. This parameter corresponds to the number of operations that any adversary has in breaking one of the security guarantees in the Privacy Pass protocol [I-D.ietf-privacypass-protocol].

10.2. Example parameterization

Using the specification above, we can give some example parameterizations. For example, the current Privacy Pass browser extension [PPEXT] has nearly 300000 active users (from Chrome and Firefox). As a result, $\log_2(U/A)$ is approximately 6 and so the maximum value of I should be 3.

If the value of U is much bigger (e.g. 5 million) then this would permit $I = (\log_2(5000000/5000)-1)/2 \approx 4$ servers.

10.3. Allowing more servers

Using the recommendations in Section 8.2.1, it is possible to tolerate larger number of servers if clients in the ecosystem ensure that they only store tokens for a small number of them. In particular, if clients limit their storage of redemption tokens to the bound implied by I , then prevents a malicious verifier from triggering redemptions for all servers in the ecosystem.

11. Extension integration policy

The Privacy Pass protocol and ecosystem are both intended to be receptive to extensions that expand the current set of functionality. As specified in [I-D.ietf-privacypass-protocol], all extensions to the Privacy Pass protocol SHOULD be specified as separate documents that modify the content of this document in some way. We provide guidance on the type of modifications that are possible in the following.

Any such extension should also come with a detailed analysis of the privacy impacts of the extension, why these impacts are justified, and guidelines on changes to the parametrization in Section 10. Similarly, extensions MAY also add new server running modes, if applicable, to those that are documented in Section 5.

Any extension to the Privacy Pass protocol must adhere to the guidelines specified in Section 4 for managing server public key data.

12. Existing applications

The following is a non-exhaustive list of applications that currently make use of the Privacy Pass protocol, or some variant of the underlying functionality.

12.1. Cloudflare challenge pages

Cloudflare uses an implementation of the Privacy Pass protocol for allowing clients that have previously interacted with their Internet challenge protection system to bypass future challenges [PPSRV]. These challenges can be expensive for clients, and there have been cases where bugs in the implementations can severely degrade client accessibility.

Clients must install a browser extension [PPEXT] that acts as the Privacy Pass client in an exchange with Cloudflare's Privacy Pass server, when an initial challenge solution is provided. The client extension stores the issued tokens and presents a valid redemption

token when it sees future Cloudflare challenges. If the redemption token is verified by the server, the client passes through the security mechanism without completing a challenge.

12.2. Trust Token API

The Trust Token API [TrustTokenAPI] has been devised as a generic API for providing Privacy Pass functionality in the browser setting. The API is intended to be implemented directly into browsers so that server's can directly trigger the Privacy Pass workflow.

12.3. Zero-knowledge Access Passes

The PrivateStorage API developed by Least Authority is a solution for uploading and storing end-to-end encrypted data in the cloud. A recent addition to the API [PrivateStorage] allows clients to generate Zero-knowledge Access Passes (ZKAPs) that the client can use to show that it has paid for the storage space that it is using. The ZKAP protocol is based heavily on the Privacy Pass redemption mechanism. The client receives ZKAPs when it pays for storage space, and redeems the passes when it interacts with the PrivateStorage API.

12.4. Basic Attention Tokens

The browser Brave uses Basic Attention Tokens (BATs) to provide the basis for an anonymity-preserving rewards scheme [Brave]. The BATs are essentially Privacy Pass redemption tokens that are provided by a central Brave server when a client performs some action that triggers a reward event (such as watching an advertisement). When the client amasses BATs, it can redeem them with the Brave central server for rewards.

12.5. Token Based Services

Similarly to BATs, a more generic approach for providing anonymous peers to purchase resources from anonymous servers has been proposed [OpenPrivacy]. The protocol is based on a variant of Privacy Pass and is intended to allow clients purchase (or pre-purchase) services such as message hosting, by using Privacy Pass redemption tokens as a form of currency. This is also similar to how ZKAPs are used.

13. References

13.1. Normative References

- [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/info/rfc2119>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

13.2. Informative References

- [Brave] "Brave Rewards", n.d., <<https://brave.com/brave-rewards/>>.
- [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/wp-content/uploads/2021/01/PrivateStats-De-Identified-Authenticated-Logging-at-Scale_final.pdf>.
- [I-D.ietf-privacypass-http-api] Valdez, S., "Privacy Pass HTTP API", Work in Progress, Internet-Draft, draft-ietf-privacypass-http-api-00, 5 January 2021, <<http://www.ietf.org/internet-drafts/draft-ietf-privacypass-http-api-00.txt>>.
- [I-D.ietf-privacypass-protocol] Celi, S., Davidson, A., and A. Faz-Hernandez, "Privacy Pass Protocol Specification", Work in Progress, Internet-Draft, draft-ietf-privacypass-protocol-00, 5 January 2021, <<http://www.ietf.org/internet-drafts/draft-ietf-privacypass-protocol-00.txt>>.
- [KLOR20] Kreuter, B., Lepoint, T., Orrù, M., and M. Raykova, "Anonymous Tokens with Private Metadata Bit", DOI 10.1007/978-3-030-56784-2_11, Advances in Cryptology - CRYPTO 2020 pp. 308-336, 2020, <https://doi.org/10.1007/978-3-030-56784-2_11>.
- [OpenPrivacy] "Token Based Services - Differences from PrivacyPass", n.d., <<https://openprivacy.ca/assets/towards-anonymous-prepaid-services.pdf>>.
- [PPEXT] "Privacy Pass Browser Extension", n.d., <<https://github.com/privacypass/challenge-bypass-extension>>.

- [PPSRV] Sullivan, N., "Cloudflare Supports Privacy Pass", n.d., <<https://blog.cloudflare.com/cloudflare-supports-privacy-pass/>>.
- [PrivateStorage] Steininger, L., "The Path from S4 to PrivateStorage", n.d., <<https://medium.com/least-authority/the-path-from-s4-to-privatestorage-ae9d4a10b2ae>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", RFC 6962, DOI 10.17487/RFC6962, June 2013, <<https://www.rfc-editor.org/info/rfc6962>>.
- [TrustTokenAPI] Google, ., "Getting started with Trust Tokens", n.d., <<https://web.dev/trust-tokens/>>.

Appendix A. Contributors

- * Alex Davidson (alex.davidson92@gmail.com)
- * Christopher Wood (caw@heapingbits.net)

Authors' Addresses

Alex Davidson
LIP
Lisbon
Portugal

Email: alex.davidson92@gmail.com

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

Email: caw@heapingbits.net

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 8 September 2022

A. Davidson
LIP
J. Iyengar
Fastly
C. A. Wood
Cloudflare
7 March 2022

Privacy Pass Architectural Framework
draft-ietf-privacypass-architecture-03

Abstract

This document specifies the architectural framework for constructing secure and anonymity-preserving instantiations of the Privacy Pass protocol. It provides recommendations on how the protocol ecosystem should be constructed to ensure the privacy of clients, and the security of all participating entities.

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 8 September 2022.

Copyright Notice

Copyright (c) 2022 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	3
2. Terminology	4
3. Architecture	5
3.1. Redemption Protocol	5
3.2. Issuance Protocol	6
3.2.1. Attester Role	7
3.2.2. Issuer Role	8
3.2.3. Metadata	10
3.2.4. Issuance Protocol Extensibility	10
4. Deployment Considerations	11
4.1. Shared Origin, Attester, Issuer	11
4.2. Joint Attester and Issuer	12
4.3. Joint Origin and Issuer	13
4.4. Split Origin, Attester, Issuer	14
5. Privacy Considerations	15
5.1. Metadata Privacy Implications	15
5.2. Issuer Key Rotation	15
5.3. Large Number of Issuers	16
5.3.1. Allowing More Issuers	17
6. Security Considerations	18
6.1. Double-spend Protection	18
6.2. Token Exhaustion	19
7. Protocol Parameterization	19
7.1. Justification	20
7.2. Example parameterization	21
7.3. Allowing more Issuers	21
8. References	21
8.1. Normative References	21
8.2. Informative References	22
Appendix A. Acknowledgements	23
Authors' Addresses	23

1. Introduction

Privacy Pass is a protocol for authorization based on anonymous-credential authentication mechanisms. Typical approaches for authorizing clients, such as through the use of long-term 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., whether or not the client is an authorized user or has completed some prior challenge, clients present unlinkable proofs that attest to this information.

The most basic Privacy Pass protocol provides a set of cross-origin authorization tokens that protect the client's anonymity during interactions with a server. This allows clients to communicate an attestation of a previously authenticated server action, without having to reauthenticate manually. The tokens retain anonymity in the sense that the act of revealing them cannot be linked back to the session where they were initially issued.

At a high level, Privacy Pass is composed of two protocols: issuance and redemption.

The issuance protocol runs between a Client and two network functions in the Privacy Pass architecture: Attestation and Issuance. These two network functions can be implemented by the same protocol participant, but can also be implemented separately. The Issuer is responsible for issuing tokens in response to requests from Clients. The Attester is responsible for attesting properties about the Client for which tokens are issued. The Issuer needs to be trusted by the server that later redeems the token. Attestation can be performed by the Issuer or by an Attester that is trusted by the Issuer. Clients might prefer to select different Attesters, separate from the Issuer, to be able to use preferred authentication methods or improve privacy by not directly communicating with an Issuer. Depending on the attestation, Attesters can store state about a Client, such as the number of overall tokens issued thus far. As an example of an Issuance protocol, in the original Privacy Pass protocol [PPSRV], tokens were only issued to Clients that solved CAPTCHAs. In this context, the Attester attested that some client solved a CAPTCHA and the resulting token produced by the Issuer was proof of this fact.

The redemption protocol runs between Client and Origin (server). It allows Origins to challenge Clients to present one or more tokens for authorization. Depending on the type of token, e.g., whether or not it is cross-origin or per-origin, and whether or not it can be cached, the Client either presents a previously obtained token or invokes the issuance protocol to acquire one for authorization.

The issuance and redemption protocols operate in concert as shown in the figure below.




Figure 1: Privacy Pass Architectural Components

This document describes requirements for both issuance and redemption protocols. This document also describes ecosystem considerations that impact the stated privacy and security guarantees of the protocol.

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.
- * **Origin:** An entity that challenges Clients for tokens.
- * **Issuer:** An entity that issues tokens to Clients for properties attested to by the Attester.
- * **Attester:** An entity that attests to properties of Client for the purposes of token issuance.

3. Architecture

The Privacy Pass architecture consists of four logical entities -- Client, Origin, Issuer, and Attester -- that work in concert as shown in Section 1 for token issuance and redemption. This section describes the purpose of token issuance and redemption and the requirements therein on the relevant participants.

3.1. Redemption Protocol

The redemption protocol is a simple challenge-response based authorization protocol between Client and Origin. Origins prompt Clients with a token challenge and, if possible, Clients present a valid token for the challenge in response. The context in which an Origin challenges a Client for a token is referred to as the redemption context. This context includes all information associated with the redemption event, such as the timestamp of the event, Client visible information (including the IP address), and the Origin name.

The challenge controls the type of token that the Origin will accept for the given resource. As described in [HTTP-Authentication], there are a number of ways in which the token may vary, including:

- * Issuance protocol. The token 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. Tokens identify which issuers are trusted for a given issuance protocol.
- * Interactive or non-interactive. Tokens can either be interactive or not. An interactive token is one which requires a freshly issued token based on the challenge, whereas a non-interactive token can be issued proactively and cached for future use.
- * Per-origin or cross-origin. Tokens can be constrained to the Origin for which the challenge originated, or can be used across Origins.

Depending on the use case, Origins may need to maintain state to track redeemed tokens. For example, Origins that accept non-interactive, cross-origin tokens SHOULD track which tokens have been redeemed already, since these tokens can be issued and then spent multiple times in response to any such challenge. See Section 6.1 for discussion.

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. If tokens protected with resources are unique to a single Origin, then said Origin **MUST NOT** admit cross-origin tokens for authorization.

3.2. Issuance Protocol

The issuance protocol embodies the core of Privacy Pass. It takes as input a challenge from the redemption protocol and produces a token, as shown in the figure below.

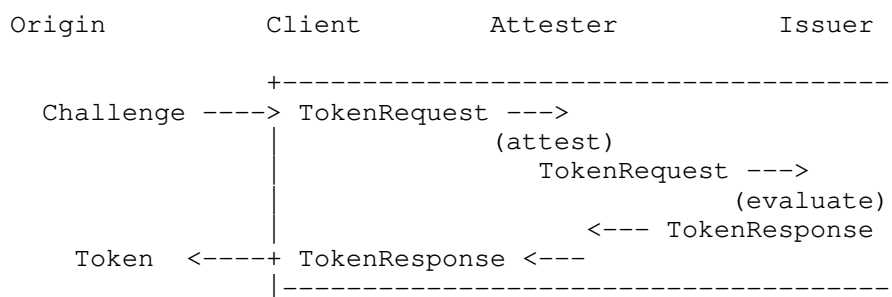


Figure 2: Issuance Overview

Clients interact with the Attester and Issuer to produce a token in response to 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 Token issuance protocol using the challenge, a randomly generated nonce, and public key for the Issuer. The Token issuance protocol itself can be any interactive protocol between Client, Issuer, or other parties that produces a valid authenticator over the Client's 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. The issuance protocol can reveal the Issuer public key for the purposes of

determining which private key to use in producing the issuance protocol. A result of this property is that the redemption flow is unlinkable from the issuance flow.

2. One-more forgery security. The issuance protocol MUST NOT allow malicious Clients or Attesters (acting as Clients) to forge tokens without interacting with the Issuer directly.
3. Concurrent security. The issuance protocol MUST be safe to run concurrently with arbitrarily many Clients.

Each Issuance protocol MUST come with a detailed analysis of the privacy impacts of the protocol, why these impacts are justified, and guidelines on changes to the parametrization in Section 7.

The mechanism by which clients obtain the Issuer public key is not specified. Clients may be configured with this key or they may discover it via some other form. See [CONSISTENCY].

Depending on the use 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, e.g., through the use of an IP-hiding proxy service like Tor. In general, Clients SHOULD minimize or remove identifying information where possible when invoking the issuance protocol.

Issuers MUST NOT issue tokens for Clients through untrusted Attesters. This is important because the Attester's role is to vouch for trust in privacy-sensitive Client information, such as account identifiers or IP address information, to the Issuer. Tokens produced by an Issuer that admits issuance for any type of attestation cannot be relied on for any specific property. See Section 3.2.1 for more details.

3.2.1. Attester Role

Attestation is an important part of the issuance protocol. Attestation is the process by which an Attester bears witness to, confirms, or authenticates a Client so as to verify a property about the Client that is required for Issuance. Examples of attestation properties include, though are not limited to:

- * Capable of solving a CAPTCHA. Clients that solve CAPTCHA challenges can attest to this capability for the purposes of being ruled out as a bot or otherwise automated Client.

- * Client state. Clients can be associated with state and the attester can attest to this state. Examples of state include the number of issuance protocol invocations, the client's geographic region, and whether the client has a valid application-layer account.
- * Trusted device. Some Clients run on trusted hardware that are capable of producing device-level attestation statements.

Each of these attestation types have different security properties. For example, attesting to having a valid account is different from attesting to be running on trusted hardware. In general, Attesters should accept a limited form of attestation formats.

Each attestation format also has an impact on the overall system privacy. For example, the number of users in possession of a single class of trusted device might be lesser than the number of users that can solve CAPTCHAs. Similarly, requiring a conjunction of attestation types could decrease the overall anonymity set size. For example, the number of Clients that have solved a CAPTCHA in the past day, have a valid account, and are running on a trusted device is lesser than the number of Clients that have solved a CAPTCHA in the past day. Attesters should not admit attestation types that result in small anonymity sets.

3.2.2. Issuer Role

Issuers MUST be uniquely identifiable by all Clients with a consistent identifier. In a web context, this identifier might be the Issuer host name. As discussed later in Section 5, ecosystems that admit a large number of Issuers can lead to privacy concerns for the Clients in the ecosystem. Therefore, in practice, the number of Issuers should be bounded. The actual Issuers can be replaced with different Issuers as long as the total never exceeds these bounds. Moreover, Issuer replacements also have an effect on client anonymity that is similar to when a key rotation occurs. See Section 5 for more details about maintaining privacy with multiple Issuers.

3.2.2.1. Key Management

To facilitate issuance, the Issuer MUST hold an Issuance key pair at any given time. The Issuer public key MUST be made available to all Clients in such a way that key rotations and other updates are publicly visible. The key material and protocol configuration that an Issuer uses to produce tokens corresponds to a number of different pieces of information.

- * The issuance protocol in use; and

- * The public keys that are active for the Issuer.

The way that the Issuer publishes and maintains this information impacts the effective privacy of the clients; see Section 5 for more details. The fundamental requirement for key management and discovery is that Issuers must be unable to target specific clients with unique keys without detection. There are a number of ways in which this might be implemented:

- * Servers use a verifiable, tamper-free registry from which clients discover keys. Similar to related mechanisms and protocols such as Certificate Transparency [RFC6962], this may require external auditors or additional client behavior to ensure the registry state is consistent for all clients.
- * Clients use an anonymity-preserving tool such as Tor to discover keys from multiple network vantage points. This is done to ensure consistent keys to seemingly different clients.
- * Clients embed Issuer keys into software.

As above, specific mechanisms for key management and discovery are out of scope for this document.

3.2.2.2. Key Rotation

Token issuance associates all issued tokens with a particular choice of key. If an Issuer issues tokens with many keys, then this may harm the anonymity of the Client. For example, they would be able to map the Client's access patterns by inspecting which key each token they possess has been issued under.

To prevent against this, Issuers MUST only use one private key for issuing tokens at any given time. Servers MAY use one or more keys for redemption to allow Issuers for seamless key rotation.

Servers may rotate keys as a means of revoking tokens issued under old or otherwise expired keys. Alternatively, Issuers may include expiration information as metadata alongside the token; See Section 3.2.3 for more discussion about metadata constraints. Both techniques are equivalent since they cryptographically bind expiration to individual tokens.

Key rotations should be limited in frequency for similar reasons. See Section 7 for guidelines on what frequency of key rotations are permitted.

3.2.3. 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. The total amount of metadata bits included in a token is the sum of public and private metadata bits. See Section 7 for discussion about metadata limits.

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. In normal circumstances, clients cannot determine if this value is correct or otherwise a tracking vector.

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

Metadata may also be arbitrarily long or bounded in length. The amount of permitted metadata may be determined by application or by the underlying cryptographic protocol.

3.2.4. Issuance Protocol Extensibility

The Privacy Pass protocol and ecosystem are both intended to be receptive to extensions that expand the current set of functionalities through new issuance protocols. Each issuance protocol SHOULD come with a detailed analysis of the privacy impacts of the extension, why these impacts are justified, and guidelines on changes to the parametrization in Section 7. Any extension to the Privacy Pass protocol MUST adhere to the guidelines specified in Section 3.2.2 for managing Issuer public key data.

4. Deployment Considerations

Client uses Privacy Pass to separate attestation context and redemption context. Linking or combining these contexts can reveal sensitive information about the Client, including their identity or browsing history. Depending on the deployment model, separating these contexts can take different forms. The Origin, Attester, and Issuer portrayed in Figure 1 can be instantiated and deployed in a number of different ways. This section covers some expected deployment models and their corresponding security and privacy considerations. The discussion below assumes non-collusion between entities when operated by separate parties. 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 the figure below.

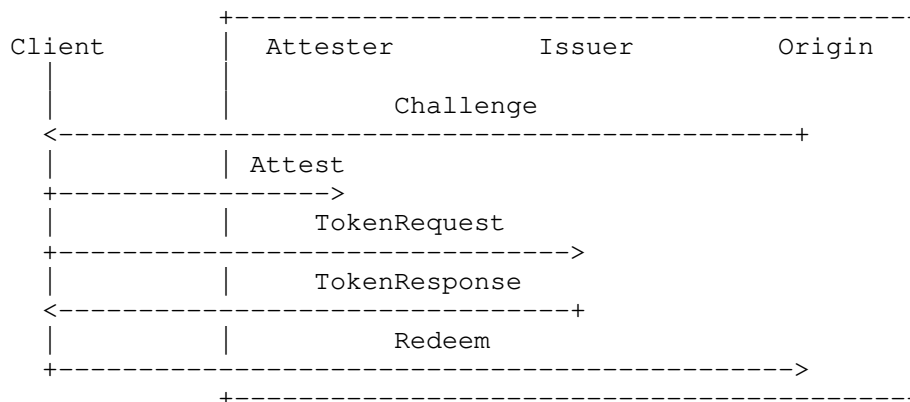


Figure 3: Shared Deployment Model

This model represents the initial deployment of Privacy Pass, as described in [PPSRV]. In this model, the Attester, Issuer, and Origin share the attestation 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 be used to learn or reconstruct a Client's browsing history.

Attestation and redemption context unlinkability requires that these events be separated over time, e.g., through the use of non-interactive tokens that can be issued without a fresh Origin challenge, or over space, e.g., through the use of an anonymizing proxy 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, as shown in the figure below.

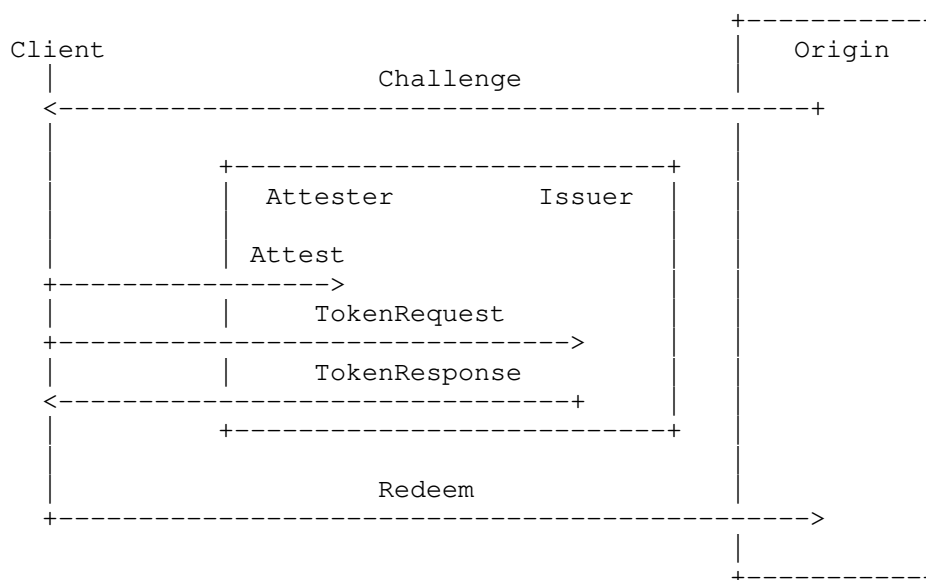


Figure 4: Joint Attester and Issuer Deployment Model

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 attestation context for the Client, which can be separate from the Origin's redemption context.

For certain types of issuance protocols, this model separates attestation and redemption contexts. However, 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 link attestation and redemption contexts 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.

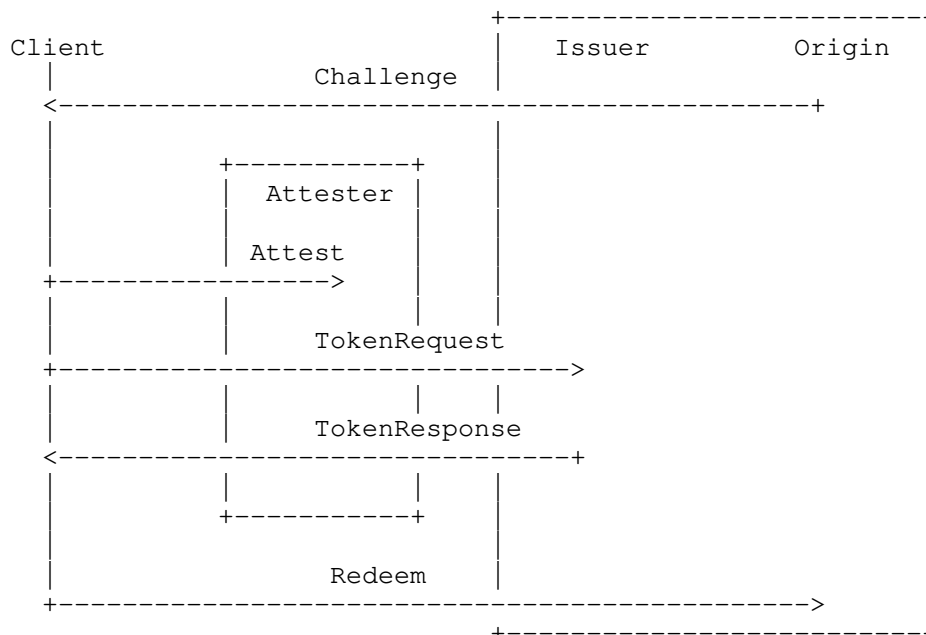


Figure 5: Joint Origin and Issuer Deployment Model

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 and redemption contexts are separate. 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. If there is only a single Origin, then per-Origin tokens are not appropriate in this model, since the Attester can learn the redemption context. (Note, however, that the Attester does not learn whether a token 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 shown in the figure below.

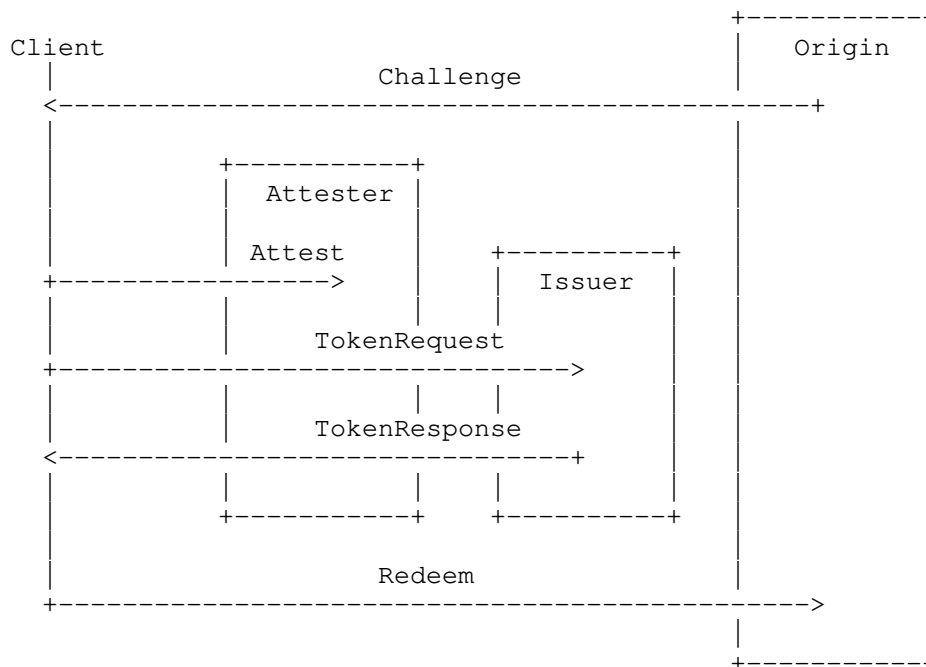


Figure 6: Split Deployment Model

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 and redemption contexts are separate, and therefore any type of token challenge is suitable in this model assuming there is more than a single Origin. As in the Joint Origin and Issuer model in Section 4.3, if there is only a single Origin, then per-Origin tokens are not appropriate.

5. Privacy Considerations

Client uses Private Pass to separate attestation context and redemption context. Depending on the deployment model, this can take different forms. For example, any Client can only remain private relative to the entire space of other Clients using the protocol. Moreover, by owning tokens for a given set of keys, the Client's anonymity set shrinks to the total number of clients controlling tokens for the same keys.

In the following, we consider the possible ways that Issuers can leverage their position to try and reduce the anonymity sets that Clients belong to (or, user segregation). For each case, we provide mitigations that the Privacy Pass ecosystem must implement to prevent these actions.

5.1. Metadata Privacy Implications

Any 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, the amount of metadata used by an Issuer in creating redemption tokens must be taken into account -- together with the bits of information that Issuer's 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. In general, Issuers should permit no more than 32 bits of metadata, as this can uniquely identify each possible user. We discuss this more in Section 7.

5.2. Issuer Key Rotation

Techniques to introduce Client "segregation" can be used to reduce Client anonymity. Such techniques are closely linked to the type of key schedule that is used by the Issuer. When an Issuer rotates their key, any Client that invokes the issuance protocol in this key cycle will be part of a group of possible clients owning valid tokens for this key. To mechanize this attack strategy, an Issuer could introduce a key rotation policy that forces Clients into small key cycles. Thus, reducing the size of the anonymity set for these Clients.

Issuers SHOULD invoke key rotation for a period of time between 1 and 12 weeks. Key rotations represent a trade-off between Client privacy and continued Issuer security. Therefore, it is still important that key rotations occur on a regular cycle to reduce the harmfulness of an Issuer key compromise.

With a large number of Clients, a minimum of one week gives a large enough window for Clients to participate in the issuance protocol and thus enjoy the anonymity guarantees of being part of a larger group. A maximum of 12 weeks limits the damage caused by a key compromise. If an Issuer realizes that a key compromise has occurred then the Issuer should generate a new key and make it available to Clients. If possible, it should invoke any revocation procedures that may apply for the old key.

5.3. Large Number of Issuers

Similarly to the Issuer rotation dynamic that is raised above, if there are a large number of Issuers, and Origins accept all of them, segregation can occur. For example, if Clients obtain tokens from many Issuers, and Origins later challenge a Client for a token from each Issuer, Origins can learn information about the Client. Each per-Issuer token that a Client holds essentially corresponds to a bit of information about the Client that Origin learn. Therefore, there is an exponential loss in anonymity relative to the number of Issuers that there are.

For example, if there are 32 Issuers, then Origins learn 32 bits of information about the Client if a valid token is presented for each one. If the distribution of Issuer trust is anything close to a uniform distribution, then this is likely to uniquely identify any Client amongst all other Internet users. Assuming a uniform distribution is clearly the worst-case scenario, and unlikely to be accurate, but it provides a stark warning against allowing too many Issuers at any one time.

In cases where clients can hold tokens for all Issuers at any given time, a strict bound SHOULD be applied to the active number of Issuers in the ecosystem. We propose that allowing no more than 4 Issuers at any one time is highly preferable (leading to a maximum of 64 possible user segregations). However, as highlighted in Section 7, having a very large user base (> 5 million users), could potentially allow for larger values. Issuer replacements should only occur with the same frequency as config rotations as they can lead to similar losses in anonymity if clients still hold redemption tokens for previously active Issuers.

In addition, we RECOMMEND that trusted registries indicate at all times which Issuers are deemed to be active. If a Client is asked to invoke any Privacy Pass exchange for an Issuer that is not declared active, then the client SHOULD refuse to retrieve the Issuer public key during the protocol.

5.3.1. Allowing More Issuers

The bounds on the numbers of Issuers that this document proposes above are very restrictive. This is because this document considers a situation where a Client could be challenged (and asked to redeem) tokens for any Issuer.

An alternative system is to ensure a robust strategy for ensuring that Clients only possess redemption tokens for a similarly small number of Issuers at any one time. This prevents a malicious verifier from being able to invoke redemptions for many Issuers since the Client would only be holding redemption tokens for a small set of Issuers. When a Client is issued tokens from a new Issuer and already has tokens from the maximum number of Issuers, it simply deletes the oldest set of redemption tokens in storage and then stores the newly acquired tokens.

For example, if Clients ensure that they only hold redemption tokens for 4 Issuers, then this increases the potential size of the anonymity sets that the Client belongs to. However, this doesn't protect Clients completely as it would if only 4 Issuers were permitted across the whole system. For example, these 4 Issuers could be different for each Client. Therefore, the selection of Issuers they possess tokens for is still revealing. Understanding this trade-off is important in deciding the effective anonymity of each Client in the system.

5.3.1.1. Redemption Partitions

Another option to allow a large number of Issuers in the ecosystem, while preventing the joining of a number of different tokens is for the Client to maintain sharded "redemption partitions". This would allow the Client to redeem the tokens it wishes to use in a particular context, while still allowing the Client to maintain a large variety of tokens from many Issuers. Within a redemption partition, the Client limits the number of different Issuers used to a small number to maintain the privacy properties the Client requires. As long as each redemption partition maintains a strong privacy boundary with each other, the verifier will only be able to learn a number of bits of information up to the limits within that "redemption partitions".

To support this strategy, the client keeps track of a partition which contains the set of Issuers that redemptions have been attempted against. An empty redemption is returned when the limit has been hit:

```

Client(partition, issuer)                                Issuer(skS, pkS)
-----
if issuer not in partition {
  if partition.length > REDEEM_LIMIT {
    Output {}
    return
  }
  partition.push(issuer)
}
token = store[issuer.id].pop()
req = Redeem(token, info)

                                req
                                ----->

                                if (dsIdx.includes(req.data)) {
                                  raise ERR_DOUBLE_SPEND
                                }
                                resp = Verify(pkS, skS, req)
                                if resp.success {
                                  dsIdx.push(req.data)
                                }

                                resp
                                <-----
Output resp

```

6. Security Considerations

We present a number of security considerations that prevent malicious Clients from abusing the protocol.

6.1. Double-spend Protection

When applicable for non-interactive tokens, all Origins SHOULD implement a robust storage-query mechanism for checking that tokens sent by clients have not been spent before. Such tokens only need to be checked for each Origin individually. But all Origins must perform global double-spend checks to avoid clients from exploiting the possibility of spending tokens more than once against distributed token checking systems. For the same reason, the global data storage must have quick update times. While an update is occurring it may be possible for a malicious client to spend a token more than once.

6.2. Token Exhaustion

When a Client holds tokens for an Issuer, it is possible for any verifier to invoke that client to redeem tokens for that Issuer. This can lead to an attack where a malicious verifier can force a Client to spend all of their tokens from a given Issuer. To prevent this from happening, tokens can be scoped to single Origins such that they can only be redeemed within for a single Origin.

If tokens are cross-Origin, Clients should use alternate methods to prevent many tokens from being redeemed at once. For example, if the Origin requests an excess of tokens, the Client could choose to not present any tokens for verification if a redemption had already occurred in a given time window.

7. Protocol Parameterization

This section provides a summary of the parameters used in the Privacy Pass protocol ecosystem. These parameters are informed by both privacy and security considerations that are highlighted in Section 5 and Section 6, respectively. These parameters are intended as a single reference point for those implementing the protocol.

Firstly, let U be the total number of Clients (or users), I be the total number of Issuers. We let M be the total number of metadata bits that are allowed to be added by any given Issuer. Assuming that each user accept tokens from a uniform sampling of all the possible Issuers, as a worst-case analysis, this segregates Clients into a total of 2^I buckets. As such, we see an exponential reduction in the size of the anonymity set for any given user. This allows us to specify the privacy constraints of the protocol below, relative to the setting of A .

parameter	value
Minimum anonymity set size (A)	5000
Recommended key lifetime (L)	2 - 24 weeks
Recommended key rotation frequency (F)	$L/2$
Maximum additional metadata bits (M)	1
Maximum allowed Issuers (I)	$(\log_2(U/A) - 1) / 2$
Maximum active issuance keys	1
Maximum active redemption keys	2
Minimum cryptographic security parameter	128 bits

Table 1

7.1. Justification

We make the following assumptions in these parameter choices.

- * Inferring the identity of a user in a 5000-strong anonymity set is difficult.
- * After 2 weeks, all Clients in a system will have rotated to the new key.

In terms of additional metadata, the only concrete applications of Privacy Pass that use additional metadata require just a single bit. Therefore, we set the ceiling of permitted metadata to 1 bit for now, this may be revisited in future revisions.

The maximum choice of I is based on the equation $1/2 * U/2^{(2I)} = A$. This is derived from the fact that permitting I Issuers lead to 2^I segregations of the total user-base U. Moreover, if we permit $M = 1$, then this effectively halves the anonymity set for each Issuer, and thus we incur a factor of 2I in the exponent. By reducing I, we limit the possibility of performing the attacks mentioned in Section 5.

We must also account for each user holding issued data for more than one possible active keys. While this may also be a vector for monitoring the access patterns of Clients, it is likely to

unavoidable that Clients hold valid issuance data for the previous key epoch. This also means that the Issuer can continue to verify redemption data for a previously used key. This makes the rotation period much smoother for Clients.

For privacy reasons, it is recommended that key epochs are chosen that limit Clients to holding issuance data for a maximum of two keys. By choosing $F = L/2$ then the minimum value of F is a week, since the minimum recommended value of L is 2 weeks. Therefore, by the initial assumption, then all users should only have access to only two keys at any given time. This reduces the anonymity set by another half at most.

Finally, the minimum security parameter size is related to the cryptographic security offered by the protocol that is run. This parameter corresponds to the number of operations that any adversary has in breaking one of the security guarantees in the Privacy Pass protocol [I-D.ietf-privacypass-protocol].

7.2. Example parameterization

Using the specification above, we can give some example parameterizations. For example, the current Privacy Pass browser extension [PPEXT] has nearly 300000 active users (from Chrome and Firefox). As a result, $\log_2(U/A)$ is approximately 6 and so the maximum value of I should be 3.

If the value of U is much bigger (e.g. 5 million) then this would permit $I = (\log_2(5000000/5000)-1)/2 \approx 4$ Issuers.

7.3. Allowing more Issuers

Using the recommendations in Section 5.3.1, it is possible to tolerate larger number of Issuers if Clients in the ecosystem ensure that they only store tokens for a small number of them. In particular, if Clients limit their storage of redemption tokens to the bound implied by I , then prevents a malicious verifier from triggering redemptions for all Issuers in the ecosystem.

8. References

8.1. Normative References

[HTTP-Authentication]

"The Privacy Pass HTTP Authentication Scheme", n.d.,
<<https://datatracker.ietf.org/doc/html/draft-pauly-privacypass-auth-scheme-00>>.

- [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>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.

8.2. Informative References

- [CONSISTENCY] Davidson, A., Finkel, M., Thomson, M., and C. A. Wood, "Key Consistency and Discovery", Work in Progress, Internet-Draft, draft-wood-key-consistency-02, 4 March 2022, <<https://datatracker.ietf.org/doc/html/draft-wood-key-consistency-02>>.
- [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/wp-content/uploads/2021/01/PrivateStats-De-Identified-Authenticated-Logging-at-Scale_final.pdf>.
- [I-D.ietf-privacypass-protocol] Celi, S., Davidson, A., Faz-Hernandez, A., Valdez, S., and C. A. Wood, "Privacy Pass Issuance Protocol", Work in Progress, Internet-Draft, draft-ietf-privacypass-protocol-02, 31 January 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-protocol-02>>.
- [KLOR20] Kreuter, B., Lepoint, T., Orrù, M., and M. Raykova, "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>.
- [PPEXT] "Privacy Pass Browser Extension", n.d., <<https://github.com/privacypass/challenge-bypass-extension>>.

- [PPSRV] Sullivan, N., "Cloudflare Supports Privacy Pass", n.d.,
<<https://blog.cloudflare.com/cloudflare-supports-privacy-pass/>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", RFC 6962, DOI 10.17487/RFC6962, June 2013,
<<https://www.rfc-editor.org/rfc/rfc6962>>.

Appendix A. Acknowledgements

The authors would like to thank Scott Hendrickson, Tommy Pauly, Benjamin Schwartz, Steven Valdez and other members 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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 9 July 2021

S. Valdez
Google LLC
5 January 2021

Privacy Pass HTTP API
draft-ietf-privacypass-http-api-00

Abstract

This document specifies an integration for Privacy Pass over an HTTP API, along with recommendations on how key commitments are stored and accessed by HTTP-based consumers.

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 9 July 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
 - 1.1. Terminology
 - 1.2. Layout
 - 1.3. Requirements
2. Privacy Pass HTTP API Wrapping
3. Server key registry
 - 3.1. Key Registry
 - 3.2. Server Configuration Retrieval
4. Key Commitment Retrieval
5. Privacy Pass Issuance

6.	Privacy Pass Redemption
6.1.	Generic Token Redemption
6.2.	Direct Redemption
6.3.	Delegated Redemption
7.	Security Considerations
8.	IANA Considerations
8.1.	Well-Known URI
9.	Normative References
	Author's Address

1. Introduction

The Privacy Pass protocol as described in [draft-davidson-pp-protocol] can be integrated with a number of different settings, from server to server communication to browsing the internet.

In this document, we will provide an API to use for integrating Privacy Pass with an HTTP framework. Providing the format of HTTP requests and responses needed to implement the Privacy Pass protocol.

1.1. Terminology

We use the same definition of server and client that is used in [draft-davidson-pp-protocol] and [draft-davidson-pp-architecture].

We assume that all protocol messages are encoded into raw byte format before being sent. We use the TLS presentation language [RFC8446] to describe the structure of protocol messages.

1.2. Layout

- * Section 2: Describes the wrapping of messages within HTTP requests/responses.
- * Section 3: Describes how HTTP clients retrieve server configurations and key commitments.
- * Section 5: Describes how issuance requests are performed via a HTTP API.
- * Section 6: Describes how redemption requests are performed via a HTTP API.

1.3. Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Privacy Pass HTTP API Wrapping

Messages from HTTP-based clients to HTTP-based servers are performed as GET and POST requests. The messages are sent via the "Sec-Privacy-Pass" header.

"Sec-Privacy-Pass" is a Dictionary Structured Header [draft-ietf-httpbis-header-structure-15]. The dictionary has two keys:

- * "type" whose value is a String conveying the function that is being performed with this request.

* "body" whose value is a byte sequence containing a Privacy Pass protocol message.

Note that the requests may contain addition Headers, request data and URL parameters that are not specified here, these extra fields should be ignored, though may be used by the server to determine whether to fulfill the requested issuance/redemption.

3. Server key registry

A client SHOULD fetch a server's current public key information prior to performing issuance and redemption. This configuration is accessible via a "CONFIG_ENDPOINT", either provided by the server or by a global registry that provides consistency and anonymization guarantees.

3.1. Key Registry

To ensure that a server isn't providing different views of their public key material to different users, servers are expected to write their commitments to a verifiable data structure.

Using a verifiable log-backed map ([verifiable-data-structures]), the server can publish their commitments to the log in a way that clients can detect when the server is attempting to provide a split-view of their key commitments to different clients.

The key to the map is the "server_origin", with the value being:

```
struct {
    opaque public_key<1..2^16-1>;
    uint64 expiry;
    uint8 supported_methods; # 3:Issue/Redeem, 2:Redeem, 1:Issue
    opaque signature<1..2^16-1>;
} KeyCommitment;

struct {
    opaque server_id<1..2^16-1>;
    uint16 ciphersuite;
    opaque verification_key<1..2^16-1>;
    KeyCommitment commitments<1..2^16-1>;
}
```

The addition to the log is made via a signed message to the log operator, which verifies the authenticity against a public key associated with that server origin (either via the Web PKI or a out-of-band key). The signature should be computed under a long-term signing key that is associated with the server identity.

The server SHOULD then store an inclusion proof of the current key commitment so that it can present it when delivering the key commitment directly to the client or when the key commitment is being delivered by a delegated party (other registries/preloaded configuration lists/etc).

The client can then perform a request for the key commitment against either the global registry or the server as described in Section 4. Note that the signature should be verified by the client to ensure that the key material is owned by the server. This requires that the client know the public verification key that is associated with the server.

To avoid user segregation as a result of server configuration/commitment rotation, the log operator SHOULD enforce limits on how many active commitments exist and how quickly the commitments are being rotated. Clients SHOULD reject configurations/commitments that violate their requirements for avoiding user segregation. These considerations are discussed as part of [draft-davidson-pp-architecture].

3.2. Server Configuration Retrieval

Inputs: - "server_origin": The origin to retrieve a server configuration for.

No outputs.

1. The client makes an anonymous GET request to "CONFIG_ENDPOINT"/.well-known/privacy-pass with a message of type "fetch-config" and a body of:

```
struct {  
    opaque server_origin<1..2^16-1>;  
}
```

1. The server looks up the configuration associated with the origin "server_origin" and responds with a message of type "config" and a body of:

```
struct {  
    opaque server_id<1..2^16-1>;  
    uint16 ciphersuite;  
    opaque commitment_id<1..2^8-1>;  
    opaque verification_key<1..2^16-1>;  
}
```

1. The client then stores the associated configuration state under the corresponding "server_origin".

(TODO: This might be mergable with key commitment retrieval if server_id = server_origin)

4. Key Commitment Retrieval

The client SHOULD retrieve server key commitments prior to both an issuance and redemption to verify the consistency of the keys and to monitor for key rotation between issuance and redemption events.

Inputs: - "server_origin": The origin to retrieve a key commitment for.

No outputs.

1. The client fetches the configuration state "server_id", "ciphersuite", "commitment_id" associated with "server_origin".
2. The client makes an anonymous GET request to "CONFIG_ENDPOINT"/.well-known/privacy-pass with a message of type "fetch-commitment" and a body of:

```
struct {  
    opaque server_id<1..2^16-1> = server_id;  
    opaque commitment_id<1..2^8-1> = commitment_id;
```

```
}
```

1. The server looks up the current configuration, and constructs a list of commitments to return, noting whether a key commitment is valid for issuance or redemption or both.
2. The server then responds with a message of type "commitment" and a body of:

```
struct {  
    opaque public_key<1..2^16-1>;  
    uint64 expiry;  
    uint8 supported_methods; # 3:Issue/Redeem, 2:Redeem, 1:Issue  
    opaque signature<1..2^16-1>;  
} KeyCommitment;
```

```
struct {  
    opaque server_id<1..2^16-1>;  
    uint16 ciphersuite;  
    opaque verification_key<1..2^16-1>;  
    KeyCommitment commitments<1..2^16-1>;  
    opaque inclusion_proofs<1..2^16-1>;  
}
```

1. The client then verifies the signature for each key commitment and stores the list of commitments to the current scope. The client SHOULD NOT cache the commitments beyond the current scope, as new commitments should be fetched for each independent issuance and redemption request. The client SHOULD verify the "inclusion_proofs" to confirm that the key commitment has been submitted to a trusted registry. Once the client receives the "ciphersuite" for the server, it should implement all Privacy Pass API functions (as detailed in [draft-davidson-pp-protocol]) using this ciphersuite.

5. Privacy Pass Issuance

Inputs: - "server_origin": The origin to request token issuance from.
- "count": The number of tokens to request issuance for.

Outputs: - "tokens": A list of tokens that have been signed via the Privacy Pass protocol.

1. When a client wants to request tokens from a server, it should first fetch a key commitment from the server via the process described in Section 4 and keep the result as "commitment".
2. The client should then call the "Generate" function requesting "count" tokens storing the resulting "input" data.
3. The client then makes a POST request to <"server_origin">/well-known/privacy-pass with a message of type "request-issuance" and a body of:

```
enum { Normal(0) } IssuanceType;
```

```
struct {  
    IssuanceType type = 0;  
    opaque msg<0..2^16-1> = input.msg;  
}
```

1. The server, upon receipt of the "request" should call the "Issue"

function with the "public_key", "secret_key" and the value of "msg" with a result of "resp".

2. The server should then respond to the POST request with a message of type "issue" and a body of:

```
struct {  
    IssuanceType type = request.type;  
    IssuanceResp resp = resp;  
}
```

1. The client should then should call the "Process" function with the "public_key", stored "inputs" and resulting "resp", to extract a list of "redemption_tokens".
2. The client should store the "public_key" associated with these tokens and the elements of "redemption_tokens" under storage partitioned by the "server_origin", accessible only via the Privacy Pass API.

6. Privacy Pass Redemption

There are two forms of Privacy Pass redemption that could function under the HTTP API. Either passing along a token directly to the target endpoint, which would perform its own redemption Section 6.1, or the client redeeming the token and passing the result along to the target endpoint. These two methods are described below.

6.1. Generic Token Redemption

Inputs: - "server_id": The server ID to redeem a token against. - "ciphersuite": The ciphersuite for this token. - "public_key": The public key associated with this token. - "redemption_token": A Privacy Pass token. - "info": Additional data to bind to this token redemption.

Outputs: - "result": The result of the redemption from the server.

1. The client should call the "Redeem" function with "redemption_token" and additional data of "info" storing the resulting "data" and "tag".
2. The client makes a POST request to <"server_origin">/well-known/privacy-pass with a message of type "token-redemption" and a body of:

```
struct {  
    opaque server_id<1..2^16-1> = server_id;  
    opaque data<1..2^16-1> = data;  
    opaque tag<1..2^16-1> = tag;  
    opaque info<1..2^16-1> = info;  
}
```

1. The server, upon receipt of "request" should call the "Verify" interface with "public_key", "secret_key" and the received "data", "tag", "info" storing the resulting "resp".
2. The server should then respond to the POST request with a message of type "redemption-result" and a signed body of:

```
struct {  
    opaque info<1..2^16-1> = info;
```

```

uint8 result = resp;
// signature of info and result using
// the server's verification key.
opaque signature<1..2^16-1>;
}

```

1. The client upon receipt of this message should verify the "signature" using the "verification_key" from the configuration and return the "result".

6.2. Direct Redemption

Inputs: - "server_origin": The server origin to redeem a token for. -
 "target": The target endpoint to send the token to. -
 "additional_data": Additional data to bind to this redemption request.

1. When a client wants to redeem tokens for a server, it should first fetch a key commitment from the server via the process described in Section 4 and keep the result as "commitment".
2. The client should then look up the storage partition associated with "server_origin" and fetch a "redemption_token" and "public_key".
3. The client should verify that the "public_key" is in the current "commitment". If not, it should discard the token and fail the redemption attempt.
4. As part of the request to "target", the client will include the token as part of the request in the "Sec-Privacy-Pass" header along with whatever other parameters are being passed as part of the request to "target". The header will contain a message of type "token-redemption" with a body of:

```

struct {
  opaque server_id<1..2^16-1> = server_id;
  uint16 ciphersuite = ciphersuite;
  opaque public_key<1..2^16-1> = public_key;
  RedemptionToken token<1..2^16-1> = redemption_token;
  opaque additional_data<1..2^16-1> = additional_data;
}

```

At this point, the "target" can perform a generic redemption as described in Section 6.1 by forwarding the message included in the request to "target".

6.3. Delegated Redemption

Inputs: - "server_origin": The server origin to redeem a token for. -
 "target": The target endpoint to send the token to. -
 "additional_data": Additional data to bind to this redemption request.

1. When a client wants to redeem tokens for a server, it should first fetch a key commitment from the server via the process described in Section 4 and keep the result as "commitment".
2. The client should then look up the storage partition associated with "server_origin" and fetch a "redemption_token" and "public_key".

3. The client should verify that the "public_key" is in the current "commitment". If not, it should discard the token and fail the redemption attempt.
4. The client constructs a bytestring "info" made up of the "target", the current "timestamp", and "additional_data":

```
struct {  
    opaque target<1..2^16-1>;  
    uint64 timestamp;  
    opaque additional_data<0..2^16-1>;  
}
```

1. The client then performs a token redemption as described in Section 6.1. Storing the resulting "redemption-result" message.
2. As part of the request to "target", the client will include the redemption result as part of the request in the "Sec-Privacy-Pass" header along with whatever other parameters are being passed as part of the request to "target". The header will contain a message of type "signed-redemption-result" with a body of:

```
struct {  
    opaque server_origin<1..2^16-1>;  
    opaque target<1..2^16-1>;  
    uint64 timestamp;  
    opaque additional_data<1..2^16-1> = additional_data;  
    opaque signed_redemption<1..2^16-1>;  
}
```

At this point, the "target" can verify the integrity of "signed_redemption.info" based on the values of "target", "timestamp", and "additional_data" and verify the signature of the redemption result by querying the current configuration of the Privacy Pass server. The inclusion of "target" and "timestamp" proves that the server attested to the validity of the token in relation to this particular request.

7. Security Considerations

Security considerations for Privacy Pass are discussed in [draft-davidson-pp-architecture].

8. IANA Considerations

8.1. Well-Known URI

This specification registers a new well-known URI.

URI suffix: "privacy-pass"

Change controller: IETF.

Specification document(s): this specification

9. Normative References

[draft-davidson-pp-architecture]
Davidson, A., "Privacy Pass: Architectural Framework",
n.d., <<https://tools.ietf.org/html/draft-davidson-pp-architecture-00>>.

- [draft-davidson-pp-protocol]
Davidson, A., "Privacy Pass: The Protocol", n.d.,
<<https://tools.ietf.org/html/draft-davidson-pp-protocol-00>>.
- [draft-ietf-httpbis-header-structure-15]
Nottingham, M. and P-H. Kamp, "Structured Headers for HTTP", n.d., <<https://tools.ietf.org/html/draft-ietf-httpbis-header-structure-15>>.
- [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/info/rfc2119>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [verifiable-data-structures]
"Verifiable Data Structures", n.d.,
<<https://github.com/google/trillian/blob/master/docs/papers/VerifiableDataStructures.pdf>>.

Author's Address

Steven Valdez
Google LLC

Email: svaldez@chromium.org

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 13 January 2022

S. Valdez
Google LLC
12 July 2021

Privacy Pass HTTP API
draft-ietf-privacypass-http-api-01

Abstract

This document specifies an integration for Privacy Pass over an HTTP API, along with recommendations on how key commitments are stored and accessed by HTTP-based consumers.

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 13 January 2022.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	2
1.2. Layout	2
1.3. Requirements	3
2. Privacy Pass HTTP API Wrapping	3
3. Server key registry	3
3.1. Key Registry	4
3.2. Server Configuration Retrieval	5
4. Key Commitment Retrieval	5
5. Privacy Pass Issuance	7
6. Privacy Pass Redemption	8
6.1. Generic Token Redemption	8
6.2. Direct Redemption	9
6.3. Delegated Redemption	10
7. Security Considerations	11
8. Privacy considerations	11
9. IANA Considerations	11
9.1. Well-Known URI	11
10. Normative References	11
Author's Address	12

1. Introduction

The Privacy Pass protocol as described in [draft-ietf-privacypass-protocol] can be integrated with a number of different settings, from server to server communication to browsing the internet.

In this document, we will provide an API to use for integrating Privacy Pass with an HTTP framework. Providing the format of HTTP requests and responses needed to implement the Privacy Pass protocol.

1.1. Terminology

We use the same definition of server and client that is used in [draft-ietf-privacypass-protocol] and [draft-ietf-privacypass-architecture].

We assume that all protocol messages are encoded into raw byte format before being sent. We use the TLS presentation language [RFC8446] to describe the structure of protocol messages.

1.2. Layout

- * Section 2: Describes the wrapping of messages within HTTP requests/responses.

- * Section 3: Describes how HTTP clients retrieve server configurations and key commitments.
- * Section 5: Describes how issuance requests are performed via a HTTP API.
- * Section 6: Describes how redemption requests are performed via a HTTP API.

1.3. Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Privacy Pass HTTP API Wrapping

Messages from HTTP-based clients to HTTP-based servers are performed as GET and POST requests. The messages are sent via the "Sec-Privacy-Pass" header.

"Sec-Privacy-Pass" is a Dictionary Structured Header [draft-ietf-httpbis-header-structure-15]. The dictionary has two keys:

- * "type" whose value is a String conveying the function that is being performed with this request.
- * "body" whose value is a byte sequence containing a Privacy Pass protocol message.

Note that the requests may contain addition Headers, request data and URL parameters that are not specified here, these extra fields should be ignored, though may be used by the server to determine whether to fulfill the requested issuance/redemption.

3. Server key registry

A client SHOULD fetch a server's current public key information prior to performing issuance and redemption. This configuration is accessible via a "CONFIG_ENDPOINT", either provided by the server or by a global registry that provides consistency and anonymization guarantees.

3.1. Key Registry

To ensure that a server isn't providing different views of their public key material to different users, servers are expected to write their commitments to a verifiable data structure.

Using a verifiable log-backed map ([verifiable-data-structures]), the server can publish their commitments to the log in a way that clients can detect when the server is attempting to provide a split-view of their key commitments to different clients.

The key to the map is the "server_origin", with the value being:

```
struct {
    opaque public_key<1..2^16-1>;
    uint64 expiry;
    uint8 supported_methods; # 3:Issue/Redeem, 2:Redeem, 1:Issue
    opaque signature<1..2^16-1>;
} KeyCommitment;

struct {
    opaque server_id<1..2^16-1>;
    uint16 ciphersuite;
    opaque verification_key<1..2^16-1>;
    KeyCommitment commitments<1..2^16-1>;
}
```

The addition to the log is made via a signed message to the log operator, which verifies the authenticity against a public key associated with that server origin (either via the Web PKI or a out-of-band key). The signature should be computed under a long-term signing key that is associated with the server identity.

The server SHOULD then store an inclusion proof of the current key commitment so that it can present it when delivering the key commitment directly to the client or when the key commitment is being delivered by a delegated party (other registries/preloaded configuration lists/etc).

The client can then perform a request for the key commitment against either the global registry or the server as described in Section 4. Note that the signature should be verified by the client to ensure that the key material is owned by the server. This requires that the client know the public verification key that is associated with the server.

To avoid user segregation as a result of server configuration/commitment rotation, the log operator SHOULD enforce limits on how many active commitments exist and how quickly the commitments are being rotated. Clients SHOULD reject configurations/commitments that violate their requirements for avoiding user segregation. These considerations are discussed as part of [draft-ietf-privacypass-architecture].

3.2. Server Configuration Retrieval

Inputs: - "server_origin": The origin to retrieve a server configuration for.

No outputs.

1. The client makes an anonymous GET request to "CONFIG_ENDPOINT"/.well-known/privacy-pass with a message of type "fetch-config" and a body of:

```
struct {  
    opaque server_origin<1..2^16-1>;  
}
```

1. The server looks up the configuration associated with the origin "server_origin" and responds with a message of type "config" and a body of:

```
struct {  
    opaque server_id<1..2^16-1>;  
    uint16 ciphersuite;  
    opaque commitment_id<1..2^8-1>;  
    opaque verification_key<1..2^16-1>;  
}
```

1. The client then stores the associated configuration state under the corresponding "server_origin".

(TODO: This might be mergable with key commitment retrieval if server_id = server_origin)

4. Key Commitment Retrieval

The client SHOULD retrieve server key commitments prior to both an issuance and redemption to verify the consistency of the keys and to monitor for key rotation between issuance and redemption events.

Inputs: - "server_origin": The origin to retrieve a key commitment for.

No outputs.

1. The client fetches the configuration state "server_id", "ciphersuite", "commitment_id" associated with "server_origin".
2. The client makes an anonymous GET request to "CONFIG_ENDPOINT"/.well-known/privacy-pass with a message of type "fetch-commitment" and a body of:

```
struct {  
    opaque server_id<1..2^16-1> = server_id;  
    opaque commitment_id<1..2^8-1> = commitment_id;  
}
```

1. The server looks up the current configuration, and constructs a list of commitments to return, noting whether a key commitment is valid for issuance or redemption or both.
2. The server then responds with a message of type "commitment" and a body of:

```
struct {  
    opaque public_key<1..2^16-1>;  
    uint64 expiry;  
    uint8 supported_methods; # 3:Issue/Redeem, 2:Redeem, 1:Issue  
    opaque signature<1..2^16-1>;  
} KeyCommitment;
```

```
struct {  
    opaque server_id<1..2^16-1>;  
    uint16 ciphersuite;  
    opaque verification_key<1..2^16-1>;  
    KeyCommitment commitments<1..2^16-1>;  
    opaque inclusion_proofs<1..2^16-1>;  
}
```

1. The client then verifies the signature for each key commitment and stores the list of commitments to the current scope. The client SHOULD NOT cache the commitments beyond the current scope, as new commitments should be fetched for each independent issuance and redemption request. The client SHOULD verify the "inclusion_proofs" to confirm that the key commitment has been submitted to a trusted registry. Once the client receives the "ciphersuite" for the server, it should implement all Privacy Pass API functions (as detailed in [draft-ietf-privacypass-protocol]) using this ciphersuite.

5. Privacy Pass Issuance

Inputs: - "server_origin": The origin to request token issuance from.
- "count": The number of tokens to request issuance for.

Outputs: - "tokens": A list of tokens that have been signed via the Privacy Pass protocol.

1. When a client wants to request tokens from a server, it should first fetch a key commitment from the server via the process described in Section 4 and keep the result as "commitment".
2. The client should then call the "Generate" function requesting "count" tokens storing the resulting "input" data.
3. The client then makes a POST request to <"server_origin">/well-known/privacy-pass with a message of type "request-issuance" and a body of:

```
enum { Normal(0) } IssuanceType;
```

```
struct {  
    IssuanceType type = 0;  
    opaque msg<0..2^16-1> = input.msg;  
}
```

1. The server, upon receipt of the "request" should call the "Issue" function with the "public_key", "secret_key" and the value of "msg" with a result of "resp".
2. The server should then respond to the POST request with a message of type "issue" and a body of:

```
struct {  
    IssuanceType type = request.type;  
    IssuanceResp resp = resp;  
}
```

1. The client should then should call the "Process" function with the "public_key", stored "inputs" and resulting "resp", to extract a list of "redemption_tokens".
2. The client should store the "public_key" associated with these tokens and the elements of "redemption_tokens" under storage partitioned by the "server_origin", accessible only via the Privacy Pass API.

6. Privacy Pass Redemption

There are two forms of Privacy Pass redemption that could function under the HTTP API. Either passing along a token directly to the target endpoint, which would perform its own redemption Section 6.1, or the client redeeming the token and passing the result along to the target endpoint. These two methods are described below.

In the HTTP ecosystem, redemption contexts should generally be keyed by the same privacy boundary used for cookies and other local storage. Generally this is the top-level origin. Any redemption context should be built following the principles outlined in [draft-ietf-privacypass-architecture] and later in Section 8.

6.1. Generic Token Redemption

Inputs: - "context": The request context to use. - "server_id": The server ID to redeem a token against. - "ciphersuite": The ciphersuite for this token. - "public_key": The public key associated with this token. - "redemption_token": A Privacy Pass token. - "info": Additional data to bind to this token redemption.

Outputs: - "result": The result of the redemption from the server.

1. The client should check whether the "server_id" is present in the "context". If it isn't and the size of the "context" is beneath the client's limit, it should be added.
2. The client should call the "Redeem" function with "redemption_token" and additional data of "info" storing the resulting "data" and "tag".
3. The client makes a POST request to <"server_origin">/well-known/privacy-pass with a message of type "token-redemption" and a body of:

```
struct {  
    opaque server_id<1..2^16-1> = server_id;  
    opaque data<1..2^16-1> = data;  
    opaque tag<1..2^16-1> = tag;  
    opaque info<1..2^16-1> = info;  
}
```

1. The server, upon receipt of "request" should call the "Verify" interface with "public_key", "secret_key" and the received "data", "tag", "info" storing the resulting "resp".

2. The server should then respond to the POST request with a message of type "redemption-result" and a signed body of:

```
struct {  
    opaque info<1..2^16-1> = info;  
    uint8 result = resp;  
    // signature of info and result using  
    // the server's verification key.  
    opaque signature<1..2^16-1>;  
}
```

1. The client upon receipt of this message should verify the "signature" using the "verification_key" from the configuration and return the "result".

6.2. Direct Redemption

Inputs: - "context": The request context to use. - "server_origin": The server origin to redeem a token for. - "target": The target endpoint to send the token to. - "additional_data": Additional data to bind to this redemption request.

1. When a client wants to redeem tokens for a server, it should first fetch a key commitment from the server via the process described in Section 4 and keep the result as "commitment".
2. The client should then look up the storage partition associated with "server_origin" and fetch a "redemption_token" and "public_key".
3. The client should verify that the "public_key" is in the current "commitment". If not, it should discard the token and fail the redemption attempt.
4. As part of the request to "target", the client will include the token as part of the request in the "Sec-Privacy-Pass" header along with whatever other parameters are being passed as part of the request to "target". The header will contain a message of type "token-redemption" with a body of:

```
struct {  
    opaque server_id<1..2^16-1> = server_id;  
    uint16 ciphersuite = ciphersuite;  
    opaque public_key<1..2^16-1> = public_key;  
    RedemptionToken token<1..2^16-1> = redemption_token;  
    opaque additional_data<1..2^16-1> = additional_data;  
}
```

At this point, the "target" can perform a generic redemption as described in Section 6.1 by forwarding the message included in the request to "target".

6.3. Delegated Redemption

Inputs: - "context": The request context to use. - "server_origin": The server origin to redeem a token for. - "target": The target endpoint to send the token to. - "additional_data": Additional data to bind to this redemption request.

1. When a client wants to redeem tokens for a server, it should first fetch a key commitment from the server via the process described in Section 4 and keep the result as "commitment".
2. The client should then look up the storage partition associated with "server_origin" and fetch a "redemption_token" and "public_key".
3. The client should verify that the "public_key" is in the current "commitment". If not, it should discard the token and fail the redemption attempt.
4. The client constructs a bytestring "info" made up of the "target", the current "timestamp", and "additional_data":

```
struct {  
    opaque target<1..2^16-1>;  
    uint64 timestamp;  
    opaque additional_data<0..2^16-1>;  
}
```

1. The client then performs a token redemption as described in Section 6.1. Storing the resulting "redemption-result" message.
2. As part of the request to "target", the client will include the redemption result as part of the request in the "Sec-Privacy-Pass" header along with whatever other parameters are being passed as part of the request to "target". The header will contain a message of type "signed-redemption-result" with a body of:


```
struct {  
    opaque server_origin<1..2^16-1>;  
    opaque target<1..2^16-1>;  
    uint64 timestamp;  
    opaque additional_data<1..2^16-1> = additional_data;  
    opaque signed_redemption<1..2^16-1>;  
}
```

At this point, the "target" can verify the integrity of "signed_redemption.info" based on the values of "target", "timestamp", and "additional_data" and verify the signature of the redemption result by querying the current configuration of the Privacy Pass server. The inclusion of "target" and "timestamp" proves that the server attested to the validity of the token in relation to this particular request.

7. Security Considerations

Security considerations for Privacy Pass are discussed in [draft-ietf-privacypass-architecture].

8. Privacy considerations

General privacy considerations for Privacy Pass are discussed in [draft-ietf-privacypass-architecture].

In order to implement this API with redemption contexts, a client needs to maintain strong privacy boundaries between different redemption contexts to avoid privacy leakage from redemptions across them. Notably in the web/HTTP world, cross-site tracking and fingerprinting will need to be considered and mitigated in order to maintain these privacy boundaries.

9. IANA Considerations

9.1. Well-Known URI

This specification registers a new well-known URI.

URI suffix: "privacy-pass"

Change controller: IETF.

Specification document(s): this specification

10. Normative References

- [draft-ietf-httpbis-header-structure-15]
Nottingham, M. and P-H. Kamp, "Structured Headers for HTTP", n.d., <<https://tools.ietf.org/html/draft-ietf-httpbis-header-structure-15>>.
- [draft-ietf-privacypass-architecture]
Davidson, A., "Privacy Pass: Architectural Framework", n.d., <<https://tools.ietf.org/html/draft-ietf-privacypass-architecture-00>>.
- [draft-ietf-privacypass-protocol]
Davidson, A., "Privacy Pass: The Protocol", n.d., <<https://tools.ietf.org/html/draft-ietf-privacypass-protocol-00>>.
- [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>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/rfc/rfc8446>>.
- [verifiable-data-structures]
"Verifiable Data Structures", n.d., <<https://github.com/google/trillian/blob/master/docs/papers/VerifiableDataStructures.pdf>>.

Author's Address

Steven Valdez
Google LLC

Email: svaldez@chromium.org

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 23 August 2021

S. Celi
Cloudflare
A. Davidson
LIP
A. Faz-Hernandez
Cloudflare
19 February 2021

Privacy Pass Protocol Specification
draft-ietf-privacypass-protocol-01

Abstract

This document specifies the Privacy Pass protocol. This protocol provides anonymity-preserving authorization of clients to servers. In particular, client re-authorization events cannot be linked to any previous initial authorization. Privacy Pass is intended to be used as a performant protocol in the application-layer.

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 23 August 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Background	4
3.1. Motivating use-cases	4
3.2. Anonymity and security guarantees	5
3.3. Basic assumptions	5
4. Protocol description	5
4.1. Server setup	6
4.2. Client setup	6
4.3. Issuance phase	6
4.4. Redemption phase	8
4.4.1. Client info	8
4.4.2. Double-spend protection	9
4.5. Handling errors	9
5. Functionality	9
5.1. Data structures	9
5.1.1. Ciphersuite	9
5.1.2. Keys	9
5.1.3. CommitRequest	10
5.1.4. CommitResponse	10
5.1.5. IssuanceInput	10
5.1.6. IssuanceResponse	10
5.1.7. RedemptionToken	11
5.1.8. RedemptionRequest	11
5.1.9. RedemptionResponse	11
5.2. API functions	12
5.2.1. Prepare	12
5.2.2. Commit	12
5.2.3. Generate	12
5.2.4. Issue	13
5.2.5. Process	13
5.2.6. Redeem	14
5.2.7. Verify	14
5.3. Error types	14
6. Security considerations	15
6.1. Unlinkability	15
6.2. One-more unforgeability	16

6.3.	Double-spend protection	16
6.4.	Additional token metadata	17
6.5.	Maximum number of tokens issued	17
7.	VOPRF instantiation	17
7.1.	Recommended ciphersuites	17
7.2.	Protocol contexts	18
7.3.	Functionality	18
7.3.1.	Generate	18
7.3.2.	Issue	19
7.3.3.	Process	19
7.3.4.	Redeem	19
7.3.5.	Verify	19
7.4.	Security justification	20
8.	Protocol ciphersuites	20
8.1.	PP(OPRF2)	20
8.2.	PP(OPRF4)	21
8.3.	PP(OPRF5)	21
9.	Extensions framework policy	21
10.	References	22
10.1.	Normative References	22
10.2.	Informative References	22
Appendix A.	Document contributors	23
Authors' Addresses	23

1. Introduction

A common problem on the Internet is providing an effective mechanism for servers to derive trust from clients that they interact with. Typically, this can be done by providing some sort of authorization challenge to the client. But this also negatively impacts the experience of clients that regularly have to solve such challenges.

To mitigate accessibility issues, a client that correctly solves the challenge can be provided with a cookie. This cookie can be presented the next time the client interacts with the server, instead of performing the challenge. However, this does not solve the problem of reauthorization of clients across multiple domains. Using current tools, providing some multi-domain authorization token would allow linking client browsing patterns across those domains, and severely reduces their online privacy.

The Privacy Pass protocol provides a set of cross-domain authorization tokens that protect the client's anonymity in message exchanges with a server. This allows clients to communicate an attestation of a previously authenticated server action, without having to reauthenticate manually. The tokens retain anonymity in the sense that the act of revealing them cannot be linked back to the session where they were initially issued.

This document lays out the generic description of the protocol, along with the data and message formats. We detail an implementation of the protocol functionality based on the description of a verifiable oblivious pseudorandom function [I-D.irtf-cfrg-voprf].

This document DOES NOT cover the architectural framework required for running and maintaining the Privacy Pass protocol in the Internet setting. In addition, it DOES NOT cover the choices that are necessary for ensuring that client privacy leaks do not occur. Both of these considerations are covered in a separate document [draft-davidson-pp-architecture]. In addition, [draft-svaldez-pp-http-api] provides an instantiation of this protocol intended for the HTTP setting.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The following terms are used throughout this document.

- * Server: A service that provides the server-side functionality required by the protocol. May be referred to as the issuer.
- * Client: An entity that seeks authorization from a server that supports interactions in the Privacy Pass protocol.
- * Key: The secret key used by the server for authorizing client data.

We assume that all protocol messages are encoded into raw byte format before being sent. We use the TLS presentation language [RFC8446] to describe the structure of protocol data types and messages.

3. Background

We discuss the core motivation behind the protocol along with the guarantees and assumptions that we make in this document.

3.1. Motivating use-cases

The Privacy Pass protocol was originally developed to provide anonymous authorization of Tor users. In particular, the protocol allows clients to reveal authorization tokens that they have been issued without linking the authorization to the actual issuance event. This means that the tokens cannot be used to link the browsing patterns of clients that reveal tokens.

Beyond these uses-cases, the Privacy Pass protocol is used in a number of practical applications. See [DGSTV18], [TrustTokenAPI], [PrivateStorage], [OpenPrivacy], and [Brave] for examples.

3.2. Anonymity and security guarantees

Privacy Pass provides anonymity-preserving authorization tokens for clients. Throughout this document, we use the terms "anonymous", "anonymous-preserving" and "anonymity" to refer to the core security guarantee of the protocol. Informally, this guarantee means that any token issued by a server key and subsequently redeemed is indistinguishable from any other token issued under the same key.

Privacy Pass also prohibits clients from forging tokens, as otherwise the protocol would have little value as an authorization protocol. Informally, this means any client that is issued "N" tokens under a given server key cannot redeem more than "N" valid tokens.

Section 6 elaborates on these protocol anonymity and security requirements.

3.3. Basic assumptions

We make only a few minimal assumptions about the environment of the clients and servers supporting the Privacy Pass protocol.

- * At any one time, we assume that the server uses only one configuration containing their ciphersuite choice along with their secret key data. This ensures that all clients are issued tokens under the single key associated with any given epoch.
- * We assume that the client has access to a global directory of the current public parts of the configurations used the server.

The wider ecosystem that this protocol is employed in is described in [draft-davidson-pp-architecture].

4. Protocol description

The Privacy Pass protocol is split into two phases that are built upon the functionality described in Section 5 later.

The first phase, "issuance", provides the client with unlinkable tokens that can be used to initiate re-authorization with the server in the future. The second phase, "redemption", allows the client to redeem a given re-authorization token with the server that it interacted with during the issuance phase. The protocol must satisfy two cryptographic security requirements known as "unlinkability" and "unforgeability". These requirements are covered in Section 6.

4.1. Server setup

Before the protocol takes place, the server chooses a ciphersuite and generates a keypair by running " $(pkS, skS) = \text{KeyGen}()$ ". This configuration must be available to all clients that interact with the server (for the purpose of engaging in a Privacy Pass exchange). We assume that the server has a public (and unique) identity that the client uses to retrieve this configuration.

4.2. Client setup

The client initialises a global storage system "store" that allows it store the tokens that are received during issuance. The storage system is a map of server identifiers ("server.id") to arrays of stored tokens. We assume that the client knows the server public key "pkS" ahead of time. The client picks a value "m" of tokens to receive during the issuance phase. In [draft-davidson-pp-architecture] we discuss mechanisms that the client can use to ensure that this public key is consistent across the entire ecosystem.

4.3. Issuance phase

The issuance phase is a two-round protocol that allows the client to receive "m" anonymous authorization tokens from the server. The first round sees the server generate a commitment. The second round sees the server issue a token to the client.


```

Client(pkS, m, info)                                Server(skS, pkS)
-----

commit_req = Prepare(info)

                                commit_req
                                ----->

                                commit_resp = Commit(skS, pkS, commit_req)

                                commit_resp
                                <-----

cInput = Generate(m, commit_resp)
req = cInput.req

                                req
                                ----->

                                issueResp = Issue(pkS, skS, req)

                                serverResp
                                <-----

tokens = Process(pkS, cInput, serverResp)
store[server.id].push(tokens)

```

Note that the first round of the protocol is only necessitated for certain ciphersuites that require client and servers commit to some value. When such commitment "commit_resp" is generated and sent to the client, the client returns "commit_resp" with the "IssuanceRequest" message. The server MUST check that the commitment corresponds to "commit_resp" that was previously committed. This requires the commitment to either be a reference to some commitment on the server, or the commitment be an encrypted (and authenticated) blob that the server can use to recover commitment. The mechanism by which servers handle this commitment is implementation specific, and similar to how TLS session resumption state is managed; see [RFC8446] for details. In addition, the "Commit" function is implementation-specific and MUST be defined by the underlying ciphersuite.

When the server does not need to generate this commitment, the client instead DOES NOT send the "CommitRequest" message, and runs:

```
cInput = Generate(m, "")
```

A server that is expecting some non-empty "commit_resp" to be passed must abort the protocol on receiving a request containing an empty "commit_resp" value.

Note: currently, no ciphersuites are supported that support working with empty commitment messages.

4.4. Redemption phase

The redemption phase allows the client to anonymously reauthenticate to the server, using data that it has received from a previous issuance phase.

```

Client(info)                                     Server(skS, pkS)
-----
token = store[server.id].pop()
req = Redeem(token, info)

                                req
                                ----->

                                if (dsIdx.includes(req.data)) {
                                    raise ERR_DOUBLE_SPEND
                                }
                                resp = Verify(pkS, skS, req)
                                if (resp.success) {
                                    dsIdx.push(req.data)
                                }

                                resp
                                <-----
Output resp

```

4.4.1. Client info

The client input "info" is arbitrary byte data that is used for linking the redemption request to the specific session. We RECOMMEND that "info" is constructed as the following concatenated byte-encoded data:

```
len(aux) || aux || len(server.id) || server.id || current_time()
```

where "len(x)" is the length of "x" in bytes, and "aux" is arbitrary auxiliary data chosen by the client. The usage of "current_time()" allows the server to check that the redemption request has happened in an appropriate time window.

4.4.2. Double-spend protection

To protect against clients that attempt to spend a value "req.data" more than once, the server uses an index, "dsIdx", to collect valid inputs it witnesses. Since this store needs to only be optimized for storage and querying, a structure such as a Bloom filter suffices. The storage should be parameterized to live as long as the server keypair that is in use. See Section 6 for more details.

4.5. Handling errors

It is possible for the API functions from Section 5.2 to return one of the errors indicated in Section 5.3 rather than their expected value. In these cases, we assume that the entire protocol aborts.

5. Functionality

This section details the data types and API functions that are used to construct the protocol in Section 4.

We provide an explicit instantiation of the Privacy Pass API in Section 7.3, based on the public API provided in [I-D.irtf-cfrg-voprf].

5.1. Data structures

The following data structures are used throughout the Privacy Pass protocol and are written in the TLS presentation language [RFC8446]. It is intended that any of these data structures can be written into widely-adopted encoding schemes such as those detailed in TLS [RFC8446], CBOR [RFC7049], and JSON [RFC7159].

5.1.1. Ciphersuite

The "Ciphersuite" enum provides identifiers for each of the supported ciphersuites of the protocol. Some initial values that are supported by the core protocol are described in Section 8. Note that the list of supported ciphersuites may be expanded by extensions to the core protocol description in separate documents.

5.1.2. Keys

We use the following types to describe the public and private keys used by the server.

```
opaque PublicKey<1..216-1>  
opaque PrivateKey<1..216-1>
```

5.1.3. CommitRequest

The "CommitRequest" struct is simply a fixed message allowing opaque metadata.

```
struct {  
    opaque info<1..2^16-1>  
} CommitRequest;
```

5.1.4. CommitResponse

The "CommitResponse" struct is contains an opaque set of bytes that correspond to some commitment that the server has generated. The structure and format of this value is implementation specific depending on whether the server is stateful.

```
struct {  
    opaque commitment<1..2^16-1>  
} CommitResponse;
```

5.1.5. IssuanceInput

The "IssuanceInput" struct describes the data that is initially generated by the client during the issuance phase.

Firstly, we define sequences of bytes that partition the client input.

```
opaque Internal<1..2^16-1>  
opaque IssuanceRequest<1..2^16-1>
```

These data types represent members of the wider "IssuanceInput" data type.

```
struct {  
    Internal data[m]  
    IssuanceRequest req[m]  
} IssuanceInput;
```

Note that a "IssuanceInput" contains equal-length arrays of "Internal" and "IssuanceRequest" types corresponding to the number of tokens that should be issued.

5.1.6. IssuanceResponse

Firstly, the "IssuedToken" type corresponds to a single sequence of bytes that represents a single issued token received from the server.

opaque IssuedToken<1..2¹⁶-1>

Then an "IssuanceResponse" corresponds to a collection of "IssuedTokens" as well as a sequence of bytes "proof".

```
struct {  
    IssuedToken tokens[m]  
    opaque proof<1..216-1>  
}
```

The value of "m" is equal to the length of the "IssuanceRequest" vector sent by the client.

5.1.7. RedemptionToken

The "RedemptionToken" struct contains the data required to generate the client message in the redemption phase of the Privacy Pass protocol.

```
struct {  
    opaque data<1..216-1>;  
    opaque issued<1..216-1>;  
} RedemptionToken;
```

5.1.8. RedemptionRequest

The "RedemptionRequest" struct consists of the data that is sent by the client during the redemption phase of the protocol.

```
struct {  
    opaque data<1..216-1>;  
    opaque tag<1..216-1>;  
    opaque info<1..216-1>;  
} RedemptionRequest;
```

5.1.9. RedemptionResponse

The "RedemptionResponse" struct corresponds to a boolean value that indicates whether the "RedemptionRequest" sent by the client is valid. It can also contain any associated data.

```
struct {  
    boolean success;  
    opaque ad<1..216-1>;  
} RedemptionResponse;
```

5.2. API functions

The following functions wrap the core of the functionality required in the Privacy Pass protocol. For each of the descriptions, we essentially provide the function signature, leaving the actual contents to be defined by specific instantiations or extensions of the protocol.

5.2.1. Prepare

A function run by the client to prepare for a commitment will be used during the issuance flow of the Privacy Pass protocol.

Inputs:

"info": An opaque byte application-specific byte string.

Outputs:

"commit_req": A "CommitRequest" struct.

This function should be implemented by any ciphersuites that require a two-phase issuance protocol ("COMMIT=true").

5.2.2. Commit

A function run by the server that generates a commitment in the first phase of the issuance protocol.

Inputs:

* "skS": A server "PrivateKey".

* "pkS": A server "PublicKey".

* "commit_req": A "CommitRequest" struct

Outputs:

* "commit_resp": A "CommitResponse" struct.

This function should be implemented by any ciphersuites that require a two-phase issuance protocol ("COMMIT=true").

5.2.3. Generate

A function run by the client to generate the initial data that is used as its input in the Privacy Pass protocol.

Inputs:

- * "m": A "uint8" value corresponding to the number of Privacy Pass tokens to generate.

Outputs:

- * "input": An "IssuanceInput" struct.

5.2.4. Issue

A function run by the server to issue valid redemption tokens to the client.

Inputs:

- * "pkS": A server "PublicKey".
- * "skS": A server "PrivateKey".
- * "req": An "IssuanceRequest" struct.

Outputs:

- * "resp": An "IssuanceResponse" struct.

Throws:

- * "ERR_FAILED_COMMITMENT" (Section 5.3)

5.2.5. Process

Run by the client when processing the server response in the issuance phase of the protocol.

Inputs:

- * "pkS": An server "PublicKey".
- * "input": An "IssuanceInput" struct.
- * "resp": An "IssuanceResponse" struct.

Outputs:

- * "tokens": A vector of "RedemptionToken" structs, whose length is equal to length of the internal "ServerEvaluation" vector in the "IssuanceResponse" struct.

Throws:

- * "ERR_PROOF_VALIDATION" (Section 5.3)

5.2.6. Redeem

Run by the client in the redemption phase of the protocol to generate the client's message.

Inputs:

- * "token": A "RedemptionToken" struct.
- * "info": An "opaque<1..2¹⁶-1>" type corresponding to data that is linked to the redemption. See Section 4.4.1 for advice on how to construct this.

Outputs:

- * "req": A "RedemptionRequest" struct.

5.2.7. Verify

Run by the server in the redemption phase of the protocol. Determines whether the data sent by the client is valid.

Inputs:

- * "pkS": An server "PublicKey".
- * "skS": An server "PrivateKey".
- * "req": A "RedemptionRequest" struct.

Outputs:

- * "resp": A "RedemptionResponse" struct.

5.3. Error types

- * "ERR_PROOF_VALIDATION": Error occurred when a client attempted to verify the proof that is part of the server's response.
- * "ERR_DOUBLE_SPEND": Error occurred when a client has attempted to redeem a token that has already been used for authorization.

- * "ERR_FAILED_COMMITMENT": Error occurs during issuance phase if non-empty commitment does not match the commitment generated in the first round.

6. Security considerations

We discuss the security requirements that are necessary to uphold when instantiating the Privacy Pass protocol. In particular, we focus on the security requirements of "unlinkability", and "unforgeability". Informally, the notion of unlinkability is required to preserve the anonymity of the client in the redemption phase of the protocol. The notion of unforgeability is to protect against an adversarial client that may look to subvert the security of the protocol.

Both requirements are modelled as typical cryptographic security games, following the formats laid out in [DGSTV18] and [KLOR20].

Note that the privacy requirements of the protocol are covered in the architectural framework document [draft-davidson-pp-architecture].

6.1. Unlinkability

Formally speaking the security model is the following:

- * The adversary runs the server setup and generates a keypair "(pkS, skS)".
- * The adversary specifies a number "Q" of issuance phases to initiate, where each phase "i in range(Q)" consists of "m_i" Issue evaluations.
- * The adversary runs "Issue" using the keypair that it generated on each of the client messages in the issuance phase.
- * When the adversary wants, it stops the issuance phase, and a random number "l" is picked from "range(Q)".
- * A redemption phase is initiated with a single token with index "i" randomly sampled from "range(m_l)".
- * The adversary guesses an index "l_guess" corresponding to the index of the issuance phase that it believes the redemption token was received in.
- * The adversary succeeds if "l == l_guess".

The security requirement is that the adversary has only a negligible probability of success greater than $1/Q$.

6.2. One-more unforgeability

The one-more unforgeability requirement states that it is hard for any adversarial client that has received m valid tokens from the issuance phase to redeem $m+1$ of them. In essence, this requirement prevents a malicious client from being able to forge valid tokens based on the Issue responses that it sees.

The security model roughly takes the following form:

- * The adversary specifies a number Q of issuance phases to initiate with the server, where each phase i in $\text{range}(Q)$ consists of m_i server evaluation. Let $m = \sum(m_i)$ where i in $\text{range}(Q)$.
- * The adversary receives Q responses, where the response with index i contains m_i individual tokens.
- * The adversary initiates m_{adv} redemption sessions with the server and the server verifies that the sessions are successful (return true), and that each request includes a unique token. The adversary succeeds in $m_{\text{succ}} \leq m_{\text{adv}}$ redemption sessions.
- * The adversary succeeds if $m_{\text{succ}} > m$.

The security requirement is that the adversarial client has only a negligible probability of succeeding.

Note that [KLOR20] strengthens the capabilities of the adversary, in comparison to the original work of [DGSTV18]. In [KLOR20], the adversary is provided with oracle access that allows it to verify that the server responses in the issuance phase are valid.

6.3. Double-spend protection

All issuing servers should implement a robust, global storage-query mechanism for checking that tokens sent by clients have not been spent before. Such tokens only need to be checked for each server individually. This prevents clients from "replaying" previous requests, and is necessary for achieving the unforgeability requirement.

6.4. Additional token metadata

Some use-cases of the Privacy Pass protocol benefit from associating a limited amount of metadata with tokens that can be read by the server when a token is redeemed. Adding metadata to tokens can be used as a vector to segment the anonymity of the client in the protocol. Therefore, it is important that any metadata that is added is heavily limited.

Any additional metadata that can be added to redemption tokens should be described in the specific protocol instantiation. Note that any additional metadata will have to be justified in light of the privacy concerns raised above. For more details on the impacts associated with segmenting user privacy, see [draft-davidson-pp-architecture].

Any metadata added to tokens will be considered either "public" or "private". Public metadata corresponds to unmodifiable bits that a client can read. Private metadata corresponds to unmodifiable private bits that should be obscured to the client.

Note that the instantiation in Section 7 provides randomized redemption tokens with no additional metadata for an server with a single key.

6.5. Maximum number of tokens issued

Servers SHOULD impose a hard ceiling on the number of tokens that can be issued in a single issuance phase to a client. If there is no limit, malicious clients could abuse this and cause excessive computation, leading to a Denial-of-Service attack.

7. VOPRF instantiation

In this section, we show how to instantiate the functional API in Section 5 with the VOPRF protocol described in [I-D.irtf-cfrg-voprf]. Moreover, we show that this protocol satisfies the security requirements laid out in Section 6, based on the security proofs provided in [DGSTV18] and [KLOR20].

7.1. Recommended ciphersuites

The RECOMMENDED server ciphersuites are as follows: detailed in [I-D.irtf-cfrg-voprf]:

- * OPFRF(dec448, SHA-512) (ID = 0x0002);
- * OPFRF(P-384, SHA-512) (ID = 0x0004);

* OPRF(P-521, SHA-512) (ID = 0x0005).

We deliberately avoid the usage of smaller ciphersuites (associated with P-256 and ristretto255) due to the potential to reduce security to unfavourable levels via static Diffie Hellman attacks. See [I-D.irtf-cfrg-voprf] for more details.

7.2. Protocol contexts

Note that we must run the verifiable version of the protocol in [I-D.irtf-cfrg-voprf]. Therefore the "server" takes the role of the "Server" running in "modeVerifiable". In other words, the "server" runs "ctxtS = SetupVerifiableServer(suite, skS, pkS)"; where "suite" is one of the ciphersuites in Section 7.1, "skS" and "pkS" is the server's secret and public key respectively (generated by calling "KeyGen"). It returns "ctxtS", which is the Server context. Likewise, run "ctxtC = SetupVerifiableClient(suite, pkS)" to generate the Client context.

7.3. Functionality

We instantiate each functions using the API functions in [I-D.irtf-cfrg-voprf]. Note that we use the framework mentioned in the document to allow for batching multiple tokens into a single VOPRF evaluation. For the explicit signatures of each of the functions, refer to Section 5.

7.3.1. Generate

The generate functionality generates an initial set of tokens and blinded representation on the client-side. The function also takes an optional (possibly empty) value for a commitment "com" committed to by the server.

```
def Generate(m, com):
    tokens = []
    blindedTokens = []
    for i in range(m):
        x = random_bytes()
        (token, blindedToken) = Blind(x, com)
        tokens[i] = token
        blindedTokens[i] = blindedToken
    return IssuanceInput {
        data: tokens,
        req: blindedTokens,
    }
```

7.3.2. Issue

For this functionality, note that we supply multiple tokens in "req" to "Evaluate". This allows batching a single proof object for multiple evaluations. While the construction in [I-D.irtf-cfrg-vopr] only permits a single input, we follow the advice for providing vectors of inputs.

```
def Issue(pkS, skS, req):
    elements, proof = Evaluate(skS, pkS, req)
    return IssuanceResponse {
        tokens: elements,
        proof: proof,
    }
```

7.3.3. Process

Similarly to "Issue", we follow the advice for providing vectors of inputs to the "Unblind" function for verifying the batched proof object.

```
Process(pkS, input, resp):
    unblindedTokens = Unblind(input.data, resp.elements,
                               input.req, pkS, resp.proof)

    redemptionTokens = []
    for bt in unblindedTokens:
        rt = RedemptionToken { data: input.data, issued: bt }
        redemptionTokens[i] = rt
    return redemptionTokens
```

7.3.4. Redeem

```
def Redeem(token, info):
    tag = Finalize(token.data, token.issued, info)
    return RedemptionRequest {
        data: token.data,
        tag: tag,
        info: info,
    }
```

7.3.5. Verify

```
def Verify(pkS, skS, req):
    resp = VerifyFinalize(skS, req.data, req.info, req.tag)
    Output RedemptionResponse {
        success: resp
    }
```

7.4. Security justification

The protocol devised in Section 4, coupled with the API instantiation in Section 7.3, are equivalent to the protocol description in [DGSTV18] and [KLOR20] from a security perspective. In [DGSTV18], it is proven that this protocol satisfies the security requirements of unlinkability (Section 6.1) and unforgeability (Section 6.2).

The unlinkability property follows unconditionally as the view of the adversary in the redemption phase is distributed independently of the issuance phase. The unforgeability property follows from the one-more decryption security of the ElGamal cryptosystem [DGSTV18]. In [KLOR20] it is also proven that this protocol satisfies the stronger notion of unforgeability, where the adversary is granted a verification oracle, under the chosen-target Diffie-Hellman assumption.

Note that the existing security proofs do not leverage the VOPRF primitive as a black-box in the security reductions. Instead, it relies on the underlying operations in a non-black-box manner. Hence, an explicit reduction from the generic VOPRF primitive to the Privacy Pass protocol would strengthen these security guarantees.

8. Protocol ciphersuites

The ciphersuites that we describe for the Privacy Pass protocol are derived from the core instantiations of the protocol (such as in Section 7).

In each of the ciphersuites below, the maximum security provided corresponds to the maximum difficulty of computing a discrete logarithm in the group. Note that the actual security level MAY be lower. See the security considerations in [I-D.irtf-cfrg-voprf] for examples.

The COMMIT parameter refers to whether the first round of the issuance phase of the protocol is necessary. When this is false, the client ignores the first message and uses an empty value for the commitment parameter to "Generate".

8.1. PP(OPRF2)

- * OPRF2 = OPRF(decaf448, SHA-512)
- * ID = 0x0001
- * COMMIT = false

- * Maximum security provided: 224 bits

8.2. PP (OPRF4)

- * $\text{OPRF4} = \text{OPRF}(\text{P-384}, \text{SHA-512})$
- * ID = 0x0002
- * COMMIT = false
- * Maximum security provided: 192 bits

8.3. PP (OPRF5)

- * $\text{OPRF5} = \text{OPRF}(\text{P-521}, \text{SHA-512})$
- * ID = 0x0003
- * COMMIT = false
- * Maximum security provided: 256 bits

9. Extensions framework policy

The intention with providing the Privacy Pass API in Section 5 is to allow new instantiations of the Privacy Pass protocol. These instantiations may provide either modified VOPRF constructions, or simply implement the API in a completely different way.

Extensions to this initial draft SHOULD be specified as separate documents taking one of two possible routes:

- * Produce new VOPRF-like primitives that use the same public API provided in [I-D.irtf-cfrg-voprf] to implement the Privacy Pass API, but with different internal operations.
- * Implement the Privacy Pass API in a different way to the proposed implementation in Section 7.

If an extension requires changing the generic protocol description as described in Section 4, then the change may have to result in changes to the draft specification here also.

Each new extension that modifies the internals of the protocol in either of the two ways MUST re-justify that the extended protocol still satisfies the security requirements in Section 6. Protocol extensions MAY put forward new security guarantees if they are applicable.

The extensions MUST also conform with the extension framework policy as set out in the architectural framework document. For example, this may concern any potential impact on client anonymity that the extension may introduce.

10. References

10.1. Normative References

- [draft-davidson-pp-architecture]
Davidson, A., "Privacy Pass: Architectural Framework", n.d., <<https://github.com/alxdavids/privacy-pass-ietf/tree/master/drafts/draft-davidson-pp-architecture>>.
- [draft-svaldez-pp-http-api]
Valdez, S., "Privacy Pass: HTTP API", n.d., <<https://github.com/alxdavids/privacy-pass-ietf/tree/master/drafts/draft-davidson-pp-architecture>>.
- [I-D.irtf-cfrg-voprf]
Davidson, A., Faz-Hernandez, A., Sullivan, N., and C. Wood, "Oblivious Pseudorandom Functions (OPRFs) using Prime-Order Groups", Work in Progress, Internet-Draft, draft-irtf-cfrg-voprf-05, 2 November 2020, <<http://www.ietf.org/internet-drafts/draft-irtf-cfrg-voprf-05.txt>>.
- [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/info/rfc2119>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.

10.2. Informative References

- [Brave] "Brave Rewards", n.d., <<https://brave.com/brave-rewards/>>.
- [DGSTV18] "Privacy Pass, Bypassing Internet Challenges Anonymously", n.d., <<https://petsymposium.org/2018/files/papers/issue3/popets-2018-0026.pdf>>.
- [KLOR20] "Anonymous Tokens with Private Metadata Bit", n.d., <<https://eprint.iacr.org/2020/072>>.

[OpenPrivacy]

"Token Based Services - Differences from PrivacyPass",
n.d., <<https://openprivacy.ca/assets/towards-anonymous-prepaid-services.pdf>>.

[PrivateStorage]

Steininger, L., "The Path from S4 to PrivateStorage",
n.d., <<https://medium.com/least-authority/the-path-from-s4-to-privatestorage-ae9d4a10b2ae>>.

[RFC7049] Bormann, C. and P. Hoffman, "Concise Binary Object
Representation (CBOR)", RFC 7049, DOI 10.17487/RFC7049,
October 2013, <<https://www.rfc-editor.org/info/rfc7049>>.

[RFC7159] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data
Interchange Format", RFC 7159, DOI 10.17487/RFC7159, March
2014, <<https://www.rfc-editor.org/info/rfc7159>>.

[TrustTokenAPI]

WICG, ., "Trust Token API", n.d.,
<<https://github.com/WICG/trust-token-api>>.

Appendix A. Document contributors

- * Alex Davidson (alex.davidson92@gmail.com)
- * Sofia Celi (cherenkov@riseup.net)
- * Christopher Wood (caw@heapingbits.net)

Authors' Addresses

Sofía Celi
Cloudflare
Lisbon
Portugal

Email: sceli@cloudflare.com

Alex Davidson
LIP
Lisbon
Portugal

Email: alex.davidson92@gmail.com

Armando Faz-Hernandez
Cloudflare
101 Townsend St
San Francisco,
United States of America

Email: armfazh@cloudflare.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 7 October 2022

S. Celi
Cloudflare
A. Davidson
Brave Software
A. Faz-Hernandez
Cloudflare
S. Valdez
Google LLC
C. A. Wood
Cloudflare
5 April 2022

Privacy Pass Issuance Protocol
draft-ietf-privacypass-protocol-04

Abstract

This document specifies two variants of the two-message issuance protocol for Privacy Pass tokens: one that produces tokens that are privately verifiable, and another that produces tokens that are publicly verifiable. The privately verifiable issuance protocol optionally supports public metadata during the issuance flow.

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 7 October 2022.

Copyright Notice

Copyright (c) 2022 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
2. Terminology	3
3. Configuration	3
4. Token Challenge Requirements	4
5. Issuance Protocol for Privately Verifiable Tokens with Public Metadata	4
5.1. Client-to-Issuer Request	5
5.2. Issuer-to-Client Response	6
5.3. Finalization	7
5.4. Issuer Configuration	8
6. Issuance Protocol for Publicly Verifiable Tokens	8
6.1. Client-to-Issuer Request	9
6.2. Issuer-to-Client Response	10
6.3. Finalization	11
6.4. Issuer Configuration	11
7. Security considerations	11
8. IANA considerations	11
8.1. Token Type	12
8.2. Media Types	12
8.2.1. "message/token-request" media type	12
8.2.2. "message/token-response" media type	13
9. Normative References	14
Appendix A. Acknowledgements	15
Appendix B. Test Vectors	15
B.1. Issuance Protocol 1 - VOPRF(P-384, SHA-384)	15
B.2. Issuance Protocol 2 - Blind RSA, 4096	16
Authors' Addresses	20

1. Introduction

The Privacy Pass protocol provides a privacy-preserving authorization mechanism. In essence, the protocol allows clients to provide cryptographic tokens that prove nothing other than that they have been created by a given server in the past [I-D.ietf-privacypass-architecture].

This document describes the issuance protocol for Privacy Pass. It specifies two variants: one that is privately verifiable based on the oblivious pseudorandom function from [OPRF], and one that is publicly verifiable based on the blind RSA signature scheme [BLINDRSA].

This document DOES NOT cover the architectural framework required for running and maintaining the Privacy Pass protocol in the Internet setting. In addition, it DOES NOT cover the choices that are necessary for ensuring that client privacy leaks do not occur. Both of these considerations are covered in [I-D.ietf-privacypass-architecture].

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 provides authorization tokens to services across the Internet, in return for authorization.
- * Issuer: A service produces Privacy Pass tokens to clients.
- * Private Key: The secret key used by the Issuer for issuing tokens.
- * Public Key: The public key used by the Issuer for issuing and verifying tokens.

We assume that all protocol messages are encoded into raw byte format before being sent across the wire.

3. Configuration

Issuers MUST provide one parameter for configuration:

1. Issuer Request URI: a token request URL for generating access tokens. For example, an Issuer URL might be `https://issuer.example.net/example-token-request`. This parameter uses resource media type "text/plain".

The Issuer parameters can be obtained from an Issuer via a directory object, which is a JSON object whose field names and values are raw values and URLs for the parameters.

Field Name	Value
issuer-request-uri	Issuer Request URI resource URL as a JSON string

Table 1

As an example, the Issuer's JSON directory could look like:

```
{
  "issuer-request-uri": "https://issuer.example.net/example-token-request"
}
```


Issuer directory resources have the media type "application/json" and are located at the well-known location /.well-known/token-issuer-directory.

4. Token Challenge Requirements

Clients receive challenges for tokens, as described in [AUTHSCHEME]. The basic token issuance protocols described in this document can be interactive or non-interactive, and per-origin or cross-origin.

5. Issuance Protocol for Privately Verifiable Tokens with Public Metadata

The Privacy Pass issuance protocol is a two message protocol that takes as input a challenge from the redemption protocol and produces a token, as shown in the figure below.



Issuers provide a Private and Public Key, denoted *skI* and *pkI*, respectively, used to produce tokens as input to the protocol. See Section 5.4 for how this key pair is generated.

Clients provide the following as input to the issuance protocol:

- * Issuer name, identifying the Issuer. This is typically a host name that can be used to construct HTTP requests to the Issuer.

- * Issuer Public Key `pkI`, with a key identifier `key_id` computed as described in Section 5.4.
- * Challenge value `challenge`, an opaque byte string. For example, this might be provided by the redemption protocol in [HTTP-Authentication].

Given this configuration and these inputs, the two messages exchanged in this protocol are described below. This section uses notation described in [OPRF], Section 4, including `SerializeElement` and `DeserializeElement`, `SerializeScalar` and `DeserializeScalar`, and `DeriveKeyPair`.

5.1. Client-to-Issuer Request

The Client first creates a context as follows:

```
client_context = SetupVOPRFClient(0x0004, pkI)
```

Here, `0x0004` is the two-octet identifier corresponding to the OPRF(P-384, SHA-384) ciphersuite in [OPRF]. `SetupVOPRFClient` is defined in [OPRF], Section 3.2.

The Client then creates an issuance request message for a random value nonce using the input challenge and Issuer key identifier as follows:

```
nonce = random(32)
context = SHA256(challenge)
token_input = concat(0x0001, nonce, context, key_id)
blind, blinded_element = client_context.Blind(token_input)
```

The `Blind` function is defined in [OPRF], Section 3.3.2. If the `Blind` function fails, the Client aborts the protocol. Otherwise, the Client then creates a `TokenRequest` structured as follows:

```
struct {
    uint16_t token_type = 0x0001;
    uint8_t token_key_id;
    uint8_t blinded_msg[Ne];
} TokenRequest;
```

The structure fields are defined as follows:

- * "`token_type`" is a 2-octet integer, which matches the type in the challenge.
- * "`token_key_id`" is the least significant byte of the `key_id`.

- * "blinded_msg" is the Ne-octet blinded message defined above, computed as `SerializeElement(blinded_element)`. Ne is as defined in [OPRF], Section 4.

The values `token_input` and `blinded_element` are stored locally and used later as described in Section 5.3. The Client then generates an HTTP POST request to send to the Issuer, with the `TokenRequest` as the body. The media type for this request is "message/token-request". An example request is shown below.

```
:method = POST
:scheme = https
:authority = issuer.example.net
:path = /example-token-request
accept = message/token-response
cache-control = no-cache, no-store
content-type = message/token-request
content-length = <Length of TokenRequest>
```

<Bytes containing the `TokenRequest`>

Upon receipt of the request, the Issuer validates the following conditions:

- * The `TokenRequest` contains a supported `token_type`.
- * The `TokenRequest.token_key_id` corresponds to a key ID of a Public Key owned by the issuer.
- * The `TokenRequest.blinded_request` is of the correct size.

If any of these conditions is not met, the Issuer MUST return an HTTP 400 error to the client.

5.2. Issuer-to-Client Response

Upon receipt of a `TokenRequest`, the Issuer tries to deseralize `TokenRequest.blinded_msg` using `DeserializeElement` from Section 2.1 of [OPRF], yielding `blinded_element`. If this fails, the Issuer MUST return an HTTP 400 error to the client. Otherwise, if the Issuer is willing to produce a token token to the Client, the Issuer completes the issuance flow by computing a blinded response as follows:

```
server_context = SetupVOPRFServer(0x0004, skI, pkI)
evaluate_element, proof = server_context.Evaluate(skI, blinded_element)
```


SetupVOPRFServer is in [OPRF], Section 3.2 and Evaluate is defined in [OPRF], Section 3.3.2. The Issuer then creates a TokenResponse structured as follows:

```
struct {  
    uint8_t evaluate_msg[Nk];  
    uint8_t evaluate_proof[Ns+Ns];  
} TokenResponse;
```

The structure fields are defined as follows:

- * "evaluate_msg" is the Ne-octet evaluated messaged, computed as `SerializeElement(evaluate_element)`.
- * "evaluate_proof" is the (Ns+Ns)-octet serialized proof, which is a pair of Scalar values, computed as `concat(SerializeScalar(proof[0]), SerializeScalar(proof[1]))`, where Ns is as defined in [OPRF], Section 4.

The Issuer generates an HTTP response with status code 200 whose body consists of TokenResponse, with the content type set as "message/token-response".

```
:status = 200  
content-type = message/token-response  
content-length = <Length of TokenResponse>
```

<Bytes containing the TokenResponse>

5.3. Finalization

Upon receipt, the Client handles the response and, if successful, deserializes the body values `TokenResponse.evaluate_response` and `TokenResponse.evaluate_proof`, yielding `evaluated_element` and `proof`. If deserialization of either value fails, the Client aborts the protocol. Otherwise, the Client processes the response as follows:

```
authenticator = client_context.Finalize(token_input, blind, evaluated_element, bl  
inded_element, proof)
```

The Finalize function is defined in [OPRF], Section 3.3.2. If this succeeds, the Client then constructs a Token as follows:

```
struct {  
    uint16_t token_type = 0x0001  
    uint8_t nonce[32];  
    uint8_t challenge_digest[32];  
    uint8_t token_key_id[32];  
    uint8_t authenticator[Nk];  
} Token;
```

Otherwise, the Client aborts the protocol.

5.4. Issuer Configuration

Issuers are configured with Private and Public Key pairs, each denoted `skI` and `pkI`, respectively, used to produce tokens. Each key pair MUST be generated as follows:

```
seed = random(Ns)  
(skI, pkI) = DeriveKeyPair(seed, "PrivacyPass")
```

The key identifier for this specific key pair, denoted `key_id`, is computed as follows:

```
key_id = SHA256(0x0001 || SerializeElement(pkI))
```

6. Issuance Protocol for Publicly Verifiable Tokens

This section describes a variant of the issuance protocol in Section 5 for producing publicly verifiable tokens. It differs from the previous variant in two important ways:

1. The output tokens are publicly verifiable by anyone with the Issuer public key; and
2. The issuance protocol does not admit public or private metadata to bind additional context to tokens.

Otherwise, this variant is nearly identical. In particular, Issuers provide a Private and Public Key, denoted `skI` and `pkI`, respectively, used to produce tokens as input to the protocol. See Section 6.4 for how this key pair is generated.

Clients provide the following as input to the issuance protocol:

- * Issuer name, identifying the Issuer. This is typically a host name that can be used to construct HTTP requests to the Issuer.
- * Issuer Public Key `pkI`, with a key identifier `key_id` computed as described in Section 6.4.

- * Challenge value challenge, an opaque byte string. For example, this might be provided by the redemption protocol in [HTTP-Authentication].

Given this configuration and these inputs, the two messages exchanged in this protocol are described below.

6.1. Client-to-Issuer Request

The Client first creates an issuance request message for a random value nonce using the input challenge and Issuer key identifier as follows:

```
nonce = random(32)
context = SHA256(challenge)
token_input = concat(0x0002, nonce, context, key_id)
blinded_msg, blind_inv = rsabssa_blind(pkI, token_input)
```

The rsabssa_blind function is defined in [BLINDRSA], Section 5.1.1.. The Client then creates a TokenRequest structured as follows:

```
struct {
    uint16_t token_type = 0x0002
    uint8_t token_key_id;
    uint8_t blinded_msg[Nk];
} TokenRequest;
```

The structure fields are defined as follows:

- * "token_type" is a 2-octet integer, which matches the type in the challenge.
- * "token_key_id" is the least significant byte of the key_id.
- * "blinded_msg" is the Nk-octet request defined above.

The Client then generates an HTTP POST request to send to the Issuer, with the TokenRequest as the body. The media type for this request is "message/token-request". An example request is shown below, where Nk = 512.

```
:method = POST
:scheme = https
:authority = issuer.example.net
:path = /example-token-request
:accept = message/token-response
:cache-control = no-cache, no-store
:content-type = message/token-request
:content-length = <Length of TokenRequest>
```

<Bytes containing the TokenRequest>

Upon receipt of the request, the Issuer validates the following conditions:

- * The TokenRequest contains a supported token_type.
- * The TokenRequest.token_key_id corresponds to a key ID of a Public Key owned by the issuer.
- * The TokenRequest.blinded_msg is of the correct size.

If any of these conditions is not met, the Issuer MUST return an HTTP 400 error to the Client, which will forward the error to the client.

6.2. Issuer-to-Client Response

If the Issuer is willing to produce a token token to the Client, the Issuer completes the issuance flow by computing a blinded response as follows:

```
blind_sig = rsabssa_blind_sign(skI, TokenRequest.blinded_rmsg)
```

This is encoded and transmitted to the client in the following TokenResponse structure:

```
struct {
    uint8_t blind_sig[Nk];
} TokenResponse;
```

The rsabssa_blind_sign function is defined in [BLINDRSA], Section 5.1.2.. The Issuer generates an HTTP response with status code 200 whose body consists of TokenResponse, with the content type set as "message/token-response".

```
:status = 200
content-type = message/token-response
content-length = <Length of TokenResponse>

<Bytes containing the TokenResponse>
```

6.3. Finalization

Upon receipt, the Client handles the response and, if successful, processes the body as follows:

```
authenticator = rsabssa_finalize(pkI, nonce, blind_sig, blind_inv)
```

The `rsabssa_finalize` function is defined in [BLINDRSA], Section 5.1.3.. If this succeeds, the Client then constructs a Token as described in [HTTP-Authentication] as follows:

```
struct {
    uint16_t token_type = 0x0002;
    uint8_t nonce[32];
    uint8_t challenge_digest[32];
    uint8_t token_key_id[32];
    uint8_t authenticator[Nk];
} Token;
```

Otherwise, the Client aborts the protocol.

6.4. Issuer Configuration

Issuers are configured with Private and Public Key pairs, each denoted `skI` and `pkI`, respectively, used to produce tokens. Each key pair MUST be generated as a valid 4096-bit RSA private key according to [TODO]. The key identifier for a keypair (`skI`, `pkI`), denoted `key_id`, is computed as `SHA256(encoded_key)`, where `encoded_key` is a DER-encoded `SubjectPublicKeyInfo` object carrying `pkI`.

7. Security considerations

This document outlines how to instantiate the Issuance protocol based on the VOPRF defined in [OPRF] and blind RSA protocol defined in [BLINDRSA]. All security considerations described in the VOPRF document also apply in the Privacy Pass use-case. Considerations related to broader privacy and security concerns in a multi-Client and multi-Issuer setting are deferred to the Architecture document [I-D.ietf-privacypass-architecture].

8. IANA considerations

8.1. Token Type

This document updates the "Token Type" Registry with the following values.

Value	Name	Publicly Verifiable	Public Metadata	Private Metadata	Nk	Reference
0x0001	VOPRF (P-384, SHA-384)	N	N	N	48	Section 5
0x0002	Blind RSA, 4096	Y	N	N	512	Section 6

Table 2: Token Types

8.2. Media Types

This specification defines the following protocol messages, along with their corresponding media types:

- * TokenRequest: "message/token-request"
- * TokenResponse: "message/token-response"

The definition for each media type is in the following subsections.

8.2.1. "message/token-request" media type

Type name: message

Subtype name: token-request

Required parameters: N/A

Optional parameters: None

Encoding considerations: only "8bit" or "binary" is permitted

Security considerations: see Section 7

Interoperability considerations: N/A

Published specification: this specification

Applications that use this media type: N/A

Fragment identifier considerations: N/A

Additional information: Magic number(s): N/A

Deprecated alias names for this type: N/A

File extension(s): N/A

Macintosh file type code(s): N/A

Person and email address to contact for further information: see Authors' Addresses section

Intended usage: COMMON

Restrictions on usage: N/A

Author: see Authors' Addresses section

Change controller: IESG

8.2.2. "message/token-response" media type

Type name: message

Subtype name: access-token-response

Required parameters: N/A

Optional parameters: None

Encoding considerations: only "8bit" or "binary" is permitted

Security considerations: see Section 7

Interoperability considerations: N/A

Published specification: this specification

Applications that use this media type: N/A

Fragment identifier considerations: N/A

Additional information: Magic number(s): N/A

Deprecated alias names for this type: N/A

File extension(s): N/A

Macintosh file type code(s): N/A

Person and email address to contact for further information: see Authors' Addresses section

Intended usage: COMMON

Restrictions on usage: N/A

Author: see Authors' Addresses section

Change controller: IESG

9. Normative References

[AUTHSCHEME]

Pauly, T., Valdez, S., and C. A. Wood, "The Privacy Pass HTTP Authentication Scheme", Work in Progress, Internet-Draft, draft-pauly-privacypass-auth-scheme-00, 31 January 2022, <<https://datatracker.ietf.org/doc/html/draft-pauly-privacypass-auth-scheme-00>>.

[BLINDRSA] Denis, F., Jacobs, F., and C. A. Wood, "RSA Blind Signatures", Work in Progress, Internet-Draft, draft-irtf-cfrg-rsa-blind-signatures-03, 2 February 2022, <<https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-rsa-blind-signatures-03>>.

[HTTP-Authentication]

"The Privacy Pass HTTP Authentication Scheme", n.d., <<https://datatracker.ietf.org/doc/html/draft-pauly-privacypass-auth-scheme-00>>.

[I-D.ietf-privacypass-architecture]

Davidson, A., Iyengar, J., and C. A. Wood, "Privacy Pass Architectural Framework", Work in Progress, Internet-Draft, draft-ietf-privacypass-architecture-03, 7 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-architecture-03>>.

[OPRF]

Davidson, A., Faz-Hernandez, A., Sullivan, N., and C. A. Wood, "Oblivious Pseudorandom Functions (OPRFs) using Prime-Order Groups", Work in Progress, Internet-Draft, draft-irtf-cfrg-voprf-09, 8 February 2022, <<https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-voprf-09>>.

- [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>>.

Appendix A. Acknowledgements

The authors of this document would like to acknowledge the helpful feedback and discussions from Benjamin Schwartz, Joseph Salowey, Sofia Celi, and Tara Whalen.

Appendix B. Test Vectors

This section includes test vectors for the two basic issuance protocols specified in this document. Appendix B.1 contains test vectors for token issuance protocol 1 (0x0001), and Appendix B.2 contains test vectors for token issuance protocol 2 (0x0002).

B.1. Issuance Protocol 1 - VOPRF(P-384, SHA-384)

The test vector below lists the following values:

- * skS: The encoded OPRF private key, serialized using `SerializeScalar` from Section 2.1 of [OPRF] and represented as a hexadecimal string.
- * pkS: The encoded OPRF public key, serialized using `SerializeElement` from Section 2.1 of [OPRF] and represented as a hexadecimal string.
- * challenge: A random challenge, represented as a hexadecimal string.
- * nonce: The 32-byte client nonce generated according to Section 5.1, represented as a hexadecimal string.
- * blind: The blind used when computing the OPRF blinded message, serialized using `SerializeScalar` from Section 2.1 of [OPRF] and represented as a hexadecimal string.
- * token_request: The `TokenRequest` message constructed according to Section 5.1, represented as a hexadecimal string.

- * token_request: The TokenResponse message constructed according to Section 5.2, represented as a hexadecimal string.
- * token: The output Token from the protocol, represented as a hexadecimal string.

```

skS: 0177781aeced893dcccdf80713d318a801e2a0498240fdcf650304bbbfd0f8d3b5c0
cf6cfee457aaa983ec02ff283b7a9
pkS: 022c63f79ac59c0ba3d204245f676a2133bd6120c90d67afa05cd6f8614294b7366
c252c6458300551b79a4911c2590a36
challenge:
a5d46383359ef34e3c4a7b8d1b3165778bffc9b70c9e6a60dd14143e4c9c9fbd
nonce: 5d4799f8338ddc50a6685f83b8ecd264b2f157015229d12b3384c0f199efe7b8
blind: 0322fec505230992256296063d989b59cc03e83184eb6187076d264137622d202
48e4e525bdc007b80d1560e0a6f49d9
token_request: 00011a02861fd50d14be873611cfff0131d2c872c79d0260c6763498a2
a3f14ca926009c0f247653406eld52b68d61b7ed2bac9ea
token_response: 038e3625b6a769668a99680e46cf9479f5dc1e86d57164ab3b4a569d
dfc486bf1485d4916a5194fdc0518d3e8444968421ba36e8144aa7902705ff0f3cf40586
3d69451a2a7ba210cc45760c2f1a6045134d877b39e8bcbbf920e5de4a3372557debf211
765cd969976860bc039f9082d6a3e03f8e891246240173d2cf3d69a4613b0f8415979029
22e74c7a1f2e4639e4
token: 00015d4799f8338ddc50a6685f83b8ecd264b2f157015229d12b3384c0f199efe
7b8742cdfb0ed756ea680868ef109a280a393e001d2fa56b1be46ecb31fa25e76731a5b1
d698ea7ab843b8e8a71ed9b2fffa70457a43a8fc687939424b29a7554b40fde130ab7a82
2715909cb73f99a45b640ca1c85180ba9ca1a40bab8b664406a34bcbcb63b5e2e5c455cea
00001a968f7

```

B.2. Issuance Protocol 2 - Blind RSA, 4096

The test vector below lists the following values:

- * skS: The PEM-encoded PKCS#8 RSA private key used for signing tokens, represented as a hexadecimal string.
- * pkS: The DER-encoded SubjectPublicKeyInfo object carrying the public key corresponding to skS, as described in Section 6.4, represented as a hexadecimal string.
- * challenge: A random challenge, represented as a hexadecimal string.
- * nonce: The 32-byte client nonce generated according to Section 6.1, represented as a hexadecimal string.
- * blind: The blind used when computing the blind RSA blinded message, represented as a hexadecimal string.

- * salt: The randomly generated 48-byte salt used when encoding the blinded token request message, represented as a hexadecimal string.
- * token_request: The TokenRequest message constructed according to Section 6.1, represented as a hexadecimal string.
- * token_response: The TokenResponse message constructed according to Section 6.2, represented as a hexadecimal string.
- * token: The output Token from the protocol, represented as a hexadecimal string.

```
skS: 2d2d2d2d2d2d424547494e2050524956415445204b45592d2d2d2d0a4d49494a517
7494241444414e42676b71686b6947397730424151454641415343435330776767b70416
74541416f4943415144584d4d364c5073637a6130696b0a39596c315a4e683868324a796
d504962704c383035354d52516f6463446c4c4672764d694c57666f59535241506b4e744
1546272324162654e334d454a2f67550a392f424958717a704d4256484852546c627a676
b364c7866766b54505a2b4a427832517177364f5555714b442f34426f5464765761536d6
a63644165555037760a4836376e396a383836307463367375546e67616c574d4b6d764d2
f5a6237644262543763345647386577706c6776444f6d616641353956354f78505545704
9510a586c454f4771414f43436c316f54686d33363836436c48413352374c5a4d78567a4
742386f594537583548396a70793532786a2f3242724a6861625033567a430a6f4c696c3
54f6e36487158785363466d4956573742787956495979756d75537659544e52757652777
3566c3775595446366f4d6d2b59722f5a506372594c7a510a4e666135634f747533532b6
649594456716f6f58375541415671382f716c4945747a794a744f7261616756393039327
0324f474c4f6d306a52384543666963520a3868363332572f76554d7739724c4c317a4d4
e7554424f4b74316c546c307265644a677668626b66515a5a55303541336a69366c71754
d2f47307250553030470a3369385253732f64694e625843505453746b616f45537350345
946496d526269756c786a544e2b626c5859744868494261556774306e2b566b544a77554
86e380a5367447534664e547142776567517265494e427732446b6e63576e676b5644416
867577635383669727351644c345843723376765856647a51375134586978730a465a6b6
d7763676d665142444f3469363078536c336337316954595362783359326e74584b4e6f5
a4c31537a424f595051496a6c734749454250677157546e5a0a64655a7852464f6c4d384
679794b6d30746471586271594b57716b663777494441514142416f4943414364314d707
044773645424467763558654168777252710a32726c717042492f6a6a50304e6e7057755
a302b6e787a53624a4361784d2f4f61436844676e654e586e573259655144524e7242505
84d5331344e646f4e554e0a56516c36497166445567637169636741696e7542622f4a687
a6c4d74466d5350466d2b667650726a4d2b6c4831544f384863354253633274414a52574
36468770a794a7663446349656d74646378437877754b6746485251704a5071356368526
56431535077766f50494c7831686959356c2f5175423478717a764149483873530a34673
949705a2f766169443558573541335847734a4f2b696a78717250706756655236474e4f5
33763515551716a44446967665a626a5864354a322b734d79460a4c435a6e514d6771577
03731756437366a4f4a445571563537454154534e6b56472f52633279537a554b4d6e354
c335a314a796d6d31694f584a5136535744760a6f5375426b435652436645635a746b546
f436c76646f5456666b63414a394e51343839686b3050754e466f593675315671614c5a4
a7764657a3931764f5966300a4b676f665a4a6e774976547733754b786f7559685033334
43644574a4f31763265487658767a4c744f4a54477054446f726f4c7a2f475459316b682
```

f484d50510a416c4f333758772f526b527752684b4d6f39545555347766642b2b5348325
4346e6730584875355254764d5339496b4266724455337a75557675434b7270356b0a377
942364473763433517932364f6367367151584a563650676b6c5770696164445778326f4
73935746578436f3346563434764a6d49327a706d5748514257630a48395331336b6d7a3
06873506274576a537535494c443061694f5630764d4c61353841685a767a32774562763
750546a494d48364d77446736326d6450744d6b0a534969517548644a38326a69316e757
8646f5678416f4942415144773431456757476e437644496239727a73704b6e61486a654
3574650385737664a6b446c4c0a61786c724a6f574e614d654c33306b646873576569765
7616730436f457656556631543455374936735a33705054675653685079374b6f4c645a4
86f757a5a390a4d774679716d694652726e574135467138544b756f4354647a3836542f5
37859753330625a476c6d6a4c432f6664453279556141486f6f5177375a2f34646b6d0a4
1693031696a36484d654e39486e5a425737656c436c4173443242436d4b5279655554793
7754b472b6b3732704630336e45694f7a42615565354d6c7a7436620a3867586b63715a5
735495478454b41726b6d486a77454379765178346a7867324573326a7653654e3041554
b624a63534a736e7a35434a65667a697561676f650a36482b3947676f62386e496a78546
34e784c7a58316a676e5236772f41302f6e657837714e6b57357773485a46644c58416f4
9424151446b734d6e75503452460a7558474a6e34316559696c43716b694e4e6b4a53753
9637456556b7772364570744230506d34556e3770535238684866767a5a43782f6b65766
b33553336546e0a67556c72416a4b72474a6a77437231457a7252726e417153466f61395
04851395a69697264626336726f474932576f63795a585859664b784a56646a4f754c440
a793947564d72575133665a3571696d7a6a394b4163304852455765784b6875514259465
554542f44323169584b4b65497568647834796b795435323857714b310a66786f4c72584
e6d5936726b77787778763334556779685a54564a75454879506e51422f54743242377a2
b342f763033665a46347361635646337738514a306c0a692b466879555a34326b63674a3
52b654c6d5558726b6d5a5959314a534332417a722f336449516a4b554655515935326a3
65775667338392f576a72516669470a31732b6c51385a7a6d4a4370416f49424151432b3
84c43686e764e574e4b37546b365431507943546b466758726351457950374a657453766
6316c4b6f654a430a304d6337692b58387a5a4e6674473478353941636163716c43376c6
96a5a55394351564f6d415159652f754d4679524371524c62453271426d79694f70357a7
00a3538487562693261513034564e554f44767644555256346468364148556e52706f534
f49356b59723079646137746f7070376a4662565165324b4c56535a742b0a746f4448384
a6c7a2f5374345774427033465a4538354747573748586a70746f756f6855344c777a464
7492f4c6d37486935783733357038716a385a6366642f0a384f756632626e6354382f674
938676b35633034307451794b483177534d4e4e6d5a496c54535943635653725369346b4
5567677684955354d726e75507647380a6256556b48584d694b7377316d63777739704d4
6373634716f6d464337586f66594d30664d6a6c4a416f494241465079565632676354534
b2b784e7976786b4c0a58577638522f2b57454569416256395674445572387a503079736
f6b34333869412b57432f32367271515a676b36446d61486d673073367356622f7a495a6
84f0a7769307a4e41446a41375751705179314f6961533333522b594b56333435656c345
35454386a433543736a79536e306559504b71396679376636614e343770570a304267664
4346d37584b454d4c66664a744d2b437a6e56536f41504c433449677257646e5a4141376
c306d57416c52576832645275664a3377703751763943770a2b315659446178785235334
e2b327930686e4b696d4b613745696970555952567834566f444a6c6d2f5a525a5769545
335796253375279514f563645334e71560a2f592f6647366563446a33674732497a50674
c4e664f3651646b556d7679364e4155386c645638754962707045446749497042684f684
a394865486a664343490a75326b4367674542414d68686e6f69312b78454a365339636e5
a327977527737433057544962662b54557230436e374f344a4d76637843544b484d62773
76a680a574c6247392f314732595966354c5673326b572f34472b2f446d586b6d4b68715

4746d6f324e50463666504b73694d64477877354149677039557a50543168550a2f6c777
26a75474d322f2b77434b70316d3645374e4d4d5659684b5841534f673473652b3676586
455356a56436f634343576162657a6c5835513262796c51480a2b2b624f6d4661504d535
5683761616d6657357573614553627366612f45506932446d3545574f456339567577525
671526143527552534632507043627279410a6b376e6b6b6746474b5a523777353053465
5366f7a776667627a2b7a33637256315535766d3076346c63794b6d524b6c4b575a51554
9624b782f5070583737640a395057536e3569594343704c432f316245372f566f4c70467
7757631656a773d0a2d2d2d2d2d454e442050524956415445204b45592d2d2d2d2d0a
pkS: 30820252303d06092a864886f70d01010a3030a00d300b060960864801650304020
2a11a301806092a864886f70d010108300b0609608648016503040202a20302013003820
20f003082020a0282020100d730ce8b3ec7336b48a4f5897564d87c87627298f21ba4bf3
4e7931142875c0e52c5aef3222d67e86124403e436d0136ebd806de37730427f814f7f04
85eace93015471d14e56f3824e8bc5f5be44cf67e241c7642ac3a39452a283ff80684ddb
66929a371d01e50fee7f1faee7f63f3ceb4b5ceacb939e06a558c2a6bccfd96fb7416d3ed
ce151bc7b0a6582f0ce99a7c0e7d5793b13d41292105e510e1aa00e082975a13866dfaf3
a0a51c0dd1ecb64cc55cc607ca1813b5f91fd8e9cb9db18ffd81ac985a6cfdd5cc2a0b8a
5e4e9falea5f149c1662155bb071c95218cae9ae4af613351baf470b1597bb984c5ea832
6f98aff64f72b60bcd035f6b970eb6edd2f9f2180d5aa8a17ed400056af3faa5204b73c8
9b4eada6a057dd3dda9d8e18b3a6d2347c1027e2711f21eb7d96fef50cc3dacb2f5ccc36
e4c138ab75953974ade74982f85b91f419654d390378e2ea5aae33f1b4acf534d06de2f1
14acfd88d6d708f4d2b646a8112b0fe181489916e2ba5c634cdf9b95762d1e120169482
dd27f959132705079fc4a00eeef353a81c1e810ade20d070d839277169e09150c08605a
fe7cea2aec41d2f85c2af7bef5d577343b4385e2c6c159926c1c8267d00433b88bad314a
5ddcef58936126f1dd8da7b5728da192f54b304e60f4088e5b0620404f82a5939d975e67
14453a533c172c8a9b4b5da976ea60a5aa91fef0203010001
challenge:
83ce743dcdadd5fc4aeb0357977bb8426635c390a15b88947f0b1c62e4a87c22
nonce: 7e0da97bfcdc4365a5f40e69262f78b81bcd2f92daf885358d9831874e3dd9d22
blind: cd6d03e332386d0166eb76b8e78522510e5cbdcf49aaac83191ea948a7719e914
0ccb6701f7301b7d445ede7adbc5e582b35edd9ac45bc4b8f794e150b2e3e407b7b7624b
6f90b33845bc255174cee0c570aa781c203dce8563afe9f48e2b49c773bba1031987fb48
d981d131876f53e264ec0609a3ea628cf2042005ed3071aeb6657472c7e7df947915b8cd
333e3f5078e456e65e5edef8f892c4f21d25a18dcd80628ed6c7d55b0b9433bc67760be0
8a4eacbdb16a4be4c5b8cab26b478fa6a36ea3c3dd1fffb420bf69feef52aab4892c9e60a
df18347b4e8256b5a0e8cbf55fc97ac62af2e7349ba98ca7462cb6a41d70b0217814a06e
1b257289c3b345be652b87d5820b06a80500880b40b8772140bf431f11497114b20fee7e
5ffc1af5cf874cc293a0c8df65d52814bcd55ae6d3701f73d140ca82c6528627129ea389
f3cbd6058f4f80b7df3818f36dd3489259b6b95df4511930ff02b5cbe643fea44306e7c4
e3d9b02f1b0559aa238b8882a6e8791bfbfd366ef4fe433fd42e5c5db208c9fceb74def
11663ce5f793c7013116995b3fef392a8633b08179a9c8309fb69359fd8486a7a8febb42
4d0726c2516b11e8b19a55fa54e9be606de6811059976473de8f9adb25af7e2862932bed
c7764b4dc50bfc9d724a4aba356a7677b5aceef21876e56b4f1b65adef0fbf8bf1636815
be01b372727e79aa6c47f41
salt: d13a47fa6466a37203e51ac34f7319831b3f04202ff74c98ab18e78088b7ac3014
06189783503227153871405c6alda0
token_request: 0002013a370077e8259098e741dcaac8184838b7c995cd82966419064
90205bf28e6745396dd9ec1761c4676e9fec3272588194c48bc60fc77c3fff19219a6b65
96523044d07d9d9e4dd88f2db9d9c369597363a12a65d244d69a1b743b365f8f5bdf8d6

```
52f2d2e249f417e9fd7da7db2626d6499d7d3856d8f9277385dfd776ffa4ffa74d7a7b17
01d87f40e525bb258a7f5b4d6c134b3b242ef46d93b32f106c84c396b1fc2772796b2473
64c48c364537708f3a8a87fb870a299ac08107a5dd3f467733d76c2359e346ee3ebcff68
c3c7e10f0f01a2b9bbdb26ffea14f81f036a71c47725b5c9f81e0320b85abfd77d5eb1c
ccfdd8eb0cf2735ea297e2a07c82df9e9b0a4d21baa81a2cfb2a72983107555035386c33
973d48f04257dd8298116d2c93298810cb6b82ad033f5b16f9f7a65d8f74b7bdcae000db
1411b40f46cb373cb69c8ab58552f98904b78229b63838ab40e833fab4ec47acf00a00db
3290c9c74ed982c64ddbaecae16fd73976ffe7da7b3b1a8e0190e95b7ce7900295f3c8c94f
2d3d85cd1825fe27073aad63fa4c530907402dc4e3a748edb300f05ce7ac5b8c1d9aeb21
66002b05dee582aeafa4f503f13bae1a51be9420e3cdcb685169bdc5ae2ebca7713ec16
666b55b097e56e5719d1b0324ecaca613af76b9f367775a90dba5e7fdd21a8da73bd80b1
31e6531117ef709ad8c7b2b3182255235acea
token_response: 061780e09bc9b851fe81e7022ee2d55b043198bcb1aa33f761d213a9
8d831abaef5417d30904a7c9d7ec531278cd9655c4b7d72f3a4e66e26565b73e5f1b9271
b541f28556543d2473c363e104a11c6ba1a1d8f99b32d3cd8f74ae07b465afdaabd21977
d6f114ea9484cd592f461e5dc4a97b86fe1463b1c69171cf734f49c240760dc2555d7cc8
9d7f882d3e1b99388bc3b561d418a7fc5770ccc66e3dd41f4a74e8267492f48e8b6aabce
592c8dc83826b1f4528f2497902a0ce4baacd4216623274057592e77091102452de115bb
d0c98ee22b14fe30df3b1277ab17bc948b62b8adb56b67e44dce74860647718c8f4bd10e
7b022ad35e2996dd497f4a6c765689931fc11b56b06cca921fb1205e8b2302540a48da
38ed7a5e20f7aa80b55f1de9f5a6c9589f4fc23b6ed2c53ad2a6a64abb8fba67aef48e66f
f77647aa05ce96527e2d1199b7553a3b900ec510f5b765211c14d0f6cbe82f28be4d8ac
8110f13d7aee3741fc68f4abf3ca33444f790a72ef4ddec72c1d92938ee5d303f914ff12
8d9a73e867d9aa5d327934d4b2b299d7ffa32350d72d35103c027ac76c5417823e6790ee
174b9761e086bfa445f1725b20cff6a94d51abdfcfe46ff7c0f69685fd38639bdbd3917f
424949027b5a322fc217364748e544257d11f0cdcab9aal3d3282f6ac2811dfc8db5b2f99
c1ba00c366057748cabe6975fc73ed60
token: 00027e0da97bfdbc4365a5f40e69262f78b81bcd2f92daf885358d9831874e3dd9
d22895c9d3fff72e71d22cdbc11706d350ab772b0820be9f33a02e003652cb00a31501000
00000000000000000000000000000000000000000000000000000000000000000004c9e19f2a624e
e7d44ac35846749b29b1d3a784f13ea1005a2c87b9d3f939f795877d5fded823f45fd399
dc0b8730cb46c66102740c679338b7093c47e8f586e48a8b042cb0ea2b901f6981e797c4
a614f52f02ffe3ae7f83fa4f9a243e8b59621975abff94f82b3fadfb4cb305cc1c1b677
42a673204e5a0aa0c25423c604430597d0332e30dddab8855de42bcf49668410df38bb3b
fa4370df28ec59316bc1c6f3c9afc8ccfdb93f4ca60365683988f649201e1c6d6bb73d14
d0dc9a0f596a9f76502ebadc0b248985f9bb66d9d99a5aca08527aa11d555b26489299ba
5b400157a9fd47b6b4fa74315eade2b22624b29d53eb84126f64b98ea5ba45914d1fa14b
1525e2327856565054a1db9b0d778871fa6ed4d0d4c26641bf3f4faa33efaa0f5b8cec80
8d52ed3f1378273d5b7b0b0b812bfc128ef5e4924a60aebd124659d31661e9ec89f8bcb9
a51bab6a5711187639c24fdd31f14abf7d80827df91f31bfe7c4916ec4d1927ca138c5ba
9a595a9e83b5055148d19ad005c523eda76ea94006ce6315e20ed0d637fb1211b541e9ea
12c9b641d48fd2cc5f0c7f479672a4e2bf7469267c8526d734df41f2c30fb62c2aea4033
214df44a53353dc683cf72dee7b1ba39ef668478958935a0e8c9a880ae85712c401d7f09
b66fbdad05cfd69d615b229bce8818c6a6472e07a8793456f19f4f4015c507ab5c1357881
68b
```

Authors' Addresses

Sofía Celi
Cloudflare
Lisbon
Portugal
Email: sceli@cloudflare.com

Alex Davidson
Brave Software
Lisbon
Portugal
Email: alex.davidson92@gmail.com

Armando Faz-Hernandez
Cloudflare
101 Townsend St
San Francisco,
United States of America
Email: armfazh@cloudflare.com

Steven Valdez
Google LLC
Email: svaldez@chromium.org

Christopher A. Wood
Cloudflare
101 Townsend St
San Francisco,
United States of America
Email: caw@heapingbits.net

cfrg
Internet-Draft
Intended status: Informational
Expires: 26 August 2021

S. Iyengar
A. Raghunathan
Facebook
22 February 2021

Verifiable Oblivious Pseudo-Random Functions with Public Metadata
draft-iyengar-cfrg-voprfmetadata-00

Abstract

This document describes a verifiable mechanism to bind public metadata to an existing Verifiable oblivious Pseudo-Random function [I-D.irtf-cfrg-voprf] (VOPRF). Using zero knowledge proofs a receiver can verify that, for an input x , a $VOPRF(k, x, \text{metadata})$, is generated from a secret key k , as well as the given metadata.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Crypto Forum Research Group mailing list (cfrg@ietf.org), which is archived at https://mailarchive.ietf.org/arch/search/?email_list=cfrg.

Source for this draft and an issue tracker can be found at <https://github.com/siyengar/verifiable-attribute-based-key-derivation>.

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 26 August 2021.

Copyright Notice

Copyright (c) 2021 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements	3
1.2. Terminology	3
2. Preliminaries	4
2.1. Prime-Order Group Dependency	4
2.2. Other Conventions	4
2.3. Discrete log proofs	4
3. Protocol	5
3.1. Overview	5
3.2. Pre-Setup	6
3.3. Evaluate VOPRF	7
4. Application considerations	9
4.1. Metadata bits	9
4.2. Encoding metadata	9
5. Comparison with other approaches	9
5.1. Pairings	9
5.2. Partially oblivious PRF	9
6. Security Considerations	10
6.1. Cryptographic security	10
6.1.1. n-Diffie Hellman exponent assumption	10
6.1.2. Selective security vs full security	10
7. IANA Considerations	11
8. References	11
8.1. Normative References	11
8.2. Informative References	11
Acknowledgments	12
Authors' Addresses	12

1. Introduction

A VOPRF allows a client and server to evaluate a psuedo-random function " $F(k, x)$ ", with secret key " k ", and input " x " without the client learning the key " k " and the server learning the input " x ". Additionally in a VOPRF, the client can verify that the output was computed using the key " k ".

One challenge in VOPRFs is to be able to bind public metadata to the output of the VOPRF while keeping the VOPRF both verifiable and oblivious. Unlike the input x to the VOPRF, public metadata is not meant to be secret to either the client or the server. This public metadata is useful in applications where being able to bind application context to a VOPRF output is critical to the security of the application.

In this draft we describe a mechanism to bind public metadata to a VOPRF by deriving the public-private keypair that is used by the VOPRF from the metadata [PrivateStats]. This method allows the use of existing elliptic curve VOPRF ciphers while only changing the way the secret key is derived. Additionally, the key derivation mechanism of the public key can be verified by a client using non-interactive zero-knowledge proofs to prove that the metadata specific key is derived from a master secret.

The draft does not describe how metadata is used, but that left to specific application protocols that use this public metadata mechanism.

1.1. Requirements

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.

1.2. Terminology

The following terms are used throughout this document.

- * PRF: Pseudorandom Function.
- * VOPRF: Verifiable Oblivious Pseudorandom Function.
- * Client: Protocol initiator. Learns pseudorandom function evaluation as the output of the protocol.

- * **Server:** Computes the pseudorandom function over a secret key. Learns nothing about the client's input.
- * **NIZK:** Non-interactive zero knowledge.
- * **DLEQ:** Discrete Logarithm Equality.

2. Preliminaries

The document defines extensions to the VOPRF required to support metadata. This document depends on the following:

- * **"GG":** A prime-order group implementing the API described in [I-D.irtf-cfrg-voprf] as well as the additional APIs defined below in Section 2.1.
- * **"Public Metadata":** The public metadata is defined as an "n" bit vector. To represent "b" values, an application could use "log b" bits.

2.1. Prime-Order Group Dependency

We define new member functions on the prime-order group "GG" defined in [I-D.irtf-cfrg-voprf]:

- * **ScalarMult(point, scalar):** A member function of "GG" that multiplies an element in the group with a "scalar" from "GF(p)".
- * **NewGenerator():** A member function of "GG" that samples a new generator for the group.

2.2. Other Conventions

All algorithm descriptions are written in a Python-like pseudocode. All scalar multiplications are performed modulo "GF(p)".

2.3. Discrete log proofs

Zero knowledge proofs for statements on discrete-logs were summarized by [Camenisch97]. We describe two algorithms used in this draft on "GG" to prove discrete log statements.

"DLEQProve(k, A, B, C, D)" proves that " $B = k * A$ " and " $D = k * C$ " without revealing the value of "k". This type of proof is used when "k" is a secret value that should not be revealed to a verifier.

```
def DLEQProve(k, A, B, C, D):
    r = GG.RandomScalar()
    E = r * A
    F = r * C
    hashInput = A || B || C || D || E || F
    cbytes = Hash(hashInput)
    c = GG.HashToGroup(cbytes)
    z = r + k * c
    return (z, E, F)
```

"DLEQVerify(A, B, C, D, proof)" verifies that the proof generated by "DLEQProve" is valid.

```
def DLEQVerify(A, B, C, D, proof):
    hashInput = A || B || C || D || proof.E || proof.F
    cbytes = Hash(hashInput)
    c = GG.HashToGroup(cbytes)
    cBE = cB + proof.E
    cDF = cD + proof.F
    zA = proof.z * A
    zC = proof.z * C
    return zA == cBE && zC == cDF
```

3. Protocol

3.1. Overview

A server first generates a main key pair "(skM, pkM)", where "skM" is the servers main secret key and "pkM" is the servers main public key. Given public metadata "t", the server generates a keypair specific to the metadata "t", "(skT, pkT) = PKGen(t, skM)", where "skT" is the secret key for metadata "t" and "pkT" is its public key. Once a metadata specific keypair is available, the server can be used to evaluate a "VOPRF(skT, x)", where "x" is the input for the user. When the VOPRF is in verifiable mode, the client also receives a NIZK proof that "skT" and "pkT" are generated from "skM" and "pkM" (in verifiable mode).

```

Client(pkM, input, metadata)      Server(skM, pkM, metadata)
-----
blind, blindedElement = Blind(input)

                                blindedElement
                                ----->
                                skT, pkT, pkProofs = PKGen(skM, metadata)

evaluatedElement, proof = Evaluate(skT, pkT, blindedElement)

                                evaluatedElement, pkT, proof, pkProofs
                                <-----

pkVerified = PKVerify(pkM, pkT, pkProofs)

output = Finalize(input, blind, evaluatedElement, blindedElement, pkT, proof)

```

In the following sections we describe modifications to the VOPRF scheme in [I-D.irtf-cfrg-voprf] to be able to augment an existing VOPRF with public metadata.

3.2. Pre-Setup

We augment the offline context setup phase phase of the VOPRF in [I-D.irtf-cfrg-voprf]. In this phase, both the client and server create a context used for executing the online phase of the protocol.

Prior to this phase, the key pair ("skM", "pkM") should be generated by using "MasterKeyGen(metadataBits)". This keypair is used as the master key for VOPRFs. This master key is not used directly within the VOPRF, however, public metadata is used to generate attribute specific keys that are used in the VOPRF evaluation.

"metadataBits" here is the number of bits of metadata that are required for the application of the VOPRF. "MasterKeyGen" samples "n" scalar elements "a0, a1, ... an" from the group and a new generator "h". "ai" is a group element associated with the "i"th bit of metatadata. Public parameters are calculated by performing scalar multiplication of "h" with each "ai".

```
def MasterKeyGen(metadataBits):
    ais = []
    his = []
    h = GG.NewGroupGenerator()
    a0 = GG.RandomScalar()
    for i in range(metadataBits):
        ai = GG.RandomScalar()
        ais.append(ai)
    for i in range(metadataBits):
        hi = GG.ScalarMult(h, ais[i])
        his.append(hi)
    P0 = GG.ScalarBaseMult(a0)
    skM = (a0, ais)
    pkM = (GG.g, h, metadataBits, P0, his)
    return (skM, pkM)
```

3.3. Evaluate VOPRF

When client and server have agreed on the metadata to use for the protocol, the server first executes "`PKGen(skM, metadata)`" to generate "`skT`" and the proof that "`skT`" is derived from "`skM`". This draft does not discuss how the client and server agree on the metadata to use, and that is left to the application.

Note that "`skM`" has one group element for each bit of the metadata "`t`", as well as the extra group element "`a0`". Given metadata "`t`", "`PKGen`" calculates the attribute specific key by performing a scalar multiplication of all the group elements in "`skM`" for each bit of "`t`" that is set to "1".

To prove that "`skT`" is derived from "`skM`", "`GenProofs`" generates upto "`n`" discrete log proofs, one for each bit of the metadata. Each proof proves that " $hi = ai * h$ " and " $Pi = ai * Pi-1$ ". This proves that "`ai`" was correctly used for bit "`i`".

```

def PKGen(t, skM, pkM):
    pis = []
    pi = skM.a0
    keyBits = len(metadata)
    for i in range(keyBits):
        if t[i] == 0:
            pis.append(None)
            continue
        pi = pi * skM[i]
        pis.append(pi)
    skT = pi
    pkT = GG.ScalarMultBase(skT)
    pkProofs = GenProofs(metadata, pis, skM, pkM)
    return (skT, pkT, pkProofs)

def GenProofs(t, pis, skM, pkM):
    proofs = []
    numProofs = len(pis)
    previousPi = pkM.P0
    for i in range(numProofs):
        if t[i] == 0:
            continue
        Pi = GG.ScalarBaseMult(pis[i])
        proofi = DLEQProve(skM.ais[i], pkM.h, pkM.his[i], previousPi, Pi)
        proofs.append((Pi, proofi))
        previousPi = Pi
    return proofs

```

Once "PKGen" has generated a public key for a set of "metadata" bits, the client can verify that "skT" is derived from "skM", using "PKVerify(pkM, pkT, pkProofs)". This verifies the sequence of discrete-log proofs generated by "PKGen".

```

def PKVerify(pkM, pkT, t, pkProofs):
    previousPi = pkM.P0
    proofVerified = True
    for proof in pkProofs:
        if t[i] == 0:
            continue
        Pi = proof.Pi
        verified = DLEQVerify(pkM.h, pkM.his[i], previousPi, Pi, proof)
        proofVerified = proofVerified & verified
        previousPi = Pi
    return proofVerified

```

A server can use "skT" generated from "PKGen" as the private key for the VOPRF mechanism in [I-D.irtf-cfrg-vopr].

4. Application considerations

4.1. Metadata bits

Applications must choose the maximum size in bits of the metadata that they would like to support before setup of the protocol. The size of the metadata impacts the following - Size of the public key - Computation time for attribute and proof generation

For " b " being the number of metadata values needed for an application, the size of the public key scales as " $O(\log b)$ ". Computation also scales as " $O(\log b)$ " number of scalar multiplications for generating a public key and number of discrete log proof generations and verifications required.

4.2. Encoding metadata

Applications must choose the number of bits of metadata required in order to be able to represent all possible values for the application's metadata. They MUST define their own mechanism encode metadata into bits.

5. Comparison with other approaches

5.1. Pairings

It is possible to construct VOPRFs with public metadata using pairing-friendly curves [I-D.draft-irtf-cfrg-pairing-friendly-curves] with an approach in [Pythia15].

However this approach has some disadvantages. Pairings are not widely available in cryptographic libraries and are also not compatible with existing deployed VOPRFs like in [I-D.irtf-cfrg-voprf]. The approach described here allows applications to use existing deployed VOPRFs while only changing the mechanism of key derivation.

5.2. Partially oblivious PRF

Another approach that could be used to bind metadata to a VOPRF evaluation is to use a similar method in [pOPRF18] which uses a regular " $PRF(k, \text{metadata})$ " to derive a secret key based on the metadata which is then used in the VOPRF.

The verifiability of public key could be achieved by publishing every public key for each metadata value in a central repository, which could be checked by the client. For large number of values of metadata " b ", this approach generates " $O(b)$ " keys, which can be

difficult for clients and servers to manage. In contrast, the approach described in this document, the size of the master public key is $O(\log b)$, and the public keys of each attribute can be verified against the master key later.

6. Security Considerations

6.1. Cryptographic security

The security properties of a VOPRF with public metadata are derived from the proof in [PrivateStats] that the VOPRF defined here is a PRF even after giving an adversary access to proofs from "PKGen". The VOPRF defined in [I-D.irtf-cfrg-voprf] when combined with attributes results in a PRF output of $\text{PRF}(\text{skM}, t, x) = a_0^{t_1} * a_1^{t_2} \dots * a_n^{t_n} * H(x)$.

6.1.1. n-Diffie Hellman exponent assumption

There are several variants of the Diffie-Hellman assumption and the proof of the VOPRF with public metadata is based on the n-Diffie Hellman exponent assumption. The n-DHE problem requires an adversary to distinguish the $n+1$ -th power of a secret "a" hidden in the exponent from a random element in "GG".

Sample uniformly at random "d" in $\{0,1\}$, and a random "r" from "GF(p)": - Given "G" is a generator in "GG" - Given "G", " $a * G$ ", " $(a^2) * G$ ", ..., " $(a^n) * G$ " - if "d" == 0: " $C = a^{(n+1)} * G$ " else: " $C = r * a$ "

Output " $d' == d$ "

6.1.2. Selective security vs full security

The security properties of the VOPRF with public metadata described in this draft is based on the proof in [PrivateStats] that the VOPRF is a selectively-secure VRF. Selective-security is a weaker notion of security that requires an adversary to commit to the challenge input (in this case, the metadata and value x) before trying to break the PRF.

In practice, if target inputs are independent of the system parameters, there should not be an advantage to allowing the attacker to choose the target after seeing system parameters. To convert our VOPRF with public metadata to one satisfying a full security notion in the random oracle model, we require that the metadata be hashed with a collision-resistant hash function with sufficiently large output (≥ 256 -bits). For smaller metadata sets therefore, the selectively-secure VRF is much more efficient.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. Normative References

- [I-D.draft-irtf-cfrg-pairing-friendly-curves]
Sakemi, Y., Kobayashi, T., Saito, T., and R. S. Wahby,
"Pairing-Friendly Curves", Work in Progress, Internet-
Draft, draft-irtf-cfrg-pairing-friendly-curves-09, 16
November 2020, <<https://tools.ietf.org/html/draft-irtf-cfrg-pairing-friendly-curves-09>>.
- [I-D.irtf-cfrg-voprf]
Davidson, A., Faz-Hernandez, A., Sullivan, N., and C. A.
Wood, "Oblivious Pseudorandom Functions (OPRFs) using
Prime-Order Groups", Work in Progress, Internet-Draft,
draft-irtf-cfrg-voprf-06, 21 February 2021,
<<https://tools.ietf.org/html/draft-irtf-cfrg-voprf-06>>.
- [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>>.

8.2. Informative References

- [Camenisch97]
"Proof Systems for General Statements about Discrete
Logarithms",
<<https://crypto.ethz.ch/publications/files/CamSta97b.pdf>>.
- [pOPRF18] "Threshold Partially-Oblivious PRFs with Applications to
Key Management", <<https://eprint.iacr.org/2018/733>>.
- [PrivateStats]
"PrivateStats, De-Identified Authenticated Logging at
Scale", <<https://research.fb.com/privatestats>>.
- [Pythia15] "The Pythia PRF Service",
<<https://eprint.iacr.org/2015/644.pdf>>.

Acknowledgments

TODO acknowledge.

Authors' Addresses

Subodh Iyengar
Facebook

Email: subodh@fb.com

Ananth Raghunathan
Facebook

Email: ananthr@fb.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 5 September 2022

A. Davidson
Brave Software
M. Finkel
The Tor Project
M. Thomson
Mozilla
C. A. Wood
Cloudflare
4 March 2022

Key Consistency and Discovery
draft-wood-key-consistency-02

Abstract

This document describes the key consistency and correctness requirements of protocols such as Privacy Pass, Oblivious DoH, and Oblivious HTTP for user privacy. It discusses several mechanisms and proposals for enabling user privacy in varying threat models. In concludes with discussion of open problems in this area.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the mailing list (), which is archived at .

Source for this draft and an issue tracker can be found at <https://github.com/chris-wood/key-consistency>.

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 5 September 2022.

Copyright Notice

Copyright (c) 2022 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
1.1. Requirements	3
2. Terminology	3
3. Core Requirements	3
4. Consistency and Correctness at Key Acquisition	4
4.1. Direct Discovery	4
4.2. Single Proxy Discovery	5
4.3. Multi-Proxy Discovery	6
4.4. Database Discovery	7
5. Minimum Validity Periods	9
6. Separate Consistency Verification	9
6.1. Independent Verification	9
6.2. Key-Based Encryption	10
7. Future Work	10
8. Security Considerations	10
9. References	10
9.1. Normative References	10
9.2. Informative References	11
Authors' Addresses	12

1. Introduction

Several proposed privacy-enhancing protocols such as Privacy Pass [PRIVACY-PASS], Oblivious DoH [ODOH], and Oblivious HTTP [OHTTP] require clients to obtain and use a public key for execution. For example, Privacy Pass public keys are used by clients for validating privately issued tokens for anonymous session resumption. Oblivious DoH and HTTP both use public keys to encrypt messages to a particular server.

User privacy in these systems depends on users receiving a key that many, if not all, other users receive. If a user were to receive a public key that was specific to them, or restricted to a small set of users, then use of that public key could be used to learn targeted information about the user. Users also need to receive the correct public key.

In this document, we elaborate on these core requirements, and survey various system designs that might be used to satisfy them. The purpose of this document is to highlight challenges in building and deploying solutions to this problem.

1.1. Requirements

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.

2. Terminology

This document defines the following terms:

Key Consistency and Correctness System (KCCS): A mechanism for providing clients with a consistent view of cryptographic key material within a period of time.

Reliant System: A system that embeds one or more key consistency and correctness systems.

The KCCS's consistency model is dependent on the implementation and reliant system's threat model.

3. Core Requirements

Privacy-focused protocols which rely on widely shared public keys typically require keys be consistent and correct. Informally, key consistency is the requirement that all users who communicate with an entity share the same view of the key associated with that entity; key correctness is that the key's secret information is controlled by the intended entity and is not known to be available to an external attacker.

Some protocols depend on large sets of users with consistent keys for privacy reasons. Specifically, all users with a consistent key represent an anonymity set wherein each user of the key in that set is indistinguishable from the rest. An attacker that can actively cause inconsistent views of keys can therefore compromise user privacy.

An attacker that can cause a user to use an incorrect key will likely compromise the entire protocol, not just privacy.

Reliant systems must also consider agility when trying to satisfy these requirements. A naive solution to ensuring consistent and correct keys is to only use a single, fixed key pair for the entirety of the system. Users can then embed this key into software or elsewhere as needed, without any additional mechanics or controls to ensure that other users have a different key. However, this solution clearly is not viable in practice. If the corresponding key is compromised, the system fails. Rotation must therefore be supported, and in doing so, users need some mechanism to ensure that newly rotated keys are consistent and correct.

Operationally, servers rotating keys may likely need to accommodate distributed system state-synchronization issues without sacrificing availability. Some systems and protocols may choose to prioritize strong consistency over availability, but this document assumes that availability is preferred to total consistency.

4. Consistency and Correctness at Key Acquisition

There are a variety of ways in which reliant systems may build key consistency and correct systems (KCCS), ranging in operational complexity to ease-of-implementation. In this section, we survey a number of possible solutions. The viability of each varies depending on the applicable threat model, external dependencies, and overall reliant system's requirements.

We do not include the fixed public key model from Section 3, as this is likely not a viable solution for systems and protocols in practice. In all scenarios, the server corresponding to the desired key is considered malicious.

4.1. Direct Discovery

In this model, users would directly query servers for their corresponding public key, as shown below.




Figure 1: Direct Discovery Example

The properties of this solution depend on external mechanisms in place to ensure consistency or correctness. Absent any such mechanisms, servers can produce unique keys for users without detection. External mechanisms to ensure consistency here might include, though are not limited to:

- * Presenting a signed assertion from a trusted entity that the key is correct.
- * Presenting proof that the key is present in some tamper-proof log, similar to Certificate Transparency ([RFC6962]) logs.
- * User communication or gossip ensuring that all users have a shared view of the key.

The precise external mechanism used here depends largely on the threat model. If there is a trusted external log for keys, this may be a viable solution.

4.2. Single Proxy Discovery

In this model, there exists a proxy that fetches keys from servers on behalf of multiple users, as shown below.




Figure 2: Single Proxy Discovery Example

If this proxy is trusted, then all users which request a key from this server are assured they have a consistent view of the server key. However, if this proxy is not trusted, operational risks may arise:

- * The proxy can collude with the server to give per-user keys to clients.
- * The proxy can give all users a key owned by the proxy, and either collude with the server to use this key or retroactively use this key to compromise user privacy when users later make use of the key.

Mitigating these risks may require tamper-proof logs as in Section 4.1, or via user gossip protocols.

4.3. Multi-Proxy Discovery

In this model, users leverage multiple, non-colluding proxies to fetch keys from servers, as shown below.




Figure 4: Database Discovery Example

The database is expected to have a table that asserts mappings between server names and keys. Examples of such databases are as follows:

- * An append-only, audited table similar to that of Certificate Transparency [RFC6962]. The log is operated and audited in such a way that the contents of the log are consistent for all users. Any reliant system which depends on this type of KCCS requires the log be audited or users have some other mechanism for checking their view of the log state (gossiping). However, this type of system does not ensure proactive security against malicious servers unless log participants actively check log proofs. This requirement may impede deployment in practice. Experience with Certificate Transparency shows that most implementations have chosen not to check SignedCertificateTimestamps before using (that is, accepting as valid) a corresponding TLS certificate.
- * A consensus-based table whose assertions are created by a coalition of entities that periodically agree on the correct binding of server names and key material. In this model the agreement is achieved via a consensus protocol, but the specific consensus protocol is dependent on the implementation.

For privacy, users should either download the entire database and query it locally, or remotely query the database using privacy-preserving queries (e.g., a private information retrieval (PIR) protocol). In the case where the database is downloaded locally, it

should be considered stale and re-fetched periodically. The frequency of such updates can likely be infrequent in practice, as frequent key updates or rotations may affect privacy; see Section 5 for details. Downloading the entire database works best if there are a small number of entries, as it does not otherwise impose bandwidth costs on each client that may be impractical.

5. Minimum Validity Periods

In addition to ensuring that there is one key at any time, or a limited number keys, any system needs to ensure that a server cannot rotate its keys too often in order to divide clients into smaller groups based on when keys are acquired. Such considerations are already highlighted within the Privacy Pass ecosystem, more discussion can be found at [PRIVACY-PASS-ARCH]. Setting a minimum validity period limits the ability of a server to rotate keys, but also limits the rate of key rotation.

6. Separate Consistency Verification

The other schemes described here all attempt to directly limit the number of keys that a client might accept. However, by changing how keys are used, clients can impose costs on servers that might discourage key diversity.

Protocols that have distinctly separate processes for acquiring and using keys might benefit from moving consistency checks to the usage part of the protocol. Correctness might be guaranteed through a relatively simple process, such obtaining keys directly from a server. A separate correctness check is then applied before keys are used.

6.1. Independent Verification

Anonymous queries to verify key consistency can be used prior to use of keys. A request for the current key (or limited set of keys) will reveal if the key that was acquired is different than the original. If the key that was originally obtained is not included, the client can abort any use of the key.

It is important that any validation process not carry any information that might tie it to the original key discovery process or that the system providing verification be trusted. A proxy (see Section 4.2) might be sufficient for providing anonymity, though more robust anonymity protections (see Section 4.3) could provide stronger guarantees. Querying a database (see Section 4.4) might provide independent verification if that database can be trusted not to provide answers that change based on client identity.

6.2. Key-Based Encryption

Key-based encryption has a client encrypt the information that it sends to a server, such as a token or signed object generated with the server keys. This encryption uses a key derived from the key configuration, specifically not including any form of key identifier along with the encrypted information. If key derivation for the encryption uses a pre-image resistant function (like HKDF), the server can only decrypt the information if it knows the key configuration. As there is no information the server can use to identify which key was used, it is forced to perform trial decryption if it wants to use multiple keys.

These costs are only linear in terms of the number of active keys. This doesn't prevent the use of multiple keys, it only makes their use incrementally more expensive. Trial decryption costs can be increased by choosing a time- or memory-hard function such as [ARGON2] to generate keys.

Encrypting this way could provide better latency properties than a separate check.

7. Future Work

The model in Section 4.3 seems to be the most lightweight and easy-to-deploy mechanism for ensuring key consistency and correctness. However, it remains unclear if there exists such an anonymity network that can scale to the widespread adoption of and requirements of protocols like Privacy Pass, Oblivious DoH, or Oblivious HTTP. Existing infrastructure based on technologies like Certificate Transparency or Key Transparency may work, but there is currently no general purpose system for transparency of opaque keys (or other application data).

8. Security Considerations

This document discusses several models that systems might use to implement public key discovery while ensuring key consistency and correctness. It does not make any recommendations for such models as the best model depends on differing operational requirements and threat models.

9. References

9.1. Normative References

- [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>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", RFC 6962, DOI 10.17487/RFC6962, June 2013, <<https://www.rfc-editor.org/rfc/rfc6962>>.
- [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

- [ARGON2] Biryukov, A., Dinu, D., Khovratovich, D., and S. Josefsson, "Argon2 Memory-Hard Function for Password Hashing and Proof-of-Work Applications", Work in Progress, Internet-Draft, draft-irtf-cfrg-argon2-13, 11 March 2021, <<https://datatracker.ietf.org/doc/html/draft-irtf-cfrg-argon2-13>>.
- [ODOH] Kinnear, E., McManus, P., Pauly, T., Verma, T., and C. A. Wood, "Oblivious DNS Over HTTPS", Work in Progress, Internet-Draft, draft-pauly-dprive-oblivious-doh-11, 17 February 2022, <<https://datatracker.ietf.org/doc/html/draft-pauly-dprive-oblivious-doh-11>>.
- [OHTTP] Thomson, M. and C. A. Wood, "Oblivious HTTP", Work in Progress, Internet-Draft, draft-ietf-ohai-ohttp-01, 15 February 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-ohai-ohttp-01>>.
- [PRIVACY-PASS] Celi, S., Davidson, A., Faz-Hernandez, A., Valdez, S., and C. A. Wood, "Privacy Pass Issuance Protocol", Work in Progress, Internet-Draft, draft-ietf-privacypass-protocol-02, 31 January 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-protocol-02>>.
- [PRIVACY-PASS-ARCH] Davidson, A., Iyengar, J., and C. A. Wood, "Privacy Pass Architectural Framework", Work in Progress, Internet-Draft, draft-ietf-privacypass-architecture-02, 31 January 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-architecture-02>>.

Authors' Addresses

Alex Davidson
Brave Software
Email: alex.davidson92@gmail.com

Matthew Finkel
The Tor Project
Email: sysrq@torproject.org

Martin Thomson
Mozilla
Email: mt@lowentropy.net

Christopher A. Wood
Cloudflare
101 Townsend St
San Francisco,
United States of America
Email: caw@heapingbits.net