Privacy Pass
Internet Draft
Intended status: Informational
Expires: November 4, 2021

M. McFadden
internet policy advisors, llc

May 4, 2021

Privacy Pass: Centralization Problem Statement
draft-mcfadden-pp-centralization-problem-01.txt

Status of this Memo

Copyright Notice

Abstract

This document discusses the problems associated with strict upper bounds on the number of Privacy Pass servers in the proposed Privacy Pass ecosystem. It documents a proposed problem statement.

Table of Contents

1. Introduction

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.

The protocol itself in defined in [ID.davidson-pp-protocol-01] and the architectural framework is in [ID.davidson-pp-architecture-01].

The architecture document leaves for a later time the issue of server centralization. This document is a discussion of the problems related to server centralization in Privacy Pass, the impact of centralization on the protocol's privacy goals, and some potential mitigations for the problem.

An important feature of the Privacy Pass Architecture is the concept of the anonymity set of each individual client. The Privacy Pass

ecosystem has a set of servers which issue tokens to clients which can then be redeemed at the application layer for authentication.

Trust is an important component in Privacy Pass. The servers have to publish their public keys and details of the ciphersuite they are using. It is necessary to publish these in a globally consistent, tamper-proof data structure. Clients that use the same registry of server information need to coordinate in some way to validate that they have the same view of the registry and its data.

Four server running modes are discussed in [ID.davidson-pp-architecture-01]. Common to all four is a discussion of the need to set an upper limit on the number of servers that are allowed. The motivation for limiting the number of servers is that the is a correlation between larger numbers of servers and dilution of privacy.

2. Potential Privacy Concerns

When a client redeems a token in Privacy Pass, there is very little information in the token itself other than the key that was used to sign the token. A key feature of the protocol is that any client can only remain private relative to the entire space of users using the protocol.

In three of the four server running modes, a Privacy Pass verifier is able to trigger redemption for any of the available servers. The greater the number of servers, the greater the loss in anonymity.

The architecture document, [ID.davidson-pp-architecture-01], provides an example where, if there are 32 servers, then the verifier learns 32 bits of information about the client. In certain circumstances, having that much information about the client can lead to the client being uniquely identified and the goals of Privacy Pass thwarted. As a result, the architecture document supplies the following mitigation:

"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. [ID.davidson-pp-architecture-01]."

Putting restrictions on the number of redemption tokens at the client is considered. However, establishing control of the client, and the number of tokens it has, is far more difficult than restricting the number of active servers.

3. Centralization in Privacy Pass - Problem Statement

For Privacy Pass to succeed clients must be able to acquire tokens that they can later redeem with greater privacy and anonymity. This document does not discuss the goals of privacy or anonymity. Instead, it identifies a problem related to the upper bound in number of servers that affects the Privacy Pass ecosystem.

For the purposes of this draft, "server centralization" is the strict limit or upper bound in the number of servers available from which a client can acquire a token for later redemption.

The architecture draft specifies an upper limit of four for this upper bound.

The problem statement for Privacy pass can be summarized: an upper bound to available Privacy Pass servers creates architectural, engineering and practical problems for the deployment of the protocol. Any successful deployment of Privacy Pass must find mitigations for these problems.

3.1. Architectural Problems

Centralization is a problem space that has been exhaustively explored by others; not least of which in the IETF itself. The now expired IAB draft, [I-D.arkko-arch-infrastructure-centralisation-00], discussed six separate issues related to centralization and several of them appear to apply to Privacy Pass.

Having a very limited number of servers available creates an architectural strain on avoiding single points of failure. While the Privacy Pass architecture document does specify up to four servers, this is a very small number for, potentially, billions of possible users. And this assumes that the protocol is only used in "human-to-server" applications and not in situations where the client is not a human but some other device - either acting on behalf oa human or autonomously. Strict limitations on the number of servers poses the question of how the Privacy Pass architecture can scale in the presence of a large user base.

The Privacy Pass architecture, by limiting the number of servers, also concentrates information and potentially limits the ability for other competing providers of the token generating services. By concentrating the information in a small number of servers, a problem appears when there are machine learning opportunities to collect and process data about clients requesting tokens.

A side effect of limiting the number of servers is that a significant amount of information ends up being in the control of a small number of entities. A client may trust a Privacy Pass server as send it information about itself in order to request tokens. However, the protocol itself can make no guarantee about the data handling practices of the server operator. Situations outside the control of the protocol may make it so there are pressures to misuse the data concentrated at the small number of servers.

3.2. Engineering Problems

In the event that a very limited number of servers can be provided while still supporting the goals of the protocol, there is clearly a global scaling problem that needs to be solved. Each server must publish a global, consistent and protected view of its published key and the cryptosystem in use. Without access to that view, the system appears to have no failure mode.

With a small number of servers, the ecosystem would likely be dominated by a few providers. With a dominant position in the market these Privacy Pass server operators would have a significant impact on default connectivity parameters in operating systems and browsers. As a result, a change to the way the access mechanism works for a variety of applications would have broad impacts to a wide variety of users. The relationship between engineering and how it affects a broad community of users has a recent example in DNS over HTTP.

3.3. Practical Problems

Limits to the number of server operators also results in practical problems outside the protocol. In the event that a small number of server operators appear in the Privacy Pass ecosystem, and a large number of clients enter into trust relationships with those operators, what happens when those operators are acquired by other organizations that have different data handling and privacy policies than the original operator?

With the requirement for a small number of operators, the architecture also doesn't consider the possibility that an organization or government could require Privacy Pass and the use of a particular set of servers. Such a requirement could potentially turn the goals of Privacy Pass against itself.

4. Problem Statement and Potential for Mitigations

4.1. Problem Statement

An upper bound to available Privacy Pass servers creates architectural, engineering and practical problems for the deployment of the protocol. Any successful deployment of Privacy Pass must find mitigations for these problems.

4.2. Potential Mitigations

The motivation for having an upper bound to available Privacy Pass servers is to limit the amount of information that could be gather because a client could be forced to redeem tokens for any issuing key. A large number of keys, means a greater about of information exposed.

One alternative to limiting the number of servers is to constrain the clients so that they only possess redemption tokens for a small number of servers. This potential mitigation doesn't address how the tokens might be cached, but it does discuss how the limitation might be implemented. However, there is much engineering experience to suggest that making a limitation work in a very large number of clients is a much greater engineering and deployment problem than placing the restriction in the server.

If the motivation for restricting the number of servers is essential for Privacy Pass - and the mitigations at either the server or client are difficult to overcome - it is hard to understand where the mitigations for the problem statement will emerge.

4.3. Redemption Contexts as a Mitigation

Contexts are groupings of resources that have shared anonymity and privacy properties. The current architecture statement has a single, global context for redemption. It is this feature that causes the problem outlined in section 4.1 above: with N issuers in the global ecosystem, there are 2^N possible anonymity sets. Adding additional metadata bits increases the number of anonymity sets.

The global redemption context results in a requirement of less than ten total issuers in order to maintain anonymity sets of 5,000.

One possible mitigation is to limit redemptions to a specific, shared context. Such an approach could limit the information available - and the potential for leakage - to a specific context. This type of solution would rely, in part, on strong

security/privacy boundaries between contexts. While information about redemptions in one context wouldn't affect information in another context, this solution depends upon there being no leakage of information between those contexts.

While this potential mitigation is not reflected in the Privacy Pass architecture, it is unclear whether it should be a part of the protocol design or it should be left to the application layer to implement. If left to the application layer, there is potential for the anonymity sets to be very small and not meet the privacy goals of the protocol.

5. Security Considerations

This document is all about security considerations for Privacy Pass. In particular it addresses the very specific problem associated with centralization of Privacy Pass servers.

6. IANA Considerations

This memo contains no instructions or requests for IANA. The authors continue to appreciate the efforts of IANA staff in support of the IETF.

7. References

7.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

[2] 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>.

[3] [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>.

8. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Mark McFadden
Internet policy advisors, ltd
Chepstow, Wales, United Kingdom

Email: mark@internetpolicyadvisors.com