

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 14 July 2024

M. Kühlewind
Ericsson Research
T. Pauly
Apple
C. A. Wood
Cloudflare
11 January 2024

Partitioning as an Architecture for Privacy
draft-iab-privacy-partitioning-05

Abstract

This document describes the principle of privacy partitioning, which selectively spreads data and communication across multiple parties as a means to improve privacy by separating user identity from user data. This document describes emerging patterns in protocols to partition what data and metadata is revealed through protocol interactions, provides common terminology, and discusses how to analyze such models.

Discussion Venues

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

Discussion of this document takes place on the Internet Architecture Board Internet Engineering Task Force mailing list (iab@iab.org), which is archived at <https://www.iab.org/>.

Source for this draft and an issue tracker can be found at <https://github.com/intarchboard/draft-obliviousness>.

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 14 July 2024.

Copyright Notice

Copyright (c) 2024 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.

Table of Contents

1.	Introduction	3
2.	Privacy Partitioning	4
2.1.	Privacy Contexts	5
2.2.	Context Separation	7
2.3.	Approaches to Partitioning	8
3.	A Survey of Protocols using Partitioning	9
3.1.	CONNECT Proxying and MASQUE	9
3.2.	Oblivious HTTP and DNS	12
3.3.	Privacy Pass	14
3.4.	Privacy Preserving Measurement	15
4.	Applying Privacy Partitioning	16
4.1.	User-Identifying Information	16
4.2.	Selecting Client Identifiers	17
4.3.	Incorrect or Incomplete Partitioning	18
4.4.	Selecting Information Within a Context	18
5.	Limits of Privacy Partitioning	19
5.1.	Violations by Collusion	19
5.2.	Violations by Insufficient or Incorrect Partitioning	20
5.2.1.	Violations from Application Information	20
5.2.2.	Violations from Network Information	20
5.2.3.	Violations from Side Channels	21
5.2.4.	Identifying Partitions	21
6.	Partitioning Impacts	21
7.	Security Considerations	23
8.	IANA Considerations	23
9.	Informative References	23
	Acknowledgments	26
	Authors' Addresses	26

1. Introduction

Protocols such as TLS and IPsec provide a secure (authenticated and encrypted) channel between two endpoints over which endpoints transfer information. Encryption and authentication of data in transit are necessary to protect information from being seen or modified by parties other than the intended protocol participants. As such, this kind of security is necessary for ensuring that information transferred over these channels remains private.

However, a secure channel between two endpoints is insufficient for the privacy of the endpoints themselves. In recent years, privacy requirements have expanded beyond the need to protect data in transit between two endpoints. Some examples of this expansion include:

- * A user accessing a service on a website might not consent to reveal their location, but if that service is able to observe the client's IP address, it can learn something about the user's location. This is problematic for privacy since the service can link user data to the user's location.
- * A user might want to be able to access content for which they are authorized, such as a news article, without needing to have which specific articles they read on their account being recorded. This is problematic for privacy since the service can link user activity to the user's account.
- * A client device that needs to upload metrics to an aggregation service might want to be able to contribute data to the system without having their specific contributions attributed to them. This is problematic for privacy since the service can link client contributions to the specific client.

The commonality in these examples is that clients want to interact with or use a service without exposing too much user-specific or identifying information to that service. In particular, separating the user-specific identity information from user-specific data is necessary for privacy. Thus, in order to protect user privacy, it is important to keep identity (who) and data (what) separate.

This document defines "privacy partitioning," sometimes also referred to as the "decoupling principle" [DECOUPLING], as the general technique used to separate the data and metadata visible to various parties in network communication, with the aim of improving user privacy. Although privacy partitioning cannot guarantee there is no link between user-specific identity and user-specific data, when applied properly it helps ensure that user privacy violations become more technically difficult to achieve over time.

Several IETF working groups are working on protocols or systems that adhere to the principle of privacy partitioning, including Oblivious HTTP Application Intermediation (OHAI), Multiplexed Application Substrate over QUIC Encryption (MASQUE), Privacy Pass, and Privacy Preserving Measurement (PPM). This document summarizes work in those groups and describes a framework for reasoning about the resulting privacy posture of different endpoints in practice.

Privacy partitioning is particularly relevant as a tool for data minimization, which is described in Section 6.1 of [RFC6973]. [RFC6973] provides guidance for privacy considerations in Internet protocols, along with a set of questions on how to evaluate the data minimization of a protocol in Section 7.1 of [RFC6973]. Protocols that employ privacy partitioning ought to consider the questions in that section when evaluating their design, particularly with regard to how identifiers and data can be correlated by protocol participants and observers in each context that has been partitioned. Privacy partitioning can also be used as a way to separate identity providers from relying parties (see Section 6.1.4 of [RFC6973]), as in the case of Privacy Pass (see Section Section 3.3).

Privacy partitioning is not a panacea; applying it well requires holistic analysis of the system in question to determine whether or not partitioning as a tool, and as implemented, offers meaningful privacy improvements. See Section 5 for more information.

2. Privacy Partitioning

For the purposes of user privacy, this document focuses on user-specific information. This might include any identifying information that is specific to a user, such as their email address or IP address, or data about the user, such as their date of birth. Informally, the goal of privacy partitioning is to ensure that each party in a system beyond the user themselves only have access to one type of user-specific information.

This is a simple application of the principle of least privilege, wherein every party in a system only has access to the minimum amount of information needed to fulfill their function. Privacy partitioning advocates for this minimization by ensuring that protocols, applications, and systems only reveal user-specific information to parties that need access to the information for their intended purpose.

Put simply, privacy partitioning aims to separate who someone is from what they do. In the rest of this section, we describe how privacy partitioning can be used to achieve this goal.

2.1. Privacy Contexts

Each piece of user-specific information exists within some context, where a context is abstractly defined as a set of data, metadata, and the entities that share access to that information. In order to prevent the correlation of user-specific information across contexts, partitions need to ensure that any single entity (other than the client itself) does not participate in more than one context where the information is visible.

[RFC6973] discusses the importance of identifiers in reducing correlation as a way of improving privacy:

Correlation is the combination of various pieces of information related to an individual or that obtain that characteristic when combined... Correlation is closely related to identification. Internet protocols can facilitate correlation by allowing individuals' activities to be tracked and combined over time.

Pseudonymity is strengthened when less personal data can be linked to the pseudonym; when the same pseudonym is used less often and across fewer contexts; and when independently chosen pseudonyms are more frequently used for new actions (making them, from an observer's or attacker's perspective, unlinkable).

Context separation is foundational to privacy partitioning and reducing correlation. As an example, consider an unencrypted HTTP session over TCP, wherein the context includes both the content of the transaction as well as any metadata from the transport and IP headers; and the participants include the client, routers, other network middleboxes, intermediaries, and server. Middleboxes or intermediaries might simply forward traffic, or might terminate the traffic at any layer (such as terminating the TCP connection from the client and creating another TCP connection to the server). Regardless of how the middlebox interacts with the traffic, for the purposes of privacy partitioning, it is able to observe all of the data in the context.




Figure 1: Diagram of a basic unencrypted client-to-server connection with middleboxes

Adding TLS encryption to the HTTP session is a simple partitioning technique that splits the previous context into two separate contexts: the content of the transaction is now only visible to the client, TLS-terminating intermediaries, and server; while the metadata in transport and IP headers remain in the original context. In this scenario, without any further partitioning, the entities that participate in both contexts can allow the data in both contexts to be correlated.

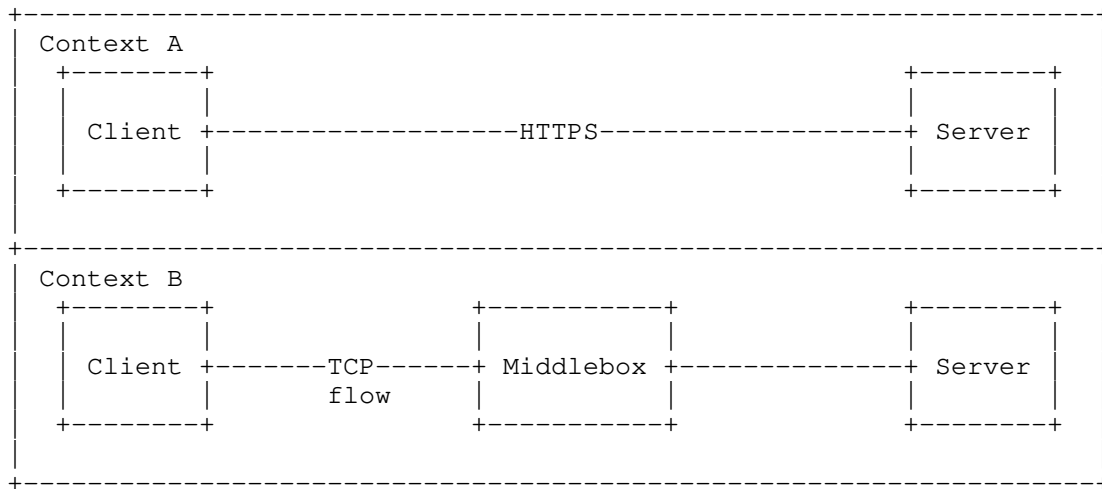


Figure 2: Diagram of how adding encryption splits the context into two

Another way to create a partition is to simply use separate connections. For example, to split two separate HTTP requests from one another, a client could issue the requests on separate TCP connections, each on a different network, and at different times; and avoid including obvious identifiers like HTTP cookies across the requests.




Figure 3: Diagram of making separate connections to generate separate contexts

Using separate connections to create separate contexts can reduce or eliminate the ability of specific parties to correlate activity across contexts. However, any identifier at any layer that is common across different contexts can be used as a way to correlate activity. Beyond IP addresses, many other factors can be used together to create a fingerprint of a specific device (such as MAC addresses, device properties, software properties and behavior, application state, etc).

2.2. Context Separation

In order to define and analyze how various partitioning techniques work, the boundaries of what is being partitioned need to be established. This is the role of context separation. In particular, in order to prevent the correlation of user-specific information across contexts, partitions need to ensure that any single entity (other than the client itself) does not participate in contexts where both identifiers are visible.

Context separation can be achieved in different ways, for example, over time, across network paths, based on (en)coding, etc. The privacy-oriented protocols described in this document generally involve more complex partitioning, but the techniques to partition communication contexts still employ the same techniques:

1. Cryptographic protection, such as the use of encryption to specific parties, allows partitioning of contexts between different parties (those with the ability to remove cryptographic protections, and those without).
2. Connection separation across time or space to allow partitioning of contexts for different application transactions over the network.

These techniques are frequently used in conjunction for context separation. For example, encrypting an HTTP exchange using TLS between client and TLS-terminating server might prevent a network middlebox that sees a client IP address from seeing the user account identifier, but it doesn't prevent the TLS-terminating server from observing both identifiers and correlating them. As such, preventing correlation requires separating contexts, such as by using proxying to conceal a client's IP address that would otherwise be used as an identifier.

2.3. Approaches to Partitioning

While all of the partitioning protocols described in this document create separate contexts using cryptographic protection and/or connection separation, each one has a unique approach that results in different sets of contexts. Since many of these protocols are new, it is yet to be seen how each approach will be used at scale across the Internet, and what new models will emerge in the future.

There are multiple factors that lead to a diversity in approaches to partitioning, including:

- * Adding privacy partitioning to existing protocol ecosystems places requirements and constraints on how contexts are constructed. CONNECT-style proxying is intended to work with servers that are unaware of privacy contexts, requiring more intermediaries to provide strong separation guarantees. Oblivious HTTP, on the other hand, assumes servers that cooperate with context separation, and thus reduces the overall number of elements in the solution.

- * Whether or not information exchange needs to happen bidirectionally in an interactive fashion determines how contexts can be separated. Some use cases, like metrics collection for PPM, can occur with information flowing only from clients to servers, and can function even when clients are no longer connected. Privacy Pass is an example of a case that can be either interactive or not, depending on whether tokens can be cached and reused. CONNECT-style proxying and Oblivious HTTP often requires bidirectional and interactive communication.
- * The degree to which contexts need to be partitioned depends in part on the client's threat models and level of trust in various protocol participants. For example, in Oblivious HTTP, clients allow relays to learn that clients are accessing a particular application-specific gateway. If clients do not trust relays with this information, they can instead use a multi-hop CONNECT-style proxy approach wherein no single party learns whether specific clients are accessing a specific application. This is the default trust model for systems like Tor, where multiple hops are used to drive down the probability of privacy violations due to collusion or related attacks.

3. A Survey of Protocols using Partitioning

The following section discusses current on-going work in the IETF that is applying privacy partitioning.

3.1. CONNECT Proxying and MASQUE

HTTP forward proxies, when using encryption on the connection between the client and the proxy, provide privacy partitioning by separating a connection into multiple segments. When connections to targets via the proxy themselves are encrypted, the proxy cannot see the end-to-end content. HTTP has historically supported forward proxying for TCP-like streams via the CONNECT method. More recently, the Multiplexed Application Substrate over QUIC Encryption (MASQUE) working group has developed protocols to similarly proxy UDP [CONNECT-UDP] and IP packets [CONNECT-IP] based on tunneling.

In a single-proxy setup, there is a tunnel connection between the client and proxy and an end-to-end connection that is tunnelled between the client and target. This setup, as shown in the figure below, partitions communication into:

- * a Client-to-Target encrypted context, which contains the end-to-end content within the TLS session to the target, such as HTTP content;

- * a Client-to-Target proxied context, which is the end-to-end data to the target that is also visible to the proxy, such as a TLS session;
- * a Client-to-Proxy context, which contains the transport metadata between the client and the target, and the request to the proxy to open a connection to the target;
- * and a Proxy-to-Target context, which for TCP and UDP proxying contains any packet header information that is added or modified by the proxy, e.g., the IP and TCP/UDP headers.

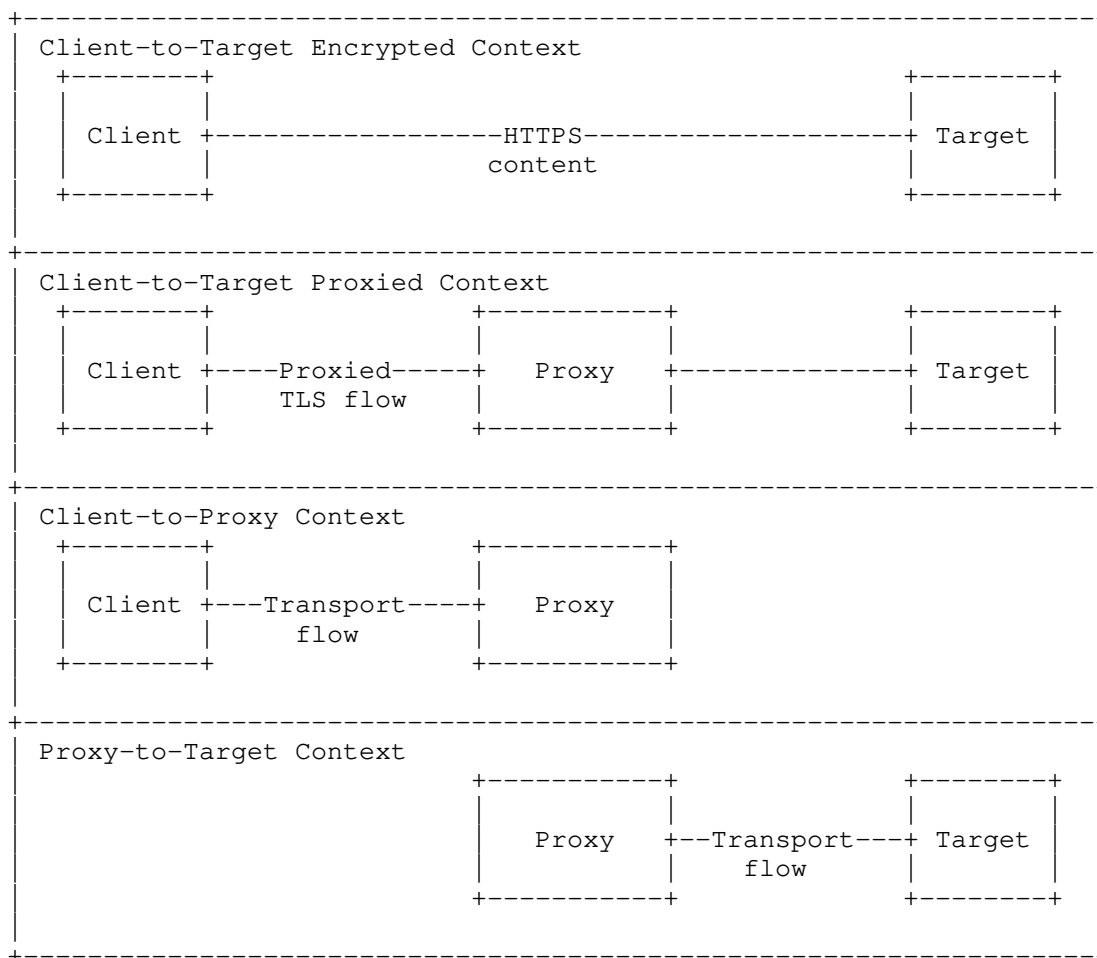
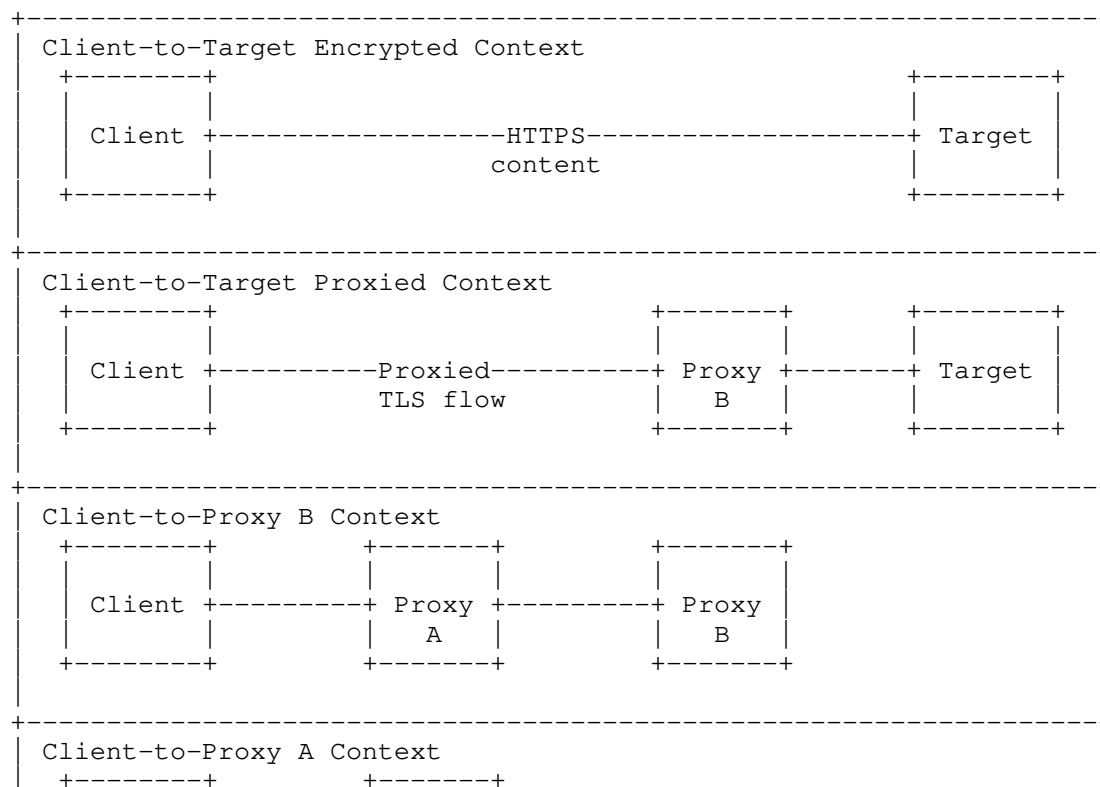


Figure 4: Diagram of one-hop proxy contexts

Using two (or more) proxies provides better privacy partitioning. In particular, with two proxies, each proxy sees the Client metadata, but not the Target; the Target, but not the Client metadata; or neither.

In addition to the contexts described above for the single proxy case, the two-hop proxy case shown in the figure below changes the contexts in several ways:

- * the Client-to-Target proxied context only includes the second proxy (referred to here as "Proxy B");
- * a new Client-to-Proxy B context is added, which is the TLS session from the client to Proxy B that is also visible to the first proxy (referred to here as "Proxy A");
- * the contexts that see transport data only (TCP or UDP over IP) are separated out into three separate contexts, a Client-to-Proxy A context, a Proxy A-to-Proxy B context, and a Proxy B-to-Target context.



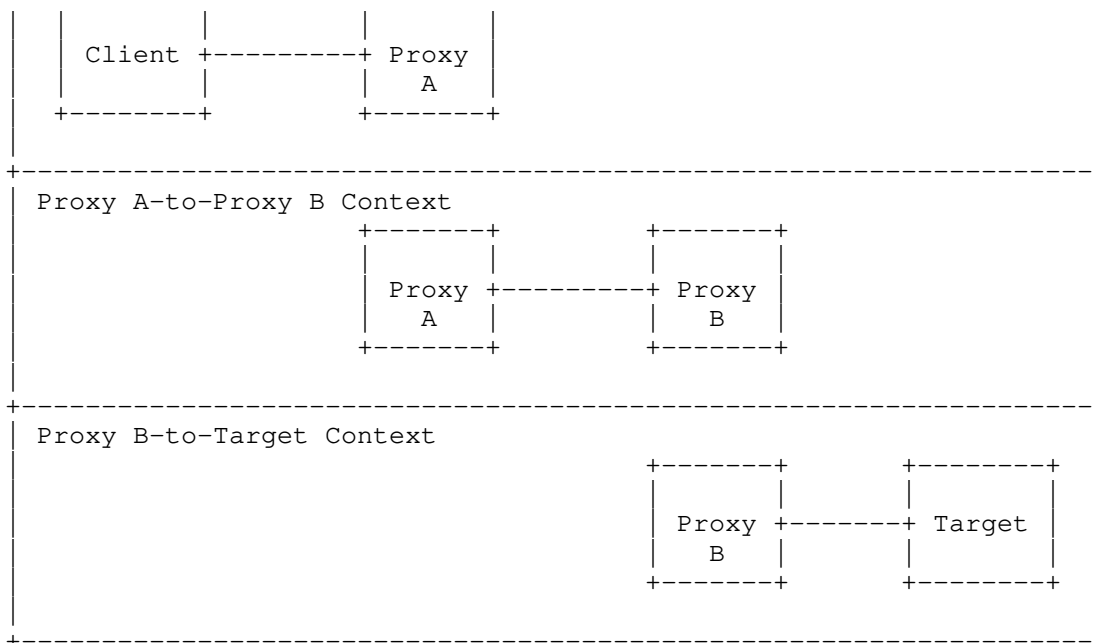


Figure 5: Diagram of two-hop proxy contexts

Forward proxying, such as the protocols developed in MASQUE, uses both encryption (via TLS) and separation of connections (via proxy hops that see only the next hop) to achieve privacy partitioning.

3.2. Oblivious HTTP and DNS

Oblivious HTTP [OHTTP], developed in the Oblivious HTTP Application Intermediation (OHAI) working group, adds per-message encryption to HTTP exchanges through a relay system. Clients send requests through an Oblivious Relay, which cannot read message contents, to an Oblivious Gateway, which can decrypt the messages but cannot communicate directly with the client or observe client metadata like IP address. Oblivious HTTP relies on Hybrid Public Key Encryption [HPKE] to perform encryption.

Oblivious HTTP uses both encryption and separation of connections to achieve privacy partitioning.

- * End-to-end messages are encrypted between the Client and Gateway. The content of these inner messages are visible to the Client, Gateway, and Target. This is the Client-to-Target context.

- * The encrypted messages exchanged between the Client and Gateway are visible to the Relay, but the Relay cannot decrypt the messages. This is the Client-to-Gateway context.
- * The transport (such as TCP and TLS) connections between the Client, Relay, and Gateway form two separate contexts: a Client-to-Relay context and a Relay-to-Gateway context. It is important to note that the Relay-to-Gateway connection can be a single connection, even if the Relay has many separate Clients. This provides better anonymity by making the pseudonym presented by the Relay to be shared across many Clients.




Figure 6: Diagram of Oblivious HTTP contexts

Oblivious DNS over HTTPS [ODOH] applies the same principle as Oblivious HTTP, but operates on DNS messages only. As a precursor to the more generalized Oblivious HTTP, it relies on the same HPKE cryptographic primitives, and can be analyzed in the same way.

3.3. Privacy Pass

Privacy Pass is an architecture [PRIVACYPASS] and a set of protocols being developed in the Privacy Pass working group that allows clients to present proof of verification in an anonymous and unlinkable fashion, via tokens. These tokens originally were designed as a way to prove that a client had solved a CAPTCHA, but can be applied to other types of user or device attestation checks as well. In Privacy Pass, clients interact with an attester and issuer for the purposes of issuing a token, and clients then interact with an origin server to redeem said token.

In Privacy Pass, privacy partitioning is achieved with cryptographic protection (in the form of blind signature protocols or similar) and separation of connections across two contexts: a "redemption context" between clients and origins (servers that request and receive tokens), and an "issuance context" between clients, attestation servers, and token issuance servers. The cryptographic protection ensures that information revealed during the issuance context is separated from information revealed during the redemption context.




Figure 7: Diagram of contexts in Privacy Pass

Since the redemption context and issuance context are separate connections that involve separate entities, they can also be further decoupled by running those parts of the protocols at different times. Clients can fetch tokens through the issuance context early, and cache the tokens to later use in redemption contexts. This can aid in partitioning identifiers and data.

[PRIVACYPASS] describes different deployment models for which entities operate origins, attesters, and issuers; in some models, they are all separate entities, but in others, they can be operated by the same entity. The model impacts the effectiveness of partitioning, and some models (such as when all three are operated by the same entity) only provide effective partitioning when the timing of connections on the two contexts are not correlated, and when the client uses different identifiers (such as different IP addresses) on each context.

3.4. Privacy Preserving Measurement

The Privacy Preserving Measurement (PPM) working group is chartered to develop protocols and systems that help a data aggregation or collection server (or multiple, non-colluding servers) compute aggregate values without learning the value of any one client's individual measurement. Distributed Aggregation Protocol (DAP) is the primary working item of the group.

At a high level, DAP uses a combination of cryptographic protection (in the form of secret sharing amongst non-colluding servers) to establish two contexts: an "upload context" between clients and non-colluding aggregation servers (in which the servers are separated into "Helper" and "Leader" roles) wherein aggregation servers possibly learn client identity but nothing about their individual measurement reports, and a "collect context" wherein a collector learns aggregate measurement results and nothing about individual client data.




Figure 8: Diagram of contexts in DAP

4. Applying Privacy Partitioning

Applying privacy partitioning to an existing or new system or protocol requires the following steps:

1. Identify the types of information used or exposed in a system or protocol, some of which can be used to identify a user or correlate to other contexts.
2. Partition data to minimize the amount of user-identifying or correlatable information in any given context to only include what is necessary for that context, and prevent the sharing of data across contexts wherever possible.

The most impactful types of information to partition are (a) user-identifying information, such as user identifiers (including account names or IP addresses) that can be linked and (b) non-user-identifying information (including content a user generates or accesses), which can be often sensitive when combined with a user identifier.

In this section, we discuss considerations for partitioning these types of information.

4.1. User-Identifying Information

User data can itself be user-identifying, in which case it should be treated as an identifier. For example, Oblivious DoH and Oblivious HTTP partition the client IP address and client request data into separate contexts, thereby ensuring that no entity beyond the client can observe both. Collusion across contexts could reverse this partitioning, but can also promote non-user-identifying information to user-identifying. For example, in CONNECT proxy systems that use QUIC, the QUIC connection ID is inherently non-user-identifying since it is generated randomly ([QUIC], Section 5.1). However, if combined with another context that has user-identifying information such as the client IP address, the QUIC connection ID can become user-identifying information.

Some information is innate to client user-agents, including details of implementation of protocols in hardware and software, and network location. This information can be used to construct user-identifying information, which is a process sometimes referred to as fingerprinting. Depending on the application and system constraints, users may not be able to prevent fingerprinting in privacy contexts. As a result, fingerprinting information, when combined with non-user-identifying user data, could promote user data to user-identifying information.

4.2. Selecting Client Identifiers

The selection of client identifiers used in the contexts used for privacy partitioning has a large impact on the effectiveness of partitioning. Identifier selection can either undermine or improve the value of partitioning. Generally, each context involves some form of client identifier, which might be directly associated with a client identity, but can also be a pseudonym or a random one-time identifier.

Using the same client identifier across multiple contexts can partly or wholly undermine the effectiveness of partitioning, by allowing the various contexts to be linked back to the same client. For example, if a client uses proxies as described in Section 3.1 to separate connections, but uses the same email address to authenticate to two servers in different contexts, those actions can be linked back to the same client. While this does not undermine all of the partitioning achieved through proxying (the contexts along the network path still cannot correlate the client identity and what servers are being accessed), the overall effect of partitioning is diminished.

When possible, using per-context unique client identifiers provides better partitioning properties. For example, a client can use a unique email address as an account identifier with each different server it needs to log into. The same approach can apply across many layers, as seen with per-network MAC address randomization [I-D.ietf-madinas-mac-address-randomization], use of multiple temporary IP addresses across connections and over time [RFC8981], and use of unique per-subscription identifiers for HTTP Web Push [RFC8030].

4.3. Incorrect or Incomplete Partitioning

Privacy partitioning can be applied incorrectly or incompletely. Contexts may contain more user-identifying information than desired, or some information in a context may be more user-identifying than intended. Moreover, splitting user-identifying information over multiple contexts has to be done with care, as creating more contexts can increase the number of entities that need to be trusted to not collude. Nevertheless, partitions can help improve the client's privacy posture when applied carefully.

Evaluating and qualifying the resulting privacy of a system or protocol that applies privacy partitioning depends on the contexts that exist and the types of user-identifying information in each context. Such evaluation is helpful for identifying ways in which systems or protocols can improve their privacy posture. For example, consider DNS-over-HTTPS [DOH], which produces a single context which contains both the client IP address and client query. One application of privacy partitioning results in ODoH, which produces two contexts, one with the client IP address and the other with the client query.

4.4. Selecting Information Within a Context

Recognizing potential applications of privacy partitioning requires identifying the contexts in use, the information exposed in a context, and the intent of information exposed in a context. Unfortunately, determining what information to include in a given context is a non-trivial task. In principle, the information contained in a context should be fit for purpose. As such, new systems or protocols developed should aim to ensure that all information exposed in a context serves as few purposes as possible. Designing with this principle from the start helps mitigate issues that arise if users of the system or protocol inadvertently ossify on the information available in contexts. Legacy systems that have ossified on information available in contexts may be difficult to change in practice. As an example, many existing anti-abuse systems depend on some client identifier such as client IP address, coupled with client data, to provide value. Partitioning contexts in these systems such that they no longer determine the client identity requires new solutions to the anti-abuse problem.

5. Limits of Privacy Partitioning

Privacy partitioning aims to increase user privacy, though as stated, it is merely one of possibly many architectural tools that help manage privacy risks. Understanding the limits of its benefits requires a more comprehensive analysis of the system in question. Such analysis also helps determine whether or not the tool has been applied correctly. In particular, the value of privacy partitioning depends on numerous factors, including, though not limited to:

- * Non-collusion across contexts; and
- * The type of information exposed in each context.

We elaborate on each below.

5.1. Violations by Collusion

Privacy partitions ensure that only the client, i.e., the entity which is responsible for partitioning, can independently link all user-specific information. No other entity individually knows how to link all the user-specific information as long as they do not collude with each other across contexts. Thus, non-collusion is a fundamental requirement for privacy partitioning to offer meaningful privacy for end-users. In particular, the trust relationships that users have with different parties affect the resulting impact on the user's privacy.

As an example, consider OHTTP, wherein the Oblivious Relay knows the client identity but not the client data, and the Oblivious Gateway knows the client data but not the client identity. If the Oblivious Relay and Gateway collude, they can link client identity and data together for each request and response transaction by simply observing requests in transit.

It is not currently possible to guarantee with technical protocol measures that two entities are not colluding. Even if two entities do not collude directly, if both entities reveal information to other parties, it will not be possible to guarantee that the information won't be combined. However, there are some mitigations that can be applied to reduce the risk of collusion happening in practice:

- * Policy and contractual agreements between entities involved in partitioning to disallow logging or sharing of data, along with auditing to validate that the policies are being followed. For cases where logging is required (such as for service operation), such logged data should be minimized and anonymized to prevent it from being useful for collusion.

- * Protocol requirements to make collusion or data sharing more difficult.
- * Adding more partitions and contexts, to make it increasingly difficult to collude with enough parties to recover identities.

5.2. Violations by Insufficient or Incorrect Partitioning

Insufficient or incorrect application of privacy partitioning can lessen or negate benefits to users. In particular, it is possible to apply partitioning in a way that is either insufficient or incorrect for meaningful privacy. For example, partitioning at one layer in the stack can fail to account for linkable information at different layers in the stack. Privacy violations can stem from partitioning failures in a multitude of ways, some of which are described below.

5.2.1. Violations from Application Information

Partitioning at the network layer can be insufficient when the application layer fails to properly partition. As an example, consider OHTTP used for the purposes of hiding client-identifying information for a browser telemetry system. It is entirely possible for reports in such a telemetry system to contain both client-specific telemetry data, such as information about their specific browser instance, as well as client-identifying information, such as the client's email address, location, or IP address. Even though OHTTP separates the client IP address from the server via a relay, the server can still learn this directly from the client's telemetry report.

5.2.2. Violations from Network Information

It is also possible to inadequately partition at the network layer. As an example, consider both TLS Encrypted Client Hello (ECH) [I-D.ietf-tls-esni] and VPNs. ECH uses cryptographic protection (encryption) to hide information from unauthorized parties, but both clients and servers (two entities) can link user-specific data to user-specific identifier (IP address). Similarly, while VPNs hide identifiers from end servers, the VPN server can still see the identifiers of both the client and server. Applying privacy partitioning would advocate for at least two additional entities to avoid revealing both identity (who) and user actions (what) from each involved party.

5.2.3. Violations from Side Channels

Beyond the information that is intentionally revealed by applying privacy partitioning, it is also possible for the information to be unintentionally revealed through side channels. For example, in the two-hop proxy arrangement described in Section 3.1, Proxy A sees and proxies TLS data between the client and Proxy B. While it does not directly learn information that Proxy B sees, it does learn information through metadata, such as the timing and size of encrypted data being proxied. Traffic analysis could be exploited to learn more information from such metadata, including, in some cases, application data that Proxy A was never meant to see. Although privacy partitioning does not obviate such attacks, it does increase the cost necessary to carry them out in practice. See Section 7 for more discussion on this topic.

5.2.4. Identifying Partitions

While straightforward violations of user privacy that stem from insufficient partitioning may seem straightforward to mitigate, it remains an open problem to rigorously determine what information needs to be partitioned for meaningful privacy, and to implement it in a way that achieves the desired properties. In essence, it is difficult to determine whether a certain set of information reveals "too much" about a specific user, and it is similarly challenging to determine whether or not an implementation of partitioning works as intended. There is ample evidence of data being assumed "private" or "anonymous" but, in hindsight, winds up revealing too much information such that it allows one to link back to individual clients; see [DataSetReconstruction] and [CensusReconstruction] for more examples of this in the real world.

6. Partitioning Impacts

Applying privacy partitioning to communication protocols leads to a substantial change in communication patterns. For example, instead of sending traffic directly to a service, essentially all user traffic is routed through a set of intermediaries, possibly adding more end-to-end round trips in the process (depending on the system and protocol). This has a number of practical implications, described below.

1. Service operational or management challenges. Information that is traditionally passively observed in the network or metadata that has been unintentionally revealed to the service provider cannot be used anymore for e.g., existing security procedures such as application rate limiting or DDoS mitigation. However, network management techniques deployed at present often rely on information that is exposed by most traffic but without any guarantees that the information is accurate.

Privacy partitioning provides an opportunity for improvements in these management techniques by enabling active exchange of information with each entity in a privacy-preserving way and requesting exactly the information needed for a specific task or function rather than relying on the assumption that are derived from a limited set of unintentionally revealed information which cannot be guaranteed to be present and may disappear at any time in future.

2. Varying performance effects and costs. Depending on how context separation is done, privacy partitioning may affect application performance. As an example, Privacy Pass introduces an entire end-to-end round trip to issue a token before it can be redeemed, thereby decreasing performance. In contrast, while systems like CONNECT proxying may seem like they would regress performance, oftentimes the highly optimized nature of proxy-to-proxy paths leads to improved performance.

Performance may also push back against the desire to apply privacy partitioning. For example, HTTPS connection reuse [HTTP2], Section 9.1.1 allows clients to use an existing HTTPS session created for one origin to interact with different origins (provided the original origin is authoritative for these alternative origins). Reusing connections saves the cost of connection establishment, but means that the server can now link the client's activity with these two or more origins together. Applying privacy partitioning would prevent this, while typically at the cost of less performance.

In general, while performance and privacy tradeoffs are often cast as a zero-sum game, in practice this is often not the case. The relationship between privacy and performance varies depending on a number of related factors, such as application characteristics, network path properties, and so on.

3. Increased attack surface. Even in the event that information is adequately partitioned across non-colluding parties, the resulting effects on the end-user may not always be positive. For example, using OHTTP as a basis for illustration, consider a

hypothetical scenario where the Oblivious Gateway has an implementation flaw that causes all of its decrypt requests to be inappropriately logged to a public or otherwise compromised location. Moreover, assume that the Target Resource for which these requests are destined does not have such an implementation flaw. Applications which use OHTTP with this flawed Oblivious Gateway to interact with the Target Resource risk their user request information being made public, albeit in a way that is decoupled from user identifying information, whereas applications that do not use OHTTP to interact with the Target Resource do not risk this type of disclosure.

4. Centralization. Depending on the protocol and system, as well as the desired privacy properties, the use of partitioning may inherently force centralization to a selected set of trusted participants. As an example, the impact of OHTTP on end-user privacy generally increases proportionally to the number of users that exist behind a given Oblivious Relay. That is, the probability of an Oblivious Gateway determining the client associated with a request forwarded through an Oblivious Relay decreases as the number of possible clients behind the Oblivious Relay increases. This tradeoff encourages the centralization of the Oblivious Relays.

7. Security Considerations

Section 5 discusses some of the limitations of privacy partitioning in practice, and advocates for holistic analysis to understand the extent to which privacy partitioning offers meaningful privacy improvements. Applied correctly, partitioning helps improve an end-user's privacy posture, thereby making violations harder to do via technical, social, or policy means. For example, side channels such as traffic analysis [I-D.irtf-pearg-website-fingerprinting] or timing analysis are still possible and can allow an unauthorized entity to learn information about a context they are not a participant of. Proposed mitigations for these types of attacks, e.g., padding application traffic or generating fake traffic, can be very expensive and are therefore not typically applied in practice. Nevertheless, privacy partitioning moves the threat vector from one that has direct access to user-specific information to one which requires more effort, e.g., computational resources, to violate end-user privacy.

8. IANA Considerations

This document has no IANA actions.

9. Informative References

[CensusReconstruction]

"The Census Bureau's Simulated Reconstruction-Abetted Re-identification Attack on the 2010 Census", n.d., <<https://www.census.gov/data/academy/webinars/2021/disclosure-avoidance-series/simulated-reconstruction-abetted-re-identification-attack-on-the-2010-census.html>>.

[CONNECT-IP]

Pauly, T., Schinazi, D., Chernyakhovsky, A., Kühlewind, M., and M. Westerlund, "Proxying IP in HTTP", Work in Progress, Internet-Draft, draft-ietf-masque-connect-ip-13, 28 April 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-masque-connect-ip-13>>.

[CONNECT-UDP]

Schinazi, D. and L. Pardue, "HTTP Datagrams and the Capsule Protocol", RFC 9297, DOI 10.17487/RFC9297, August 2022, <<https://www.rfc-editor.org/rfc/rfc9297>>.

[DataSetReconstruction]

Narayanan, A. and V. Shmatikov, "Robust De-anonymization of Large Sparse Datasets", IEEE, 2008 IEEE Symposium on Security and Privacy (sp 2008), DOI 10.1109/sp.2008.33, May 2008, <<https://doi.org/10.1109/sp.2008.33>>.

[DECOUPLING]

Schmitt, P., Iyengar, J., Wood, C., and B. Raghavan, "The decoupling principle: a practical privacy framework", ACM, Proceedings of the 21st ACM Workshop on Hot Topics in Networks, DOI 10.1145/3563766.3564112, November 2022, <<https://doi.org/10.1145/3563766.3564112>>.

[DOH]

Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/rfc/rfc8484>>.

[HPKE]

Barnes, R., Bhargavan, K., Lipp, B., and C. Wood, "Hybrid Public Key Encryption", RFC 9180, DOI 10.17487/RFC9180, February 2022, <<https://www.rfc-editor.org/rfc/rfc9180>>.

[HTTP2]

Thomson, M., Ed. and C. Benfield, Ed., "HTTP/2", RFC 9113, DOI 10.17487/RFC9113, June 2022, <<https://www.rfc-editor.org/rfc/rfc9113>>.

[I-D.ietf-madinas-mac-address-randomization]

Zúñiga, J. C., Bernardos, C. J., and A. Andersdotter, "Randomized and Changing MAC Address state of affairs", Work in Progress, Internet-Draft, draft-ietf-madinas-mac-

address-randomization-10, 11 January 2024,
<<https://datatracker.ietf.org/doc/html/draft-ietf-madinas-mac-address-randomization-10>>.

[I-D.ietf-tls-esni]

Rescorla, E., Oku, K., Sullivan, N., and C. A. Wood, "TLS Encrypted Client Hello", Work in Progress, Internet-Draft, draft-ietf-tls-esni-17, 9 October 2023,
<<https://datatracker.ietf.org/doc/html/draft-ietf-tls-esni-17>>.

[I-D.irtf-pearg-website-fingerprinting]

Goldberg, I., Wang, T., and C. A. Wood, "Network-Based Website Fingerprinting", Work in Progress, Internet-Draft, draft-irtf-pearg-website-fingerprinting-01, 8 September 2020, <<https://datatracker.ietf.org/doc/html/draft-irtf-pearg-website-fingerprinting-01>>.

[ODOH]

Kinnear, E., McManus, P., Pauly, T., Verma, T., and C.A. Wood, "Oblivious DNS over HTTPS", RFC 9230, DOI 10.17487/RFC9230, June 2022,
<<https://www.rfc-editor.org/rfc/rfc9230>>.

[OHTTP]

Thomson, M. and C. A. Wood, "Oblivious HTTP", Work in Progress, Internet-Draft, draft-ietf-ohai-ohttp-10, 25 August 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-ohai-ohttp-10>>.

[PRIVACYPASS]

Davidson, A., Iyengar, J., and C. A. Wood, "The Privacy Pass Architecture", Work in Progress, Internet-Draft, draft-ietf-privacypass-architecture-16, 25 September 2023,
<<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-architecture-16>>.

[QUIC]

Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021,
<<https://www.rfc-editor.org/rfc/rfc9000>>.

[RFC6973]

Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013,
<<https://www.rfc-editor.org/rfc/rfc6973>>.

- [RFC8030] Thomson, M., Damaggio, E., and B. Raymor, Ed., "Generic Event Delivery Using HTTP Push", RFC 8030, DOI 10.17487/RFC8030, December 2016, <<https://www.rfc-editor.org/rfc/rfc8030>>.
- [RFC8981] Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", RFC 8981, DOI 10.17487/RFC8981, February 2021, <<https://www.rfc-editor.org/rfc/rfc8981>>.

Acknowledgments

We would like to thank Martin Thomson, Eliot Lear, Mark Nottingham, Niels ten Oever, Vittorio Bertola, Antoine Fressancourt, Cullen Jennings, and Dhruv Dhody for their reviews and feedback.

Authors' Addresses

Mirja Kühlewind
Ericsson Research
Email: mirja.kuehlewind@ericsson.com

Tommy Pauly
Apple
Email: tpauly@apple.com

Christopher A. Wood
Cloudflare
Email: caw@heapingbits.net