Network Working Group
Internet-Draft
Intended status: Informational
Expires: 8 November 2026

D. Smullen
B. Scriber
CableLabs

7 May 2026

Privacy Preference Declaration for Home Networks
draft-dsmullen-ppd-architecture-05

Abstract

This document describes an architecture for signaling household privacy preferences to devices in home networks through Privacy Preference Declarations (PPDs). The architecture enables a PPD participant to discover a PPD service endpoint, establish trust in that endpoint through the applicable protocol and security profile, retrieve the applicable household policy instance, and acknowledge receipt of that policy instance. The acknowledgment establishes that a specific policy instance was made available to the participant; it does not, by itself, assert anything about the participant’s subsequent behavior.

About This Document

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

The latest revision of this draft can be found at https://drspangle.github.io/draft-dsmullen-ppd-architecture/draft-dsmullen-ppd-architecture.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-dsmullen-ppd-architecture/.

Source for this draft and an issue tracker can be found at https://github.com/drspangle/draft-dsmullen-ppd-architecture.

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/.

Copyright Notice

Table of Contents

1. Introduction

The rapid growth of Internet-connected devices in the home has introduced new and often overwhelming challenges to personal privacy. While many of these devices collect sensitive data by design, the tools offered to users to understand or control that collection are fragmented, confusing, or entirely absent. When privacy settings do exist, they are often buried in obscure menus, expressed in legal or technical jargon, and lack the contextual clarity needed to support meaningful decision-making.

The result is a fragmented operational model. Users must manage privacy through device-specific controls that vary widely in quality and semantics, while device vendors and service providers often implement isolated mechanisms with no common way to convey household privacy preferences across devices. This lack of a shared signaling model makes it difficult for households to understand which devices have been presented with which privacy expectations, and it makes interoperable deployment harder for implementers.

[RFC7258] frames mass data collection as a technical threat, urging protocol designers to limit exposure through encryption and data minimization. While this principle is crucial in adversarial, internet-scale contexts, the model proposed in this document takes a

different approach: rather than hiding data flows, it seeks to govern them. Privacy here is not achieved by making devices blind, but by making user-defined preferences visible to devices and associated services.

This use of privacy is related to, but distinct from, the privacy guidance in [RFC6973], which emphasizes reduced observability, linkability, and identifiability in protocol design. Those properties remain important, but PPD focuses on a different home-network problem: a user needs a consistent way to express household privacy preferences and to know that those preferences were made available to participating devices or associated services. This document is also aligned with the user-agency goals described in [RFC8280], but it is narrower and more operational. It describes an architecture for privacy-preference signaling and recordkeeping, not a general framework for human-rights analysis or for constraining device behavior. Home networks are a significant and operationally important IoT environment. They commonly place a local administrative boundary around large numbers of devices, many with limited or no end-user interface, making them a concrete target for a privacy-preference signaling architecture. In this architecture, discovery identifies candidate participant-facing service endpoints. Trust in a selected endpoint, and in the policy instances it presents, is established separately through the applicable protocol and security mechanisms rather than by discovery alone. This also addresses an asymmetry common in current deployments: the household user is often required to acknowledge device- or vendor-defined terms, while the household has no comparable way to record that a participating device or associated service was presented with the household’s privacy policy. PPD introduces a reciprocal signaling path in which presentation and acknowledgment of a household policy instance can be recorded by the household domain. The objective is to provide a coherent architectural basis for devices and associated services to retrieve, acknowledge, and keep current with household privacy preferences within that administrative domain.

2. Conventions and Definitions

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 terms below define both protocol roles and core concepts used by this architecture. The definitions of privacy, transparency, and user control are included here because they describe the conceptual scope of PPD rather than separate protocol mechanisms.

2.1. Terminology

Privacy: In this document, the ability of users to understand and shape how data about them, their household, or their home environment is collected, used, retained, and shared by devices and associated services.

Transparency: The property that data practices are made visible and understandable to the user, including what data is collected, how it is processed, where it is shared, and what policy preferences apply.

User control: The ability of a user or household to define privacy preferences and make those preferences visible to devices or associated services in a consistent and actionable way.

Privacy Preference Declaration (PPD): A structured expression of household privacy preferences that can be discovered, retrieved, and acknowledged by PPD participants.

PPD service endpoint: A participant-facing service, and the baseline discovery target for participants, through which a PPD participant discovers, retrieves, and acknowledges applicable policy instances.

Policy authority: The authoritative source of household policy state and of any inputs used to derive an effective policy for a participant. The policy authority may be local or remote. Participants are not required to discover or address the policy authority directly in the baseline architecture.

Household policy: A policy selected or maintained for a home network that represents the household’s privacy preferences.

Effective policy derivation: The logical function, performed by or on behalf of the policy authority, that determines the effective policy instance for a participant.

Effective policy: The policy instance that applies to a particular PPD participant at a particular time, after effective policy derivation has resolved household policy state and any applicable participant-specific inputs.

PPD participant: A device, or a backend service acting on behalf of a device, that participates in PPD by retrieving and acknowledging an applicable policy instance.

Policy instance: A specific version or representation of an

effective policy that can be identified for acknowledgment and recordkeeping.

Association: The state established when a PPD participant has retrieved the current applicable effective policy and acknowledged receipt of that specific policy instance.

Current association: Association state that still corresponds to the latest applicable effective policy for the participant and remains fresh according to the renewal model enforced by the PPD service endpoint.

Association freshness: The property that an association remains within the bounded interval, or before the renewal deadline, accepted by the PPD service endpoint for treating that association as current.

Stale association: Association state that still refers to the latest applicable effective policy instance, but whose freshness can no longer be confirmed because renewal did not occur within the bounded interval accepted by the PPD service endpoint.

Needs reassociation: A state in which current association cannot be confirmed because the applicable effective policy changed, participant state relevant to effective policy derivation changed, or enough state was lost that the existing association no longer applies reliably.

Reassociation: The process by which a PPD participant recovers from stale association or a needs-reassociation state and re-establishes current association.

Broken association: A state in which stored or reported information is contradictory or incomplete enough that current association cannot be determined reliably.

Policy acknowledgment: A signal that a PPD participant has received a specific effective policy instance. A policy acknowledgment is not a statement that the device is compatible with every policy term or that the device will behave in a particular way.

Network-observed device: A device that is visible to the local network through ordinary network observation but that has not established association through PPD.

Unmanaged device: A network-observed device that is not known to participate in PPD or is not currently manageable through PPD association.

3. Limitations of Existing Mechanisms

Current mechanisms for managing data privacy within the home environment exhibit limitations.

3.1. Device-specific Configurations

Individual devices often employ unique privacy settings, thereby complicating user management of privacy across the entire network. This complexity can inadvertently lead to unintended data sharing.

3.2. Ineffective and Unusable User Interfaces

Navigating and configuring privacy settings on individual devices can be a time-consuming and frustrating experience for users. These ineffective interfaces often lead users to habitually agree to relax their privacy preferences without fully understanding the implications of their decisions. This fosters a general resignation towards privacy management, making it difficult for users to exert meaningful control over their personal data and ultimately compromising their privacy expectations.

3.3. DNT and P3P

Protocols like Do Not Track (DNT) and Platform for Privacy Preferences Project (P3P) have not achieved widespread adoption and have proven inadequate for addressing nuanced privacy needs. These protocols lack the granularity required to express fine-grained user preferences and may not effectively prevent unintended data collection or sharing. For instance, consider a smart security camera within a home network. While DNT or P3P can signal a user’s general preference not to be tracked or share data, they do not enable users to specify detailed privacy settings, such as allowing the camera to record video but not audio, or restricting data sharing to only within the home network. Users need more precise control options to manage their privacy effectively and communicate their preferences across different devices and contexts. These limitations, coupled with the increasing complexity of the IoT ecosystem, hinder the effective exercise of user control over their data within the home environment and pose a significant threat to user privacy.

4. Operational Scenarios

This section describes representative operational cases for the architecture in home-network environments. The scenarios focus on discovery, association, reassociation, and mixed-participant visibility rather than on user-interface details.

4.1. Initial Discovery and Association

A PPD participant joins the home network and obtains one or more candidate PPD service endpoints through configuration or local network discovery. In a common home deployment model, the PPD service endpoint is hosted by a residential gateway or equivalent home-network service. Discovery identifies reachability, not authority. The participant establishes a secure connection to a selected endpoint, confirms that endpoint through the applicable trust mechanism, retrieves the applicable effective policy instance, and acknowledges receipt of that policy instance. The PPD service endpoint may present policy derived from a local or remote policy authority without exposing that internal topology to the participant. At the end of this process, the participant has established association if the current applicable effective policy has been delivered and acknowledged. The PPD service endpoint also determines the initial freshness state of that association.

4.2. Policy Update and Reassociation

The household policy, or the participant’s effective policy, changes. The PPD service endpoint immediately invalidates current association for the participant. The participant enters a needs-reassociation state until it retrieves and acknowledges the updated effective policy instance. This scenario illustrates that association state is tied to a specific policy instance and not to prior acknowledgments alone. Reassociation re-establishes current association by confirming that the participant has seen the latest applicable policy instance.

4.3. Association Freshness Expiry and Renewal

The applicable effective policy instance is unchanged, but the participant does not renew within the bounded interval accepted by the PPD service endpoint. The association becomes stale even though no policy change occurred. The participant no longer has current association until it completes the required renewal procedure. This scenario distinguishes stale association from a needs-reassociation state caused by a changed policy instance.

4.4. Participant State Change

A participant changes in a way that can affect the applicable effective policy instance, such as a declaration update, capability change, or other state change relevant to effective policy derivation. The PPD service endpoint determines that current association can no longer be confirmed using existing state alone. The participant then retrieves and acknowledges the newly applicable

effective policy instance. This scenario keeps the architecture focused on policy signaling and recordkeeping without assuming that every state change requires the same local handling or transport behavior.

4.5. Mixed-Participant Network Visibility

A home network contains both PPD participants and devices that do not participate in PPD. The household can still use local management functions to distinguish associated participants, participants whose current association cannot be confirmed, and network-observed or unmanaged devices. This scenario illustrates that non-participating devices are an expected operational reality. Their presence can inform transparency and local management decisions, but it does not create association or change the baseline signaling role of PPD.

5. Goals

5.1. Enhance User Control

* Support a household’s ability to define privacy preferences that can be made available consistently across participating devices and associated services.

* Provide an architectural basis for recording whether the current applicable policy was made available to a participant.

* Create a reciprocal acknowledgment model in which the household can retain a record that a participant or associated service was presented with, and acknowledged, a specific household policy instance.

5.2. Promote Interoperability

* Establish a standardized mechanism for devices from diverse manufacturers to discover PPD service endpoints, retrieve applicable privacy policies, and acknowledge policy instances.

* Support consistent association and reassociation behavior across heterogeneous participants.

5.3. Enable Flexibility

* Allow deployments to place policy storage and effective-policy derivation locally or remotely without changing the baseline participant-facing contract.

* Leave room for deployment-specific protocol profiles where constrained environments or different operational models require them.

5.4. Facilitate Transparency

* Provide a basis for local management functions to distinguish currently associated participants, stale or reassociation-needed participants, and non-participating devices.

* Improve visibility into which participants have been presented with the current applicable policy instance, without implying enforcement of device behavior.

6. Scope

This document focuses on defining a high-level architectural framework for a Privacy Preference Declaration (PPD) protocol specifically designed for home network environments. This document concentrates on the conceptual framework, key architectural components, and fundamental principles for enabling users to express privacy preferences and signal those preferences to devices within their home networks.

This document does not delve into specific implementation details, such as message formats, data structures, security algorithms, or user interface design. Furthermore, this document does not define mechanisms that modify device behavior, legal and regulatory considerations, or specific security protocols. Where this document discusses recordkeeping, that recordkeeping is limited to signaling and recording that an applicable household policy was made available to and acknowledged by a PPD participant. That recordkeeping can provide a basis for later accountability, audit, or dispute analysis, but this document does not define enforcement behavior or prove subsequent compliance.

Specific implementation details and message formats are expected to be addressed in companion specifications. This document aims to be complementary to existing and future standards related to home networking, IoT security, and data privacy.

This document provides the foundation for subsequent work, including:

* Privacy Preference Declaration Taxonomy: This document will define a taxonomy of privacy preference categories and attributes, including a mechanism for registration and management of these categories.

* Privacy Preference Declaration Protocol Specification: This document is expected to specify the message formats, data structures, and communication procedures for the PPD protocol, including mechanisms for PPD service endpoint discovery, policy retrieval, policy acknowledgment, participant resynchronization, and optional participant status reporting.

7. Architecture Overview

7.1. Assumptions

This document makes the following assumptions:

* User Control: It is assumed that users have a reasonable level of control over their home network infrastructure. This includes the ability to configure routers, install software updates, and manage device access to the network. This control is essential for users to effectively manage their privacy preferences and make them available to devices within their home network.

* Resource Constraints: It is assumed that the home network environment and devices operating therein have resource limitations, such as limited processing power and bandwidth. We limit this assumption by considering that the PPD protocol and its associated mechanisms should be designed with these constraints in mind, minimizing overhead and ensuring efficient operation even on resource-constrained devices.

* Security Considerations: It is assumed that home networks in scope of this document are susceptible to typical security threats, including insider threats (or non-malicious misconfiguration) and vulnerability to local attacks. We limit this assumption by considering specific security threats to protect user privacy and the integrity of the privacy policy. This includes considerations for secure policy dissemination, device authentication, and protection against unauthorized access and modification of privacy preferences.

* Single User Policy: This document assumes that each device implementing the protocol is governed by a single, unified privacy policy defined by its primary user. While other individuals within the same physical environment (e.g., household) may have different privacy preferences, the protocol is designed with the expectation that a device or associated service can receive the policy established by its primary user. Future extensions could explore mechanisms for managing and reconciling multiple user-defined policies on a single device, particularly in shared or multi-user environments.

* Endpoint Discovery and Trust: It is assumed that configuration or local network mechanisms can identify one or more candidate PPD service endpoints for a participant. Discovery alone does not establish that an endpoint is authoritative for household policy. The applicable protocol profile needs a separate way to authenticate the selected endpoint and confirm that policy presented through that endpoint is authoritative for the participant's household context.

* Policy Signaling: It is assumed that PPD participants can retrieve the applicable household privacy policy through a PPD service endpoint and acknowledge receipt of that policy instance. This acknowledgment forms the basis of association. It is a receipt signal only; it does not assert that the participant is compatible with every policy term or that it will behave in a particular way. Current association also depends on association freshness as determined by the PPD service endpoint.

* Association Freshness: It is assumed that current association expires unless the participant renews within a bounded interval accepted by the PPD service endpoint. Participant-initiated exchanges provide the renewal or recovery path, but the PPD service endpoint remains the source of truth for whether association is current, stale, or in a needs-reassociation state.

* Local Recordkeeping: At a minimum, the architecture enables the household to know which PPD participants have acknowledged the current applicable policy. Deployment-specific responses to participants that do not acknowledge policy, or to devices that do not participate in PPD, are local management decisions and are outside the baseline signaling function defined here.

7.2. Association State and Freshness

This architecture treats the PPD service endpoint as the source of truth for participant association state. A participant establishes association when it retrieves and acknowledges a specific applicable effective policy instance.

Current association exists only when both of the following are true:

* the acknowledged policy instance still corresponds to the latest applicable effective policy for that participant; and

* the association remains fresh according to the renewal model enforced by the PPD service endpoint.

If the applicable effective policy instance is unchanged but the freshness interval expires before renewal, the participant enters stale association. If the applicable effective policy changes, if participant state relevant to effective policy derivation changes, or if enough state is lost that the prior association can no longer be trusted, the participant enters a needs-reassociation state. In either case, the participant no longer has current association.

Participant-initiated exchanges provide the renewal or recovery path, but they are not the source of truth for whether association is current. The PPD service endpoint determines whether a participant is current, stale, or in needs reassociation.

7.3. Discovery and Policy-Authority Boundary

This architecture separates discovery of a participant-facing service endpoint from trust establishment. A participant may learn one or more candidate PPD service endpoints through configuration or local network mechanisms, but discovery alone does not make any candidate authoritative. Before treating a policy instance as authoritative, the participant needs the applicable protocol profile to authenticate the selected endpoint and confirm that it is authorized to present policy for the household context.

The participant-facing contract is the PPD service endpoint, not direct access to the policy authority. A deployment may place storage, policy combination, and effective policy derivation behind that service. When the PPD service endpoint and policy authority are distinct, the deployment needs to preserve at least:

* authenticity of the effective policy instance presented to the participant;

* integrity of policy-instance identifiers and association-freshness metadata;

* an unambiguous binding between the selected PPD service endpoint and the policy authority on whose behalf it presents policy.

These are architectural invariants. The specific transport, metadata confirmation, and cryptographic mechanisms are left to future protocol specifications.

7.4. Key Components

User Interface: A user-friendly interface (e.g., mobile app, web portal) for creating and managing privacy preferences.

PPD Service Endpoint: A participant-facing service through which PPD participants discover, retrieve, and acknowledge applicable policy instances. In a common home deployment model, this service is hosted by a residential gateway or equivalent home-network service. A participant may learn candidate PPD service endpoints through configuration or local network discovery, but it treats a selected endpoint as authoritative only after the applicable trust mechanism succeeds.

Policy Authority: The authoritative source of household policy state and any inputs used for effective policy derivation. The policy authority may be local or remote. A PPD service endpoint can obtain policy from a policy authority without exposing internal storage or computation topology to participants. Participants are not required to discover or communicate with the policy authority directly in the baseline architecture.

Effective Policy Derivation: The logical function, performed by or on behalf of the policy authority, that determines the applicable policy instance for a participant.

Participant Declarations and Consent Requests: Optional participant inputs that can disclose data-handling declarations or request consent for uses not covered by baseline policy. These inputs are distinct from the minimal path of policy retrieval and policy acknowledgment.

Recordkeeping and Management Mechanisms: Deployment-specific mechanisms for presenting association state, participant status, effective policy views, and network-observed devices to the household. Such mechanisms are not device-behavior requirements in the baseline PPD architecture.

7.5. Data Flows

This section outlines the high-level data interactions between users, PPD participants, the PPD service endpoint, and the policy authority in the Privacy Preference Declaration (PPD) framework. It describes how privacy preferences are defined by users, made available to participants, and used as the basis for signaling and recordkeeping in a home network environment.

The process begins when a user defines a set of privacy preferences that apply to their household. These preferences may express rules such as which types of data may be collected, under what conditions data may be processed or shared, or which retention practices are acceptable. The design of the user interface used to author these preferences, including its presentation, usability, or input

modalities, is out of scope for this document, and will be addressed separately. Likewise, the underlying vocabulary and structure of the privacy preferences, including data categories and associated constraints, is specified in (Privacy Preference Declaration Taxonomy).

Once created, the user’s preferences are maintained by a policy authority, which may reside locally on a networked controller or be accessible through other trusted infrastructure. The policy authority may include storage, effective policy derivation, or both. When a new device joins the home network, it initiates an onboarding process during which it obtains one or more candidate PPD service endpoints through configuration or local network mechanisms. Discovery identifies reachable candidates, but does not by itself establish that any candidate is authoritative for household policy. The participant then establishes a secure channel to a selected endpoint and authenticates that endpoint according to the applicable protocol profile before retrieving policy. Following onboarding, the PPD participant performs association, which involves retrieving the household privacy policy and acknowledging receipt of the applicable policy instance. In some deployments, the participant is a backend service associated with the device rather than the local device itself. The PPD service endpoint may present policy derived from a local or remote policy authority without exposing that internal topology to the participant. The participant-facing contract ends at the PPD service endpoint; any split between that service and the policy authority is internal to the deployment. Where those components are distinct, the deployment preserves the authenticity and integrity of the effective policy instance, policy-instance identifier, and freshness metadata presented through the service endpoint. Devices may optionally report their data handling declarations to the PPD service endpoint at this stage. The PPD service endpoint also determines a freshness interval or renewal deadline for the resulting association state.

If a device seeks to perform actions not permitted under the baseline policy, for example, collecting or sharing data beyond what the user has authorized, it may initiate a consent request workflow. However, the design and behavior of this consent mechanism is explicitly out of scope for this document. Inappropriate or poorly designed consent flows, such as those that involve excessive prompting, ambiguous language, or misleading options, can inadvertently pressure users into accepting data practices that conflict with their preferences. Even without malicious intent, these experiences may degrade trust and lead to outcomes inconsistent with user expectations. Future specifications should describe consent interactions that are clear, proportionate, and respectful, helping users make informed decisions without friction or fatigue.

Current association is not indefinite. If the participant does not renew before the freshness interval expires, the PPD service endpoint treats the association as stale even if the applicable effective policy instance is unchanged. Reassociation is required when current association can no longer be confirmed for a participant. This can occur because the applicable effective policy changed, because participant state relevant to effective policy derivation changed, because the association became stale, or because enough state was lost that the prior association can no longer be trusted. Reassociation re-establishes current association by retrieving and acknowledging the latest applicable effective policy instance, or by completing a renewal procedure defined by the applicable protocol profile when the policy instance is unchanged. Devices are not expected to re-collect consent for data uses already covered by existing, valid consent.

To support straightforward implementation and debugging, future protocol work should define a simple machine-readable representation for privacy policies, declarations, and acknowledgment state. JSON is a practical baseline encoding for the architecture and for many constrained home-network deployments. More compact encodings can be considered later if a specific deployment profile demonstrates a need for them.

Future protocol specifications also need to define how association freshness is conveyed, including whether the protocol uses bounded renewal intervals, explicit renewal deadlines, or equivalent lease-style semantics. Those specifications need to distinguish stale association from needs-reassociation states caused by policy or participant-state changes.

It is important to note that the baseline requirement under this architecture is limited to discovery, retrieval, acknowledgment, and any renewal needed to maintain current association for the user’s privacy policy. These actions provide a signaling and recordkeeping mechanism for establishing that the current applicable policy was made available to the participant. However, this document does not define how device behavior is changed by the policy, nor does it specify how to handle cases where a device cannot fully satisfy a given policy. These aspects, including optional status reporting, conflict resolution, or auditing, may be addressed in future work.

Finally, while this document defines the overall data flow and interaction sequence, it does not define message formats, communication protocol details, or consent interface specifications. These elements will be specified in a companion document, (Privacy Preference Declaration Protocol Specification).

7.6. Non-PPD and Network-Observed Devices

Home networks commonly include devices that do not implement PPD, cannot be updated to implement PPD, or are visible only through local network observation. The architecture treats these devices as expected operational cases rather than exceptional failures.

A local management function can classify such devices as network-observed or unmanaged based on information available within the home network. That classification can improve household transparency by showing that a device is present even though it has not established association through PPD. Network observation does not create association, does not imply that the device has received a household policy, and does not imply anything about the device’s behavior.

Any local response to unmanaged devices, such as notification, inventory display, or other network management action, is a deployment decision outside the baseline PPD signaling architecture.

8. Policy Language

The specific details of the Privacy Policy Language, including its syntax, structure, and extensibility mechanisms, are considered out of scope for this document, which focuses on the overall framework. The Privacy Policy Language, along with a taxonomy of privacy concepts and attributes, is expected to be defined in a separate document, the "Privacy Preference Declaration Taxonomy", allowing for more detailed exploration and development of this component of the PPD framework.

8.1. Language Requirements

* Human-readable: Policies should be easily understandable by users.

* Machine-readable: Policies should be machine-processable for automated interpretation and signaling.

* Extensible: The language should be flexible enough to accommodate evolving privacy needs and technologies.

* Internationalization-compatible: Policies and identifiers used within them may need to support multilingual environments and non-ASCII characters.

To ensure consistent interpretation and comparison of string-based policy elements, such as device names, labels, or category identifier string handling practices should align with the guidelines defined in [RFC7564]. This is particularly important when identifiers or user-facing labels are created, stored, or matched across vendors or systems that operate in different locales or character encodings.

8.2. Proposed Approach

Consider leveraging existing privacy policy languages (e.g., P3P) or drawing lessons from privacy labeling systems used in modern application ecosystems--such as Apple’s App Privacy Labels and Google’s Data Safety section for Android apps. While these approaches are not protocols per se, they demonstrate how structured, declarative privacy metadata can be communicated to users and systems in a consistent way.

Alternatively, a new, concise, and user-friendly privacy policy language may be developed specifically for the PPD framework. One possibility is to define an intermediate representation--similar in spirit to the intermediate representation used in compilers such as LLVM--that captures the fundamental privacy constraints and regulatory considerations (e.g., from GDPR, CCPA) in a machine-readable form. This representation would support automated interpretation while being straightforward to translate into human-readable language.

Future specifications should also define guidance for how string identifiers--such as device roles, policy tags, or consent status labels--are formatted, compared, and stored, to avoid ambiguities across systems. In contexts where internationalized strings are involved, alignment with [RFC7564] should be considered to ensure interoperability and consistency.

9. Future Work

This document defines an architectural framework for enabling users to express privacy preferences and signal those preferences within home network environments. Several aspects critical to a fully operational implementation are intentionally left out of scope here and are expected to be addressed in future specifications or companion documents.

9.1. Policy Taxonomy and Semantics

A separate document, tentatively titled "Privacy Preference Declaration Taxonomy", will define:

* A common vocabulary and set of categories for expressing privacy preferences.

* Attributes and semantics for data types, sharing constraints, and processing conditions.

* An extensibility model for incorporating future data types and policy dimensions.

This taxonomy is foundational for consistent policy interpretation across heterogeneous devices and vendors.

9.2. Protocol Specification and Message Formats

A companion document, "Privacy Preference Declaration Protocol Specification", is expected to define:

* Message formats for device onboarding, policy retrieval, policy acknowledgment, participant resynchronization, and optional participant status reporting.

* Optional mechanisms for consent request flows.

* Discovery profiles, lightweight metadata confirmation, and trust-establishment bindings for PPD service endpoints.

* Transport-layer considerations, service authentication, and policy-authority trust expectations.

* Association freshness semantics, including how renewal intervals or deadlines are conveyed and how stale association is distinguished from needs-reassociation states.

* Baseline encoding expectations for structured data, with JSON as a practical starting point and more compact encodings reserved for deployment profiles that need them.

9.3. Consent Request Workflow Design Specifications

The mechanism by which devices request additional user consent for data uses not covered by the baseline policy is out of scope. However, future specifications should:

* Define clear constraints to prevent manipulative or fatiguing consent flows (e.g., dark patterns).

* Describe consent interactions that are transparent, infrequent, proportionate, and user-respecting.

* Explore user interface standards or API affordances to preserve meaningful choice.

This is a particularly sensitive area and must balance user experience, privacy expectations, and implementation feasibility.

9.4. Recordkeeping and Local Management

This architecture does not define how devices act on privacy policies or how departures from policy are detected or remediated. The baseline function is signaling: a participant can receive an applicable household policy and acknowledge that policy instance. Future work may include:

* Optional participant status reporting models and device-side implementation expectations.

* Recordkeeping mechanisms for correlating policy delivery and acknowledgment records.

* State models that distinguish current, stale, and needs-reassociation participant status.

* Deconfliction strategies for devices unable to meet all user-defined constraints.

* Deployment-local management options, such as notifications or inventory display.

9.5. User Interface Design Specifications

The user-facing interface used to author, modify, and review privacy preferences is out of scope. Future design guidance may address:

* User experience design principles for presenting privacy concepts clearly and accessibly.

* Models for progressive disclosure of policy impact.

* Multi-user and household-role-specific control models (e.g., parental vs. administrative roles).

9.6. Interoperability Testing and Reference Implementations

Future work may also include:

* Development of reference implementations of the PPD protocol, PPD service endpoint, and policy-authority components.

* Interoperability testing across devices and vendors.

* Conformance guidelines and self-certification procedures.

10. IANA Considerations

This document has no IANA actions.

11. Internationalization Considerations

In contexts where privacy preferences or taxonomy elements involve user-facing or vendor-defined string identifiers, additional work may be required to:

* Define string normalization and comparison rules, particularly for internationalized text.

* Support identifier consistency across diverse vendors and locales.

* Consider alignment with [RFC7564] for handling Unicode-aware identifiers in a secure and interoperable way.

12. Security Considerations

For a privacy framework to be effective, it needs to support the expression of user preferences and protect those preferences during transmission, retrieval, and acknowledgment. This section outlines safeguards for confidentiality, authenticity, integrity, and metadata minimization during PPD operations.

12.1. Secure Policy Dissemination

Communication between PPD participants and the PPD service endpoint needs protection against unauthorized access and tampering. When the PPD service endpoint and policy authority are distinct, deployments also need to preserve policy authenticity and integrity across that boundary. Discovery mechanisms can identify candidate PPD service endpoints, but discovery alone is not sufficient to establish that an endpoint is authorized to present household policy. Future protocol specifications need to identify appropriate cryptographic mechanisms, such as encryption and mutual authentication, so that legitimate participants can retrieve privacy policies and detect modification. Those specifications also need to protect the binding between the authenticated participant-facing service endpoint and the policy state it presents.

12.2. Anonymity and Metadata Protection

Even when privacy policies themselves do not contain sensitive personal information, the act of retrieving or acknowledging a policy can reveal characteristics about the household, such as the types of devices in use, specific user preferences, or behavioral patterns over time. [RFC7258] cautions against protocol designs that expose unnecessary metadata, treating the accumulation of such information as a legitimate technical threat. This framework takes that warning seriously: metadata exposure during policy retrieval and device onboarding needs to be minimized to avoid turning privacy infrastructure into a new source of privacy leakage. Concepts from [RFC9577] may help inform this effort. [RFC9577] introduces techniques for authorization without identification, enabling a client to prove it is authorized without revealing who it is. While [RFC9577] is optimized for pseudonymous web authentication over the public internet and assumes a centralized token issuer model, its core ideas, particularly around unlinkable token presentation, could be adapted to the PPD protocol to reduce metadata correlation and minimize household identifiability during policy exchanges. However, this needs careful analysis, as the assumptions of [RFC9577] do not fully align with the goals or context of a local, user-governed home network.

12.3. Policy Integrity

Devices need assurance that the policy retrieved is authentic and unaltered. Integrity protections, such as digital signatures, are necessary to ensure that users' preferences cannot be tampered with in transit or at rest by other devices, malicious actors, or misconfigurations. If policy is obtained through a participant-facing service from a distinct policy authority, integrity protections also need to cover the policy-instance identifier and any freshness metadata presented through that service.

12.4. Device Authentication

Devices participating in the privacy framework need an authentication model before accessing the PPD service endpoint. This limits policy dissemination to known, authorized participants and helps users maintain trust in the integrity of their home network's privacy relationships. If the PPD service endpoint and policy authority are distinct, the deployment also needs a way to preserve the authenticity of policy state presented through the participant-facing service. By aligning with the concerns raised in [RFC7258] and incorporating ideas from [RFC9577] where appropriate, this framework seeks to protect users not only from overt data collection, but also from silent inference and passive metadata surveillance. At the same

time, it avoids treating anonymity as an end in itself. The goal is to support privacy with recordkeeping, where user-defined preferences are signaled consistently, devices are identifiable only as much as necessary for the exchange, and the user retains visibility into what occurs within their domain.

12.5. Policy Acknowledgment and Recordkeeping

PPD participants need a way to acknowledge receipt of the applicable privacy policy instance. This acknowledgment should be recorded and verifiable so that the household can determine which participants have seen the current policy. The record needs to bind the participant identity, the acknowledged policy instance, and the time or sequence context in which the acknowledgment was made. For devices that rely on a backend service, the record also needs to distinguish between acknowledgment by the local device and acknowledgment by the backend service acting on behalf of that device. This record is important because it creates a reciprocal acknowledgment path. In many current deployments, the household user is asked to acknowledge device or vendor policy terms, but there is no comparably strong household-controlled record that the participant was presented with the household's own privacy policy. An authenticated and integrity-protected acknowledgment record allows the household to show that presentation and acknowledgment occurred, which can support later accountability or review even when the architecture does not define automated enforcement. Future protocol specifications need to define how acknowledgments are protected against forgery, replay, and stale-policy confusion while still being practical for constrained home-network devices. At minimum, the selected mechanism needs to provide:

* participant authentication sufficient to bind the acknowledgment to the device or backend service that made it;

* policy-instance integrity so that the acknowledged policy can be identified unambiguously;

* freshness or sequencing so that an old acknowledgment cannot be replayed as evidence of current association;

* verifiability sufficient for the acknowledgment record to function as a protected receipt of policy presentation and acknowledgment; and

* a way to retain or export the acknowledgment record without exposing more household metadata than necessary.

A policy acknowledgment is not, by itself, an assertion about subsequent device behavior. Any local response to non-participation or other local observations is outside the baseline signaling mechanism defined by this architecture.

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/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>.

13.2. Informative References

[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>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/rfc/rfc7258>.

[RFC7564] Saint-Andre, P. and M. Blanchet, "PRECIS Framework: Preparation, Enforcement, and Comparison of Internationalized Strings in Application Protocols", RFC 7564, DOI 10.17487/RFC7564, May 2015, <https://www.rfc-editor.org/rfc/rfc7564>.

[RFC8280] ten Oever, N. and C. Cath, "Research into Human Rights Protocol Considerations", RFC 8280, DOI 10.17487/RFC8280, October 2017, <https://www.rfc-editor.org/rfc/rfc8280>.

[RFC9577] Pauly, T., Valdez, S., and C. A. Wood, "The Privacy Pass HTTP Authentication Scheme", RFC 9577, DOI 10.17487/RFC9577, June 2024, <https://www.rfc-editor.org/rfc/rfc9577>.

Authors' Addresses

Daniel Smullen
CableLabs
Email: d.smullen@cablelabs.com

Brian Scriber
CableLabs
Email: brian.scriber@computer.org