INTERNATIONAL TELECOMMUNICATION UNION
TELECOMMUNICATION STANDARDIZATION SECTOR
STUDY PERIOD 2025-2028

SG17-Cn
STUDY GROUP 17
Original: English

Question(s): 1/17 Geneva, 30 March – 2 April 2026

CONTRIBUTION

Source: Broadcom Europe Ltd., Vodafone Group

Title: Proposed first revision of the baseline text of X.PARCEP: Interoperable Parental Control Enforcement Principles (PARCEP) for Child Online Protection

Contact: Gianpaolo Angelo Scalone
Vodafone Group
United Kingdom
Tel: +393480001806
E-mail: gianpaolo-angelo.scalone@vodafone.com

Contact: Arnaud Taddei
Broadcom Europe Ltd.
United Kingdom
Tel: +41795061120
E-mail: Arnaud.Taddei@broadcom.com

Abstract: This document provides the proposed first revision of the baseline text of X.PARCEP: Interoperable Parental Control Enforcement Principles (PARCEP) for Child Online Protection

Editor's notes:

- The feedback received during SG17 2nd plenary meeting 3-11 of December was reviewed and processed, leaving 4 out of the 9 actions to followup on.

Notes based on the discussions during Q1 incubation meeting on C399:

#	Remark/Discussion	Proposed fix or comment
~~1~~	~~There are inconsistencies in the title of this Recommendation at different locations in the text, e.g. the acronym PARCEP and "for Child Online Protection" were missed~~	~~Fixed in this text as per revision mark~~
9	During WP Closing, the title of the new work item was discussed and the term Policies is problematic and after debate it was proposed to change it with Principles	The term Enforcement Policies is a technical term (see point 2 above) but in the title there is a risk that it is interpreted as an international Policy term. Several options were tested: - A) Add Principles after Policies - B) Replace Policies by Principle - C) Put Guidelines - D) Replace Policies by Configuration - Other alternatives o Add 'technical' in front of Policies

		It was agreed to go B) but with the remark that this Work Item will revise its title during its development.
2	The term enforcement is not used consistently	The term enforcement refers to the term that lead to "policy enforcement point" which is the technical term to represent the architectural component in the system that enforces the policy to be effective in the system. ~~There were a few locations in the text where this term was missing.~~ The term enforcement could be either referred from other standards or redefined here if needed. After searching in the ITU terms and definitions database, the term enforcement policy is used in other networking contexts (DPI, etc.) but not with the right context for this Recommendation. ~~For future development of this draft Recommendation we note these existing definitions:~~ The IETF defined terminology in the context of networking deployments as: **RFC 3198 – *Terminology for Policy-Based Management* (Nov. 2001)** defines both *policy enforcement* and *Policy Enforcement Point (PEP)*: • **policy enforcement** – "The execution of a policy decision." IETF Datatracker • **Policy Enforcement Point (PEP)** – "A logical entity that enforces policy decisions [RFC2753]." IETF Datatracker • policy decision (P) Two perspectives of "policy decision" exist: - A "process" perspective that deals with the evaluation of a policy rule's conditions - A "result" perspective that deals with the actions for enforcement, when the conditions of a policy rule are TRUE • Policy Decision Point (PDP) (P) A logical entity that makes policy decisions for itself or for other network elements that request such decisions [RFC2753]. (See also "policy decision".)

Formatted: Font: (Default) Times New Roman

Formatted: Bulleted + Level: 1 + Aligned at: 0.63 cm + Indent at: 1.27 cm

Formatted: Font: (Default) Times New Roman

Formatted: Font: (Default) Times New Roman

Formatted: Bulleted + Level: 1 + Aligned at: 0.63 cm + Indent at: 1.27 cm

Formatted: Font: (Default) Times New Roman

		ISO as well has definitions such as: **ISO/IEC 29146:2024 – *Information technology — Security techniques — A framework for access management*** gives a precise definition of **policy enforcement point**: • **policy enforcement point (PEP)** (3.11) – "service that enforces the access decision by the policy decision point (3.10)" Iteh Standards Notes further clarify that: o The PEP *receives authorization decisions made by the PDP and implements them in order to control access by entities to resources*. Iteh Standards o It corresponds to the *access enforcement function (AEF)* in ISO/IEC 10181-3. Iteh Standards **OMA-AD-Policy_Evaluation_Enforcement_Management V1.0 (PEEM)**, the glossary includes a definition of **policy enforcement** itself: • **Policy Enforcement** – "The process of executing actions, which may be performed as a consequence of the output of the policy evaluation process or during the policy evaluation process." openmobilealliance.org In the same document, *Policy Enforcement Point (PEP)* is defined via abbreviation, and the architecture follows the classic IETF PDP/PEP model and explicitly references **RFC 3198** and **RFC 2753**. openmobilealliance.org This gives you a more operational/process-oriented definition of *enforcement* that is still grounded in standards work.

Formatted: Indent: Left: 0 cm

		The **NIST CSRC glossary**, referencing **CNSSI-4009-2015** and various NIST SPs (800-162, 800-207, etc.), aggregates several normative definitions of **policy enforcement point (PEP)**, for example: • "A system entity that requests and subsequently enforces authorization decisions." csrc.nist.gov • "Enforces policy decisions in response to a request from a subject requesting access to a protected object; the access control decisions are made by the policy decision point." csrc.nist.gov
		MEF (Metro Ethernet Forum – industry SDO-like forum) **MEF 118 – *Zero Trust Framework for MEF Services* (2022)** includes a terminology table: • **Policy Enforcement Point (PEP)** – "An entity that implements Policy decisions that were made by the PDP." Editors considerations: - Many such definitions are related to specific network deployment contexts - IETF definition relates to an ambiguity unresolved in the term 'policy decision' and even if many text refer to this RFC, this is an informational RFC, so, non normative in nature - ISO will force to a paywall - OMA is significantly mobile biased - MEF is significantly Ethernet biased - And NIST definition is ZT biased So the editors propose that X.Parcep creates its own definition even if inspired by some of the above
~~3~~	~~There is a question about "across operating systems" in the scope of the A.1 on the basis that if this Recommendation tries to go down to~~	~~This is recognized that this Recommendation cannot list each specification for each operating system and other vendors but should make it a general requirement towards all targeted devices.~~

Formatted: Font: (Default) Times New Roman
Formatted: Font: (Default) Times New Roman
Formatted: Font: (Default) Times New Roman
Formatted: Font: (Default) Times New Roman
Formatted: Font: (Default) Times New Roman
Formatted: Font: (Default) Times New Roman

	~~each specific operating systems, this sounds like a too ambitious task~~	
4	There are a number of trademarks in the initial document and it is proposed to keep the document out of trademarks and more general	This was fixed and further considerations will be discussed on **how to make it more generalised**
5	There is a remark on how this capability that parents could have with this Recommendation be transferred to other parties	This is a good suggestion that will be developed in the future of this document. This is a delegation of administration feature that is likely required.
~~6~~	~~There is demand for a workshop in one of the Content Week on the matter including all stakeholders (operating systems vendors, including CPE vendors, etc.)~~	~~This is a good suggestion to be brought in the report of Q1 and then in the report of WP3 for discussion at SG17 level~~
~~7~~	~~There is a question regarding why ITU-T is in the Relations in A.1 table~~	~~This was indeed removed~~
~~8~~	~~There are editorial nits~~	~~A number of editorial fixes were applied~~
~~9~~	~~During WP Closing, the title of the new work item was discussed and the term Policies is problematic and after debate it was proposed to change it with Principles~~	~~The term Enforcement Policies is a technical term (see point 2 above) but in the title there is a risk that it is interpreted as an international Policy term. Several options were tested:~~ ~~- A) Add Principles after Policies~~ ~~- B) Replace Policies by Principle~~ ~~- C) Put Guidelines~~ ~~- D) Replace Policies by Configuration~~ ~~It was agreed to go B) but with the remark that this Work Item will revise its title during its development.~~

Formatted: Font: Bold

Formatted

Draft new Recommendation ITU-T X.PARCEP

Interoperable PARental Control Enforcement Principles (PARCEP) for Child Online Protection

Summary

~~<Mandatory>~~Recommendation ITU-T X.PARCEP specifies an interoperable framework for Child Online Protection (COP) covering policy authoring, distribution, and enforcement across heterogeneous Policy Enforcement Points (PEPs) deployed on devices, applications, set-top-boxes, gaming consoles, home gateways, access networks, and online services. The framework defines a vendor-neutral data model expressing the household policy authored by a guardian (or co-guardian set) on behalf of a dependant, a federated trust model in which Policy Information Points (PIPs) issue and attest the credentials, classifier-model packages, content taxonomies, and jurisdictional overrides on which PEPs depend, and the requirements of the synchronization protocol that the IETF is expected to standardize. It defines roles (guardian, dependant, co-guardian), security and privacy requirements, transparency obligations to the dependant, AI-related threat-and-defense capability requirements, and a set of category-specific conformance profiles. It does not require content scanning, the weakening of end-to-end encryption, or the creation of a centralized database of dependants' data.

Keywords

~~<Mandatory>~~Age assurance; Child Online Protection (COP); conformance profile; co-guardian; data model; dependant; end-to-end encryption preservation; federated trust; guardian; household policy; interoperability; on-device classification; parental control; policy administration point (PAP); policy decision point (PDP); policy enforcement point (PEP); policy information point (PIP); policy synchronization; transparency to the dependant.

Introduction

~~<Optional – This clause should appear only if it contains information different from that in Scope and Summary>~~The parental-control market is broad and feature-rich but structurally fragmented: every device operating system, every social platform, every set-top-box vendor, every Internet service provider, and every age-assurance vendor exposes its own data model, its own configuration surface, and its own enforcement primitives. Guardians who want to express a single COP policy across the heterogeneous estate that a contemporary dependant uses must today configure each surface independently, accept feature gaps at every boundary, and rely on network controls that the rise of TLS 1.3 with Encrypted Client Hello (ECH), encrypted DNS, and end-to-end-encrypted messaging has rendered porous. Meanwhile, the threat landscape has materially worsened: industrial-scale generative-AI synthesis of child sexual abuse material (CSAM), AI-driven grooming impersonating peers, deepfake sextortion targeting real children, voice cloning, and recommender-amplified self-harm content are documented in 2025 evidence from the Internet Watch Foundation, the National Center for Missing & Exploited Children, WeProtect Global Alliance, and Thorn.

This Recommendation responds by specifying, for the first time at an international standards level, the architecture, data schema, and synchronization-protocol requirements through which a guardian-authored COP policy can be distributed once and consumed consistently across compliant PEPs at the device, application, set-top-box, gaming, home-gateway, access-network, and online-service layers. The architecture is federated by design (multiple jurisdictional trust anchors; no centralized database of dependants' data) and is built on, and intended to interoperate with, the established work of the IETF (MLS, RFC 9420; MIMI; DNSOP) and the eIDAS 2 / EUDI Wallet age-verification building block of the European Union.

Draft new Recommendation ITU-T X.PARCEP

Interoperable PARental Control Enforcement Principles (PARCEP) for Child Online Protection

1 Scope

This Recommendation specifies a vendor-neutral shared data model and policies for interoperable parental control policy exchange, synchronization, and enforcement requirements across devices, operating systems, applications, 3rd party online safety tools, and online services. It defines roles (guardians, dependants, co-guardians), trust and authorization models, security and privacy requirements, minimal data collection principles, transparency indicators, and a conformance profile enabling cross-vendor interoperability. The Recommendation does not prescribe UI/UX, business models, content taxonomies, or jurisdictional rating schemes and does not require content scanning or backdoors.

The Recommendation also covers integration with network-based enforcement at CPE and ISP level as policy consumers within PARCEP. These apply household rules across all connected devices using privacy-preserving methods (e.g., DNS-bound enforcement, time-based access control) without traffic decryption or content scanning.

This Recommendation extends the prior text with the following additional scope: (a) a federated trust model in which jurisdictional Policy Information Points (PIPs) issue and attest the credentials and signed artifacts (guardianship credentials, age credentials, classifier-model packages, content taxonomies, jurisdictional overrides) on which the Policy Decision Points depend; (b) a normative reference architecture decomposing the system into Policy Administration Point (PAP), Policy Decision Point (PDP), Policy Enforcement Point (PEP), and Policy Information Point (PIP) as defined in IETF RFC 3198 and ISO/IEC 29146; (c) requirements for transparency to the dependant aligned with the obligations of regulators including the European Commission under Article 28(1) of the Digital Services Act, the United Kingdom Information Commissioner's Office Children's Code, and the United Kingdom Office of Communications Highly Effective Age Assurance guidance; (d) requirements for the integration of identity-grade and estimation-based age assurance, including compatibility with the European Digital Identity Wallet age-verification building block; (e) requirements for AI-related threats and defenses, including AI-generated CSAM, AI-driven grooming, deepfake sextortion, and AI voice cloning, with corresponding on-device classifier interfaces; and (f) category-specific conformance profiles for device-OS, third-party device add-on, set-top-box/console/TV, ISP/MNO/managed-DNS, CPE/router, social platform/application, and age-assurance product categories.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

~~[ITU-T X.yyy]~~ ~~Recommendation ITU-T X.yyy (date), *Title*.~~[ITU-T X.1080] Recommendation ITU-T X.1080 (in force), Series on cybersecurity for child online protection.

[ITU-T X.1283] Recommendation ITU-T X.1283, Security guidelines for the management of children's digital identities.

[ISO/IEC 29146] ISO/IEC 29146:2024, Information technology — Security techniques — A framework for access management.

[IETF RFC 3198] IETF RFC 3198 (November 2001), Terminology for Policy-Based Management.

[IETF RFC 9420] IETF RFC 9420 (July 2023), The Messaging Layer Security (MLS) Protocol.

[IETF RFC 9162] IETF RFC 9162 (December 2021), Certificate Transparency Version 2.0.

[IETF RFC 8949] IETF RFC 8949 (December 2020), Concise Binary Object Representation (CBOR).

[IETF RFC 9052] IETF RFC 9052 (August 2022), CBOR Object Signing and Encryption (COSE): Structures and Process.

[IETF RFC 9460] IETF RFC 9460 (November 2023), Service Binding and Parameter Specification via the DNS (SVCB and HTTPS Resource Records).

[W3C VC 2.0] W3C Recommendation, Verifiable Credentials Data Model 2.0.

[ISO/IEC 18013-5] ISO/IEC 18013-5:2021, Personal identification — ISO-compliant driving licence — Part 5: Mobile driving licence (mDL) application.

3 Definitions

<Check in the ITU-T terms and definitions database at www.itu.int/go/terminology-database whether the term has already been defined in another Recommendation. It would be more consistent to refer to such a definition rather than to redefine the term>

3.1 Terms defined elsewhere

<Normally, terms defined elsewhere will simply refer to the defining document. In certain cases, it may be desirable to quote the definition to allow for a stand-alone document. Before defining a new term, verify whether it has already been defined in the official ITU terminology database, at www.itu.i.nt/go/terms. >

This Recommendation uses the following terms defined elsewhere:

~~**3.1.1 <Term 1>** [Reference]: <optional quoted definition>.~~**3.1.1 policy enforcement** [IETF RFC 3198]: "The execution of a policy decision."

~~**3.1.2 <Term 2>** [Reference]: <optional quoted definition>.~~**3.1.2 policy enforcement point (PEP)** [ISO/IEC 29146]: "Service that enforces the access decision by the policy decision point."

3.1.3 policy decision point (PDP) [IETF RFC 3198]: A logical entity that makes policy decisions for itself or for other network elements that request such decisions.

3.1.4 verifiable credential [W3C VC 2.0]: A tamper-evident credential whose authorship can be cryptographically verified.

3.1.5 Messaging Layer Security (MLS) group [IETF RFC 9420]: An asynchronously secured group with forward secrecy and post-compromise security; used in this Recommendation as the canonical transport for policy distribution.

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

~~**3.2.1 <Term 3>**: <definition>.~~**3.2.1 guardian**: A responsible adult acting on behalf of a dependant for the purposes of authoring and administering a household policy in this Recommendation. The term covers parents, legal guardians, kinship caregivers, foster carers, and educators acting in loco parentis.

3.2.2 dependant: A child or adolescent under the age of 18 (or the higher protected age set by the applicable jurisdiction) for whom a guardian authors a household policy under this Recommendation.

3.2.3 co-guardian: A guardian sharing policy-administration rights with at least one other guardian for the same dependant. The co-guardian construct accommodates separated households, kinship and foster carers, schools acting in loco parentis, and other time-bounded delegated administrators.

3.2.4 household policy: The set of rules that a guardian (or a co-guardian set) authors under this Recommendation, expressed in the data schema specified in Annex B, and binding one or more dependants to one or more devices and accounts.

3.2.5 policy administration point (PAP): The guardian-controlled component, implemented as a mobile application, a web dashboard, or a guardian-controlled cloud service, where the household policy is authored, signed, and persisted under the guardian's authority.

3.2.6 policy information point (PIP): A federated trust anchor that issues and attests credentials and signed artifacts on which Policy Decision Points depend, including guardianship credentials, age credentials, classifier-model packages, content taxonomies, and jurisdictional override descriptors. Multiple independent PIPs operate per jurisdictional function.

3.2.7 transparency notice: A visible, age-appropriate indication rendered to the dependant whenever the household policy actively monitors or restricts the dependant; the notice identifies the nature of the restriction, the role enforcing it, and the appeal channel.

3.2.8 conformance profile: A defined subset of the requirements and data schema of this Recommendation, applicable to a specific category of Policy Enforcement Point (see Annex D for the seven category profiles defined by this Recommendation).

3.2.9 age credential: A verifiable credential expressing an age-band attribute and/or the boolean answers to age-threshold predicates (e.g., "holder is at least 13"), issued by a PIP and supporting selective disclosure so that the verifier learns the answer without learning the holder identity.

3.2.10 classifier-model credential: A verifiable credential issued by a PIP binding a signed classifier-model package to a classifier class (e.g., nudity, grooming, deepfake, voice-clone) and to per-demographic performance disclosures.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

~~<abbr><expansion>~~AEAD Authenticated Encryption with Associated Data

CBOR Concise Binary Object Representation [IETF RFC 8949]

COP Child Online Protection

COSE CBOR Object Signing and Encryption [IETF RFC 9052]

CPE Customer Premises Equipment

CSAM Child Sexual Abuse Material

DoH DNS over HTTPS [IETF RFC 8484]

DoT DNS over TLS [IETF RFC 7858]

DSA Digital Services Act (European Union)

ECH Encrypted Client Hello (TLS 1.3)

EUDI European Digital Identity

HEAA Highly Effective Age Assurance (Ofcom)

ISP Internet Service Provider

MDM Mobile Device Management

MIMI More Instant Messaging Interoperability (IETF Working Group)

MLS Messaging Layer Security [IETF RFC 9420]

MNO Mobile Network Operator

PAP Policy Administration Point

PDP Policy Decision Point

PEP Policy Enforcement Point

PIP Policy Information Point

SD-JWT-VC Selective Disclosure JWT Verifiable Credential (IETF)

SVCB Service Binding DNS resource record [IETF RFC 9460]

VC Verifiable Credential [W3C VC 2.0]

~~<Include all abbreviations and acronyms used in this Recommendation>~~

5 Conventions

~~<Mandatory clause. Describe any particular notation, style, presentation, etc. used within the Recommendation, if any. If none, write "None.">~~This Recommendation uses the conformance keywords SHALL, SHOULD, MAY, SHALL NOT, and SHOULD NOT in the sense in which they are used in ITU-T Author's Guide and consistent with the meanings defined in IETF RFC 2119 and RFC 8174. SHALL denotes a mandatory requirement for conformance; SHOULD denotes a recommended requirement that may be omitted if a documented justification is provided; MAY denotes a permitted but optional capability.

Requirements in this Recommendation are identified by a stable identifier of the form CATEGORY-NNN where CATEGORY is one of ARCH (architecture), FR-A through FR-M (functional requirements), AI-T / AI-D / AI-S (AI-related requirements), NFR-S / NFR-P / NFR-Perf / NFR-U / NFR-R / NFR-C (non-functional requirements), or CONF-A through CONF-G (conformance profiles). Identifier stability is preserved across revisions of this Recommendation; deprecated identifiers SHALL NOT be reused.

6 Why Now: Encryption and Fragmentation

End to end encryption is increasingly embedded across consumer messaging, storage and platform services. While essential for privacy and security, it reduces the feasibility of legacy safeguards that relied on platform side visibility or network interception. In parallel, families now manage a larger mix of devices and services than ever before. Together, encryption and fragmentation make it harder to implement consistent household rules, leaving gaps that children can unintentionally or deliberately exploit. In many countries this difficulty is contributing to the growth of smartphone free childhood movements, where the perceived complexity of managing devices drives some parents and schools towards prohibition or removal rather than effective management. This trend underlines that solutions must be made simpler and more accessible if parents are to be empowered rather than overwhelmed.

PARCEP addresses this gap by enabling user side, standards-based coordination between parental controls on devices and services, without weakening encryption or creating backdoors, so that protections remain effective in modern, privacy preserving environments.

Network-based enforcement remains relevant as part of a layered approach. By integrating PARCEP policies with CPE and ISP-level controls using privacy-preserving methods (e.g., DNS-bound enforcement), families gain consistent protection without weakening encryption.

7 Problem Statement

Today's families often rely on multiple, non interoperable parental control systems, including those for devices, applications, games, platforms routers and third-party tools. These tools are essential for managing screen time, content access, and healthy digital use. Yet, because they are proprietary and fragmented, there is no standard way for parents to manage their children's online experiences across the multi device households that are now the norm. This fragmentation reduces parental visibility, forces families to navigate multiple inconsistent systems, and increases the likelihood of gaps that can expose children to harm. For many families, especially those with lower digital literacy, this represents an almost insurmountable challenge.

Fragmentation also affects network-level controls; without integration, CPE and ISP solutions cannot align with device policies, leaving gaps in enforcement.

Technologies such as the Encrypted Client Hello (ECH) extension to TLS 1.3, and the rise of VPN's and other tunnelling protocols has rendered ISP network level controls ineffective and easily circumvented, meaning more control has to be asserted over devices by parents to protect children. In a fragmented ecosystem this is technical, complicated, and far from foolproof. Giving a base level of control back to the network will ensure that a single configuration and enforcement policy provides a comprehensive safety net for any network interactions by those in our care, augmenting existing parental control capabilities.

8 Technical Proposal: What PARCEP Is and What It Does

PARCEP is a proposed interoperability protocol that enables coordination between different parental control tools and platforms. Modelled on successful IETF efforts like MIMI (messaging interoperability) and MLS (secure group messaging), it aims to provide the foundational technical standard for cross vendor family device management. Key features include:

1. Shared Data Model: A vendor agnostic schema for screen time settings, application restrictions, content filters, age ratings, and purchase permissions.

2. Policy Synchronisation: A standard protocol for issuing, updating, and revoking family control rules across all compliant devices and apps,~~.~~ APIs, (Editor's note: investigate is Matter matters?)

2bis Policy enforcement: interoperable cross device time quota management, interoperable cross device runtime aspects, feedback loop (actions, logs, errors, etc.) (Editor's note: investigate is Matter matters?)

3. Role Delegation: Support for multi household families, enabling roles and permissions for parents, carers, and guardians.

4. Privacy and Security by Design: Built in safeguards for data minimisation, consent, and protection of children's personal information.

5. Transparency and Auditability: Interfaces for guardians to review usage data, override requests, and receive alerts, with visible signals on children's devices when restrictions apply.

6. Device Management: leveraging existing device management tools (eg MDM) to lock down key settings on devices (eg the DNS resolver) and to control the ability to load new applications.

PARCEP should also support network-based enforcement by allowing CPE and ISP systems to consume and apply policies using privacy-preserving methods such as DNS-bound enforcement.

9 Scope and Non Goals:

PARCEP defines how parental control systems interoperate; it does not prescribe a single user interface or business model. It explicitly does not require content scanning, weakening of encryption, or any form of universal backdoor. It does not create a centralised database of children's data. The standard is designed for data minimisation, consent-based operation and vendor autonomy. Content taxonomies, rating schemes and filtering lists remain vendor or jurisdiction specific; PARCEP supplies the "plumbing" that allows compliant systems to express and enforce policies consistently across devices and services.

Integration with network-based enforcement is limited to policy consumption and application at CPE/ISP level; it does not involve traffic inspection, content scanning, or centralized data storage.

10 Impact

10.1 Practical Benefits for Families:

From a parent's perspective, the standard would enable a much more coherent and straightforward experience. Instead of having to configure different rules on every device or app their child uses, families would be able to set clear expectations once and have them applied consistently across all compliant systems. For example, a rule that no social media should be accessible after 9 p.m. would automatically apply across a child's phone, tablet, laptop, and games console. Screen time limits, purchase permissions, or content restrictions would no longer require multiple dashboards with conflicting terminology, but would operate in harmony regardless of platform. Parents and carers would also be able to share management roles across households, ensuring continuity when children split time between parents, guardians, or extended family. By simplifying management, PARCEP empowers parents to spend less time navigating settings and more time supporting their child's safe and positive engagement with technology.

Network-based enforcement extends these benefits to all devices on the home network, ensuring consistent application of PARCEP policies even when device-level controls are limited.

10.2 Illustrative scenario

Emma sets weeknight "homework hours" from 18:00 to 20:00 with messaging and social media paused, educational apps allowed, and YouTube permitted for 30 minutes total. Through PARCEP, that single rule is enforced across her son's iPhone, school Chromebook and games console. On alternate weekends, the co-parent can grant a one-off extension from their own dashboard; the child's devices display a clear on screen reason when an app is blocked ("Paused during homework hours") and the same usage log appears in both households.

If Emma's home router enforces PARCEP policies, the same "homework hours" rule applies to specified devices on the network, even those without native parental control apps, ensuring consistent protection.

10.3 Benefits for Device and Service Providers

For device manufacturers and service providers, PARCEP reduces the pressure to build fully comprehensive parental control suites in isolation. Instead, vendors can integrate their existing tools with the shared protocol, ensuring compatibility and enhancing user satisfaction without sacrificing their unique user experience. This interoperability strengthens trust with families, making it more likely that parents will recommend or remain loyal to a device or service that aligns with their wider household controls. Trust is particularly important as parents seek consistent, ubiquitous solutions across multi vendor platforms.

Public debates in many countries around smartphone free childhood movements, already noted earlier in this document, further highlight the difficulty parents face in managing access and screen time, and reinforce the need for simpler, more effective parental empowerment. By making it easier for parents to manage technology confidently, PARCEP delivers a direct advantage to families and an indirect but significant advantage to device and service providers whose offerings are seen as safe, reliable and family friendly. It also positions providers ahead of regulatory trends that increasingly emphasise consistency, transparency, and child protection.

By adopting PARCEP, vendors demonstrate leadership in online safety, contribute to levelling the playing field across the industry, and benefit from reduced reputational risk when safeguarding expectations are met through common standards. Early adopters can differentiate on trust, reduce support burden from complex multi device households, and minimise churn by aligning with families' real world needs.

PARCEP does not replace existing tools but acts as a neutral bridge that allows them to work together. This ensures that protections are consistent regardless of device or platform, that parents can manage rules more easily, and that vendors can maintain their own user experiences while aligning with a shared interoperability framework.

Network-based enforcement offers additional value for ISPs and CPE vendors by enabling PARCEP integration at the network edge, applying parental controls even on legacy devices without requiring apps, updates, or draining battery, while ensuring consistent, privacy-preserving protections.

11 11 Architecture

This clause specifies the architecture of a PARCEP-conformant system. It restates the architectural decision recorded in Annex A in normative form for the main body.

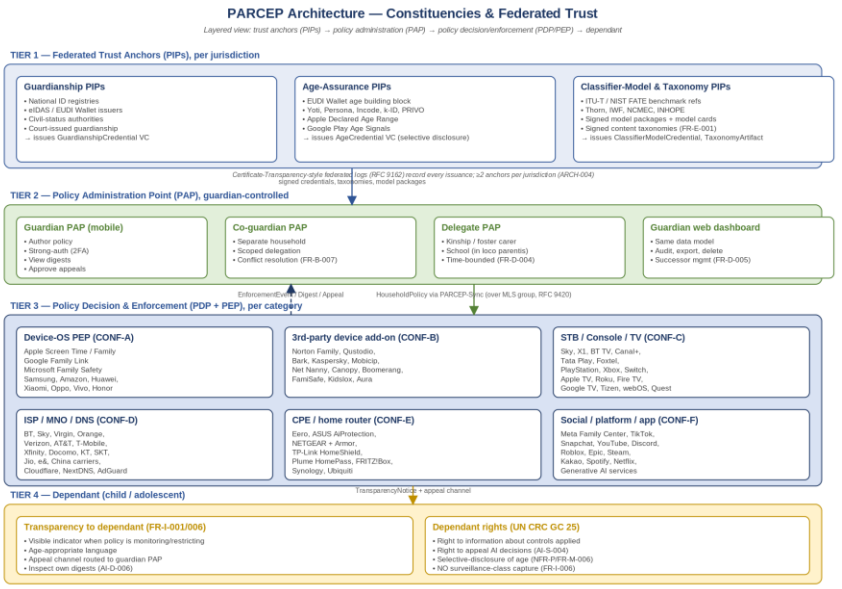

Figure 1 — PARCEP constituencies and federated trust topology (Tier 1 federated PIPs; Tier 2 guardian-controlled PAP; Tier 3 PDP and PEP per category; Tier 4 dependant with transparency notice).

11.1 11.1 Architectural decisions

Three architectural decisions anchor this Recommendation.

ARCH-001 SHALL — A PARCEP-conformant system shall implement the policy administration / decision / enforcement / information decomposition defined in clause 11.2, consistent with IETF RFC 3198 and ISO/IEC 29146 terminology.

ARCH-002 SHALL — A PARCEP-conformant system shall not require, and shall not depend on, any central database in which dependants' personal data, age signals, or household policies are aggregated across households. *Note: Honors the non-goal of this Recommendation.*

ARCH-003 SHALL — The Policy Administration Point shall be guardian-controlled and shall persist the canonical household policy under the guardian's authority. Vendor-side caches and replicas are permitted but shall be subordinate to the guardian PAP.

ARCH-004 SHALL — Policy Information Points shall be federated: at least two independent Policy Information Points shall operate per jurisdictional function so that no single Policy Information Point failure or compromise blocks enforcement.

ARCH-005 SHOULD — Policy Information Points should be operated under transparency-log obligations modeled on IETF RFC 9162 (Certificate Transparency v2), so that the issuance of household-affecting artifacts (e.g., classifier models, taxonomy updates) is publicly auditable. *Note: Defense against silent policy mutation.*

ARCH-006 SHALL — Policy distribution between Policy Administration Point and Policy Decision Point shall use authenticated, integrity-protected channels and shall not require the policy to traverse any third party in cleartext.

ARCH-007 SHALL — PARCEP shall be transport-neutral above the security boundary; the wire protocol shall not mandate a specific underlying transport.

ARCH-008 SHOULD — Policy distribution should be modeled on the IETF MIMI architecture (messaging interoperability) so that a household policy can be expressed once and consumed by multiple, heterogeneous Policy Enforcement Points without per-vendor adaptation.

ARCH-009 SHALL — Network-side enforcement (CPE, ISP, MNO) shall act as a policy consumer of the household policy and shall not require deep packet inspection, traffic decryption, or content scanning to enforce it. *Note: Compatible with TLS 1.3, ECH, QUIC, DoH/DoT.*

ARCH-010 SHOULD — Network-side enforcement should use DNS-bound primitives (encrypted DNS resolver pinning per IETF RFC 9462 and RFC 9463 and DNS structured error data per the IETF DNSOP draft) to communicate enforcement reasons to the dependant device, with no per-user data egress to the network operator.

ARCH-011 SHALL — PARCEP shall support cross-network policy continuity: when the dependant's device changes Wi-Fi or moves to a cellular network, the household policy shall remain enforced by the device Policy Enforcement Point and, where available, by the new network-side Policy Enforcement Point.

ARCH-012 SHOULD — PARCEP should support graceful degradation: if a Policy Enforcement Point cannot reach its Policy Decision Point or Policy Information Point, the most recent valid policy shall continue to apply, with explicit transparency events to the guardian after a configurable grace period.

ARCH-013 MAY — PARCEP may support federated audit logs that record policy-enforcement decisions in a privacy-preserving form for regulatory or research audit, without exposing dependant identity.

11.2 11.2 Functional decomposition (PAP, PDP, PEP, PIP)

The functional decomposition follows the IETF RFC 3198 and ISO/IEC 29146 policy-management terminology. The Policy Administration Point (PAP) is the guardian-controlled authoring surface, implemented as a mobile application, a guardian web dashboard, or a guardian-controlled cloud service. The PAP authenticates the guardian using strong (two-factor) authentication, persists the canonical household policy under the guardian's authority, exposes an age-graded preset library, supports policy versioning and rollback, and accommodates co-guardians, time-bounded delegates, and successor management. The PAP is the source of all PUBLISH, UPDATE, REVOKE, and RESOLVE messages defined in clause 13.

The Policy Decision Point (PDP) is, by default, co-located with the Policy Enforcement Point. It receives the signed household policy, validates the artifacts referenced therein (classifier-model packages, taxonomies, jurisdictional overrides), caches the policy locally for offline operation, and evaluates per-request decisions. For ISP and MNO operators the PDP may be remote relative to the PEP.

The Policy Enforcement Point (PEP) blocks, allows, paces, delays, blurs, or warns. It is implemented per category in accordance with clause 18 (conformance profiles) and decomposes into five reusable sub-modules: content and conduct enforcement (including on-device AI classifiers), the time manager, the contact and communication guard, the commerce approver, and the transparency and telemetry emitter. The PEP is the source of all EVENT messages and the entry point for APPEAL messages from the dependant.

The Policy Information Points (PIPs) are federated trust anchors operating per jurisdiction. They issue and attest credentials and signed artifacts on which Policy Decision Points depend: guardianship credentials, age credentials, classifier-model packages, content taxonomies, jurisdictional override descriptors. Every issuance is recorded in a public transparency log per IETF RFC 9162.

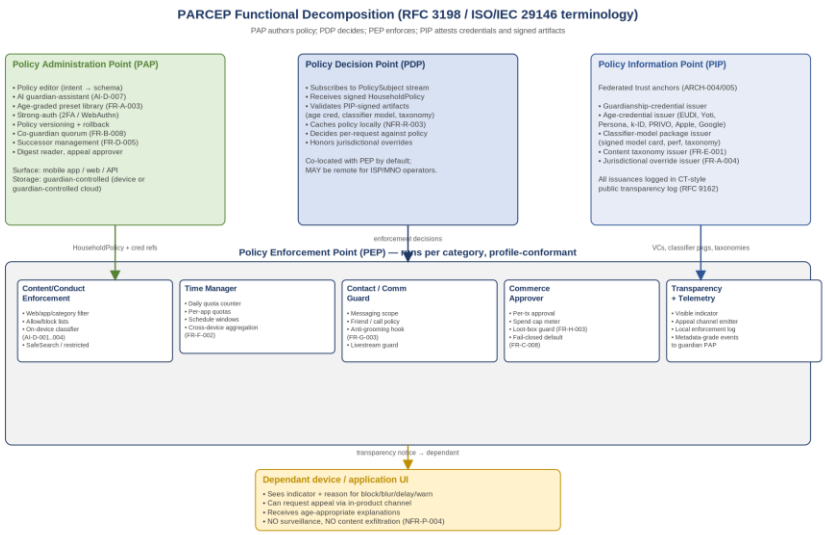

Figure 2 — Functional decomposition (PAP / PDP / PEP / PIP) per IETF RFC 3198 and ISO/IEC 29146.

11.3 11.3 Federated trust model

The trust model follows Internet-PKI conventions extended for selective-disclosure verifiable credentials. The root trust set is, per jurisdiction, a small federation of Policy Information Points. The Policy Information Points issue verifiable credentials in W3C VC 2.0 and SD-JWT-VC format and, where applicable, ISO/IEC 18013-5 mDL format, that Policy Enforcement Points verify offline against pinned Policy Information Point roots.

Three credential chains converge at the Policy Enforcement Point: (a) the guardianship chain (Policy Information Point to GuardianshipCredential to guardian signature on the HouseholdPolicy); (b) the age chain (Policy Information Point to AgeCredential to dependant's selective-disclosure presentation at the PEP, for example 'holder is at least 13 in jurisdiction X'); (c) the classifier-model chain (Policy Information Point to ClassifierModelCredential to model package, signature verified at load time).

Every credential shall carry a revocation endpoint. Two revocation mechanisms are supported: short-lived credential validity (where the credential is re-issued frequently and revocation is implicit by non-renewal); and status-list endpoints (W3C VC Status List 2021 or IETF Token Status List). PEPs shall check revocation at policy load time and at a policy-defined refresh interval.

Cross-jurisdictional verification: a Policy Enforcement Point operating in jurisdiction A may receive a credential issued in jurisdiction B. The PEP shall accept a credential from any Policy Information Point whose root is on a PEP-recognized trust list. The trust list is, by default, the union of the PEP's home jurisdiction trust list and any additional trust lists explicitly imported by the guardian via the Policy Administration Point.

12 12 Data schema

The data schema is split into five layers: identity, credentials, policy, operational, and discovery. The schema is expressed in a vendor-neutral abstract syntax; concrete encodings shall include both JSON-LD compatible with W3C VC 2.0 for human-readable and web-compatible exchanges, and CBOR with COSE signatures (IETF RFC 8949 / 9052) for compact, embedded-friendly exchanges. The two encodings shall produce semantically equivalent decisions for the same policy.

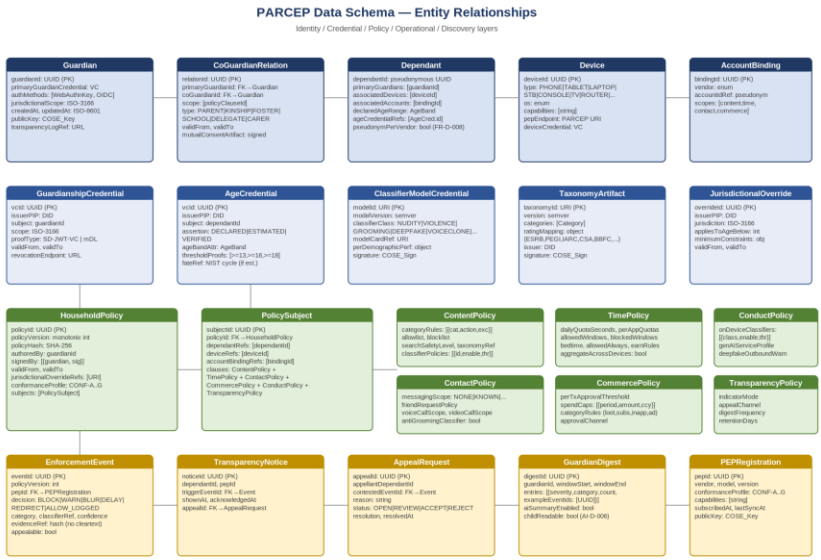

Figure 3 — Data schema entity-relationship diagram (identity, credentials, policy, operational, discovery layers).

12.1 12.1 Identity layer

The identity layer defines: Guardian (with guardianId, primary GuardianshipCredential, authentication methods, jurisdictional scope, public key in COSE_Key format, transparency-log reference, creation and update timestamps); CoGuardianRelation (with relationId, primary and co-guardian identifiers, scope expressed as a subset of policy-clause identifiers, type drawn from PARENT, KINSHIP, FOSTER, SCHOOL, DELEGATE, CARER, validity period, mutual-consent COSE_Sign1 artifact); Dependant (with pseudonymous dependantId, associated guardians, associated devices, associated account bindings, declared age range, age-credential references, per-vendor pseudonym flag); Device (with deviceId, type drawn from PHONE, TABLET, LAPTOP, DESKTOP, STB, CONSOLE, TV, ROUTER, SMART_SPEAKER, XR_HEADSET, OTHER, operating-system enumeration, capability tags, Policy Enforcement Point endpoint URI, optional DeviceCredential); AccountBinding (with bindingId, vendor, per-vendor account pseudonym, scope drawn from CONTENT, TIME, CONTACT, COMMERCE, CONDUCT).

12.2 12.2 Credentials layer

The credentials layer defines: GuardianshipCredential (issued by a Policy Information Point, with subject identifying the guardian, jurisdictional scope, proof format drawn from SD-JWT-VC, ISO/IEC 18013-5 mDL, or JSON-LD VC, validity period, revocation endpoint); AgeCredential (selective-disclosure capable, with assertion class drawn from DECLARED, ESTIMATED, VERIFIED, age-band attribute, threshold proofs for at-least-13, at-least-16, at-least-18, NIST FATE benchmark reference when ESTIMATED, Ofcom HEAA conformance flag when VERIFIED, validity period, revocation endpoint); ClassifierModelCredential (binding a signed model package identified by modelId URI and SemVer version to a classifier class drawn from NUDITY, VIOLENCE, SELF_HARM, HATE, GROOMING, DEEPFAKE, VOICECLONE, MALWARE, OTHER, with model-card reference, per-demographic performance object including skin tone, age band, and gender breakdowns, adversarial-robustness disclosure, and COSE_Sign1 signature); TaxonomyArtifact (signed content-category taxonomy with cross-regional rating-scheme mapping

covering ESRB, PEGI, IARC, CSA / Arcom, BBFC, CERO, USK, ACB, and equivalents); JurisdictionalOverride (regulator-issued minimum policy constraints applicable to dependants below a defined age in a given ISO-3166-1-alpha-2 jurisdiction, covering time, content, contact, and commerce constraints).

12.3 12.3 Policy layer

The policy layer defines the HouseholdPolicy as the signed root object the Policy Administration Point publishes. The HouseholdPolicy carries: policyId (immutable UUID), policyVersion (monotonic integer per policyId), policyHash (SHA-256), authoring-guardian reference, signedBy array including a guardian identifier and a COSE_Sign1 signature for each co-guardian endorsement, validity period, jurisdictional-override references, conformance-profile selector drawn from CONF-A through CONF-G defined in clause 18, subjects array, and a namespaced extensions array. Each PolicySubject binds a clause-set to a set of dependant, device, and account-binding identifiers.

Each PolicySubject contains six policy clauses. The ContentPolicy carries a taxonomy reference, an array of categoryRules each comprising category, action drawn from ALLOW, BLOCK, WARN, and an exceptions list, an allowlist and a blocklist of URI patterns, a search-safety level drawn from OFF, MODERATE, STRICT, an array of classifier-policy bindings each carrying modelId, enable flag, and threshold, and a region-specific rating-scheme selector.

The TimePolicy carries a daily quota in seconds, per-app or per-category quotas, allowed windows expressed as days-of-week and time-of-day ranges with time-zone, blocked windows, a bedtime window, an allowed-always exception list, earn rules each comprising a trigger and bonus-seconds, and an aggregate-across-devices flag.

The ContactPolicy carries a messaging scope drawn from NONE, KNOWN, MUTUAL, APPROVED, ALL, a friend-request policy drawn from BLOCK, APPROVE, ALLOW, voice and video call scopes, livestream broadcast and receive flags, an anti-grooming-classifier enable flag, and a blocked-identities list.

The CommercePolicy carries a per-transaction approval threshold (amount and currency, or 'ALWAYS'), an array of spend caps each over a period, an array of category rules covering in-app purchases, subscriptions, loot boxes, advertising clicks, and real-money gambling, and an approval channel drawn from PAP push notification, email, or both.

The ConductPolicy carries an array of on-device classifier bindings, a Generative-AI Service Profile defined in clause 12.4, a deepfake-outbound-warning flag, and a voice-clone-inbound-warning flag.

The TransparencyPolicy carries an indicator-mode selector drawn from PROMINENT and DISCREET, an appeal-channel URI or enumeration, a digest-frequency selector drawn from NONE, DAILY, WEEKLY, a digest-detail selector drawn from CATEGORY_ONLY and INSTANCE_REFS, a child-readable-digest flag, a retention period in days, and an age-graded-vocabulary locale tag.

12.4 12.4 Generative-AI Service Profile

The Generative-AI Service Profile, a sub-object of the ConductPolicy, defines the minimum control surface that any generative-AI service marketed to or accessible by a dependant shall expose to be referenceable by a household policy. The profile carries: consent-required flag; content-category toggles for sexual content, violence, self-harm content, weapons content, and romantic role-play; distress-signal escalation to the guardian; quiet-hours windows; time bounds expressed as a maximum daily minutes; and a model-attestation-required flag. The default for romantic role-play under the age of 16 shall be off.

Example 1 — Wire format: HouseholdPolicy v6 in JSON-LD with a co-guardian signature, jurisdictional override, and Content/Time/Transparency clauses

```
{
  "@context": [
    "https://www.w3.org/ns/credentials/v2",
    "https://parcep.itu.int/ns/v1"
  ],
  "type": ["HouseholdPolicy"],
  "policyId":       "urn:uuid:c1d2e3f4-...",
  "policyVersion": 6,
  "validFrom":      "2026-05-18T00:00:00Z",
  "validTo":        "2027-05-18T00:00:00Z",
  "conformanceProfile": "CONF A",
  "signedBy": [
    { "guardianId": "urn:uuid:guardian-...",
      "proof":      { "type": "DataIntegrityProof", "cryptosuite": "eddsa-2022",
                      "created": "...", "proofValue": "..." } }
  ],
  "jurisdictionalOverrideRefs": [
    "https://pip.fr.example/overrides/under-18-arcom-v3"
  ],
  "subjects": [
    {
      "subjectId": "urn:uuid:subject-...",
      "dependantRefs": ["urn:uuid:dependant-anon-1"],
      "deviceRefs":    ["urn:uuid:device-phone-1", "urn:uuid:device-tablet-1"],
      "clauses": {
        "content": {
          "taxonomyRef":  "https://pip.eu.example/taxonomies/v4",
          "categoryRules":[
            {"category":"adult_sexual",  "action":"BLOCK"},
            {"category":"violence_extreme","action":"BLOCK"}
          ],
          "classifierPolicies":[
            {"modelId":"https://pip.example/models/nudity-v1.2.0",
             "enable":true, "threshold":0.85}
          ]
        },
        "time": {
          "dailyQuotaSeconds": 7200,
          "aggregateAcrossDevices": true,
          "bedtime": {"start":"21:30","end":"07:00"}
        },
        "transparency": {
          "indicatorMode":"PROMINENT",
          "appealChannel":"app://parcep/appeal",
          "digestFrequency":"WEEKLY",
          "childReadableDigest": true,
          "retentionDays": 30
        }
      }
    }
  ]
}
```

12.5 12.5 Operational and discovery layers

The operational layer defines: EnforcementEvent (metadata-grade, with eventId, policyVersion, pepId, timestamp, decision class drawn from BLOCK, WARN, BLUR, DELAY, REDIRECT, ALLOW_LOGGED, category, classifier reference where applicable, confidence in the range zero to one, evidence hash, appealable flag, and appeal identifier where applicable); TransparencyNotice (the dependant-visible artifact emitted by a PEP per enforcement event, with noticeId, dependantId, pepId, trigger-event identifier, shown-at timestamp, acknowledged-at timestamp, and appeal identifier); AppealRequest (the dependant-initiated artifact contesting a decision, with appealId, appellant-dependant identifier, contested-event identifier, free-text reason, status drawn from

OPEN, UNDER_REVIEW, ACCEPTED, REJECTED, resolution, and resolved-at timestamp); GuardianDigest (the periodic aggregation surfaced to the guardian, with digestId, guardianId, window start and end, entries array each comprising severity, category, count, example-event identifiers, an AI-summary object, and a child-readable flag).

The discovery layer defines: PEPRegistration (with pepId, vendor, model, version, conformance profile, capability tags, subscribed-at and last-sync-at timestamps, and public key in COSE_Key format); PIPDescriptor (with pipId expressed as a Decentralized Identifier, ISO-3166-1-alpha-2 jurisdiction, root key, certificate chain, issuance capabilities drawn from GUARDIANSHIP, AGE, CLASSIFIER, TAXONOMY, OVERRIDE, endpoints object, and transparency-log reference).

Example 2 – Wire format: EnforcementEvent in CBOR diagnostic notation, signed COSE Sign1 by the PEP

```
// CBOR diagnostic notation; the wire form is canonical CBOR (RFC 8949),
// signed with COSE Sign1 (RFC 9052) by the PEP's deviceCredential key.

{
  1:  h'a1b2c3...',                  // eventId (UUID, 16 bytes)
  2:  6,                             // policyVersion
  3:  h'pep1...',                    // pepId
  4:  1747512000,                    // timestamp (epoch sec)
  5:  "BLOCK",                       // decision
  6:  "adult_sexual",                // category
  7:  "https://pip.example/models/nudity-v1.2.0",  // classifierRef
  8:  0.91,                          // confidence
  9:  h'sha256(evidence)...',        // evidenceHash
  10: true                           // appealable
}
```

13 13 Synchronization protocol requirements

This clause specifies the requirements that any synchronization-protocol implementation that bridges Policy Administration Points, Policy Decision Points, Policy Enforcement Points, and Policy Information Points shall satisfy. The wire protocol itself is expected to be developed in the IETF as the PARCEP-Sync protocol; the requirements herein constitute the normative input from this Recommendation.

13.1 13.1 Transport

PROT-001 SHALL — The canonical transport for the synchronization protocol between the Policy Administration Point and the Policy Enforcement Points of one household shall be a Messaging Layer Security group (IETF RFC 9420) with the Policy Administration Point as administrator and the Policy Enforcement Points as members.

PROT-002 MAY — An mTLS-over-HTTPS profile may be defined as a peer alternative for legacy or constrained Policy Enforcement Points that cannot run MLS. The peer profile shall provide equivalent authentication, integrity, and confidentiality properties.

PROT-003 SHALL — Every protocol exchange between PAP, PDP, PEP, and PIP shall be authenticated, integrity-protected, and replay-resistant.

13.2 13.2 Method set

PROT-010 SHALL — The synchronization protocol shall define at minimum the following methods: ENROLL (PEP-to-PAP registration), WELCOME (PAP-to-PEP MLS Welcome message), PUBLISH (PAP-to-group new HouseholdPolicy version), UPDATE (PAP-to-group policy delta), REVOKE (PAP-to-group policy revocation), ACK (PEP-to-PAP receipt and applied status), EVENT (PEP-to-PAP metadata-grade EnforcementEvent emission), DIGEST (PAP-to-guardian periodic aggregated digest), APPEAL (PEP-to-PAP dependant appeal routing), RESOLVE (PAP-to-PEP appeal resolution), ATTEST (PEP-to-PIP conformance attestation),

FETCH (PEP-to-PIP signed-artifact retrieval), NOTIFY (MLS group-level notifications), and a namespaced extension primitive for vendor-specific operations.

PROT-011 SHALL — PUBLISH, UPDATE, and REVOKE shall be idempotent under repeat of the same (policyId, policyVersion) tuple.

PROT-012 SHALL — Policy revocation shall propagate to all subscribed Policy Enforcement Points within a propagation window of no greater than 300 seconds when all PEPs are reachable.

PROT-013 SHALL — Concurrent co-guardian updates shall be resolved by a deterministic conflict-resolution algorithm. The default algorithm shall be last-writer-wins keyed by signed timestamp, with both the winning and the losing update preserved in the audit log and surfaced to the co-guardian set.

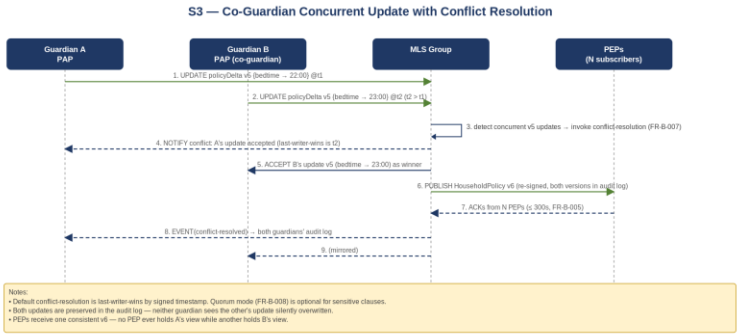

Figure 6 — Co-guardian concurrent update with deterministic conflict resolution.

PROT-014 MAY — A quorum-based conflict-resolution algorithm may be configured for sensitive policy clauses.

13.3 13.3 Discovery and bootstrap

PROT-020 SHALL — Discovery shall support at minimum the following two mechanisms: (a) DNS SVCB records per IETF RFC 9460 under a household-issued domain; (b) a per-Policy-Enforcement-Point .well-known/parcep endpoint per IETF RFC 8615.

PROT-021 SHALL — Bootstrap shall use a guardian-mediated out-of-band pairing artifact (for example, a QR code) carrying the Policy Enforcement Point's public key, conformance profile, and a fresh nonce; the Policy Administration Point then issues an MLS Welcome message that adds the PEP to the household MLS group.

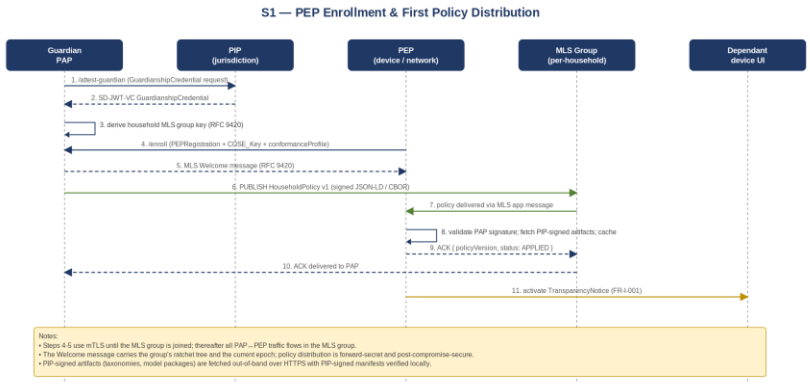

Figure 4 — PEP enrollment and first policy distribution.

13.4 13.4 Wire format

PROT-030 SHALL — The synchronization protocol shall support a dual wire format: JSON-LD with W3C Verifiable Credentials for human-readable policy authoring, web dashboards, and regulatory audit; CBOR with COSE signatures per IETF RFC 8949 and RFC 9052 for event telemetry, network signaling, and IoT-class Policy Enforcement Point exchanges. The two encodings shall be deterministically inter-convertible.

13.5 13.5 Error model

PROT-040 SHALL — Every response shall carry a status drawn from OK, REJECTED, DEFERRED, ERROR and a structured reasonCode drawn from at least the following set: BAD_SIGNATURE, EXPIRED_CREDENTIAL, REVOKED_CREDENTIAL, UNSUPPORTED_PROFILE, INVALID_SCHEMA, JURISDICTIONAL_BAR, CAPABILITY_MISSING, OUT_OF_ORDER, RATE_LIMITED, TEMPORARY_UNAVAILABLE, INTERNAL.

13.6 13.6 Cryptographic agility

PROT-050 SHALL — The synchronization protocol shall be cryptographically agile. Default ciphersuites at version 1 shall include Ed25519 and X25519 for signing and key exchange, ChaCha20-Poly1305 for AEAD, and SHA-256 for hashing.

PROT-051 SHOULD — A post-quantum suite (for example, ML-KEM for key exchange and ML-DSA for signing) should be specified at version 1 as a SHOULD and at version 2 as a SHALL.

14 14 Functional requirements

This clause specifies the functional requirements applicable to PARCEP-conformant implementations. The requirements are organized in thirteen sub-clauses corresponding to the requirement-identifier prefixes FR-A through FR-M.

14.1 14.1 Policy authoring (FR-A)

FR-A-001 SHALL — The Policy Administration Point shall allow a guardian to author a household policy that covers, at minimum, the six feature families defined in clause 12.3: content controls, time controls, contact controls, commerce controls, conduct controls including AI safety, and transparency-to-dependant controls.

FR-A-002 SHALL — The PAP shall allow a guardian to bind a policy to (a) one or more dependant identities and (b) one or more devices or accounts owned or used by those dependants. The binding shall survive a device change with explicit guardian re-binding.

FR-A-003 SHALL — The PAP shall expose age-graded preset profiles (at minimum: under 6, 6 to 9, 10 to 12, 13 to 15, 16 to 17) that act as default policy bundles a guardian may accept, customize, or replace.

FR-A-004 SHALL — The PAP shall accept and respect jurisdictional override descriptors issued by a recognized Policy Information Point that constrain minimum or maximum policy values for dependants under a given age in a given jurisdiction.

FR-A-005 SHOULD — The PAP should support a natural-language policy intent that is deterministically rendered into the formal policy schema. The rendering shall be presented for guardian review before activation.

FR-A-006 SHALL — The PAP shall support policy versioning, with a rollback path to a previous valid policy for at least the last five guardian-issued versions.

FR-A-007 SHALL — The PAP shall require strong authentication of the guardian for any change to a policy, where strong authentication is at minimum two-factor and proportionate to the sensitivity of the change.

FR-A-008 MAY — The PAP may allow a dependant to propose changes to their own policy. Proposed changes shall be subject to guardian approval and shall not be silently applied.

14.2 14.2 Policy distribution and synchronization (FR-B)

FR-B-001 SHALL — PARCEP shall define a normative wire-format for household policy expressed in the vendor-neutral schema of clause 12. The schema shall be machine-readable, signed by the Policy Administration Point, and uniquely identified by a policy version identifier.

FR-B-002 SHALL — PARCEP shall define a normative protocol for policy publication, update, and revocation between the PAP and each subscribed Policy Decision Point. The protocol shall be authenticated, integrity-protected, and replay-resistant.

FR-B-003 SHOULD — Policy distribution should be modeled on a publish-subscribe pattern so that the addition of a new Policy Enforcement Point is handled without changing the policy itself.

FR-B-004 SHALL — Policy distribution shall be tolerant of intermittent connectivity. A Policy Decision Point that has not contacted the PAP for at least 72 hours shall continue to enforce the last valid policy and shall surface a transparency event to the guardian no later than the next contact.

FR-B-005 SHALL — Policy revocation shall propagate to all subscribed PDPs within a bounded propagation window. The bound shall be no greater than 300 seconds for the case where all PDPs are reachable.

FR-B-006 SHOULD — PARCEP should support partial policy updates (delta synchronization) so that a small change does not require redistributing the entire policy.

FR-B-007 SHALL — PARCEP shall define the conflict-resolution algorithm when two co-guardians issue conflicting policy updates within the propagation window. The default algorithm shall be last-writer-wins keyed by signed timestamp, with both the winning and the losing update preserved in the audit log and surfaced to the co-guardian set.

FR-B-008 MAY — PARCEP may support a quorum-based conflict-resolution algorithm for co-guardian sets above two, where a configurable threshold of co-guardian signatures is required for policy changes above a configurable sensitivity level.

14.3 14.3 Policy enforcement (FR-C)

FR-C-001 SHALL — Each Policy Enforcement Point shall enforce the most recent valid policy from its subscribed Policy Decision Point. The PEP shall not silently relax the policy in the absence of a connection to the PDP; the last valid policy continues to apply.

FR-C-002 SHALL — PEPs that consume the same policy shall produce semantically equivalent decisions for the same request. PARCEP shall publish a reference test suite to verify cross-PEP semantic equivalence.

FR-C-003 SHALL — A PEP shall emit a transparency event (see clause 14.9) for every enforcement decision that materially affects the dependant (block, blur, delay, warn, redirect).

FR-C-004 SHALL — A PEP shall implement at least one fail-safe state for each enforcement category. The fail-safe behavior shall be defined by the policy and shall default to the more restrictive option when the PEP cannot reach the PDP and no fail-safe is configured.

FR-C-005 SHOULD — PEPs should expose machine-readable enforcement decisions to a guardian-readable feedback channel so that the guardian can audit, correct, or appeal a decision without leaving the PAP.

FR-C-006 SHALL — Network-side PEPs (CPE, ISP, MNO) shall enforce policy using mechanisms that do not require traffic decryption, content scanning, or downgrade of encryption protections.

FR-C-007 SHOULD — Device-side PEPs should be able to enforce policy when the dependant is offline, by caching the policy locally with integrity protection.

FR-C-008 SHALL — Enforcement decisions related to commerce (purchase approval, spend cap) shall block the transaction by default when the PEP cannot reach the PDP, regardless of fail-safe settings.

14.4 14.4 Identity, roles, and delegation (FR-D)

FR-D-001 SHALL — PARCEP shall define a dependant identity that is independent of network addressing (MAC, IP, IMEI, IMSI) and survives device hand-off, SIM swap, and network change.

FR-D-002 SHALL — PARCEP shall define a guardian identity that supports proof-of-guardianship credentials. The credential shall be revocable and shall carry an explicit jurisdictional scope.

FR-D-003 SHALL — PARCEP shall support co-guardianship as a first-class role. The co-guardian model shall accommodate (a) separated households for the same dependant, (b) short-term carers including kinship caregivers and youth-serving organizations, and (c) schools acting in loco parentis for the school-relevant subset of the policy.

FR-D-004 SHALL — Delegation of guardian authority shall be time-bounded by default. A delegation shall carry an explicit start, an explicit end, and an explicit scope (which subset of the policy the delegate can administer).

FR-D-005 SHALL — Successor management shall be supported: PARCEP shall provide a defined procedure for guardian-of-record transfer in the event of death, custody change, dependant aging-out, or guardianship revocation by competent authority.

FR-D-006 SHOULD — The dependant identity should be expressible as a verifiable credential compatible with IETF and W3C credential formats (SD-JWT-VC, ISO/IEC 18013-5 mDL) so that it interoperates with the EU Digital Identity Wallet.

FR-D-007 SHALL — PARCEP shall require minimum disclosure: a PEP shall receive only the information required to evaluate the policy, and shall not receive the full dependant identity by default.

FR-D-008 SHOULD — PARCEP should support pseudonymous dependant identifiers per-vendor so that the same dependant cannot be cross-correlated across vendors without explicit guardian consent.

14.5 14.5 Content controls (FR-E)

FR-E-001 SHALL — PARCEP shall define a normative minimum content category taxonomy. The taxonomy shall include at least: adult sexual content, violence, gambling, drugs, self-harm, weapons, hate speech, age-inappropriate marketing, malware and phishing, and CSAM-indicative content.

FR-E-002 SHOULD — PARCEP should define an extensibility mechanism that allows vendors and jurisdictions to extend the taxonomy without forking the schema. Vendor extensions shall be namespaced.

FR-E-003 SHALL — PARCEP shall define an age-rating mapping that maps the normative taxonomy onto the principal regional rating schemes (ESRB, PEGI, IARC, CSA/Arcom, BBFC, CERO, USK, ACB, Indian U/UA/A) so that a household policy expressed in one regional scheme can be consumed by a PEP that natively understands another.

FR-E-004 SHALL — Content classification at the PEP shall be capable of operating on-device. Cloud classification may be used additionally but shall not be required for conformance.

FR-E-005 SHALL — Allowlist and blocklist primitives shall be supported and shall take precedence over category-based decisions when both apply, in the direction set by the policy.

FR-E-006 SHALL — Search-engine and recommender-system filtering primitives (SafeSearch, restricted-mode equivalents) shall be addressable from the policy schema.

14.6 14.6 Time controls (FR-F)

FR-F-001 SHALL — PARCEP shall define a normative time-control schema covering at least: daily total quota; per-app or per-category quota; daily schedule (allowed and blocked windows); bedtime; downtime; allowed-always overrides for educational apps and emergency contacts.

FR-F-002 SHALL — Quota consumption shall be aggregated across devices for the same dependant when the policy specifies a shared quota.

FR-F-003 SHOULD — PARCEP should support time-quota carry-over and earning rules (for example, bonus time for completing a chore or homework session).

FR-F-004 SHALL — The time-control PEP shall warn the dependant in age-appropriate language at least once before enforcing a quota limit.

14.7 14.7 Contact and communication controls (FR-G)

FR-G-001 SHALL — PARCEP shall define normative contact-control primitives covering at least: messaging-permitted contact set (default-friends-only / mutual-followers / parent-approved / off); friend-request approval policy; voice/video-call permitted contact set; livestream broadcast/receive policy.

FR-G-002 SHALL — Contact-control policies shall be expressible without requiring the PEP to disclose the dependant's social graph to the PAP.

FR-G-003 SHALL — PARCEP shall support anti-grooming linguistic classification as an optional PEP capability. When enabled, the classifier shall run on-device or within the dependant's trust boundary on the platform, and shall not exfiltrate message content.

FR-G-004 SHOULD — PARCEP should integrate with platform safety signals where the platform already runs server-side anti-grooming detection. Integration shall be by classifier output (signal) rather than content.

14.8 14.8 Commerce controls (FR-H)

FR-H-001 SHALL — PARCEP shall define normative commerce-control primitives covering at least: per-transaction approval; daily, weekly, or monthly spend cap; category-level approval (in-app purchases, subscriptions, gambling-like loot mechanics); approval-channel (PAP push notification, guardian email).

FR-H-002 SHALL — Commerce PEPs shall block the transaction by default when the PEP cannot reach the PDP.

FR-H-003 SHOULD — PARCEP should expose distinct treatment for loot-box and chance-based monetization mechanics independent of the underlying age rating, in line with regulatory direction in Belgium, the Netherlands, the United Kingdom, and Australia.

14.9 14.9 Transparency to the dependant (FR-I)

FR-I-001 SHALL — Every Policy Enforcement Point that enforces a policy on a dependant device or platform shall emit a visible indicator to the dependant whenever the policy is actively monitoring or restricting the dependant.

FR-I-002 SHALL — The indicator shall identify (a) the nature of the restriction, (b) the guardian role enforcing it (without revealing the personal identity of the guardian beyond what is necessary), and (c) the appeal channel.

FR-I-003 SHALL — The dependant shall have an appeal mechanism that routes a request to the guardian Policy Administration Point within a bounded time.

FR-I-004 SHALL — Indicators and appeal flows shall be expressed in age-appropriate language. PARCEP shall define a minimum age-graded vocabulary.

FR-I-005 SHOULD — Indicators should be locale-aware and should be reviewed by a child-rights expert at standardization time.

FR-I-006 SHALL — Surveillance-class monitoring (continuous screen recording, keystroke logging, ambient audio capture) shall not be allowed as a PARCEP-conformant feature. Implementations that include such capabilities shall not claim PARCEP conformance.

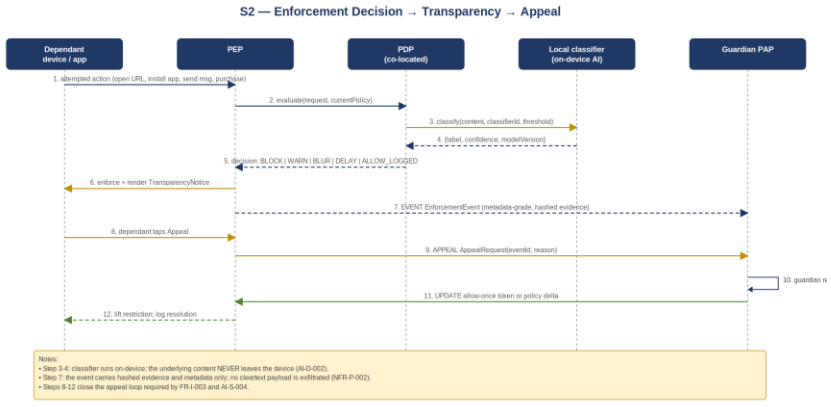

Figure 5 — Enforcement decision, transparency notice, and dependant appeal loop.

14.10 14.10 Telemetry, logging, and feedback (FR-J)

FR-J-001 SHALL — Each PEP shall maintain a local enforcement log retained for a configurable, policy-defined window.

FR-J-002 SHALL — Logs visible to the guardian shall be metadata-grade by default: who, when, what category, what decision. They shall not include the content of communications and shall not include precise location unless the policy explicitly authorizes that disclosure.

FR-J-003 SHALL — Logs shall be readable in the PAP, exportable in a documented format, and deletable by the guardian at any time.

FR-J-004 SHOULD — Logs should be aggregated into guardian-readable digests that summarize without overwhelming. Severity-categorized highlights are preferred over raw enumeration.

FR-J-005 SHALL — Telemetry sent off-device shall be authenticated, integrity-protected, and minimized. PARCEP shall define the maximum telemetry payload classes that may leave the dependant's device.

14.11 14.11 Device management integration (FR-K)

FR-K-001 SHALL — PARCEP shall define a normative integration profile with established mobile device management protocols (MDM, including Apple MDM, Android Enterprise, Microsoft Intune, Samsung Knox) so that a household policy can lock down critical device settings (DNS resolver, profile installation, sideloading) without requiring vendor-specific PAP code.

FR-K-002 SHOULD — The MDM integration should support school-issued devices that operate in a dual-management context (school during school hours, household otherwise), with explicit precedence rules.

FR-K-003 MAY — PARCEP may define a Matter integration profile for smart-home enforcement points where the protocol applies (smart speakers, smart TVs, IoT cameras).

14.12 14.12 Network-based enforcement (FR-L)

FR-L-001 SHALL — PARCEP shall support a network-side enforcement profile that allows a CPE (home router) or an ISP/MNO to act as a PEP for traffic from the dependant's device on that network.

FR-L-002 SHALL — Network-side enforcement shall not require deep packet inspection, traffic decryption, or content scanning.

FR-L-003 SHALL — Network-side enforcement shall use authenticated, encrypted DNS (DoH per IETF RFC 8484, DoT per IETF RFC 7858, or DoQ) to enforce content policy, with structured DNS error data (IETF DNSOP draft) to convey enforcement reasons.

FR-L-004 SHALL — Network-side enforcement shall pin the encrypted DNS resolver via discovery mechanisms compliant with IETF RFC 9462 (Discovery of Designated Resolvers) and IETF RFC 9463 (Discovery of Network-designated Resolvers), so that ECH and DoH-in-browser cannot silently bypass the policy.

FR-L-005 SHALL — Network-side enforcement shall not aggregate per-dependant telemetry centrally. Operators may retain coarse-grained category-level statistics for service operation but shall not retain per-dependant browsing history beyond the bounded retention required by applicable law.

FR-L-006 SHOULD — When the dependant's device leaves the operator's network for a different access network, the household device-side policy should remain in force. Operators should not depend on the network as the sole PEP.

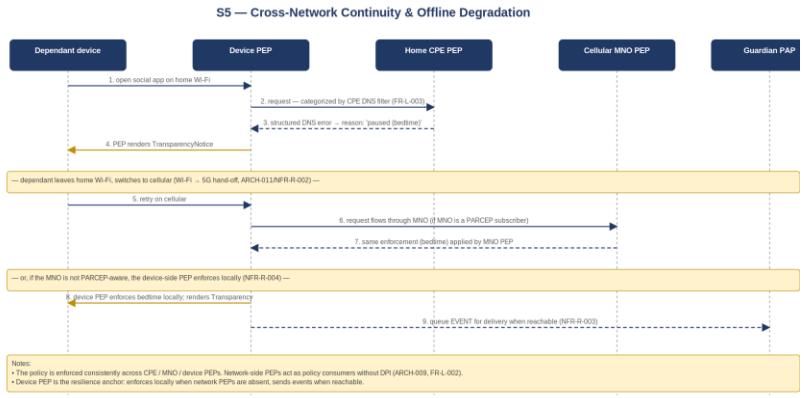

Figure 8 — Cross-network continuity (Wi-Fi to cellular hand-off) and offline degradation.

14.13 14.13 Age-assurance integration (FR-M)

FR-M-001 SHALL — PARCEP shall define a slot in the policy schema for a dependant age signal. The slot shall distinguish at least: declared (parent-asserted), estimated (face or behavioral), and verified (identity-grade).

FR-M-002 SHALL — Age signals consumed by PARCEP shall be expressed as verifiable credentials issued by a recognized PIP and shall be revocable.

FR-M-003 SHOULD — PARCEP should accept age signals expressed in the European Digital Identity Wallet (eIDAS 2) age-verification building block, in Apple's Declared Age Range API (iOS 26), and in Google Play's Age Signals API.

FR-M-004 SHALL — When a jurisdiction requires identity-grade age verification for a specific service category (for example, France/ARCOM for adult content, UK/Ofcom HEAA), PARCEP shall pass the jurisdictional requirement through to the PEP so that the PEP can enforce the higher bar.

FR-M-005 SHALL — Facial-age estimation, when used as part of a PARCEP-conformant flow, shall meet the demographic-fairness criteria defined by Ofcom HEAA and shall publish per-demographic accuracy on at least the NIST FATE benchmark cycle.

FR-M-006 SHALL — Identity-grade age verification shall preserve user privacy via selective disclosure (for example, age-threshold attribute disclosure without identity disclosure), consistent with the EUDI Wallet double-anonymity model required by ARCOM.

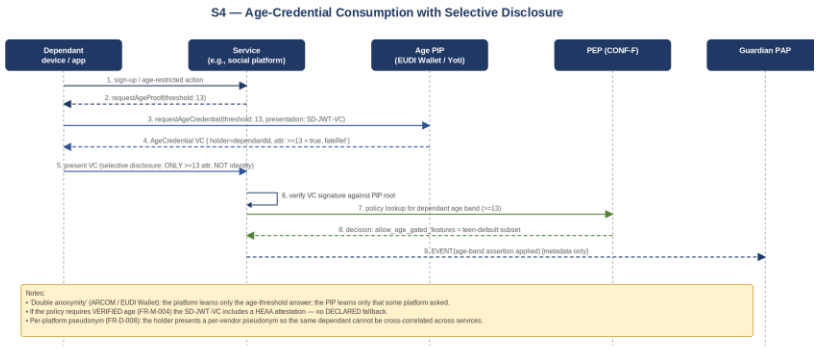

Figure 7 — Age-credential consumption with selective disclosure (double-anonymity flow).

15 15 AI-related requirements

This clause addresses the two faces of AI in child online protection: AI as a threat and AI as an enabler. It also specifies AI safety requirements that any AI subsystem must meet to claim PARCEP conformance.

15.1 15.1 AI threat-defense capability requirements (AI-T)

AI-T-001 SHALL — A PARCEP-conformant device-OS or platform PEP shall expose a normative interface for on-device generative-image and generative-video safety classification, so that newly synthesized content (including CSAM and deepfake imagery) can be evaluated before it is sent, received, or rendered.

AI-T-002 SHALL — A PARCEP-conformant platform shall expose a normative interface for known-CSAM perceptual-hash signaling (PhotoDNA, CSAI Match, Thorn Safer hash sets) at the platform's content-acceptance boundary; PARCEP shall not mandate on-device hash matching for client-side communications.

AI-T-003 SHALL — A PARCEP-conformant messaging or social platform shall expose an anti-grooming classifier output to its in-app safety pipeline, so that guardian-relevant signals can be surfaced without disclosing message content.

AI-T-004 SHALL — A PARCEP-conformant device-OS or platform shall provide a normative interface for deepfake-imagery detection on outbound content (for example, warning the dependant before sharing AI-synthesized intimate imagery of a peer).

AI-T-005 SHALL — A PARCEP-conformant device-OS shall provide a normative interface for AI-voice-cloning detection on inbound real-time audio (for example, warning the dependant during a call when synthetic-voice indicators are present).

AI-T-006 SHALL — A PARCEP-conformant recommender system that serves a dependant shall expose feed-composition analytics to the guardian PAP at a normative granularity (concentration of harm-classified topics per category, per time window), without exposing the underlying recommendation algorithm.

AI-T-007 SHOULD — A PARCEP-conformant browser or messaging client should warn a dependant before submitting credentials or personal data to a URL whose newness and reputation signals exceed a configurable risk threshold.

15.2 15.2 AI-enabler requirements (AI-D)

AI-D-001 SHALL — AI classifiers used inside a PARCEP-conformant PEP shall operate on-device by default. Server-side classification may be additionally used by the upstream platform but shall not be a precondition for PARCEP conformance on the dependant device.

AI-D-002 SHALL — An on-device classifier shall not exfiltrate the content it classifies. Classifier outputs (label, confidence, decision class) may be communicated to the PDP; the underlying content shall not be.

AI-D-003 SHALL — Classifier model packages consumed by a PARCEP-conformant PEP shall be signed by a recognized Policy Information Point. Versioning, integrity, and revocation shall follow the trust-model requirements of clause 11.3.

AI-D-004 SHALL — Classifier model packages shall declare their training-data provenance, evaluation methodology, and per-demographic performance to a depth sufficient for guardian-readable disclosure (model card).

AI-D-005 SHOULD — PARCEP should specify a Generative-AI Service Profile (see clause 12.4) that any generative AI service marketed to or accessible by a dependant must implement to be referenceable by a PARCEP household policy.

AI-D-006 SHALL — Guardian digests generated by AI summarization shall be redaction-controlled, retention-bounded, and dependant-readable. The dependant shall have the right to inspect the digest about themselves.

AI-D-007 MAY — PARCEP may specify an AI guardian-assistant capability that turns a natural-language policy intent into a draft policy. The draft shall be presented for guardian review and shall not be activated without explicit guardian consent.

AI-D-008 SHOULD — AI-driven age estimation, when used in a PARCEP-conformant flow, should reference the NIST FATE Age Estimation benchmark and report submission identifiers in the model card.

15.3 15.3 AI safety, fairness, and auditability (AI-S)

AI-S-001 SHALL — Any AI subsystem used in a PARCEP-conformant flow shall publish per-demographic accuracy on at least skin tone (using a documented scale such as Monk Skin Tone), age band, and gender presentation; degradation greater than a defined floor relative to overall accuracy shall be treated as a conformance failure.

AI-S-002 SHALL — Confidence thresholds for safety decisions (block, blur, warn, escalate) shall be documented in the model card and adjustable by the guardian within bounds defined by jurisdictional override.

AI-S-003 SHALL — Adversarial robustness shall be evaluated and documented; the classifier shall be tested against the perturbation, paraphrase, transformation, and prompt-injection attack classes relevant to its modality.

AI-S-004 SHALL — AI decisions that affect the dependant (block, blur, warn, escalate to guardian) shall be appealable. The appeal path shall be documented, observable to the dependant, and resolvable within a bounded time.

AI-S-005 SHALL — AI decision logs shall include the model identifier, the model version, the confidence score, and a reference to the evidence (without disclosing the evidence itself in cleartext to a third party). The log entry shall be sufficient to support an appeal.

AI-S-006 SHOULD — PARCEP should align AI auditability requirements with the European Union Artificial Intelligence Act Article 86 (right to explanation) and equivalent regional obligations as they emerge.

AI-S-007 SHALL — AI subsystems shall not be used to circumvent or weaken end-to-end encryption. Client-side classification shall operate above the encryption layer on the sender or recipient device only, and the classification shall not cause cleartext to leave the endpoint.

AI-S-008 SHOULD — AI subsystems should publish model behavior in a child-appropriate transparency notice, so that an older dependant can understand what the classifier is looking for, what it is not, and how to appeal a mistake.

16 16 Non-functional requirements

16.1 16.1 Security (NFR-S)

NFR-S-001 SHALL — PARCEP shall not mandate or facilitate any backdoor, key escrow, or cryptographic downgrade. Any conformant implementation shall preserve end-to-end encryption properties at the transport and application layers.

NFR-S-002 SHALL — PAP, PDP, PEP, and PIP communications shall use modern authenticated, integrity-protected cryptographic protocols (TLS 1.3 or successors; MLS where applicable).

NFR-S-003 SHALL — Trust anchors (PIPs) shall publish their root credentials and certificate revocation mechanisms in a manner verifiable by every conformant PDP.

NFR-S-004 SHOULD — PARCEP cryptography should be specified in a post-quantum-aware manner, with a migration path to ML-KEM, ML-DSA, or equivalent NIST PQC primitives as they are standardized.

NFR-S-005 SHALL — Strong authentication of the guardian shall be required for any policy change. The minimum authentication strength shall be two-factor and shall be proportionate to the sensitivity of the change.

NFR-S-006 SHALL — PARCEP shall publish a threat model (see clause 17) and conformance shall include adversarial-robustness testing against the threat model.

16.2 16.2 Privacy (NFR-P)

NFR-P-001 SHALL — PARCEP shall practice data minimization: a PEP shall receive only the policy attributes required for its enforcement decisions; a PDP shall request only the information required to evaluate the policy.

NFR-P-002 SHALL — Personal data of the dependant shall not leave the dependant's trust boundary except as required to enforce the policy. The boundary is, by default, the dependant's device; the policy may extend the boundary to a guardian-controlled household scope.

NFR-P-003 SHALL — PARCEP shall be compatible with GDPR-K, COPPA, the UK Children's Code, the EU DSA Article 28(1), the EU AI Act, ARCOM (France), and analogous frameworks. Where these frameworks set incompatible requirements, the conformant implementation shall apply the stricter of the two for the dependant's jurisdiction.

NFR-P-004 SHALL — Surveillance-class capabilities (continuous screen recording, keystroke logging, ambient audio capture, full message-content disclosure to a guardian) shall not be PARCEP-conformant.

NFR-P-005 SHALL — PARCEP shall require explicit purpose limitation on every data class. Telemetry collected for safety purposes shall not be repurposed for marketing or advertising.

NFR-P-006 SHOULD — PARCEP-conformant systems should support pseudonymous dependant identities to limit cross-vendor correlation.

16.3 16.3 Performance and scalability (NFR-Perf)

NFR-Perf-001 SHALL — Policy distribution from PAP to a freshly subscribed PDP shall complete within 30 seconds under normal network conditions.

NFR-Perf-002 SHALL — Enforcement decisions at a device PEP shall add no more than a defined budget of latency to the dependant's flow (the budget shall be specified per category: for example, 100 ms for content categorization, 500 ms for age-verification credential check, 1 second for cross-device quota aggregation).

NFR-Perf-003 SHOULD — On-device classifiers should achieve their stated accuracy at a compute and energy budget compatible with mid-range devices in the dependant's region (i.e., no requirement that creates a digital divide for less-affluent dependants).

NFR-Perf-004 SHALL — PARCEP shall scale to a household with at least eight dependant identities and at least 25 dependant-owned or dependant-used devices and accounts.

16.4 16.4 Usability (NFR-U)

NFR-U-001 SHALL — The PAP shall expose an age-graded preset library so that a guardian with no technical expertise can produce a usable policy in under five minutes.

NFR-U-002 SHALL — The dependant-facing transparency surface shall be expressed in age-appropriate language and shall be reviewed at standardization time by a child-rights expert.

NFR-U-003 SHOULD — The PAP should support locale-specific language and conventions (date, time, age, jurisdictional terms) for at least the six UN official languages plus the language of the dependant's jurisdiction.

NFR-U-004 SHOULD — The PAP should be accessible to guardians with disabilities (WCAG 2.2 AA or successor).

16.5 16.5 Resilience (NFR-R)

NFR-R-001 SHALL — PARCEP shall function in the presence of TLS 1.3, Encrypted Client Hello, DNS-over-HTTPS, DNS-over-TLS, and QUIC.

NFR-R-002 SHALL — PARCEP shall function across a Wi-Fi-to-cellular hand-off without loss of enforcement and without policy degradation.

NFR-R-003 SHALL — PARCEP-conformant device PEPs shall continue to enforce a cached policy when offline, with explicit transparency events to the guardian no later than the next connectivity event.

NFR-R-004 SHOULD — PARCEP should specify a behavior for the dependant device behind a third-party VPN: at minimum, the device PEP shall continue to enforce its policy locally even if network-side enforcement is bypassed.

NFR-R-005 SHALL — PARCEP shall not depend on a single trust anchor: a single PIP failure shall not prevent enforcement.

16.6 16.6 Compliance (NFR-C)

NFR-C-001 SHALL — PARCEP conformance shall be expressible in a published, jurisdictional compliance map covering at least GDPR-K, COPPA, the UK Children's Code and Online Safety Act, the EU DSA, the EU AI Act, ARCOM (France), Ofcom HEAA, eIDAS 2 / EUDI Wallet, NIST FATE, the US REPORT Act 2024, Australia OSA and AATT findings, Japan's Filtering Act, Korea's Telecommunications Business Act parental-control obligations, and China's 2024-2025 CAC Minor Mode regulations.

NFR-C-002 SHALL — Conformance shall be testable. A reference test suite covering each requirement in this catalogue shall be published alongside this Recommendation.

NFR-C-003 SHOULD — Conformance levels should be reported in a machine-readable form so that a guardian's PAP can verify, before subscribing a PEP, which conformance profile and which optional requirements the PEP implements.

17 17 Threat model

PARCEP publishes an explicit threat model. The threat model captures the adversaries the system is designed to resist and explicitly excludes adversaries that lie outside the system's scope (for example, physical seizure of an unencrypted device).

17.1 17.1 In-scope adversaries

The technically sophisticated dependant attempting to bypass the policy via VPN, alternative DNS, alternative operating system, sideloading, rooted or jailbroken device, guest accounts, or alternative network access (school Wi-Fi, public Wi-Fi, friend's hotspot, cellular).

The malicious app or service that misrepresents its age suitability, exfiltrates dependant data outside its declared purpose, or persists despite a deletion request.

The adversarial content producer using generative AI to evade classifiers (perturbation, paraphrase, embedding-space attacks, prompt injection).

The social-engineering adult predator using LLM tooling to scale grooming, deepfake imagery for sextortion, or voice-cloning impersonation.

The network-layer attacker attempting to observe or modify policy traffic between PAP, PDP, PEP, and PIP.

The coercive adult inside the household who would misuse PARCEP to surveil, stalk, or control the dependant beyond legitimate guardianship; this adversary is one of the principal reasons PARCEP excludes surveillance-class capabilities.

The custody-conflict adversary: a co-guardian acting in bad faith against another co-guardian; an estranged guardian denied guardianship by competent authority; an unauthorized adult attempting to register as a guardian.

17.2 17.2 Out-of-scope adversaries

Physical-access adversaries with the dependant's device unlocked: PARCEP cannot defend an unlocked device against a physically present attacker.

Lawful intercept by competent state authority operating under judicial authorization: PARCEP does not opine on this and does not enable it.

Universal man-in-the-middle adversaries who control the dependant's root trust anchors: PARCEP relies on PKI assumptions consistent with the wider Internet.

18 18 Category-specific conformance profiles

A single PARCEP conformance profile would under-specify some product categories and over-specify others. This Recommendation defines seven category-specific profiles. Each profile inherits the architecture, identity, security, privacy, and resilience requirements; each profile specifies the functional requirements applicable to its category. Profile-level conformance shall be reported in a machine-readable form.

18.1 18.1 CONF-A — Device-OS-native

Applicability: vendor of an operating system that ships a built-in family or supervised-account system. Mandatory requirements: ARCH-001 through ARCH-013; FR-A-001 through FR-A-008; FR-B-001 through FR-B-007; FR-C-001 through FR-C-008; FR-D-001 through FR-D-007; FR-E-001 through FR-E-006; FR-F-001 through FR-F-004; FR-G-001 and FR-G-003; FR-H-001 and FR-H-002; FR-I-001 through FR-I-004 and FR-I-006; FR-J-001 through FR-J-005; FR-K-001 and FR-K-002; FR-M-001 through FR-M-004; AI-T-001, AI-T-004, AI-T-005; AI-D-001 through AI-D-004; AI-S-001 through AI-S-007; NFR-S-001 through NFR-S-006; NFR-P-001 through NFR-P-005; NFR-R-001 through NFR-R-005. Recommended: FR-A-005; FR-D-006 and FR-D-008; AI-D-005 (Generative-AI Service Profile).

18.2 18.2 CONF-B — Third-party device add-on

Applicability: vendor of a parental-control suite that runs on top of a device operating system. Mandatory requirements: ARCH-001 through ARCH-013; FR-A, FR-B, FR-C, FR-E, FR-F, FR-G, FR-I, FR-J groups; AI-D-001 through AI-D-004 when AI classification is offered; AI-S-001 through AI-S-007; the NFR-S, NFR-P, and NFR-R groups. Explicit exclusion: surveillance-class features (FR-I-006, NFR-P-004) — products offering these shall not claim PARCEP conformance.

18.3 18.3 CONF-C — Set-top-box, console, smart-TV, streaming device

Applicability: pay-TV set-top-box, IPTV box, gaming console, streaming device, smart-TV operating system. Mandatory requirements: ARCH-001 through ARCH-013; FR-A-001 through FR-A-003 and FR-A-007; FR-B-001 through FR-B-005; FR-C-001 through FR-C-004; FR-D-001, FR-D-002, FR-D-003; FR-E-001 through FR-E-006; FR-F-001 through FR-F-004; FR-H-001 and FR-H-002; FR-I-001 through FR-I-004; FR-J-001 through FR-J-003; FR-M-001 and FR-M-002; the NFR-S, NFR-P, and NFR-R groups. Recommended: AI-T-006 (recommender-system feed-composition analytics for smart-TV recommenders).

18.4 18.4 CONF-D — ISP, MNO, managed-DNS

Applicability: Internet service provider, mobile network operator, managed-DNS service offering household or per-line parental controls. Mandatory requirements: ARCH-009 through ARCH-012; FR-B-001 through FR-B-005; FR-C-001, FR-C-006, FR-C-007; FR-D-001 and FR-D-007; FR-E-001 through FR-E-005; FR-L-001 through FR-L-006; FR-M-001 through FR-M-004; the NFR-S, NFR-P, and NFR-R groups; NFR-C-001. Out of scope: surveillance-style content scanning; deep packet inspection.

18.5 18.5 CONF-E — CPE, home router

Applicability: home router, mesh Wi-Fi system, edge security appliance. Mandatory requirements: ARCH-009 through ARCH-012; FR-B-001 through FR-B-005; FR-C-001 and FR-C-006; FR-D-001; FR-E-001 through FR-E-005; FR-L-001 through FR-L-005; the NFR-S, NFR-P, and NFR-R groups. Recommended: AI-D-001 if the CPE offers an AI-driven content classifier.

18.6 18.6 CONF-F — Social platform, application

Applicability: social-network family hub, gaming platform with parental supervision, streaming service with kids profiles, messaging platform with parental supervision. Mandatory requirements: ARCH-001 through ARCH-013; FR-A-001 through FR-A-007; FR-B-001 through FR-B-007; FR-C-001 through FR-C-003; FR-D-001 through FR-D-007; FR-E-001 through FR-E-003; FR-G-001 through FR-G-004; FR-H-001 through FR-H-003; FR-I-001 through FR-I-006; FR-J-001 through FR-J-005; AI-T-001, AI-T-003, AI-T-004, AI-T-006; AI-D-001 through AI-D-005 (Generative-AI Service Profile mandatory for AI services); AI-S-001 through AI-S-008; the NFR-S, NFR-P, and NFR-R groups.

18.7 18.7 CONF-G — Age-assurance

Applicability: identity and age-assurance provider. Mandatory requirements: ARCH-001, ARCH-004, ARCH-005; FR-D-001, FR-D-002, FR-D-006, FR-D-007; FR-M-001 through FR-M-006; AI-S-001 (per-demographic accuracy); the NFR-S group; NFR-P with double-anonymity for verification flows consistent with ARCOM and EUDI Wallet guidance; NFR-C-001.

~~Annex A~~

~~<Annex Title>~~Annex A

Reference Architecture and Federated Trust Model

(This annex forms an integral part of this Recommendation.)

~~<Body of annex A>~~**A.1 Architectural decisions**

A.1.1 A purely centralized architecture (a single per-jurisdiction registry of dependants' identities, ages, household policies, and enforcement events) SHALL NOT be adopted. Such an architecture would concentrate the most sensitive class of data covered by this Recommendation in a single repository, would create a surveillance honeypot incompatible with the rights anchored in the UN Convention on the Rights of the Child and General Comment No. 25 (2021), and would directly conflict with the non-goal of this Recommendation stated in clause 7.

A.1.2 A purely peer-to-peer decentralized architecture (no trust anchors at all) SHALL NOT be adopted. Such an architecture cannot anchor content taxonomies, classifier-model packages, or identity-grade age credentials, and cannot satisfy the regulatory consistency required by Article 28(1) of the European Union Digital Services Act, the United Kingdom Children's Code, the Cyberspace Administration of China Minor Mode regulations, and equivalent national frameworks.

A.1.3 The architecture SHALL be federated. Multiple independent jurisdictional Policy Information Points (PIPs) issue and attest credentials and signed artifacts. The household policy SHALL be authored and stored under the guardian's authority at the Policy Administration Point (PAP). Policy distribution between the PAP and the household's Policy Decision and Enforcement Points SHALL flow peer-to-peer over an authenticated, encrypted channel as defined in Annex C. The federation pattern follows established prior art: DNS PKI, the EU eIDAS 2 / EUDI Wallet architecture, Certificate Transparency [IETF RFC 9162], and the IETF Messaging Layer Security and More Instant Messaging Interoperability work.

A.2 Four-tier reference architecture

Tier 1: federated PIPs per jurisdictional function. The functions SHALL include at minimum (a) guardianship-credential issuance; (b) age-credential issuance compatible with the EUDI Wallet age-verification building block, Apple Declared Age Range, and Google Play Age Signals; (c) classifier-model package issuance with signed model card and per-demographic performance; (d) content-taxonomy issuance with cross-regional rating-scheme mapping; (e) jurisdictional-override descriptor issuance. At least two independent PIPs SHALL operate per jurisdictional function.

Tier 2: Policy Administration Point (PAP). Guardian-controlled. Implemented as a mobile application, a guardian web dashboard, or a guardian-controlled cloud service. The PAP SHALL authenticate the guardian using strong (two-factor) authentication; SHALL persist the canonical household policy under the guardian's authority; SHALL expose an age-graded preset library; SHALL support policy versioning and rollback; SHALL accommodate co-guardians, time-bounded delegates, and successor management.

Tier 3: Policy Decision Point (PDP) and Policy Enforcement Point (PEP). The PDP SHALL receive the signed household policy, validate the PIP-signed artifacts referenced therein, cache the policy locally for offline operation, and evaluate per-request decisions against the policy. The PEP SHALL implement the enforcement actions (block, warn, blur, delay, redirect, allow-logged) appropriate to its category. The PDP and PEP are, by default, co-located. For ISP and MNO operators the PDP MAY be remote relative to the PEP.

Tier 4: the dependant. The dependant SHALL receive a visible transparency notice when the policy actively monitors or restricts the dependant, expressed in age-appropriate language, identifying the

nature of the restriction, the role enforcing it, and the appeal channel; this requirement is consistent with the obligations of UK ICO Children's Code clause 11 and the EU DSA Article 28(1) Guidelines.

A.3 Trust model

A.3.1 Every PIP issuance SHALL be recorded in a public transparency log following the model of IETF RFC 9162 (Certificate Transparency v2), so that silent issuance or rogue PIP behavior is detectable by independent auditors. PEPs SHALL be able to verify the inclusion proof of a credential offline.

A.3.2 Every credential issued by a PIP SHALL carry a revocation endpoint. Implementations SHALL support either short-lived credentials (implicit revocation by non-renewal) or status-list endpoints (W3C VC Status List 2021 or equivalent). PEPs SHALL check revocation at policy load time and at a policy-defined refresh interval.

A.3.3 Cross-jurisdictional verification. A PEP operating in jurisdiction A SHALL accept a credential issued in jurisdiction B if the issuing PIP's root is on the PEP's recognized trust list. The trust list SHALL by default be the union of the PEP's home jurisdiction trust list and any additional trust lists explicitly imported by the guardian via the PAP.

Annex B

Reference Data Schema

(This annex forms an integral part of this Recommendation.)

B.1 Schema layers

The reference data schema SHALL be expressed in five layers: (a) identity, (b) credentials, (c) policy, (d) operational, and (e) discovery. Each layer is expressed in a vendor-neutral abstract syntax; concrete encodings SHALL include both JSON-LD (W3C Verifiable Credentials Data Model 2.0) for human-readable and web-compatible exchanges, and CBOR with COSE signatures (IETF RFC 8949 / 9052) for compact, embedded-friendly exchanges. The two encodings SHALL produce semantically equivalent decisions for the same policy.

B.2 Identity layer

The identity layer SHALL define at minimum the following entities: Guardian (with guardianId, primary credential, authentication methods, jurisdictional scope, public key, transparency-log reference); CoGuardianRelation (binding two guardians with scope, type, validity, mutual consent artifact); Dependant (with pseudonymous dependantId, associated guardians, devices, accounts, declared age range, age-credential references, per-vendor pseudonym flag); Device (with deviceId, type, OS, capabilities, PEP endpoint, optional device credential); AccountBinding (per-vendor account pseudonym with scope).

B.3 Credentials layer

The credentials layer SHALL define: GuardianshipCredential (issued by a PIP, with subject, scope, proof format such as SD-JWT-VC or ISO/IEC 18013-5 mDL, validity, revocation endpoint); AgeCredential (selective-disclosure capable, with assertion class DECLARED / ESTIMATED / VERIFIED, age-band attribute, threshold proofs for at-least-13, at-least-16, at-least-18, NIST FATE benchmark reference when ESTIMATED, Ofcom HEAA conformance flag when VERIFIED); ClassifierModelCredential (binding a signed model package to a classifier class such as NUDITY, GROOMING, DEEPFAKE, VOICECLONE, with model-card reference and per-demographic performance disclosures); TaxonomyArtifact (signed content-category taxonomy with cross-regional rating-scheme mapping including ESRB, PEGI, IARC, CSA, BBFC, CERO, USK, ACB); JurisdictionalOverride (regulator-issued minimum policy constraints applicable to dependants below a defined age in a given jurisdiction).

B.4 Policy layer

The policy layer SHALL define: HouseholdPolicy (signed root object with policyId, monotonic policyVersion, policy hash, authoring guardian, signature array including co-guardian signatures, validity, jurisdictional-override references, conformance profile, subjects array, namespaced extensions); PolicySubject (binding a clause set to a dependant set, a device set, and an account-binding set); and the six clause types: ContentPolicy, TimePolicy, ContactPolicy, CommercePolicy, ConductPolicy, TransparencyPolicy.

The six clause types SHALL cover, respectively: content controls (category rules with allow/block/warn action, allow- and blocklists, search safety level, classifier-policy bindings, rating scheme); time controls (daily quota, per-app quotas, allowed and blocked windows, bedtime, allowed-always exceptions, earn rules, cross-device aggregation flag); contact controls (messaging scope, friend-request policy, voice and video call scope, livestream broadcast and receive, anti-grooming classifier enable, blocked-identities list); commerce controls (per-transaction approval threshold, period spend caps, category rules including loot-box and real-money-gambling, approval channel); conduct controls (on-device classifier bindings, Generative-AI Service Profile defined in clause B.4.1, deepfake outbound warning, voice-clone inbound warning); transparency controls

(indicator mode, appeal channel, digest frequency and detail, child-readable-digest flag, retention period, age-graded vocabulary locale).

B.4.1 Generative-AI Service Profile. The ConductPolicy SHALL include a Generative-AI Service Profile that any AI service marketed to or accessible by a dependant must implement to be referenceable by a household policy. The profile SHALL include: a consent model, content-category toggles (sexual content, violence, self-harm, weapons, romantic role-play), distress-signal escalation to the guardian, quiet hours, time bounds, model-attestation transparency, and a no-romantic-role-play default for dependants under the age of 16.

B.5 Operational and discovery layers

The operational layer SHALL define: EnforcementEvent (metadata-grade, with eventId, policyVersion, pepId, timestamp, decision class, category, classifier reference where applicable, confidence, evidence hash, appealable flag); TransparencyNotice (the dependant-visible artifact emitted by a PEP per enforcement event); AppealRequest (the dependant-initiated artifact contesting a decision); GuardianDigest (the periodic aggregation surfaced to the guardian, optionally AI-summarised, with child-readable flag). The discovery layer SHALL define: PEPRegistration (with pepId, vendor, model, version, conformance profile, capabilities, public key); PIPDescriptor (with pipId expressed as a Decentralized Identifier, jurisdiction, root key, certificate chain, issuance capabilities, endpoints, transparency-log reference).

Annex C

PARCEP-Sync Protocol Overview (referenced for IETF development)

(This annex forms an integral part of this Recommendation.)

The synchronization protocol between the PAP, the PEPs, and the PIPs (designated "PARCEP-Sync" for the purposes of this Recommendation) is specified at the level of requirements herein and is expected to be developed as a wire-protocol Internet-Draft in the IETF. This annex captures the requirements that any conformant PARCEP-Sync protocol SHALL satisfy.

C.1 Transport

C.1.1 The canonical transport for PARCEP-Sync between the PAP and the PEPs of one household SHALL be an MLS group (IETF RFC 9420) with the PAP as administrator and the PEPs as members. The choice is motivated by the four properties of MLS required by this Recommendation: cryptographically-proven group membership, forward secrecy and post-compromise security, asynchronous group operations, and a clean key-rotation primitive used when a PEP is removed.

C.1.2 An mTLS-over-HTTPS profile MAY be defined as a peer alternative for legacy or constrained PEPs that cannot run MLS. The peer profile SHALL provide equivalent authentication, integrity, and confidentiality properties.

C.2 Methods

PARCEP-Sync SHALL define at minimum the following methods: ENROLL (PEP→PAP registration); WELCOME (PAP→PEP MLS Welcome message); PUBLISH (PAP→group, new HouseholdPolicy version); UPDATE (PAP→group, policy delta); REVOKE (PAP→group, revoke a policy); ACK (PEP→PAP, acknowledge a policy version with status); EVENT (PEP→PAP, emit metadata-grade EnforcementEvent); DIGEST (PAP→guardian, periodic aggregated digest); APPEAL (PEP→PAP, route dependant appeal); RESOLVE (PAP→PEP, resolve an appeal); ATTEST (PEP→PIP, attest conformance); FETCH (PEP→PIP, fetch signed artifact); NOTIFY (MLS group-level notifications); plus a namespaced extension primitive for vendor-specific operations.

C.3 Discovery and bootstrap

Discovery SHALL support at minimum (a) DNS SVCB records (IETF RFC 9460) under a household-issued domain, and (b) a per-PEP .well-known/parcep endpoint per IETF RFC 8615. Bootstrap SHALL use a guardian-mediated out-of-band pairing artifact (e.g., a QR code) carrying the PEP's public key, conformance profile, and a fresh nonce; the PAP then issues an MLS Welcome message that adds the PEP to the household MLS group.

C.4 Wire format

The PARCEP-Sync protocol SHALL support a dual wire format: JSON-LD with W3C Verifiable Credentials for human-readable policy authoring, web dashboards, and regulatory audit; CBOR with COSE signatures (IETF RFC 8949 / 9052) for event telemetry, network signaling, and IoT-class PEP exchanges. The two encodings SHALL be deterministically inter-convertible.

C.5 Idempotency, ordering, conflict resolution

PUBLISH, UPDATE, REVOKE SHALL be idempotent under repeat of the same (policyId, policyVersion) tuple. Policy revocation SHALL propagate to all subscribed PEPs within a propagation window of no greater than 300 seconds when all PEPs are reachable. Concurrent co-guardian updates SHALL be resolved by a deterministic conflict-resolution algorithm; the default algorithm SHALL be last-writer-wins keyed by signed timestamp, with both the winning and the losing update preserved in the audit log and surfaced to the co-guardian set. A quorum-based alternative MAY be configured for sensitive policy clauses.

C.6 Liaison with the IETF

It is recommended that the wire protocol be developed in the IETF following a Birds-of-a-Feather session and the formation of a dedicated Working Group, with liaison to the MLS, MIMI, DNSOP, OAUTH and DICE Working Groups. The expected initial deliverables are: an architecture and terminology Internet-Draft; the PARCEP-Sync wire-protocol Internet-Draft referencing IETF RFC 9420; an MLS profile Internet-Draft for the household-group ciphersuites and member attestation; and a discovery Internet-Draft for the SVCB and .well-known/parcep mechanisms.

Annex D

Category-Specific Conformance Profiles

(This annex forms an integral part of this Recommendation.)

This Recommendation defines seven category-specific conformance profiles. Each profile inherits the architecture, identity, security, privacy, and resilience requirements; each profile specifies the functional requirements applicable to its category. Conformance level SHALL be reported in a machine-readable form attached to the PEPRegistration of clause B.5.

D.1 CONF-A — Device-OS-native. Applicability: vendor of an operating system that ships a built-in supervised-account / screen-time / family system. Mandatory: architecture and identity layers, policy authoring, distribution, enforcement, content, time, contact, commerce, transparency, telemetry, device-management integration, age-assurance integration; on-device classifier interfaces for deepfake and voice-clone defence; AI safety, security, privacy, resilience.

D.2 CONF-B — Third-party device add-on. Applicability: a parental-control suite that runs on top of a device OS. Mandatory: architecture, policy authoring/distribution/enforcement, content, time, contact, transparency, telemetry, AI safety, security, privacy, resilience. Surveillance-class features (continuous screen recording, keystroke logging, ambient audio, full message-content disclosure to guardian) SHALL NOT be implemented; products offering these SHALL NOT claim conformance.

D.3 CONF-C — Set-top-box, console, smart-TV, streaming device. Mandatory: architecture, policy authoring (subset), distribution, enforcement, identity (subset), content, time, commerce, transparency, telemetry, age-assurance (subset).

D.4 CONF-D — ISP / MNO / managed-DNS. Mandatory: network architecture clauses, policy distribution, enforcement (without DPI, decryption, or content scanning), identity (minimum disclosure), content, network enforcement, age-assurance, security, privacy, resilience, compliance.

D.5 CONF-E — CPE / home router. Mandatory: network architecture clauses, policy distribution, enforcement (DNS-bound, structured DNS error data), identity (minimum disclosure), content, network enforcement, security, privacy, resilience.

D.6 CONF-F — Social platform / application. Mandatory: full architecture, policy authoring, distribution, enforcement, identity, content (subset), contact, commerce, transparency, telemetry, AI threat-and-defence interfaces (including anti-grooming classifier hook, deepfake outbound warning, recommender feed-composition analytics), AI safety, security, privacy. The Generative-AI Service Profile of clause B.4.1 SHALL be implemented by any generative-AI service marketed to or accessible by a dependant.

D.7 CONF-G — Age-assurance. Mandatory: architecture (subset), identity, age-assurance, AI safety (per-demographic accuracy disclosure per NIST FATE), security, privacy with double-anonymity for verification flows consistent with ARCOM and EUDI Wallet guidance, compliance.

~~Appendix I~~

~~<Appendix Title>~~Appendix I

AI-Related Threats and Defences for Child Online Protection

(This appendix does not form an integral part of this Recommendation.)

~~<Body of appendix I>~~

I.1 Threats

The AI-related threat landscape facing dependants has materially worsened between 2023 and 2026. The Internet Watch Foundation reported a 260-fold year-over-year increase in AI-generated CSAM videos in 2025; the National Center for Missing & Exploited Children recorded a 1,325 % rise in generative-AI-related exploitation reports between 2023 and 2024; deepfake sextortion using synthetic nude imagery of real adolescents has emerged as the dominant new harm vector, with documented suicides; companion-LLM products have been documented grooming minors and have been the subject of wrongful-death litigation; AI voice cloning is increasingly used in coercion targeting minors; recommender systems amplify self-harm and pro-anorexia content to vulnerable adolescents.

I.2 Defence capabilities

On-device AI classifiers (illustrated by Apple Communication Safety on iOS and SafeToNet HarmBlock+ in the HMD Fuse device) provide nudity, deepfake, and outbound-share classification at the dependant's trust boundary, with classifier outputs (label, confidence, decision class) communicated to the PDP without exfiltrating the underlying content. Anti-grooming linguistic classification (illustrated by Thorn Safer Predict, which scores conversations line-by-line for early-stage exploitation signals) integrates at the messaging platform's in-app safety pipeline. Known-CSAM perceptual hash matching (PhotoDNA, CSAI Match, Thorn Safer) operates upstream at the content-hosting platform under the obligations of the US REPORT Act and equivalents and is referenced here as complementary infrastructure. Age estimation referenced to the NIST FATE benchmark provides the verifiability floor required by clause B.3 AgeCredential and by Ofcom HEAA guidance.

I.3 Encryption-preserving design

Following the late-2025 convergence of the European Union on the position that mandatory scanning of end-to-end encrypted communications is not technically defensible, this Recommendation requires that any AI classification used in a conformant flow operate above the encryption layer on the sender or recipient device and that no cleartext leave the endpoint. The classifier output may inform a policy decision; the content itself never leaves the dependant's trust boundary.

Bibliography

~~[b-ITU-T X.yyy]~~ ~~Recommendation ITU-T X.yyy (date), *Title*.~~[b-ITU PP-179] ITU Plenipotentiary Resolution 179 (Rev. Bucharest, 2022), ITU's role in child online protection.

[b-ITU-COP] ITU Child Online Protection Guidelines for Parents, Carers, Guardians and Educators (2020).

[b-ITU-UNICEF] ITU / UNICEF Guidelines for Industry on Child Online Protection.

[b-OECD-2021] OECD Recommendation on Children in the Digital Environment, OECD/LEGAL/0389 (2021).

[b-UN-CRC-GC25] UN Committee on the Rights of the Child, General Comment No. 25 (2021), Children's rights in relation to the digital environment.

[b-EU-DSA-Min] European Commission, Guidelines on the protection of minors under the Digital Services Act, 14 July 2025.

[b-UK-ICO-11] UK Information Commissioner's Office, Age-Appropriate Design Code (Children's Code), clause 11 — Parental controls.

[b-Ofcom-HEAA] Office of Communications (UK), Guidance on Highly Effective Age Assurance, April 2025.

[b-COPPA] US Federal Trade Commission, Children's Online Privacy Protection Rule (COPPA), 16 CFR Part 312.

[b-ARCOM] ARCOM (France), Technical guidelines for age verification for the protection of persons under 18 from online pornography.

[b-EUDI-ARF] European Commission, EU Digital Identity Wallet Architecture and Reference Framework, including the age-verification building block.

[b-NIST-FATE] NIST, Face Analysis Technology Evaluation — Age Estimation and Verification (FATE-AEV) benchmark.

[b-IWF-2026] Internet Watch Foundation, How AI is being abused to create child sexual abuse imagery, 2026 report.

[b-WeProtect-2025] WeProtect Global Alliance, Global Threat Assessment 2025.

[b-IETF-MIMI] IETF More Instant Messaging Interoperability (MIMI) Working Group.

[b-IETF-DNSOP-SDE] IETF DNSOP Working Group, Structured DNS Error Data (Internet-Draft).

[b-SD-JWT-VC] IETF, Selective Disclosure JWT Verifiable Credential (draft-ietf-oauth-sd-jwt-vc).

[b-RFC-9162] IETF RFC 9162, Certificate Transparency Version 2.0.

[b-Thorn-Safer] Thorn Inc., Safer Predict (AI-driven CSAM/CSE detection).

_____

_____

_____