Network Working Group
Internet-Draft
Intended status: Informational
Expires: 9 May 2026

M. Knodel
Social Web Foundation
G. A. Scalone
Vodafone Group
T. Newton
Qoria

5 November 2025

Age Verification Architecture
draft-knodel-age-arch-00

Abstract

This document describes solution-agnostic and technology-neutral schema for how various intermediaries can gate content and services based on age. The analysis of the architecture is done based on the effectiveness of permitting or restricting access based on age. The document concludes with recommendations as well as critical privacy, security and human rights considerations.

About This Document

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

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-knodel-age-arch/.

Source for this draft and an issue tracker can be found at https://github.com/mallory/draft-knodel-age-arch.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 9 May 2026.

Table of Contents

1. Background

Our goal is to describe the technical difficulties in any age-gating mechanism such that it is effective and does not introduce security and privacy risks as well as contravene human rights. We also hope to show that age verification mechanisms are wholly technical solutions that are separate from, albeit often motivated by, the means of protecting young people online.

Prior focus on child safety has led to robust trust and safety measures taken by large user platforms whereby content and behaviour moderation became industry standard. Its phases: Define, detect, evaluate, enforce, appeal, educate. This cycle ensures that criminal and unlawful content is taken down. Platforms and services with user generated content can also use this cycle to ensure content or behavior that violates platform terms of service.

In parallel, network operators have for many years implemented forms of age-based access control that rely on content categorization or DNS-level filtering rather than on the collection of personal data. In several jurisdictions, ISPs are required to offer or enforce filtering systems that restrict access to adult or otherwise unsuitable content for minors. These systems typically work by classifying destination domains or content types into broad categories—such as adult, gambling, or violence—and allowing or blocking them according to the subscriber's or guardian's chosen policy. Because they do not require identity documents or individual profiling, such network-assisted methods can provide a baseline of child protection with substantially lower privacy risk. These approaches cannot replace service-level moderation or legal accountability, nor should they be standalone solutions outside of parental controls or family management settings.

Relatively new are proposals to protect children via age gate. Age gating in the analog world is typically enforced at the point of action, sale or entry through hard document checks at liquor stores, as one example. Replicating age gating in the online world has encountered consistent tensions between accuracy and privacy, often reducing the problem to one of feasibility. Often, online age verification has taken the form of requiring websites to show a prompt to the user to self-declare their age or birthdate before they can gain access to content or services.

Some services, websites, or apps require uploading hard documents to verify identity, just as some services, websites or apps require payment with a credit card. These often create an illusion for policy makers that users could be required by any service to require sharing of the same hard documents or credit card details in order to

verify age at scale, both for all users and for a wide variety of platforms and services. However any such system is expensive, difficult to scale, and introduces data protection liability and privacy risks to users— including potential data breaches or the exclusion of users who do not possess traditional identity documents but who would otherwise be lawful users. Hard document identity review in the context of general-use platforms like social media unnecessarily scales the risk of privacy, access, and equity harms from requisite age gating for all users on all apps all the time.

Many age verification methods conflict with data-protection frameworks and data minimization principles and pose serious safety, data security, and privacy risks. As such risks will be present in nearly all alternative methods of age verification, an overview of this category of risk will be of some benefit, while we leave method-specific commentary to the security, privacy and human rights considerations sections. Requiring all users on all platforms to submit verifiable credentials can create large, sensitive data troves in centralized intermediaries that are vulnerable to breaches, fraud, or misuse. Once compromised, this information is difficult—if not impossible—to secure again. Compounding these concerns, precise location—required to determine compliance with jurisdiction-specific laws—is implicitly inferred in all methods, adding another layer of sensitive data to the name and age information being collected, processed, and stored.

From an operational and architectural perspective, centralizing or repeatedly exchanging such verification data may also create systemic risks to resilience, security, and interoperability. Large-scale credential exchanges or cross-border look-ups introduce new attack surfaces and potential points of failure within authentication or content-delivery paths.

1.1. Terminology

Language matters because it shapes how we think about these systems. "Assurance" is less official than "verification" and "estimation" admits inherent imprecision. [EFF-Age-Terminology]

Age Verification Process of confirming a user's age using an authoritative or verifiable source (e.g., government ID, credential, or trusted record).

Age Assurance Broader set of methods providing confidence about a user's age or age range without requiring formal identity verification.

Age-gating - refers to age-based restrictions on access to online services. Age gating can be required by law or voluntarily imposed as a corporate decision.

Age Estimation - Instead of asking you directly, the system _guesses_ your age based on data it collects about you.

Self-Attestation User-declared statement of age or date of birth without third-party validation.

Guardian Attestation Verification performed or consent provided by a parent or legal guardian on behalf of a minor.

Intermediary Any actor between the user and a service, including device OS, application platform, content-delivery system, resolver, or network operator.

Service-Level Age Control Mechanisms enforced by online platforms or applications (e.g., login requirements, content gating, or age tokens).

Network-Assisted Age Control Privacy-preserving measures implemented at or near the connectivity layer, such as DNS or policy-based content filtering, typically based on user or guardian preference and without identity collection.

On-Device Assurance Local verification or estimation methods executed on the user's device (e.g., age-range inference, parental settings) where data does not leave the device.

Credential Provider Entity that issues, stores, or validates user credentials used for age verification or assurance.

Data Minimization Principle that personal data should be limited to what is necessary for a specific, explicit purpose (per frameworks such as GDPR).

Privacy-Preserving Signaling Exchange of metadata or tokens between system components to enforce policy without exposing personal identifiers.

Shared-Responsibility Model Architectural approach in which service, device, and network layers each contribute to age-appropriate access control within their scope, reducing centralization and single points of failure.

Circumvention Any attempt by users or systems to bypass or falsify age-related controls or policies.

Jurisdictional Compliance Ability of a system to align with local or national laws defining age-related content or access restrictions.

2. Analysis of age gating methods

In this section, we analyse purely based on the efficacy of providing and restricting access based on age. Efficacy means that it works, but also we consider feasibility, circumvention and accuracy.

Efficacy should also reflect cooperation across layers—service, device, and network—so that no single actor bears full responsibility or control of user data. In/Feasibility also includes operational scalability—methods that add latency or centralize look-ups may weaken both reliability and privacy. Circumvention may occur not only by users but also through weak or misaligned intermediaries; distributing enforcement across independent layers limits large-scale bypass. In/accuracy of the method, false positives and negatives, and what users can do about it. Recourse and remedy, e.g. who is responsible and what happens when: the determination of age could be wrong; age-verification is inaccurate; no one adopts it and users don’ t use it; people abuse it; companies abuse it; there is a breach; it can be circumvented.

For all of these approaches, there are also a number of overarching issues that make them less desirable. For example, as noted above, many would require collection of precise location information to comply with jurisdiction-by-jurisdiction privacy laws or other requirements. Important to note that geolocating users remains an unsolved problem. Inaccurate geolocation means people will be unnecessarily excluded or have an incorrect set of laws applied to them. This problem is likely to get worse as IPv4 fragments further.

These sorts of efforts would also create complex questions of how to obtain accurate age information from long-standing, existing users. Difficult questions also arise around how users can “age up” in platforms over time, retaking control of their accounts and data without oversight.

Beyond these threshold issues, specific proposals for age gating or age verification may or may not be effective. We take effective to mean that they work for both those permitted and those restricted and that they are feasible, durable against circumvention, and accurate.

Additionally, age gating naturally puts a barrier to entry on a given site. Publishers spend a lot of time optimising their sites for ease of use, and are now being asked to pay a third party to bounce away a lot of their customers - both legitimate and otherwise. As such the economic argument for publishers to ignore the law is strong,

particularly for laws enacted in states where that publisher has no legal entity. Inevitably, in the often used example of pornography, there will be many more non-compliant sites than compliant sites, and these non-compliant sites are likely to be non-compliant in other ways - eg effectively age gating their content creators.

2.1. Age credentials

The State issues the credentials that are "ground truth," however none provide age as the singular datum, thereby disclosing more data about the user than is specific to the mandate to verify age. Examples of the provisioning of credentials include: Issuance of a birth certificate; Issuance of a passport; Teens: Issuance of a driver's licence (though only in the US would this be for a u18); Issuance of a social insurance number / SSN / generic international name (no birthdate).

Other hard documents that risk being less accurate include credit card (no birthdate but demonstrates banked eligibility) or student ID (not standardized, often no birthdate but gives approximation through education level).

Even in digital form, such as national eID or mobile driving-licence systems, these credentials expose more information than necessary.

2.2. Age verification

Age verification is performed when a guardian, the service or third party has direct, verified knowledge of the credential and assures a service of age. Hard document review would require the disclosure of documents containing a wide variety of sensitive information. This information would not be limited to a users' age and would include data not necessary to determine whether or not a user should be permitted to use a social media platform (e.g., address, credit card number, etc.). Requiring such disclosures of all users would create substantial security and privacy concerns (e.g., risk of data breaches, exposure of additional personal information to platforms themselves).

For credentials both physical and digital, unless they are coupled with selective-disclosure or verifiable zero-knowledge proof (ZKP) techniques. Using intermediaries at the device, service, or network layer to handle only minimal attributes ("over 16," "over 18") can preserve regulatory trust while limiting personal-data exposure. Verification frameworks could also rely on distributed attestations or policy-based tokens issued by trusted intermediaries, enabling services to confirm age eligibility without persistent identifiers or central repositories.

Moreover, determined users may still falsify credentials or exploit systemic workarounds, undermining the effectiveness of these measures. Minors and adults alike may migrate to other extra jurisdictional platforms for comparable features rather than share the credentials required to clear a higher bar for age verification.

Considerations for various intermediaries of age include the following.

A guardian who is accountable for the child assures the child's age. This can be limiting for those with nontraditional family structures, wards of the state, and also assumes guardianship duties are always performed in the best interest of the child.

A jurisdiction or a governmental agency to be the arbiter of whether a citizen or resident would be allowed to have access to the internet or access specific websites and services. This would have both process and political consequences.

KYC (know-your-customer laws in countries that define certain industries as sensitive, such as banks, telecoms, etc) could theoretically be extended to any industry including social media, education platforms or blogging websites. However, doing so would blur the line between financial compliance and social regulation, embedding surveillance logics into general internet use. It could also make access to general-use and basic communication services contingent on a financial transaction framework, which not all internet-accessible platforms and services provide.

Operating systemlevel verification shifts responsibility to device vendors and app ecosystems. While this potentially removes redundancy for users who already have a trust relationship with these intermediaries, this approach would concentrate power in a few private actors. The few OS or app ecosystems could incentivize and or be required to facilitate wider-spread age data requests from apps. They could also be in a position to determine or enforce access across wider categories of apps or content beyond characteristics such as age. Such centralization would raise competition, transparency, and accountability concerns, particularly where OS providers operate globally but respond to local regulatory pressures.

Other third-party intermediaries or age verification services could offer modular, privacy-preserving verification as a service layer. Yet without strong oversight and interoperability standards, these actors could create new forms of data brokerage, fragmentation, or lock-in. Their economic incentives—to monetize verification or analytics—may conflict with users' social and cultural rights, especially if pricing or access varies across regions.

In many jurisdictions, regulated entities such as banks or telecommunications operators already perform age or identity checks under audited privacy frameworks. Extending such existing infrastructures with additional privacy-preserving layers could bridge regulatory and technical feasibility.

Verifiers should be held accountable in systems where they are disclosed information (no matter how privacy preserving). Those presenting proofs should have a way to report or contest a request or determination made within these systems. Especially in the case of denial of access. Treating the user as the only potentially hostile party would ignore the power dynamics in jurisdictions where age verification is mandatory. Also, verifiers should be held to a standard of reporting and registration of their scope of collection. Especially in countries where age verification systems are being developed in tandem with digital ID systems.

2.3. Age assurance

Age assurance is an umbrella term often used to describe the various methods whereby the user, a guardian, the service, third party, or age verification service assures the service of age to the direct, or verified knowledge of the credential to various degrees of confidence.

Age assurance can be performed by any of the age verification intermediaries listed in the previous section.

The user can self-attest, which is status quo for almost every service at this time.

A guardian can attest and be a legal guarantor to their children like Facebook Messenger Kids. This mechanism also serves to ensure a user is in fact a child, and not an adult impersonating a child.

Connecting new users to parents' accounts, COPPA "requires those operators to obtain verifiable consent from the children's parents before collecting, using, or disclosing children's personal information" and under GDPR they must "obtain this consent from a parent and make reasonable efforts to verify the identity of that

parent." (Epic Games, Kids Web Services (KWS) at "Welcome to Kids Web Services," https://dev.epicgames.com/docs/kids-web-services/welcome-to-kws (https://dev.epicgames.com/docs/kids-web-services/welcome-to-kws)).

An additional concern is that this measure proliferates rather severely user data, risking feasibility. Rather than minimizing data not just about the child but for the parent, this requires all parents to first verify their age (invoke recursion analysis). Verifiable consent is then required in addition to verifying a parent-child relationship between users, all of which invoke hard document review.

Non-state institutional or contract-bound intermediary like a court-appointed guardian, a school or an employer. The linking of undergraduate and professional emails as methods by which social media platforms could increase the accuracy of their age verification processes fall short because these are typically only given to individuals of undergraduate or professional age, eg o18. At the same time requiring an undergraduate or professional email is likely to exclude a substantial number of adults. Users aged 13-17 who have not begun attending undergraduate institutions or working in professional environment will be unable to create accounts on the platform as a result Additionally, users from lower-income backgrounds, homeschool settings, or international markets where educational institutions do not provide email addresses will be irrationally excluded. Such an arbitrary and discriminatory outcome would amount to an unreasonable limit on access.

The service itself could perform age verification, age estimation or age assurance via the user directly or a third party, eg guardian.

2.4. Age estimation

Some examples of age estimation or inference of age have been proposed to use a variety of behavioural and context-specific signals, not all of which users explicitly opt in to. There are general concerns that the use of user data for the purposes of age estimation can contravene data privacy frameworks that limit data usage to specific consent structures. Nonetheless we elaborate a couple of mechanisms here:

The behaviour of a user on a platform, including who they are friends with or what kind of content they engage in can provide some clues as to their age. However this data is limited to what service providers know about these users and is not as accurate for younger users since statistically there are fewer ways for the platform to know that its estimation is accurate.

Biometric signals are considered age estimation rather than age assurance or verification because it's not rooted in authoritative or ground truth. Biometric methods such as image or video facial scans are accompanied by a variety of flaws that prevent it from being a reasonable alternative. First, facial analysis technology is notoriously unreliable in estimating age, especially for teenagers, whose facial features change rapidly and vary widely. (Ngan, M. and P. Grother, Face Recognition Vendor Test (FRVT) Performance of Automated Age Estimation Algorithms, NIST Interagency Report 7995, Nat'l Inst. of Standards & Tech. (Mar. 20, 2014), https://nvlpubs.nist.gov/nistpubs/ir/2014/NIST.IR.7995.pdf) These tools can easily misclassify users depending on lighting, ethnicity, or facial expression. (Ganel, T., Sofer, C., & Goodale, M. A. (2022). Biases in human perception of facial age are present and more exaggerated in current AI technology. Scientific reports, 12(1), 22519. https://doi.org/10.1038/s41598-022-27009-w (https://doi.org/10.1038/s41598-022-27009-w))

2.5. Other

There are perhaps other more indirect ways of achieving age appropriate use of online content and services. Perhaps broadly we could refer to these as market and content-based solutions.

Market: The prohibition of advertising, marketing, or subscriptions targeting children under the allowed age supports harm reduction while preserving privacy by disincentivizing children from joining platforms. This is realizable through regulatory action and addresses risk without compelling platforms to surveil or profile young or any users by disallowing the acceptance of payments for advertisements to children under the allowed ages.

At the infrastructure level, network or payment intermediaries can also enforce such prohibitions through category-based or transaction-type filtering that does not require user profiling. These methods are enforceable through consumer protections and guardrails that exist in the financial sector.

Rather than continuous geolocation, minimal-exposure signaling between networks and services could meet jurisdictional needs without tracking individuals. Networks maintain awareness of user location and applicable jurisdiction through routing, addressing, and service-delivery functions. This contextual knowledge could support privacy-preserving determination of the user-to-jurisdiction link, avoiding the need for continuous device-level geolocation. However there are implications for enforcement at the network level given the broad surface available for network level censorship capability on a per-user basis.

Content moderation more broadly is the commitment of services and platforms to ensure online experiences are fit for purpose, given the intended user base. For example certain k-12 educational websites with user generated content are certainly engaged in some strong degree of content moderation because the platform is for learning in schools, and not, say, popular culture or entertainment.

3. Enforcement

Once age has been determined to a satisfactory degree, it's important to interrogate the means by which content or a service is either accessible or not. This happens solely at the service level. In this section we expand on what happens once the service has obtained age assurance, or age verification if it has been party to hard documentation directly.

It is generally important to note that the very definition of "adult content" is not universal. Material that one jurisdiction or culture classifies as adult—such as partial nudity, depictions of smoking, or the use of firearms—may be considered artistic, commercial, or sporting content elsewhere. As a result, even a compliant service might not be aware that some of its material would qualify as adult under another legal framework. This variability underscores the need for proportional, interoperable signaling systems rather than rigid or global classifications.

Platforms and services enforce by limiting access or visibility of certain categories of content based on the verified age attribute of a user or account. This can take several forms: hard blocks on sign-up, "shadow" restrictions such as hiding posts from underage accounts, or algorithmic filtering that reduces exposure to sensitive material. Enforcement at the service level can also include requiring a logged-in session tied to an attested age, but this raises issues around traceability, persistent identifiers, and the erosion of anonymous access to lawful content.

Devices enforce by interpreting or acting on signals from applications, browsers, or operating systems to restrict or permit access. For instance, parental-control settings or OS-level content ratings can automatically block apps or sites flagged as adult. Device-level enforcement is often marketed as a privacy-preserving solution because it happens locally, but it can also centralize control in a few proprietary ecosystems and create dependency on vendor-defined "age ratings." This can entrench cultural biases or commercial interests rather than reflect nuanced or localized standards.

Some considerations for networks that enforce age gating may be instructive for any system design as well, though content moderation at this lower network layer tends to be objectionable [[RFC 9505]]:

* There exist internet standards for any service to signal to the user that it is adult-only. Any device, including web browsers, can elect to confirm this signal. It would be recommended to require all services that disallow users under a specific age to use this standard to use the “restricted to adults” label and that the OS, browser and app levels heed this signal. (RTA Label, available at https://www.rtalabel.org/ (https://www.rtalabel.org/)) This should also be paired with self-attestation methods. Enforced by an ISP would be network level censorship.

* Similar signaling could extend to DNS or resolver functions, where privacy-preserving flags indicate that a domain hosts age-restricted material. [[cite http://blog.cloudflare.com/introducing-1-1-1-1-for-families/ (http://blog.cloudflare.com/introducing-1-1-1-1-for-families/)]] Such information can guide optional filtering at the user or guardian’s request without content inspection or identity disclosure.

There is a view that common, interoperable approaches across service, device, and network layers could lower compliance costs and align incentives for safer adoption. However this lowering of barriers to implementation would be considered a negative if age-gating is considered a form of disenfranchisement and censorship, or if the age-gating capabilities are applied to additional attributes of end users.

4. Types of web-based platforms and internet services

We have described a range of methods for age gating, however we must acknowledge that not all content, platforms, and services need to take the same approach. Enforcement need not be uniform across all Internet services. Core infrastructure and access networks may implement category-level controls; devices and browsers can interpret standardized labels; and services can apply context-specific age checks. Aligning these layers reduces redundancy and minimizes risk.

It is our hope that in this section we try to describe the type of platforms and services such that “no age assurance”, “self-attestation” and “hard document age verification” can vary in their approach, and also vary across age groups as well: under 13, teens, over 18/adult.

Core Internet infrastructure — Physical connectivity, routing, DNS, and encryption layers. These should remain neutral and privacy-preserving, but can optionally support privacy-preserving signaling (e.g., domain-category metadata or parental-policy tokens) chosen by users or guardians.

Governmental and intergovernmental services — Sites such as taxation, immigration, or public-sector portals (e.g., irs.gov (http://irs.gov)). These normally require verified identity and age by design and operate under statutory data-protection frameworks.

Essential social, cultural, and economic rights (ESCR) services — Platforms providing access to banking, education, news, or health information. These typically must be broadly accessible regardless of age and should avoid unnecessary barriers.

General-purpose platforms — Social-media, messaging, gaming, or e-commerce services serving mixed audiences. These are best suited to layered age-assurance approaches combining user, guardian, and optional network-assisted controls, depending on jurisdiction.

Context-dependent or borderline content — Material whose classification varies by culture or law (e.g., artistic nudity, tobacco, weapons). Here, enforcement may rely on localized signaling or policy mapping at device or network level rather than global content bans.

Adult-only services — Platforms distributing explicit sexual, gambling, or similar content (e.g., pornhub). These require high-assurance age verification but should implement it with minimal data disclosure and interoperable signaling for users and networks.

5. Types of Platforms and Services

While age assurance can apply across many kinds of online services, the majority of user exposure and regulatory debate concerns general-use platforms— social-media, messaging, gaming, and app-distribution ecosystems that mix adult and minor audiences. These platforms collect large volumes of user data and operate globally, so privacy, feasibility, and jurisdictional diversity become the most acute design challenges.

In some jurisdictions, regulations already mandate the blocking or restriction of specific platforms for users below a certain age— for instance, national orders prohibiting access to social-media services such as TikTok for minors. Such measures highlight both the policy urgency and the architectural complexity of enforcing age-based restrictions at scale, since these platforms often serve mixed audiences and operate across borders.

General-use platforms can be further grouped according to how users interact and how moderation and access controls are applied:

Social-interaction platforms (e.g., social networks, video-sharing, live streaming) — rely heavily on user-generated content and recommendation algorithms; enforcement usually combines self-attestation, parental tools, and service-level moderation.

Communication platforms (e.g., messaging, voice, forums) — involve private or semi-private exchanges where content is not always visible to the provider; here, age assurance must operate without content inspection, typically through account-level or device-level signaling.

Gaming and virtual-world platforms — include interactive environments with in-app purchases and chat; suitable controls combine guardian consent, payment-based age hints, and optional network-assisted filtering for external links.

App stores and distribution platforms — act as aggregation points enforcing developer compliance and can apply uniform age labels or assurance tokens to downstream services.

Beyond these general-use categories, several specialized domains require tailored approaches:

Adult-only or restricted-commerce services — high-assurance verification with minimal data disclosure.

Governmental and public-sector services — identity-bound by default within existing eID frameworks.

Essential-rights services (banking, health, news, education) — broad accessibility and minimal friction.

Core Internet infrastructure (DNS, routing, encryption) — privacy-preserving and neutral, supporting only optional signaling chosen by the user or guardian.

Focusing the analysis on general-use platforms helps clarify where architectural guidance is most needed. Specialized or regulated sectors already have established compliance channels, whereas mixed-audience platforms face the hardest trade-offs between safety, usability, and privacy.

6. Concluding recommendations

* Reducing harm to children on the internet requires an incremental, all-hands approach and cannot be solved by age verification alone. A holistic approach would embrace privacy by design and data minimization principles that protect children as well as adults from platform overreach.

* If lawmakers are in a position to outlaw internet services for users of ages under 18 years old, they are in a position to define new credentialing systems that are fit for purpose rather than rely on hard documentation meant for operating automobiles, crossing national borders or social services entitlements.

* Introducing additional age-based signup requirements would risk harm to user privacy and free expression for all users of the web, not just to children, but especially to children.

* To date, jurisdictions around the world are considering various potential age-gating alternatives to mitigate potential safety risks to children without negatively impacting all social media users and without unduly compromising user privacy have reported evidence that advanced technical approaches to verify and assure age are currently infeasible or have significant downsides as compared to status quo approaches of self-reporting.

* Any mandate of age verification effectively regulates all users, rather than companies, to clear a compliance bar in which they must verify their age to the service to use an app. This approach does not address the central thrust of the problem statement, which seems to be that social media companies build platforms that are inclusive to children. Age assurance also regulates all users but has a lower bar and reduced friction for compliance making it a more inclusive choice overall.

* Content moderation of user generated content by platforms and services continues to be an established and effective way to ensure unlawful and disallowed content and behaviour is detected or reported and actioned with proper recourse and remedy mechanisms in the case of overreach.

* A more resilient approach may rely on a plurality of mechanisms operating at different layers of the Internet architecture— service, device, and network— each limited in scope and aligned with privacy-by-design principles. In this model, no single actor holds or processes all user information; rather, complementary methods (for example, self-attestation, trusted-service assurance, or privacy-preserving network-assisted filtering) can contribute to age-appropriate access control according to local regulation and user choice. Such diversity of methods can improve overall robustness and inclusiveness while reducing dependence on any single trust anchor. Contradicting data-protection principles like minimization means that if widely implemented without such safeguards, age-verification systems could still result in mass data collection on both adults and children, with far-reaching implications for user privacy and safety.

7. Security considerations

In general the cross-platform and over-the-wire exchange of information described in nearly all of the architectural choices above implicate security risks due to the complexity of the requirements, cooperation between several different parties and the expectation that this would be done at scale, for all users, not perhaps naively assumed just for youth.

When security tools are considered services that need age-gating such as in proposals to not allow youth to use end-to-end encryption this puts them at great risk and would never be supported by security considerations. Nor would age-gating of encryption be possible without some kind of intervention akin to backdooring encryption.

Hard document review If instantiated, such measures would require the collection, processing, storage and securing of sensitive personal data from all users, including minors, which increases the risk of harm in the case of data breaches. (Privacy International, The Sustainable Development Goals, Identity, and Privacy: Does their implementation risk human rights? (Aug. 29, 2018), https://privacyinternational.org/long-read/2237/sustainable-development-goals-identity-and-privacy-does-their-implementation-risk) This processing and storage are vectors for mass and targeted surveillance by any State jurisdiction party to the UN Convention on Cybercrime.

Guardian and parental controls This approach leaves the responsibility for age verification to parents, which can be fraught in some of the most acute cases of child abuse. [[Lily Hay Newman, Apple Kills Its Plan to Scan Your Photos for CSAM. Here’s What’s

Next, WIRED (Dec. 7, 2022), https://www.wired.com/story/apple-photo-scanning-csam-communication-safety-messages/ (https://www.wired.com/story/apple-photo-scanning-csam-communication-safety-messages/)]] In short, outside the context of mandatory age verification, these mechanisms can be used as tools for abusive parental or guardian surveillance.

If the surveillance power is given to the state instead, encouraging oversight by a potentially abusive regime is of real concern. Neither state nor parental control addresses in-person abuse of a child and technology can not solve that societal issue.

8. Privacy considerations

Risks previously mentioned fundamentally challenge data-minimization efforts and compliance with established data-protection frameworks such as the EU General Data Protection Regulation (GDPR), which requires that personal data be collected only for specific purposes and limited to what is strictly necessary for those purposes.

General Cross-platform tracking with methods that require all intermediary methods described above, spreading this information around rather than minimizing it.

In general, any mandatory age verification will technically enforce limitations of being anonymous online and the right to access resources on the web without being bound to a general or long term identification process over time, which have implications for human rights.[[cite https://undocs.org/A/HRC/2932]]

With third-party verification the providers involved notwithstanding, the centralized risk of data being resold or mishandled increases.

Guardian and parental controls The main concern with these parental controls features is that they enable use but potentially, depending on how they are designed, surveilled use, which may harm children and teens by creating confusion about their privacy and autonomy. On the one hand they may feel a false sense of privacy and that their activities are insulated from the platform, and on the other hand they might self-censor out of concern for the oversight that the parental controls provide their caretakers. These proposals take a narrow view of parent-child relationships and fail to consider the harms as described briefly by CDT: “In particular, LGBTQ youth and children in abusive homes are especially vulnerable to injury and reprisals, including from their parents or guardians, and may inadvertently expose sensitive information about themselves or their friends to adults, with disastrous consequences.” [[cite CDT: Apple’s

Changes to Messaging and Photo Services Threaten Users' Security and Privacy, (Aug. 5, 2021), https://cdt.org/press/cdt-apples-changes-to-messaging-and-photo-services-threaten-users-security-and-privacy (https://cdt.org/press/cdt-apples-changes-to-messaging-and-photo-services-threaten-users-security-and-privacy)]]

9. Human rights considerations

There is also significant agreement amongst the civil liberties and human rights communities that age verification poses more peril than the promise because of, "the ways in which they are often inaccurate; can be circumvented; present privacy and security risks; and may be entirely inaccessible to certain groups, including undocumented immigrants, unbanked individuals, people with disabilities, and others who either do not have access to government ids or who might be more commonly misidentified by biometric technology." [[cite Ruane, Branum, Doty, Jain, CDT Files Amicus Brief in Free Speech Coalition v. Paxton, Challenging TX Age Verification Law, Center for Democracy and Technology (Sept. 23, 2024), https://cdt.org/insights/cdt-files-amicus-brief-in-free-speech-coalition-v-paxton-challenging-tx-age-verification-law/ (https://cdt.org/insights/cdt-files-amicus-brief-in-free-speech-coalition-v-paxton-challenging-tx-age-verification-law/)]]

The Universal Declaration of Human Rights is fundamental to designing technical means of age-gating but also whether and how these means are implemented. [[RFC 8280]] Prior, we have addressed privacy in the previous section. Additionally are considerations for free expression and free association. Economic, Social and Cultural rights are also important to consider as these include the right to personhood eg hard documents issued by the state; but also how age-gating might impact a variety of aspects of life for young people in the digital age.

9.1. Free expression

Any content gating risks limiting lawful access to information, disproportionately affecting most marginalized people's ability to engage in political, educational, and artistic discourse. Overbroad implementation risks chilling participation in online spaces that are essential for learning, advocacy, and identity formation, undermining Article 19 of the UDHR.

9.2. Free association

Mandatory identity checks for access to online services can deter participation in communities and movements that rely on pseudonymity for safety— such as youth networks, LGBTQ+ forums, or activist groups. This threatens the rights to assembly and association under Articles 20 and 23 of the UDHR by forcing users to trade anonymity for access.

9.3. Social Rights

Right to social security and social protection — Systems may require digital identity or proof-of-age to access benefits, which can exclude those without ID or those in marginalized groups (ICESCR art. 9). Think about kids "in the system" or as "wards of the state" here.

Right to work and just conditions — Young workers might face barriers if age-verification systems are used in hiring or platform work, e.g., gig apps that require invasive ID scans. Kids 14-16 and up can often work in most places in the world, where bossware and other measures are increasingly in place.

Right to health — Age-gating can affect access to online sexual and reproductive health information, mental-health support forums, harm-reduction services, or LGBTQ+ youth resources.

9.4. Cultural Rights

Right to participate in cultural life and access information — Content filters and strict age gates may overblock art, literature, or cultural expression, especially where sexuality or gender diversity is part of culture (ICESCR art. 15).

Scientific progress and its benefits — Excessive identification hurdles can chill participation in online learning, open science communities, or software sharing.

9.5. Economic Rights

Right to education and vocational training — Age verification that requires credit cards, government IDs, or costly processes can exclude minors or low-income students from MOOCs or online courses (ICESCR art. 13).

Right to enjoy the benefits of one’ s own creative work — Young creators and small businesses can be locked out of platforms if compliance costs are high or verification is inaccessible. There are a lot of famous kids on the internet making a lot of money from sponsorships and ads and this should be democratized and equally accessible globally.

Non-discrimination in economic life — Systems that assume everyone has a passport, bank account, or biometric record can indirectly discriminate against migrants, refugees, undocumented people, or low-income families.

10. IANA Considerations

This document has no IANA actions.

11. Informative References

[EFF-Age-Terminology]
Electronic Frontier Foundation (EFF), "Age Assurance, Estimation, Verification— Oh My! A Guide to the Terminology", October 2025, <https://www.eff.org/deeplinks/2025/10/age-verification-estimation-assurance-oh-my-guide-terminology>.

Acknowledgments

TODO acknowledge.

Authors’ Addresses

Mallory Knodel
Social Web Foundation
Email: mallory.knodel@nyu.edu

Gianpaolo Angelo Scalone
Vodafone Group
Email: gianpaolo-angelo.scalone@vodafone.com

Tom Newton
Qoria
Email: tom.newton@qoria.com