

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 20 August 2026

M. Nottingham

M. Thomson
16 February 2026

Report from the IAB/W3C Workshop on Age-Based Restrictions on Content
Access
draft-iab-agews-report-02

Abstract

The Workshop on Age-Based Restrictions on Content Access was convened by the Internet Architecture Board (IAB) and World Wide Web Consortium (W3C) in October 2025. This report summarizes its significant points of discussion and identifies topics that may warrant further consideration and work.

Note that this document is a report on the proceedings of the workshop. The views and positions documented in this report are those of the workshop participants and do not necessarily reflect IAB or W3C views and positions.

About This Document

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

The latest revision of this draft can be found at <https://intarchboard.github.io/draft-iab-agews-report/draft-iab-agews-report.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-iab-agews-report/>.

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

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 20 August 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	3
1.1. Views Expressed in this Report	3
1.2. Chatham House Rule	4
2. Overview of the Workshop	4
3. Key Takeaways	6
3.1. There is a Need for Cross-Cutting Collaboration	7
3.2. Identifying the Roles Involved is Important	7
3.3. A Common Vocabulary is Necessary	8
3.4. Privacy and Trust Expectations Need Further Discussion	9
3.5. More than One Approach will be Required	10
3.6. Mapping the Risks for Architectures is a Useful Next Step	11
3.7. Safety Requires More than a Technical Solution	14
4. Security Considerations	14
5. IANA Considerations	15
6. Informative References	15
Appendix A. Workshop Agenda	15
A.1. Topic: Introduction	15
A.2. Topic: Setting the Scene	16
A.3. Topic: Guiding Principles	16
A.4. Topic: Potential Impacts	16
A.5. Topic: Where Enforcement Happens	17
A.6. Topic: Available Techniques	17
A.7. Discussion	17
A.8. Summary and Reflection	17
A.9. Outcomes	17
Appendix B. Workshop Attendees	17
Appendix C. Potential Impacts	19

C.1. Impact on Children	19
C.2. Ecosystem Impact	20
C.3. Implementation and Deployment Difficulties	20
C.4. Security and Privacy	21
C.5. Equity	21
C.6. Societal Impacts	21
Appendix D. Desirable and Essential Properties of a Solution . .	22
D.1. Functional	22
D.2. Accountability and Transparency	22
D.3. Privacy and Security	23
D.4. Equity	23
D.5. Jurisdiction and Geopolitical	23
D.6. Usability	24
D.7. Implementation and Deployment	24
D.8. General / Other	24
Authors' Addresses	25

1. Introduction

Regulators and legislators around the world are increasingly restricting what can be made available to young people on the Internet, in an effort to reduce the online harms that they encounter.

In October 2025, the Internet Architecture Board and the World Wide Web Consortium convened the Workshop on Age-Based Restrictions on Content Access. It brought together technologists, civil society advocates, business interests, and government stakeholders to discuss the nuances of the introduction of such measures.

The primary focus was "to perform a thorough examination of the technical and architectural choices that are involved in solutions for age-based restrictions on access to content", with a goal of "build[ing] a shared understanding of the properties of various proposed approaches."

See the workshop announcement [ANNOUNCE] for details; papers and presentation materials are linked from the announcement. This report summarises the proceedings of the workshop.

1.1. Views Expressed in this Report

This document is a report on the proceedings of the workshop. The views and positions documented in this report were expressed during the workshop by participants and do not necessarily reflect the IAB's or W3C's views and positions, nor those of all participants.

Furthermore, the content of the report comes from presentations given by workshop participants and notes taken during the discussions, without interpretation or validation. Thus, the content of this report follows the flow and dialogue of the workshop but does not attempt to capture a consensus.

1.2. Chatham House Rule

Participants agreed to conduct the workshop under the Chatham House Rule [CHATHAM-HOUSE], so this report does not attribute statements to individuals or organizations without express permission. Most submissions to the workshop were public and thus attributable; they are used here to provide substance and context.

Appendix B lists the workshop participants, unless they requested that this information be withheld.

2. Overview of the Workshop

The IAB/W3C workshop on Age-Based Restrictions on Content Access brought together a diverse group of participants from technical, policy, regulatory, and research communities to examine how the Internet might accommodate demands for age-based access controls. Over three days, discussions traversed the intersection of technology, governance, human rights, and social expectations, with a recurring emphasis on privacy, accountability, and the preservation of the open architecture of the Internet.

The workshop began with a framing session that emphasized the Internet's original design as a universal, non-segmented space. Participants observed that the web does not natively distinguish between adult and child users, and that governments are creating regulatory environments that shift responsibility from parents and individuals to service providers. The scope of discussion was tightly defined: not the morality or policy of age restrictions, but the technical, architectural, and human-rights implications of enforcing them. The challenge, many participants agreed, lay in building mechanisms that are accurate, respect privacy, maintain global interoperability, and avoid creating infrastructure that could be repurposed for censorship or surveillance.

Early exchanges focused on terminology and scope: whether "age verification" should be understood narrowly as identity checking or more broadly as "age assurance." The conversation also touched on the diversity of cultural expectations about parental authority and the variety of legal frameworks emerging across jurisdictions. Some participants warned of "slippery slope" effects, where mechanisms designed for age checks might evolve into tools for broader identity

enforcement. Several noted that while liability drives many policy decisions, technical design should aim to minimize harm and avoid over-centralization. The question of who bears responsibility for child safety -- platforms, regulators, or device manufacturers -- surfaced repeatedly.

Human-rights principles were foregrounded as a basis for evaluation. Privacy was discussed not only in terms of data protection law (including techniques like minimization) but as protection from unwanted exposure or interaction. Freedom of expression and opinion was considered, particularly how adults and children both have rights to communicate, access information and associate, free from the chilling effects of surveillance or discrimination. The group revisited long-standing Internet design tenets, such as decentralization and the end-to-end principle, asking how they should inform modern architectures that could easily drift toward central control. Some argued that successful systems must remain open, interoperable, and reversible, while others cautioned that any solution -- even a well-intentioned one -- would inevitably reshape the Internet's social and economic balance.

Technical sessions explored a spectrum of enforcement models: service-based, network-based, and device-based. Service-enforced systems place the compliance burden on websites, risking fragmentation and user fatigue from repeated verification flows. Network-based filtering -- already common in some jurisdictions -- offers broad coverage but limited accuracy and significant privacy trade-offs. Device-enforced models, in which operating systems mediate access based on a one-time verification, were praised for their potential usability and consistency but criticized for potential concentration of power among major vendors. Many participants noted that a pluralistic approach is more likely to be successful, recognizing that no single architecture can meet all requirements equally across jurisdictions.

Privacy-enhancing technologies (PETs) such as anonymous credentials and zero-knowledge proofs were discussed as promising, though not necessarily sufficient, tools. In particular, PETs don't address all privacy concerns, and likewise don't address wider issues around access to underlying sources of truth. Furthermore, some participants cautioned that PETs cannot prevent circumvention or censorship, are relatively untested, and that open-sourcing code does not automatically make systems trustworthy. A recurring concern was that while credential-based verification may work well in countries with unified ID systems, it risks excluding people without access to such credentials and entrenching inequalities.

Discussions on parental controls and network operator roles highlighted practical tensions between effectiveness, usability, and user rights. Although some participants saw value in layered approaches combining device, service, and network measures, others noted the high complexity and low adoption of parental-control tools even where available. The workshop also revisited the ethical dimension: whether designing better tools might unintentionally legitimize over-broad or intrusive regulation.

By the third day, participants reflected on the need for collaboration across disciplines and institutions. Many acknowledged that while complete solutions are unlikely in the short term, articulating shared vocabulary, architectural roles, and evaluation properties was an essential foundation. There was broad agreement that future work should map risks against possible architectures, document trade-offs in neutral terms, and communicate clearly with policymakers to prevent outcomes that could undermine Internet openness.

The meeting closed with reflections on what process might be followed to take proposed solutions through a standards process. Both IETF and W3C representatives outlined how exploratory work might proceed within their respective frameworks, stressing that standardization would require consensus, open participation, and time.

While this workshop would not provide specific standards proposals or take positions on the advisability of regulatory proposals, it was suggested that leadership bodies, including the Internet Architecture Board and Technical Architecture Group, could be places to make such statements.

While the current status quo -- piecemeal, opaque, and often privacy-eroding -- was unsatisfactory to most participants, many cautioned that hasty solutions could entrench worse problems. This led to growing recognition that protecting children online must not come at the expense of the Internet's foundational freedoms, and that sustained, multi-stakeholder collaboration is the only viable path forward.

3. Key Takeaways

This section highlights aspects of discussion at the workshop that appeared to be most impactful.

3.1. There is a Need for Cross-Cutting Collaboration

Many participants remarked that the workshop allowed them to appreciate perspectives that they had not fully considered previously. Although several substantial efforts have included industry, civil society, government, and technologists, collaboration across all stakeholders appears to be rare.

This was especially evident when considering the involvement of the technical community. Although there have been a number of consultations by governments and other bodies, involvement of the technical community is often limited to participation by the policy representatives of tech companies. This can lead to an underappreciation of the architectural impact and related harms of the design decisions made.

Architectures effective for the goals and less likely to have profound harmful consequences may require the cooperation of multiple actors fulfilling different Section 3.2. To that end, standardization may be especially important for interoperable, collaborative development of architectures involving both servers and clients.

Some participants also noted that approaches where liability rests only on one party -- for example, a content or platform provider -- are unlikely to lead to the desired results, because this creates disincentives for the cooperation that is necessary for meaningful reduction of harms. An approach that considers the roles of the young, their parents, device manufacturers, operating system vendors, content providers, and society overall was believed to be more likely to succeed.

3.2. Identifying the Roles Involved is Important

One of the more substantive discussions on architecture involved presentations on the functional roles involved in any system [HANSON].

Four key roles were identified:

Verifier: The verifier role determines whether a person falls into a target age range.

Enforcer: The enforcer is responsible for ensuring that a person who does not satisfy the verifier is unable to access age-restricted content or services.

Policy selector: The policy selector is responsible for determining

which policies should apply to the user, based on their jurisdiction, status or preferences.

Rater: The rater is responsible for determining whether content or services require age restrictions and the age ranges that apply.

In addition, it was noted that ratings and laws are often limited by geography or jurisdiction, so it is often necessary for services to first identify the applicable jurisdiction. It was generally accepted that this function often uses IP geolocation mappings, despite acknowledged limitations around accuracy and susceptibility to circumvention using VPNs.

3.3. A Common Vocabulary is Necessary

Early discussions highlighted how not all participants used the same terminology when referring to different activities or functions. There was a recognition of the value of shared language, and some participants pointed to [ISO-IEC-27566-1], which establishes key terms, including:

Age assurance: Age assurance is an umbrella term for technology that provides some entity with information about the age of a person. This is understood to encompass multiple classes of specific methods, including age verification, age estimation, and age inference. Age assurance does not need to result in a specific age; age ranges are often preferred as these can have better privacy properties.

Age verification: Age verification refers to gaining high assurance that a person is within a given age range. Strong assurances are often tied to official or governmental documentation, so age verification can involve the use of government-issued digital credentials.

Age estimation: Age estimation uses statistical processes that process physical or behavioral characteristics of a person to produce a probabilistic value for how old someone is or whether their age is in a target range. A variety of techniques are used, the most common being facial age estimation, which uses machine learning models to estimate how old a person is based on still or moving images of their face.

Age inference: Age inference draws on data sources to determine

whether a person fits a given age range. This method can require identification information, such as an email address or phone number, to find relevant records. For example, evidence of online activity prior to a certain date in the past might support the view that a person is older than a target threshold.

Age gating: Age gating is the process of restricting access to something based on the age of the person requesting access.

Relating these functions to the roles described in Section 3.2, all age assurance types fit the "verifier" role, where age gating applies to the "enforcer" role.

3.4. Privacy and Trust Expectations Need Further Discussion

Privacy was a recurrent theme at the workshop, but it was clear that there are multiple considerations at play when talking about privacy. The question of privacy was often caught up in discussions of trust, where approaches each depend on different sorts of trust between the different actors.

Participants identified privacy as important to maintaining trust in any system that involves age assurance or age gating.

Where private information is used by the actors in a proposed architecture, those actors might need to be trusted to handle that private information responsibly. In that approach, the importance of different safeguards on personal information, such as the prompt disposal of any personal information -- a practice that many age verification providers promise -- becomes a core part of what might allow people to trust that system.

Several people observed that the sort of trust that is asked from people might not correspond with the role that certain entities play in people's lives. This will depend on context, where "adult" content providers generally serve anonymous users, whereas social media often already has a lot of personal information on users.

In either case, users might have no prior knowledge of -- or trust in -- providers that are contracted to provide age assurance functions. It was observed that one likely consequence of some arrangements is to train people to become more trusting of strange sites that ask for personal information.

Alternatively, it might be that trust in the system is not vested in actors, but instead the system as a whole. This is possible if no information is made available to different actors, removing the need to trust their handling of private information. For this to be

achievable, the use of zero-knowledge proofs or similar cryptographic techniques was seen as a way to limit what each entity learns. Some participants noted, however, that these techniques do not address circumvention or censorship risks, still introduce new information into the ecosystem, and may concentrate trust in particular software implementations.

Other aspects of trust were considered equally important from different perspectives. Services that rely on an independent age assurance provider need to trust that the provider makes an accurate determination of age, at least to the extent that they might be held liable in law. They also need to trust that the service respects privacy, lest the use of a low-quality provider could create other forms of liability or drive away potential customers.

3.5. More than One Approach will be Required

A recurrent theme in discussion was the insufficiency of any particular age assurance technique in ensuring that people are not unjustifiably excluded. All age assurance methods discussed fail to correctly classify some subset of people:

- * Age verification that depends on government-issued credentials will fail when people do not hold accepted credentials. This includes people who do not hold credentials and those who hold credentials, but not those that are recognized.
- * Age estimation produces probabilistic information about age that can be wrong by some number of years, potentially excluding people near threshold ages. This manifests as both false acceptance (people who are outside the target age range being accepted) and false rejection (people who are in the target age range being rejected). Where there is a goal of minimizing the false acceptance rate, that increases the number of false rejections.
- * Age inference techniques can fail due to lack of information.

Discussion often came back to an approach that is increasingly recommended for use in age verification, where multiple methods are applied in series. Checks with lower friction -- those that require less active participation from people -- or that are less invasive of privacy are attempted first. Successive checks are only used when a definitive result cannot be achieved.

Some participants noted that inconsistent friction and invasiveness create a different kind of discrimination, one that can exacerbate existing adverse discrimination. For example, the accuracy of age estimation for people with African heritage is often significantly

lower than for those with European ancestry [FATE]. This is attributed to the models used being trained and validated using datasets that have less coverage of some groups. People who are affected by this bias are more likely to need to engage with more invasive methods.

One consequence of having multiple imperfect techniques is the need to recognize that any system will be imperfect. That creates several tensions:

- * Some people will never be able to satisfy age assurance checks and will therefore be excluded by strict assurance mandates. Here, discussions acknowledged that purely technical systems are likely inadequate.
- * Some people who should be blocked from accessing content or services will find ways to circumvent restrictions. In this context, the term "advanced persistent teenager" was recognized as characterizing the nature of the "adversary": individuals who are considered too young to access content, but who are highly motivated, technically sophisticated, and have time to spare.
- * Offering more choices to people can improve privacy because they get to choose the method that suits them. However, when a chosen method fails, having to engage with additional methods has a higher privacy cost.

Some participants argued that accepting these risks is necessary in order to gain any of the benefits that age-based restrictions might confer. Other participants were unwilling to accept potential impositions on individual rights in light of the insufficiency of restrictions in providing meaningful protection; see Section 3.7.

3.6. Mapping the Risks for Architectures is a Useful Next Step

How the identified roles (see Section 3.2) are arranged into architectures was some of the more substantive discussion. [JACKSON] describes some of the alternatives, along with some of the implications that arise from different arrangements.

Throughout this discussion, it was acknowledged that active deployments tended to fall into a common pattern, where content providers are required to age-gate access and contract a third party to interpose that service. Several participants noted that this is a somewhat natural consequence of some of the constraints that actors are subject to. Figure 1 shows the typical deployment model for age-gated content and services, along with the roles from Section 3.2.

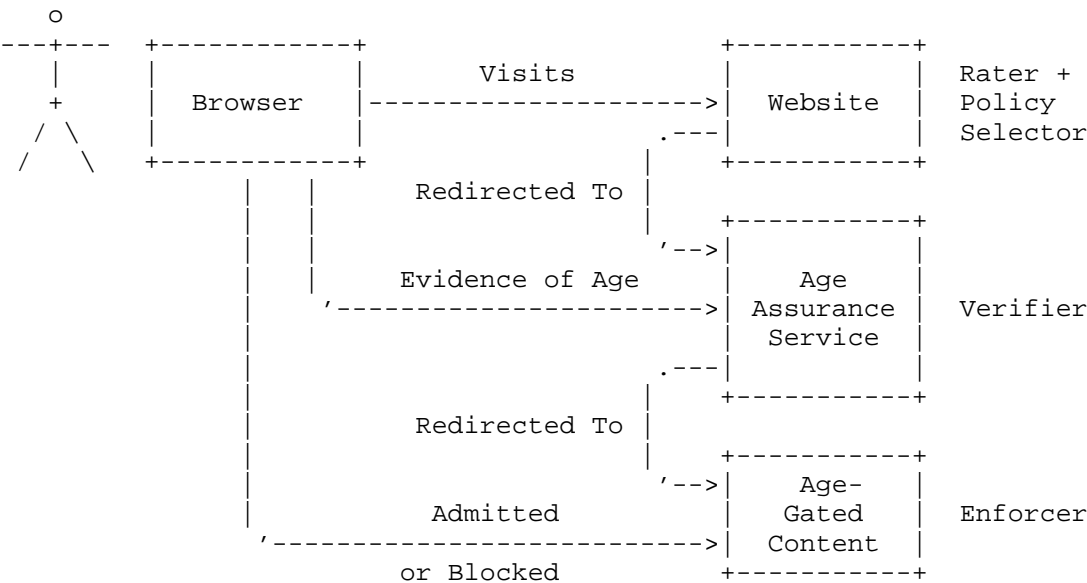


Figure 1: Typical Deployment Model

Some participants also noted that certain approaches may carry higher path-dependence risk once widely deployed, even if they remain theoretically possible to withdraw or replace. This can arise from accumulated architectural dependencies, operational integration with third-party services, and evolving expectations among users and service providers. As a result, architectures that tightly couple functionality to external verification services or embed assumptions about routine age signalling may increase the practical cost of transition, should alternative approaches later emerge that address privacy, equity, or effectiveness concerns more effectively.

Figure 2 shows a deployment model for parental control software, showing how the roles from Section 3.2 might apply. Here, parental controls do any verification of age necessary and select policies; content ratings might be performed by websites or the parental control software on the device; enforcement is performed on-device.

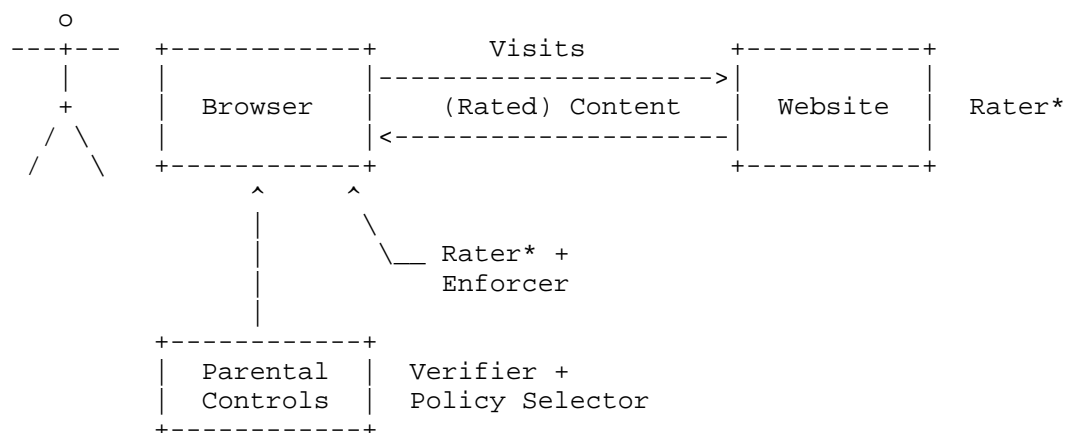


Figure 2: Parental Control Deployment Model

An observation was made that laws often seek to designate a single entity as being responsible for ensuring that age restrictions are effective. That lawmakers feel the need to designate a responsible entity is due to constraints on how laws function, but one that creates other constraints.

Another constraint identified was the need for specialist expertise in order to administer all of the multiple different age assurance techniques; see Section 3.5. This means that there is a natural tendency for services to contract with specialist age assurance services.

Some of the proposed architectures were better able to operate under these constraints. Others required greater amounts of coordination, further emphasizing the importance of collaboration identified in Section 3.1.

In discussion of the constraints on different architectures, it was common for participants to point to a particular aspect of a given approach as carrying risks. Indeed, the final reckoning of risks produced a long list of potential issues that might need mitigation (see Appendix C).

Architectures are not equally vulnerable to different risks, so a more thorough analysis is needed to identify how each risk applies to a different approach. An analysis that considers the constraints and assumptions necessary to successfully deploy different architectures is a contribution that would likely be welcomed by the community.

3.7. Safety Requires More than a Technical Solution

Experts in child safety frequently acknowledged that restricting access to selected content cannot be assumed to be sufficient. The task of ensuring that children are both kept appropriately safe, while preparing them for the challenges they will face in their lifetimes, is a massively complex task.

A recurrent theme was the old maxim, "it takes a village to raise a child". This concept transcends cultural boundaries and was recognized. The role of parents, guardians, educators, governments, and online services in creating an environment in which children can thrive and grow.

Content and service restrictions are likely only a small part of a suite of actions that combine to provide children with protection, but also support and encouragement. This theme was raised several times, despite the goal of the discussion being to explore technical and architectural questions.

Restrictions are necessarily binary and lacking in nuance. Though questions of what to restrict were out of scope for the workshop, discussions often identified subject matter that highlighted the challenges inherent in making simplistic classifications. Participants acknowledged the importance of the role of the adults who support children in their life journey. For example, on the subject of eating disorders, which can be challenging to classify, participants pointed to the importance of being able to recognize trends and inform and engage responsible adults. Ultimately, each child has their own challenges and the people around them are in the best position to provide the support that best suits the child.

The concept of age-appropriate design was raised on several occasions. This presents significant privacy challenges in that it means providing more information about age to services. However, it was recognized that there are legal and moral obligations on services to cater to the needs of children of different age groups. This is a more complex problem space than binary age restrictions, as it requires a recognition of the different needs of children as they get older.

4. Security Considerations

Age verification has a significant potential security impact upon the Internet; see Section 3.4.

5. IANA Considerations

This document has no IANA actions.

6. Informative References

[ANNOUNCE] Internet Architecture Board, "IAB/W3C Workshop on Age-Based Restrictions on Content Access (agews)",
<<https://datatracker.ietf.org/group/agews/about/>>.

[CHATHAM-HOUSE]
Chatham House, "Chatham House Rule",
<<https://www.chathamhouse.org/about-us/chatham-house-rule>>.

[FATE] Ngan, M., Grother, P., and A. Hom, "Face Analysis Technology Evaluation (FATE)",
<<https://nvlpubs.nist.gov/nistpubs/ir/2023/NIST.IR.8491.pdf>>.

[HANSON] Hanson, J., "Where Enforcement Happens",
<<https://datatracker.ietf.org/doc/slides-agews-slides-where-enforcement-happens/>>.

[ISO-IEC-27566-1]
ISO, "Information security, cybersecurity and privacy protection — Age assurance systems Part 1: Framework",
<<https://www.iso.org/standard/88143.html>>.

[JACKSON] Jackson, D., "Where Enforcement Happens",
<<https://datatracker.ietf.org/doc/slides-agews-where-enforcement-happens/>>.

Appendix A. Workshop Agenda

The following sections cover each of the main topics of discussion sessions.

A.1. Topic: Introduction

We will launch the workshop with a greeting, a round of introductions, and an explanation of the terms of engagement, background, goals and non-goals of the workshop.

A.2. Topic: Setting the Scene

Successfully deploying age restrictions at Internet scale has many considerations and constraints. We will explore them at a high level in order. The goal is to discuss within the group about the scope of topics that the workshop will seek to address.

A.3. Topic: Guiding Principles

Architectural principles give us a framework for evaluating additions and changes to the Internet. Technical principles are subject to a number of other considerations, in particular human rights principles. We will review the principles that might apply to age-based restrictions, explain their function, impact, and how they are applied. Including human rights impacts, such as:

- * Privacy and Security
- * Safety and efficacy
- * Censorship and Access
- * Access to the Internet
- * Freedom of Expression

And effects on the internet and web architecture, such as:

- * Deployment, Extensibility and Evolution
- * Avoid Centralization
- * End-to-End
- * One Global Internet/Web
- * Layering and Modularity

A.4. Topic: Potential Impacts

We now want to look at some of the higher-level considerations that apply regardless of approach. We will look at some different perspectives on how to think of the overall problem. Discussion will seek to find how those perspectives can be shaped to guide choices.

A.5. Topic: Where Enforcement Happens

The Internet standards community is in the unique position to make controlled changes to the architecture of the Internet, and so there are multiple ways and places to deploy age restrictions. We will examine the options, with an eye to the deployment properties of each location and configuration, as related to the architectural principles. In particular, it will consider the establishment of new roles as well as the use of existing ones.

A.6. Topic: Available Techniques

There are several active and proposed systems for age restriction on the Internet. We will review them from the perspective of their interaction with the architectural principles, potential impacts, and with consideration of the enforcement options. Including:

- * Age verification: including server-side solutions using government identity systems / ZKP
- * Age estimation: including biometrics and data analysis
- * Age "inference" approaches
- * In-network solutions
- * Classification and On-device / parental control designs

A.7. Discussion

We will follow up on incomplete discussions and revisit architectural learnings.

A.8. Summary and Reflection

We will summarise what we have discussed and learned thus far.

A.9. Outcomes

We will outline the potential outcomes, further actions and next steps.

Appendix B. Workshop Attendees

Attendees of the workshop are listed with their primary affiliation. Attendees from the program committee (PC), the Internet Architecture Board (IAB), and W3C Technical Architecture Group (TAG) are also marked.

- * Steve Bellovin
- * Hadley Beeman, TAG (PC)
- * Matthew Bocci, IAB (Observer)
- * Christian Bormann, SPRIND
- * Marcos Cceres, TAG (Observer)
- * Andrew Campling, 419 Consulting
- * Sofa Celi, Brave
- * David Cooke, Aylo
- * Iain Corby, Age Verification Providers Association
- * Dhruv Dhody, IAB (Observer)
- * Nick Doty, Center for Democracy and Technology (PC)
- * Sarah Forland, New America Open Technology Institute
- * Jrme Gorin, Ecole Polytechnique
- * Alexis Hancock, Electronic Freedom Foundation
- * Julia Hanson, Apple
- * Wes Hardaker, University of Southern California Information Sciences Institute
- * Kyle den Hartog, Brave
- * Dennis Jackson, Mozilla
- * Leif Johansson, SIROS Foundation
- * Mallory Knodel, Article 19
- * Mirja Khlewind, IAB (Observer)
- * Jonathan Langley, Ofcom UK
- * Veronica Lin, Carnegie Mellon University
- * Thibault Meunier, Cloudflare

- * Tom Newton, Qoria
- * Mark Nottingham, IAB (PC Co-Chair)
- * Georgia Osborn, Ofcom UK
- * Tommy Pauly, IAB (PC)
- * John Perrino, Internet Society
- * Eric Rescorla, Knight-Georgetown Institute
- * Beatriz Rocha, Ceweb.br
- * Omari Rodney, Yoti
- * Gianpaolo Scalone, Vodafone
- * Sarah Scheffler, Carnegie Mellon University
- * Andrew Shaw, UK National Cyber Security Centre
- * Aline Sylla, German Federal Commissioner for Data Protection and Freedom of Information
- * Martin Thomson, TAG (PC Co-Chair)
- * Carmela Troncoso, EPFL, the Swiss Federal Institute of Technology in Lausanne
- * Benjamin VanderSloot, Mozilla
- * Tara Whalen, World Wide Web Consortium (PC)

Appendix C. Potential Impacts

During the workshop, participants were asked to name potential impacts -- whether positive or negative -- that could be seen in association with the introduction of an age control mechanism. This list is not exhaustive, focuses largely on the challenges surrounding the introduction of such a mechanism, and does not imply that all points were agreed to by all participants'

C.1. Impact on Children

1. Children encounter online harms
2. Pushing kids to less safe resources

3. Kids lose the ability to explore on their own
4. Diminishing children's rights

C.2. Ecosystem Impact

1. Centralization
2. Fragmentation of the Internet
3. Increased costs for running a Web site
4. Chilling effects on use of the Internet
5. VPNs proliferate
6. Chilling effects on the publication of borderline content
7. Less content being available online
8. Restricting people to a few platforms / services
9. More use/utility of the Internet due to a perception of safety
10. More (or all) online services require a verified login

C.3. Implementation and Deployment Difficulties

1. Device compatibility
2. "Advanced Persistent Teenagers"
3. Difficulties regarding jurisdiction checking
4. Spillover to other software (e.g., VPNs)
5. Displacing users from compliant to non-compliant sites
6. False sense of addressing the problem
7. Dealing with conflict of laws
8. Operators pulling out of territories
9. Increasing the footprint of the deep web
10. Imposition of cultural norms on other jurisdictions

11. Technical solutions are reused for other purposes (scope creep)
12. Dealing with obsolete and non-compliant systems

C.4. Security and Privacy

1. Increased cybersecurity risks
2. Fingerprinting risk
3. Ad targeting could get creepier
4. Needing to trust someone on their word without evidence
5. Normalizing online identity requests -- increase to phishing risk
6. Data breaches

C.5. Equity

1. Lack of access (e.g. due to lack of device support)
2. Refugees, stateless people, people without identity
3. Harm to vulnerable people
4. Not addressing other vulnerable groups (i.e., not age-based)
5. Lack of availability of redress mechanisms
6. Users' rights to restitution
7. Loss of control over and access to data
8. Risk to anonymity
9. Loss of ability to run software of your choice

C.6. Societal Impacts

1. Air cover for blocking the Internet
2. User control of the content they see online
3. Costs to society (e.g., regulatory overhead)
4. Increased online tracking and state surveillance

5. Use as a censorship mechanism
6. Advancing foreign policy goals with censorship
7. Abuse of guardians who don't cut off their wards

Appendix D. Desirable and Essential Properties of a Solution

During the workshop, participants were asked to nominate the properties that they believed would be advantageous or even essential for a solution in this space to have. This set of requirements and desiderata was recognised as not all being achievable, as some goals are in tension with others.

D.1. Functional

1. Underage don't access content that's inappropriate
2. Not trivially by-passable
3. Flexible enough to be provided through different means
4. Bound to the user
5. Reliable
6. Handles user-generated content
7. Enables differential experiences or age-appropriate design (not just blocking)
8. Agile by design -- assume adversarial engagement
9. Difficult to bypass
10. Accurate

D.2. Accountability and Transparency

1. Transparency and accountability regarding what is blocked
2. Minimises the need for trust decisions
3. Can be independently / publicly verifiable and tested
4. Auditability
5. Appeal mechanism for incorrect labelling of content

D.3. Privacy and Security

1. Issuer-Verifier and Verifier-Verifier unlinkability
2. Unlinkability across components
3. Purpose limitation of the data processed
4. Security of data processed
5. Phishing-resistant
6. Doesn't process or transfer any more data than is necessary
7. Avoids becoming a tracking vector

D.4. Equity

1. Inclusive
2. Fair -- avoids or minimises bias
3. Does not create inequalities (e.g., across education, other properties)
4. Discriminates solely upon age, not other properties
5. Works on open devices
6. Device independence
7. Usable by people of all ages to increase their safety online
8. User choice in who verifies their age, and how
9. No clear losers
10. Accessible to people with disabilities
11. Includes appeal mechanisms for incorrect age determinations

D.5. Jurisdiction and Geopolitical

1. Able to handle arbitrary composition of different jurisdictional requirements (possibly down to school level)
2. Applicable globally

3. Applies the rule of law in the jurisdiction where it applies universally
4. No concentration of power in any one entity (or small group of them)
5. No concentration of power in any country
6. Aligned to legal duties
7. Based upon a valid legal basis

D.6. Usability

1. Economically sustainable
2. Low friction for adults
3. Fast
4. Comprehensible by users

D.7. Implementation and Deployment

1. Low dependency on a single root of trust
2. Enforceable by a good mix of technology and law
3. Broad deployability -- not expensive or complex
4. Decentralized
5. Future-proof
6. Ability to report / learn when there are issues in the system / telemetry

D.8. General / Other

1. Not perfect
2. Technically robust
3. Not a single, sole solution
4. Stable -- resilient
5. Alignment of incentives among participants

6. Simple to implement
7. Resistance to repurposing for censorship
8. Unable to be used for surveillance
9. Addresses risk of verification becoming over-prevalent
10. Accountable governance
11. Open Standards-based

Authors' Addresses

Mark Nottingham
Email: mnot@mnot.net

Martin Thomson
Email: mt@lowentropy.net