

Workgroup:	Network Working Group
Internet-Draft:	draft-nottingham-iab-age-restrictions-00
Published:	4 August 2025
Intended Status:	Informational
Expires:	5 February 2026
Author:	M. Nottingham

Architectural Considerations of Age Restriction

Abstract

Around the world, policymakers are considering (and in some cases implementing) age restriction systems for Internet content. This document explores the unwanted impacts on the Internet that these systems are likely to have, and makes recommendations to increase their chances of becoming a successful part of the Internet infrastructure.

The RFC Editor will remove this note

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-nottingham-iab-age-restrictions/>.

information can be found at <https://mnot.github.io/I-D/>.

Source for this draft and an issue tracker can be found at <https://github.com/mnot/I-D/labels/age>.

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 5 February 2026.

Copyright Notice

Copyright (c) 2025 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. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- 1. Introduction
- 2. Risks of Age Restriction
 - 2.1. Centralization
 - 2.2. Privacy and Security
 - 2.3. Barriers to Access
 - 2.4. Fragmentation
 - 2.5. An Age-Gated Internet
- 3. IANA Considerations
- 4. Security Considerations
- 5. Informative References
- Author's Address

1. Introduction

Increasingly, policymakers are proposing and implementing regulation that restricts what content young people can access online. A recurring theme in these efforts is that it is no longer considered sufficient to rely on self-assertions of age, and so stronger guarantees are deemed necessary.

Age restrictions are already deployed on the Internet: for example, some Web sites already require proof of age to create an account. However, when such deployments become more prevalent, they tend to have greater impact upon the Internet architecture, thereby endangering other properties that we depend upon for a healthy online ecosystem. Systems that are designed for deployment in a single, homogenous domain rarely are suitable for the diversity of requirements and considerations that apply to Internet-scale systems.

[Section 2](#) explores the unwanted impacts on the Internet that these systems are likely to have, and makes recommendations to increase their chances of becoming a successful part of the Internet infrastructure.

2. Risks of Age Restriction

This section catalogues risks that age restriction systems might encounter, expressed in terms of the Internet's architectural principles.

2.1. Centralization

One of the defining characteristics of the Internet's architecture is its avoidance of "gatekeeper" roles where a single entity (or small number of them) has unavoidable access to or control over communications. This architecture makes the Internet more resilient and avoids problematic concentrations of power.

Inherently, any age restriction system involves controlling access to content and services on the Internet. When that system's use is legally required, it can be characterized as being centralized. However, that centralization is mitigated (to some degree) by the governance mechanisms that legitimize the legal basis of the system.

However, other forms of centralization can occur in age verification systems that are more avoidable. For example, many proposed age restriction systems require verification by a single party. This is not how Internet-scale services are designed, because of the many risks incurred:

First, that single operator will have access to a significant amount of sensitive information about people's activities on the Internet. Such a system is vulnerable to cyberattacks by organized criminals and nation-state actors. Centralizing so much sensitive data in one place creates an irresistible 'honey pot' for attackers.

Second, that single operator will also have extraordinary access to that same data. Without very strong controls, they will be tempted to leverage that access (e.g., by reselling it, or 'insights' into the data collected). This is especially true if their incentives are not aligned with end users and services (e.g., if they have a profit motive and little opportunity for additional revenue).

Third, any failure -- whether technical or business -- by that single operator creates an outage for all users dependant upon it. If the problem is significant, that outage may extend for a considerable amount of time.

In all of these cases, having only one entity operating the system increases the associated risks dramatically.

Even when multiple parties provide verification services, centralization can emerge if there is too much friction against user switching between them. For example, if Web sites rather than end users select the verification service used, this does not create a market that respects end user preferences; it only respects the self-interest of sites.

Finally, age restriction systems can also have secondary effects that lead to centralization. For example, if an age restriction system requires use of a particular Web browser (or a small number of them), that effectively distorts the market for Web browsers.

Therefore, age restriction systems that are intended to become part of Internet infrastructure MUST:

- Avoid reliance on a single party to provide age verification services
- Provide some mechanism for easy switching between verification services by end users

- Avoid requiring use of an arbitrarily limited set of operating systems, Web browsers, client programs, or other software or hardware

[[CENTRALIZATION](#)] explores these issues greater detail.

2.2. Privacy and Security

It is technically challenging to design an age restriction system that exhibits the security and privacy properties expected of Internet standards. Exposing information about Internet users to third parties -- whether verifiers, services, or others -- creates powerful incentives. In particular, commercial interests desire access to it to be able to track activity across the Internet so that this can be sold (to advertisers, insurers and others), and so they will use (and abuse) any facility that offers such information to learn about what people are doing.

In this context, age restriction systems can introduce several new risks to the Internet.

Most immediately, a system that reveals the age (or birthdate) of Internet users to services on it is a privacy risk. A person's age is an attribute that can be used to discriminate against them without justification, and is legally protected by privacy law in many jurisdictions.

Beyond that immediate risk, verifying services' potential access to personal information creates a powerful incentive for misuse -- whether as part of the business model of the verifying service, or by third parties (such as nation-state attackers).

This is the case when verifying services over-collect such information (for example, age estimation services that use photos and biometrics), and it is also the case when users' activity is exposed to the verifying service when age restriction takes place. The latter risk is similar to the risk of tracking and profiling seen on the Web, which the Internet standards community has expended considerable effort to mitigate (see e.g., [[RFC7258](#)]).

Furthermore, exposing information beyond age to services creates additional privacy and security risks. For example, an age verification system that also exposes the country a person is a citizen of allows sites to discriminate against that attribute, which is beyond the purpose of age restriction.

Finally, even on its own a simple attribute like 'age in years' or 'birthdate' can be used to add entropy to an identifier for the end user (in combination with other information, such as IP addresses, browser characteristics, and so on), creating a new tracking vector when exposed to services that collect such information. See [[TRACKING](#)].

In all cases, the privacy and security of an age restriction system needs to be proven: considerable experience has shown that merely trusting assertions of these properties is ill-founded. Likewise, because data that is not collected cannot be shared or stolen, a system that uses technical means to limit data collection and usage is preferable over one that uses legal means (such as contractual terms).

Therefore, age restriction systems that are intended to become part of Internet infrastructure MUST:

- Avoid over-collection of information by age verifiers
- Avoid sharing information about service usage with age verifiers
- Avoid sharing information other than age information with services
- Minimise the amount of age information shared with services (e.g., using age brackets)

- Provide technical limitations to the collection and use of data
- Be based upon publicly available specifications that have had adequate security and privacy review to the level that Internet standards are held to

See also [[PRIVACY](#)].

2.3. Barriers to Access

Many existant proposals for age restriction systems require its users to have certain capabilities -- for example, a personal smartphone, a government credential, or a camera with certain resolution.

Imposing these requirements means that some number of people will be disenfranchised from full use of the Internet – especially if age restriction becomes pervasive across many services. At the scale of the entire Internet (or even in a single country), this can be a large number of disenfranchised people.

For example, many people only have Internet access from public computers (such as those in libraries), and do not have exclusive or reliable access to a smartphone. Others lack government-issued identity documents that some schemes rely upon.

While such restrictions may be palatable in a closed system (such as on a single platform or in a single jurisdiction), they are not suitable for Internet-wide deployment.

Therefore, age restriction systems that are intended to become part of Internet infrastructure MUST:

- Avoid requiring hardware capabilities not widely available in desktop and mobile computers globally, both in terms of overall performance and specific features
- Avoid relying on a single mechanism for proving age

2.4. Fragmentation

The likelihood of incompatible age restriction systems being deployed in different jurisdictions around the world introduces a risk of fragmentation -- i.e., that the Internet will not work the same way in different places.

Fragmentation is a growing concern for the Internet: various local requirements are creating friction against global deployment of new applications, protocols, and capabilities. As the Internet fragments, the benefits of having a single, globe-spanning networking technology are correspondingly lessened. Although a single factor (such as diverging approaches to age restriction) is unlikely to fragment the Internet on its own, the sum of such divergences increases the risk of fragmentation greatly, risking the viability of the Internet itself.

In the context of age restriction, fragmentation is most concerning if someone were to need to understand and interact with (possibly after some onboarding procedure) a new system for each jurisdiction they visit. This would represent a significant barrier for users who travel, and would also present increased complexity and regulatory burden for businesses, potentially leading to further lack of competitiveness in some industries by increasing costs.

Fragmentation is best addressed by adoption of common technical standards across jurisdictions.

However, it is important to recognise that the mere existence of an international standard does not imply that it is suitable for deployment: experience has shown that voluntary adoption by implementers is important to prove their viability.

Furthermore, standards do not necessarily enforce interoperability and the architectural goals that they espouse: if they allow too much interpretation of their application (for example, through optional-to-implement features, or too-generous extensibility mechanisms), they can fall short of these goals and only provide "air cover" for those who wish to claim standards compliance, without actually serving public interest goals -- thereby leading to fragmentation.

Therefore, age restriction systems that are intended to become part of Internet infrastructure MUST:

- Be based upon internationally recognised, open technical standards
- Be based upon technical standards that are voluntarily adopted by implementers
- Be specified in a manner that provides interoperability and ensures architectural alignment
- Be coordinated across jurisdictions wherever feasible

2.5. An Age-Gated Internet

The Internet is designed to be used without permission, both by servers and clients. Easy-to-use age restriction mechanisms risk creating a 'papers please' Internet, where a credential is required to access large portions of the Internet's services. Such an outcome would amplify the other harms listed.

This risk is heightened if there are incentives for sites to deploy it, such as increased access to non-age data.

Access to more granular age information also heightens many risks, because it makes a restriction system simultaneously useful in a broader variety of cases, and more attractive for misuse, because it offers more information about users.

Therefore, age restriction systems that are intended to become part of Internet infrastructure MUST:

- Make the use of age restrictions visible to end users
- Have a structural disincentive for indiscriminate use of age restriction

3. IANA Considerations

This document has no instructions for IANA.

4. Security Considerations

Age restriction systems undoubtedly have numerous security considerations, should they be deployed.

5. Informative References

[CENTRALIZATION] Nottingham, M., "Centralization, Decentralization, and Internet Standards", RFC 9518, DOI 10.17487/RFC9518, December 2023, <<https://www.rfc-editor.org/rfc/rfc9518>>.

[PRIVACY]

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

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

[TRACKING] W3C TAG, "Unsanctioned Web Tracking", July 2015, <<https://www.w3.org/2001/tag/doc/unsanctioned-tracking/>>.

Author's Address

Mark Nottingham

Melbourne

Australia

Email: mnot@mnot.net

URI: <https://www.mnot.net/>