Device-based Age Verification

Benjamin VanderSloot
Mozilla

August 8, 2025

Websites typically ask visitors to **assert** they are of legal age before serving them adult content. As the EFF identified last year[1], numerous efforts are underway to require websites to instead **verify** their visitor's age. The mechanisms vary, but typically require users to present credentials derived from digitized identity documents or upload selfies for automated age estimation.

Legal requirements for websites to verify visitor ages come with substantial risks for user privacy and open access to the web. Unlike self-assertion, the methods put forward for age verification entail leaking information about the user to third parties, impacting privacy. For any given jurisdiction enacting such requirements, a substantial fraction of websites may be unable to participate and so either choose to self-censor through geo-blocking or be censored by regulators, affecting open access to the web, even between adults.

These issues can be mitigated with an architectural change, moving the enforcement from the website to the device, where it can remain under the control of the device owner. We propose such a system for web browsers, which achieves the same goal of restricting access to inappropriate content with fewer externalities.

We do not take a stance how the verification of age should be done, instead focusing on which party enforces the verification. We take the question of **where** the verification occurs as the critical one, separating it from the question of **how** that verification is performed.

The core of this proposal is to place a responsibility on the browser to respect existing signals from the operating system about the user's content preferences and from websites about the content they contain.

1 Key Points

This proposal is practical today. Many adult websites already label themselves to be Restricted to Adults[2] through an HTML meta tag. Contemporary operating systems like Microsoft Windows, iOS and macOS already expose signals about the user's content preferences to third-party applications through dedicated APIs. This proposal only requires that browsers consume and act on these existing content labels and content preferences.

This proposal is compatible with all verification mechanisms. In contexts where self-asserted age verification is sufficient, this proposal does not introduce the necessity of providing strong age verification. Nor does this proposal preclude cryptographic systems with government

[1] https://www.eff.org/deeplinks/2024/12/global-age-verification-measures-2024-year-review
[2] https://www.rtalabel.org/?content=howto

issuers as an age verification mechanism. Discussions of what mechanism is best for each use case and in each jurisdiction are independent of this proposal.

Website-enforced age verification comes with substantial risks. Performing age verification over the network is costly for websites, time-consuming for end-users, and entails substantial risk to individual's privacy and equitable access to the Internet in all systems used to date.

Device-enforced age verification avoids these risks. Devices already come with systems for managing multiple users and access to content. Handling age verification locally avoids website cost to integrate complex age verification systems for each jurisdiction they serve, doesn't associate private information with website visits, and can be transparent to users.

Device-enforced age verification is more effective. Although an underage user with administrative access to their device can easily defeat this proposal, this is equally true of website-enforced approaches given the popularity of widely-available tools like VPNs. Device-enforced age verification can also cover content beyond the web, e.g. Apple's Sensitive Content Warning[3].

Device-enforced age verification can be improved iteratively. This proposal, while practical today, would not be perfect. However, there are multiple directions for improvement that each provide device owners more control over content displayed on their devices through a one-time configuration.

2 Proposal

We propose a system of content classification that is enforced on user devices according to the owner's preferences, configured at the OS level and based on content labels embedded in websites.

2.1 <meta name="rating">

We propose to standardize the extension[4] of the `meta` tag pioneered by the ASACP[5]. This tag sees use in practice. Of the top million domains observed by HTTPArchive in March 2025, 5,383 domains voluntarily label some content[6].

This requires only the following change to the HTML specification[7]:

> When inserting a `meta` tag to the document, if the `"name"` attribute is `"rating"` and the `"content"` attribute is `"RTA-5042-1996-1400-1577-RTA"` or `"adult"`, and if the browser is configured by the host platform to restrict adult content, the operation must fail and the browser may redirect the user to an error page.

Some detail and extensions may be added, however this is sufficient to convey our intent. This elides how host platform configuration to restrict adult content is performed, including what age verification technique is used, but states clearly that the browser must respect platform signals in combination with self-reported adult content.

[3] https://support.apple.com/en-gb/105071
[4] https://wiki.whatwg.org/wiki/MetaExtensions#:~:text=restricted%20to%20adults
[5] https://www.asacp.org/
[6] https://github.com/bvandersloot-mozilla/restricted-to-adults/blob/main/rta-domains.txt
[7] https://html.spec.whatwg.org

2.2 Potential Extensions

There are additional features that may add value, but are not essential to this proposal:

- Supporting templated error pages that the user navigates to, e.g. `"https://blockpage.example.com?url=$url"`, in order to support device-owner supervision of the failed connections.
- Extending the `"content"` attribute to define lower levels of restriction, e.g. `"teen"`.
- Defining labels descriptive of the content to allow varying age limits in different jurisdictions for the same content, e.g. `"nudity"`, `"gambling"`, `"pornography"`.

2.3 Platform Signals

This proposal puts significant weight on the **existing** platform signals indicating the users' desire for web content to be restricted.

MacOS and iOS, provide WebContentSettings's FilterPolicy[8] for blocking adult web content. On Windows, Microsoft have laid out a principled approach[9] that includes the Content Restrictions API[10] which includes desired restrictions of web content.

On Android, however, there is currently no OS-provided content preferences API for apps to integrate with despite previous requests[11].

Existing platform signals exist on many devices. If they are insufficient in utility, accuracy, or usability when compared to other solutions to age restriction it is due to under-investment, rather than a fundamental difference in user experience constraints. Given the current interest in age restriction, the degree of investment by platforms may be sufficient to provide much more powerful user controls.

2.4 Signal Accuracy

This proposal's efficacy hinges upon the accuracy of these platform signals to reflect the capability of the **device user** to view restricted content as defined by the **device owner** as they are permitted by law. The constraints placed upon the device owner by the law are outside of the scope of this proposal, but are discussed in Section 3. Worth technical consideration is how to enable a device owner to manage the access of other users on the device.

Desktop computing has the paradigm of user accounts, which maps neatly onto this proposal. An administrator account can configure the device, including the creation and age restrictions of other users. Mobile does not have such a universal paradigm, however Guided Access introduces the concept of a reduced privilege state meant for children[12].

Another approach that can enable differentiation between device owners and users is requiring the system passcode to enter an elevated state that signifies the device administrator, another existing paradigm for device configuration. Placing the content restrictions configuration or a temporary relaxation of content restrictions behind such a state allows for a variety of user experiences with age restriction depending on the device owners' needs, all mediated by the platform.

[8] https://developer.apple.com/documentation/managedsettings/webcontentsettings/filterpolicy
[9] https://learn.microsoft.com/en-us/windows/win32/parcon/windows-family-safety-solution
[10] https://learn.microsoft.com/en-us/uwp/api/windows.media.contentrestrictions?view=winrt-26100
[11] https://issuetracker.google.com/issues/302210616
[12] https://support.apple.com/en-us/111795

3 Relationship with legal and policy initiatives

This proposal does not rely on any new legal or policy initiatives. It is already implementable on a broad range of devices. While it may work on a voluntary basis, lawmakers could take any number of actions from encouraging to requiring its use by devices and websites, weighing the potential benefits and harms.

3.1 Compatible legal and policy initiatives

This proposal opens up the possibility of more diverse policy initiatives for age restriction online by distributing the responsibility across browsers, operating systems, and sites.

Performing cryptographic age verification on the device is still compatible with this proposal. A wallet app could prove to the operating system that the user at hand is of age. Similarly, the age verification could be performed at the time of sale of the device, providing the device owner an unlocked state.

Device-enforced age verification also allows lawmakers to tailor requirements for categories of devices according to the likelihood for harm. For example, they could place more stringent requirements on devices marketed as suitable for children and less stringent or no requirements on devices not marketed directly to consumers.

Another direction that lawmakers may consider is mandating that hardware-backed attestation is used to prove that age verification and content filtering is enforced on the device. However, this denies users agency and ownership of their own devices, excludes older devices that do not have suitably modern hardware and does not improve effectiveness over schemes that do not rely on attestation.

Orthogonal to the initiatives focusing specifically on age verification are initiatives focused on getting websites to label their content correctly. Initiatives here can be driven voluntarily (through industry groups), technically (by identifying non-compliant sites and attempting to restrict them), legally (by enforcing compliance through regulation), or a combination thereof. While this toes closely toward censorship, there is an important distinction that sites are still free to publish as they were, so long as it is appropriately labeled. As long as the age verification mechanism on the device is not so onerous that any adult can bypass it, speech between adults is preserved.

Finally, inaction by lawmakers is entirely compatible with this technical approach. If the only material change is that browsers and operating systems plumb these signals together, a voluntary system of self-labeling and content filtering orchestrated by browsers, operating systems, and sites is a better world than we have today.

3.2 Enforceability

The proposal places the difficult task of rendering whether or not the user can view age restricted content as defined by applicable laws, whatever age verification mechanism the user is able to present, and the device owner's preference squarely at the feet of the device's operating system. This leaves a much simpler task, labeling content, for the much larger number of sites that may need to do so. This makes regulation more enforceable in some contexts because the burden of compliance is so much lower, especially on sites meant for adults only or social media sites that already have reporting mechanisms.

This has the consequence of adding an enforcement requirement onto operating system vendors or device sellers. However, regulating these may be easier as they are involved in the sale of a

physical good in the regulator's jurisdiction, particularly when compared to websites that may have no physical or business presence in the regulator's jurisdiction. Exceptions likely should be made for open source operating systems, e.g. on the grounds of freedom of speech.

4 Alternatives considered

While not exhaustive of our considerations, the leading alternatives are laid out below.

4.1 Website-enforced age verification

This proposal is an alternative to existing proposals for network-based access controls. The Center for Democracy and Technology addresses the numerous issues with these proposals at length in their amicus brief in Free Speech Coalition, Inc. v. Paxton last year[13].

One example is the use of the Digital Credentials API for age verification. Presenting credentials on the web carries significant risks[14] for user privacy and equitable access to web content which cannot be effectively mitigated.

This proposal enjoys several advantages over website-enforced approaches:

- It does not inherently require users to obtain credentials, and therefore doesn't disadvantage those who cannot.
- It does not add privacy risks from presenting government IDs online to numerous websites or age verification providers.
- It is cheaper and easier for websites to adopt, as it does not require costly integrations with third-party verification services, nor do websites have to identify which prospective users come under screening requirements.
- Device owners can more effectively customize their content preferences, providing greater safety and reducing the impact of unwanted blocking.
- Location arbitrage via VPN is not feasible in device-enforced age verification, no longer incentivizing restriction of access to VPNs.

4.2 Prefer:Safe

The 'safe' HTTP Preference (RFC 8674[15]) defined a mechanism by which browsers could indicate to websites that they did not want to receive objectionable content (as interpreted by the website). It received a mixed response from the community[16].

This proposal is the inverse of that approach, the website indicates the characteristics of its content and the browser decides whether to display it.

5 Acknowledgments

This work was significantly impacted by discussions with Nick Doty, Dennis Jackson, John Schanck, and Martin Thompson.

[13] https://cdt.org/wp-content/uploads/2024/09/20240920161551554_23-1122-Amicus-Brief.pdf
[14] https://github.com/w3c/credential-considerations/blob/main/credentials-considerations.md
[15] https://datatracker.ietf.org/doc/html/rfc8674
[16] https://prefersafe.github.io/