AGE VERIFICATION

Is Coming For the Internet.

EFF Built You a Resource Guide to Fight Back.

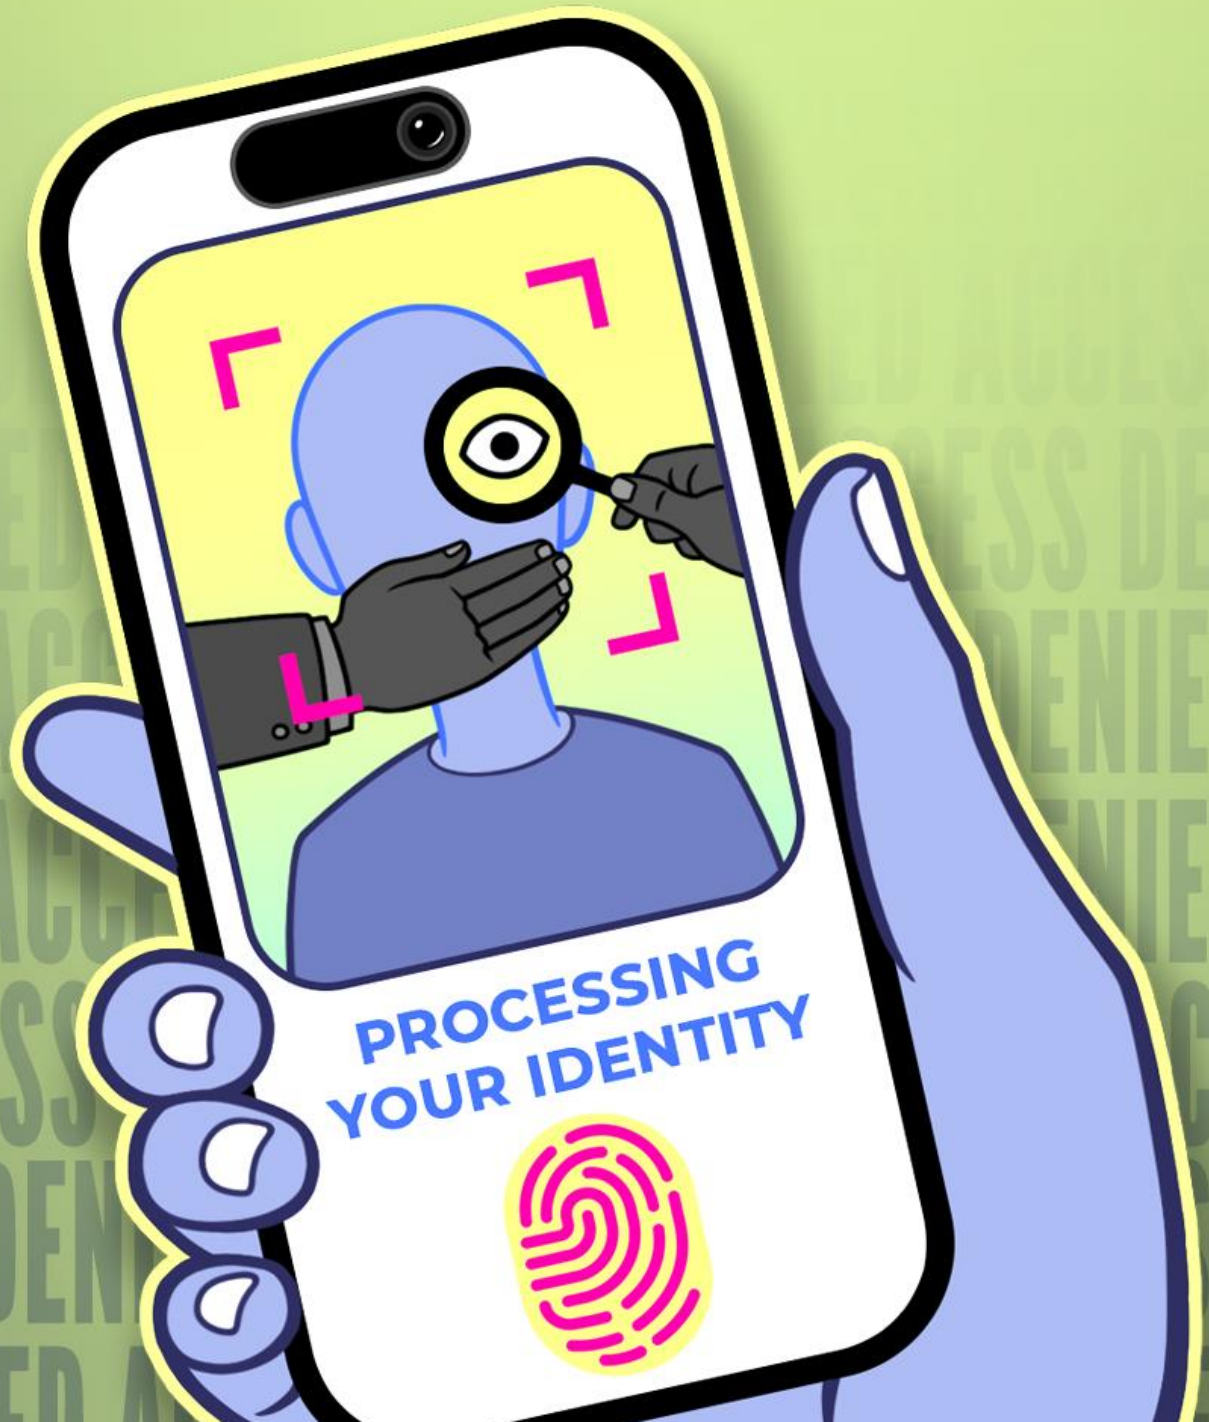

AGE VERIFICATION AND AGE GATING RESOURCE GUIDE

Why Join the Fight?

Age verification laws are proliferating fast across the United States and around the world, creating a dangerous and confusing tangle of rules about what we're all allowed to see and do online. Though these mandates claim to protect children, in practice they create harmful censorship and surveillance regimes that put everyone—adults and young people alike—at risk.

The term "age verification" is colloquially used to describe a wide range of age assurance technologies, from age verification systems that force you to upload government ID, to age estimation tools that scan your face, to systems that infer your age by making you share personal data. While different laws call for different methods, one thing remains constant: *every method out there* collects your sensitive, personal information and creates barriers to accessing the internet. We refer to all of these requirements as age verification, age assurance, or age-gating.

If you're feeling overwhelmed by this onslaught of laws and the invasive technologies behind them, you're not alone. It's a *lot.* But understanding how these mandates work and who they harm is critical to keeping yourself and your loved ones safe online. Age verification is lurking around every corner these days, so we must fight back to protect the internet that we know and love.

That's why we created this Age Verification Resource Guide and Resource Hub: a one-stop shop for users seeking to understand what these laws actually do, what's at stake, why EFF opposes all forms of age verification, how to protect yourself, and how to join the fight for a free, open, private, and yes—safe—internet.

Why Age Verification Mandates Are a Problem

In the U.S., more than half of all states have now passed laws imposing age-verification requirements on online platforms. Congress is considering even more at the federal level, with a recent House hearing weighing nineteen distinct proposals relating to young people's online safety—some sweeping, some contradictory, and each one more drastic and draconian than the last.

The rest of the world is moving in the same direction. We saw the UK's Online Safety Act go into effect last summer, Australia's new law barring access to social media for anyone under 16 went live shortly after, and a slew of other countries are currently considering similar restrictions.

We all want young people to be safe online. However, age verification is not the silver bullet that lawmakers want you to think it is. In fact, age-gating mandates will do more harm than good—*especially* for the young people they claim to protect. They undermine the fundamental speech rights of adults and young people alike; create new barriers to accessing vibrant, lawful, even life-saving content; and needlessly jeopardize all internet users' privacy, anonymity, and security.

If legislators want to meaningfully improve online safety, they should pass a strong, comprehensive federal privacy law instead of building new systems of surveillance, censorship, and exclusion.

What's Inside the Resource Guide and Resource Hub

Our Resource Guide and Resource Hub is built to answer the questions we hear from users every day, such as:

- How do age verification laws actually work?
- What's the difference between age verification, age estimation, age assurance, and all the other confusing technical terms I'm hearing?
- What's at stake for me, and who else is harmed by these systems?
- How can I keep myself, my family, and my community safe as these laws continue to roll out?
- What can I do to fight back?

- And if not age verification, what *else* can we do to protect the online safety of our young people?

We've got explainers, user-friendly guides, technical breakdowns, and advocacy tools. And this is just the start, so keep checking back as we continue to build out the site with new resources and answers to more of your questions on all things age verification.

How to Spot an Age Verification Mandate

EFF opposes all mandatory age verification legislation because it enables government censorship and burdens access to online speech. Though age verification mandates are often touted as "online safety" measures for kids, the laws actually do more harm than good. They undermine the fundamental speech rights of adults and young people alike, create new barriers to internet access, and put at risk all internet users' privacy, anonymity, and security.

So What Counts as an Age Verification Bill?

Age verification bills generally require online services to verify all users' ages—often through invasive tools like ID checks, biometric scans, and other dubious "age estimation" methods—before granting them access to certain online content or services. Some state bills mandate the age verification explicitly, including Texas's H.B. 1181, Florida's H.B. 3, and Indiana's S.B. 17. Other state bills *claim* not to require age verification, but still threaten platforms with liability for showing certain content or features to minor users. These bills—including Mississippi's H.B. 1126, Ohio's Parental Notification by Social Media Operators Act, and the federal Kids Online Safety Act—raise the question: how are platforms to know which users are minors *without* imposing age verification?

EFF's answer: they can't. We call these bills "implicit age verification mandates" because, though they might expressly deny requiring age verification, they still force platforms to either impose age verification measures or, worse, to censor whatever content or features deemed "harmful to minors" for all users—not just young people—in order to avoid liability.

No matter how they show up in legislation, age verification requirements are the wrong approach to protecting young people online. No one should have to hand over their most sensitive personal information or submit to invasive biometric surveillance just to access lawful online speech.

Age Verification, Estimation, Assurance, Oh My! A Guide to the Terminology

If you've been following the wave of age-gating laws sweeping across the country and the globe, you've probably noticed that lawmakers, tech companies, and advocates all seem to be using different terms for what sounds like the same thing. Age verification, age assurance, age estimation, age gating—they get thrown around interchangeably, but they technically mean different things. And those differences matter a lot when we're talking about your rights, your privacy, your data, and who gets to access information online.

So let's clear up the confusion. Here's your guide to the terminology that's shaping these laws, and why you should care about the distinctions.

Age Gating: “No Kids Allowed”

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 gating does not necessarily refer to any specific technology or manner of enforcement for estimating or verifying a user’s age. It simply refers to the fact that a restriction exists. Think of it as the concept of “you must be this old to enter” without getting into the details of how they’re checking.

Age Assurance: The Umbrella Term

Think of **age assurance** as the catch-all category. It covers any method an online service uses to figure out how old you are with *some level* of confidence. That's intentionally vague, because age assurance includes everything from the most basic check-the-box systems to full-blown government ID scanning.

Age assurance is the big tent that contains all the other terms we're about to discuss below. When a company or lawmaker talks about "age assurance," they're not being specific about *how* they're determining your age—just that they're trying to. For decades, the internet operated on a “self-attestation” system where you checked a box saying you were 18, and that was it. These new age-verification laws are specifically designed to replace that system. When lawmakers say they want "robust age assurance," what they really mean is "we don't trust self-attestation anymore, so now you need to prove your age beyond just swearing to it."

Age Estimation: Letting the Algorithm Decide

Age estimation is where things start getting creepy. Instead of asking you directly, the system *guesses* your age based on data it collects about you.

This might include:

- Analyzing your face through a video selfie or photo
- Examining your voice
- Looking at your online behavior—what you watch, what you like, what you post
- Checking your existing profile data

Companies like Instagram have partnered with services like Yoti to offer facial age estimation. You submit a video selfie, an algorithm analyzes your face, and spits out an estimated age range. Sounds convenient, right?

Here's the problem, "estimation" is exactly that: it's a guess. And it is inherently imprecise. Age estimation is notoriously unreliable, especially for teenagers—the exact group these laws claim to protect. An algorithm might tell a website you're somewhere between 15 and 19 years old. That's not helpful when the cutoff is 18, and what's at stake is a young person's constitutional rights.

And it gets worse. These systems consistently fail for certain groups:

- People of color are routinely misidentified (even Yoti's own research admits higher error rates for darker skin tones)
- Trans and nonbinary people are frequently misclassified
- People with disabilities that affect their appearance fall outside the algorithm's training parameters and anyone who doesn't fit the algorithmic "norm" gets flagged

When estimation fails (and it often does), users get kicked to the next level: actual verification. Which brings us to...

Age Verification: "Show Me Your Papers"

Age verification is the most invasive option. This is where you have to prove your age to a certain date, rather than, for example, prove that you have crossed some age threshold (like 18 or 21 or 65). EFF generally refers to most age gates and mandates on young people's access to online

information as “age verification,” as most of them typically require you to submit hard identifiers like:

- Government-issued ID (driver's license, passport, state ID)
- Credit card information
- Utility bills or other documents
- Biometric data

This is what a lot of new state laws are actually requiring, even when they use softer language like "age assurance." Age verification doesn't just confirm you're over 18, it reveals your full identity. Your name, address, date of birth, photo—everything.

Here's the critical thing to understand: age verification is really identity verification. You're not just proving you're old enough—you're proving exactly who you are. And that data has to be stored, transmitted, and protected by every website that collects it.

We already know how that story ends. Data breaches are inevitable. And when a database containing your government ID tied to your adult content browsing history gets hacked—and it will—the consequences can be devastating.

Why This Confusion Matters

Politicians and tech companies love using these terms interchangeably because it obscures what they're actually proposing. A law that requires "age assurance" sounds reasonable and moderate. But if that law defines age assurance as requiring government ID verification, it's not moderate at all—it's mass surveillance. Similarly, when Instagram says it's using "age estimation" to protect teens, that sounds privacy-friendly. But when their estimation fails and forces you to upload your driver's license instead, the privacy promise evaporates.

Here's the uncomfortable truth: most lawmakers writing these bills have no idea how any of this technology actually works. They don't know that age estimation systems routinely fail for people of color, trans individuals, and people with disabilities. They don't know that verification systems have error rates. They don't even seem to understand that the terms they're using mean different things. The fact that their terminology is all over the place—using "age assurance," "age verification," and "age estimation" interchangeably—makes this ignorance painfully clear, and leaves the onus on platforms to choose whichever option best insulates them from liability.

Language matters because it shapes how we think about these systems. "Assurance" sounds gentle. "Verification" sounds official. "Estimation" sounds technical and impersonal, and also admits its inherent imprecision. But they all involve collecting your data and create a metaphysical age gate to the internet. The terminology is deliberately confusing, but the stakes are clear: it's your privacy, your data, and your ability to access the internet without constant identity checks. Don't let fuzzy language disguise what these systems really do.

Why We Are Against Age Verification Mandates

At its core, age-restrictive legislation is a harmful censorship regime masquerading as an online safety measure. But it's important to know that this kind of mandate doesn't just affect young people. After all, in order to restrict access for one select age group, platforms must require *every* user to prove they are old enough to enter.

There are millions of U.S. adults who could be effectively locked out of much of the internet by draconian age verification measures because they lack valid government ID. Many more, including undocumented people, unhoused people, trans and nonbinary individuals, and survivors of abuse, have good reasons to keep their real-world identities separate from their online lives. Age gates risk turning huge parts of the web into closed spaces for the "verified," cutting off access to protected speech and vital information.

These laws also hand governments the power to arbitrarily decide what speech is "harmful." These definitions are often vague, subjective, and unscientific, causing platforms to over-censor user content (or shut down entirely) to avoid punishment. Indeed, we've already seen platforms respond to age-verification mandates in the U.S. and abroad by censoring or blocking

things like LGBTQ+ resources, reproductive health information, global news, and political organizing spaces.

In the U.S., anonymity is a fundamental cornerstone of free expression. Around the world, it has protected journalists, whistleblowers, and activists. **But age verification laws threaten to tie your most sensitive, immutable data—your name, face, birthday, and home address—to your online activity, ruining your anonymity in the process**.

Blocking entire communities or resources because of their subject matter does not make the internet safer; rather, it silences the people who rely on those online spaces for life-saving support, education, or safety.

Age Verification Systems Are Surveillance Systems

Here's the kicker: the technology being used to implement these mandates isn't even up for the task. Age estimation and ID verification tech is clunky at its best, and discriminatory at its worst. And every age-verification system is, at its core, a surveillance system.

While there are several different age-checking methods on the market, there is no technology available that is entirely privacy-protective, fully accurate, and that guarantees complete coverage of the population.

Facial-recognition systems are notoriously unreliable for marginalized communities, and commonly misjudge age based on skin tone, gender, and disability. And ID- or data-based systems exclude those who either don't have ID or whose appearance doesn't match their current documents.

The most common method of age verification—uploading a photo of your government ID along with a "live" selfie or video—is far more invasive than an in-person ID check.

Online, there's no way for users to verify that their private information has been deleted, or to ensure that it won't be copied, sold, or stolen. Companies that claim to delete personal data such as IDs—including age verification companies—have already experienced data breaches, and the more data a company collects, the more likely the chance of a data breach.

Alternative approaches, like facial age estimation or behavior tracking, aren't safe either.

They rely on algorithms or AI systems with high error rates, particularly when it comes to estimating age across race and gender lines. Other schemes link users to financial databases or digital ID systems, worsening discrimination and excluding people in the process. Finally, as EFF has shown, even promising technologies like zero-knowledge proofs or VPNs can't solve the underlying problem when the law itself is flawed.

These mandates create irresistible targets for hackers and governments alike, normalizing constant ID checks across the

web and defying long-standing internet safety norms in the process.

No one should have to sacrifice their privacy or anonymity in order to exercise their free speech rights online.

Online Age Verification Isn’t Just Like Showing Your ID In Person

One of the most common refrains we hear from age verification proponents is that online ID checks are nothing new. After all, you show your ID at bars and liquor stores all the time, right? And it’s true that many places age-restrict access in-person to various goods and services, such as tobacco, alcohol, firearms, lottery tickets, and even tattoos and body piercings.

But the comparison falls apart under scrutiny. There are fundamental differences between flashing your ID to a bartender and uploading government documents or biometric data to websites and third-party verification companies. Online age-gating is more invasive, affects far more people, and poses serious risks to privacy, security, and free speech that simply don't exist when you buy a six-pack at the corner store.

Online age verification burdens many more people.

Online age restrictions are imposed on many, many more users than in-person ID checks. Because of the sheer scale of the internet, regulations affecting online content sweep in an enormous number of adults and youth alike, forcing them to disclose sensitive personal data just to access lawful speech, information, and services.

Additionally, age restrictions in the physical world affect only a limited number of transactions: those involving a narrow set of age-restricted products or services. Typically this entails a bounded interaction about one specific purchase.

Online age verification laws, on the other hand, target a broad range of internet activities and general purpose platforms and services, including social media sites and app stores. And these laws don’t just wall off specific content deemed harmful to minors (like a bookstore would); they age-gate access to websites wholesale. This is akin to requiring ID every time a customer walks into a convenience store, regardless of whether they want to buy candy or alcohol.

There are significant privacy and security risks that don’t exist offline.

In offline, in-person scenarios, a customer typically provides their physical ID to a cashier or clerk directly. Oftentimes, customers need only flash their ID

for a quick visual check, and no personal information is uploaded to the internet, transferred to a third-party vendor, or stored. Online age-gating, on the other hand, forces users to upload—not just momentarily display—sensitive personal information to a website in order to gain access to age-restricted content.

This creates a cascade of privacy and security problems that don't exist in the physical world. Once sensitive information like a government-issued ID is uploaded to a website or third-party service, there is no guarantee it will be handled securely. You have no direct control over who receives and stores your personal data, where it is sent, or how it may be accessed, used, or leaked outside the immediate verification process.

Data submitted online rarely just stays between you and one other party. All online data is transmitted through a host of third-party intermediaries, and almost all websites and services also host a network of dozens of private, third-party trackers managed by data brokers, advertisers, and other companies that are constantly collecting data about your browsing activity. The data is shared with or sold to additional third parties and used to target behavioral advertisements. Age verification tools also often rely on third parties just to complete a transaction: a single instance of ID verification might involve two or three different third-party partners, and age estimation services often work directly with data brokers to offer a complete product. Users' personal identifying data then circulates among these partners.

All of this increases the likelihood that your data will leak or be misused. Unfortunately, data breaches are an endemic part of modern life, and the sensitive, often immutable, personal data required for age verification is just as susceptible to being breached as any other online data. Age verification companies can be—and already have been—hacked. Once that personal data gets into the wrong hands, victims are vulnerable to targeted attacks both online and off, including fraud and identity theft.

Troublingly, many age verification laws don't even protect user security by providing a private right of action to sue a company if personal data is breached or misused. This leaves you without a direct remedy should something bad happen.

Some proponents claim age estimation is a privacy-preserving alternative to ID-based verification. But age estimation tools still require biometric data collection, often demanding users submit a photo or video of their face to access a site. And again, once submitted, there's no way for you to verify how that data is processed or stored. Requiring face scans also normalizes pervasive biometric surveillance and creates infrastructure that could easily

be repurposed for more invasive tracking. Once we've accepted that accessing lawful speech requires submitting our faces for scanning, we've crossed a threshold that's difficult to walk back.

Online age verification creates even bigger barriers to access.

Online age gates create more substantial access barriers than in-person ID checks do. For those concerned about privacy and security, there is no online analog to a quick visual check of your physical ID. Users may be justifiably discouraged from accessing age-gated websites if doing so means uploading personal data and creating a potentially lasting record of their visit to that site.

Given these risks, age verification also imposes barriers to remaining anonymous that typically don't exist in-person. Anonymity is essential for those wishing to access sensitive, personal, or stigmatized content online. And users have a right to anonymity, which is "an aspect of the freedom of speech protected by the First Amendment." Even if a law requires data deletion, users must still be confident that every website and online service with access to their data will, in fact, delete it—something that is in no way guaranteed.

In-person ID checks are additionally less likely to wrongfully exclude people due to errors. Online systems that rely on facial scans are often incorrect, especially when applied to users near the legal age of adulthood. These tools are also less accurate for people with Black, Asian, Indigenous, and Southeast Asian backgrounds, for users with disabilities, and for transgender individuals. This leads to discriminatory outcomes and exacerbates harm to already marginalized communities. And while in-person shoppers can speak with a store clerk if issues arise, these online systems often rely on AI models, leaving users who are incorrectly flagged as minors with little recourse to challenge the decision.

In-person interactions may also be less burdensome for adults who don't have up-to-date ID. An older adult who forgets their ID at home or lacks current identification is not likely to face the same difficulty accessing material in a physical store, since there are usually distinguishing physical differences between young adults and those older than 35. A visual check is often enough. This matters, as a significant portion of the U.S. population does not have access to up-to-date government-issued IDs. This disproportionately affects Black Americans, Hispanic Americans, immigrants, and individuals with disabilities, who are less likely to possess the necessary identification.

We’re talking about First Amendment-protected speech.

It's important not to lose sight of what’s at stake here. The good or service age-gated by these laws isn't alcohol or cigarettes—it’s First Amendment-protected speech. Whether the target is social media platforms or any other online forum for expression, age verification blocks access to constitutionally-protected content.

Access to many of these online services is also necessary to participate in the modern economy. While those without ID may function just fine without being able to purchase luxury products like alcohol or tobacco, requiring ID to participate in basic communication technology significantly hinders people’s ability to engage in economic and social life.

This is why it’s wrong to claim online age verification is equivalent to showing ID at a bar or store. This argument handwaves away genuine harms to privacy and security, dismisses barriers to access that will lock millions out of online spaces, and ignores how these systems threaten free expression. Ignoring these threats won’t protect children, but it will compromise our rights and safety.

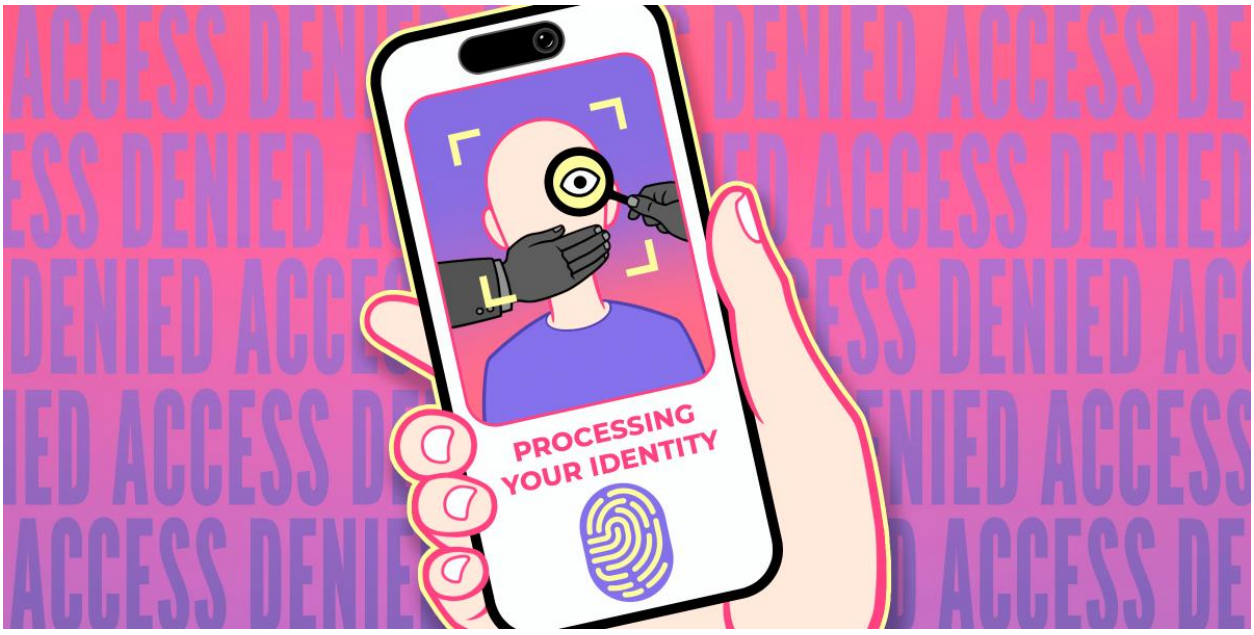

Who's Harmed by Age Verification Mandates?

Age-verification laws may sound straightforward to some: protect young people online by making everyone prove their age. But in reality, these mandates force users into one of two flawed systems—mandatory ID checks or biometric scans—and *both* are deeply discriminatory. These proposals burden everyone's right to speak and access information online, and structurally excludes the very people who rely on the internet most. In short, although these laws are often passed with the intention to protect children from harm, the reality is that these laws harm both adults and children.

Here's who gets hurt, and how:

1. Adults Without IDs Get Locked Out

Document-based verification assumes everyone has the right ID, in the right name, at the right address. About 15 million adult U.S. citizens don't have a driver's license, and 2.6 million lack any government-issued photo ID at all. Another 34.5 million adults don't have a driver's license or state ID with their current name and address.

Specifically:

- 18% of Black adults don't have a driver's license at all.
- Black and Hispanic Americans are disproportionately less likely to have current licenses.
- Undocumented immigrants often cannot obtain state IDs or driver's licenses.
- People with disabilities are less likely to have current identification.
- Lower-income Americans face greater barriers to maintaining valid IDs.

Some laws allow platforms to ask for financial documents like credit cards or mortgage records instead. But they still overlook the fact that nearly 35% of U.S. adults also don't own homes, and close to 20% of households don't have credit cards. Immigrants, regardless of legal status, may also be unable to obtain credit cards or other financial documentation.

2. Communities of Color Face Higher Error Rates

Platforms that rely on AI-based age-estimation systems often use a webcam selfie to guess users' ages. But these algorithms don't work equally well for everyone. Research has consistently shown that they are less accurate for people with Black, Asian, Indigenous, and Southeast Asian backgrounds; that they often misclassify those adults as being under 18; and sometimes take longer to process, creating unequal access to online spaces. This mirrors the well-documented racial bias in facial recognition technologies. The result is that technology's inherent biases can block people from speaking online or accessing others' speech.

3. People with Disabilities Face More Barriers

Age-verification mandates most harshly affect people with disabilities. Facial recognition systems routinely fail to recognize faces with physical differences, affecting an estimated 100 million people worldwide who live with facial differences, and "liveness detection" can exclude folks with limited mobility. As these technologies become gatekeepers to online spaces, people with disabilities find themselves increasingly blocked from essential services and platforms with no specified appeals processes that account for disability.

Document-based systems also don't solve this problem—as mentioned earlier, people with disabilities are also less likely to possess current driver's licenses, so document-based age-gating technologies are equally exclusionary.

4. Transgender and Non-Binary People Are Put At Risk

Age-estimation technologies perform worse on transgender individuals and cannot classify non-binary genders at all. For the 43% of transgender Americans who lack identity documents that correctly reflect their name or gender, age verification creates an impossible choice: provide documents with dead names and incorrect gender markers, potentially outing themselves in the process, or lose access to online platforms entirely—a risk that no one should be forced to take just to use social media or access legal content.

5. Anonymity Becomes a Casualty

Age-verification systems are, at their core, surveillance systems. By requiring identity verification to access basic online services, we risk creating an internet where anonymity is a thing of the past. For people who rely on anonymity for safety, this is a serious issue. Domestic abuse survivors need to stay anonymous to hide from abusers who could track them through their online activities. Journalists, activists, and whistleblowers regularly use anonymity to protect sources and organize without facing retaliation or government surveillance. And in countries under authoritarian rule, anonymity is often the only way to access banned resources or share information without being silenced. Age-verification systems that demand government IDs or biometric data would strip away these protections, leaving the most vulnerable exposed.

6. Young People Lose Access to Essential Information

Because state-imposed age-verification rules either block young people from social media or require them to get parental permission before logging on, they can deprive minors of access to important information about their health, sexuality, and gender. Many U.S. states mandate "abstinence only" sexual health education, making the internet a key resource for education and self-discovery. But age-verification laws can end up blocking young people from accessing that critical information. And this isn't just about porn, it's about sex education, mental health resources, and even important literature. Some states and countries may start going after content they deem "harmful to minors," which could include anything from books on sexual health to art, history, and even award-winning novels. And let's be clear: these laws often get used to target anything that challenges certain political or cultural narratives, from diverse educational materials to media that simply includes themes of sexuality or gender diversity. What begins as a "protection" for kids could easily turn into a full-on censorship movement, blocking content that's actually vital for minors' development, education, and well-being.

This is also especially harmful to homeschoolers, who rely on the internet for research, online courses, and exams. For many, the internet is central to their education and social lives. The internet is also crucial for homeschoolers' mental health, as many already struggle with isolation. Age-verification laws would restrict access to resources that are essential for their education and well-being.

7. LGBTQ+ Youth Are Denied Vital Lifelines

For many LGBTQ+ young people, especially those with unsupportive or abusive families, the internet can be a lifeline. For young people facing family rejection or violence due to their sexuality or gender identity, social media platforms often provide crucial access to support networks, mental health resources, and communities that affirm their identities. Age verification systems that require parental consent threaten to cut them from these crucial supports.

When parents must consent to or monitor their children's social media accounts, LGBTQ+ youth who lack family support lose these vital connections. LGBTQ+ youth are also disproportionately likely to be unhoused and lack access to identification or parental consent, further marginalizing them.

8. Youth in Foster Care Systems Are Completely Left Out

Age verification bills that require parental consent fail to account for young people in foster care, particularly those in group homes without legal guardians who can provide consent, or with temporary foster parents who cannot prove guardianship. These systems effectively exclude some of the most vulnerable young people from accessing online platforms and resources they may desperately need.

9. All of Our Personal Data is Put at Risk

An age-verification system also creates acute privacy risks for adults and young people. Requiring users to upload sensitive personal information (like government-issued IDs or biometric data) to verify their age creates serious privacy and security risks. Under these laws, users would not just momentarily display their ID like one does when accessing a liquor store, for example. Instead, they'd submit their ID to third-party companies, raising major concerns over who receives, stores, and controls that data. Once uploaded, this personal information could be exposed, mishandled, or even breached, as we've seen with past data hacks. Age-verification systems are no strangers to being compromised—companies like AU10TIX and platforms

like Discord have faced high-profile data breaches, exposing users' most sensitive information for months or even years.

The more places personal data passes through, the higher the chances of it being misused or stolen. Users are left with little control over their own privacy once they hand over these immutable details, making this approach to age verification a serious risk for identity theft, blackmail, and other privacy violations. Children are already a major target for identity theft, and these mandates perversely increase the risk that they will be harmed.

10. All of Our Free Speech Rights Are Trampled

The internet is today's public square—the main place where people come together to share ideas, organize, learn, and build community. Even the Supreme Court has recognized that social media platforms are among the most powerful tools ordinary people have to be heard.

Age-verification systems inevitably block some adults from accessing lawful speech and allow some young people under 18 users to slip through anyway. Because the systems are both over-inclusive (blocking adults) and under-inclusive (failing to block people under 18), they restrict lawful speech in ways that violate the First Amendment.

The Bottom Line

Age-verification mandates create barriers along lines of race, disability, gender identity, sexual orientation, immigration status, and socioeconomic class. While these requirements threaten everyone's privacy and free-speech rights, they fall heaviest on communities already facing systemic obstacles.

The internet is essential to how people speak, learn, and participate in public life. When access depends on flawed technology or hard-to-obtain documents, we don't just inconvenience users, we deepen existing inequalities and silence the people who most need these platforms. As outlined, every available method—facial age estimation, document checks, financial records, or parental consent—systematically excludes or harms marginalized people. The real question isn't whether these systems discriminate, but how extensively.

Age Verification Won't "Protect the Children"

We all want young people to be safe, but age verification is decidedly *not* the silver bullet solution to young people's online safety that regulators and corporations want you to think it is.

First, age verification mandates diminish the safety that comes with online privacy.

For decades, we've rightfully taught young people this simple rule: don't share personal information with strangers. But age verification laws fly in the face of that longstanding internet safety norm by requiring users to hand over their *real* names, *real* faces, *real* proof of their *real life* identities—and in some cases, their parents' identities too—just to use basic online services. These mandates will make it harder to distinguish between legitimate services and scams, creating new opportunities for fraud and identity theft. If adults struggle to tell the difference, how can we expect our children—whose data is even more valuable to data thieves—to fare any better?

Second, age verification mandates threaten the safety that comes with empowering young people's autonomy and allowing them broad access to important information.

Young people across the world have the right to freedom of expression. In the United States, that means their right to access information and to communicate with others is protected by the First Amendment. Cutting off young people's access to the internet will not only erase their important voices and perspectives, but could also impact their development, ability to form offline relationships, and autonomy. After all, social media sites are not just for entertainment; they provide space for young people to explore their identities—whether by creating and sharing art, practicing religion, or engaging in politics. Blocking our youth from these diverse spaces robs them of opportunities to develop as individuals and participate in public life, and to find safety in supportive online communities that they can't always access in the physical world.

Third, it is impossible to define what's "appropriate" for everyone under 18.

Each young person is different, shaped by their family, cultural background, and maturity level. It makes little sense to treat all people under 18 the same—and what's "appropriate" for a 6-year-old is not the same for a 17-year-old. Yet age verification mandates commonly treat all young people the same, erasing that critical nuance.

In the end, it's families—not lawmakers or tech companies—who are best positioned to meaningfully guide their young people's online lives.

Open, honest conversations with young people about privacy, safety, and digital literacy are far more effective safety measures than blanket age-gating restrictions that hand those important decisions over to the government. And for those parents who do want some digital assistance guiding their children's internet use, there are already plenty of existing parental controls they can use to customize nearly every level of the user experience—without forcing the entire internet to show ID at the door.

The Impact of Age Verification Measures Goes Beyond Porn Sites

As age verification bills pass across the world under the guise of “keeping children safe online,” governments are increasingly giving themselves the authority to decide what topics are deemed “safe” for young people to access, and forcing online services to remove and block anything that may be deemed “unsafe.” This growing legislative trend has sparked significant concerns and numerous First Amendment challenges in the US and abroad.

These challenges keep arising because this isn’t just about safety—**it’s censorship**. Age verification laws target a slew of broadly-defined topics. Some block access to websites that contain some "sexual material harmful to minors," but define the term so loosely that “sexual material” could encompass anything from sex education to R-rated movies; others simply list a variety of vaguely-defined harms. In either instance, lawmakers and regulators could use the laws to target LGBTQ+ content online.

This risk is especially clear given what we already know about platform content policies. These policies, which claim to "protect children" or keep sites “family-friendly,” often label LGBTQ+ content as “adult” or “harmful,” while similar content that doesn't involve the LGBTQ+ community is left untouched. Sometimes, this impact—the censorship of LGBTQ+ content—is implicit, and only becomes clear when the policies (and/or laws) are actually implemented. Other times, this intended impact is explicitly spelled out in the text of the policies and bills.

In either case, **it is critical to recognize that age verification bills could block far more than just pornography.**

Take Oklahoma’s bill, SB 1959, for example. This state age verification law aims to prevent young people from accessing content that is “harmful to minors” and went into effect last November 1st. It incorporates definitions from another Oklahoma statute, Statute 21-1040, which defines material “harmful to minors” as any description or exhibition, in whatever form, of nudity and “sexual conduct.” That same statute then defines “sexual conduct” as including acts of “homosexuality.” Explicitly, then, SB 1959 requires a site to verify someone’s age before showing them content about homosexuality—a vague enough term that it could potentially apply to content from organizations like GLAAD and Planned Parenthood.

This vague definition will undoubtedly cause platforms to over-censor content relating to LGBTQ+ life, health, or rights out of fear of liability. Separately, bills such as SB 1959 might also cause users to self-police their own speech for the same reasons, fearing de-platforming. The law leaves platforms unsure and unable to precisely exclude the minimum amount of content that fits the bill's definition, leading them to over censorship of content that may just also include this very blog post.

Beyond Individual States: Kids Online Safety Act (KOSA)

Laws like the proposed federal Kids Online Safety Act (KOSA) make government officials the arbiters of what young people can see online and will lead platforms to implement invasive age verification measures to avoid the threat of liability. If KOSA passes, it will lead to people who make online content about sex education, and LGBTQ+ identity and health, being persecuted and shut down as well. All it will take is one member of the Federal Trade Commission seeking to score political points, or a state attorney general seeking to ensure re-election, to start going after the online speech they don't like. These speech burdens will also affect regular users as platforms mass-delete content in the name of avoiding lawsuits and investigations under KOSA.

Senator Marsha Blackburn, co-sponsor of KOSA, has expressed a priority in "protecting minor children from the transgender [sic] in this culture and that influence." KOSA, to Senator Blackburn, would address this problem by limiting content in the places "where children are being indoctrinated." Yet these efforts all fail to protect children from the actual harms of the online world, and instead deny vulnerable young people a crucial avenue of communication and access to information.

LGBTQ+ Platform Censorship by Design

While the censorship of LGBTQ+ content through age verification laws can be represented as an "unintended consequence" in certain instances, barring access to LGBTQ+ content is part of the platforms' design. One of the more pervasive examples is Meta suppressing LGBTQ+ content across its platforms under the guise of protecting younger users from "sexually suggestive content." According to a recent report, Meta has been hiding posts that reference LGBTQ+ hashtags like #lesbian, #bisexual, #gay, #trans, and #queer for users that turned the sensitive content filter on, as well as showing users a blank page when they attempt to search for LGBTQ+ terms. This leaves teenage users with no choice in what content they see, since the sensitive content filter is turned on for them by default.

This policy change came on the back of a protracted effort by Meta to allegedly protect teens online. In January last year, the corporation announced a new set of "sensitive content" restrictions across its platforms (Instagram, Facebook, and Threads), including hiding content which the platform no longer considered age-appropriate. This was followed later by the introduction of Instagram For Teens to further limit the content users under the age of 18 could see. This feature sets minors' accounts to the most restrictive levels by default, and teens under 16 can only reverse those settings through a parent or guardian.

Meta has apparently now reversed the restrictions on LGBTQ+ content after calling the issue a "mistake." This is not good enough. In allowing pro-LGBTQ+ content to be integrated into the sensitive content filter, Meta has aligned itself with those that are actively facilitating a violent and harmful removal of rights for LGBTQ+ people—all under the guise of keeping children and teens safe. Not only is this a deeply flawed strategy, it harms everyone who wishes to express themselves on the internet. These policies are written and enforced discriminatorily and at the expense of transgender, gender-fluid, and nonbinary speakers. They also often convince or require platforms to implement tools that, using the laws' vague and subjective definitions, end up blocking access to LGBTQ+ and reproductive health content.

The censorship of this content prevents individuals from being able to engage with such material online to explore their identities, advocate for broader societal acceptance and against hate, build communities, and discover new interests. With corporations like Meta intervening to decide how people create, speak, and connect, a crucial form of engagement for all kinds of users has been removed and the voices of people with less power are regularly shut down.

And at a time when LGBTQ+ individuals are already under vast pressure from violent homophobic threats offline, these online restrictions have an amplified impact.

LGBTQ+ youth are at a higher risk of experiencing bullying and rejection, often turning to online spaces as outlets for self-expression. For those without family support or who face the threat of physical or emotional abuse at home because of their sexual orientation or gender identity, the internet becomes an essential resource. A report from the Gay, Lesbian & Straight Education Network (GLSEN) highlights that LGBTQ+ youth engage with the internet at higher rates than their peers, often showing greater levels of civic engagement online compared to offline. Access to digital communities and resources is critical for LGBTQ+ youth, and restricting access to them poses unique dangers.

Call to Action: Digital Rights Are LGBTQ+ Rights

These laws have the potential to harm us all—including the children they are designed to protect.

As more U.S. states and countries pass age verification laws, it is crucial to recognize the broader implications these measures have on privacy, free speech, and access to information. This conglomeration of laws poses significant challenges for users trying to maintain anonymity online and access critical content—whether it's LGBTQ+ resources, reproductive health information, or otherwise. These policies threaten the very freedoms they purport to protect, stifling conversations about identity, health, and social justice, and creating an environment of fear and repression.

The fight against these laws is not just about defending online spaces; it's about safeguarding the fundamental rights of all individuals to express themselves and access life-saving information.

We need to stand up against these age verification laws—not only to protect users' free expression rights, but also to safeguard the free flow of information that is vital to a democratic society. Reach out to your state and federal legislators, raise awareness about the consequences of these policies, and support organizations like the LGBT Tech, ACLU, the Woodhull Freedom Foundation, and others that are fighting for digital rights of young people alongside EFF.

The fight for the safety and rights of LGBTQ+ youth is not just a fight for visibility—it's a fight for their very survival. Now more than ever, it's essential for allies, advocates, and marginalized communities to push back against these dangerous laws and ensure that the internet remains a space where all voices can be heard, free from discrimination and censorship.

Age Gates Are A Windfall for Big Tech—And A Death Sentence For Smaller Platforms

Lawmakers often sell age-verification mandates as a silver bullet for Big Tech's harms, but in practice, these laws do nothing to rein in the tech giants. Instead, they end up crushing smaller platforms that can't absorb the exorbitant costs. We've seen this play out in several places in the U.S. and across the globe, and the pattern is clear: age verification laws entrench Big Tech's dominance, while pushing smaller communities like Bluesky and Dreamwidth offline altogether.

Case Study: Sorry Mississippians, We Can't Afford You

In the Summer of 2025, when Mississippi's wide-sweeping age-verification law, H.B. 1126, became effective, social platforms Bluesky and Dreamwidth both decided to block all users in Mississippi from their services rather than risk hefty fines under the state's oppressive age verification mandate.

Bluesky was the first platform to make the announcement. In a public blogpost, Bluesky condemned H.B. 1126's broad scope, barriers to innovation, and privacy implications, explaining that the law forces platforms to "make *every* Mississippi Bluesky user hand over sensitive personal information and undergo age checks to access the site—or risk massive fines." As Bluesky noted, "This dynamic entrenches existing big tech platforms while stifling the innovation and competition that benefits users." Instead, Bluesky made the decision to cut off Mississippians entirely until the courts consider whether to overturn the law.

About a week later, we saw a similar announcement from Dreamwidth, an open-source online community similar to LiveJournal where users share creative writing, fanfiction, journals, and other works. In its post, Dreamwidth shared that it too would have to resort to blocking the IP addresses of all users in Mississippi because it could not afford the hefty fines.

Dreamwidth wrote: "Even a single $10,000 fine would be rough for us, but the per-user, per-incident nature of the actual fine structure is an existential threat." The service also expressed fear that being involved in the lawsuit against Mississippi left it particularly vulnerable to retaliation—a clear illustration of the chilling effect of these laws. For Dreamwidth, blocking Mississippi users entirely was the only way to survive.

Age Verification Mandates Don't Rein In Big Tech—They Entrench It

Proponents of age verification claim that these mandates will hold Big Tech companies accountable for their outsized influence, but really the opposite is true. As we can see from Mississippi, age verification mandates concentrate and consolidate power in the hands of the largest companies—the only entities with the resources to build costly compliance systems and absorb potentially massive fines. While megacorporations like Google (with YouTube) and Meta (with Instagram) are already experimenting with creepy new age-estimation tech on their social platforms, smaller sites like Bluesky and Dreamwidth simply cannot afford the risks.

We've already seen how this plays out in the UK. When the Online Safety Act came into force recently, platforms like Reddit, YouTube, and Spotify implemented broad (and extremely clunky) age verification measures while smaller sites, including forums on parenting, green living, and gaming on Linux, were forced to shutter. Take, for example, the Hamster Forum, "home of all things hamstery," which announced in March 2025 that the OSA would force it to shut down its community message boards. Instead, users were directed to migrate over to Instagram with this wistful disclaimer: "It will not be the same by any means, but . . . We can follow each other and message on there and see each others [sic] individual posts and share our hammy photos and updates still."

This perfectly illustrates the market impact of online age verification laws. When smaller platforms inevitably cave under the financial pressure of these mandates, users will be pushed back to the social media giants. These huge companies—those that can afford expensive age verification systems and aren't afraid of a few $10,000 fines while they figure out compliance—will end up getting *more* business, *more* traffic, and *more* power to censor users and violate their privacy.

This consolidation of power is a dream come true for the Big Tech platforms, but it's a nightmare for users. While the megacorporations get more traffic and a whole lot more user data (read: profit), users are left with far fewer community options and a bland, corporate surveillance machine instead of a vibrant public sphere. The internet we all fell in love with is a diverse and colorful place, full of innovation, connection, and unique opportunities for self-expression. That internet—*our* internet—is worth defending.

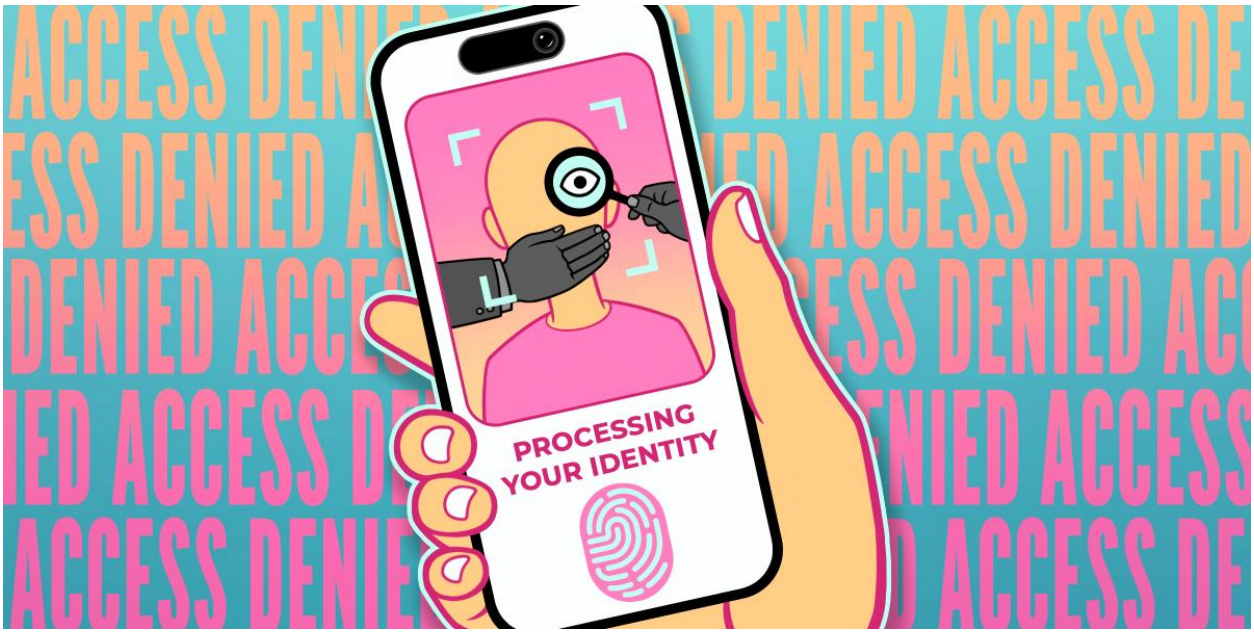

Is This Legal?

The Supreme Court's June, 2025 decision in *Free Speech Coalition v. Paxton* was a direct blow to the free speech rights of adults. The Court ruled that "no person—adult or child—has a First Amendment right to access speech that is obscene to minors without first submitting proof of age." This ruling allows states to enact onerous age-verification rules that will block adults from accessing lawful speech, curtail their ability to be anonymous, and jeopardize their data security and privacy. **These are real and immense burdens on adults, and the Court was wrong to ignore them in upholding Texas' law.**

Importantly, the Court's reasoning applies only to age-verification rules for certain sexual material, and not to age limits in general. We will continue to fight against age restrictions on online access more broadly, such as on social media and specific online features.

Still, the decision has immense consequences for internet users in Texas and in other states that have enacted similar laws. The Texas law forces adults to submit personal information over the internet to access entire websites that hold some amount of sexual material, not just the pages or portions of sites that specifically contain sexual materials. Many sites that cannot reasonably implement age verification measures for reasons such as cost or technical requirements will likely block users living in Texas and other states with similar laws wholesale.

Many users will not be comfortable sharing private information to access sites that do implement age verification, for reasons of privacy or concern for data breaches. Many others do not have a driver's license or photo ID to complete the age verification process. This decision will, ultimately, deter adult users from speaking and accessing lawful content, and will endanger the privacy of those who choose to go forward with verification.

What the Court Said

In the 6-3 decision, the Court ruled that Texas' HB 1181 is constitutional. This law requires websites that Texas decides are composed of "one-third" or more of "sexual material harmful to minors" to confirm the age of users by collecting age-verifying personal information from all visitors—even to access the other two-thirds of material that is not adult content.

In 1997, the Supreme Court struck down a federal online age-verification law in *Reno v. American Civil Liberties Union*. In that case the court ruled that many elements of the Communications Decency Act violated the First Amendment, including part of the law making it a crime for anyone to engage in online speech that is "indecent" or "patently offensive" if the speech could be viewed by a minor. Like HB 1181, that law would have resulted in many users being unable to view constitutionally protected speech, as many websites would have had to implement age verification, while others would have been forced to shut down.

In *Reno* and in subsequent cases, the Supreme Court ruled that laws that burden adults' access to lawful speech are subjected to the highest level of review under the First Amendment, known as strict scrutiny. This level of scrutiny requires a law to be very narrowly tailored or the least speech-restrictive means available to the government.

That all changed with the Supreme Court's decision on.

The Court now says that laws that burden adults' access to sexual materials that are obscene to minors are subject to less-searching First Amendment review, known as intermediate scrutiny. And under that lower standard, the Texas law does not violate the First Amendment. The Court did not have to respond to arguments that there are less speech-restrictive ways of reaching the same goal—for example, encouraging parents to install content-filtering software on their children's devices.

The court reached this decision by incorrectly assuming that online age verification is functionally equivalent to flashing an ID at a brick-and-mortar

store. As we explained in our amicus brief, this ignores the many ways in which verifying age online is significantly more burdensome and invasive than doing so in person. As we and many others have previously explained, unlike with in-person age-checks, the only viable way for a website to comply with an age verification requirement is to require all users to upload and submit—not just momentarily display—a data-rich government-issued ID or other document with personal identifying information.

This leads to a host of serious anonymity, privacy, and security concerns—all of which the majority failed to address. A person who submits identifying information online can never be sure if websites will keep that information or how that information might be used or disclosed. This leaves users highly vulnerable to data breaches and other security harms. Age verification also undermines anonymous internet browsing, even though courts have consistently ruled that anonymity is an aspect of the freedom of speech protected by the First Amendment.

This Supreme Court broke a fundamental agreement between internet users and the state that has existed since its inception

The Court sidestepped its previous online age verification decisions by claiming the internet has changed too much to follow the precedent from *Reno* that requires these laws to survive strict scrutiny. Writing for the minority, Justice Kagan disagreed with the premise that the internet has changed: "the majority's claim—again mistaken—[is] that the internet has changed too much to follow our precedents' lead."

But the majority argues that past precedent does not account for the dramatic expansion of the internet since the 1990s, which has led to easier and greater internet access and larger amounts of content available to teens online. The majority's opinion entirely fails to address the obvious corollary: the internet's expansion also has benefited adults. Age verification requirements now affect exponentially more adults than they did in the 1990s and burden vastly more constitutionally protected online speech. The majority's argument actually demonstrates that the burdens on adult speech have grown dramatically *larger* because of technological changes, yet the Court bizarrely interprets this expansion as justification for weaker constitutional protection.

What It Means Going Forward

This Supreme Court broke a fundamental agreement between internet users and the state that has existed since its inception: the government will not stand in the way of people accessing First Amendment-protected material.

There is no question that multiple states will now introduce similar laws to Texas. As of June 2025, when the *Paxton* decision came down, two dozen already had, though they were not all in effect. Even worse, at least three of those states have no limit on the percentage of sexual material required before the age verification law applies—a sweeping restriction on every site that contains *any* material that the state believes the law includes. These laws will force U.S.-based adult websites to implement age-verification or block users in those states, as many have in the past when similar laws were in effect.

Rather than submit to verification, research has found that people will choose a variety of other paths: using VPNs to indicate that they are outside of the state, accessing similar sites that don't comply with the law, often because the site is operating in a different country. While many users will simply not access the content as a result, others may accept the risk, at their peril.

We expect some states to push the envelope in terms of what content they consider "harmful to minors," and to expand the type of websites that are covered by these laws, either through updated language or threats of litigation. Even if these attacks are struck down, operators of sites that involve sexual content of any type may be under threat, especially if that information is politically divisive. We worry that the point of some of these laws will be to deter queer folks and others from accessing lawful speech and finding community online by requiring them to identify themselves. We will continue to fight to protect against the disclosure of this critical information and for people to maintain their anonymity.

EFF Will Continue to Fight for All Users' Free Expression and Privacy

That said, the ruling does not give states or Congress the green light to impose age-verification regulations on the broader internet. The majority's decision rests on the fact that minors do not have a First Amendment right to access sexual material that would be obscene to minors. In short, adults have a First Amendment right to access those sexual materials, while minors do not. Although it was wrong, the majority's opinion ruled that because Texas is blocking minors from speech they have no constitutional right to access, the age-verification requirement only incidentally burdens adult's First Amendment rights.

But the same rationale does not apply to general-audience sites and services, including social media. Minors and adults have coextensive rights to both speak and access the speech of other users on these sites because the vast majority of the speech is not sexual material that would be obscene to

minors. Lawmakers should be careful not to interpret this ruling to mean that broader restrictions on minors' First Amendment rights, like those included in the Kids Online Safety Act, would be deemed constitutional.

Free Speech Coalition v. Paxton will have an effect on nearly every U.S. adult internet user for the foreseeable future. It marks a worrying shift in the ways that governments can restrict access to speech online. But that only means we must work harder than ever to protect privacy, security, and free speech as central tenets of the internet.

Why We Need Comprehensive Data Privacy Laws Instead

Comprehensive data privacy legislation is the best way to hold tech companies accountable in our surveillance age, including for harm they do to children. Well-written privacy legislation has the added benefit of being constitutional—unlike the flurry of laws that restrict content behind age verification requirements that courts have recently blocked. Such misguided laws do little to protect kids while doing much to invade everyone's privacy and speech.

The answer is to re-focus attention on comprehensive data privacy legislation, which would address the massive collection and processing of personal data that is the root cause of many problems online. Just as important, it is far easier to write data privacy laws that are constitutional. Laws that lock online content behind age gates can almost never withstand First Amendment scrutiny because they frustrate all internet users' rights to access information and often impinge on people's right to anonymity.

It Is Comparatively Easy to Write Data Privacy Laws That Are Constitutional

EFF has long pushed for strong comprehensive commercial data privacy legislation and continues to do so. Data privacy legislation has many components. But at its core, it should minimize the amount of personal data that companies process, give users certain rights to control their personal data, and allow consumers to sue when the law is violated.

EFF has argued that privacy laws pass First Amendment muster when they have a few features that ensure the law reasonably fits its purpose. First, they regulate the commercial processing of personal data. Second, they do not impermissibly restrict the truthful publication of matters of public concern. And finally, the government's interest and law's purpose is to protect data privacy; expand the free expression that privacy enables; and protect the security of data against insider threats, hacks, and eventual government surveillance. If so, the privacy law will be constitutional if the government shows a close fit between the law's goals and its means.

EFF made this argument in support of the Illinois Biometric Information Privacy Act (BIPA), and a law in Maine that limits the use and disclosure of personal data collected by internet service providers. BIPA, in particular, has proved wildly important to biometric privacy. For example, it led to a

settlement that prohibits the company Clearview AI from selling its biometric surveillance services to law enforcement in the state. Another settlement required Facebook to pay hundreds of millions of dollars for its policy (since repealed) of extracting faceprints from users without their consent.

Courts have agreed. Privacy laws that have been upheld under the First Amendment, or cited favorably by courts, include those that regulate biometric data, health data, credit reports, broadband usage data, phone call records, and purely private conversations.

The Supreme Court, for example, has cited the federal 1996 Health Insurance Portability and Accountability Act (HIPAA) as an example of a "coherent" privacy law, even when it struck down a state law that targeted particular speakers and viewpoints. Additionally, when evaluating the federal Wiretap Act, the Supreme Court correctly held that the law cannot be used to prevent a person from publishing legally obtained communications on matters of public concern. But it otherwise left in place the wiretap restrictions that date back to 1934, designed to protect the confidentiality of private conversations.

Support Comprehensive Privacy Legislation That Will Stand the Test of Time

Instead of requiring online services to collect, analyze, and store our most sensitive and immutable data—exposing us to massive harms, both online and off—lawmakers should focus on passing laws that will have a lasting impact for adults and young people alike: strong, well-written comprehensive data privacy.

So, You've Hit an Age Gate. What Now?

EFF is against age gating and age verification mandates, and we hope we'll win in getting existing ones overturned and new ones prevented. But mandates are already in effect, and every day many people are asked to verify their age across the web, despite prominent cases of sensitive data getting leaked in the process.

At some point, you may have been faced with the decision yourself: should I continue to use this service if I have to verify my age? And if so, how can I do that with the least risk to my personal information? This is our guide to navigating those decisions, with information on what questions to ask about the age verification options you're presented with, and answers to those questions for some of the top most popular social media sites. Even though there's no way to implement mandated age gates in a way that fully protects speech and privacy rights, our goal here is to help you minimize the infringement of your rights as you manage this awful situation.

Follow the Data

Since we know that leaks happen despite the best efforts of software engineers, we generally recommend submitting the absolute least amount of

data possible. Unfortunately, that's not going to be possible for everyone. Even facial age estimation solutions where pictures of your face never leave your device, offering some protection against data leakage, are not a good option for all users: facial age estimation works less well for people of color, trans and nonbinary people, and people with disabilities. There are some systems that use fancy cryptography so that a digital ID saved to your device won't tell the website anything more than if you meet the age requirement, but access to that digital ID isn't available to everyone or for all platforms. You may also not want to register for a digital ID and save it to your phone, if you don't want to take the chance of all the information on it being exposed upon request of an over-zealous verifier, or you simply don't want to be a part of a digital ID system

If you're given the option of selecting a verification method and are deciding which to use, we recommend considering the following questions for each process allowed by each vendor:

- **Data:** What info does each method require?
- **Access:** Who can see the data during the course of the verification process?
- **Retention:** Who will hold onto that data after the verification process, and for how long?
- **Audits:** How sure are we that the stated claims will happen in practice? For example, are there external audits confirming that data is not accidentally leaked to another site along the way? Ideally these will be in-depth, security-focused audits by specialized auditors like NCC Group or Trail of Bits, instead of audits that merely certify adherence to standards.
- **Visibility:** Who will be aware that you're attempting to verify your age, and will they know which platform you're trying to verify for?

We attempt to provide answers to these questions below. To begin, there are two major factors to consider when answering these questions: the tools each platform uses, and the overall system those tools are part of.

In general, most platforms offer **age estimation options like face scans** as a first line of age assurance. These vary in intrusiveness, but their main problem is inaccuracy, particularly for marginalized users. Third-party age verification vendors Private ID and k-ID offer on-device facial age estimation, but another common vendor, Yoti, sends the image to their servers during age checks by some of the biggest platforms. This risks leaking the images themselves, and also the fact that you're using that particular website, to the third party.

Then, there's the **document-based verification services**, which require you to submit a hard identifier like a government-issued ID. This method thus requires you to prove both your age and your identity. A platform can do this in-house through a designated dataflow, or by sending that data to a third party. We've already seen examples of how this can fail. For example, Discord routed users' ID data through its general customer service workflow so that a third-party vendor could perform manual review of verification appeals. No one involved ever deleted users' data, so when the system was breached, Discord had to apologize for the catastrophic disclosure of nearly 70,000 photos of users' ID documents. Overly long retention periods expose documents to risk of breaches and historical data requests. Some document verifiers have retention periods that are needlessly long. This is the case with Incode, which provides ID verification for Tiktok. Incode holds onto images forever by default, though TikTok should automatically start the deletion process on your behalf.

Some platforms offer alternatives, like proving that you own a credit card, or asking for your email to check if it appears in databases associated with adulthood (like home mortgage databases). These tend to involve less risk when it comes to the sensitivity of the data itself, especially since credit cards can be replaced, but in general still undermine anonymity and pseudonymity and pose a risk of tracking your online activity. We'd prefer to see more assurances across the board about how information is handled.

Each site offers users a menu of age assurance options to choose from. We've chosen to present these options in the rough order that we expect most people to prefer.

Meta – Facebook, Instagram, WhatsApp, Messenger, Threads

Inferred Age

If Meta can guess your age, you may never even see an age verification screen. Meta, which runs Facebook, Threads, Instagram, Messenger, and WhatsApp, first tries to use information you've posted to guess your age, like looking at "Happy birthday!" messages. It's a creepy reminder that they already have quite a lot of information about you.

If Meta cannot guess your age, or if Meta infers you're too young, it will next ask you to verify your age using either facial age estimation, or by uploading your photo ID.

Face Scan

If you choose to use facial age estimation, you'll be sent to Yoti, a third-party verification service. Your photo will be uploaded to their servers during this process. Yoti claims that "as soon as an age has been estimated, the facial image is immediately and permanently deleted." Though it's not as good as not having that data in the first place, Yoti's security measures include a bug bounty program and annual penetration testing. Researchers from Mint Secure found that Yoti's app and website are filled with trackers, so the fact that you're verifying your age could be not only shared to Yoti, but leaked to third-party data brokers as well.

You may not want to use this option if you're worried about third parties potentially being able to know you're trying to verify your age with Meta. You also might not want to use this if you're worried about a current picture of your face accidentally leaking—for example, if elements in the background of your selfie might reveal your current location. On the other hand, if you consider a selfie to be less sensitive than a photograph of your ID, this option might be better. If you do choose (or are forced to) use the face check system, be sure to snap your selfie without anything you'd be concerned with identifying your location or embarrassing you in the background in case the image leaks.

Upload ID

If Yoti's age estimation decides your face looks too young, or if you opt out of facial age estimation, your next recourse is to send Meta a photo of your ID. Meta sends that photo to Yoti to verify the ID. Meta says it will hold onto that ID image for 30 days, then delete it. Meanwhile, Yoti claims it will delete the image immediately after verification. Of course, bugs and process oversights exist, such as accidentally replicating information in logs or support queues, but at least they have stated processes. Your ID contains sensitive information such as your full legal name and home address. Using this option not only runs the (hopefully small, but never nonexistent) risk of that data getting leaked through errors or hacking, but it also lets Meta see the information needed to tie your profile to your identity—which you may not want. If you don't want Meta to know your name and where you live, or rely on both Meta and Yoti to keep to their deletion promises, this option may not be right for you.

Google – Gmail, YouTube

Inferred Age

If Google can guess your age, you may never even see an age verification screen. Your Google account is typically connected to your YouTube account, so if (like mine) your YouTube account is old enough to vote, you may not need to verify your Google account at all. Google first uses information it already knows to try to guess your age, like how long you've had the account and your YouTube viewing habits. It's yet another creepy reminder of how much information these corporations have on you, but at least in this case they aren't likely to ask for even more identifying data.

If Google cannot guess your age, or decides you're too young, Google will next ask you to verify your age. You'll be given a variety of options for how to do so, with availability that will depend on your location and your age.

Google's methods to assure your age include ID verification, facial age estimation, verification by proxy, and digital ID. To prove you're over 18, you may be able to use facial age estimation, give Google your credit card information, or tell a third-party provider your email address.

Face Scan

If you choose to use facial age estimation, you'll be sent to a website run by Private ID, a third-party verification service. The website will load Private ID's verifier within the page—this means that your selfie will be checked without any images leaving your device. If the system decides you're over 18, it will let Google know that, and only that. Of course, no technology is perfect—should Private ID be mandated to target you specifically, there's nothing to stop it from sending down code that does in fact upload your image, and you probably won't notice. But unless your threat model includes being specifically targeted by a state actor or Private ID, that's unlikely to be something you need to worry about. For most people, no one else will see your image during this process. Private ID will, however, be told that your device is trying to verify your age with Google and Google will still find out if Private ID thinks that you're under 18.

If Private ID's age estimation decides your face looks too young, you may next be able to decide if you'd rather let Google verify your age by giving it your credit card information, photo ID, or digital ID, or by letting Google send your email address to a third-party verifier.

Email Usage

If you choose to provide your email address, Google sends it on to a company called VerifyMy. VerifyMy will use your email address to see if you've done things like get a mortgage or paid for utilities using that email

address. If you use Gmail as your email provider, this may be a privacy-protective option with respect to Google, as Google will then already know the email address associated with the account. But it does tell VerifyMy and its third-party partners that the person behind this email address is looking to verify their age, which you may not want them to know. VerifyMy uses "proprietary algorithms and external data sources" that involve sending your email address to "trusted third parties, such as data aggregators." It claims to "ensure that such third parties are contractually bound to meet these requirements," but you'll have to trust it on that one—we haven't seen any mention of who those parties are, so you'll have no way to check up on their practices and security. On the bright side, VerifyMy and its partners do claim to delete your information as soon as the check is completed.

Credit Card Verification

If you choose to let Google use your credit card information, you'll be asked to set up a Google Payments account. Note that debit cards won't be accepted, since it's much easier for many debit cards to be issued to people under 18. Google will then charge a small amount to the card, and refund it once it goes through. If you choose this method, you'll have to tell Google your credit card info, but the fact that it's done through Google Payments (their regular card-processing system) means that at least your credit card information won't be sitting around in some unsecured system. Even if your credit card information happens to accidentally be leaked, this is a relatively low-risk option, since credit cards come with solid fraud protection. If your credit card info gets leaked, you should easily be able to dispute fraudulent charges and replace the card.

Digital ID

If the option is available to you, you may be able to use your digital ID to verify your age with Google. In some regions, you'll be given the option to use your digital ID. In some cases, it's possible to only reveal your age information when you use a digital ID. If you're given that choice, it can be a good privacy-preserving option. Depending on the implementation, there's a chance that the verification step will "phone home" to the ID provider (usually a government) to let them know the service asked for your age. It's a complicated and varied topic that you can learn more about by visiting EFF's page on digital identity.

Upload ID

Should none of these options work for you, your final recourse is to send Google a photo of your ID. Here, you'll be asked to take a photo of an

acceptable ID and send it to Google. Though the help page only states that your ID "will be stored securely," the verification process page says ID "will be deleted after your date of birth is successfully verified." Acceptable IDs vary by country, but are generally government-issued photo IDs. We like that it's deleted immediately, though we have questions about what Google means when it says your ID will be used to "improve [its] verification services for Google products and protect against fraud and abuse." No system is perfect, and we can only hope that Google schedules outside audits regularly.

TikTok

Inferred Age

If TikTok can guess your age, you may never even see an age verification notification. TikTok first tries to use information you've posted to estimate your age, looking through your videos and photos to analyze your face and listen to your voice. By uploading any videos, TikTok believes you've given it consent to try to guess how old you look and sound.

If TikTok decides you're too young, appeal to revoke their age decision before the deadline passes. If TikTok cannot guess your age, or decides you're too young, it will automatically revoke your access based on age—including either restricting features or deleting your account. To get your access and account back, you'll have a limited amount of time to verify your age. As soon as you see the notification that your account is restricted, you'll want to act fast because in some places you'll have as little as 23 days before the deadline passes.

When you get that notification, you're given various options to verify your age based on your location.

Face Scan

If you're given the option to use facial age estimation, you'll be sent to Yoti, a third-party verification service. Your photo will be uploaded to their servers during this process. Yoti claims that "as soon as an age has been estimated, the facial image is immediately and permanently deleted." Though it's not as good as not having that data in the first place, Yoti's security measures include a bug bounty program and annual penetration testing. However, researchers from Mint Secure found that Yoti's app and website are filled with trackers, so the fact that you're verifying your age *could* be leaked not only to Yoti, but to third-party data brokers as well.

You may not want to use this option if you're worried about third parties potentially being able to know you're trying to verify your age with TikTok. You also might not want to use this if you're worried about a current picture of your face accidentally leaking—for example, if elements in the background of your selfie might reveal your current location. On the other hand, if you consider a selfie to be less sensitive than a photograph of your ID or your credit card information, this option might be better. If you do choose (or are forced to) use the face check system, be sure to snap your selfie without anything you'd be concerned with identifying your location or embarrassing you in the background in case the image leaks.

Credit Card Verification

If you have a credit card in your name, TikTok will accept that as proof that you're over 18. Note that debit cards won't be accepted, since it's much easier for many debit cards to be issued to people under 18. TikTok will charge a small amount to the credit card, and refund it once it goes through. It's unclear if this goes through their regular payment process, or if your credit card information will be sent through and stored in a separate, less secure system. Luckily, these days credit cards come with solid fraud protection, so if your credit card gets leaked, you should easily be able to dispute fraudulent charges and replace the card. That said, we'd rather TikTok provide assurances that the information will be processed securely.

Credit Card Verification of a Parent or Guardian

Sometimes, if you're between 13 and 17, you'll be given the option to let your parent or guardian confirm your age. You'll tell TikTok their email address, and TikTok will send your parent or guardian an email asking them (a) to confirm your date of birth, and (b) to verify their own age by proving that they own a valid credit card. This option doesn't always seem to be offered, and in the one case we could find, it's possible that TikTok never followed up with the parent. So it's unclear how or if TikTok verifies that the adult whose email you provide is your parent or guardian. If you want to use credit card verification but you're not old enough to have a credit card, and you're ok with letting an adult know you use TikTok, this option may be reasonable to try.

Photo with a Random Adult?

Bizarrely, if you're between 13 and 17, TikTok claims to offer the option to take a photo with literally any random adult to confirm your age. Its help page says that any trusted adult over 25 can be chosen, as long as they're holding a piece of paper with the code on it that TikTok provides. It also

mentions that a third-party provider is used here, but doesn't say which one. We haven't found any evidence of this verification method being offered. Please do let us know if you've used this method to verify your age on TikTok!

Photo ID and Face Comparison

If you aren't offered or have failed the other options, you'll have to verify your age by submitting a copy of your ID and matching photo of your face. You'll be sent to Incode, a third-party verification service. In a disappointing failure to meet the industry standard, Incode itself doesn't automatically delete the data you give it once the process is complete, but TikTok does claim to "start the process to delete the information you submitted," which should include telling Incode to delete your data once the process is done. If you want to be sure, you can ask Incode to delete that data yourself. Incode tells TikTok that you met the age threshold without providing your exact date of birth, but then TikTok wants to know the exact date anyway, so it'll ask for your date of birth even after your age has been verified.

TikTok itself might not see your actual ID depending on its implementation choices, but Incode will. Your ID contains sensitive information such as your full legal name and home address. Using this option not only runs the (hopefully small, but never nonexistent) risk of that data getting accidentally leaked through errors or hacking. If you don't want TikTok or Incode to know your name, what you look like, and where you live—or if you don't want to rely on both TikTok and Incode to keep to their deletion promises—then this option may not be right for you.

Everywhere Else

We've covered the major providers here, but age verification is unfortunately being required of many other services that you might use as well. While the providers and processes may vary, the same general principles will apply. If you're trying to choose what information to provide to continue to use a service, consider the "follow the data" questions mentioned above, and try to find out how the company will store and process the data you give it. The less sensitive information, the fewer people have access to it, and the more quickly it will be deleted, the better. You may even come to recognize popular names in the age verification industry: Spotify and OnlyFans use Yoti (just like Meta and Tiktok), Quora and Discord use k-ID, and so on.

Unfortunately, it should be clear by now that **none of the age verification options are perfect in terms of protecting information, providing access to everyone, and safely handling sensitive data**. That's just one of the

reasons that EFF is against age-gating mandates, and is working to stop and overturn them across the United States and around the world.

VPNs Are Not a Solution to Age-Gating Mandates

Age verification laws introduce surveillance systems that threaten everyone's rights to speech and privacy, and introduce more harm than they seek to combat.

Within this context, it is no surprise that Google searches for Virtual Private Networks (VPNs) have skyrocketed. But as more states and countries pass age verification laws, it is crucial to recognize the broader implications these measures have on privacy, free speech, and access to information. **While VPNs may be able to disguise the source of your internet activity, they are not foolproof—nor should they be necessary to access legally protected speech**.

A VPN routes all your network traffic through an "encrypted tunnel" between your devices and the VPN server. The traffic then leaves the VPN to its ultimate destination, masking your original IP address. From a website's point of view, it appears your location is wherever the VPN server is. A VPN should not be seen as a tool for anonymity. While it can protect your location from some companies, a disreputable VPN service might deliberately collect personal information or other valuable data. There are many other ways companies may track you while you use a VPN, including GPS, web cookies, mobile ad IDs, tracking pixels, or fingerprinting.

With varying mandates across different regions, it will become increasingly difficult for VPNs to effectively circumvent these age verification requirements because each state or country may have different methods of enforcement and different types of identification checks, such as government-issued IDs, third-party verification systems, or biometric data. As a result, VPN providers will struggle to keep up with these constantly changing laws and ensure users can bypass the restrictions, especially as more sophisticated detection systems are introduced to identify and block VPN traffic.

The ever-growing conglomeration of age verification laws poses significant challenges for users trying to maintain anonymity online, and have the potential to harm us all—including the young people they are designed to protect.

Can VPNs Be Banned?

Unsurprisingly, politicians have now discovered that people are using Virtual Private Networks (VPNs) to protect their privacy and bypass these invasive laws. Their solution so far has been: to ban uses of VPNs.

For example, in 2025, Wisconsin lawmakers escalated their war on privacy by targeting VPNs in the name of "protecting children" in A.B. 105/S.B. 130. It's an age verification bill that requires all websites distributing material that could conceivably be deemed "sexual content" to both implement an age verification system and also to block the access of users connected via VPN. Another proposed Michigan bill requires "An internet service provider providing internet service in this state [to] actively monitor and block known circumvention tools." Circumvention tools being: VPNs.

This follows a notable pattern: As we've explained previously, lawmakers, prosecutors, and activists in conservative states have worked for years to aggressively expand the definition of "harmful to minors" to censor a broad swath of content: diverse educational materials, sex education resources, art, and even award-winning literature.

Here's Why This Is A Terrible Idea

VPNs mask your real location by routing your internet traffic through a server somewhere else. When you visit a website through a VPN, that website only sees the VPN server's IP address, not your actual location. It's like sending a letter through a P.O. box so the recipient doesn't know where you really live.

So when a state demands that websites "block VPN users from X State," they're asking for something that's technically impossible. Websites have no way to tell if a VPN connection is coming from Milwaukee, Michigan, or Mumbai. The technology just doesn't work that way.

Websites subject to this proposed law are left with this choice: either cease operation in the state, or block *all* VPN users, everywhere, just to avoid legal liability in the state. One state's terrible law could break VPN access for the entire internet, and the unintended consequences of this provision would far outweigh any theoretical benefit.

It's A Privacy Nightmare

Here's what happens if VPNs get blocked: everyone has to verify their age by submitting government IDs, biometric data, or credit card information directly to websites—without any encryption or privacy protection.

We already know how this story ends. Companies get hacked. Data gets breached. And suddenly your real name is attached to the websites you visited, stored in some poorly-secured database, waiting for the inevitable leak. This has already happened, and is not a matter of *if* but *when*. And when it does, the repercussions will be huge.

Forcing people to give up their privacy to access legal content is the exact opposite of good policy. It's surveillance dressed up as safety.

It Won't Even Work

Let's say a state somehow manages to pass a VPN ban. Here's what will actually happen:

People who want to bypass it will use non-commercial VPNs, open proxies, or cheap virtual private servers that the law doesn't cover. They'll find workarounds within hours. The internet always routes around censorship.

Even in a fantasy world where every website successfully blocked all commercial VPNs, people would just make their own. You can route traffic through cloud services like AWS or DigitalOcean, tunnel through someone else's home internet connection, use open proxies, or spin up a cheap server for less than a dollar.

Meanwhile, everyone else (businesses, students, journalists, abuse survivors, regular people who just want privacy) will have their VPN access impacted. The law will accomplish nothing except making the internet less safe and less private for users.

Nonetheless, as we've mentioned previously, while VPNs may be able to disguise the source of your internet activity, they are not foolproof—nor should they be necessary to access legally protected speech. Like the larger age verification legislation they are a part of, VPN-blocking provisions simply don't work. They harm millions of people and they set a terrifying precedent for government control of the internet. More fundamentally, legislators need to recognize that age verification laws themselves are the problem. They don't work, they violate privacy, they're trivially easy to circumvent, and they create far more harm than they prevent.

A False Dilemma

People have (predictably) turned to VPNs to protect their privacy as they watch age verification mandates proliferate around the world. Instead of taking this as a sign that maybe mass surveillance isn't popular, lawmakers

have decided the real problem is that these privacy tools exist at all and are trying to ban the tools that let people maintain their privacy.

Let's be clear: lawmakers need to abandon this entire approach. The answer to "how do we keep kids safe online" isn't "destroy everyone's privacy." It's not "force people to hand over their IDs to access legal content." And it's certainly not "ban access to the tools that protect journalists, activists, and abuse survivors."

If lawmakers genuinely care about young people's well-being, they should invest in education, support parents with better tools, and address the actual root causes of harm online. What they shouldn't do is wage war on privacy itself. Attacks on VPNs are attacks on digital privacy and digital freedom. And this battle is being fought by people who clearly have no idea how any of this technology actually works.

What Can You Do?

If you are navigating protecting your privacy or want to learn more about VPNs, EFF provides a comprehensive guide on using VPNs and protecting digital privacy–a valuable resource for anyone looking to use these tools.

No one should have to hand over their driver's license just to access free websites. EFF has long fought against mandatory age verification laws, from the U.S. to Canada and Australia. And under the context of weakening rights for already vulnerable communities online, politicians around the globe must acknowledge these shortcomings and explore less invasive approaches to protect all people from online harms.

Dozens of bills currently being debated by state and federal lawmakers could result in dangerous age verification mandates. We will resist them. We must stand up against these types of laws, not just for the sake of free expression, but to protect the free flow of information that is essential to a free society. Contact your state and federal legislators, raise awareness about the unintended consequences of these laws, and support organizations that are fighting for digital rights and privacy protections alongside EFF, such as the ACLU, Woodhull Freedom Foundation, and others.

Zero-Knowledge Proofs Are Not A Solution to Age-Gating Mandates

For a lot of people, having physical government documentation like a driver's license, passport, or other ID is not a simple binary of having it or not. Physical ID systems involve hundreds of factors that impact their accuracy and validity, and everyday situations occur where identification attributes can change, or an ID becomes invalid or inaccurate or needs to be reissued: addresses change, driver's licenses expire or have suspensions lifted, or temporary IDs are issued in lieu of obtaining permanent identification.

The digital ID systems currently being introduced potentially solve *some* problems like identity fraud for business and government services, but leave the holder of the digital ID vulnerable to the needs of the companies collecting such information. State and federal embrace of digital ID is based on claims of faster access, fraud prevention, and convenience. But with digital ID being proposed as a means of online verification, it is just as likely to block claims of public assistance and other services as facilitate them. That's why legal protections are as important as the digital IDs themselves. To add to this, in places that lack comprehensive data privacy legislation, verifiers are not heavily restricted in what they can and can't ask the holder. In response, some privacy mechanisms have been suggested and few have been made mandatory, such as the promise that a feature called Zero Knowledge Proofs (ZKPs) will easily solve the privacy aspects of sharing ID attributes.

Zero Knowledge Proofs: The Good News

The biggest selling point of modern digital ID offerings, especially to those seeking to solve mass age verification, is being able to incorporate and share something called a Zero Knowledge Proof (ZKP) for a website or mobile application to verify ID information, and not have to share the ID itself or information explicitly on it. ZKPs provide a cryptographic way to not give something away, like your exact date of birth and age from your ID, instead offering a "yes-or-no" claim (like above or below 18) to a verifier requiring a legal age threshold. More specifically, two properties of ZKPs are "soundness" and "zero knowledge." Soundness is appealing to verifiers and governments to make it hard for an ID holder to present forged information (the holder won't know the "secret"). Zero-Knowledge can be beneficial to the holder, because they don't have to share explicit information like a birth date, just cryptographic proof that said information exists and is valid. There have been recent announcements from major tech companies like Google who

plan to integrate ZKPs for age verification and “where appropriate in other Google products”.

Zero Knowledge Proofs: The Bad News

What ZKPs don’t do is mitigate verifier abuse or limit their requests, such as over-asking for information they don’t need or limiting the number of times they request your age over time. They don’t prevent websites or applications from collecting other kinds of observable personally identifiable information like your IP address or other device information while interacting with them.

ZKPs are a great tool for sharing less data about ourselves over time or in a one time transaction. But this doesn’t do a lot about the data broker industry that already has massive, existing profiles of data on people. We understand that this was not what ZKPs for age verification were presented to solve. But it is still imperative to point out that utilizing this technology to share even more about ourselves online through mandatory age verification establishes a wider scope for sharing in an already saturated ecosystem of easily linked, existing personal information online. Going from presenting your physical ID maybe 2-3 times a week to potentially proving your age to multiple websites and apps every day online is going to render going online itself as a burden at minimum and a barrier entirely at most for those who can’t obtain an ID.

Protecting The Way Forward

Mandatory age verification takes the potential privacy benefits of mobile ID and proposed ZKPs solutions, then warps them into speech chilling mechanisms.

Until the hard questions of power imbalances for potentially abusive verifiers and prevention of phoning home to ID issuers are addressed, **these systems should not be pushed forward without proper protections in place**. A more private, holder-centric ID is more than just ZKPs as a catch all for privacy concerns. The case of safety online is not solved through technology alone, and involves multiple, ongoing conversations. Yes, that sounds harder to do than age checks online for everyone. Maybe, that’s why this is so tempting to implement. However, we encourage policy and law makers to look into what is best, and not what is easy.

Help Us Fight Back

Age verification laws give governments and corporations the power to decide what we can see, say, and share—and to track who we are while we do it.

EFF is fighting back by challenging these unconstitutional laws in court, pushing lawmakers to reject them before they pass, and helping communities protect their rights to privacy, anonymity, and free expression online.

Ready to join us? Urge your state lawmakers to reject harmful age-verification laws. Call or email your representatives to oppose KOSA and any other proposed federal age-checking mandates.

Beyond that, the most important thing you can do right now is to educate yourself on the risks and harms of these bills, and make sure your community does the same. Make your voice heard by talking to your friends and family about what we all stand to lose if the age-gated internet becomes a global reality. And get loud in your opposition to these oppressive digital regimes. Because the fight for a free internet starts with us.

Reading List

This reading list includes a selection of recommended blogs, articles, and materials from external organizations, designed to deepen your understanding and expand your knowledge on age verification. These resources complement the content shared in this Resource Guide and Resource Hub and offers diverse perspectives and expert insights:

Basics:

- *Woodhull Freedom Foundation:* Fact Checked - Age Verification
- *Free Speech Coalition:* Age Verification Bill Tracker
- *Manshuya Foundation:* Statement on Age Gating
- *UNICEF*: Age Restrictions Alone Won't Keep Children Safe Online
- *IAPP:* Mind the gap: Understanding age verification and assurance
- *Fight For the Future:* StopOnlineIDChecks.org
- *EFF:* "Privacy First: A Better Way to Address Online Harms"
- *EFF:* "Behind the One-Way Mirror: A Deep Dive Into the Technology of Corporate Surveillance"
- *MonkeyExplains* on YouTube: "Discord's Data Breach in 11 Minutes & 56 Seconds"

Further Reading:

- Eric Goldman, *Stanford Technology Law Review*: "The "Segregate-and-Suppress" Approach to Regulating Child Safety Online"
- Steven M. Bellovin, *Georgetown University Institute of Technology Law & Policy*: "Privacy-Preserving Age Verification—And Its Limitations"
- *Open Rights Group:* Age Verification Facts
- *Collaborative Research Center for Resilience:* Public Infrastructure, Technology and Democracy page
- Ella Dorn, *New Statesman:* "Big Tech is the only winner of the Online Safety Act"
- Jasmine Mithani, *19th news*: "Age Verification on Porn Sites is Putting Queer Adult Industry Workers at Risk"
- *Center for Democracy & Technology:* "Teen and Parent Perspectives on Approaches to Age Verification"
- *Sunday Times:* "Online Clampdown Puts Sites Like Mine at Agitators' Mercy"
- Taylor Lorenz, *UserMag:* "We Must Fight Age Verification With All We Have"

- Jon Severs, *TES Magazine*: "Is Jonathan Haidt Right About Smartphones?"
- Sarah Scheffler, *Communications of the ACM:* "Age Verification Systems Will Be a Personal Identifiable Information Nightmare"
- *New America:* "Challenges with Age Verification"
- Shawn Musgrave, *Intercept:* "Project 2025 Co-Author Caught Admitting the Secret Conservative Plan to Ban Porn"
- Cory Doctorow, *Medium*: "'Privacy preserving age verification' is Bullshit"
- Adi Robertson, *Verge*: "The Right to Anonymity is Powerful, and America Is Destroying It"
- Dia Kayyali & Jasmine Mithani, *Tech Policy Press:* "Age Verification Is Locking Trans People Out of the Internet"
- The TBOTE Project: "Age Verification Lobbying: Dark Money, Model Legislation & Institutional Capture"