

Authors

Karen Abrams

Salima Bacchus-Hinds

Education Before Regulation

An Evidence-Based Assessment of Guyana's Proposed Social Media Age Verification Regime

Executive Summary

In June 2026 the Government of Guyana launched a national consultation toward a Child Online Safety and Ethical Technology Framework. The proposals under consideration include a minimum social media age of sixteen, parental consent for younger users, mandatory age verification, stronger platform accountability, and digital literacy programmes in schools, together with a Child Online Protection Authority.

Protecting children online is an important and legitimate objective. Children face genuine risks, including cyberbullying, online exploitation, privacy violations, exposure to harmful content, and digital fraud, and governments worldwide are working out how best to respond. This brief does not question that aim. It examines whether the evidence required to justify a nationwide age verification mandate has yet been placed before the public, and whether age verification should be the centrepiece of Guyana's response.

It is worth stating plainly where Abrams & Associates agrees with the Government. The consultation already pairs age restrictions with digital literacy, parental engagement, and platform accountability — the very tools international child-protection bodies most strongly endorse. The disagreement is narrower, and therefore more precise: it concerns sequencing and proportionality. A mandatory, identity-linked age verification regime is the single most intrusive and least reversible element of the package, and it is the element for which the supporting evidence is currently the thinnest.

This brief identifies three concerns.

First, the evidence base is thin in both directions. No publicly available Guyana-specific data has been presented establishing the prevalence of social-media-related harm among Guyanese children — but neither does the country yet have data showing such harm is rare. The honest conclusion is not that regulation is unwarranted, but that Guyana should measure before it mandates.

Second, many Guyanese children remain digitally excluded. For a substantial share of the child population the pressing problem is too little connectivity, not too much. A policy centred on restriction risks addressing a condition that, for many children, does not yet describe their reality, while doing little to advance the broader goal of preparing them for a digital economy.

Third, age verification carries real and now-observable costs. Australia became the first country to enforce an under-16 ban in December 2025, giving Guyana a live test case rather than a hypothetical one. Its experience shows both what such a regime can do and the privacy, accuracy, and circumvention problems it raises — problems Guyana should weigh deliberately before committing to the same path.

The recommendation is an evidence-first sequence: establish the baseline, assess the privacy risks, expand literacy and child-protection capacity, and pilot before any nationwide verification mandate takes effect.

1. The Evidence Question

Good policy begins with a clear account of the problem it intends to solve. Before introducing measures that could affect how hundreds of thousands of citizens access information and one another, policymakers should be able to answer a set of basic questions: how many Guyanese children use social media, and on which platforms; how many encounter cyberbullying, exploitation, or harmful content; and what measurable effect, if any, this has on their wellbeing and development.

To date, no comprehensive national study addressing these questions appears to have been published. This absence should be read carefully. It does not show that the problem is small — it shows that its scale has not yet been measured. That cuts in two directions at once, and it is important to be candid about both.

It would be inconsistent for this brief to fault the Government for relying on assumptions and then rest its own case on the assumption that harm is negligible. The point is not that Guyanese children are safe online; it is that no one has yet demonstrated, with local data, how safe or unsafe they are. A regime as consequential as mandatory age verification should rest on more than inference imported from abroad.

And importation is the real risk. Research from the United States, Europe, and Australia cannot be transplanted to Guyana without adjustment: usage patterns, connectivity, household structure, schooling, and cultural norms all differ. The international literature itself counsels against sweeping conclusions. The National Academies of Sciences, Engineering, and Medicine found in 2024 that social media's effects vary considerably across individuals and contexts, depending on platform design, family environment, peer relationships, and individual vulnerability. Orben and Przybylski (2019), analysing data on more than 350,000 adolescents, found that digital-technology use accounted for well under one percent of the variation in adolescent wellbeing — an effect they noted was comparable in size to wearing glasses or eating potatoes. The evidence is real, but it is neither uniform nor large, which is precisely why local measurement matters.

Before regulating access, Guyana should first establish, with its own data, the problem it is regulating.

2. Guyana's Digital Reality

Much of the global debate assumes near-universal, always-on internet access. That assumption does not fully describe Guyana. Connectivity has expanded markedly in recent years, but substantial disparities remain, and for many children the day-to-day experience is shaped by limited household internet, shared devices, high data costs, unreliable connections, and regional inequality.

UNICEF and the International Telecommunication Union have estimated that roughly two-thirds of school-age children worldwide lack internet access at home. That is a global figure, and — consistent with the argument of Section 1 — it should be treated as a prompt to gather Guyanese data, not as a substitute for it. The conceptual point holds regardless of the exact local number: for a meaningful share of Guyanese children, the binding constraint is access, not excess.

This matters because a policy should address the conditions children actually face. A child who cannot reliably reach educational resources online faces a different set of risks from one spending six hours a day on social media, and a framework built mainly around restriction speaks more directly to the second child than the first. UNESCO's Global Education Monitoring Report makes the related point that technology policy should advance educational opportunity, digital literacy, and equitable access — developmental goals that a restriction-led approach can unintentionally crowd out.

3. The Australian Test Case

Guyana is not legislating in the dark. Australia's Online Safety Amendment (Social Media Minimum Age) Act, passed in November 2024, took effect on 10 December 2025 and made Australia the first country to enforce a nationwide under-16 social media ban. It is the closest real-world analogue to what Guyana is now considering, and its early record deserves close study rather than a passing reference.

Several features are instructive. The Australian law places the obligation on platforms, not on children or parents, with penalties reaching tens of millions of dollars for failures to take reasonable steps. By mid-December 2025 the national regulator reported that covered platforms had removed access to roughly 4.7 million under-16 accounts. To comply, major platforms turned to a mix of facial age estimation and identity-document checks — the same verification methods, and therefore the same privacy questions, that any Guyanese regime would have to confront. Within months the regulator was already amending the rules to better target the obligation, an early signal of how much implementation detail such a scheme requires.

Two cautions emerge that bear directly on Guyana. First, enforcement is technically and administratively demanding even for a wealthy state with a well-resourced regulator; replicating it in a smaller administration is a serious undertaking that should be costed honestly. Second, the open questions in Australia — the accuracy of age-estimation technology, the privacy implications

of document checks, and the ease with which determined teenagers circumvent the rules — remain unresolved. Guyana has the advantage of being able to watch that evidence develop before committing to the same approach. It should use it.

4. Can It Work? Definition, Evasion, and Enforcement

A framework should be judged not only by its intentions but by whether it can function once it meets reality. Three design features decide that: how the law defines its target, how easily that target is evaded, and how it is enforced. On each, the international experience gives Guyana reason for caution.

The definition cannot hold still

Australia defines an age-restricted platform functionally — a service whose sole or significant purpose is to enable online social interaction, that lets users link to and interact with others, and that lets them post material. That sounds tidy until it meets the real world, at which point the government must carve out whole categories by hand: messaging, gaming, business collaboration, education, and health are all exempt. The regulator then maintains a non-exhaustive list of in-scope services and leaves every other provider to self-assess whether it is caught, while the Minister retains discretion to add or exempt platforms. The result is a boundary that is arbitrary, contested, and constantly moving — and it has already been revised more than once. A small administration would inherit the standing burden of policing that line.

Why is gaming exempt?

The exemption for gaming and messaging is the regime’s sharpest inconsistency, and it is worth asking the Government to justify it directly. Australia excluded Roblox, Discord, and WhatsApp from its under-16 ban — yet these are among the most socially interactive spaces children occupy, and they are where a great deal of documented harm actually occurs. Roblox reports more than 150 million daily users, most of them under eighteen. In 2025 both Roblox and Discord faced lawsuits alleging their platforms served as a hunting ground for predators who groomed minors, and several jurisdictions have pursued them over child-exploitation risks. Australia’s own eSafety Commissioner placed Roblox “on notice” and pressed it to switch off direct chat for unverified users and block adults from contacting under-16s, observing that platforms popular with children inevitably become popular with adult predators.

This is revealing. The Government had to chase the exempted platforms through other regulatory powers — effectively conceding that its own age-restriction line does not track where the harm is. A definition that exempts the venues of the very harm it targets is sorting by category label, not by risk. The line between “social media” and “social gaming” has largely dissolved: both offer chat, user-generated content, contact with strangers, algorithmic surfacing, and in-game spending. If Guyana’s rationale is harm, it should be able to explain why gaming receives a pass — and if it cannot, that is a sign the definition is doing something other than protecting children.

Evasion — and the rise of build-your-own

The most basic form of evasion is simply lying, and children do it at scale. The UK regulator Ofcom found that roughly 22 percent of eight-to-seventeen-year-olds present themselves as eighteen or over, and when journalists tested major platforms they opened adult accounts with no proof requested at all. A law that rests on a self-declared birth year is therefore largely cosmetic; to bite, it must demand real verification, which returns us to the privacy problem of the sections that follow.

But evasion is no longer confined to lying or to VPNs. With AI-assisted development — now widely called “vibe coding,” the practice of building working software by describing it in plain language to an AI assistant, with little or no traditional programming skill — a single motivated teenager can stand up a functional chat app, forum, or self-hosted social network in an afternoon. What recently required a development team now requires a prompt. This is fatal to the regime’s enforcement architecture, because that architecture rests entirely on fining a corporate provider. There is no company to fine when the “platform” is a hobby project run by a sixteen-year-old, quite possibly hosted outside Guyana’s jurisdiction. The supply of new social spaces is now effectively unbounded, and it can originate from the very minors the law is meant to restrict.

This also collides with the brief’s own digital-economy argument. A definition broad enough to capture these home-grown services would also capture exactly the entrepreneurial, creative software-building that Guyana says it wants its young people to learn. A definition narrow enough to spare young builders leaves the evasion route wide open. The regime risks being over-inclusive and under-inclusive at the same time — sweeping up the students it should be celebrating while missing the spaces to which children actually migrate.

The penalty trap

Once it is accepted that children will evade, the law has only two ways to respond: verify everyone, or penalise the evaders. The first is the privacy problem. The second raises the question of who, exactly, is to be punished.

Across every regime examined, the answer is the same: penalties fall on platforms, never on families. Australia fines providers up to A\$49.5 million and imposes no consequences on the parents or children who breach the rules; United States state laws vest enforcement in attorneys general with civil penalties of roughly five to fifty thousand dollars per violation against companies; and lying about one’s age is not in itself a crime, but a breach of a platform’s terms of service, answered with account deletion rather than legal sanction.

So Guyana should ask itself plainly: is it willing to criminalise parents and children? No government has been, and for three sound reasons. It would be practically unenforceable across hundreds of thousands of household decisions. It would be regressive, falling hardest on the less digitally literate and lower-income families already identified in Section 2 as being on the wrong side of the digital divide. And it would be politically and morally untenable to make offenders of teenagers and their parents over a social media account. Yet declining to penalise families — the correct choice — is precisely what throws the entire enforcement burden back onto population-

wide verification. The humane path and the intrusive path are bound together: closing the first forces you onto the second.

The regime is therefore caught in a vice. Rely on self-declaration and it is ineffective. Mandate real verification and it is intrusive, imposed on the whole population to catch a minority, and still evadable by anyone able to build or migrate to an alternative. Penalise families and it is unenforceable and unjust. None of these is a comfortable place to rest, and a small state should be confident it can live with the trade-offs before it commits. This is the practical case for the sequencing this brief recommends: demonstrate, at pilot scale, that the design can actually work before imposing it on the nation.

5. How Verification Would Actually Work

The framework requires age verification, which raises a practical question the public deserves to see answered: how would a platform distinguish a fifteen-year-old from a sixteen-year-old in Guyana? Broadly, three families of method exist, and each carries trade-offs.

Document-based verification — national ID cards, passports, birth certificates, or driver’s licences — is the most accurate but the most privacy-intensive, because it links a verified identity to online activity.

Facial age estimation infers age from an image and can avoid storing identity documents, but its accuracy falls precisely at the ages that matter most here, near the fifteen-to-sixteen boundary, and its performance can vary across demographic groups.

Third-party and privacy-preserving services — including “double-blind” or tokenised systems designed so that neither the platform nor the verifier holds a complete identity-plus-activity record — are the most protective in principle. But they shift the risk onto external vendors whose governance, security, and longevity Guyana would need to oversee, and they remain comparatively unproven at national scale.

None of these is a free option, and it would be a mistake to assume that age verification must mean a centralised identity database. The point is the opposite: the choice of method determines the privacy cost, and the public is entitled to know which method is intended, and on what evidence it was chosen, before a mandate is enacted.

6. Privacy and Civil Liberties

Privacy is not a secondary consideration here; it is central. Any verification system necessarily collects, processes, stores, transmits, and must protect personal information, and citizens are entitled to clear answers before that machinery is built. What information will be collected, and by whom? Where will it be stored, and for how long? Can it be accessed by government agencies or compelled through legal process? What happens in the event of a breach?

These questions matter all the more because identity-linked systems have a documented tendency to expand beyond their original purpose — a pattern privacy scholars call “function creep.” A database built to confirm that a user is sixteen can, over time, become a tool for other forms of

monitoring unless its limits are fixed in law from the outset. For that reason the privacy implications of any proposed scheme should be assessed independently, and published, before implementation rather than after.

7. Why Wait? Answering the Precautionary Case

A reasonable objection runs as follows: if children may be at risk and the evidence is incomplete, surely caution argues for protecting them now rather than waiting for studies. The instinct is understandable, and it deserves a direct answer rather than silence.

The answer is that the precautionary principle applies to the costs of acting as well as the costs of waiting. Mandatory, identity-linked age verification is not a low-cost, easily reversed measure. It builds standing infrastructure that touches every internet user, raises civil-liberties questions that are difficult to unwind once normalised, and — on the Australian evidence — is far from certain to achieve its aim, given circumvention and the limits of age-estimation technology. Where an intervention is intrusive, durable, and of uncertain effectiveness, the burden of justification should be correspondingly high.

Crucially, caution does not require inaction. The genuinely low-risk, high-value measures — digital literacy, stronger reporting and child-safe design obligations on platforms, better counselling and exploitation-response services — can begin immediately and protect children now, without waiting for any study and without building an identity database. The argument of this brief is therefore not “do nothing until we have data.” It is “do the safe and reversible things now, and reserve the intrusive and irreversible thing until the evidence justifies it.”

8. International Guidance: Prepare, Not Only Restrict

A striking feature of the international literature is that the major child-focused organisations generally do not advocate age restrictions as a standalone solution. UNICEF emphasises children’s rights, safer platform design, platform accountability, parental engagement, and digital literacy; its Australian arm argued during that country’s debate that the better fix is improving platform safety rather than simply delaying access. UNESCO emphasises media literacy, digital citizenship, critical thinking, equitable access, and teacher preparedness. Both proceed from the same premise: children will participate in digital life regardless, so the priority is to prepare them for it, not only to fence them out of it. Guyana’s consultation already reflects much of this thinking; the recommendations below are intended to strengthen that side of the framework relative to the verification mandate.

9. An Alternative Centre of Gravity

Rather than making age verification the centrepiece, Guyana could organise its framework around four pillars that can advance in parallel and largely without privacy risk.

Digital literacy

Embed online safety, privacy awareness, misinformation and AI literacy, and cybersecurity basics into the school curriculum, with materials adapted to Guyanese contexts and languages.

Parent education

Provide practical monitoring guidance, conversation tools, and risk-awareness resources, recognising that many parents are themselves new to these platforms.

Platform accountability

Press for stronger reporting tools, child-safe default settings, and transparency reporting — obligations that fall on well-resourced companies rather than on families or children.

Child protection services

Strengthen counselling, school-based support, and dedicated mechanisms for responding to online exploitation, so that children already being harmed have somewhere to turn today.

Recommendations

The following sequence is ordered so that protective, low-risk measures begin immediately while the case for a verification mandate is tested.

1. **Conduct a National Child Online Safety and Connectivity Baseline Study** commissioned through an independent institution with published methodology, on a defined timeline, to measure usage, harm, and connectivity among Guyanese children. This is the evidentiary foundation for everything else.
2. **Commission and publish a Privacy Impact Assessment** of any proposed verification mechanism, conducted independently and published in full, before — not after — a mandate is enacted.
3. **Begin nationwide expansion of digital-literacy instruction** now, as a no-regret measure that does not depend on the study's findings.
4. **Launch parental-education programmes** now, with materials suited to local contexts.
5. **Strengthen child-protection services** now, including counselling and a clear exploitation-response pathway.
6. **Pilot any age-verification mechanism** — limited in geography and duration, with pre-defined success measures — before any nationwide rollout, and study Australia's implementation experience throughout.

7. **Build in independent oversight and hard legal limits** with statutory limits on data retention and secondary use, and an independent oversight body, written into law before any system goes live.
8. **Publish annual outcome reports** on the outcomes of each measure, so that the framework can be adjusted on evidence rather than assumption.

Conclusion

Protecting children online is a legitimate and urgent objective, and the Government is right to take it up. The question this brief raises is not whether to act, but in what order and with what proportionality.

At present the scale of social-media-related harm among Guyanese children has not been publicly measured, a significant share of those children remain under-connected rather than over-connected, and the one country that has enforced a comparable regime is still working through its privacy, accuracy, and enforcement problems in real time. Those facts argue for a particular sequence: begin immediately with the protective measures that carry little risk and can help children now, and reserve the intrusive, hard-to-reverse step of mandatory age verification until a baseline study and an independent privacy assessment have shown it to be both necessary and the most effective option available.

The strongest child-protection strategy for Guyana may not begin with age verification. It may begin with measurement, literacy, parental engagement, stronger child-protection services, and the discipline to test before it mandates.

References

- Australian Government, eSafety Commissioner (2026). Social Media Age Restrictions — implementation updates.
- National Academies of Sciences, Engineering, and Medicine (2024). Social Media and Adolescent Health.
- Ofcom (2025). Findings on children’s use of false ages and age assurance on social media (reported via BBC News).
- Online safety and child-exploitation litigation and regulatory action concerning Roblox and Discord (Google search, 2025–2026), including Australian eSafety Commissioner undertakings and US state actions.
- Office of the Surgeon General (2023). Social Media and Youth Mental Health Advisory.
- Online Safety Amendment (Social Media Minimum Age) Act 2024 (Australia).
- Orben, A., & Przybylski, A. K. (2019). “The association between adolescent well-being and digital technology use.” *Nature Human Behaviour*.
- UNESCO (2023). Global Education Monitoring Report: Technology in Education.

- UNICEF & ITU (2020). How Many Children and Young People Have Internet Access at Home?
- Valkenburg, P. M., Meier, A., & Beyens, I. (2022). “Social media use and its impact on adolescent mental health.” *Current Opinion in Psychology*.