Minors Protection Online Requires More Than Verification

DIDAS position paper | Why privacy preserving verification, governed infrastructure and platform accountability must develop together

AUTHORS

Axel Schild (Adnovum), Denis Roio (Jaromil, Plan-B), Raphael Guye (Swisscom), Darrell O'Donnell (Ayra Association), Daniel Säuberli (DIDAS)

As of August 2026 · Submitted to GDC26, Geneva, 1–3 September 2026

DIDAS

Management Summary

The debate over age verification runs two policy questions together: who may access a platform, and what a platform owes the people already on it. Age verification can answer the first question. It says almost nothing about the second. It has nonetheless become the default answer to harmful content, addictive design and the commercial exploitation of minors, which are problems for other regulatory tools.

DIDAS does not reject age verification in principle. Where access must be restricted, verification is required. It runs in parallel with platform and product regulation. What we reject is any architecture that links age to identity at the moment of presentation and hands that link to platforms or third parties. Such architectures produce surveillance regardless of intent, contract or data protection law.

Where verification is unavoidable, a technically defensible path exists: selective disclosure through national identity schemes such as the Swiss e-ID, secured by privacy enhancing technology. That path only works under conditions that do not yet exist. We therefore call for:

- No age verification mandate for social media before a privacy preserving, universally accessible and governance anchored solution is actually available.
- A categorical prohibition on treating minors as a product: no monetisation of their data and no engagement optimising algorithms applied to them.
- Binding platform and product regulation as the universal baseline of protection, independent of age and of parental capability.
- Active Swiss contribution to an emerging global framework of principles and minimal governance for interoperable, privacy preserving verification, through GDC26 and the UN Global Digital Compact.
- Governance frameworks enacted alongside the technical standards, as a precondition for any deployed solution.

Why age verification is the wrong instrument for most of these goals

When politicians call for age verification, they pursue several legitimate goals at once: keeping children away from adult content, curbing addictive design, restoring parental control, and sending a signal. These goals require different instruments, and age verification is the wrong one for most of them.

Deciding which content reaches which user is a moderation and ranking problem. On device filtering and configurable recommendation systems handle it, and none of that needs identity infrastructure. Protecting children through their parents is a delegation problem, solved by permissions enforced inside the platform, which needs no centralised verification either. Age verification is the right tool only in the residual case, where access itself has to be restricted because the service is reserved for adults and no moderation alternative applies: online

gambling, alcohol sales, platforms given over entirely to adult content. Services that mix content can usually deal with the risk through moderation, filtering and product design instead of identity-based access control. Two or more constraints apply where verification is the right tool: the intrusiveness of the check must match the risk of the service. The check must be enforceable at the scale of the service.

Age verification is one instance of an old and unsolved Internet problem: proving something true about yourself without revealing who you are. The Internet was built without a trust layer, and the things that filled the gap, certificate authorities and later the commercial identity providers, filled it by centralising control and monetising the data that resulted. A verification mandate built on that legacy inherits its failure mode. It adds another layer of potential surveillance, whatever the stated purpose.

The methods now being piloted show the same pattern. Document copies, video identification, facial age estimation and credit card checks all collect identity linked data beyond what an age check needs, and all of them build a central dataset that is worth something to attackers, advertisers and governments. This is not a hypothetical consequence. Where users upload identity documents for an age check, those documents accumulate with verification vendors and their support systems, and breaches at such third party providers have already exposed government identity documents collected for age verification. Facial estimation has error rates that fall unevenly across sex and region of birth, and that comes from the models themselves rather than from any setting anyone can correct. Requirements built around documents and smartphones exclude refugees, undocumented people, people without a smartphone and people with limited digital literacy, who are often the people for whom online community matters most.

The lobbying record shows what happens when a mandate is poorly governed: the platforms it is meant to regulate capture it. One platform's federal lobbying spending reached a company record of 26.3 million US dollars in 2025, and part of that effort pushed legislation that shifts the verification burden onto app stores. Weak mandate design produces exactly this: a compliance burden engineered to land on competitors rather than on the platform doing the lobbying, and a verification architecture chosen for its competitive effect rather than its privacy properties. The same risk attaches to any mandate, including one built on the e-ID, if it is enacted without the governance floor and the open, auditable implementation this paper asks for. Without those, a privacy preserving credential is simply available for capture by the next well resourced lobbying effort.

Principles that cannot be traded off

Any verification architecture deployed at scale has to satisfy six principles at once. They cannot be traded against one another: an architecture that satisfies five and fails the sixth becomes a data collection system in practice, whatever its stated purpose. These are recognised

requirements for trustworthy credential infrastructure, anchored in the Trust Over IP stack and in Swiss law, in particular the E-ID Act and the revised FADP.

Privacy by design. A platform receiving an age assertion must be technically incapable of linking it to an identity: no verification logging, no persistent session identifiers, no path to correlate a check with a profile.

Privacy by default. The most protective configuration is always the default, available without any user action. For age verification that means no fallback that requires identity disclosure.

Unlinkability. No verifier, issuer or intermediary should be able to link two verification events to the same person. A contractual promise is not enough; the architecture has to enforce it.

Data minimisation. Only the minimum needed for the purpose may be processed. For age verification that minimum is a single binary value. Anything more serves a different purpose.

Open source and auditability. The full deployment stack, including cryptography, logging, session management, interfaces and retention, must be open and independently auditable. Closed systems cannot be trusted to enforce the guarantees they claim.

Decentralisation. No single issuer, verifier or intermediary may be in a position to observe or aggregate verification across users and sessions. Shared global technical standards are needed, and so is legitimate, accountable governance at the level where the decisions affect people.

Technical trust and human trust are different things, however inseparable in Digital Trust Frameworks. Cryptography answers whether a credential is intact. It cannot answer whether the issuer was authorised to make the statement, whether the attribute was established correctly, or whether the verifier was entitled to ask. Those questions are settled by governance: mandate, accountability and redress.

The principles also explain why well-meaning mandates fail. A law that requires "privacy preserving" verification, when no such verification is technically available, constrains none of the participants in an ecosystem. A cryptographically sound system deployed without rules defeats its own privacy goals, as they cannot be enforced legally. Neither side works without the other, so technology and governance have to be designed, evaluated and enforced together. The same holds internationally across jurisdictions: shared protocols are useless without shared and mutually recognised thin global governance layers as well as standards.

What the Swiss e-ID makes possible

The Swiss e-ID is built on a different architecture. Its infrastructure, swiyu, is open source and publicly auditable, and it rests on selective disclosure: credentials in the SD-JWT VC format, issued and presented over the OpenID4VCI and OpenID4VP protocols, with issuers and verifiers identified by decentralised identifiers using did:webvh. Selective disclosure works at the

level of the individual claim. The holder reveals only the attributes the verifier asked for, and everything else stays an opaque signed value the verifier never sees.

For age verification that is the right primitive. The e-ID is designed to transmit a single binary assertion: over or under a given threshold. It transfers no name, no date of birth, no document number, no biometric and no reusable session identifier. The verifier learns one bit and gains no handle that ties the holder to any other interaction.

The e-ID is also designed to be unlinkable across presentations. The Federal Council has determined that it must meet this requirement, and the mechanism is batch issuance: each holder is provisioned with many single-use credentials, so that no two presentations share a credential or a key pair. A verifier therefore cannot correlate a person across services or sessions, and two verifiers cannot combine their records to do so without help from the issuer. That protection sits in the architecture rather than in a contractual promise between platforms, though it still depends on the issuing authority being bound by governance rules against logging and correlation. It is the kind of safeguard our governance floor is meant to secure. The Swiss Profile issuance specification makes batch issuance normative: batch_credential_issuance is RECOMMENDED for any privacy relevant use cases, with a batch size of at least ten. The swiyu stack is open source, so the delivered system can be audited against the specification.

What is needed beyond the e-ID

DIDAS helped shape the e-ID architecture and supports it. We also want to be clear about what it does not yet solve.

The credential alone is not enough. International platforms are under no obligation to accept the e-ID, and absent that obligation they will default to the methods that serve their data interests. A nationally issued credential, however good, stays irrelevant to those platforms until there is a globally recognised framework of acceptable verification practices. There are also access gaps. The wallet currently requires a mainstream smartphone, the e-ID is rightly voluntary so part of the population may never hold one, and the governance framework for deployment is still being written. A mandate premised on the e-ID today would therefore work either as a covert e-ID requirement that shuts out everyone without one, or as an opening for less protective alternatives. The mandate and the readiness have to arrive together.

Governance has to be interoperable too. Shared protocols do not help much if each jurisdiction applies its own incompatible rules to them. A structurally neutral, nonprofit registry of registries, of the kind the Ayra Association coordinates, would let a credential issued under Swiss governance be validated against another jurisdiction's requirements without bilateral agreements. Only a neutral nonprofit can win the trust of sovereign and competing actors. UNCITRAL adopted its Model Law on the Use and Cross-border Recognition of Identity Management and Trust Services (MLIT) on 7 July 2022, endorsed by the General Assembly. Extending that instrument, or its Guide to Enactment to privacy preserving attribute verification,

would define what qualifies as privacy preserving, what governance obligations any mandate must carry and the basis for mutual recognition. Switzerland has a privacy by design infrastructure in public beta and a record in international standard setting, which puts it in a good position to propose this and to align its own regulation with the EU Digital Services Act by design. That would let it take the benefit of the Act's enforcement leverage as a market rather than as a member.

The stronger lever is product regulation. The debate mostly stops at who is allowed on. The harder question is what platforms may do to the people already there, and age verification says nothing about that. The dominant business model converts attention into advertising value, which makes the user the product. DIDAS holds as an absolute that minors may never be treated as a product, and neither consent nor parental permission can license it. In practice that means no personal data of minors collected, retained or processed for commercial purposes, including profiling, ad targeting, brokerage and the training of recommendation models. It also means that engagement optimising mechanisms, infinite scroll, variable rewards, streaks and the rest, may not be applied to minors, with compliance demonstrable through independent audit. This baseline binds every user whatever their age, and whether or not a parent has the capacity to configure anything. It is the universal floor, and it is what makes a lighter touch elsewhere defensible. Product safety is the closer analogy: a version of the service that children may use must not contain these mechanisms. We want a safe product, not a product that children are kept out of.

Guardianship is a limited, voluntary complement to that baseline: a parent can grant permissions for a child through a delegation credential bound once, at setup, to the child's account. The child then presents nothing per session, and the platform learns the policy without learning either party's identity. This is a legitimate addition, but it cannot carry the universal protection, because it only reaches families with the literacy, the devices and the wallets to use it. On its own it would leave the least advantaged children the least protected. The product regulation baseline above is what prevents that, and it is why guardianship can stay optional.

The credential format is only part of what guardianship needs. A delegation works across services only if every digital system exposes a standard interface through which a guardianship policy can be invoked, read and revoked, so that a parent sets the conditions once across the services consumed by the minor. What exists today is proprietary per service. Open specification on a protocol that binds a delegation to an account, a shared vocabulary for what a permission means and an interface through which the policy can be read by the parent, by the child and by an auditor, would be a welcome starting point. GDC26 convenes the bodies that could carry this work, and its Age Assurance Task Force is the obvious home.

Primary responsibility for protection also belongs with the platform rather than with the conduit. Controls at the operating system, browser and network level are useful complements. They help parents block known categories of content, enforce screen-time limits and apply age-appropriate filtering across services. What they cannot reach is the harm that drives much of

this debate. A router or an operating system can block a domain, but it cannot make a feed non-manipulative, stop behavioural profiling inside an application, or check whether a recommendation system is optimising for engagement. These controls are also imperfect: they depend on broad implementation across vendors, and they can usually be bypassed with another device, another network or another service.

Measures at the conduit level therefore cannot substitute for platform accountability. They should stay optional tools that support parental choice rather than become the primary mechanism of protection. Making the conduit responsible for enforcement also carries a further risk: it rebuilds the surveillance and blocking infrastructure that the six principles exist to prevent. Responsibility for harmful design belongs first with the party that designs it and profits from it, the platform, and not with the neutral infrastructure that carries the traffic.

Prioritised demands

Immediate, to prevent deterioration

- No general age verification mandate for social media before a privacy preserving, universally accessible and governance anchored solution is actually available. Readiness is assessed against criteria published in advance, with public reporting against them.
- No age verification mandate at the app store level, which would make two corporations the gatekeepers of digital participation.

Short term, to set direction

- Platform and product regulation now: prohibit addictive and dark patterns; require recommendation systems for minors that neither profile nor optimise for engagement; prohibit advertising targeted at minors; prohibit any monetisation based on minors' data. All of it demonstrable through independent audit.
- Regulatory design in Switzerland that is compatible with the EU Digital Services Act.
- A near-term implementation pathway: deploy e-ID based age verification capability through a fedpol-issued credential with a reduced set of non-identifying attributes.
- Open specification of a guardianship interface: one standard way to invoke, read and revoke a delegation policy across services, pursued through GDC26 and its Age Assurance Task Force.
- Begin the governance floor: pursue, through UNCITRAL and GDC26, an extension of the 2022 Model Law on Identity Management and Trust Services (MLIT), or of its Guide to Enactment, to privacy preserving attribute verification and mutual recognition across jurisdictions.

Medium term, to establish e-ID readiness

- A legal obligation to accept the e-ID as a verification method, contingent on simultaneous technical maturity, governance adequacy, and international recognition.
- Extension of the e-ID to devices that are not smartphones and to alternative operating systems.
- Active Swiss contribution to the emerging global framework through GDC26, the UN Global Digital Compact, the Ayra trust registry coordination network, Trust Over IP governance, and relevant standards bodies, with particular attention to governance interoperability as the missing precondition.

Foundational, non negotiable

- The six principles apply to every deployed solution: privacy by design, privacy by default, unlinkability, data minimisation, open source and auditability, and decentralisation.
- Strict separation: verification must never transmit identity data to platforms or third parties.
- Open source implementation and independent audit of the full deployment stack.
- No transfer of verification data to law enforcement or intelligence services without a judicial order.
- Governance before mandate: no verification obligation enacted without the governance framework that gives it effect being enacted at the same time.

References and further reading

- Swiss e-ID Act (BGEID) and SWIYU ecosystem
  https://eid.admin.ch
- First Person Project White Paper V1.2 (2026)
  https://www.firstperson.network/white-paper
- Trust Over IP Foundation - ToIP Stack and Trust Spanning Protocol:
  https://trustoverip.org
- Jaromil (Dyne.org): "Do Not Turn Child Protection Into Internet Access Control" (2026)
  https://news.dyne.org/child-protection-is-not-access-control/
- Jaromil (Dyne.org): "Age Verification for Humans" (2025)
  https://news.dyne.org/age-verification-for-humans/
- Jaromil (Dyne.org): "Privacy in EUDI"
  https://news.dyne.org/privacy-in-eudi
- Kyle den Hartog: "Decentralizing Age Verification with SSI"
  https://kyledenhartog.com/decentralized-age-verification/
- Joint Statement of Security and Privacy Scientists on Age Assurance, February 2026
  https://csa-scientist-open-letter.org/ageverif-Feb2026
- Digitale Gesellschaft: "Make Platforms Better, Not Exclusive" (May 2026)

- EU Digital Services Act and DSA Task Force on Age Verification
  https://digital-strategy.ec.europa.eu
- Global Digital Collaboration Conference (GDC26), Geneva, 1 to 3 September 2026
  https://globaldigitalcollaboration.org
- UNCITRAL Model Law on the Use and Cross-border Recognition of Identity Management and Trust Services (MLIT, 2022)
  https://uncitral.un.org/en/mlit
- UN Global Digital Compact
  https://www.un.org/global-digital-compact/en
- Nolan McKendry, "Big tech giants battle over competing age verification bills"
  https://justthenews.com/nation/states/center-square/big-tech-giants-battle-over-competing-age-verification-bills
- Data from Open Secrets on Meta Lobbying
  https://www.opensecrets.org/federal-lobbying/clients/meta?client_id=197457&cycle=2025