EXECUTIVE BRIEF

The Admission Layer

Age assurance asks who is entering. Admission asks what is allowed inside.
A proposal for African co-design of an interoperable child-safety metadata standard.

AGE ASSURANCE	ADMISSION
Who is this?	**What is this?**
Establishes whether the user is a child, and how old. Necessary. Actively being built. Contested on privacy, accuracy and inclusion grounds.	Establishes whether an item of content is eligible for the environment the child has entered. Decided once, before publication, and recorded machine-readably.

These are orthogonal predicates. Even a perfect answer to the first tells you nothing about the second.

A child correctly identified as eleven years old, admitted to an environment nobody has classified, is not protected. The identity side of age-appropriate access is under active construction across the field. The content side — a portable, machine-readable answer to whether this material is appropriate for this age in this jurisdiction — is comparatively neglected. This brief concerns that second layer, and proposes that African institutions determine what it should contain.

What the admission layer asks of a child: nothing

The admission decision is made once, upstream, by a publisher or curator who knows the material and its language, and is then carried as metadata. At the point of use it requires no identity document, no age estimate, no behavioural profiling, and no classifier running over a language it was never trained on. This matters as a rights position, not only an engineering one: it honours the child's privacy under Article 16 of the Convention on the Rights of the Child rather than trading privacy for protection.

It also matters under conditions that are ordinary across much of the continent. Classification made in advance does not depend on connectivity at the moment of access, on compute available to a low-end device, on reviewer capacity in language, or on an adult being present. Where civil registration is incomplete, age assurance alone does not fail neutrally — it bears hardest on children with the least documentation. An admission layer asks nothing of the child, and therefore introduces no additional exclusion based on a child's ability to prove identity or age.

Why Africa is where this should be designed — not where it should be tested

A standard that survives only in the conditions it was written for is not a standard. The case for African co-design is that the continent presents, together, the conditions under which a content-classification schema is hardest to get right and most valuable if it holds: extraordinary linguistic diversity; mobile-first and largely prepaid connectivity; shared household devices; the fastest growth in child connectivity anywhere; multiple legal systems with genuinely different national content norms; and a developed regional children's-rights architecture against which the work can be measured.

If a schema holds across Kiswahili, Kinyarwanda, English and French, across contrasting national thresholds, on shared handsets and intermittent connections, it has demonstrated portability. If it is built elsewhere and brought, it has demonstrated nothing. This is a substantive claim about standards engineering before it is a political one.

The diagnosis is not ours: policy is outpacing practice

Writing for CIPESA in June 2026, Patricia Ainembabazi observes that child online safety has returned to the front of African and global digital governance, but that policy commitments are multiplying while implementation continues to lag — a gap felt most acutely in Africa, where connectivity is expanding faster than child protection systems can keep pace.[1] That is a diagnosis made by African civil society, not by us, and this proposal is offered as one answer to it. A machine-readable standard is precisely an implementation instrument: it is what converts a stated principle into something a ministry can procure against, a publisher can comply with, and an auditor can test.

The instruments to be implemented already exist and this proposal is subordinate to them. The African Union Child Online Safety and Empowerment Policy, adopted in May 2024, establishes principles of children's safety, privacy and participation to guide member states in developing national strategies, and categorises risk on a four-part model of content, contact, conduct and consumer risks.[2] An admission layer is a concrete implementation instrument for the *content* element of that model — not a new framework competing with it. The Global Digital Compact commits states to prioritise national online child safety policies and *standards* by 2030;[1] standards of that kind have to be drafted, tested and maintained by somebody, in time.

The Africa Taskforce on Child Online Protection, launched by UNICEF and the GSMA at Kigali in October 2025 as a multi-stakeholder platform with formal youth participation and membership including regional telecommunications bodies,[3] convenes precisely the actors this layer would require: regulators, mobile operators, civil society and young people. We are not proposing a parallel structure. We are proposing a candidate technical artifact that such a body might choose to develop, amend, adopt, or reject.

The ground-level evidence base is also further advanced than is often assumed. The Digital Rights Alliance Africa study of child protection and safety online spans ten countries — Algeria, Botswana, Egypt, Ghana, Kenya, Nigeria, Rwanda, South Africa, Tanzania and Uganda — and its recommendations point in two directions this design has tried to follow: that telecommunications companies should submit their protection measures to independent audit, and that children should be actively engaged as participants in the co-design of digital solutions rather than consulted after the fact.[1] The conformance architecture and the participation condition set out below are attempts to build to those recommendations, not around them.

What the standard is, and what it is not tied to

The standard is a controlled vocabulary and machine-readable schema for age-appropriateness — thirty-nine terms at version 1.0, language-independent by construction, published as a draft for comment. It is closer in kind to a cataloguing standard than to a filtering product: it describes content so that many different systems can act on the description consistently.

That distinction carries weight. The schema is not synonymous with any single deployment. Once age-appropriateness is expressed in portable metadata, it is usable by mobile operators, schools, libraries, browsers, app stores, publishers, educational platforms, content repositories, parental-control systems, and AI systems serving children. A bounded child-safe environment is one implementation of the standard. It is not the standard, and adopting the standard does not commit anyone to it.

Structure is shared; thresholds are sovereign

A classification standard drafted in one jurisdiction and exported whole carries that jurisdiction's substantive norms with it. The standard therefore separates two layers. **Structure** — fields, value types, vocabulary — is shared and versioned in common. **Thresholds** — which values are acceptable at which ages in which country — are set nationally and published in the same machine-readable form. ISO 8601 fixes how a date is written without deciding which dates are holidays.

The Foundation does not determine what any country's thresholds should be, and does not intend to. Where content categories are contested — and several are — the answer is procedural rather than substantive: thresholds are set through national processes, and the shared baseline is developed through a published process with appeal rights and participation, including by children. A standards body that resolved those questions on countries' behalf would deserve the objection it received.

Who classifies, and why anyone should trust them

Moving the decision upstream does not make it self-executing. The same scarcity that starves moderation can starve admission, and a publisher under commercial pressure to publish will not reliably self-certify without consequence. The layer is therefore governed, not merely declared:

- **Tiered entry.** Low-risk categories may self-declare against the vocabulary. Higher-risk categories require independent validation before admission.
- **Sampling and audit.** Declarations are sampled against a reference set, with inter-rater reliability measured per language rather than assumed.
- **Challenge and correction.** Any party may challenge a classification; the publisher corrects or contests it.
- **Graduated consequence.** Notice, cure, suspension pending correction, and ultimately loss of admission — with appeal at every step, and enhanced audit for repeat error.

Deliberate misclassification and persistent error carry the same remedy. This is ordinary conformance governance rather than editorial judgement, and it is the part of the design that most needs African institutional input, because the reference sets and the reliability testing must be built in the languages that matter.

Why this is implementable before the law catches up — and who pays for it

Implementation lags policy for a structural reason: most of what is recommended depends on primary legislation, enforcement capacity, or cross-border cooperation, and each of those moves slowly. An admission standard is unusual in that list because its enforcement is contractual rather than statutory. The remedies are admission, suspension pending correction, and loss of admission — exercised by the body administering the standard against parties who have agreed to be bound by it. No new content-control legislation is required for a voluntary implementation to begin, no regulator is asked for powers it does not have, and no party outside the scheme is compelled to do anything. Participation is established by an admission agreement between the publisher or registrant and the independent body administering the standard; which institutional form that body should take is a question the demonstrator exists to answer, not one we presuppose the Foundation should hold.

There is a live African precedent for the institutional shape this implies. The African Organisation for Standardisation and inABLE are developing Africa's first harmonised Standard for Accessible ICT

Products and Services — a twenty-four-month participatory process, supported by GIZ, engaging youth, women and organisations of persons with disabilities, and aimed at adoption across ARSO's forty-five member states, with fourteen countries having announced support.[4] The governance shape is instructive: the standards body owns the standard and a specialist organisation supplies technical leadership. That division is the one we would propose here. The accessibility work also proceeds on the principle that the people affected are contributors and decision-makers rather than subjects consulted after decisions are made[4] — which is the principle behind the participation condition below, applied to children instead.

This does not displace public law, and is not intended to pre-empt it; it allows implementation to begin while public law develops on its own timetable. Where national statutes do follow, a published, versioned technical standard is incorporable by reference — the ordinary route by which technical standards acquire legal force without a legislature having to enumerate content categories in primary legislation and re-legislate each time technology moves. A parliament acting on the first recommendation is materially better placed if there is a maintained standard to point at. The work can therefore start now and be adopted into law later, rather than waiting on a sequence that has already taken years.

The funding mechanism is the second half of that answer. Recommendations of this kind ordinarily assume that ministries will invest, companies will comply, and civil society will monitor — without saying who pays for the monitoring, year after year. Here the recurring cost of administering the standard falls on the parties who derive commercial value from admission: recurring sources to be tested include fees from participating commercial entities for conformance, certification or admission, and operator revenues associated with child-safe service tiers. Philanthropic and multilateral finance is appropriate for two discrete things — developing the standard and running the demonstrator — and should be able to step out afterwards rather than being required in perpetuity.

Two limits on that claim, stated plainly. The operating model must never make a child's level of protection depend on household income; affordability is a design requirement, not a later concession, and a scheme that protected only the well-served would have failed rather than partly succeeded. And the model is a design intention that the demonstrator exists to test, not a result. Anyone who tells you the unit economics of this are settled has not run it.

What this answers, and what it does not

Against the recommendations set out in the Digital Rights Alliance Africa study,[1] an admission layer speaks directly to the call for independent audit of protection measures, supplies national strategies with a concrete instrument for allocating responsibility across the content chain, gives parliaments a maintained standard to legislate by reference to, generates published operating evidence for civil society to hold everyone to, and treats child co-design as a condition rather than a courtesy. It does not legislate, and it does not touch contact or conduct risks — grooming, bullying, extortion — which are separate elements of the African Union model requiring separate instruments. It is one layer among several, and it should be judged on that basis.

Six conditions, all load-bearing

Content-control infrastructure built for children has been repurposed against adults in more than one jurisdiction. The objection is legitimate and we treat it as a design constraint. The architecture is not defensible without all six of the following, and we would abandon it rather than proceed without them:

- Participation is voluntary and household-elected, never mandatory or governmental.
- No party is compelled to suppress, remove, or modify expression in the open internet.
- The open internet continues to operate unaltered. The layer is additive: it adds a described environment and removes nothing from the undescribed one.
- Policy authority sits with an institutionally independent body operating a published process with appeal rights.
- Thresholds are set nationally; no country adopts another country's content judgements.
- **Children participate in the development and periodic review of the standard, consistent with their evolving capacities** — as participants in co-design, not as consultees after the fact.[1]

Honest limits

Content that has not satisfied the admission standard is ineligible for the environment. That is an architectural property, not a guarantee of outcome: classification error is real, and the assurance architecture above exists because of it. The layer substantially reduces accidental exposure — the modal harm — and raises the cost of casual evasion. It does not stop determined egress, and it does nothing about contact and conduct risks, which are separate elements of the AU model requiring separate instruments. A motivated adolescent can leave, and we do not claim otherwise. The schema has not been validated across multiple linguistic and legal contexts, and should not be adopted as an international standard until African institutions have had the opportunity to break it.

What we propose to build together

Schema validation	African linguists, child-development specialists and reviewers test whether the thirty-nine fields survive translation and cultural interpretation. Published in English, French and at least one major regional language from the outset.
National threshold design	Two founding jurisdictions set their own thresholds through their own processes, including consultation with older adolescents.
Child participation	Young people review classification concepts, categories and interfaces, and sit in the review of the standard.
Publisher tooling	Classification made cheap enough that local publishers and education providers will actually do it.
Conformance suite	Open test suite so any implementer can verify compliance independently.
Demonstrator	Two-country implementation measuring decision volume, inter-rater reliability by language, error rates against a reference sample, translation integrity, provisioning behaviour on shared and prepaid devices, and household election patterns. Results published whether or not they are favourable.

What IFC-SiS contributes: the existing schema, offered without commercial condition; the governance and conformance model; maintenance of versioning and the change-control process; and in-kind technical work. The hypothesis under test is the usefulness of the admission layer, not the correctness of the Foundation's present vocabulary. We would regard it as a good outcome if the schema served only as a worked example and African institutions built the layer themselves. Our expectation is that a standard of this kind should end up owned by an African standards institution — continental or national — with the Foundation in a technical-lead role rather than a proprietary one, on the division of labour the accessibility standard is already using.[4] We would rather establish that arrangement at the outset than negotiate it after the work has acquired a name.

What the work requires: African institutional co-designers with authority over the scope, revision and validation criteria of the work, not an advisory role over ours; local publishers and educational-content producers, who are best placed to say whether the classification burden is realistic; child and youth participants; participation from the existing UNICEF and GSMA coordination structures; two government partners and their regulators; operator and registry partners; an independent research and evaluation partner; and demonstration capital for the standards-development and demonstrator phases described above.

1. Patricia Ainembabazi, *Protecting Children Online in Africa Must Move from Policy to Practice*, CIPESA, 11 June 2026 — and the Digital Rights Alliance Africa report cited therein, *Child Protection and Safety Online in Africa: The Law, Privacy, Challenges and Solutions*, covering Algeria, Botswana, Egypt, Ghana, Kenya, Nigeria, Rwanda, South Africa, Tanzania and Uganda.

2. African Union, *Child Online Safety and Empowerment Policy*, adopted by the Executive Council, 21 May 2024.

3. UNICEF and GSMA, launch of the Africa Taskforce on Child Online Protection, MWC25 Kigali, 22 October 2025.

4. African Organisation for Standardisation (ARSO) and inABLE, African Standard for Accessible ICT Products and Services, announced at the Zero Project Conference, Vienna, 2026; and CIPESA, *Accelerating Digital Accessibility and Solutions for Africa's Future*, 25 June 2026.

Rights instruments referenced: Convention on the Rights of the Child, Articles 5 and 16; African Charter on the Rights and Welfare of the Child, Articles 9(2), 10 and 20.

> *The purpose of the work is not to establish that this architecture is correct. It is to determine whether an admission layer is useful, governable and portable — and, if it is, what African institutions decide it should become. Critical reactions are more valuable to us than supportive ones.*

SOURCES

S. Harrison Knudson, J.D.
Director · International Foundation for Child-Safe internet Standards
shknudson@ifcsis.org · +1 202 684 6309 · standard.ifcsis.org