The Reservoir Model: Infrastructure for a child-safe internet

@Steele Harrison Knudson · Nov 26, 2026

S. Harrison Knudson, J.D. — Director, International Foundation for Child-Safe internet Standards (IFC-SiS) · ORCID 0009-0004-3600-4115

Abstract

We cannot make the whole internet safe for children, any more than we can desalinate the ocean. We can build a reservoir.

The reservoir is an opt-in delivery environment. Participating services admit material, service capabilities and contact roles against published rules before making them available to an enrolled child. Independent admitting bodies can suspend or revoke that authorization. Families choose among what has been admitted, and children keep confidential routes to help.

The model borrows nothing new. Water, food and medicine each rest on infrastructure, a label that says what is inside, and dosing by age. The internet has none of the three. The reservoir supplies the infrastructure; the label is a machine-readable code on the model of the International Building Code; the dose belongs to the family.

The reservoir is infrastructure, not a regulatory framework, and the paper keeps the two apart. It proposes a first trial that is voluntary: it would test whether independent admission and revocation improve accountability in participating services, and measure the cost, access and privacy trade-offs of doing so. Any legal duty on providers is set out separately, as legislative options for debate.

1. You can't desalinate the ocean

The internet is not a river that can be filtered at some point along its course. It is an ocean — more exactly, a gyre, like the garbage patch in the Pacific, where everything from everywhere accretes, compacts and circulates. Anything can be next to anything at any moment. There is no direction of flow, so there is no place to put the filter.

For most of history, distance did quiet work that nobody had to specify. Physical distance never acted alone: it generated a bundle of interacting frictions — travel, attendance, search, matching, coordination, translation, temporal access, and the social visibility of



being somewhere in person — and their joint decline is what changes institutions (The Geography of Zero). One thing those frictions supplied, without anyone designing it, was a child's separation from adult life. Networked access removed the frictions. It did not remove the demand for what they did.

So we have tried to put the protection back by filtering. Content moderation, takedown orders, detection and age checks all work on the ocean itself. They pull out what they can find, after a child may already have found it. Desalination is possible in a laboratory; desalinating an ocean is not.

The early evidence points the same way. Australia passed the world's most ambitious rule: from 10 December 2025, age-restricted social media platforms must take reasonable steps to keep Australians under 16 off their services. Among children aged 10 to 15, the share holding at least one account on an age-restricted platform fell from 52.4% to 42.1% after three months. Half of the children who still held accounts said the platform had never asked them to confirm their age, and eSafety identified platforms' inaction on age assurance as the main reason (eSafety, Early days, early insights, July 2026). These findings do not show which other approach would work better. They do show that a rule enforced by the platforms it restrains depends on those platforms acting.

Families have been handed the same impossible job. Weak uptake of parental controls is usually read as weak demand for child safety. But a market for safety instruments measures three things at once: willingness to be protected, willingness to perform the implementation labour, and capacity to perform it (The Demand Side). A household that declines a filter has not declared indifference to its child's safety. Nobody measures demand for safe drinking water by counting sales of home filters. Where that market is large — bottled water on an unreliable supply — we read it correctly, as institutional failure rather than consumer enthusiasm.

The question is not how to clean the ocean. It is why we are asking children, and their parents, to drink from it.

2. Where the model comes from: zoning, and new space

The reservoir comes from zoning, and zoning is the oldest and most widely understood form of local authority: every town does it, and almost no one finds it controversial. It decides what may stand where, and it protects the places where children live, learn and play. The reservoir is land-use planning applied to virtual space: a designated child-safe space governed by a published code of permitted and excluded uses (Additive Namespace Architecture).



New York's Times Square shows what zoning can do. By the early 1990s, 42nd Street was the city's red-light district. In 1995 the City Council amended the Zoning Resolution to bar adult businesses from certain commercial districts (Untapped Cities), and the city kept them at least 500 feet from schools, churches and homes (The Nation). The same year, an old skin-flick house reopened as the New Victory, a theatre for children, and in 1997 *The Lion King* opened at the restored New Amsterdam (Vital City). The street went from showing what it had been to showing what a city could choose to be.

But Times Square took decades, billions of dollars and years of litigation, and it worked by clearing an existing place. The internet offers something no city has ever had: new space, created instantly. Digital real estate does not have to be taken from anyone. The slum can stay. The Wild West can stay. Nothing has to be cleaned up, condemned or moved. We can simply build a new place, zoned from the first day for children. You do not give up the open web to join it.

Zoning also teaches where the danger lies. Enrollment asks who may enter; siting asks what a child crosses to get there. A school can check every adult in the building and still sit beside a strip club, because the exposure lives in the walk, not at the address (Siting the School). Today every child online walks through the whole ocean to reach anything. In a reservoir, content that was never admitted is not there to be circumvented. That is a property of construction, not a claim about how well a filter works.

The structure is not new to law. English law already knows it as a scheme of development: entry voluntary, area defined, rules published in advance, and holdings elsewhere untouched. Every owner inside the scheme is bound to every other, so the code holds across the whole area, not just one lot.

The internet itself was zoned from the start. When Jon Postel and Joyce Reynolds set out the first top-level domains in October 1984, each was a district with its own use: .gov for government, .edu for education, .com for commerce, .mil for the military, and .org for other organizations (RFC 920). The domain name system is a land-use map. Twice, someone has tried to add a district for children.

In 2002 the United States Congress passed the Dot Kids Implementation and Efficiency Act, which created kids.us: a space inside the US country-code domain for content suitable for children under 13 (NTIA). The ACLU urged the Senate to oppose it (ACLU). It was closed in 2012 after fewer than 1,000 registrations and about half a dozen active sites (Domain Incite). A second attempt is still running: .kids, a generic top-level domain delegated to the DotKids Foundation in April 2022 (IANA).

Why kids.us failed is not documented here, and my reading of it is a hypothesis for the trial to test. It was a destination, not a distribution system: nothing linked a child's account to it, nothing required providers to respect it, and so publishers had no reason to build there. It was also a single government's project, with no international standard behind it. A district on its own, whether kids.us or .kids, is a place to publish. The reservoir starts where a district stops — with the account, the rules and the coalition.

3. Build a reservoir

No city asks its residents to drink from the ocean, and no city expects each household to dig its own well. Cities build reservoirs. The reservoir is one of the oldest ways of protecting people in dense settlements, and it works in every culture and every century because the idea is simple: protect a supply, control what flows into it, and deliver it to the tap.

A child-safe internet can be built the same way. The reservoir is an opt-in delivery environment governed by one principle: admission before contact. Nothing reaches an enrolled child unless it has been admitted, and three different kinds of thing are admitted:

- **Material:** a specific, versioned item of content — a lesson, a book, a video — labelled against the published standard. A changed version is a new item and must be admitted again.
- **Service capabilities:** what a service is allowed to do — whether it may accept uploads, display third-party embeds, send messages, or generate personal replies.
- **Contact roles:** the people and AI systems allowed to interact with a child, each in a defined role, such as tutor or coach.

Each is verified before delivery or contact, and its eligibility stays subject to review and revocation. Admission is a different question from age assurance. Age assurance asks who is entering; admission asks what is allowed inside (The Admission Layer).

The boundary must be stated plainly. Admission is enforced by participating services within the enrolled session. It governs what those services deliver to an enrolled child; it does not govern everything a child can reach on a shared device, over another connection, or through a service outside the reservoir. Managed devices and internet-provider provisioning can extend that coverage, and the trial should measure how far, but the paper claims protection only inside the enrolled session.



The reservoir fails closed. Whatever is unlabelled, unverified, revoked or expired is not delivered. Revocation has a published maximum time to take effect; active sessions end and stored copies are withdrawn within that time. And the route to confidential help is maintained separately, so it keeps working when an account, an enrollment or an admission check fails.

Three features of a water system carry over directly.

- **The protection is at the intake.** The gatekeeper works where water comes in, testing what enters. It is not a wall around children.
- **The pipes never cross.** Drinking water, storm drains and sanitary sewers run in separate pipes, and a valve on the drinking line prevents backflow. Nothing from the ocean ever runs uphill into the reservoir.
- **It is a source, not a destination.** You don't go to the reservoir to drink. Its water is treated, tested and delivered to the places children already are: the tap at home, the fountain at school, the sink at the library. The reservoir comes to the child, on the device and in the place the family already uses.

A worked example. A nine-year-old opens a tutoring service on the family's shared tablet. The service checks for an enrolled session: her parent has enrolled her, and authorized one role — a mathematics tutor. The service checks that its own tutoring capability and today's lesson are admitted, and delivers the lesson. It cannot message her outside the lesson, accept uploads, or show her anything that has not been admitted. When her older brother picks up the same tablet without signing in, he gets the ordinary internet; the reservoir governs her enrolled session, not the device. If the admitting body revokes the tutor's authorization, her session ends within the published time, the lesson copies are withdrawn, and the help button on her screen still works.

It is not a walled garden. A walled garden is one company's private enclosure; the reservoir is a public code applied across a whole designated space. It does not replace, censor, drain or close the ocean. The open internet stays exactly as it is, and the reservoir is an additional choice, not compulsory confinement.

Nor is it meant to hold a child forever. The reservoir is a place to practise: a bounded environment in which a child exercises genuine autonomy, because a defined class of externally supplied hazard has been removed and ranging freely is no longer reckless (A Place to Practise). The closest analogy is a graduated driving licence — a learner's permit, then provisional driving, then a full licence — each stage widening as capability grows. The reservoir is built to be outgrown.



A reservoir reduces and manages risk; it does not promise perfect safety. Water is tested after it enters the reservoir as well as before, and operators remain accountable after admission. What it removes is the impossible task, now laid on every platform, carrier, regulator and parent, of making the ocean drinkable.

4. What is safe to consume has always been built this way

Every society already knows how to protect what children consume. Water, food and medicine each rest on the same three layers: infrastructure that holds and delivers a protected supply; a label that says what is inside; and dosing by age, decided by the people responsible for the child. The internet is the only thing children consume every day that has none of the three.

What is consumed	Infrastructure	The label	The dose
Water	Reservoir, treatment, separate pipes to the tap	Water quality standards and public reporting	Not needed: the supply is safe for everyone
Food	Inspected supply chains	Ingredients, allergens, nutrition panel	Portions and foods chosen by the parent
Medicine	Licensed manufacture and pharmacies	The drug label and the pharmacopoeia	Pediatric dosing by age and weight, on prescription
The internet, today	None — the ocean	None that is uniform or machine-readable	Guessed by the platform, if at all
The internet, in the reservoir model	The reservoir	A published, machine-readable code	Set by the parent, through enrollment

The comparison is not decoration. It is the model. Each of these regimes was built after people learned, at real cost, that a free-flowing supply could not be made safe after the fact. Each separated a protected supply from the open one. And each kept the decision about a particular child with the people responsible for that child, while the system made sure they had the facts to decide.

Each also followed the same path from a vague duty to a specific requirement. The backyard swimming pool is the plainest example. The common law said a landowner owed a trespasser almost nothing, except where something on the land would attract children who could not see the danger. That doctrine, attractive nuisance, required reasonable care but named no particular precaution. Then a private standards body wrote the precaution down, and governments adopted that standard by reference rather than writing their own (Build the Fence). That is why pools have fences. The internet needs the same fence, in three parts: enrollment, certification against a published standard of conduct, and continuing duties of monitoring, reporting and removal.

We do not have to invent the content of those duties. Teaching, counselling, medicine and coaching have already built tested safeguards for the risks that arise when an adult with authority holds a child's trust: secrecy, misplaced trust, and authority without oversight. Online child-safety law keeps writing those rules from scratch. It should reference them instead: the statute supplies the duty; the standard supplies the content (Reference the Professions). The internet is a new place for very old relationships.

The code that holds the label should work the way building codes work. The International Building Code is one readable text that jurisdictions adopt — in full, in part or with additions — and then enforce through their own inspectors. It is a floor, not a ceiling. A child-safe internet needs the same thing: one machine-readable code, adopted locally, so that a regulator in Brisbane and a regulator in Salt Lake City are reading the same label.

5. The label says what is inside

A food label does not tell you whether to eat the food. It tells you what is in it. The reservoir's label works the same way: it describes; it does not judge. Nothing is removed from the world. Everything admitted to the reservoir is described. The descriptor is the standard; the country and the parent set the limit.

The label is a published standard, C-SiS — the child-safe internet standard — written as a readable code in MUST, SHOULD and MAY (C-SiS). It describes two kinds of things.

- **Images.** "Nudity" is too large a word to be useful, so the standard breaks it into anatomical regions, in the terms doctors already share across countries, and records an observable coverage state for each. Garments are never named: a bikini and a thong differ only in coverage of one region. Describe the anatomy, not the garment. Two people looking at the same image should assign the same values without sharing any beliefs. Stick figures, avatars, emoji and robots with faces count as people; there is no exemption for stylization (Images).



- **Subject matter.** This is harder, because the usual labels are verdicts. So the standard borrows vocabularies that already exist: library subject headings and the book trade's codes for themes such as self-harm and body image, and forensic medicine and criminal codes for physical acts. Violence is recorded by mechanism and injured region; legal labels such as murder or self-defence are conclusions about intent, and they stay with each country. Each subject is marked as mentioned, discussed or depicted, because a recovery story and a how-to page may share a theme and treatment tells them apart. Beliefs are described neutrally and never ranked (Subject Matter).

This is what should make the label work across cultures. Existing ratings do not translate across faiths, cultures or languages, because they are verdicts. A description can. The design aim is that a progressive parent and a conservative parent can both use the same label to make opposite choices, without either feeling the label was written for the other one.

That two assessors will assign the same values is a hypothesis, not a promise, and the trial should test it. Anatomical coverage is relatively observable; context, implied conduct, and the line between depicting something and instructing in it can produce real disagreement. So every descriptor also has two honest states, **unknown** and **disputed**. The standard publishes worked examples with adjudicated labels, and the trial measures agreement between assessors and wrongful exclusions, by language and by type of content. Health education, abuse-prevention material and recovery stories are the first test cases, because a mistake there can block a child from legitimate help.

The label is machine-readable, because the reservoir admits content at the speed of the internet. The structure and the fields are fixed; the limits are local. A nation may adopt the code in full, in part or with additions. One country may set an age gate at 16 and another at 18, while the descriptors stay identical, and a parent may always choose stricter. A conformance core cannot be removed, and conformance is verified by public tests in many languages. The vocabulary is versioned and grows more precise with use; replaced terms are mapped forward, never deleted.

The label describes content. A second, separate set of descriptors records what a service can do: whether it can initiate messages, accept uploads, generate personal replies, or connect a child with an adult. The two work at different levels. Admission sets the most a service may offer; a family's or curator's choices can only narrow it.



6. The dose belongs to the parent

A pharmacist does not decide how much medicine a child takes, and neither does the manufacturer. The label says what is inside; the drug is approved by an independent regulator; and the dose is prescribed for a particular child by someone responsible for that child. The reservoir separates the same three decisions and gives each to a different party.

Decision	What is decided	Who decides
Description	What a piece of content or a service is, in machine-readable terms, with no verdicts	An open, published standard
Admission	Whether an operator, a person or an AI system is fit to be inside the reservoir, and stays fit	An independent admitting body, applying published rules, with review and appeal
Selection	What a particular child may use, and who may reach that child	The family, with curators it chooses

Admission: the adult proves fitness. At a liquor store the customer proves age. At a school the adult proves fitness and the child proves nothing (Fitness to Serve). Every adult who works with children is checked, credentialed, trained, renewed and revocable. The reservoir asks the same of everyone admitted, and of AI systems too: as a system moves toward sustained, personal interaction with an identified child, the question is fitness, not age. Parental permission and independent fitness are separate requirements. A parent may invite someone in; the gatekeeper still checks that they are fit.

Selection: the parent holds the account. Every offline institution that serves a child begins the same way: a parent hands the child over. School, camp, the doctor and the coach all start with enrollment. The internet is the only place a stranger can form a relationship with a child without a parent ever being asked. Enrollment restores the offline default: the parent's act, not the platform's guess, is what identifies the child (Enrollment Before Contact). Self-declared age never qualifies. The parent or guardian holds the account; the child uses it. Contact with a child runs through organizations the child is enrolled in — a school, a club, a congregation — and an adult may reach the child only when both belong to the same one.

The rule is specific to the decision. A child's agreement to participate does not substitute for authorization from a parent or other legally authorized person where the law requires it — and within the reservoir, the standard requires it everywhere (section 8). Within the reservoir, enrollment and independent approval are required before a provider may initiate contact in a defined role. Children retain a confidential route to help without enrollment. Data-protection law already draws the same line within its own scope: the US Children's Online Privacy Protection Act requires verifiable parental consent before an online service collects personal information from a child under 13 (COPPA); Article 8 of the EU's General Data Protection Regulation requires the consent of the holder of parental responsibility where an online service offered directly to a child relies on consent and the child is under 16, or a lower age down to 13 set by a member state (GDPR); and India requires verifiable parental consent before processing the personal data of anyone under 18.

Curators advise; parents decide. Most parents will not set every choice themselves, and they should not have to. Curators — a school, a library, a health group, a faith community, any organization a parent trusts — publish age-based lists saying whether each descriptor is allowed, held for the parent to review, or blocked. Curators need no special status. Every line is visible, a parent can change any line for any child, and a parent who disagrees can switch curators at any time. The standard never decides what a child should see (How Curation Works). Because the controls follow the child's evolving capacities, as Article 5 of the UN Convention on the Rights of the Child describes, the dose grows with the child: the child holds the rights, and the parent is the recognised agent, scaled to the child's maturity. The UN Committee on the Rights of the Child's General Comment No. 25 applies the same principle to the digital environment, including children's privacy ([General Comment No. 25](https://digitallibrary.un.org/nanna/record/3906061/files/CRC_C_GC_25-EN.pdf)).

When the parent does not act. Measures that rely on parents fail the children who get the least support at home (When the Parent Does Not Act). Inside the reservoir, the default is closed: if no one enrolls a child for a role, that role does not open. Parents can open one defined role; they cannot switch protection off. Whether providers outside the reservoir should carry a similar duty is a separate question, set out in the legislative options below.

Help never waits for a parent. A child can reach help without enrollment, without notice to a guardian, and without leaving a record a guardian can see. Every service in the reservoir must follow the reporting duties that apply to it, and where telling a guardian could endanger the child, a disclosure goes to the competent authorities, not home. The



same care applies where guardianship is disputed, where a child is in state care and the legally authorized person is a carer or an agency, and as authority passes to the young person with age. Enrollment governs who may reach a child; it never closes the child's way out.

7. The sandbox is the clinical trial

No one puts a medicine on the shelf for children before it has been tested. The reservoir should meet the same standard. No write-gated child namespace has yet been operated at scale, so its admission rates, error rates, cost, latency and sufficiency are unmeasured (Siting the School). This paper makes no performance claims. It argues for a trial.

Regulatory sandboxes and government-commissioned trials are where that testing belongs: time-limited and voluntary, under a regulator's supervision, with an independent evaluator — ideally a university — and a public report. A trial is a clinical trial for infrastructure, and like a clinical trial it should begin small and answer one question.

The first trial: does independent admission improve accountability beyond what participating services already do? It would run on an existing domain, with a small number of participating services, reviewed educational material, and narrowly defined contact roles. The same service would be compared under its ordinary safeguards and under independent admission, with every legally required protection kept in place in both. It would start with simulated cases and operational checks, and move to a bounded pilot with real families only after an independent review.

The decisions the data will support are registered in advance. The trial measures wrongful admissions and wrongful exclusions, the time it takes to revoke access, whether confidential help stays available, how many families complete enrollment on shared devices, the privacy burden on families, and the cost to operators. Failures are reported by language and by setting. A small pilot is evidence about feasibility and whether the controls perform. It cannot show a reduction in rare harms, and it cannot justify switching on wider AI or social features.

Over later stages, the reservoir has to answer two larger sets of questions.

Does the infrastructure work?

1. **Can the reservoir be built?** Can a bounded namespace be stood up in which content is admitted against the label and anything unverified is kept out? Measures: admission volumes, error rates, cost, latency, registrant behaviour.



2. **Can accounts be linked to it?** Can participating services reliably recognize an enrolled child's session and enforce admission within it? Managed devices and provisioning through the household's internet provider are optional ways of extending coverage, each measured separately, and the trial reports what happened inside the enrolled session apart from what children could reach elsewhere. Measures: leakage, and how often the restriction is bypassed, including on shared and prepaid devices.
3. **Does revocation work?** When an authorization is withdrawn, does access actually end within the published time, across sessions and stored copies?
4. **Will it hold?** Do the limits survive real use, and where are the gaps? A friend's mobile data and public Wi-Fi are known gaps, and the trial should measure them honestly.

Does it work for families and children?

5. **Can enrollment be simple, verifiable and accessible** — including for parents with little time, little money, or a different first language?
6. **Do admitted services stay within their role?** An AI tutor must stay a tutor. Admission cannot govern the output of a generative system, because that content does not exist until a child calls it into being. So AI begins with reviewed material and tightly scoped tutoring functions. Open-ended generation is evaluated in a separate stage, with simulated conversations, independent assessment of whether the system stays in its role, change control, and stop criteria agreed before any wider use with children. A successful infrastructure trial would establish neither that a tutor is safe nor that it is fit for a professional role.
7. **Do independent audits catch failures** and lead to correction, and to removal when a standard is breached?
8. **Does it make children safer without unnecessarily reducing their access, autonomy, privacy or participation?** And does it work in every language the trial runs in? A label written in one language carries that language's assumptions.

The trial must also measure demand correctly. If families have to configure the reservoir device by device, the trial will measure their spare time, not their wish to protect their children. The point is to remove that labour, not to restore it.

Nor should the trial claim more than it tests. What it can produce is something no stakeholder has today: operating data on a child-safe internet that families actually used.



8. The rules the reservoir needs, and the legislative options

The reservoir is infrastructure, not regulation. But infrastructure succeeds only when the rules around it are in place. Water works because the law requires it: plumbing codes require buildings to connect to the public supply, and they forbid cross-connections, so a pipe carrying unsafe water can never feed the drinking line. Those rules bind the builder, not the person drinking. Without them a reservoir is just one more source of water.

The rules also have to put the burden in the right place. Today it sits with parents, who are asked to filter the ocean device by device and app by app. The right question is who can most cheaply reduce the risk once the activity is happening, and that is rarely the parent. The commercial market in filtering subscriptions is a price quotation on a burden that exceeds ordinary capacity (The Misallocated Burden). Nor was that allocation ever designed: it accumulated from a cluster of decisions made in the late 1990s, when the internet was new.

Two kinds of rule are involved, and the paper keeps them apart. The first kind governs the reservoir itself and is needed for the trial. The second kind would reach providers outside the reservoir; it is a matter for legislatures, and it is offered here as options for debate.

Rules inside the reservoir. These bind every participating service through its admission agreement.

- **No cross-connection.** Nothing from the ocean may be piped into the reservoir. Whatever is delivered to an enrolled child has been admitted against the label and the service's admitted capabilities, and what is revoked or expired is not delivered.
- **Enrollment and approval before contact.** This is a requirement of the standard, and it is the same in every jurisdiction where the reservoir operates: a provider must verify enrollment by a parent or other legally authorized person and obtain independent approval before initiating contact with a child in a covered role. A child's click, signature, or agreement cannot substitute for either requirement. Providers may neither seek nor rely on a liability waiver from a child. Children may seek confidential help without enrollment and exercise any rights to consent that applicable law gives them.
- **Inspection.** Independent testing and certification, renewed and revocable, as water is tested, food is inspected and drugs are approved — with a recognizable mark that parents can trust without reading the code.
- **A limited presumption, not a blanket shield.** A service that demonstrates a specific enrollment or admission control earns a rebuttable presumption of due care for that control only (No Free Pass). The presumption falls away on evidence that the service

ignored reports, acted outside its admitted role, or failed to act after revocation.

The standard sets one rule because the underlying laws differ, and none of them makes a child's click a sound basis for a relationship. India ties general capacity to contract to the age of majority (Indian Contract Act, 1872). The law of England and Wales recognizes that minority can make a contract unenforceable against the minor or open to repudiation, a common-law rule adjusted by the Minors' Contracts Act 1987. German law generally requires a legal representative's consent when a minor's declaration does more than confer a legal benefit, with a limited exception for transactions completed with money given for that purpose or for the minor's free use (German Civil Code, sections 107 and 110). GDPR Article 8 separately regulates consent to certain data processing and expressly leaves member states' contract law untouched. The reservoir does not depend on which of these applies; it simply never treats a child's agreement as the authorization.

Legislative options. The voluntary trial does not depend on any new law, and it should not be read as proposing one. But the trial's results will inform a question legislatures are already asking: whether providers outside any reservoir should have to obtain authorization before forming a relationship with a child. The paper offers one way to frame that question, not an answer.

- **A conduct trigger.** A future duty could be triggered by what a provider does, not by guessing a user's age — for example, initiating repeated, individualized contact with a known child in a defined role, such as tutor, coach or companion. The scope of any such duty — who is covered, what triggers liability, and whether providers would have to ascertain age across the open internet — is open for debate.
- **Preserved access.** Any duty must preserve emergency, safeguarding and other lawful confidential access for children.
- **Why the duty is reasonable.** Judge Learned Hand's formula states the test: when the burden of a precaution is less than the probability of harm multiplied by its gravity, failing to take it is negligence. Whether enrollment is in fact a low-cost precaution is a question the trial should measure.
- **Related proposals.** Allowing under-16 accounts on platforms covered by Australia's minimum-age law, where those accounts are linked to the reservoir, would be a separate proposal needing its own statutory basis.

None of these rules imposes a burden on adults or removes anything from the internet. They move the cost of protection toward the parties best placed to carry it, and they give regulators what they lack today: a place where their rules can be enforced.



9. Where it is tested first: Utah, Australia, India

The reservoir should be tested the way a medicine is: at several sites at once, each under its own law, all reporting against one protocol. Three candidate sites stand out. Each could test something the others cannot, and each still needs a local partner and an identified legal pathway.

Utah. Utah created the first state Office of Artificial Intelligence Policy in 2024, with a learning laboratory that studies AI policy and a regulatory mitigation program that lets participants test under supervision (Utah Department of Commerce; Utah Code § 13-72-401). It is the natural place to test enrollment for child-facing AI in schools, and to draft the model legislation that follows from the results.

Australia. Australia already has the mandate: platforms must take reasonable steps to keep under-16s off (Department of Infrastructure). The early evidence shows that the mandate needs something to enforce against. Australia has also shown how such a trial is run: its independent Age Assurance Technology Trial assessed more than 60 technologies from 48 vendors (Minister for Communications). A first Australian trial should be scoped to services and activities permitted under current law, and it can test how accounts are linked and whether limits hold. Allowing under-16 accounts on covered platforms if they are linked to the reservoir would be a separate legal proposal, needing its own statutory basis; parental enrollment alone does not satisfy the current minimum-age obligation.

India. India has put in place, or committed to, four of the five legal pieces a working reservoir in India would rest on (India Has Written Four of the Five Laws). It requires verifiable parental consent before any personal data of a child under 18 is processed — the highest age threshold of any major data protection law (PRIVO) — with those obligations taking effect on 13 May 2027 (AZB & Partners). Its telecom rules already put a verified adult behind every connection, and commitments on connections and accounts for minors are moving through its regulators and its Supreme Court. A verified telecom subscriber, a verified parent and an enrolled child session are different things, and each duty needs its own legal instrument; where none exists yet, the reservoir is a design hypothesis. Taken together, though, India's rules point toward one principle. The fifth piece is the one that does not exist: the child's connection, drawing only from the reservoir. The parent owns the line; the child uses it. What no one has yet built is the infrastructure that lets a child use a connection the parent owns.

Australia keeps children under 16 off social media altogether. India lets children in only with a parent's consent. A reservoir that works under both laws answers both models.



A multi-site trial also needs a coordinating centre: a university that holds the protocol, compares the results across sites and publishes them. Whichever institution hosts that work will lead on child-safe internet infrastructure.

10. A public good

The reservoir is public infrastructure, and the code that labels what enters it is an open standard. No company owns it, and no government owns it. Each party has one role:

- **IFC-SiS** develops and maintains the open standard through a published, collaborative process.
- **Independent admitting bodies** apply the standard and provide review and appeal.
- **Operators** run participating services and enforce admission decisions.
- **Public authorities** exercise their own lawful oversight.
- **A separately governed coalition** may pursue a community domain, if the trial shows that doing so would improve deployment.

Any nation may adopt the standard as its floor and build on it in strictness.

The design carries specific safeguards against censorship (Additive Namespace Architecture):

1. Participation is voluntary and parental, not mandatory or governmental.
2. The household, not the state, is the point of consent.
3. No third party has to suppress or modify expression in the open internet.
4. The open internet keeps working alongside the reservoir.
5. Policy authority sits with an independent body that uses a published process and offers appeal.

It is also built to resist capture. Inside the reservoir, every admitted operator is bound by the same published code and to every other operator, the way every owner in a scheme of development is bound to every other. A company cannot turn a public standard into a walled garden. A government that wants to misuse it can only copy it inside its own borders; it cannot change the standard everyone else reads.

It must be designed by the people it serves, not only by the countries that build the most technology. Africa combines the conditions in which a standard is hardest to get right and most valuable if it holds: many languages, prepaid and mobile-first connections, and



devices shared across a household. So African institutions should hold real authority over the standard's scope, revision and validation, not an advisory seat (The Admission Layer). The same is true wherever the standard is adopted.

And it must pay for itself, the way a municipal water system does. Testing, monitoring and inspection cost money, and the reservoir should fund them from the services it carries — operator revenues from child-safe service tiers — rather than from charity or from the families it protects. That raises a fair question: will the largest operators end up financing the people who decide whether to remove them? So certification fees are set by a published schedule, not negotiated; auditors are appointed and paid by the admitting body, never by the operator being audited; and appeals go to a body independent of both.

Over time, the reservoir could be stewarded by an international coalition of professional organizations with expertise in protecting children online, consulting governments but never led by them. That coalition would be governed separately from IFC-SiS. If the trial shows that a community top-level domain would improve deployment, the coalition can apply for one. Testing comes first.

A child's right to protection cannot depend on the purchasing power of the child's school — or the size of the child's country. The standard of protection should be universal, even when the capacity to implement it differs. The reservoir is offered on those terms: as a gift, to be tested, adopted and improved by anyone who will build a child-safe internet.

11. An open discussion

This paper is an invitation, not a finished proposal. The reservoir has already been improved by people who know children's lives better than any one discipline can.

Social workers and lawyers have shaped how the model protects a child's access to reporting abuse, so that enrollment never closes a child's way out. An advocacy organization that keeps an annual naughty list of mainstream companies that facilitate sexual exploitation helped me understand the regulatory landscape. A digital-rights policy expert in Africa taught me how teachers and the institutions above them carry authority for children across Africa, and that a principle can hold everywhere while the mechanism is shaped to local institutions. Colleagues in India showed me how language barriers decide whether a standard works at all: a label written in one language carries that language's assumptions, so the code has to be tested in many.

I welcome that discussion from everyone: parents, teachers, social workers, clinicians, librarians, lawyers, engineers, regulators, and the platforms themselves. Where the model is wrong, I want to know. Where it can be tested, I want to help test it. I look forward to more trials, in more places, and to every improvement they bring to the reservoir model. Correspondence is welcome at shknudson@ifcsis.org.

Appendix A: The story behind the theory

My grounding in zoning comes from my training as an urban planner. Years before the reservoir, I prepared for a post in Saipan, in the Commonwealth of the Northern Mariana Islands, by studying the law of sexually oriented businesses across the United States. I never took the job. The Attorney General told me he could not in good conscience ask me to move my family there, because my son needed hospital care the islands could not provide. The research stayed with me.

The reservoir began on 20 April 2019, in Tysons Corner, Virginia. I was looking out at the Tysons II mall while the Metro was under construction, thinking about how communities zone sexually oriented businesses. Two ideas came clear at once: reverse the zone, and build new digital space. Instead of trying to zone the harm out of every place children already go, zone a new place for children in.

The water image came later, as a way of explaining the idea, and it sharpened over time.

1. **A river of filth.** The first picture was a polluted river, with everyone trying to get upstream of the source. It explains the danger, but it suggests a direction of flow, and so a place to filter. The name came from that picture. Explaining the river of filth in her kitchen to the founder of a nonprofit whose start-up challenges turn inventions into new companies for children's health, who was coaching me at the time, I said what the answer was: it's a reservoir.
2. **A pile, not a river.** An academic colleague in Australia saw that this was wrong. The internet has no direction. It is a pile, not a river: there is no upstream or downstream, only before admission and after exposure.
3. **The ocean and the garbage patch.** The pile is really an ocean — more exactly, the Pacific gyre, where everything from everywhere accretes, compacts and circulates, with no current that leads anywhere.
4. **You can't desalinate the ocean.** Harmful content is like salt: it is dissolved through the whole of the water. Removing it means desalinating the entire ocean, which no one can do.

5. **The reservoir.** You don't need to. The water cycle already lifts fresh water out of the sea and delivers it to the mountains. A reservoir catches it, protects it and sends it to the tap. The reservoir is not a war on the ocean; it is a natural part of the water cycle.

Appendix B: Acknowledgments

Many people have shaped the reservoir model, and the paper describes them by their role. As each gives permission to be named, they will be identified here. This appendix will be updated as the discussion continues.

Sources

The author's earlier work (IFC-SiS library)

- The Geography of Zero (WP-05)
- The Misallocated Burden (WP-03)
- The Demand Side (WP-04)
- Build the Fence (WP-07)
- Siting the School (NB-01)
- Fitness to Serve (NB-02)
- A Place to Practise (NB-03)
- No Free Pass (NB-04)
- Enrollment Before Contact (NB-05)
- When the Parent Does Not Act (NB-07)
- Reference the Professions (NB-08)
- Additive Namespace Architecture (EB-01)
- The Admission Layer: A Proposal for African Co-Design (EB-02)
- India Has Written Four of the Five Laws (EB-03)
- C-SiS: The Child-Safe Internet Standard, with Images, Subject Matter and How Curation Works

Other sources

- J. Postel and J. Reynolds, RFC 920: Domain Requirements (October 1984)
- NTIA, Notice of public meeting on kids.us (2003)

- ACLU, Letter to the Senate urging opposition to the Dot Kids Implementation and Efficiency Act (6 November 2002)
- Domain Incite, kids.us coverage
- Untapped Cities, Rudy Giuliani battles the Times Square porn shops
- The Nation, Mickey & the Peep Show
- Vital City, Times Square: good Times Square, bad
- eSafety Commissioner, Early insights from eSafety's comprehensive evaluation project (July 2026)
- Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, Social media minimum age
- Minister for Communications, Age Assurance Technology Trial findings
- Utah Department of Commerce, Office of Artificial Intelligence Policy — Policy Work
- Utah Code § 13-72-401, regulatory mitigation agreements (effective 6 May 2026)
- PRIVO, India's Digital Personal Data Protection Act: what global businesses need to know
- AZB & Partners, India's digital personal data protection framework comes into effect (19 November 2025)
- United Nations, Convention on the Rights of the Child, Article 5
- *United States v. Carroll Towing Co.*, 159 F.2d 169 (2d Cir. 1947) (Hand, J.)
- eSafety Commissioner, Early days, early insights: Understanding experiences of social media age restrictions at the three-month follow-up (July 2026)
- IANA, Delegation record for .kids
- US Federal Trade Commission, Children's Online Privacy Protection Rule (COPPA)
- Regulation (EU) 2016/679, General Data Protection Regulation, Article 8
- UN Committee on the Rights of the Child, General Comment No. 25 (2021) on children's rights in relation to the digital environment
- Indian Contract Act, 1872
- Minors' Contracts Act 1987 (England and Wales)
- German Civil Code (Bürgerliches Gesetzbuch), sections 107 and 110

Licence

© 2026 S. Harrison Knudson. Published under the Creative Commons Attribution 4.0 International licence (CC BY 4.0). Anyone may share, adapt and build on this work, including commercially, provided they credit the original.

