Before you read this

What follows describes a programme in development. The evidence on plastic pollution it cites is documented, cited, and independently verifiable. The architecture it proposes was built drawing on established work in citizen science, environmental litigation, deliberative democracy, and institutional scenario design. It has not yet been tested at the institutional scale it requires. The invitation is not to accept it as it is. It is to engage with it, challenge it where it falls short, and contribute to making it better.

Scope of this Brief. It describes the Spotters network and the app. These sit outside the CHF 1,705,000 founding round (Years 1 and 2 of the five-year CHF 3,850,000 foundation budget), which builds the institution only. The app is a later, separately funded stage. Its figures here are indicative, in euros, and will be re-based before that stage opens.

Power is the problem

Plastic entered commercial use in the 1950s on a promise of necessity. The promise was delivered. Our lives are measurably better.

The question of what to do with the material when it was no longer in use was never asked, and was never resolved. It became rubbish, thrown away.

There is no such place as away. There is only everywhere. Two billion tonnes in seventy-five years, and everywhere now includes our blood.

This was known, measured, and written about as a threat to life on earth for over fifty years. The facts are done. The research is done. The understanding is done.

What has not been done is to stop it.

For half a century, every pathway has failed. Regulation has been outpaced. Voluntary pledges have been broken. Advocacy has been absorbed. Markets have not responded. Litigation has resolved single cases while the system continued.

The reason is not ignorance, and it is not lack of will. The reason is power. The wealth created by plastic does not want the wealth to stop, and has the resources to ensure it does not.

That is the problem this programme is designed to address. The Threshold is what follows.

A verified, legally usable, institutionally endorsed evidence base, built by the people who carry the harm in their bodies. Every observation scored for reliability. Chain of custody from phone to proceeding. A cross-stakeholder deliberation that produces ranked interventions on permanent record. An evidence infrastructure built upstream of the actors who will eventually use it, ready before they need it.

Not protest. Not policy. Not litigation. The architecture that makes all three forceful enough, sooner.

The Threshold is still a proposal. Its purpose is to make the cost of continued suppression exceed the cost of accountability.

The Threshold is built to be ready. Whether and when the threshold is crossed depends on what is built, and on who joins in building it.

App Investor and Builder Brief

The world's first trust-weighted citizen intelligence platform

Building the largest structured evidence constituency for environmental policy ever assembled, layer by layer.

DOCUMENT	**App Investor and Builder Brief · Version 7**
PREPARED BY	**Andrew Henderson · Lead Architect · Marans, France**
DOMAIN	thresholdinitiative.org/app
DATE	October 2026
STATUS	Confidential · not for public distribution

The opportunity in one paragraph

"The Threshold" is a free, politically neutral citizen intelligence app for plastic pollution. It is not a campaign tool. It is an institution with an independent Oversight Board, a published mission, and legal protections for every contributor. Its core function is building the largest, most credible, most institutionally diverse evidence constituency ever assembled around a specific policy crisis, and delivering that constituency's weight into institutions that have already agreed to receive it.

What it is	What it builds	What it produces
A free app. Governance-protected. Non-political. No data sold. Available globally from a browser or app store.	Six layers of named, engaged, institutionally-weighted constituency: from two million citizen Spotters to named pension funds and city mayors, all contributing to the same evidence base.	A trust-weighted evidence record, published policy instruments, a documentary series, and a six-level institutional engagement cascade that generates real policy leverage.

Parameter	Specification
App name	"The Threshold" · thresholdinitiative.org/app
Platform	iOS · Android · Progressive web app (browser, no download required)
Price to users	Free always · optional Founding Circle paid tiers
Languages at launch	Six languages
Governance	Independent Oversight Board · published mission · succession protocol
Core features	Spotter Network · Threshold Points · The Pulse · Ledger Live · Signal/Noise · The Poll · The Archive · Founding Circle · Schools and Youth Layer
Phase 1 build cost	€103,800 (progressive web app, months 1 to 6)
Full 26-month investment	€1,193,280 (all phases, including 10% contingency, before revenue offset)
Net investment required	€1,158,280 (after conservative Founding Circle revenue)
Active user target (month 12)	5,000 verified Spotters in Phase 1 priority geographies (Rhine corridor, Mediterranean, French waterways)
Active user target (month 24)	1,200,000
Founding Circle revenue (month 24 annualised)	€280,000/year at target

A	B	C	D	E	F
Verified fact	Evidence-based	Consensus direction	Informed inference	Contested	Unsubstantiated

1 AN INSTITUTION, NOT JUST AN APP

Why governance matters to investors, and to every user

Most apps can be discontinued, acquired, or pivoted. The Threshold cannot, not without a governance process. This is not a constraint on its commercial potential. It is the source of its analytical credibility.

The platform's outputs will be used in policy processes and potentially in legal proceedings. A Spotter who contributes evidence over 18 months, a school that builds curriculum around the programme's Evidence Briefs, a pension fund that issues a formal disclosure request based on The Poll, all of them are depending on the platform being what it says it is. Institutional governance is what makes that dependability real.

What the institutional structure provides

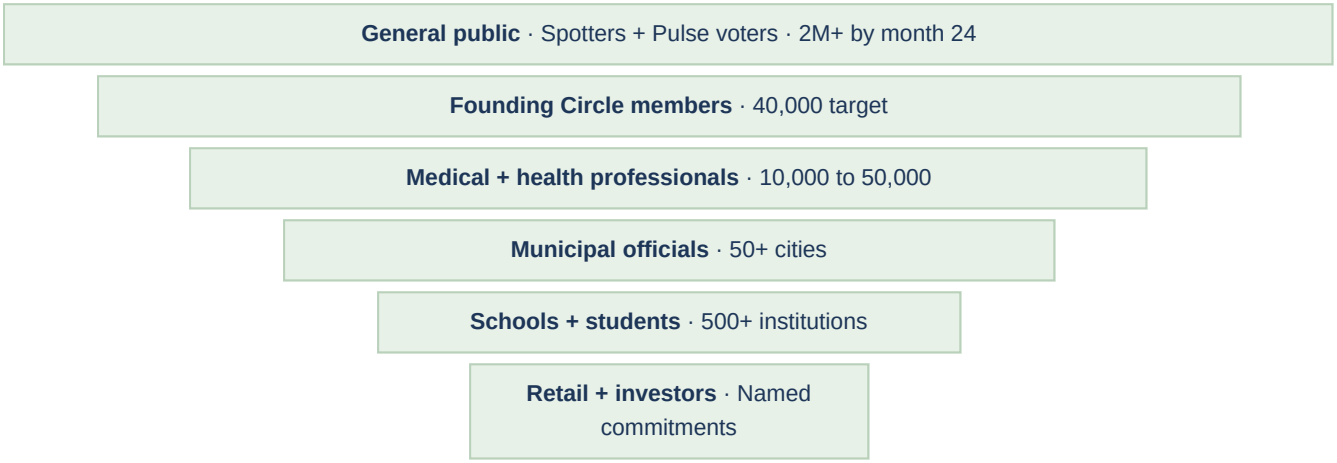
Governance element	What it means for credibility
Independent Oversight Board (seven members, seated by a panel of three outsiders)	The Lead Architect is accountable to an independent body. No funder sits on it, picks its members or holds a key. The Founder holds one permanent seat, with one vote of seven, and cannot chair. The Board oversees the operations of the Knowledge Base and the Althing Round, never their content, and for the Story it oversees operational budgets only. It rules on whether the Canons are kept, and arbitration enforces its rulings.
Published mission (binding on future operators)	The programme cannot be turned into something else. Neutrality, non-commerciality, and no data sales are permanently binding, not dependent on the current leadership's good intentions.
Succession protocol	If the Lead Architect exits, the role passes by Board election and the programme continues under the Canons. The institutional knowledge and methodology are documented and transferable.
Financial transparency	Annual financial statements published. All funder relationships disclosed. No funder receives influence over analytical findings at any investment level.
Contributor legal protections	Users retain copyright over their submissions. The platform holds a usage licence only. Payment is for usage rights, not employment. All protections survive changes in platform ownership or leadership.
Formal correction protocol	When false data is detected it is corrected publicly and completely within 72 hours: never quietly deleted. A permanent audit log is published. This is the most important single credibility mechanism for a platform whose outputs enter policy and legal processes.

2 BUILDING THE CONSTITUENCY

Six layers of engaged, named, institutionally-weighted public support

Every feature of the app is a constituency-building mechanism. The Spotter Network builds an evidence constituency. The Pulse builds a public opinion constituency. The Poll builds an institutional constituency. The Schools layer builds a future constituency. Together they construct something that has never existed for any environmental policy crisis: a layered, named, institutionally-weighted body of public support that is almost impossible to dismiss as lobby noise.

Constituency pyramid: six layers of engaged, named, institutionally-weighted public support



Why this matters to policymakers and funders

A petition can be ignored because it has no institutional weight. A report can be dismissed because it comes from a single organisation. A six-layer constituency, where the same evidence base has been independently validated by 50,000 citizen Spotters, endorsed by 10,000 named medical professionals, cited by 50 city cost audits, studied in 500 school classrooms, and formally presented to institutional investors, cannot be credibly dismissed. Each layer adds a different kind of legitimacy. Together they are what systemic policy change actually requires.

Layer 1 · Citizen Spotters and Pulse voters · 2 million+ by month 24

How engaged: Free app download. Spotter submissions. Weekly Pulse responses. Ledger Live engagement.
Leverage it creates: A live, geographically distributed, trust-weighted evidence base that no research institution can replicate. At target, 200,000 Spotter submissions per month from 40+ countries.

Layer 2 · Founding Circle members · 40,000 target

How engaged: Paid voluntary subscription (€3 to €50/month). Named in publications.
Leverage it creates: Financial sustainability independent of single funders. A publicly named community of financially committed supporters generating €280,000/year at target.

Layer 3 · Medical and health professionals · 10,000 to 50,000

How engaged: The Poll Level C: Global Medical Declaration on Plastic and Human Health. Named, verified against professional registers.
Leverage it creates: Named medical endorsement carries weight that no campaign signature can match. A Declaration co-signed by 10,000 doctors in 60 specialties across 80 countries is a document that health regulators and parliaments must respond to.

Layer 4 · Municipal officials and city governments · 50+ cities

How engaged: The Poll Level D: city plastic cost audits using the programme's published methodology.
Leverage it creates: Municipal cost audit data from 50+ cities creates a global database of plastic externalities that is legally usable in producer liability contexts and is owned by the cities themselves, not the programme.

Layer 5 · Schools and student communities · 500+ institutions

How engaged: Guided Knowledge Layer, downloadable Evidence Briefs, optional supervised class submissions.
Leverage it creates: Educational use is the highest-trust credibility signal available to the platform. If qualified teachers find the programme's evidence genuinely useful for critical thinking education, that is a credibility endorsement that no marketing budget can buy.

Layer 6 · Retailers, investors, and insurers · Named institutional commitments

How engaged: The Poll Levels E and F: retail supply chain commitments and investor plastic liability disclosure requests.
Leverage it creates: Level F is where the programme reaches capital markets. A named pension fund formally requesting producer liability disclosure from a plastic manufacturer moves markets in ways that no petition has ever moved them.

3 EIGHT FEATURES, EACH BUILDING A DIFFERENT LAYER

What the app does, and which constituency each feature builds

Spotter Network

50,000+ active Spotters at month 12 target

Citizens become amateur reporters: photos, video, first-hand stories, news events. Three-stage validation pipeline. Trust-weighted by Reliability Index and convergence scoring. Micro-payments (€5 to 25) for high-value evidence accepted into the Knowledge Base.

The primary evidence constituency layer.

The Pulse

250,000 weekly responses at month 12 target

One specific policy question every Monday. Drawn from the live analytical debate. Results published Friday with geographic heat map. Personalised to user topic interests. Generates public constituency data alongside expert analysis.

The primary daily engagement mechanic.

Ledger Live

4 real-time inaction counters on every home screen

Tonnes leaked, days without binding controls, health studies published, municipal clean-up cost. Calculated from the annual Inaction Ledger to the second. Share button on each counter. Meaning feedback for every user, every day.

Evidence literacy constituency.

Signal/Noise

A to F evidence grade before every sentence

Three daily items: a news story, a scientific publication, a policy development. One sentence each. Preceded by an A to F evidence grade. Builds evidential literacy as a daily habit. Differentiates the platform from every existing news or campaign feed.

Evidence literacy constituency.

The Archive

100% of programme outputs in mobile-first format

Every published output: policy briefs, white papers, Inaction Ledger, predictive commercial analyses, formatted for mobile with share function and a "write my representative" button that triggers The Poll Level B action.

Institutional engagement gateway.

Founding Circle

€280k annual sustainable revenue at month 24 target

Three tiers: Observer (€3/month), Supporter (€10/month), Patron (€50/month). Named in publications at Supporter and Patron levels. Annual briefing call for Patrons. No influence over findings at any tier.

Financial sustainability layer.

Schools and Youth Layer

500+ school institutions at target

Filtered, simplified, classroom-safe evidence. Downloadable Evidence Briefs (not advocacy slides; neutral observation summaries). Optional supervised student class submissions. No individual student tracking. Teacher-led interpretation throughout.

Long-term future constituency.

The Poll

6 levels of institutional engagement

From a two-minute signature to a formal disclosure request. Each level targets a different institution: citizens, doctors, city officials, retailers, investor representatives. At every stage this app shows what the contribution produced, not a like count.

The full institutional cascade.

4 THE SCHOOLS AND YOUTH LAYER

Educational access as institutional credibility: not a side feature

If qualified teachers find this programme's evidence genuinely useful for educating students in critical thinking, that is the highest-trust credibility signal available. It is also the programme's most powerful long-term constituency: a generation of young people who understand the evidence before they become voters, officials, and decision-makers.

4.1 Access architecture by age

Age group	Individual account	Spotter	Pulse	School pathway
Under 13	No individual accounts	No	No	Class group submission only, via teacher account
13 to 14	Yes, with verified parental consent	Yes, text only	Yes	School pathway available in addition
15 to 17	Yes, standard	Yes, full	Yes	School pathway available in addition
18+	Yes, full access	Yes, full	Yes	N/A: adult pathway

4.2 The Guided Knowledge Layer

A filtered version of the Knowledge Base, specifically structured for classroom use. Not a simplified version of the adult app; a distinct educational interface. Content is structured by topic (health, environment, economics, policy, law) rather than by submission date. Simplified in language without losing evidential rigour; the A to F evidence grades remain, teaching students to read evidential quality as a skill. Contextualised with what was observed, where, over what time period, what patterns emerged, and what is still uncertain.

4.3 Downloadable Evidence Briefs

What they are called	What they are not called	Why the distinction matters
Evidence Brief · Observation Summary · Field Pattern Report	Lesson plan · Campaign kit · Advocacy resource	The platform's credibility rests entirely on neutrality. Material framed as advocacy loses that neutrality the moment a teacher or journalist notices. Teachers are sophisticated critical consumers of educational resources and will trust the neutral framing.

Format: 5 to 10 slides, clean visuals, maps and cluster data, minimal text. Updated as new evidence enters the Knowledge Base.

4.4 Class group submission pathway

A teacher account aggregates observations from a class project and submits them as a group record. This carries institutional metadata (school name, teacher account) that the moderation system uses to apply expedited light review. Students are not individually tracked. No student data enters any individual Reliability Index.

4.5 The strategic value of the education layer

Benefit	Why it matters to investors
Credibility multiplier	Educational use by qualified teachers is peer validation that the evidence meets critical thinking standards. This is more valuable than any endorsement the programme can purchase.
Long-term constituency	Students who engage with the programme's evidence at age 15 are voters, officials, and decision-makers by the time the programme's policy recommendations mature into legislative proposals.
Institutional relationships	School board and university partnerships create institutional relationships that outlast any single programme season. These become the anchors for Season 2 and the methodology licensing model.
App store compliance	A clearly structured, age-gated educational pathway significantly reduces the risk of app store removal or restriction. Apple and Google both respond positively to platforms with documented child safety policies.

5 THREE-PHASE BUILD PLAN

What launches when, what each phase adds, and what comes later

Phase 1: Launch · Months 1 to 6 · Progressive web app

Core features on any browser: no download required

Feature	Phase 1 capability	Added in Phase 2
Ledger Live	Full: four real-time counters, share buttons	Nothing excluded
Signal/Noise	Full: three daily A to F graded items, archive	Fifth premium item for Founding Circle
The Pulse	Weekly question, results published Friday	Personalised questions by topic, heat maps
The Archive	All published outputs, write-to-representative	Nothing excluded
Founding Circle	All three subscription tiers, Stripe, share badges	Nothing excluded
The Poll	Levels A and B fully functional	Levels C to F (Phase 2, require verification systems)
Spotter Network	Text and link submissions, community vote	Camera, video, GPS, Reliability Index, micro-payments (Phase 2)
Schools layer	Evidence Briefs downloadable, Guided Knowledge Layer	Class group submission pathway (Phase 2)

Phase 2: Augmented · Months 6 to 18 · Native iOS and Android

Full Spotter Network · verification stack · community features

Phase 2 releases the native iOS and Android app. Camera and video capture, GPS location, push notifications, and the micro-payment pipeline all require native app infrastructure. The Phase 1 progressive web app continues operating alongside the native app.

What Phase 2 adds	Constituency it builds
Camera and video Spotter with GPS location tagging	Dramatically increases evidence quality and geographic specificity. Makes coordinated false submission significantly harder.
Reliability Index and convergence scoring (fully live)	Creates a self-improving evidence quality mechanism. High-RI Spotter submissions carry progressively more weight. The evidence base becomes more credible as the network matures.
Micro-payment system (usage rights fees €5 to 25)	Creates genuine economic relationship between platform and contributors. Incentivises sustained high-quality contribution. Distinguishes the platform from every existing environmental app.
Personalised Pulse with topic matching	Increases Pulse engagement rates by matching questions to the topics users self-identified as their concerns. Deepens constituency engagement quality.
The Poll Levels C to F (professional and institutional)	Opens the medical, municipal, retail, and investor constituency layers. This is where the platform moves from public pressure to institutional leverage.
Class group submission pathway	Opens the formal school constituency. Creates institutional relationships with education systems that amplify programme credibility in ways no other feature can.
Push notifications and community maps	Sustains daily engagement. Makes the distributed network visible. Creates the social context that turns individual contributors into a community without social media mechanics.

Phase 3: Future developments · Months 18+

Built into the architecture; pursued when resources and legal framework allow

Future development: Litigation Signal, opt-in witness pipeline

Users explicitly opt in to having their high-quality, high-RI submissions considered as potential evidence in producer liability litigation. Only Phase 3 submissions (fully validated, convergence-scored) are eligible. Two-stage consent. Legally isolated database. When: after legal architecture is designed and funded separately. Requires specialist legal counsel beyond current scope. Estimated additional cost: €80,000 to €120,000 in legal architecture.

Future development: Exposure Diary, longitudinal citizen health record

A private encrypted personal journal for health observations and plastic exposure events. With academic-standard consent, could constitute the largest longitudinal citizen science study on microplastic health exposure ever conducted. When: after academic ethics board partnership is established. EU and UK only initially. Additional cost: €40,000 to €60,000 setup plus €15,000/year.

Future development: Institutional API, data access for research partners

Structured data feeds for named institutional partners: universities, government bodies, research institutions. When: after Phase 2 data quality standards are validated in practice. Requires a legal data-sharing framework and quality specification.

6 HOW THE EVIDENCE IS VERIFIED

Three-layer AI and human system: truth through structure, not authority

The Spotter Network generates value only if its evidence is credible. Credibility requires verification. Verification at scale requires a system where AI handles speed and humans handle judgement, and where truth emerges from convergence, not from any single authority's decision.

Ingest	**AI Triage**	**Clustering**	**Light review**	**Deep review**	**Integration**
All submissions received	AI classifies and scores	Groups by location + time	Human spot-check 20 to 30%	Senior review high-value 5 to 10%	Published evidence

6.1 The Reliability Index

Every Spotter has a Reliability Index, a score from 0 to 100 that determines how much evidential weight their submissions carry. The formula is published and transparent to every user.

Formula: accuracy 40% + consistency 35% + corroboration 25%

Builds slowly (+1 to +3 per validated contribution). Decays faster if submissions fail review (-5 to -15).

RI level	What it means	Typical profile
0 to 25 New Spotter	Submissions receive minimal weight until corroborated	Recently joined. First few submissions. Building a track record.
25 to 50 Reporter	Validated submissions carry moderate weight	Consistent contributor. Several submissions have passed moderation.
50 to 75 Trusted Spotter	High-weight evidence. Eligible for full payment tiers.	Established track record. Submissions regularly corroborated.
75 to 100 Expert contributor	Maximum evidential weight. Level 8+ in Threshold Rank. Named in programme publications.	Long-term, high-quality contributor. Submission history independently verified multiple times.

6.2 Convergence scoring

The most powerful evidential mechanism in the system. When multiple Spotters independently report the same event from the same location and time period, the system clusters them and raises the confidence score of the entire cluster. A cluster of five independent Spotters with an average RI of 65 reporting the same event produces evidence that is far stronger than any individual submission.

A geographic scarcity weighting ensures that Spotters in data-sparse regions, where fewer users means fewer potential corroborators, are not structurally disadvantaged. A lone Spotter documenting industrial pollution in a region with few other Spotters receives a provisional boost that reflects the genuine value of evidence from under-represented geographies.

6.3 What AI handles and what humans decide

AI handles (automatically, at scale)	Humans decide (the critical minority)
Classifying submissions: relevant, irrelevant, duplicate	All Knowledge Base inclusions: no automated entry
Assigning initial confidence scores	All payment decisions above base tier
Grouping by location, time, and topic	All high-impact or anomaly-flagged submissions
Flagging coordinated attack patterns and velocity anomalies	All formal appeals from Spotters disputing RI or payment decisions
Generating preliminary Fact/Opinion grades for new documents	Triggering the correction protocol when false data is detected
Pre-screening for child protection violations in submitted images	Final Reliability Index adjustments for significant changes

The Intelligence Director, a dedicated senior technical role, designs and monitors all AI systems, publishes the RI formula, and is the human accountable for AI performance. AI outputs are reviewed weekly for drift, bias, and accuracy. If the system is systematically misclassifying a category of submission, the Intelligence Director adjusts the rules and documents the change.

7 INVESTMENT REQUIRED

What each phase costs, what it builds, and how costs evolve with scale

The budget reflects a professional verification operation built to the standard that a platform producing evidence for policy and legal use actually requires. The moderation and verification costs are the largest single line item, and the most important. A platform that cuts moderation costs cuts credibility.

Phase 1: Progressive web app · Months 1 to 6

Core features · six languages · age-gated access · basic Schools layer

Budget line	Cost (€)	Notes
Progressive web app build	50,000	Ledger Live, Signal/Noise, basic Pulse, Archive, Founding Circle, text Spotter, Evidence Briefs download, Guided Knowledge Layer.
Schools layer: Guided Knowledge Layer and Evidence Briefs	8,000	Filtered content interface, brief templates, teacher account system, age-gate with parental consent flow.
Legal documents: Terms of Service, Privacy Policy, age consent flow	8,000	Legal Counsel drafts. Includes GDPR Article 8 age-gate, institutional declaration, RI formula publication.
Content management system	3,000	Signal/Noise daily publishing tool
Hosting (6 months)	4,800	€800/month · cloud infrastructure
Analytics setup	1,500	Privacy-compliant engagement measurement
Founding Circle: Stripe integration	3,500	Subscription management
Launch marketing and press	12,000	App announcement, press outreach
Security audit (pre-launch)	4,000	Independent penetration test
Part-time moderation (Phase 1 scale)	9,000	0.5 FTE at €1,500/month. Sufficient at Phase 1 submission volumes.
PHASE 1 TOTAL	**103,800**	**Progressive web app including Schools layer**

Phase 2: Native iOS and Android · Months 6 to 18

Full Spotter Network · professional verification · Schools group submission

Budget line	Cost (€)	Notes
Native app build (iOS + Android)	130,000	Camera Spotter, GPS, Reliability Index, convergence scoring, micro-payments, personalised Pulse, push notifications, community maps, Poll C to F, class group submission.
Intelligence Director (12 months)	96,000	1.0 FTE at €8,000/month. Governs all AI systems, RI algorithm, anomaly detection, and correction protocol.
Threshold Points and RI engine development	18,000	Points rules, RI calculation, appeal mechanism backend, anti-gaming logic, rank badge system.
Micro-payment infrastructure (Stripe Connect)	8,000	Contributor payouts. Identity verification for users over €100/year.
Moderation admin panel	12,000	Three-lane review tool
Frontline moderation team (12 months): 6 FTE	180,000	€2,500/month average blended rate. Global, distributed. 24-hour coverage. Schools-specific moderation protocols included.
Senior Validator team (12 months): 1.5 FTE	60,000	Deep review, payment disputes, cluster validation, correction protocol execution.
AI infrastructure (12 months)	48,000	€4,000/month. Scales with submission volume.
Micro-payment fund (12 months)	45,000	500 accepted entries/month at average €7.50.
App hosting and maintenance (12 months)	38,400	€3,200/month combined.
Schools layer: class group submission and teacher accounts	12,000	Teacher account system, group submission pipeline, GDPR-compliant student anonymisation.
Data Protection Officer (contracted, 12 months)	18,000	Legally required at Phase 2 scale under GDPR. €1,500/month contracted.
PHASE 2 TOTAL	**665,400**	**Native app + full verification + Schools + data protection**

Phase 3: Operations · Months 18 to 26

Running costs · no new major build

Budget line	Cost (€)	Notes
Moderation team (8 months)	120,000	6 FTE continuing
Senior Validators (8 months)	32,000	1.5 FTE continuing
Intelligence Director (8 months)	64,000	€8,000/month
AI infrastructure (8 months)	32,000	€4,000/month
Micro-payment fund (8 months)	30,000	Continuing
App hosting and maintenance (8 months)	25,600	€3,200/month
Data Protection Officer (8 months)	12,000	€1,500/month
PHASE 3 TOTAL	**315,600**	**Months 18 to 26**

Investment summary

Investment summary	Amount (€)
Phase 1: progressive web app	103,800
Phase 2: native app + full verification	665,400
Phase 3: running costs months 18 to 26	315,600
Contingency (10%)	108,480
TOTAL APP INVESTMENT (26-month programme)	1,193,280
Founding Circle revenue offset (2-year conservative estimate)	- 35,000
NET INVESTMENT REQUIRED	**1,158,280**

The Founding Circle at target (40,000 members, €280,000/year) covers approximately 24% of annual running costs from year 2 onwards and provides financial independence from any single funder. Season 2 of the platform costs approximately 30% less because the verification infrastructure, moderation team, and schools relationships are already in place.

8 REVENUE AND RETURN

How the platform generates sustainable income over time

Revenue stream	Season 1 (2-year)	Long-term potential
Founding Circle subscriptions	€315,000	€720,000/year at 40,000 members. €5.4M/year at 200,000 members.
Spotter usage rights (secondary publication use)	€35,000	Grows with programme publications and media outputs.
Schools and institutional licensing	Not Season 1	€80,000 to €150,000/year if methodology licensed to other organisations.
Season 2 cost reduction	Not applicable	Season 2 app approximately 30% cheaper: verification infrastructure already built.

Investment structures available

Investment type	Amount	What it funds	What the investor receives
Phase 1 sponsor	€103,800	Complete progressive web app including Schools layer, legal setup, age-gating, and six-month operations	Named in all publications as founding sponsor. Quarterly reporting. First right of refusal on Phase 2.
Phase 2 sponsor	€665,400	Native app build, full verification stack including Intelligence Director and professional moderation team, class group submission pathway, Data Protection Officer	Named in all publications as Phase 2 sponsor. Named in final report. Quarterly reporting.
Full programme sponsor	€1,158,280	Complete 26-month app programme from Phase 1 through to close-out and legacy	Named as founding sponsor in app and all publications. Annual briefing call with Lead Architect. Named in final report. First right of refusal on Season 2.
Micro-payment fund only	€45,000/year	The contributor usage rights fees: cash that goes directly to Spotters whose evidence is accepted into the Knowledge Base	Named in the Knowledge Base as the fund that makes Spotter payments possible. Noted in Inaction Ledger.

9 RISKS AND HOW THEY ARE MANAGED

Proportionate assessment: most risks are structural strengths in disguise

Every platform at this scale faces challenges. The ones listed here are the ones the programme has specifically designed for, not as afterthoughts, but as first-class architectural decisions.

Risk	Structural response
Commercial retaliation (app store complaints, legal pressure, coordinated fake submissions)	Users retain copyright; the platform holds only a usage licence. Distributed network of tens of thousands of Spotters makes total suppression difficult. Reliability Index and convergence scoring make coordinated fake submissions expensive. Strict neutrality positioning reduces the legal attack surface. Pre-funded legal defence arrangement.
False data entering the Knowledge Base	Six-step public correction protocol as a first-class feature: detection, assessment, downgrade, public retraction, user notification, audit log. Never quiet deletion. Permanent transparency record. This is a credibility feature, not a weakness.
Moderation quality under volume pressure	Three-lane validation model. Senior Validators supervise frontline team. Random audit of automated decisions. Weekly quality reporting. Published training standards. Appeal mechanism for all challenged decisions.
Child protection obligations	Age-gated access by legal requirement. No individual student tracking. Teacher-mediated class submissions. Moderation protocol specifically covering minor content. Overseen by the Oversight Board and governed by the Canons.
App fails to retain users after download	Ledger Live daily mechanic makes inaction continuously personal. Weekly Pulse streak sustains engagement. Meaning feedback at every contribution stage. The platform's expert programme outputs are valuable regardless of app user count; app scale is a multiplier, not a prerequisite.
Founding Circle revenue below target	The programme budget does not depend on Founding Circle revenue. It is a sustainability multiplier. The platform functions without it at the cost of slower progress toward long-term financial independence.
Reinsurance market repricing	Swiss Re SONAR 2025 and the Lloyd's Market Association September 2025 guidance both flagged polymer/plastic producer liability as an emerging risk category. Accelerating reinsurance withdrawal from plastic producer D&O and liability cover is a structural tailwind for this programme: it signals that institutional capital is beginning to price the risk that this evidence base is built to document. The programme is positioned upstream of this shift.

10 MARKET POSITION

A new category: not an improvement on existing platforms

The global environmental app market had approximately 850 million registered users in 2025. The citizen science market had approximately 8 million active participants. The policy engagement app market had approximately 12 million active users. No existing platform operates in all three simultaneously with a trust-weighted evidence system, a professional verification operation, a six-level institutional engagement cascade, and an educational layer. This is a new category.

Nearest competitor	What it does	The gap
Litterati / eBird / iNaturalist	Citizen science observation collection	Data collection only. No policy delivery mechanism. No expert analysis backbone. No educational layer. No institutional cascade.
Avaaz / Change.org	Petition and email campaign platforms	Signatures without evidence quality. No trust weighting. No institutional engagement architecture. The Threshold builds evidence; petitions collect agreement.
Environmental NGO apps (WWF, Greenpeace)	Campaign communication, fundraising	Advocacy-aligned: cannot claim neutrality. No citizen intelligence function. No educational evidence layer.
Global Fishing Watch / similar	Satellite or sensor data monitoring	Professional data only. No public participation. No constituency-building mechanism.

> *The gap in the market is precisely defined: a platform that turns citizen observation into trust-weighted evidence, delivers that evidence into a professional expert analysis process, and simultaneously builds the institutional constituency that makes the evidence impossible to dismiss. Nothing that does this exists. This is the opportunity.*

Institutional validation: the reinsurance signal

The most sophisticated risk-pricing institutions in the world are beginning to price what this programme documents. Swiss Re SONAR 2025 identified plastic pollution and polymer producer liability among emerging risk categories requiring attention from corporate underwriters. The Lloyd's Market Association September 2025 guidance noted accelerating withdrawal of liability cover for producers in sectors facing accumulating scientific and regulatory scrutiny. Neither document names plastics producers specifically; both describe exactly the risk architecture that this programme's evidence base is built to support. This is not coincidental. It is the institutional market reflecting, slowly, what the science has been saying for two decades. The programme is positioned to accelerate that reflection.

Contact

Investment and sponsorship	info@thresholdinitiative.org · Andrew Henderson · Lead Architect
App development and builder enquiries	info@thresholdinitiative.org
Schools and educational partnerships	info@thresholdinitiative.org
Media and documentary	info@thresholdinitiative.org

The Threshold App · Investor and Builder Brief · Version 7 · October 2026 · Confidential · not for public distribution. All budget figures are estimates. This document does not constitute an offer of investment or a legal commitment of any kind. Copyright THRESHOLD 2026. All rights reserved.

What follows from reading this

You have now read the evidence. The medical evidence is in our blood and our children's. The economic case for recycling, as it has been sold, is dead. The policy interventions exist and have been costed. The architecture that would resolve a case of this shape has been mapped.

From this point you are an informed witness. That is the threshold.

The next question is not whether you care. The historical pattern is clear. Systemic harm has been resolved in every documented precedent (tobacco, asbestos, lead, opioids) not by enough people caring but by enough informed witnesses acting in ways small and large. The relevant distinction is not between those who care and those who do not. It is between those who act and those who do not.

Three civic acts cross the threshold.

Vote. Vote for political representatives who will price the externality, build the levy, and pass the producer liability framework. Plastic is a political problem before it is a chemical one.

Vote in the Threshold poll. When the Knowledge Base presents the ranked interventions, vote on them. The poll is the public record of which interventions citizens have endorsed.

Write your Pledge. If you care, show it. Write to your unborn grandchildren. Tell them what you knew. What you did. What you did not do. Seal it. Send it to the Threshold. We will hold it until the threshold is crossed.

The Threshold Effect is the pattern by which documented systemic harm crosses a threshold making it increasingly difficult to ignore, legally, medically, and politically, in documented historical cases, deliberately and systematically. The Threshold Effect is in development.

The case has crossed every threshold but yours.

Whether you cross it is now observed.