MISSION

The internet learned to distribute claims instantly. It never learned to keep the reasoning.

Every day people meet assertions they have no practical way to check. Every tool built for them hands back a verdict, a note, or a generated answer and asks for trust — and the reasoning that produced it disappears the moment it is read.

VeriStrata makes that reasoning visible and keeps it. What evidence bears on a claim, who judged it, how strongly, and what would change the conclusion — recorded once, inspectable by anyone, and reusable by the next person who asks the same question.

Making public reasoning visible: not only what conclusion was reached, but how the evidence led us there.

A working evidence-mapping platform, built and running. Twelve months to put it in its first paid classroom.

$160,000	$20,000	12 months	$2,500–$10,000
target round	closed	to first paid pilot	pilot price

Alpha-stage estimates and validation targets, not forecasts.

01 / EXECUTIVE SUMMARY

Investment Case

VeriStrata has built the evidence-mapping system. This round puts it into classrooms and turns repeat use into revenue.

Every day people encounter claims they have no practical way to check. Existing tools usually give them a verdict, a note, or a generated answer and ask them to trust it. VeriStrata exposes the reasoning: the assertions a piece of content makes, the evidence that bears on each assertion, and the human judgments connecting them.

The product is working in alpha. Content goes in; the system extracts its important assertions, retrieves candidate sources, decomposes those sources, and proposes evidentiary relationships. Evaluators accept, reject, or revise those proposals. A browser extension returns the result at the point of reading. Roughly thirty people have used the product in guided sessions.

The first market is deliberately narrow: university programs where evaluating claims is already assigned work — media literacy, journalism, data science, civic reasoning, and misinformation research. A department, lab, or civic program pays $2,500–$10,000 for a course, workshop sequence, or scoped pilot. That market has not been sold into yet: there are no signed pilots and no revenue today, which is exactly what this round is scoped and timed to change.

Already built	What this round tests	Month 12 evidence
Working alpha and browser extension	Repeat use across assignments	First paid pilot running
AI-assisted evidence workflow	Instructor value and student usability	3–5 active pilots or evaluations
Governance and audit primitives	Whether reviewed maps improve through use	A credible basis for the next round

VeriStrata is raising $160,000 on a $6 million post-money valuation cap. $20,000 is closed across two signed VeriStrata SAFEs, alongside $2,000 of assumed legacy instruments, for $22,000 of invested capital to date. The base round buys fifteen months — twelve to a first paid pilot, plus three months of margin beyond it — and funds the move from working alpha to pilot-ready product, pipeline creation, and the first paid institutional test.

The larger opportunity is an evidence layer for the internet: a reusable, inspectable record of how disputed assertions were evaluated, by whom, and against what evidence. The institutional wedge is how that gets built — every mapped classroom assignment adds to the same corpus, and every reviewed judgment makes the next extraction better.

Every evaluator session also produces a second asset alongside the evidence map: a labelled record of where a person agreed with the pipeline and where they overrode it — on assertion boundaries, on whether a source bears on a claim, on whether a link is support or context. That is training signal for exactly the task the pipeline performs, generated as a by-product of coursework an institution is paying for. It is a validation target rather than a proven moat, and section 12 sets out what would have to be true for it to become one.

02 / PROBLEM + WHY NOW

The Missing Layer

The internet can distribute claims instantly. It still cannot preserve the reasoning used to evaluate them.

Traditional fact-checkers produce valuable work, but most output remains article-centric: a reader gets a published conclusion, while the underlying evidentiary relationships are difficult to inspect or reuse. Community Notes reaches people inside the feed, but the note stays bound to a particular post and platform. General AI assistants move fast, but their reasoning is generated anew for each user and rarely becomes a durable, reviewable object.

The missing artifact is not another verdict. It is a shared evidence map that the next person can inspect, challenge, and extend instead of restarting the investigation.

Why this is timely

Change	Implication for VeriStrata
Institutional fact-checking has retreated from major feeds	The community model is expanding, but it needs inspectable evidence underneath the note.
Evidence evaluation is becoming an explicit teaching objective	Universities already assign the cognitive work VeriStrata structures.
Generative systems make plausible claims and citations nearly free	Manual, article-by-article review cannot keep up; reusable evidence relationships can compound.

The dates matter because they are checkable. Meta announced the end of its US third-party fact-checking programme in January 2025 and shut it down that March, replacing it with a crowd-sourced Community Notes model. Since 2021 a growing set of states have moved media and information literacy from optional to required — New Jersey's statewide K–12 information literacy law in January 2023, Delaware's Digital Citizenship Education Act in 2022, media literacy requirements in Illinois and Texas, and California's comprehensive law in October 2023. Those are K–12 mandates and VeriStrata's buyer is a university programme, so they are not the purchase order. They establish that evaluating evidence is now treated as a teachable, examinable skill.

These forces do not prove demand for VeriStrata. They make the validation experiment worth running now: can a structured, human-accountable workflow turn evidence evaluation from a one-off assignment into a reusable institutional product?

The buyer’s job

An instructor needs students to distinguish claims from evidence, compare sources, explain why a source bears on a claim, and expose their reasoning for review. Today that work disappears into papers, slide decks, and discussion threads. VeriStrata turns it into a visible map that can be graded, challenged, reused, and published.

03 / PRODUCT

The Product: From Claim to Evidence Map

AI assembles the raw material. People make the accountable evidentiary judgments.

A reader encounters an article, video, podcast, or post. If it has already been mapped, the VeriStrata browser extension surfaces a compact evidence signal without forcing the reader to leave the page. The reader can move from the signal to the content's key assertions, then to the evidence on both sides and the reasoning behind every link.

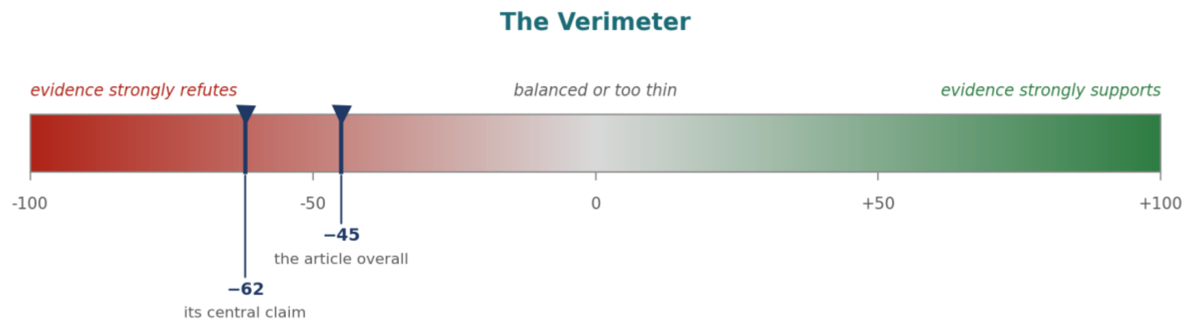

The Verimeter is a doorway into the evidence map, not a substitute for it.

If the content has not been mapped, it enters the evaluator workflow. The platform proposes the important assertions, finds evidence candidates, extracts relevant source assertions, and suggests how those assertions may bear on the original. Evaluators accept, reject, re-scope, or reverse those proposals and record their rationale and confidence.

One pass, concretely. An article argues that blue-blocking glasses meaningfully improve sleep. The extension returns a Verimeter of −45. Five assertions from the article have been evaluated and the central one sits at −62 — four supporting assertions drawn from two small lab studies, against nine refuting or qualifying assertions drawn from seven sources, including a meta-analysis the article never cites. The reader has not been told what to think. In five seconds they have been shown where the evidence sits and which source to read instead.

Three product nouns are enough

Term	Meaning
Assertion	A checkable statement extracted from content or a source.
Evidence map	The human-judged links bearing on one assertion.
Verimeter	A −100 to +100 summary of the mapped evidence, with every input inspectable.

Where the number comes from

Source reliability is not VeriStrata's own opinion of who is trustworthy. It is assembled from independent providers, each covering a different kind of source: Wikidata for publisher and author identity, Wikipedia Perennial Sources for community-maintained reliability assessments, SCImago for journal ranking, Crossref and OpenAlex for scholarly metadata and citation context, Media Bias/Fact Check and the Global Disinformation Index for news domains, GDELT for coverage context, the Internet Archive for archival footprint, RDAP for domain provenance, OpenSanctions for enforcement exposure, and public complaint, regulatory, and litigation records for commercial publishers.

Coverage is uneven and the system reports that rather than hiding it. In a live run against an archived CDC document, six providers returned no match and the only available signal was a weak archival footprint, which caps the grade rather than raising it. Where providers do not resolve a source, the link carries less weight and the gap stays visible in the map instead of being filled with a default. Appendix B sets out the full provider table and the trust model.

04 / WEDGE

The First Customer and Pilot

Start where evidence evaluation is already required—not where participation must first be invented.

The initial customer is an instructor, department, research lab, journalism program, or civic-learning organization already responsible for teaching or practicing evidence evaluation. VeriStrata supports work these buyers have to do anyway. It does not ask them to launch a new volunteer community.

Pilot proposition

An instructor runs VeriStrata across at least two assignments or workshops. Participants turn disputed content into assertions, inspect candidate sources, connect source assertions with typed links, explain their reasoning, and review one another's maps. The instructor receives an assessable record of how conclusions were reached, not only a final answer.

Pilot question	Success signal
Will participants use it more than once?	Repeat use across at least two assignments or sessions.
Does it make reasoning assessable?	Instructors can inspect and compare evidentiary choices.
Does review improve the output?	Maps materially change through peer or instructor challenge.
Does the work become reusable?	A later evaluator extends a completed map instead of restarting.
Will a budget owner pay?	$2,500–$10,000 paid pilot or renewal commitment.

Procurement strategy

The first pilots are scoped below institutional procurement thresholds: departmental or grant-funded purchases, no LMS integration, and no institutional roster data. That makes a Month 12 paid pilot plausible on this budget. It also imposes a ceiling: larger deployments will require accessibility, privacy, security, and LMS work funded in a later round.

05 / COMMERCIAL TEST

Validation and Go-to-Market

The round funds pipeline creation: demos into scoped proposals, proposals into the first paid pilot.

VeriStrata has contacts at two universities and one active conversation: a civic organization forming in the San Diego area, where meetings toward a scoped pilot are underway. No written scope, price, or budget owner has been agreed, and there is no letter of intent or paid commitment. The commercial plan begins from that position and treats outreach as a measurable operating function.

Period	Commercial work	Gate
Months 0–3	Finish onboarding and demo package; deliver 15–20 demos to named instructors or lab leads.	Pilot-ready product and a qualified prospect list.
Months 4–6	Turn the strongest conversations into written pilot scopes for the Fall 2027 term.	2–3 programs agree in writing to scoped pilots.
Months 7–9	Complete contracting and purchasing while campuses are between terms; begin next-round fundraising.	Fall commitments or purchase orders in process.
Months 10–12	Support launch and measure repeat use, map quality, instructor value, and willingness to pay.	First paid pilot running; 3–5 active pilots or evaluations.

Founder-led selling is the primary motion. Reed O’Neal remains the only full-time operator, so demos compete directly with development time. The plan therefore limits both the feature agenda and the account target. It assumes no sales capacity from advisors and no viral referral channel until a pilot demonstrates one.

What gets measured

- Demo-to-scoped-pilot conversion and time to budget-owner identification.
- Participant activation, completion, and repeat use across assignments.
- Maps created, sources opened, links revised, and peer challenges resolved.
- Instructor-reported teaching value and willingness to renew or pay.
- Inference and search cost per completed map.

06 / POSITIONING

Competition and Market

VeriStrata competes for assignment time by structuring evidence—not discussion, annotation, or a verdict.

Alternative	What it does well	Gap VeriStrata addresses
Kialo Edu	Structures pro/con argument trees.	Does not decompose sources into assertions or preserve evidentiary provenance.
Hypothes.is	Anchors collaborative annotations to documents and integrates with LMSs.	Annotations do not compose into typed, reusable evidence maps.
Community Notes	Places context at the point of reading.	Platform- and post-bound; no durable evidence structure beneath the note.
Traditional fact-checkers	Produce expert reporting and defensible published conclusions.	Output is hard to reuse as a student-facing, inspectable workflow.
General AI assistants	Return fast answers with citations.	No persistent evidence object, accountable evaluator, or review history.

Bottom-up market

Layer	Basis	Scale
Wedge TAM	2,000–3,000 relevant English-language or English-supported programs × $10,000	≈ $20M–$30M
Five-year SAM	300–500 programs reachable through founder- and partner-led selling × $10,000	≈ $3M–$5M
36-month base case	8–15 paying accounts at $5,000–$10,000	≈ $40K–$150K ARR

This is intentionally narrower than a top-down "information integrity" market. Newsrooms, NGOs, public-health organizations, trust-and-safety teams, and provenance APIs are possible expansion paths, but none is included in the base case because none has been commercially tested.

07 / REVENUE

Business Model and Economics

Near term, VeriStrata sells software and workflow—not verdicts.

The pricing experiment begins at $2,500–$10,000 for a course, lab, workshop sequence, or civic program, depending on scope and participant count. A $10,000–$35,000 departmental license is an expansion hypothesis, not part of the base case until a buyer renews or broadens a pilot.

Period	Conservative	Base	Upside
Months 0–18	1–2 paid pilots; $2.5K–$10K	3–5 paid pilots; $7.5K–$50K	6–10 paid pilots; $15K–$100K
Months 18–36	5–8 repeat accounts; $12.5K–$40K ARR	8–15 accounts; $40K–$150K ARR	15–25 accounts incl. 2–4 departmental licenses; $110K–$310K ARR

Unit economics to prove

At a $2,500 course pilot producing roughly 300 maps, revenue is about $8.33 per map. To preserve a 90% gross margin, inference and retrieval cost must remain below roughly $0.83 per completed map; at 80% gross margin the ceiling is roughly $1.67. The alpha already screens candidate assertions before full processing to prevent cost from scaling without bound. Pilot telemetry must show whether that architecture holds under classroom use.

The Months 18–36 base case assumes founder-led selling at roughly one new account per quarter plus some referral inside a small, connected discipline. Referral remains a hope, not a channel, until a successful pilot produces a case study that travels.

08 / TEAM

Founder and Operating Model

A complete system built before institutional capital — and a plan sized honestly to one operator.

Reed O’Neal built VeriStrata and leads engineering, product design, customer discovery, demos, support, and fundraising. He is a senior systems engineer and full-stack builder with more than twenty years modernizing financial and operational systems, including database migrations and GLASS, a general-ledger suite used in hedge-fund administration during growth from roughly $2 billion to $20 billion in assets under management.

Eric Forst advises on commercialization and go-to-market. His present role is limited and no selling capacity is assumed from him in the operating plan. Lane Rasberry, Wikipedian in Residence at the University of Virginia, advises on volunteer knowledge communities, editor recruitment, source standards, and governance norms.

Operating constraint

There is no second full-time operator. The budget and milestones are written on that basis. Product work is limited to pilot readiness, onboarding, reliability, and the classroom workflow. Features that do not improve conversion, activation, map quality, or cost per map are deferred.

Strength	Risk	Operating response
Founder built the working system before institutional capital.	Product, selling, and support depend on one person.	Keep burn low, document systems, and add focused contractor capacity only around pilot traction.
Remote-first, browser-based product.	Founder works from Hong Kong while selling to U.S. institutions.	Prioritize Eastern-time outreach; reserve travel for live conversion opportunities.
Advisors cover commercialization and governance.	Advisors are not full-time operators.	Do not count their capacity in the base model.

09 / FINANCING

Raise, Runway, and Use of Funds

$160,000 buys fifteen months: twelve to a first paid pilot, and three beyond it to raise on the result.

Use	$160K base	Purpose
Founder runway	$90,000	Full-time product, pilot, sales, support, and fundraising execution.
Product and UX	$18,000	Onboarding, classroom workflow, reliability, and focused contractor support.
Infrastructure + AI/API	$20,000	Hosting, monitoring, search, extraction, and pilot inference.
Academic partnerships	$9,000	Demo materials and travel tied to conversion.
Legal + IP	$8,000	SAFE, privacy, terms, trademark, and remaining cleanup.
Marketing + investor prep	$10,000	Case studies, content, demo video, and data room.
Reserve	$5,000	Unexpected product or infrastructure needs.

Total monthly burn in the base case is approximately $10,667 across fifteen months. Founder compensation is $6,000 per month, or $72,000 annualized and 56% of the round. It funds the only full-time builder, operator, seller, and support lead and remains below market compensation for the role.

That figure applies from the close. It follows roughly $148,600 of the founder's own capital drawn for living costs through 2025 while the platform was built — a year of full-time development financed entirely by the founder, with no outside investment and no salary from any other source. The $20,000 received before the close was disbursed as founder compensation during the raise, at a higher monthly rate than the funded plan, and is disclosed in the management-prepared financial statements.

Terms and current position

- YC post-money SAFE, valuation cap only, no discount, at a $6,000,000 cap.
- $20,000 closed across two signed VeriStrata SAFEs; $140,000 remains in the target round. Total invested capital to date, including $2,000 of assumed legacy instruments, is $22,000.
- Reed O'Neal owns 100% of currently issued common stock before SAFE conversion and future grants.

What the cap is priced on

The $6,000,000 cap is not priced on the base-case revenue in section 07. At 8–15 accounts and $40,000–$150,000 of ARR at 36 months, the institutional wedge alone does not support it, and presenting it otherwise would not survive an hour of diligence.

It is priced on three things. First, what already exists: roughly two years of self-financed, full-time development produced a working alpha, a browser extension, the five-stage evidence pipeline, the governance and audit primitives, and the reliability stack in Appendix B — all before institutional capital. This round does not fund building the system; it funds selling it. Second, what the instrument buys: $160,000 at a $6,000,000 post-money cap is approximately 2.67% of the company, on an unmodified Y Combinator SAFE with a valuation cap only — no discount, no side letters, no maturity date, and no non-standard terms. Third, the asset described in section 12: a corpus of attributable evidentiary judgments and correction records generated as a by-product of coursework an institution pays for.

The third is the one that has to be true, and it is a validation target rather than a proven moat. An investor who does not credit that thesis should price this company on section 07 alone and conclude the cap is too high. That is a defensible conclusion and it is the right question to press on. This round exists to test whether the paid workflow underneath it is real.

10 / THE ASK

How to Participate

$140,000 of the $160,000 target remains open.

The instrument is a Y Combinator post-money SAFE, valuation cap only, no discount, at a $6,000,000 post-money cap. $20,000 is closed across two signed VeriStrata SAFEs; with the $2,000 of assumed legacy instruments described in Appendix D, total invested capital to date is $22,000. Minimum direct investment is $2,500. A $50,000 bridge tranche closes first and funds six months, sized to reach the Month 6 written-agreement gate; the full round funds fifteen and carries the company through a first paid institutional pilot in the Fall 2027 term with three months of margin beyond it. A parallel Wefunder community round will carry its own lower minimum.

What the money buys is set out in section 09. What it must produce is set out in section 11, together with the conditions under which the company stops instead.

To invest, review the data room and contact Reed O'Neal directly. Instrument copies, the current capitalization table, the founder IP assignment, and supporting financial records are available there.

Contact: reedko@veristrata.ai · Data room and instrument copies available on request

11 / ACCOUNTABILITY

Milestones, Kill Criteria, and Risks

A validation round should make failure visible while enough capital remains to respond.

Gate	Failure condition	Action
Month 6	Fewer than two programs agree in writing to a scoped Fall pilot.	Stop selling into academia as the primary wedge; re-scope toward civic or newsroom buyers.
Month 12	Zero paid pilots running.	The institutional wedge is not the business; pivot toward the public extension product or wind down.
Month 18*	Fewer than two accounts renew for a second term.	Do not raise a seed on the institutional thesis; re-scope or stop.

*Month 18 is beyond the fifteen months funded by the base round; only the $250,000 acceleration case funds through it. It is otherwise reached with a subsequent round raised from Month 9 against written Fall commitments and early launch evidence.

Risk	Why it matters	Mitigation
Thin pipeline	One active pilot conversation; no scoped pilot, price, or budget owner exists today.	15–20 named demos by Month 3 and written-agreement gate by Month 6.
Founder concentration	Development and selling compete for the same hours.	Narrow feature agenda; conservative account targets; contractor support tied to traction.
Cognitive labor	Evidence mapping asks users to think carefully.	Begin where that work is already assigned and assessable.
AI quality	Extraction and retrieval are noisy.	AI proposes; humans decide; every link remains visible and editable.
Governance / brigading	Public ratings invite manipulation.	Separate evidence reading from crowd sentiment; review thresholds, reputation, and tamper-evident history.
Legal / reputational	Publisher-level signals create exposure.	Private pilot workspaces; named third-party reliability inputs; legal review before public publisher scoring.

12 / VISION

From Institutional Tool to Public Evidence Layer

The classroom builds what the public layer runs on: the method, the trained evaluators, and a source-level record that carries across.

VeriStrata has two horizons. The near-term business is institutional evidence mapping for classrooms, research teams, journalism programs, and civic organizations. The long-term product is a public evidence layer consulted wherever disputed assertions appear.

The bridge must be stated honestly. Assigned classroom readings will often have little topical overlap with today's public controversies. What transfers is the typed relationship vocabulary, the evaluation and peer-review workflow, a population of trained evaluators, and source- or publisher-level history. Individual evidence maps will transfer only when topics overlap.

Distribution is the harder problem. Wikipedia shows that governed contributors can build durable public infrastructure. Community Notes shows that governance is not enough: the result must appear where the reading happens. VeriStrata's browser extension is therefore a convenience in the institutional phase and the central distribution product in the public phase.

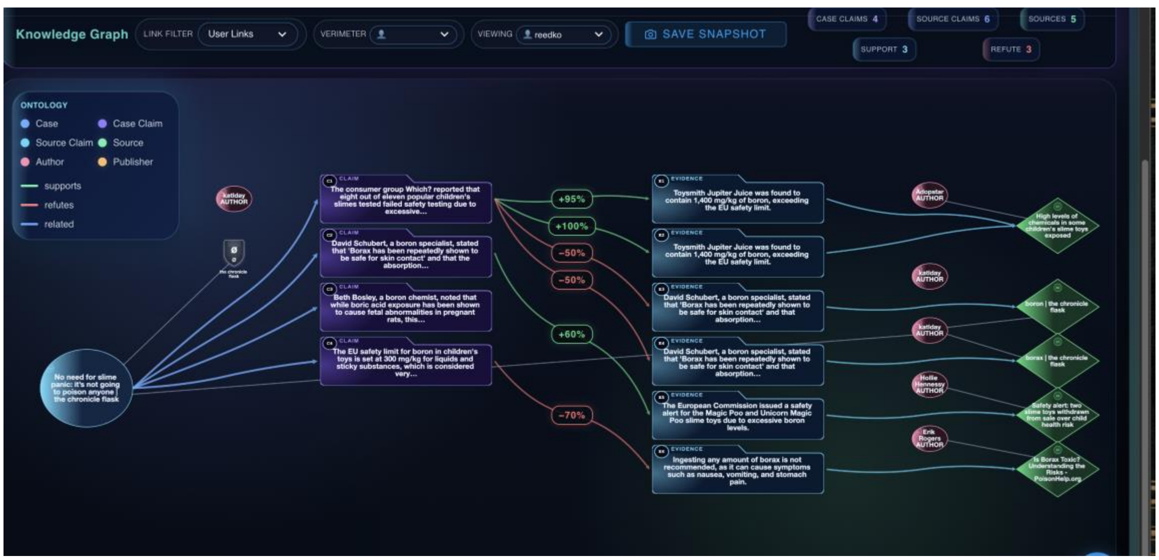

A VeriStrata knowledge graph: assertions, source assertions, evaluators, publishers, and evidentiary links remain inspectable.

Why the work compounds

Every evaluator session produces a second asset alongside the evidence map: a labelled record of where a person agreed with the pipeline and where they overrode it — on assertion boundaries, on whether a source bears on a claim, on whether a link is support or context.

That is training signal for exactly the task the pipeline performs, generated as a by-product of coursework an institution is paying for. A media literacy class produces an evidence map for the corpus and a correction record that makes the next extraction better. Both are only worth something because every judgment is attributable, which is why the governance layer was built before the demand that justifies it. It is a validation target rather than a proven moat, and it requires the pilots to happen.

The durable opportunity is a corpus of attributable evidentiary judgments and correction records that improves through use. That could become valuable to education, verification, publishing, search, and AI provenance. It is not a moat today. This round tests the paid workflow that would make the corpus possible.

Make public reasoning visible: what evidence led us here, who evaluated it, and what would change our minds?

DILIGENCE APPENDICES

Appendix A — Product and Technology

The main plan describes the customer outcome. This appendix records the alpha architecture and the technical assertions requiring pilot validation.

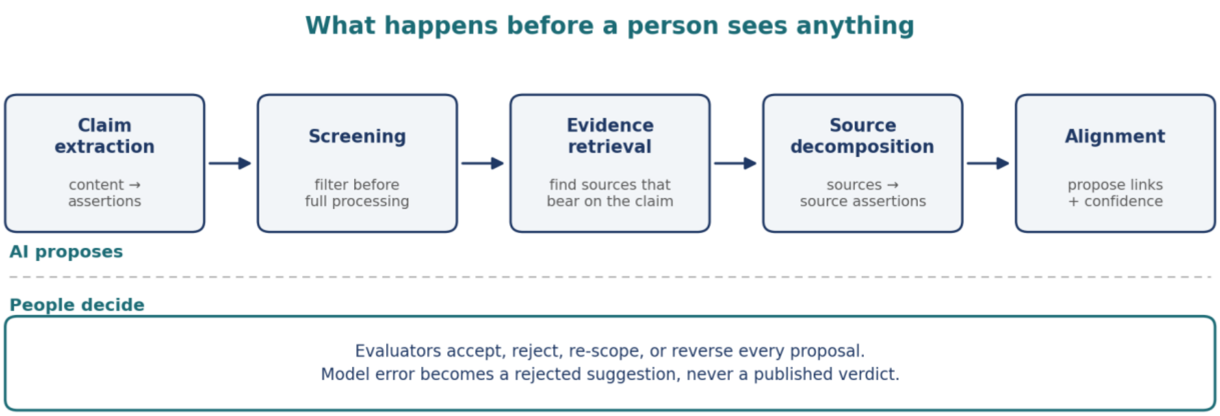

Five stages prepare evidence for review. Every output remains a proposal until a person judges it.

Stage	Function	Validation issue
Assertion extraction	Preserve the checkable statements the author actually made.	Avoid paraphrase drift and over-broad assertions.
Screening	Filter candidates before expensive processing.	Hold cost per completed map inside the margin ceiling.
Evidence retrieval	Find sources that bear on the assertion, including refuting material.	Measure relevance, diversity, and unresolved cases.
Source decomposition	Extract the germane assertions from each source.	Balance useful granularity against noise.
Alignment	Propose support, refute, context, or insufficient-evidence links.	Humans must be able to correct unstable or mistaken bearing judgments.

Current production objective: a useful set of approximately correct evidence candidates with human review, not perfectly stable model judgments. The pilot must test whether noise degrades into rejected suggestions rather than published error and whether evaluators can complete maps efficiently.

Appendix B — Verimeter and Trust Model

The Verimeter runs from −100 to +100. Negative values indicate that mapped evidence refutes an assertion; positive values indicate support. Zero can mean balanced or thin evidence, which remains visible in the underlying map.

Each link carries direction and strength, then is weighted by source reliability and evaluator standing. Reliability is not VeriStrata's own opinion of who is trustworthy: it is assembled from independent providers, each covering a different kind of source.

Provider	What it contributes
Wikidata	Entity identity — who a publisher or author actually is
Wikipedia Perennial Sources	Community-maintained reliability assessments
SCImago	Journal impact ranking for academic sources
Crossref / OpenAlex	Scholarly metadata, DOI resolution, citation context
Media Bias/Fact Check	Independent media reliability and bias assessment
Global Disinformation Index	Risk assessment for news domains
GDELT	Global coverage and event context
Internet Archive	Archival footprint. No reliability signal, but a thin footprint caps the grade an obscure source can reach
RDAP	Domain registration provenance
OpenSanctions	Sanctions and enforcement exposure
BBB, CFPB, court records	Complaint, enforcement, and litigation history for commercial publishers

Coverage is uneven and the system reports that rather than hiding it. A peer-reviewed article resolves cleanly through Crossref, OpenAlex, and SCImago. A mainstream news domain resolves through Media Bias/Fact Check and the Global Disinformation Index. A government archive page may match almost nothing: in a live run against an archived CDC document, six providers returned no match and the only signal available was a weak archival footprint, which caps the grade rather than raising it. Where providers do not resolve a source the link carries less weight and the gap stays visible in the map instead of being filled with a default.

This matters beyond accuracy. Reliability weights trace to named external authorities, evidentiary judgments trace to named evaluators with recorded reasoning, and the arithmetic between them is inspectable. That is a materially different posture from a platform issuing ratings of its own.

Two readings, never blended

Measure	What it represents	What can move it
Verimeter	The evidence reading.	Reviewed evidentiary links, source signals, and evaluator standing.
Tally	The crowd’s expressed position.	Moderated participant responses, reported as proportions.

The two can disagree, and that disagreement is information. A coordinated group may move the Tally; it cannot move the Verimeter without producing attributable links that survive review. Pilot workspaces keep publisher-level ratings private while the legal and adversarial posture remains untested.

Trust primitive	Current status	Pilot validation
Tamper-evident history	Built; not load-tested	Verify concurrent classroom use and contribution history.
Approval workflow	Built; not pilot-tested	Measure review consistency and disputed-map handling.
Reviewer eligibility / reputation	Built; not pilot-tested	Test whether weighting improves reliability without entrenching incumbents.
Anti-brigading controls	Partial / designed	Adversarial testing remains future work.
Evidence corpus moat	Validation target	Depends on repeat use, quality, reuse, and scale.

Appendix C — Procurement and Compliance

Gate	Current posture	Expansion work
LMS / LTI 1.3	Not required for scoped pilots; web login only.	Integration required for scaled departmental deployment.
Accessibility	Pilot UI requires practical accessibility review.	VPAT and WCAG 2.1 AA program before institutional scaling.
Student data / FERPA	No institutional roster data in initial pilots.	Formal data-flow and contractual posture for larger deployments.
Security review	Scoped to purchases below formal review thresholds.	HECVAT/security documentation and controls in a later round.

This strategy is a wedge, not an evasion. It makes the first paid experiment possible without pretending the alpha clears enterprise procurement. No account can expand materially until the compliance work is funded and completed.

Pilot data package

- Written pilot scope and success measures agreed before launch.
- No LMS roster import or unnecessary student data.
- Usage, cost, completion, repeat-use, and map-quality telemetry.
- Instructor interview and renewal or payment decision at completion.
- Exportable work product and clear treatment of participant intellectual contributions.

Appendix D — Capitalization and Predecessor Entity

Item	Position
Issuer	VeriStrata, Inc., Delaware corporation.
Current instrument	YC post-money SAFE; $6M cap; no discount.
Round target / closed	$160,000 target; $20,000 closed across two signed VeriStrata SAFEs. With the assumed legacy instruments below, $22,000 of invested capital to date.
Legacy instruments	Two $1,000 TruthTrollers, LLC instruments transferred to and assumed by VeriStrata, Inc. on their original $6M cap and 20% discount terms. Executed.
Founder ownership	Reed O'Neal owns 100% of currently issued common stock before conversions and grants.
Option pool	None created; future board decision.
Advisor equity	No equity issued to Eric Forst as of August 2026.
IP	Founder Confidential Information and Inventions Assignment Agreement executed 23 June 2026, assigning founder work product to VeriStrata, Inc. A confirmatory assignment covering pre-incorporation intellectual property, from TruthTrollers, LLC and from the founder individually, is drafted and will be executed before the round closes and before the LLC is dissolved. Founder is the sole member of TruthTrollers, LLC and the sole holder of VeriStrata, Inc. common stock; no third party has written code, created assets, or been granted any interest in the platform intellectual property.

TruthTrollers, LLC and VeriStrata, Inc. are separate entities. The operating intellectual property and current new-money SAFEs sit in VeriStrata. The LLC is intended to be wound down after the legacy transfers complete. The legacy documents and countersignature status should be refreshed immediately before investor distribution.

Historical development context

The LLC's 2025 records reflect founder-financed development, not a go-forward operating model: $150,005 of founder contributions, $148,611.89 returned as owner draws supporting personal runway, $1,284.81 of recorded business expense, and no revenue. The records are management-prepared and unaudited. Offering-specific financial statements should be prepared to the standard required by any Regulation Crowdfunding campaign.

Appendix E — Extended Financials and Diligence Checklist

Scenario	Runway	Monthly burn	Primary outcome
$50K bridge	6 months	$8,333	Reaches the Month 6 written-agreement gate. Hardens the alpha and carries 2–3 conversations to written scope; not a revenue path.
$160K base	15 months	$10,667	First paid pilot at Month 12 and 3–5 active pilots/evaluations, with three months of runway beyond the gate.
$250K acceleration	20 months	$12,500	Additional UX, partnership, and pilot capacity, and the only scenario funded through the Month 18 renewal gate; terms to be separately defined.

Investor diligence checklist

- Current cap table and board approvals.
- All executed VeriStrata SAFEs.
- Executed legacy transfer and assumption documents.
- Founder IP assignment and product ownership record.
- Management financials and use-of-funds model.
- Product demonstration, architecture summary, and pilot telemetry plan.
- Privacy, terms, security posture, and pilot data-flow statement.
- Named prospect list, demo log, written pilot scopes, and conversion evidence as created.

Refresh rule: dates, closed financing, transfer status, pipeline status, and milestone calendar should be updated before every distribution. The document is a validation plan, not a forecast, and should remain explicit about what is built, what is observed, and what is still only a target.