

0. About This Report

This is the LIGHT HOPE 2026 Interim Social Responsibility Report. It covers the reporting period 2026-01-01 through 2026-06-30. The data freeze point is 2026-08-11, forty-two days after period end, chosen so that subsequent events immediately following the period could be recorded rather than hidden. Every factual statement in this report traces to a single frozen evidence set; where a statement carries a two-part label, the first part is the delivery stage and the second is the evidence grade, defined in Chapter 2.

The following statement of identity is made verbatim and governs everything that follows.

LIGHT HOPE LLC is a Massachusetts limited liability company. It is not a 501(c)(3) tax-exempt organization, a charity, or an NGO. This report is a voluntary corporate disclosure. It is not an audit report, an assurance engagement, a solicitation, or investment, medical, legal, or tax advice. No contribution to LIGHT HOPE LLC is tax-deductible.

What this report does not claim. It does not claim proven outcomes. It contains zero claims at the highest combined grade of measured-outcome plus machine-traceable evidence; that count is zero by design and by honest necessity, and Chapter 2 explains why. This report has not been independently audited. Its contents are self-reported and internally cross-checked. Where a fact cannot be traced to the frozen evidence set, it was dropped rather than softened.

Scope. The report describes a bounded roster of admitted projects, selected by the admission rules in Chapter 2. Projects that were excluded — for skeleton-only status, for containing personal data that cannot be published, for political sensitivity, or for hard firewall separation — are absent by rule, not by oversight. The number of admitted projects is deliberately small relative to the wider portfolio; a report is only as honest as the evidence it can stand behind.

1. Empower Good With AI

"Empower Good With AI" is the sentence LIGHT HOPE organizes its work around. A slogan is not a commitment until it can be checked, so this chapter states what the phrase means operationally — as constraints visible in code and behavior, not as aspiration.

Good is enforced before the model runs, not after. In the evidence-navigation factory, a safety filter sits ahead of the language model, and the internal test battery records that adversarial inputs are intercepted before the model is ever invoked — 341 of 341 fixtures in the factory's own suite. This is a self-reported internal result, and it is described that way; the point is architectural placement, not the score. [SVR-C4-m320-002]

Good is expressed as structure, not policy. In the household-finance knowledge graph, there are physically no score, rank, or best_product nodes or edges. The system therefore cannot represent "product X is best for person Y" even if asked. A policy can be bypassed; an absent data structure cannot be. This is the clearest instance of the principle in the portfolio. [SVR-C5-m631-001]

Good means the crisis path is not optional. In the low-back-pain navigator, crisis hotline numbers and an emergency red-flag redirect are hard-coded into the copy and are not disableable by configuration. [SVR-C4-m321-002]

Good means discarding the number that flatters you. The AI-education review pipeline excluded a retracted 2026-04 meta-analysis whose reported effect size, $g=0.867$, would have been the most favorable figure available to it. An automatic process threw away its own best headline because the source had been retracted. That single act says more about whether the discipline is real than any statement of values could. [SVR-C6-m414-007]

Good means publishing the zero. The clearest test of the phrase is whether we report the numbers that make us look unfinished. The hiring platform had received zero applications in its first ten days live, and this report records that zero as a fact worth stating rather than a defect worth hiding. [SVR-C7-m647-002]

The rest of this report is the evidence for and against these five sentences.

2. How We Measure

This is the longest chapter because the method is the disclosure. A social responsibility report that does not explain how it decided what counts as true is asking to be believed on faith. We measure every claim on two independent axes, we admit projects by a fixed rule, we exclude by a fixed rule, and we publish our zeros. Each of those choices is described here.

Two axes, never one

Most reporting collapses "how far did the work get" and "how good is the evidence" into a single impression. We separate them, because the most dangerous error in a report like this is to write "the evidence is real" and let the reader hear "the impact has happened."

The **delivery-stage** axis has five values, each higher one implying the lower ones hold: STAGED (material on disk, not built), BUILT (builds and tests pass, no public surface), LIVE (publicly reachable, with a dated probe), USED (a real, non-team action was counted — and the count may be zero), and OUTCOME (a measurable change in someone's situation, with an attribution argument).

The **evidence-grade** axis has three values: VERIFIED (machine-traceable to a frozen snapshot field, a probe result, or a disk path plus commit), SELF-ATTESTED (an internal test, gate, or audit log, with no independent check), and UNAUDITABLE (no evidence path — stated together with the reason it cannot be audited).

A claim missing either label is dropped, not downgraded. UNAUDITABLE is not a failure label; it is an honesty label, and it always carries its reason.

The rules that bind the two axes

Five combination rules hold across the report. A LIVE claim requires a dated probe, and a probe more than a few weeks old is downgraded rather than trusted. A USED claim must name a specific action metric, and a value of zero must be written out. A summary number inherits the weakest grade among its parts. Stages may not be skipped in narration — "we built X, therefore families now have Y" is not permitted; each step needs its own evidence. And the rule that governs the whole report:

This report contains zero claims at the OUTCOME · VERIFIED level. We have not established causal attribution for any change in any person's situation. Where we discuss outcomes at all, they are labeled

UNAUDITABLE, and we state why.

The single OUTCOME-stage claim in the entire ledger — a tax platform's estimate that it cuts manual classification time by 90 percent or more — is labeled UNAUDITABLE precisely because no usage data or real users exist to support it. [SVR-C5-m548-007]

Admission, and how many did not make it

Projects entered this report through a fixed roster, not editorial choice. Of the candidate set drawn from the ecosystem's frozen snapshots, twenty-four were admitted. Two were held as conditional because they had no current snapshot entry to stand on and could not be verified in time. None was blocked outright. An earlier data cut had blocked the education chapter's intended core project and held back a longevity project for stale snapshots; a refresh at the freeze date cleared both, and they are admitted here. The admitted count is small on purpose; the alternative to a small honest roster is a large padded one.

The seven exclusion categories

Exclusion is fail-closed: a project on the exclusion list cannot appear in any form, including anonymized or aggregated counts. Seven categories were excluded, described here without naming the projects: skeleton or placeholder builds with no substance; work containing personal data that has not cleared a privacy review, including anything touching minors; identity-verification work whose human subjects have not approved publication; politically sensitive material; legal-stream work held behind a hard firewall; access-controlled internal systems; and one layered case admitted only at its public tier, with its private tier excluded.

Why we publish the zeros

A social responsibility report with no zero values and no unauditible items is not clean — it is unbelievable. The production pipeline that generated this report is built to halt if fewer than three zero-value usage claims or fewer than five unauditible claims appear, on the theory that a report which is too flattering is more likely to be fabricated than to be true. Concrete zeros in this edition include a hiring platform's zero applications [SVR-C7-m647-002], a curriculum platform's zero real external users [SVR-C6-m552-002], an expert atlas with readership at zero and an unwired analytics tag [SVR-C3-m554-003], and a completed book series with zero real readers [SVR-C6-m800-006].

The limits of our own evidence

Even the VERIFIED grade has boundaries we state openly. Several projects' snapshots contradict themselves — for example a navigator whose web-presence field reports a live 200 while its own verification flag reads false and its DNS bind is still pending — and we record the contradiction rather than pick the flattering side. [SVR-C4-m518-003]

We also test our own claims adversarially. An internal verification bridge (m620) runs negative controls against the ecosystem's cancer-evidence claims and caught three of three deliberately injected errors; its rule is that every verdict must embed the real negative-control output rather than a description of it [SVR-C2-m620-001][SVR-C2-m620-002]. That tool is honest about its own limits, too: it has produced verdicts on seven real cases but has zero downstream consumers and has been idle since May

[SVR-C2-m620-003].

This report has not been independently audited. The word "verified" here means machine-traceable to our own frozen evidence, and nothing stronger.

3. Cancer & Clinical Evidence

This chapter reports on three cancer-related projects within the LIGHT HOPE ecosystem: m511 (ClearCancer Circle), m554 (TruthAnchor expert verification), and m610 (Roll Cancer Dynamics). Where possible, we separate what is publicly reachable from what is only planned, and we state zero-values and known regressions directly.

3.1 m511 — ClearCancer Circle

The m511 public site at ccc.lighthope.ai returned HTTP 200 when probed on 2026-06-24. That probe was an HTTPS HEAD request only; it confirms the surface responds but measures neither content nor readership. m511 records a self-assessed production readiness of 78 percent, a self-reported figure that is below 100 percent and not independently confirmed. [SVR-C3-m511-001][SVR-C3-m511-002]

Two limitations bound the m511 display and outreach. The site currently caps displayed experts at 9,000 (EXPERTS_CAP=9000) as a temporary workaround for the Cloudflare Pages 20,000-file limit; the permanent pagination or lazy-load fix is unimplemented and the issue remains open. The newsletter backend holds one subscriber, the owner's own address, meaning zero non-team subscribers to date. A campaign branch (campaign/ccc-oncology-2026-05) remains unmerged pending a post-event merge.

[SVR-C3-m511-003][SVR-C3-m511-004][SVR-C3-m511-005] The site did see real inbound during the reporting period — 28 genuine form applications between 2026-02-21 and 2026-03-09 — after which intake went cold. [SVR-C3-m511-007]

A stand-down statement page has been live at ccc.lighthope.ai/statements/2026-05-13 since 2026-05-13; reply and engagement to that statement are unmeasured. [SVR-C3-m511-006]

3.2 m554 — TruthAnchor Expert Verification

The m554 static expert atlas served at ccc.lighthope.ai returned HTTP 200 (77,731 bytes) when measured on 2026-07-28. Publication is therefore live, but consumption is not measured. The published sitemap lists 7,365 expert profile pages, exceeding the current live database of 1,155 experts, because the published corpus was generated from an earlier v2 pack and its scale differs from the present database.

[SVR-C3-m554-001][SVR-C3-m554-002]

We report m554's readership plainly as zero: its Google Analytics tag is the unwired placeholder G-XXXXXXXXXX, so no traffic instrumentation exists. The API is bound to 127.0.0.1 only and shows no external consumer traffic in its logs — healthy but private-only. The knowledge-production pipeline has produced no new experts or evidence since 2026-05-10, which contradicts a stale snapshot claim that a batch worker runs continuously. [SVR-C3-m554-003][SVR-C3-m554-004][SVR-C3-m554-005]

The central regression is honest and documented: a 2026-04-21 pgdata volume loss wiped the roughly 11,000-expert v1.0.0 state, leaving 1,155 experts in the live database, while the VERSION file still reads 1.0.0

and overstates the current corpus. The batch pipeline currently covers 10 of 35 target oncology domains, with 25 remaining. [SVR-C3-m554-006][SVR-C3-m554-007]

3.3 m610 — Roll Cancer Dynamics

m610 is planned only. Its domain rcd.lighthope.ai returned HTTP 000 (no response) when probed on 2026-06-24, and port 8610 is reserved but not enabled, so no running service occupies it. m610 nonetheless records a self-assessed readiness of 98 percent, a self-reported figure that conflicts with the absence of any live surface. [SVR-C3-m610-001][SVR-C3-m610-002][SVR-C3-m610-003]

The case-harvesting workflow is not yet exercised: directories 03_cases_raw and 04_cases_cooked contain 0 files, and the core scripts scan_inbound.py and cancer_dynamics_audit.py are stubs. The m610-to-m554 NWUP envelope is not wired, its jsonschema and requests dependencies are uninstalled, and the downstream contract sign-off with m554, m511 and m590 is unconfirmed.

[SVR-C3-m610-004][SVR-C3-m610-005][SVR-C3-m610-006]

This chapter is thinner than intended because two further candidate cancer projects could not be admitted this period for lack of a current snapshot entry.

4. Health Evidence Navigation

This chapter covers six health-oriented evidence-navigation projects. All are research-use evidence navigation only; they are not a medical device and do not diagnose individuals. Test counts and readiness figures below are internal, self-run results unless stated otherwise, and several are contradicted by other fields in the same snapshots. We report those contradictions rather than resolve them.

The shared architecture is a safety filter placed before the language model, so adversarial or crisis inputs are intercepted before any generation step. The m320 factory produces 24 domain instances (m321–m344) from a single config-driven codebase, with app.py byte-identical across instances [SVR-C4-m320-001]. Across those instances the suite reports 341/341 adversarial fixtures intercepted before the LLM; this is an internal test battery described as a starter, expandable set, not external validation [SVR-C4-m320-002]. The factory itself has no public hostname and its web presence status is internal with an empty domains list [SVR-C4-m320-003], and per-domain public deploy is a three-step, owner-gated procedure that is not auto-run [SVR-C4-m320-004]. The tracked issues list is empty, which is not evidence that no defects exist [SVR-C4-m320-005].

m321, the Low Back Pain navigator, hardcodes crisis hotlines 988 (US) and 12356 (CN) plus a cauda-equina red-flag emergency redirect [SVR-C4-m321-002], and its safety battery reports 29/29 adversarial fixtures intercepted before the LLM, self-run, with fixture counts drifting from 24 to 32 across the snapshot [SVR-C4-m321-004]. The snapshot self-contradicts on live status: backpain.lighthope.ai returned HTTP 200 when probed on 2026-06-24 [SVR-C4-m321-001], while public deploy is elsewhere recorded as PENDING OWNER and not auto-run [SVR-C4-m321-005].

m518, the Cortical Target navigator, reports its G11 battery intercepting 200/200 adversarial fixtures across seven attack categories before the LLM, though the status field elsewhere cites 236/236, an internal drift [SVR-C4-m518-001]. Its web presence shows cortex.lighthope.ai at HTTP 200 (probed 2026-06-24) while live_verified is false and the domain status is RESERVED with DNS bind pending — a self-contradiction we

do not reconcile [SVR-C4-m518-003]. Public cutover is blocked on the owner supplying a Cloudflare API token and binding DNS [SVR-C4-m518-004]. A documented catalog drift records 116 cortical targets advertised in a pack README against 128 in the actual YAML extract [SVR-C4-m518-008].

m517 (NeuroDynamics) reports neuro.lighthope.ai at HTTP/2 200, verified 2026-05-25 and probed again 200 on 2026-06-24 [SVR-C4-m517-001]. Its self-assessed production readiness is 95 percent, self-reported, with the remaining 5 percent being meeting execution and sign-offs [SVR-C4-m517-006].

m551, the autism science platform, publishes 24 expert profiles under dynamic routes, though those identities are static and not cross-checked against m554 TruthAnchor [SVR-C4-m551-002]. It applies an A/B/C/D four-level evidence-label system across 15 Level-A resources, hand-annotated and not wired to the m590 vault [SVR-C4-m551-005]. Production readiness is recorded at 83 percent, unchanged since 2026-04-20 [SVR-C4-m551-004].

m504, the MA Family Health navigator, records production readiness of 42 percent, unchanged since 2026-04-12 [SVR-C4-m504-001], with scaffold-level content: the /providers endpoint returns 2 hardcoded providers [SVR-C4-m504-003] and the medical dictionary is a scaffold of 4 terms [SVR-C4-m504-004]. It has no public hostname and web presence status is none, so external HTTP function cannot be probed [SVR-C4-m504-002].

A longevity-evidence navigator, LongProof (m516), joins this chapter this period. Its research site is publicly reachable [SVR-C4-m516-001], and its discipline is built around what it refuses to say: a dedicated wall enumerates 28 English and 14 Chinese forbidden marketing phrases — anti-aging and disease-reversal claims the site must not make — and the doctrine prohibits individual prognosis and individual-lifespan prediction [SVR-C4-m516-002][SVR-C4-m516-003]. It publishes 84 claim cards across 15 hallmarks of aging, though its own metadata inconsistently lists 52 cards across 12 hallmarks — a discrepancy we record rather than resolve [SVR-C4-m516-004]. Expert sidecars carry nine industry conflict-of-interest declarations, self-reported and unverified against a registry [SVR-C4-m516-005]. The doctrine forbids strong verified-intervention claims and requires falsifier-ready framing [SVR-C4-m516-006]. A freeze audit of its own materials found the pitch deck had overstated Chinese affiliation as 30 percent against a recorded 20 percent — a self-caught overstatement, still open [SVR-C4-m516-007].

5. Household Financial Wellbeing

helping families navigate systems they were never given the vocabulary for is the stated purpose of the four projects grouped under this chapter. All four remain pre-public or auth-gated as of this reporting period, and the disclosures below describe internal build states and self-reported figures rather than externally observed outcomes.

5.1 LongProof Finance Atlas (m631)

The central design decision in m631 is what its snapshot calls "compliance by structure." The knowledge graph physically contains no score, rank, or best_product nodes or edges — recorded as the G-1 invariant and exercised by the P9 hard red-line probe. Because the graph has no vocabulary in which a ranking or a "best" designation can be represented, the system cannot structurally express a statement of the form "product X is best for client Y." [SVR-C5-m631-001]

This is an architectural constraint rather than a policy promise, and the distinction matters for how the control should be read. A policy is a rule applied on top of a system that remains capable of the prohibited behaviour; it can be bypassed, misconfigured, or forgotten. An absent node type cannot be populated, so judgment-type queries fail closed rather than returning a degraded or hedged answer. We note plainly that this guarantee currently holds inside an INTERNAL_DEMO edition running on synthetic data with zero real PII and a site-wide watermark, and the boundary is self-declared, not externally checked. [SVR-C5-m631-002]

Readiness is recorded at 96 percent, and the snapshot is candid about why it is not higher: the G1 attorney gate is unpassed, real tenant data is not connected, and DNS is not bound. Acceptance probes P1 through P11 are reported all green under a true-fire local Worker, but the same snapshot states that in dry mode P1–P5 are NOT_TESTED. The green result therefore reflects an internal acceptance run under specific conditions, not a continuously verified live state. [SVR-C5-m631-004]

Two internal inconsistencies should be disclosed rather than reconciled away. The DEEP facet states that finatlas.x1000.ai is deployed, while web_presence records no bound domain and an empty domains list — a direct self-contradiction in the project's own records. [SVR-C5-m631-005] Separately, the project reports 100 percent build readiness alongside an honest commercial-delivery rate of 25–30 percent; build completeness is not equivalent to delivery readiness, and the gap between the two figures is the more useful number for a reader assessing whether families can actually use the system today. [SVR-C5-m631-006]

5.2 Multi-Entity Tax (m548)

The classification engine holds 68 rules across 20 categories, with long-tail transactions still requiring manual handling. [SVR-C5-m548-002] The site tax.x1000.ai returned HTTP 307 then 200 on the 2026-06-24 probe, but the application is auth-gated and redirects to /login, so an unauthenticated probe cannot confirm that the application functions behind the gate. [SVR-C5-m548-001]

The one outcome-level claim — a 90 percent or greater reduction in manual transaction-classification time — is marked unauditably, because no measurement data, usage logs, or real users appear in the snapshot to support the figure. We present it as a design target, not a measured result. [SVR-C5-m548-007] We also record honestly that commit bd4e047 on 2026-07-07 untracked a committed .env secret leak: a secret had previously been committed in plaintext before being untracked. [SVR-C5-m548-006]

5.3 Bilingual Mortgage (m506)

Production readiness is recorded at 92 percent, a value unchanged since 2026-04-21 across drift refreshes. [SVR-C5-m506-001] The project has no public domain: web_presence records status none and empty domains, so no external public probe exists despite the container being reported healthy. [SVR-C5-m506-002] Eight digital-human videos were produced and seven are described as published on-site, but all youtubeId values are null and zero videos are published to any external platform. [SVR-C5-m506-004]

The most material integrity event in this project is corrective. A B6 integrity audit retracted false "real-time rates" and lending overclaims, meaning prior public-facing claims were inaccurate and had to be removed. [SVR-C5-m506-005]

5.4 Finance & Markets Review (m632)

The FMR compliance Wall prohibits advisory recommendations, return guarantees, and price-prediction-as-fact — the editorial posture the project summarises as investment intelligence, not investment advice. The Wall is a documented editorial gate rather than a physically enforced graph constraint, and should be read as weaker than m631's structural approach. [SVR-C5-m632-001]

The project is not deployed: the domain finreview.x1000.ai is unbound and the HTTP probe returned 000, with publication owner-gated. [SVR-C5-m632-002] Work to date is a staging pilot at stage S0-BIRTH with a five-gate pass; go-live remains owner-gated. [SVR-C5-m632-005] The pilot issue FMR-2026-06 narrowed 27 raw citations to 20 verified, with one citation unresolvable and the issue still in staging. [SVR-C5-m632-003]

6. Education

This chapter reports on five education initiatives. All of them are early-stage, and most have no real external users. We lead with the initiative intended as this segment's centerpiece.

6.1 EduMap curriculum intelligence (m552)

EduMap's data plane is not functioning. On 2026-07-27, edumap.x1000.ai returned HTTP 500 on its root and on all /api/v1/* routes; the process is up, but the core product is unusable and this is tracked as a critical open issue [SVR-C6-m552-001]. The /health route returns HTTP 200, but it is a static handler and is not evidence that the application works, because every database-backed route is failing [SVR-C6-m552-005]. Reflecting this deployed-but-not-firing state, readiness was cut from 90 to 55 percent on 2026-07-27, with the prior figure regarded as an overstatement [SVR-C6-m552-004]. EduMap has zero real external users: all 14 entries in the access log originate from testclient, so the first-real-user sensor has fired on no real visitor [SVR-C6-m552-002]. On content, 418 courses are authored on disk, but the wired loader serves only about 116, because a 302-course expansion is not in the load path [SVR-C6-m552-003]. The snapshot lists one CRITICAL and two HIGH open issues, counted from the internal deep-refresh audit [SVR-C6-m552-006].

6.2 RollMath free SAT tutoring (m465)

RollMath is offered at zero cost to the learner. It has no public domain and no public hostname; its web-presence status is recorded as none [SVR-C6-m465-002]. The application ships with debug=True hardcoded in app.py, which is flagged internally as a security blocker ahead of any release [SVR-C6-m465-003]. Its live test suite passes 2 of 2 tests covering health endpoints, and only those two live tests exist [SVR-C6-m465-001]. The content comprises 12 learning seeds spanning S1 to S12, though the SAT corpus remains example-level [SVR-C6-m465-004]. Production readiness has held at 60 percent since 2026-05-23 with no change across seven refreshes [SVR-C6-m465-005]. RollMath does not replace teachers, and no efficacy claim is made for it.

6.3 AI Education Review — AIER (m414)

AIER is served at aiedu.lighthope.ai, returning HTTP 200 on 2026-07-27 via a Cloudflare Worker, and its public pilot issue AIER-2026-06 also returned HTTP 200 on that date [SVR-C6-m414-001]. The planned canonical domain aiedu.x1000.ai does not resolve, yet the sitemap and homepage links still hardcode that dead host [SVR-C6-m414-003]. There is no audience instrumentation on disk, so AIER has zero readership and its

consumption has no audit path [SVR-C6-m414-004]. On the pilot issue, 26 raw references were processed down to 21 verified with 0 corrupted, and a compliance check passed with 0 violations with a human second pass confirming 19 of 19 references real, all self-attested [SVR-C6-m414-005]. As a point of editorial discipline, AIER excluded a retracted 2026-04 meta-analysis that had reported a favorable effect size of $g=0.867$ [SVR-C6-m414-007].

6.4 Global AI Education Book Series (m800)

The EduBooks core comprises 21 volumes and 170 chapters at roughly 128k Chinese lines, written but unpublished [SVR-C6-m800-001]. The 170-chapter English translation is measured as 100 percent complete with 0 stub files under 50 lines [SVR-C6-m800-002]. The RAG index holds 843 chapters and 33,575 chunks, but it serves localhost only [SVR-C6-m800-003]. The Flask container is healthy, with /health returning 200 (v1.3.0) on 2026-07-27, bound to localhost:8800 with no public hostname [SVR-C6-m800-004]. There is no ISBN and no publication [SVR-C6-m800-005]. No reader, download, or sales instrumentation exists, and real readers stand at 0 [SVR-C6-m800-006].

6.5 Rudo control plane (m8201)

Rudo K12 v0.1.3 is live at rudo.x1000.ai, returning HTTP 200 (probed 2026-06-24) via a private cloudflared tunnel rather than open public access [SVR-C6-m8201-001]. A priority-zero finding was recorded: an RLS isolation test using the lighthouse SUPERUSER+BYPASSRLS role failed 6 of 6 household-isolation cases on the development database [SVR-C6-m8201-002]. A code-layer fix (commit 186da86), using a NOSUPERUSER NOBYPASSRLS role, makes 6 of 6 isolation tests pass on dev postgres, but production deployment remains pending [SVR-C6-m8201-003]. The 28-redline compliance filter passed 15 of 15 corpus tests [SVR-C6-m8201-005]. Rudo has not reached its 30-household north-star baseline, with 2 of 5 hard indicators unmet and no user-activity data on disk [SVR-C6-m8201-004].

6.6 Scope note

The chapter's core project this period is Lighthouse Math (m510), a mathematics learning platform live at its public domain [SVR-C6-m510-001]. Its build is substantial — 58 data models and 178 route files on disk — but that depth stands in front of almost no usage [SVR-C6-m510-002]. The core learning engine has recorded zero real-learner sessions; every diagnostic and mastery record belongs to a seed or demo account [SVR-C6-m510-003]. A lead form captured 43 real contact emails, but they produced zero conversions and zero completed sessions, and the payment path was never switched to live, so there have been zero real purchases [SVR-C6-m510-004][SVR-C6-m510-005]. The project's stated purpose advertises an eight-week, 92-percent success rate for reaching an SAT score of 750 or above; we record that figure only as an unauditable claim, because there are no real learners, no control group, and no before-and-after baseline behind it [SVR-C6-m510-006]. An under-13 protection gate (COPPA plus row-level security) blocks cross-system use of user data pending an owner decision, and K-12 row-level-security enforcement is not yet in production [SVR-C6-m510-007].

7. Community & Local Access

This chapter reports on four community-facing services. We describe what is currently live, what has been used, and where we do not yet have evidence to support a claim. Statements are self-reported unless noted, and the underlying checks were internally reviewed by m562, an internal audit project, as evidence-consistency checks.

7.1 JobBridge (m647)

JobBridge is the clearest case in this chapter of a service that runs but has not yet been used, and we report it in full rather than around it. The public site returned HTTP 200 healthy on 2026-07-20 (version 0.1.0, running on a Cloudflare Worker with D1) [SVR-C7-m647-001]. Over the ten days from launch to that probe, applications received stood at zero [SVR-C7-m647-002]. The three audit events recorded on the service all originate from deployment and seeding; none is user-generated, so we cannot distinguish any real end-user activity from operational events [SVR-C7-m647-003].

We state plainly that we do not yet know why the application count is zero. It may reflect insufficient outreach, a mismatch between the listings and what local employers actually need, or a defect in the product itself. We have not yet run the analysis that would let us separate these explanations, and we decline to assert one over the others in advance of that work [SVR-C7-m647-002].

The listing inventory carries a similar caveat. Of 14 total listings, 12 are published, but only 2 are first-party employer postings; the remaining 10 are aggregated from a seed file that is not tracked in the repository [SVR-C7-m647-004]. The service has 1 real external client — a verified adult-day-health center with a driver posting; two further community-organization tenants remain inactive, awaiting logos, and there is no disk evidence that the client hired anyone through the platform [SVR-C7-m647-005]. Finally, the R1 go-live gate file is named with a `.DONE` suffix while its own header still reads DRAFT — an unresolved internal contradiction we note here rather than resolve by assertion [SVR-C7-m647-007].

7.2 CovenantBridge (m648), public tier

We report CovenantBridge only at its public tier. On 2026-07-21 the public endpoint at usa.lighthope.ai returned status ok with layer=public via a Cloudflare Worker [SVR-C7-m648-001], and the public routes returned HTTP 200, including the /zh/ page at 33,329 bytes [SVR-C7-m648-002]. The private operational tier is not exposed publicly: those routes return HTTP 404 at the public edge, all four checked on the same date [SVR-C7-m648-003]. The deployed /zh/ page was md5 byte-identical to the local build on 2026-07-21, so what is served matches what was built [SVR-C7-m648-004].

An internally-run audit of the content caught a fabricated citation and reduced mislabeled DISPUTED entries from 57 to 3; m556 did not re-verify the audit's own conclusions [SVR-C7-m648-006].

7.3 CalmFaith (m524)

CalmFaith is live at faith.lighthope.ai, went live on 2026-06-02, and returned HTTP 200 when probed on 2026-06-24 [SVR-C7-m524-001]. Its internally-run launch gate report recorded 13 of 13 gates passing, with zero proselytizing content, zero medical claims, and zero unanchored quotes; this is a self-audited internal report [SVR-C7-m524-003]. We note two limits on completeness: only 3 of 9 traditions are officially anchored (Catholic universal, Islam scoped, Judaism scoped), while the remaining 6 stand as editorial commentary without an official anchor [SVR-C7-m524-005]. The repository has been idle since 2026-03-19 on a

single-commit master, meaning no recent content iteration [SVR-C7-m524-006].

7.4 BostonCulture (m388)

BostonCulture is live, with boston.lighthope.ai returning HTTP 200 when probed on 2026-07-21 [SVR-C7-m388-001]. Beyond availability, however, a project forensic review found no real consumption. Real bookings were zero: all 58 orders share a single created_at timestamp matching hardcoded stress-test data [SVR-C7-m388-003]. Real memberships were zero (10 records all at example.com) and real payments were zero dollars (2 payments, both from a test address, 0.3 seconds apart) [SVR-C7-m388-004]. Of 15,909 hit-counter events, 15,795 are healthcheck machine probes, leaving attributable human visits near zero [SVR-C7-m388-005]. A previously owner-attested figure of 7 to 9 real visitors was retracted by the project as unverified and not for external citation, and no verifiable source for a visitor count now remains [SVR-C7-m388-006].

8. Protecting People

The projects in this report touch people who are, in various ways, exposed: patients searching for evidence, families without the language of a system, anxious people seeking steadiness, and minors. This chapter collects the protective mechanisms that are actually implemented, and marks where protection is designed but not yet independently checked.

Minors are separated at the data layer, not by a rule. The K-12 control plane's first isolation test used a database role with superuser and row-level-security bypass privileges and failed all six household-isolation cases — a result the project recorded rather than buried. A code-layer fix using a restricted role now passes all six isolation cases on the development database, with the production cutover still pending. [SVR-C6-m8201-002][SVR-C6-m8201-003]

Content aimed at minors passes a compliance filter, and an inflated safety number was corrected downward. The K-12 platform's redline filter passes fifteen of fifteen corpus tests, and an internal adversarial review found a self-reported 5.3 percent answer-leak rate to be overstated and revised it to zero with a stated boundary. The correction went against the project's own earlier, more flattering number. [SVR-C6-m8201-005][SVR-C6-m8201-006]

The crisis path is hard-wired. In the low-back-pain navigator, crisis hotline numbers for the United States and China and a cauda-equina emergency red-flag redirect are hard-coded and cannot be switched off through configuration. [SVR-C4-m321-002]

Harmful inputs are caught before the model. In the cortical-target navigator, the safety battery records interception of 200 of 200 adversarial fixtures across seven attack categories before the language model is invoked. This is a self-reported internal result and is not external validation. [SVR-C4-m518-001]

Vulnerable-audience discipline is visible in the evidence handling. The autism science platform applies a four-level evidence-labeling scheme across its Level-A resources, a deliberate constraint for an audience that is a frequent target of unsupported claims. [SVR-C4-m551-005]

Private data is walled off at the edge. In the community covenant platform, the public tier reports its layer as public, while the private operational tier returns HTTP 404 at the public edge for all of its routes, so private-tier

content is not reachable from the public surface. A public case page separately recorded a named-entity personal-information scrub with zero remaining hits. This report describes only the public tier; the private tier is on the exclusion list and its contents are not described here. [SVR-C7-m648-003][SVR-C7-m648-005]

Protection here is real where it is coded and honest where it is unfinished. None of these mechanisms has been reviewed by an outside party this period.

9. What We Did Not Achieve

This chapter is the reason the rest of the report can be trusted. It lists, without cushioning, what did not work, what we could not measure, and what we left unfinished. There is no closing reassurance in this chapter by design.

Nobody used the things we shipped. The hiring platform, m647, was live and healthy and had received zero applications in its first ten days. We do not know whether that zero reflects insufficient outreach, a mismatch with employer needs, or a product problem, and we have not run the analysis that would tell us. Its recorded activity events were all deployment and seeding events, not real user actions.

[SVR-C7-m647-002][SVR-C7-m647-003]

The pattern repeats across the portfolio. The curriculum-intelligence platform recorded zero real external users, with all fourteen access-log entries coming from a test client, while its database-backed routes returned HTTP 500. [SVR-C6-m552-002] The expert atlas had readership at zero and its analytics tag was an unwired placeholder. [SVR-C3-m554-003] The completed book series, twenty-one volumes and a finished English translation, had zero real readers and no publication. [SVR-C6-m800-006] The AI-education review had zero readership, traffic, or analytics evidence at all. [SVR-C6-m414-004] The Boston culture platform's forensic review found real bookings at zero across fifty-eight seed orders, real memberships at zero, and real payments at zero dollars, with its hit counter almost entirely machine probes. [SVR-C7-m388-003][SVR-C7-m388-004]

Several claims cannot be audited, and we say so. A set of statements in this report carry the UNAUDITABLE grade because no evidence path exists for them. The tax platform's estimate of a 90-percent-or-greater reduction in manual classification time has no usage data behind it. [SVR-C5-m548-007] The book series' reader count has no reader, download, or sales instrumentation. [SVR-C6-m800-006] The K-12 platform has not reached its thirty-household baseline and collects no user-activity data on disk to measure it. [SVR-C6-m8201-004] A previously owner-attested visitor figure for the culture platform was retracted by the project itself as unverified and unfit for external citation. [SVR-C7-m388-006] These are unauditables, and each is marked with the reason it cannot be checked.

We left process debts open past their dates. An ecosystem metadata migration is overdue and was not completed within the reporting period; an unmerged campaign branch remains outstanding pending a post-event merge. [SVR-C3-m511-005] Several of the live-reachability probes cited in earlier chapters predate the freeze window and should be re-run before any of these claims is repeated in a later report.

This report received zero external review, and zero internal review by m562. m562 is an internal audit project inside the LIGHT HOPE ecosystem; it is not an outside firm. Its six institutional axes are frozen until late 2026, so no internal review by m562 was requested or performed for this period. The internal review coverage of the claims in this report is therefore zero. Everything here is self-reported and internally

cross-checked, and where cross-checking was not possible, the claim was labeled UNAUDITABLE rather than dropped silently.

We are reporting these gaps because a report that hid them would be worth less than no report at all.

10. Commitments for the 2026 Annual Report

Each commitment below is tied to a specific state recorded in this report, so that the next edition can be checked against it rather than against a mood. We state them as verifiable conditions, not intentions.

Turn zeros into real numbers or explain why they stayed zero. The hiring platform's zero applications is the clearest test. By the annual report we commit to deploying usage instrumentation on it and reporting either a real, non-team action count or a documented account of why the count is still zero. [SVR-C7-m647-002]

Restore the platform that is down. The curriculum-intelligence platform's database-backed routes returned HTTP 500 with zero real users. We commit to reporting either a working data plane with a dated probe or its retirement. [SVR-C6-m552-002]

Land the isolation fix in production. The K-12 control plane's row-level-security fix passes all six isolation cases on the development database only. We commit to reporting the production cutover result, pass or fail. [SVR-C6-m8201-003]

Re-probe every live claim inside the freeze window. Several reachability probes in this edition predate the freeze. We commit to dated, in-window probes for every LIVE claim in the next report, or to downgrading the claim. [SVR-C7-m388-002]

Widen genuine coverage rather than restate it. The faith platform anchors only three of nine traditions to official sources. We commit to reporting the anchored count as it stands, up or down, not to asserting completeness. [SVR-C7-m524-005]

Keep the outcome line honest. We do not commit to producing OUTCOME · VERIFIED claims by the annual report; manufacturing attribution to fill that line would defeat the purpose of this report. We commit instead to reporting any genuine before-and-after measurement we can defend, and to leaving the line at zero if we cannot. [SVR-C5-m548-007]

The measure of the 2026 annual report will not be how many of these turned green. It will be whether we reported the reds as plainly as this edition reports its own.

Appendix. Claim Reference Index

Every claim identifier used in this report:

SVR-C2-m620-001
SVR-C2-m620-002
SVR-C2-m620-003
SVR-C3-m511-001
SVR-C3-m511-002
SVR-C3-m511-003
SVR-C3-m511-004
SVR-C3-m511-005

SVR-C3-m511-006
SVR-C3-m511-007
SVR-C3-m554-001
SVR-C3-m554-002
SVR-C3-m554-003
SVR-C3-m554-004
SVR-C3-m554-005
SVR-C3-m554-006
SVR-C3-m554-007
SVR-C3-m610-001
SVR-C3-m610-002
SVR-C3-m610-003
SVR-C3-m610-004
SVR-C3-m610-005
SVR-C3-m610-006
SVR-C4-m320-001
SVR-C4-m320-002
SVR-C4-m320-003
SVR-C4-m320-004
SVR-C4-m320-005
SVR-C4-m320-006
SVR-C4-m321-001
SVR-C4-m321-002
SVR-C4-m321-003
SVR-C4-m321-004
SVR-C4-m321-005
SVR-C4-m321-006
SVR-C4-m504-001
SVR-C4-m504-002
SVR-C4-m504-003
SVR-C4-m504-004
SVR-C4-m504-005
SVR-C4-m504-006
SVR-C4-m516-001
SVR-C4-m516-002
SVR-C4-m516-003
SVR-C4-m516-004
SVR-C4-m516-005
SVR-C4-m516-006
SVR-C4-m516-007
SVR-C4-m517-001
SVR-C4-m517-002
SVR-C4-m517-003
SVR-C4-m517-004
SVR-C4-m517-005
SVR-C4-m517-006
SVR-C4-m517-007
SVR-C4-m518-001
SVR-C4-m518-002
SVR-C4-m518-003
SVR-C4-m518-004
SVR-C4-m518-005
SVR-C4-m518-006
SVR-C4-m518-007
SVR-C4-m518-008
SVR-C4-m551-001
SVR-C4-m551-002
SVR-C4-m551-003
SVR-C4-m551-004
SVR-C4-m551-005
SVR-C4-m551-006
SVR-C5-m506-001
SVR-C5-m506-002
SVR-C5-m506-003
SVR-C5-m506-004

SVR-C5-m506-005
SVR-C5-m506-006
SVR-C5-m548-001
SVR-C5-m548-002
SVR-C5-m548-003
SVR-C5-m548-004
SVR-C5-m548-005
SVR-C5-m548-006
SVR-C5-m548-007
SVR-C5-m631-001
SVR-C5-m631-002
SVR-C5-m631-003
SVR-C5-m631-004
SVR-C5-m631-005
SVR-C5-m631-006
SVR-C5-m632-001
SVR-C5-m632-002
SVR-C5-m632-003
SVR-C5-m632-004
SVR-C5-m632-005
SVR-C5-m632-006
SVR-C6-m414-001
SVR-C6-m414-002
SVR-C6-m414-003
SVR-C6-m414-004
SVR-C6-m414-005
SVR-C6-m414-006
SVR-C6-m414-007
SVR-C6-m465-001
SVR-C6-m465-002
SVR-C6-m465-003
SVR-C6-m465-004
SVR-C6-m465-005
SVR-C6-m510-001
SVR-C6-m510-002
SVR-C6-m510-003
SVR-C6-m510-004
SVR-C6-m510-005
SVR-C6-m510-006
SVR-C6-m510-007
SVR-C6-m552-001
SVR-C6-m552-002
SVR-C6-m552-003
SVR-C6-m552-004
SVR-C6-m552-005
SVR-C6-m552-006
SVR-C6-m800-001
SVR-C6-m800-002
SVR-C6-m800-003
SVR-C6-m800-004
SVR-C6-m800-005
SVR-C6-m800-006
SVR-C6-m8201-001
SVR-C6-m8201-002
SVR-C6-m8201-003
SVR-C6-m8201-004
SVR-C6-m8201-005
SVR-C6-m8201-006
SVR-C7-m388-001
SVR-C7-m388-002
SVR-C7-m388-003
SVR-C7-m388-004
SVR-C7-m388-005
SVR-C7-m388-006
SVR-C7-m388-007

SVR-C7-m388-008
SVR-C7-m524-001
SVR-C7-m524-002
SVR-C7-m524-003
SVR-C7-m524-004
SVR-C7-m524-005
SVR-C7-m524-006
SVR-C7-m647-001
SVR-C7-m647-002
SVR-C7-m647-003
SVR-C7-m647-004
SVR-C7-m647-005
SVR-C7-m647-006
SVR-C7-m647-007
SVR-C7-m648-001
SVR-C7-m648-002
SVR-C7-m648-003
SVR-C7-m648-004
SVR-C7-m648-005
SVR-C7-m648-006