

OPEN SKY LEADERSHIP INSTITUTE

OpenSkyX Leadership Intelligence

Investor & Builder Brief

The organization that remembers.

Dr. Robert Hausmann · July 2026 · Confidential working draft

www.openskyx.org/intelligence · rhausmann@openskyx.org

1. Executive Summary

Organizations spend enormous sums learning about themselves — leadership programs, consulting engagements, offsites, interviews, surveys, simulations — and retain almost none of what those engagements reveal. The evidence evaporates when the room clears. Six months later they pay to rediscover what they already knew.

OpenSkyX Leadership Intelligence is a private, governed intelligence layer that fixes this. At its center is an ontology — a defined set of organizational objects (Leader, Team, Tension, Decision, Experiment, Outcome, Learning) and the relationships between them — that lets a coaching note, a simulation log, a network map, and a quarterly outcome speak the same language. Around it sits an evidence pipeline with provenance and human validation built in: AI proposes, evidence explains, humans decide.

The deliberate analogy is Palantir: ontology-first infrastructure applied to a domain that has never had it — how organizations actually learn. The difference is origin. This platform is being built from inside a working leadership practice, instrumented against live client programs and a university doctoral course, by the practitioner who runs them.

The foundation is running today: ontology v0.8, an evidence-packet intake standard, a tools registry, a local Apple-silicon AI runtime, a governance and review model, and two live pilot environments. This brief describes what exists, what is being built next, and the three kinds of partners we are inviting in: design partners, builders, and patient capital.

2. The Problem: Organizational Amnesia

Global spend on leadership development is commonly estimated in the tens of billions of dollars annually (frequently cited at \$60B+, with corporate L&D spend several times that). Nearly all of it produces a one-time experience and a binder.

The structural failure is memory, not quality. A workshop surfaces a real tension. An interview names what everyone has been circling for months. A simulation reveals a team's decision logic under pressure. Then the engagement ends and the evidence scatters into slide decks, notebooks, and hallway recollection. The organization cannot ask: What tensions keep repeating? Which assumptions have we actually tested? Where is trust blocking execution? What do we keep saying but not changing?

Every interaction creates evidence. Evidence may strengthen, challenge, or extend the organization's shared intelligence — but only if it is captured, connected, and governed.

Consultancies have weak incentives to solve this; re-diagnosis is revenue. Software vendors ship survey dashboards, not connected memory. And enterprises rightly refuse to paste confidential interviews and strategy documents into public AI tools. The gap is a trusted, private, evidence-first layer — and it is wide open.

3. The Platform

3.1 The ontology

A versioned model (currently v0.8) of shared objects and relationships: Leader, Team, Stakeholder, Goal, Tension, Assumption, Decision, Signal, Action, Experiment, Outcome, Learning — plus instruments, data sources, and lifecycle statuses. Client-facing users never need the word 'ontology'; it is what makes every other capability possible.

3.2 Evidence packets

Everything enters through one standard wrapper carrying source, context, privacy classification, consent and allowed uses, content, ontology hints, and processing status. Existing tools are not rebuilt; their outputs are translated by adapters. Claims are explicitly typed as Extracted (present in the source), Inferred (with confidence noted), or Missing (asked back to a human, never silently assumed).

3.3 The intelligence cycle

Interaction → Evidence → Interpretation → Ontology proposal → Human validation → Organizational intelligence → Decision or intervention → Outcome → New evidence. The cycle is recursive by design: outcomes become evidence for the next round of sensemaking, and the ontology changes only when human review confirms the evidence supports it.

3.4 Governance and trust architecture

- **AI proposes, humans decide.** AI-generated entities, relationships, and interpretations default to Proposed and become Accepted only after human review.
- **Provenance is mandatory.** Every insight points back to the packets, tools, dates, participant scope, privacy level, and review status behind it.

- **Confidentiality travels with the evidence.** Consent, allowed uses, excluded uses, and retention rules are attributes of each packet, per program.
- **Local or client-controlled processing.** Confidential material is processed on local infrastructure (Apple-silicon MLX runtime) or inside approved client systems — never public AI tools.
- **No hostage model.** The client keeps its data, its learning map, and its ability to move the system.

3.5 The instruments

The OpenSky Toolkit already produces structured evidence: Interview Lab, Network Mapping, Organization Simulator, Safe-to-Fail Experiments, Waypoints, the Learning Observatory, ABCD, Both/And Analysis, leadership portfolios, and more — roughly thirty registered instruments. Field Notes, the founder's own reflection practice, is the first fully instrumented intake channel: one capture becomes a published essay, a manuscript fragment, learning intelligence, and an ontology proposal. We run the system on our own work first.

4. What Exists Today

Component	Status	Notes
OpenSkyX Ontology	v0.8, running	Entities, relationships, instruments, data sources, lifecycle governance in versioned YAML
Evidence Packet standard	Defined	Source, context, privacy, consent, ontology hints, processing status
Tools registry	Running	~30 instruments registered with ontology types, privacy class, review requirements
Local AI runtime	Operational	MLX-LM on Apple silicon; no public port, no external AI dependency for confidential data
Ontology Review Console	Prototype	Data model and workflow for human validation of AI proposals
Corporate pilot	In motion	Emerging-leaders program with a longstanding Fortune 500 client; next module instrumented for August 2026
University pilot	In motion	Doctoral leadership course (University of Houston) as high-volume, lower-risk iteration

Component	Status	Notes
		environment
Marketing & distribution	Live	openskyx.org with dedicated Intelligence section; institute, academy, toolkit, publishing pipeline

This is a foundation, not a finished product. There is no production database, no multi-user access model, and the review console is not yet a web UI. We say this plainly because the honesty is the point of the architecture.

5. Why Now

- **Local AI crossed the threshold.** Capable models now run on commodity hardware, making private, on-premise evidence analysis feasible for the first time — the enabling condition for enterprise trust.
- **The evidence exhaust is already digital.** Transcripts, survey exports, workshop artifacts, simulation logs, and meeting recordings exist in every organization; nothing connects them.
- **The trust gap is the moat.** Public-model prohibitions in enterprises are hardening, not softening. A provenance-first, human-validated, local-capable architecture is structurally hard for incumbents to retrofit.
- **Practice-led development is rare.** Most attempts at this space come from software companies without leadership practice, or consultancies without engineering. This platform grows from a working practice with clients, a professorship, a toolkit, and a publishing pipeline already in place.

6. Business Model: Services-Led, Platform-Built

The model is deliberately both/and: services revenue funds and de-risks platform development, and every engagement makes the platform smarter.

Layer	Offer	Revenue character
1. Instrumented programs	Leadership programs, retreats, and engagements delivered with the intelligence layer underneath — facilitator briefings, cohort synthesis, evidence-backed design	Premium services; exists today

Layer	Offer	Revenue character
2. Intelligence subscriptions	Ongoing access to the client's own governed intelligence layer: query interface, review console, synthesis reports	Recurring; pilot stage
3. Platform & IP	The ontology, adapters, review tooling, and methodology as licensable infrastructure; books, teaching cases, and academic research as reach	Long-term; in development

The wedge is narrow and honest: make the client's next program module measurably smarter than the last one. Compounding does the rest.

7. Roadmap: The Next Four Quarters

- **Program workspaces.** Client/program separation by context, consent, privacy level, and retention (corporate pilot and doctoral course as first two workspaces).
- **Source adapters.** Turn existing workshop tools, reflections, interviews, simulations, network maps, and course artifacts into evidence packets.
- **Human review queues.** Proposed themes, tensions, experiments, and recommendations approved before becoming organizational intelligence.
- **Cohort synthesis outputs.** What we are hearing, what it means, what might be tried next, what requires care — with the evidence named.
- **Program integration surfaces.** A plain-language query interface, a facilitator console, and optional client-site embeds.
- **Proof artifacts.** Published before/after case studies from both pilots; the first evidence-derived manuscript chapters and teaching cases.

8. What We Know, What We Assume, What We Don't Know

In the spirit of the platform itself, this brief separates its own claims.

Category	Claims
Facts	The foundation components in Section 4 exist and run. Two pilot environments are live. The founder's practice, client portfolio,

Category	Claims
	professorship, toolkit, and publishing pipeline exist and generate engagements today.
Assumptions	Clients will pay a premium for programs whose learning compounds. Enterprise trust requirements favor local/private processing. Facilitator-plus-system is a durable advantage over software-only entrants. Estimated market figures are industry-standard citations, not audited.
Unknowns	Time-to-value for organizations without an existing OpenSky engagement. The right pricing for intelligence subscriptions. How quickly adapters generalize across client toolstacks. Regulatory treatment of employee-derived evidence in some jurisdictions.

9. The Invitation

We are looking for three kinds of partners, in this order of immediacy:

- **Design partners** — organizations with a live leadership program, transformation, or cohort who want the learning to compound. You get an instrumented program and early roadmap influence; we get the constraint-rich environment that makes the platform real.
- **Builders** — engineers, researchers, and designers who care about ontology design, local-first AI, evidence provenance, or learning science. The architecture is early enough that your fingerprints will be on it.
- **Investors and strategic partners** — patient capital and aligned platforms that see the same structural gap. The next step is a working session: we will walk the running system, the pilots, and the numbers behind the practice.

Contact: Dr. Robert Hausmann · rhausmann@openskyx.org ·
www.openskyx.org/intelligence

10. About the Founder

Dr. Robert Hausmann is a professor in the University of Houston EdD program in Professional Leadership and an international management and leadership consultant with twenty-five years of practice across corporate, higher-education, and non-profit environments. He is the founder of the Open Sky Leadership Institute — place-based leadership development built on the Open Sky Paradox Framework™ — the builder of the OpenSky AI toolkit, and the author of two

manuscripts in development (The Space Between and Between the Beats) that are themselves being assembled through the platform's own field-note instrument.

He is building the platform the way he advises clients to build: as a series of safe-to-fail experiments, run inside real work, with the evidence kept.

This brief describes a platform in active development. It is an invitation to a conversation, not an offer of securities. Client references anonymized where engagements are not public. © 2026 Open Sky Leadership Institute.