Five cases. One underlying system.

HSDLG makes disciplined inquiry reusable across AI-assisted work. These studies show how the same foundation can organize evidence, expose defects and preserve unresolved decisions across domains.

What the evidence supports

The portfolio combines a bounded software repair, selected local mechanics, public-data analysis, illustrative engineering calculations and prospective research design. Each case states its own execution and validation boundary.

What the evidence does not establish

The studies are not five deployed businesses or controlled measurements of HSDLG efficacy. No universal quality improvement, cost saving, model ranking or production readiness follows from their existence.

How to read the cases

Each page identifies the problem, contribution, concrete finding, execution route, limitation and next customer proof. The final page records producing-model provenance and broader prototype history.

The business direction

A reusable Epistemic OS with Bring Your Own Intelligence (BYOI), APIs for future applications, and a proposed initial customer offer in security and compliance evidence review. Founder-led work is the current delivery route; full MCP integration remains in progress.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.

AI-Native Compliance Infrastructure

Evidence must cover the actual decision

Problem and contribution

How can a reviewer determine whether policies, contracts and assessments support a particular vendor relationship and exposure? The study used a public vendor-oversight case to develop a paper, nine application standards with 45 requirements and proposed acceptance cases, and a linked expert-review workspace.

Concrete finding or repair

Review corrected statements about absent policies, contracts and assessments. The stronger question was whether the available evidence covered the relevant provider, service and decision. The work preserved uncertainty about unavailable company records.

Execution evidence

The retained R1 evidence brief reports eleven assertions in an isolated authority exercise and one reference-worker dispatch. It also reports acceptance of an acyclic plan and rejection of a cyclic successor. A separate native review preflight held without executing its payload.

Boundary

This is research and selected local mechanics. It does not establish a deployed compliance service, a legal determination, independent expert endorsement or prevention of the historical incident.

Next customer proof

Proposed first offer: a fixed-scope vendor-security evidence review, with customer-approved inputs and accountable reviewer dispositions. Compare accepted findings, missed issues and total review effort with the existing process.

Source / scope: Making AI-Native Compliance Infrastructure Real and its Research and Review Companion, as retained in the reviewed R1 evidence brief. Original case framing distinguishes allegations from settlement requirements.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.

Self-Maintaining APIs

A correct-looking result concealed a lifetime defect

Problem and contribution

Could an authored reporting application preserve required behavior across a known pandas / SQLAlchemy transition? The study exercised a real historical library incompatibility against authored behavioral requirements, then applied review and repair beyond checking whether rows were returned.

Concrete finding or repair

The retained study reports 11/12 selected primary checks on the original stack, 1/12 after the provider upgrade, 11/12 after an ordinary pandas upgrade, and 12/12 for the repaired application on both selected stacks. The old-stack rollback defect was preexisting.

Execution evidence

An initial repair passed 3/4 additional challenges. Review found that a connection context ended before iteration finished. The revised candidate retained the connection and passed 4/4. Existing local controls completed nine bounded stages and avoided redispatch on completed-stage replay. An earlier unresolved reservation remained unresolved.

Boundary

One known transition on SQLite with exposed tests and same-assistant reasoning/review. These counts are selected checks, not independent migrations or a general repair success rate. The assistant supplied semantic reasoning and code.

Next customer proof

A supervised integration-maintenance pilot: representative held-out changes, engineer-reviewed patches, acceptance evidence and measured total effort. Customer pull-request acceptance and unattended operation remain unestablished.

Source / scope: Making Self Maintaining APIs Real, final HSDLG_V2 package, as retained in the reviewed R1 evidence brief. pandas 1.5.3; SQLAlchemy 1.4.46 and 2.0.0.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.

New Operating Systems for the Physical World

Useful analysis can include a negative result

Problem and contribution

How can municipal maintenance connect asset evidence, resource constraints and responsible operating decisions? The work combined a fleet snapshot analysis, scheduling research, nine application standards, a proposed pilot and linked expert-review materials.

Concrete finding or repair

In the retained October 30 NYC fleet snapshot, 11 of 24 critical groups were below reported availability targets. Gross group shortfall was 378 vehicles and the net gap after other surpluses was 177. Two collection groups represented 66.9% of the gross gap. Review separated gross from net and corrected independence claims.

Execution evidence

Thirteen synthetic authority-control checks passed in the retained study. All 135 operational acceptance scenarios remained NOT_RUN. The mathematical adaptation showed no incremental final scheduling benefit beyond the conventional optimizer.

Boundary

Availability comparisons do not establish missed services or repair causes. No live equipment dispatch or production grant occurred. No measured service improvement or independent field validation is claimed.

Next customer proof

A bounded shadow pilot with an operating partner: validate asset data and constraints, obtain SME judgments, compare plans, and measure service outcomes before admitting operational actions.

Source / scope: Making Physical-World Operating Systems Real and Methods, Review and Human Decisions, as retained in the R1 evidence brief. NYC Daily Fleet Report snapshot dated October 30, 2026.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.

Compute at Sea

Cooling savings do not settle the investment case

Problem and contribution

Under what site and workload conditions would floating compute outperform credible land alternatives? The work compared floating water-cooled, land-air-cooled and land-water-cooled designs. It screened Casco Bay and Kempenfelt Bay, challenged assumptions, and developed a plan for obtaining decisive site evidence.

Concrete finding or repair

In one explicitly illustrative 1 MW IT scenario, annual utility savings of 158,400 currency units were smaller than a marine premium of approximately 312,745, leaving a 154,345 deficit. The land-water comparator also challenged whether offshore placement was necessary. These are assumed inputs, not quoted project economics.

Execution evidence

The R2 monograph reports bounded local mathematical-helper execution on declared or synthetic inputs. Document-guided analysis separately examined thermal margins, accepted useful work, resource limits and chronology. No full HSDLG runtime or native enforcement is claimed.

Boundary

Neither study area is an investment recommendation. Actual power, fiber, rights, marine cost and operating performance remain unresolved. Same-assistant analysis and review do not establish independent engineering assurance.

Next customer proof

Acquire actual site, utility and supplier information and domain SME input. Move from broad screening to deeper, site-specific engineering and economics. The next commitment should purchase decisive information.

Source / scope: Compute_at_Sea_HSDLG_V2_Monograph_R2 and Compute_at_Sea_Casco_Kempenfelt_Feasibility_R1. Feasibility arithmetic independently recalculated during this package review; no field verification.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.

Proving Human / Data for the Real World

A testable research program with scoped trust claims

Problem and contribution

How can physical observations and digital authenticity signals support decisions without confusing provenance with truth or credential possession with identity? The package developed a research monograph, custom standards and an operationalization paper. Physical data and authenticity share an evidence lifecycle while retaining different objects of inference and domain requirements.

Concrete finding or repair

The BPL record identifies a confounded sensing comparison: data density could vary alongside retraining, hardware, thresholds or workflow. The revised primary comparison fixes the downstream model, training data, thresholds and workflow, with broader adaptation assessed separately.

Execution evidence

The record reports two same-assistant, document-guided review cycles. It narrows predictability and control claims and separates candidate system roles from deployed capability. No native HSDLG runtime or empirical pilot was executed for this paper.

Boundary

This is a prospective research program, not a functioning proof-of-human product or demonstrated fraud reduction. Sensor provenance does not prove measurement accuracy. A key-possession signal does not by itself prove identity or intent.

Next customer proof

Choose one use case and accountable owner. Specify lawful data access, reference measurements or threat scenarios, protected evaluation, decision costs, privacy and accessibility outcomes, and conditions that would falsify the benefit.

Source / scope: Evidence-Governed Research Monograph; Making Data for the Real World Real; companion BPL Review Record. The fifth case spans two Fall 2026 prompts.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.

Model provenance and portability

Producing configurations

AI-Native Compliance Infrastructure: GPT-6 Sol, Medium effort.
Self-Maintaining APIs: GPT-6 Sol, Medium effort.
New Operating Systems for the Physical World: GPT-6 Astra, Medium effort.
Compute at Sea: GPT-6.1 Terra, Medium effort.
Proving Human / Data for the Real World: GPT-6 Luna, Medium effort.

Broader prototype experience

The founder reports prototype use on Claude and Gemini, with more limited testing on Grok and local AIs such as Gemma. Exact versions, task coverage and comparative outcomes for these additional prototypes are not established in this package.

What BYOI means

Customers would bring their chosen intelligence to a shared foundation for evidence, review, correction and continuity. The intended MCP interface and application APIs support this architecture direction. Client-specific compatibility and end-to-end behavior still require qualification.

Evidence needed for stronger claims

Use the same tasks, source access and resource accounting across selected configurations, with and without the workflow. Independently adjudicate material defects, false findings, retained constraints, accepted outputs and total effort. Different producing models alone do not establish independence or comparative performance.

Source / scope: Case model identities and effort settings supplied by Michael L. Dell on November 2, 2026. Broad cross-provider experience is founder-reported prototype history.

Additional technical documentation, standards, test records and implementation details are available for scoped diligence under an appropriate NDA or contract.