HSDLG

An Epistemic OS for AI

Disciplined inquiry, reusable across AI-assisted work.
Vector To Gold

BYOI: Bring Your Own Intelligence

Learning has recurring patterns

Across fields, useful inquiry gathers evidence, forms explanations, challenges assumptions and preserves what survives review.
HSDLG turns that observation into a reusable system for human-AI work.

Design thesis; broad effectiveness remains to be measured.

Decisions need a defensible basis

A convincing answer can still omit a constraint, misread an assessment or hide an unresolved assumption.

Teams need to see what supports a conclusion, what changed and who must decide.

Initial customer focus: security and compliance evidence review.

Bring Your Own Intelligence

Customers choose their AI. HSDLG is designed to supply shared evidence, review, correction and continuity functions.

APIs provide the intended connection points for future applications.

Full MCP integration is in progress; universal client compatibility is unproven.

A Multitude of Applications

Compliance: evidence scope and reviewer decisions
API maintenance: bounded software repair
Physical operations: asset and scheduling analysis
Compute at Sea: feasibility and competing alternatives
Proving Human / real-world data: testable research design

Research and local results; five deployed products are not claimed.

A repair that review changed

Trigger: a SQLAlchemy upgrade broke an authored reporting app.
Defect: a chunk iterator outlived its connection context.
Correction: yield chunks while the connection remains open.
Verification: the lifetime check changed from FAIL to PASS.
All four additional challenges passed after correction.

One known transition, exposed tests and same-assistant review.
Retained result excerpts: yc.vectortogold.com/Evidence_API_Repair.html

Cooling savings versus marine cost

Illustrative annual utility savings: 158,400 CU
Illustrative annual marine premium: 312,745 CU
The comparison exposed a 154,345 CU deficit and a credible land-water alternative.

Assumed 1 MW IT scenario, not site quotes. Actual site and SME inputs enable deeper analysis.

Concrete corrections across domains

Compliance: test whether evidence covers the relevant provider and exposure.
Physical operations: separate gross shortfall from net gap; retain no optimizer uplift.
Real-world data: isolate sensing changes from retraining and workflow changes.

Source records support these bounded findings, not universal causal uplift.

THE FOUNDER

Michael L. Dell

CISSP, ISSEP, ISSMP
25 years in DoD cybersecurity and risk management.
Creating and validating NIST-compliant RMF packages for formal approval.

Experience preparing evidence for accountable risk decisions; no HSDLG certification implied.

THE FIRST CUSTOMER OFFER

Security and compliance evidence review

Founder-led, fixed-scope engagements: review supplied evidence and prepare findings for accountable decisions.
Proposed buyer: security and risk teams responsible for vendor reviews.

Validate accepted findings, missed issues, total effort and repeat use. Pricing remains unvalidated.

A complete MCP delivery path

Founder-reported prototypes span Claude and Gemini, with limited Grok and local-AI use, including Gemma.
The main challenge is getting the full system functioning through MCP and qualifying selected clients.

BYOI and extensible applications are the architecture direction; end-to-end qualification remains open.

THE NEXT STAGE

A reusable foundation and a focused business

Establish paid, repeatable customer workflows.
Complete MCP integration with engineering support.
Qualify application interfaces and demonstrate retained value across repeated work.

Additional technical documentation, standards and test records available under an appropriate NDA or contract.

The results behind the example

pandas 1.5.3 / SQLAlchemy 1.4.46 and 2.0.0 / SQLite

Selected primary checks	Passed
Original application, old stack	11 / 12
Provider upgrade alone	1 / 12
First repair or ordinary pandas upgrade	11 / 12 each
Revised repair on each tested stack	12 / 12

Separate additional challenges: 3 / 4 before correction, 4 / 4 after.
C01 checked connection ownership during iteration, beyond returned rows.

Source: API evidence study, pages 1-4. One transition, authored checks, same-assistant review.
yc.vectortogold.com/Project_Studies/02_APIs_Evidence_Study.pdf