Evidence for decisions that matter.

WHAT HSDLG IS

The purpose

High-Stakes Decision-Level Governance is an Epistemic Operating System: an application and governance layer around AI models. It connects the task, evidence, review and retained result so people can inspect what supports a decision and what remains unresolved.

What it helps people do

Review reports, code and analysis against a defined purpose and supplied evidence. Correct supported defects within agreed authority. Carry source links, findings, decisions and open questions into later work.

The shared foundation

HSDLG brings together knowledge management and governance, model and theory formation, reusable representations, continuity, tool planning, work capture and selected mathematical checks.

Current delivery

Founder-led, scoped reviews and analytical work are the current engagement route. The codebase includes selected local implementations and test evidence. Reliable full-system MCP delivery and client qualification remain engineering milestones.

The ambition

Make disciplined inquiry reusable across substantive AI-assisted work, with applications building on shared functions and organizations choosing their intelligence providers.

Review with a clear purpose. Revise with a record.

HOW THE WORK PROCEEDS

1 / Establish the question

Agree the intended result, available sources, protected decisions, permission to revise, acceptance checks and stopping conditions.

2 / Plan a proportionate route

Tool Use Planning selects useful methods and dependencies. Unavailable evidence, access or implementation remains an explicit prerequisite gap.

3 / Backpass

BP reviews the specified work backward through its sources, assumptions, requirements and authority. The result distinguishes supported findings, evidence gaps and decisions needing accountable judgment.

4 / Backpass Loop

BPL adds authorized correction, validation and fresh review of the whole revised target. Work stops when complete, when evidence or authority is missing, when progress stalls, or at agreed limits.

5 / Retain the result

Preserve what changed, what survived review, the relevant evidence and unresolved questions. Later continuation requires current access and scope checks.

Human and automated roles

Models supply reasoning and generation. Deterministic checks address supported structures and calculations. Subject-matter experts and accountable owners retain consequential judgments. Same-assistant review is not independent assurance.

Shared discipline. Separate domain responsibilities.

ARCHITECTURE AND APPLICATIONS

Core

Eleven ownership areas cover knowledge acquisition and governance, theory formation, primitives and archetypes, continuity, tools, work capture and mathematical foundations.

Concepts and Extensions

The Concepts Registry supplies reviewable analytical standards and contracts. Standalone extensions implement selected methods for information value, uncertainty, claim relations, evaluation, routing and recovery within declared profiles.

TAIE / EMSM

Geological and mineral-system applications organize evidence and competing interpretations to identify useful further investigation. Source code and local tests do not establish field validity or production admission.

AAAM

A separate evidence-assessment application retrieves candidates and retains accountable reviewer determinations. Evidence retrieval does not establish control coverage, certification or an Authority to Operate.

ValidityBench

A local application candidate supports evaluating completed user work against selected criteria. Semantic-service binding and customer operational qualification remain open.

Bring Your Own Intelligence

APIs provide intended connection points for models and applications. Founder-reported prototype use spans several AI systems. This history does not establish universal model compatibility or comparative performance.

A repair that review changed.

A CONCRETE RESULT

The setting

An authored pandas reporting application encountered a real SQLAlchemy compatibility transition. The selected primary checks passed 11/12 on the original stack, 1/12 after the provider upgrade, and 12/12 for the revised repair on each tested stack. An ordinary pandas upgrade passed 11/12. The original rollback defect predates the upgrade.

The hidden defect

A chunk iterator returned from a connection context could still produce expected SQLite rows after the context ended. A separate review checked ownership during iteration and observed zero checked-out connections when one was required.

The correction and verification

The revision kept the generator inside the connection context during consumption. The lifetime challenge changed from FAIL to PASS. The four additional challenges changed from 3/4 passing to 4/4 passing. These are a separate test set from the twelve primary checks.

What the result earns

The retained records show a particular defect, a corresponding correction and selected passing checks. The study used one known transition, authored tests and same-assistant production and review. It did not isolate HSDLG’s causal contribution or establish a general repair success rate.

Further work and contact

Public evidence studies also cover compliance, physical operations, infrastructure feasibility and authenticity research. Fuller analyses, standards and selected test records are available for scoped diligence under an appropriate NDA or contract.

Michael L. Dell, CISSP, ISSEP, ISSMP
michael.dell@vectortogold.com
https://hsdlg.vectortogold.com