

AI OPERATIONS CASE STUDY

Building an evidence-governed research and production system

Independent Portfolio Project

A human-directed, AI-assisted project connecting research, claim control, workflow design, documentation, testing, and durable proof.

Evidence Level 4 — Tested

Prepared by Esse Obayando · AI Operations Specialist

esse.obayando@rarecapital.tech · [linkedin.com/in/esse-obayando-240686287](https://www.linkedin.com/in/esse-obayando-240686287)

Release Candidate v1.0.0-rc.1 · Not deployed

1. Project overview

The project tests a practical operating model for AI-assisted work: move from an ambiguous question to review-ready outputs without allowing unsupported evidence, invisible decisions, or automated approval to enter the process.

Problem

Generative systems can produce polished language before the evidence is ready. Research, drafting, production planning, review, and QA are also often fragmented, weakening traceability and recovery.

My role

- Defined requirements, research boundaries, workflow architecture, quality standards, and approval gates.
- Retained final human authority over evidence rules, professional positioning, and approval decisions.
- Directed bounded Codex support for research, implementation, generation, validation, testing, formatting, and rendering.

Independent portfolio project developed using a human-directed, AI-assisted workflow.

2. Method and workflow

The operating sequence separates questions, source collection, assessment, claims, conflict review, research lock, drafting, claim recheck, QA, and human approval. Rejected sources and unresolved claims remain visible instead of being silently discarded.




Figure 1. End-to-end project workflow. The approval gate is a control boundary, not a decorative step.

3. Human-directed AI operating model

Human responsibility and AI assistance are deliberately separated. AI can accelerate bounded implementation and structured generation; it does not set requirements, approve evidence, assign professional status, or authorize publication.




Figure 2. Human direction, bounded AI assistance, review, and preserved evidence.

Operational controls

- Approval status remains separate from claim eligibility.
- Version control and recovery files preserve exact handoffs.
- Automated tests verify structural rules; visual inspection checks rendered artifacts.
- Private review records remain outside public-facing exports.

4. Research and evidence

The locked research reviewed 24 candidate sources and evaluated 27 claims. 21 claims were cleared for drafting; rejected and unresolved claims remained excluded.




Figure 3. Research evidence narrows through explicit quality and eligibility decisions.

Evidence boundary

These are verified project-process metrics. They do not demonstrate client results, revenue impact, sales lift, conversion improvement, or external validation.

5. Outputs and capabilities demonstrated

- Source register, claims ledger, conflict log, research lock, and limitations record.
- Long-form and short-form content plans tied to approved claims.
- Decision, risk, task, approval, QA, change, status, and recovery records.
- Ten editable diagrams with accessible SVG, PNG, HTML, and PDF derivatives.
- Local review dashboard and responsive static portfolio with privacy boundaries.
- Application résumé, LinkedIn copy, interview kit, case study, and publication-readiness controls.

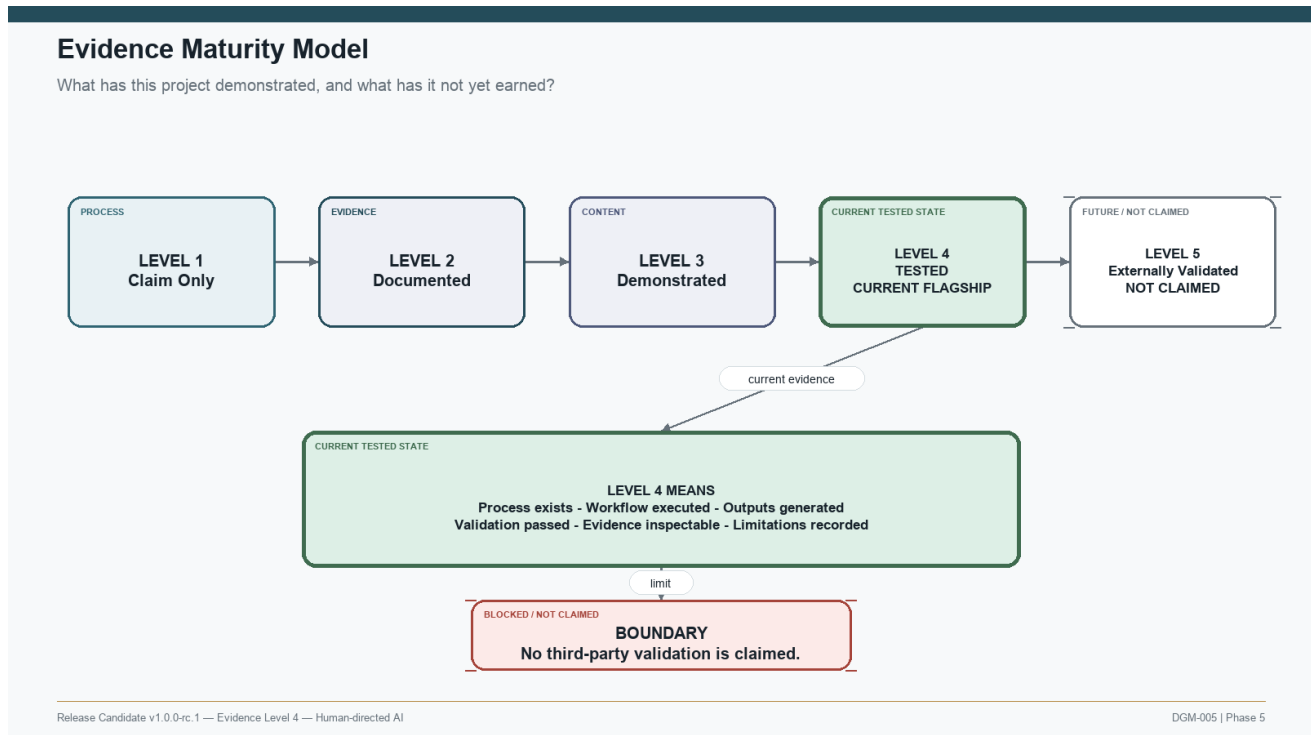


Figure 4. Evidence Level 4 represents executed and tested process evidence; Level 5 is not claimed.

6. Limitations and next proof

- The work is an independent portfolio project, not paid client employment.
- No commercial outcome, testimonial, team-management history, or external validation is claimed.
- The Phase 4 research is a preserved historical snapshot and requires freshness review before publication.
- Product features, vendor documentation, and legal guidance can change.
- Publication requires final privacy, rights, link, accessibility, security-header, and rollback verification.

What I would improve

The next evidence step would be qualified external review or controlled real-world use under a separately approved scope, followed by recorded feedback and a new evidence version. Publication itself would still require identity, privacy, rights, link, accessibility, and rollback checks.

Capabilities demonstrated

AI Operations · Human-in-the-Loop AI · AI-Assisted Workflow Design · Workflow Design · Research · Information Synthesis · Evidence Governance · Project Coordination · Quality Assurance · Documentation · Process Design · Content Operations · Knowledge Organization · Operational Planning · Prompt Development · Responsible AI

Contact

Esse Obayando · AI Operations Specialist

esse.obayando@rarecapital.tech · [linkedin.com/in/esse-obayando-240686287](https://www.linkedin.com/in/esse-obayando-240686287)

Location, GitHub, and portfolio URL intentionally omitted.