ECON 606 Study Bot

Personalized Learning in Economics with RAG-Based Generative AI

Powered by BoodleBox: Merging Multi-Agent Orchestration with FERPA-Compliant Academic Integrity

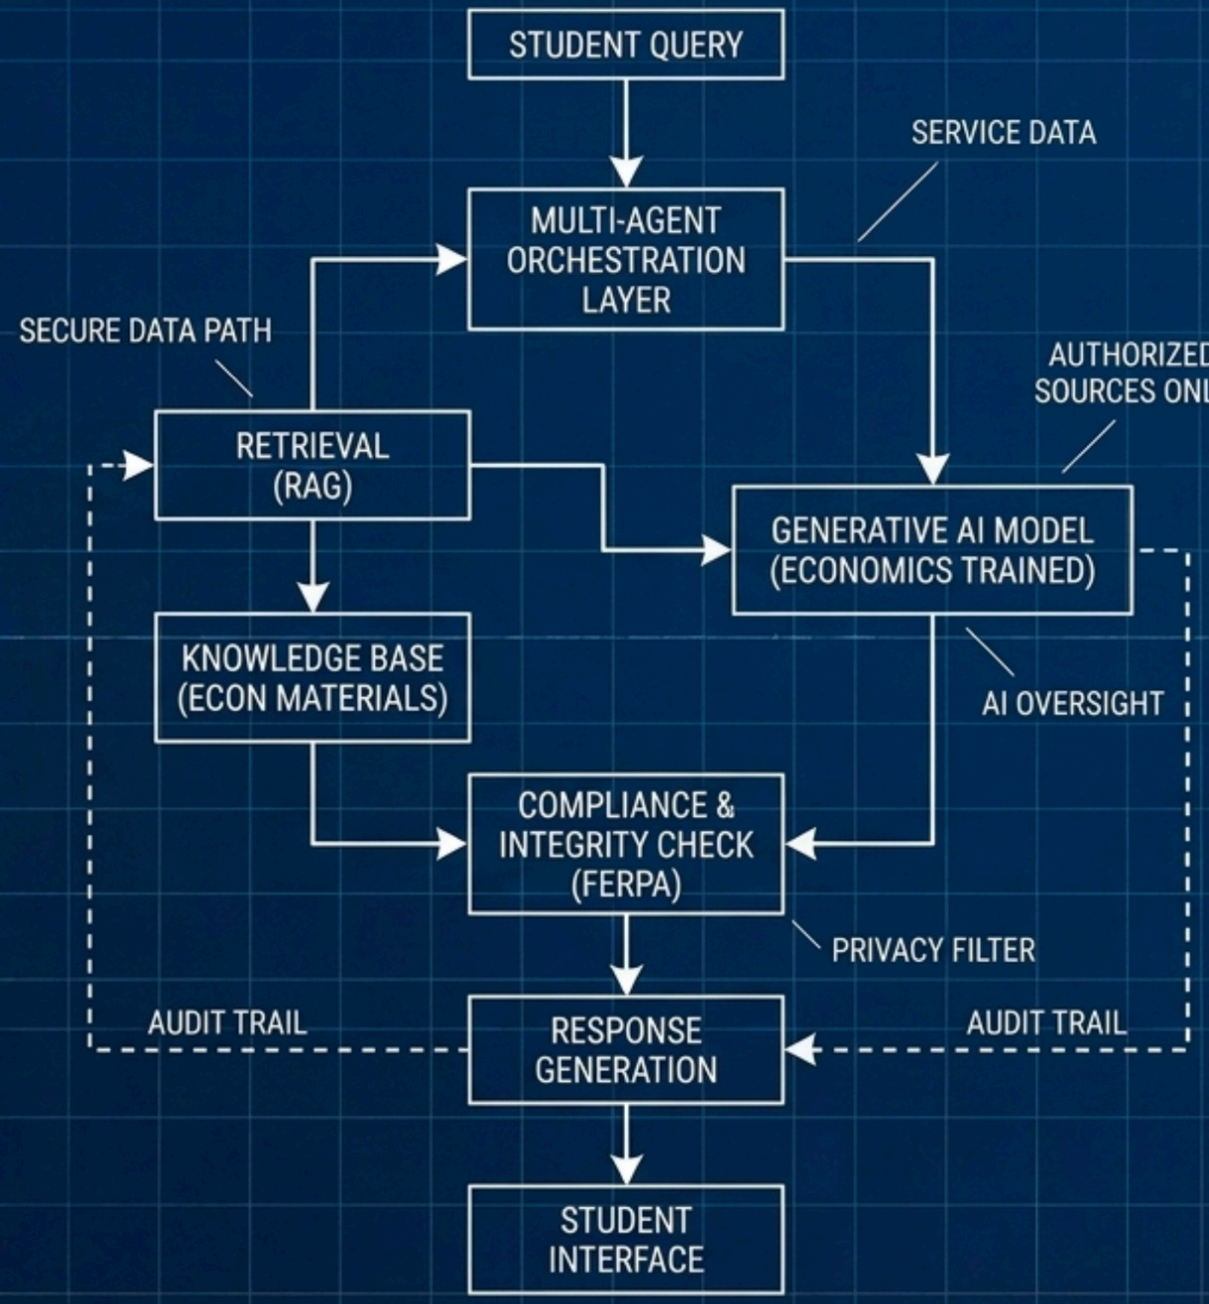

We Keep the Humanity in Technology

The Defining Question:

How might we cultivate human discernment and agency in an era where where AI makes it increasingly tempting to outsource our thinking?

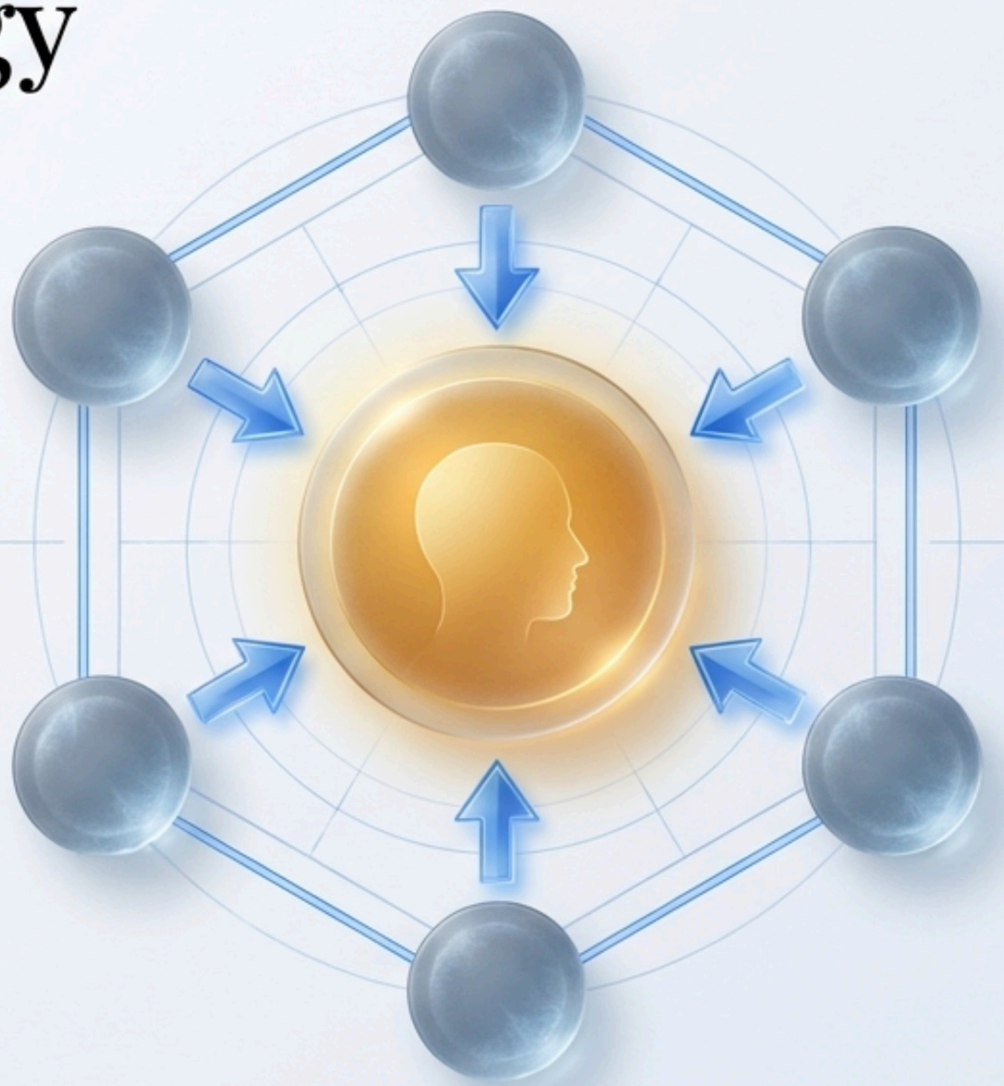

The BoodleBox Philosophy:

Collaboration Over Automation. Our platform actively resists replacing human cognition, deliberately transforming students from passive consumers into skilled AI practitioners.

The Educational AI Dilemma

Promises of General AI

Content Fidelity

Instant answers.

Quality Assurance

Automated tool creation.

Academic Integrity

Enhanced productivity.

Data Privacy

Personalized tutoring.

AI Literacy

Empowered students.

Realities of the Classroom

Content Fidelity

15-25% hallucination rates in quantitative domains; synthesis failures combining incorrect theories (Shuster et al., 2021).

Quality Assurance

Unvalidated code generation fails under the mathematical edge cases of managerial economics.

TRANSLATION FLOW

Academic Integrity

78% of faculty express concern over AI submissions, yet traditional detection methods fail and create adversarial surveillance cultures (Cotton et al., 2024).

Data Privacy

Consumer AI platforms ingest institutional intellectual property and violate student privacy boundaries.

AI Literacy

Unstructured access leads to over-reliance, reducing critical thinking and creating passive consumers.

The Diverging Paths of AI in Higher Education

Dimension	General Consumer AI	Enterprise Educational AI (BoodleBox)
Pedagogical Goal	Replace thought (Automation)	Enhance thought (Collaboration)
Knowledge Boundary	Hallucinates widely	Enforces a Zero-Tolerance Knowledge Gate
Academic Integrity	Relies on guessing-based surveillance	Utilizes a Privacy-Preserving Verification Ledger
Data Privacy	Uses student data for model training	Strictly FERPA/SOC 2 compliant with zero-training Enterprise APIs

A Fortress of Trust: Bridging Law, Privacy, and Ethics

FERPA & HIPAA Compliant

Strict adherence to student and patient privacy regulations.

SOC 2 Type II Certified

Enterprise-grade security protocols audited annually.

Enterprise API Contracts

Legal agreements explicitly prohibiting AI models from utilizing student data or queries for future training.

Ethical AI Grounding

Safe infrastructure allows models to safely process integrated global AI ethics guidelines and frameworks on Corporate responsibility in the digital economy.

Developing Metacognition Through Calibrated Interventions

Pattern Tracking

The system securely monitors usage to identify when a student might be defaulting to automation rather than active thinking.

Evidence-Based Intervention

When triggered, the Coach provides targeted prompts to redirect the student's analytical process.

The Outcome

Students develop vital metacognitive skills to leverage AI as an enhancement tool rather than a replacement for human thought.

BoodleBox Enhanced AI Coach Mode

Building AI Literacy in strict compliance with FERPA.

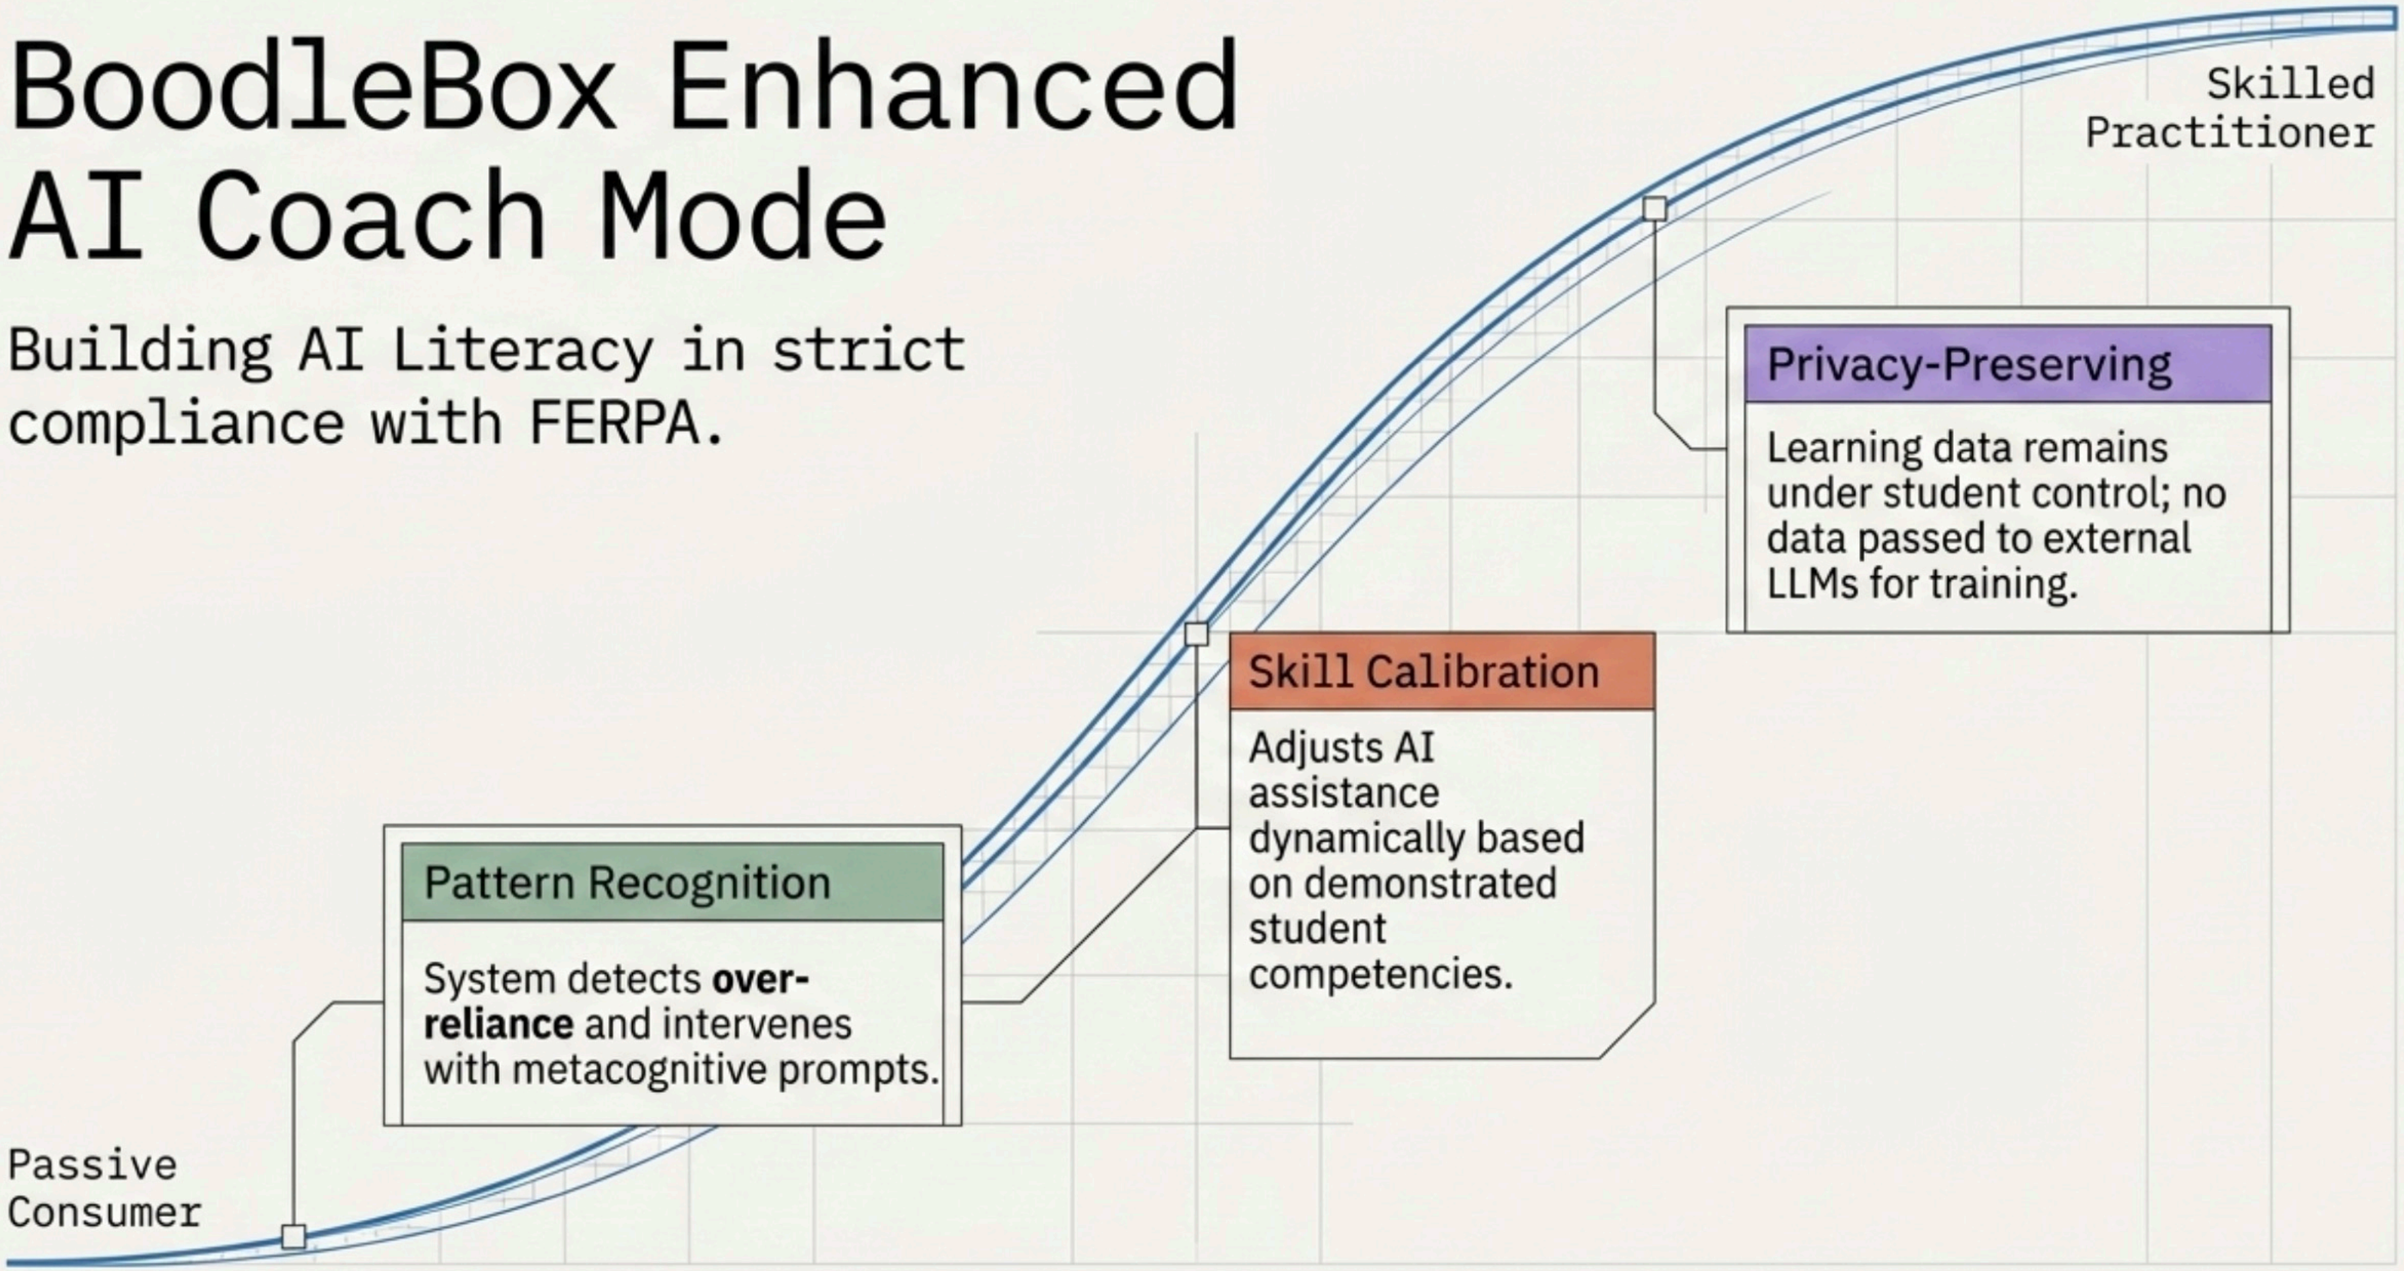

The ECON 606 Solution Paradigm

Pillar 1: Content Fidelity

Mechanism: Knowledge-gated Retrieval-Augmented Generation

Goal: Eliminate hallucinations hallucinations through exclusive domain grounding

Pillar 2: Quality Assurance

Mechanism: Validation-first code generation and adversarial stress testing

Goal: Guarantee mathematical precision in interactive tools

Pillar 3: Academic Integrity

Mechanism: Privacy-preserving cryptographic verification

Goal: Enable instructor forensics without student surveillance.

The 8-Phase Orchestration Pipeline

Tracking a complex query: ‘Generate a GDP Calculator’

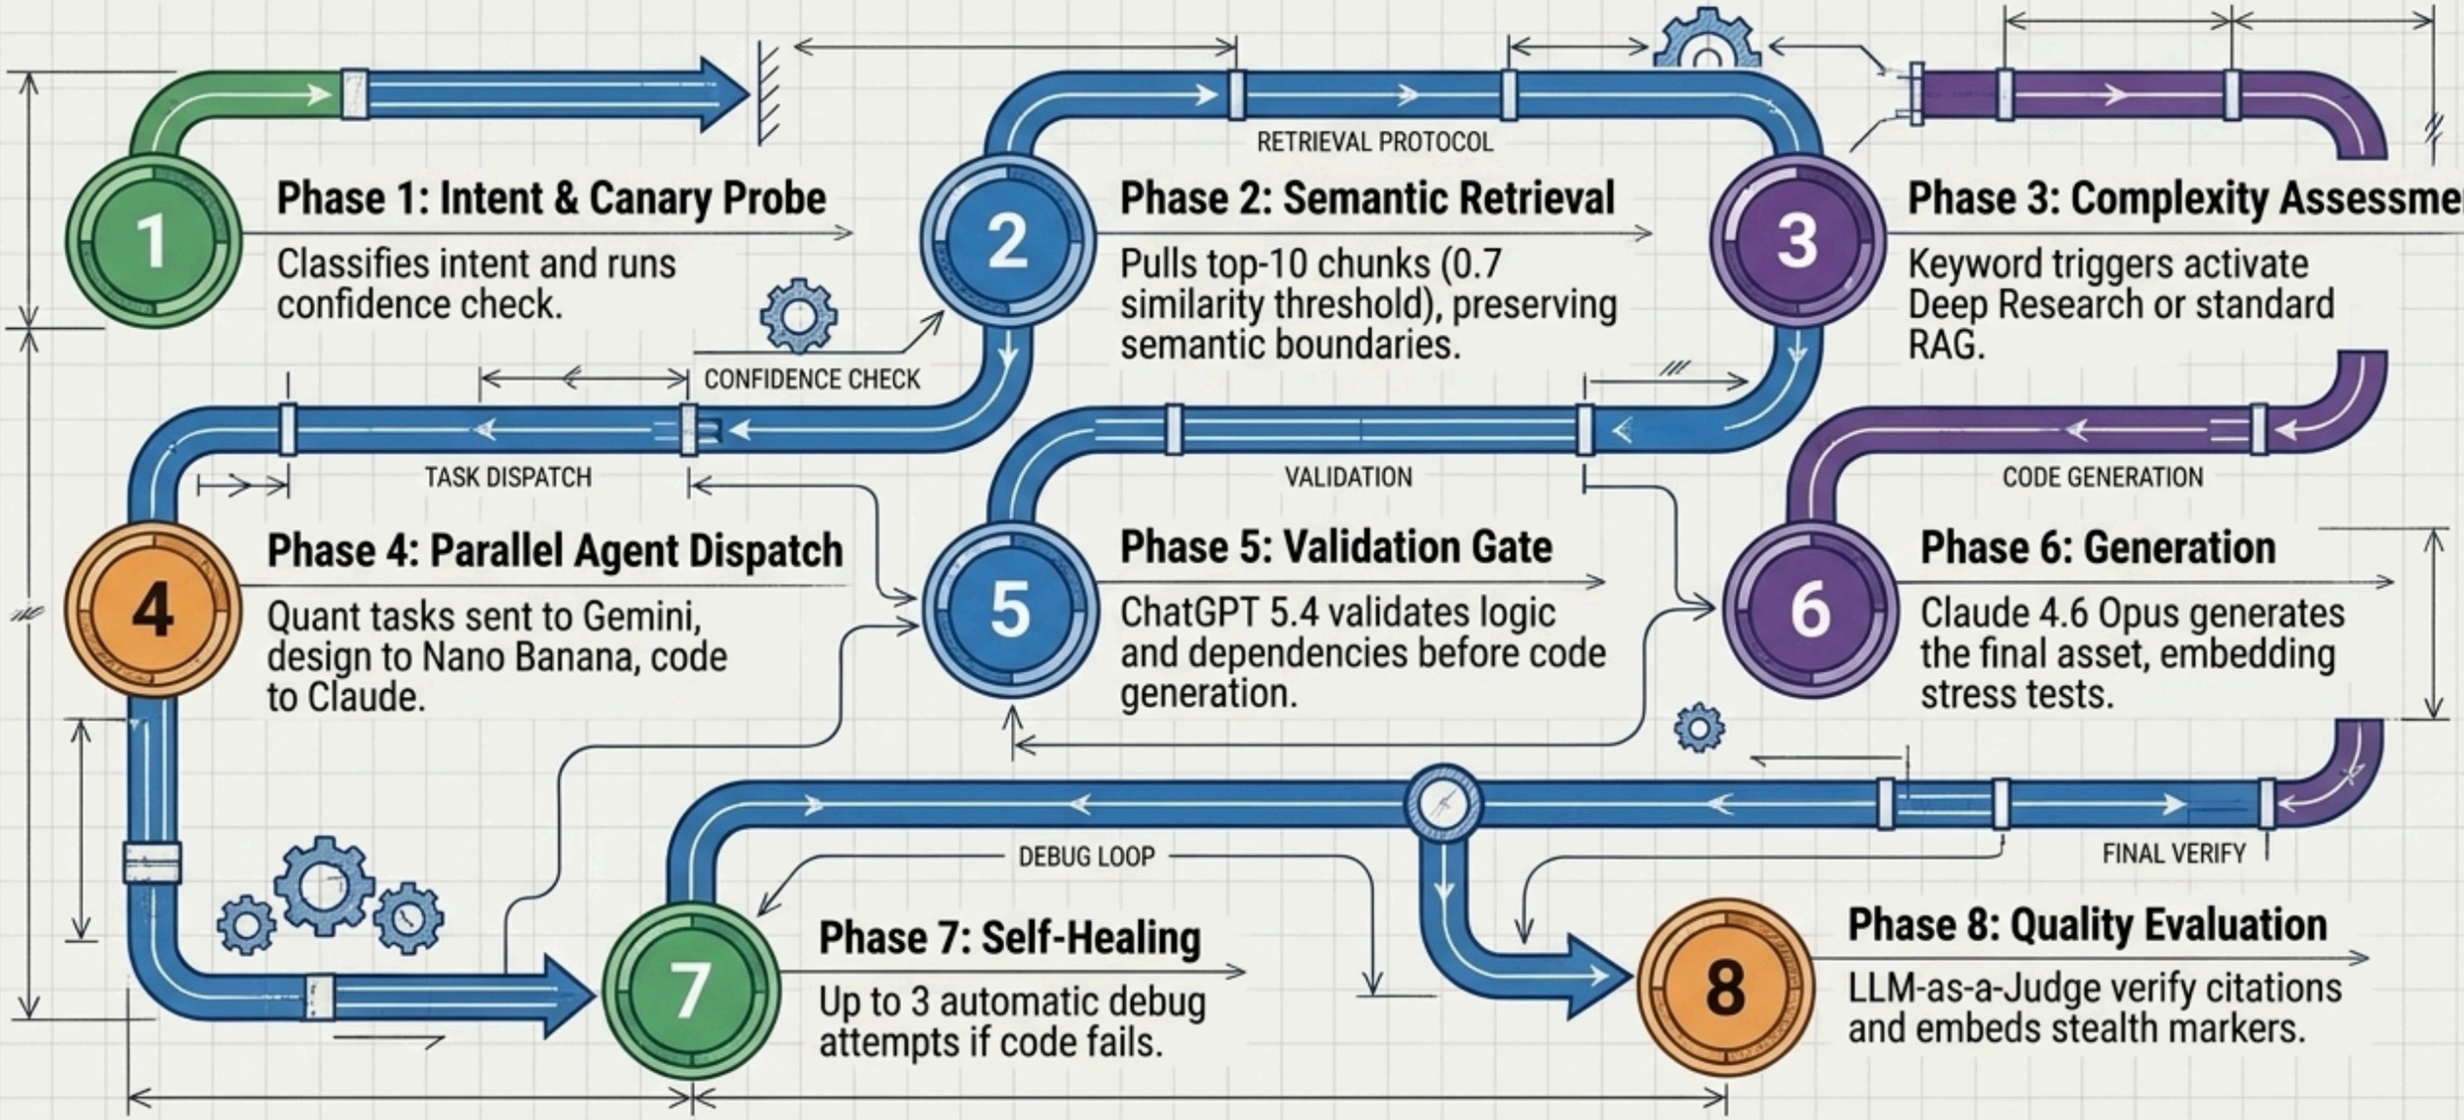

The Zero-Tolerance Knowledge Gate

The Zero-Tolerance Protocol:
ALL responses must cite specific page numbers from authorized sources (e.g., Shapiro, Perloff). Common knowledge exceptions and Wikipedia are strictly forbidden.

User Query

Step 1: Topic Extraction & Scan

Query parsed against semantic chunking index using text-embedding-3-large.

Step 2: Canary Knowledge Probe (Confidence Scoring)

Confidence >= 0.8

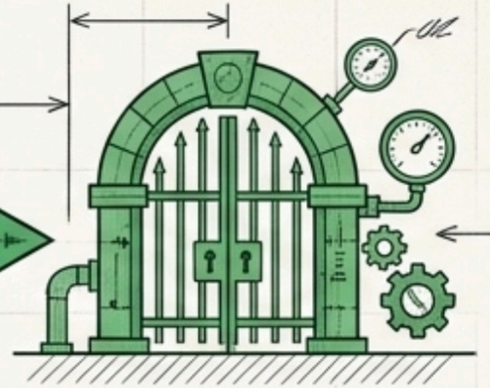

Authorized ECON 606 Corpus

Full generation proceeds using retrieved semantic chunks.

Confidence 0.5 - 0.8

Generation proceeds with explicit scope limitation.

Confidence 0.3 - 0.5

System halts generation; offers alternative related topics.

Confidence < 0.3

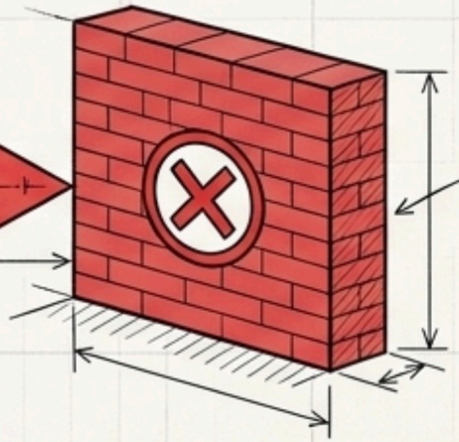

The General Knowledge Void

Hard Halt.
System explicitly states its knowledge boundary. Factual anchoring overrides creative guessing.

Mixture-of-Models (MoM) Task Routing

AI Agent Profile	Foundational Model	Primary Role & Specialization	Parameters
Master Orchestrator	Claude 4.6 Opus	Complex synthesis, code generation, curriculum mapping.	Temp: 0.3 | Context: 128K Output
Fast Responder	Claude 4.6 Sonnet	Quick concept explanations, daily Q&A.	Temp: 0.2 | Context: 100K
Quant Specialist	Gemini 3.1 Pro	Financial calculations, bond math, formula derivations.	Temp: 0.1 (Optimized for precision)
Material Generator	ChatGPT 5.4	DOCX/PPTX formatting, practice exams.	Temp: 0.2
Curriculum Researcher	Perplexity Pro/Sonar	Real-time fact-checking against course standards.	Temp: 0.2
Visual Designer	Nano Banana Pro	Financial flowcharts, decision trees.	Temp: 0.5 (Higher variance for creativity)

Agentic Deep Research Protocol

Triggered by Complexity: Activated by keywords like 'evaluate,' 'compare,' or 'synthesize.' Bypasses standard RAG limits to utilize Claude's massive 128K output window.

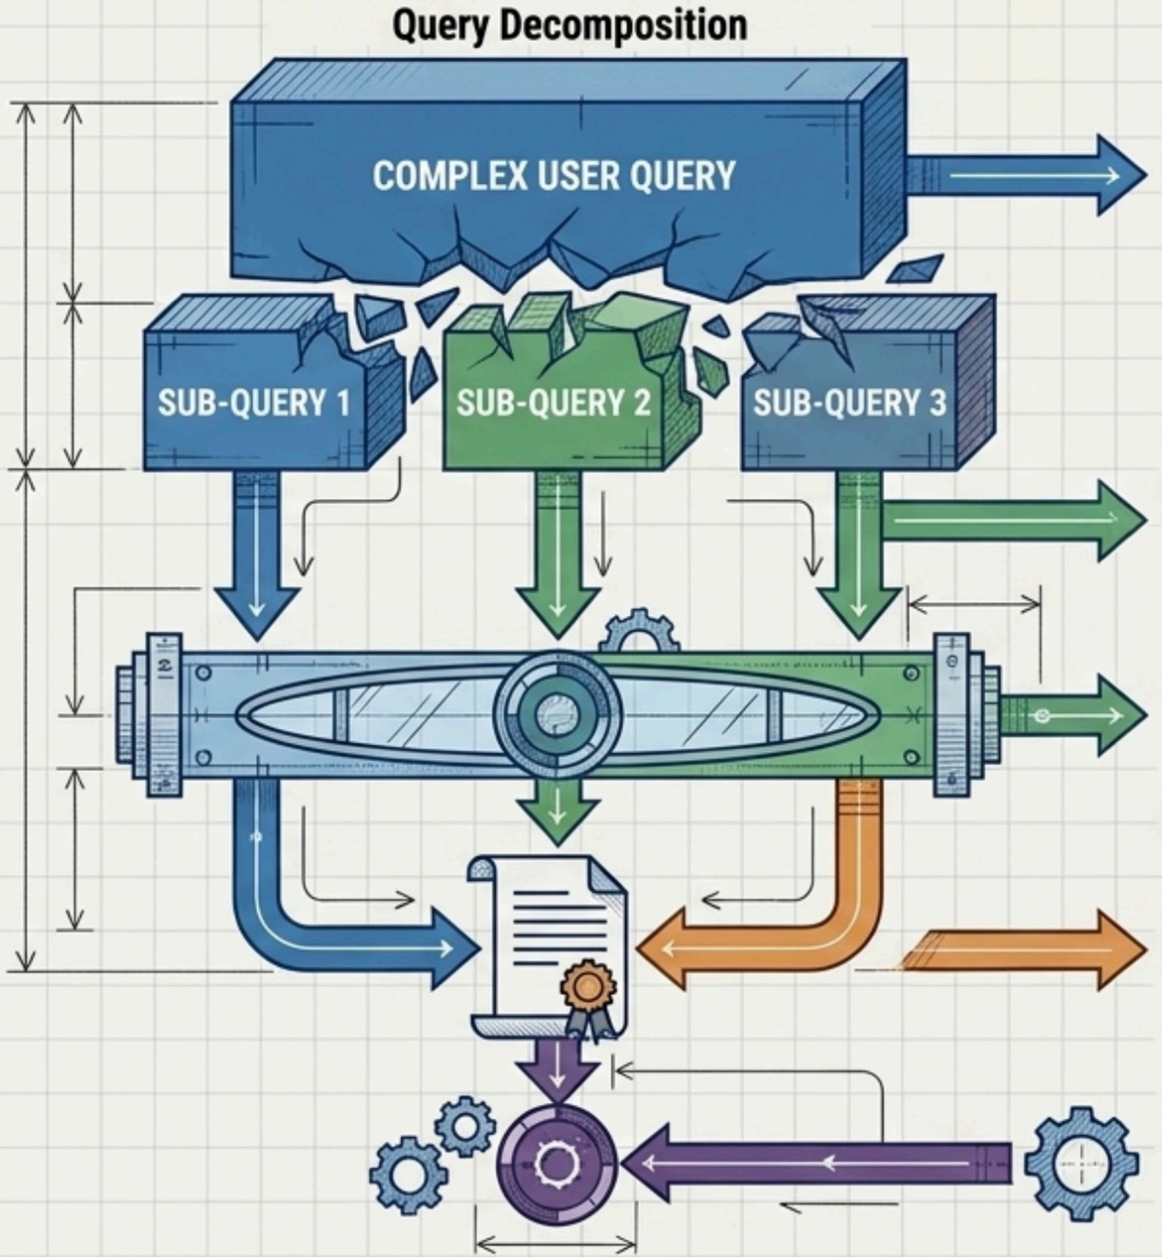

1. Decomposition: Gemini 3.1 Pro breaks the complex query into discrete, targeted sub-questions.

2. Iterative Retrieval: Each sub-question undergoes semantic search (k=5) and a dedicated canary probe.

3. Cross-Reference Validation: ChatGPT 5.4 validates findings across sources, flagging contradictions.

4. Synthesis: Claude 4.6 Opus integrates findings into a comprehensive, fully-cited response.

5. Quality Evaluation: Revision loop triggered if LLM-as-a-Judge score drops below 0.8.

Quality Assurance: The 16-Layer Recovery Cascade

A unified defense achieving 97-99% overall success, utilizing pre-validation gates, self-test block embedding, and up to 3 automatic self-healing debug attempts.

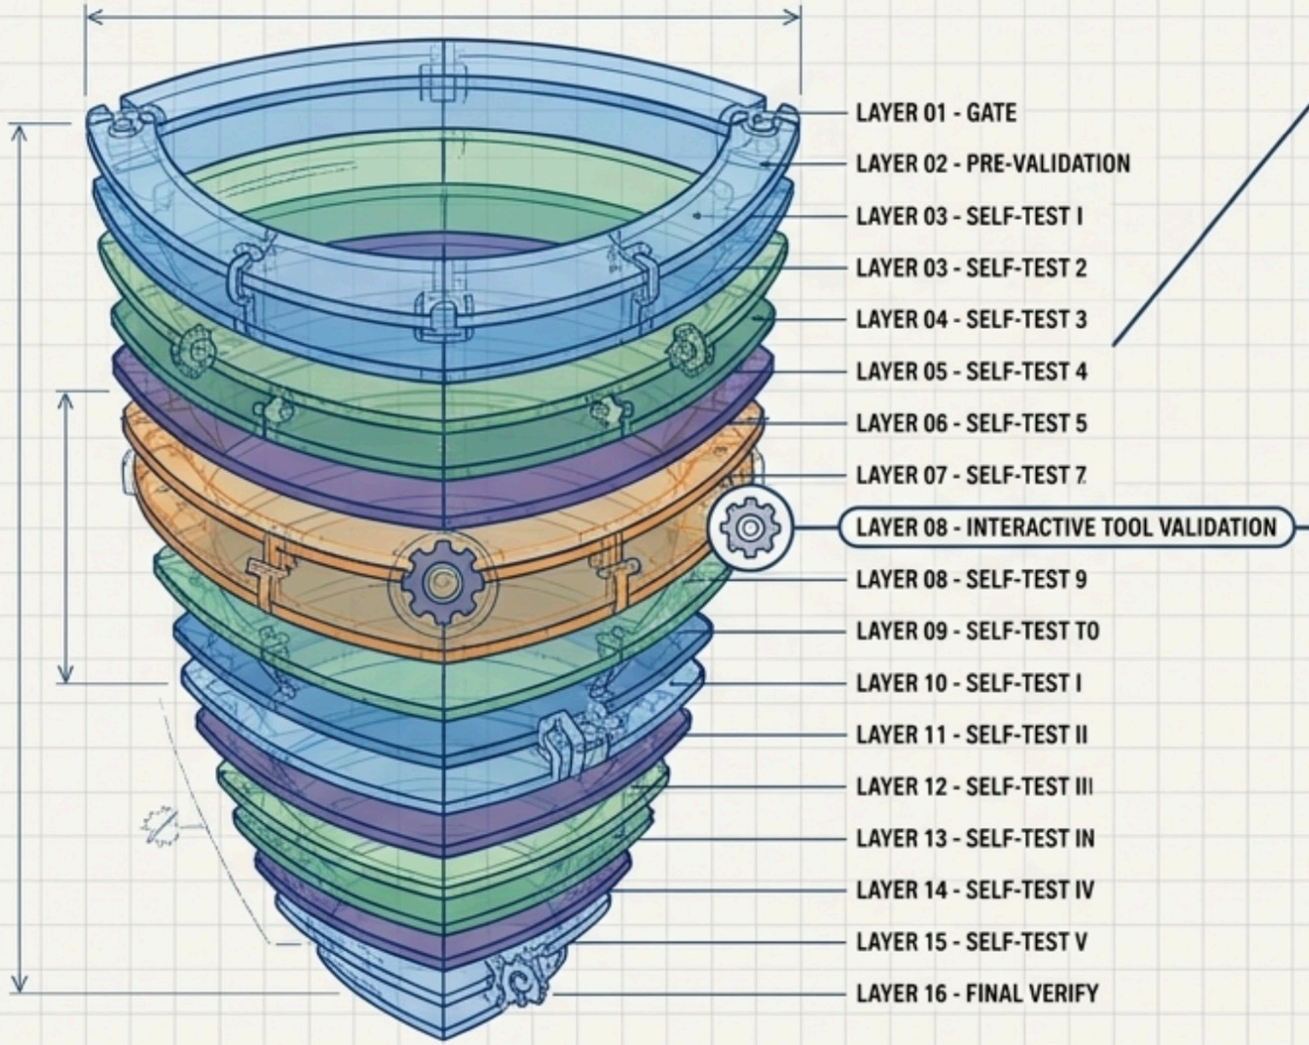

Interactive Tool Validation

Before an HTML calculator is served to a student, Claude 4.6 Opus embeds automated JavaScript self-tests targeting 8 adversarial edge cases:

- ☑ 1. Zero values
- ☑ 2. Negative inputs
- ☑ 3. Extreme large values (no overflow)
- ☑ 4. Extreme small decimals (precision maintained)
- ☑ 5. Boundary condition limits
- ☑ 6. Invalid input types (strings/nulls)
- ☑ 7. Edge case combinations
- ☑ 8. Rapid sequential inputs (race conditions)

Rethinking Academic Integrity: Stealth Verification

Accountability Without Surveillance: Moving away from unreliable AI Detectors toward a cryptographic verification model.

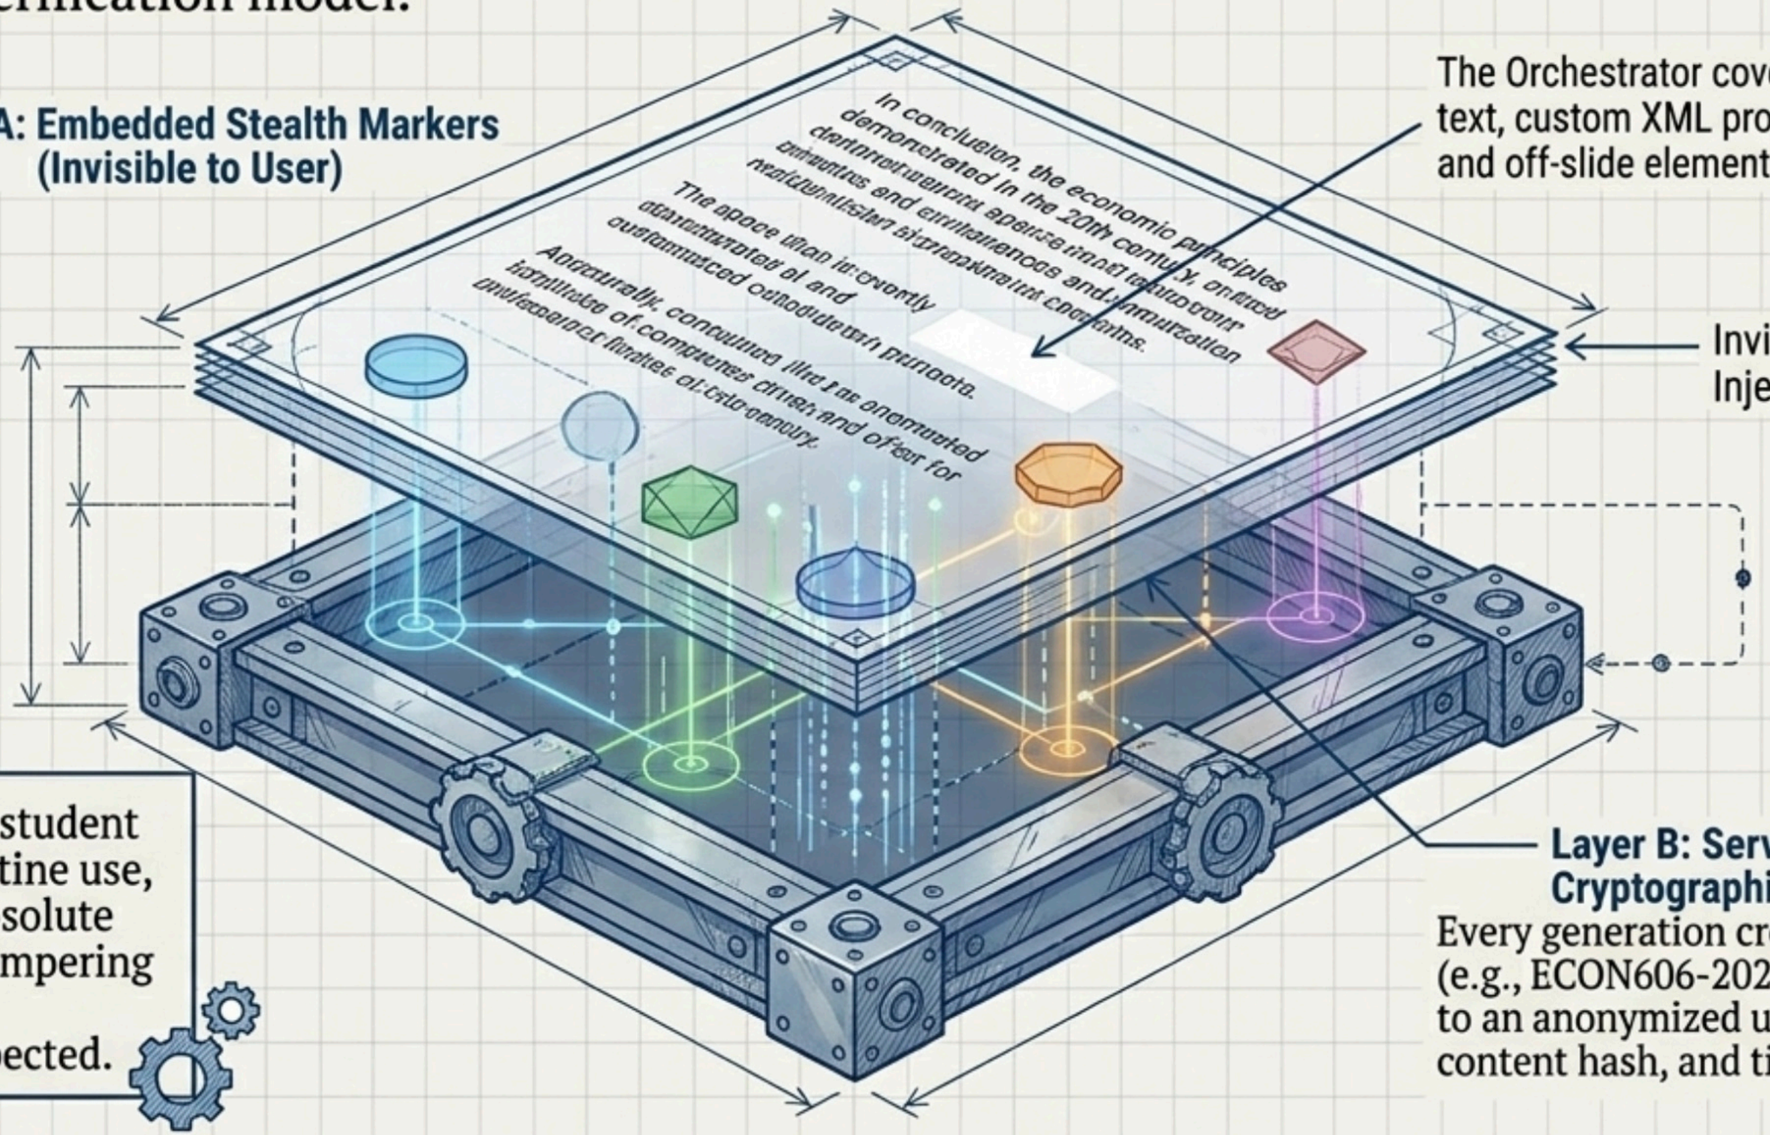

Layer A: Embedded Stealth Markers (Invisible to User)

The Orchestrator covertly injects 1pt white text, custom XML properties (core.xml), and off-slide elements in PPTX files.

Invisible Marker Injection Points

Layer B: Server-Side Cryptographic Ledger

Every generation creates a unique ID (e.g., ECON606-20260413-RV7K), tied to an anonymized user hash, SHA-256 content hash, and timestamp.

Impact: Preserves student privacy during routine use, while providing absolute forensic proof if tampering or unauthorized submission is suspected.