

Four products. One behavioral standard. Cryptographically verifiable across your entire AI footprint — from pre-deployment testing to autonomous robotic systems. Every product shares the LCSH behavioral standard and hash-chain verification, so a control proven in one product is verifiable across all four.

<p>PRODUCT 01 AVAILABLE</p> <h3>Preflight</h3> <p><i>Test It</i></p> <p>On-platform research and testing environment. Test any AI model with your actual system prompt, knowledge files, and tools before production. Run single assessments, structured Trials with A/B comparison, or scheduled recurring checks.</p> <p>Built for: AI Engineers · Research Teams · Product Managers · QA</p>	<p>PRODUCT 02 AVAILABLE</p> <h3>Certify</h3> <p><i>Prove It</i></p> <p>Runtime behavioral certification through TypeScript SDK and REST API. Every assessment produces a SHA-256 hash chain receipt anchored to Ethereum — tamper-evident, third-party-verifiable proof for auditors, regulators, and boards.</p> <p>Built for: Platform Engineers · DevOps · Compliance · Enterprise IT</p>
<p>PRODUCT 03 COMING SOON</p> <h3>Fleet</h3> <p><i>Govern It</i></p> <p>Independent conscience agent (Grillo) for multi-agent AI systems with constitutional separation of powers. Temporal Drift Index, Ethical Flight Plans, fleet-level anomaly detection — Grillo cannot be overridden by the agents it monitors.</p> <p>Built for: Fleet Operators · Enterprise AI Platforms · Autonomous Systems</p>	<p>PRODUCT 04 COMING SOON</p> <h3>Embody</h3> <p><i>Extend It</i></p> <p>Behavioral governance for physical AI systems — autonomous vehicles, industrial robots, humanoid assistants, drone fleets. The same LCSH framework and conscience architecture, adapted for safety-critical real-time control loops.</p> <p>Built for: Robotics · Autonomous Vehicles · Industrial Automation · Defense</p>

SUITE-WIDE CAPABILITIES

Behavioral Assessment™

Live in Preflight & Certify · Foundation for Fleet & Embody

120-question LCSH battery measuring agent disposition across Lying, Cheating, Stealing, and Harm. Personality classification into four archetypes via 4D Euclidean distance. Cryptographically sealed results.

Real-Time Output Evaluation™

In development for all four products

Framework-grounded runtime classification of agent outputs. Per-principle quadrant evaluation (Well Adjusted, Manipulative, Psychopath, Misguided) at the moment of action. Cross-family LLM judging prevents same-family bias.

LCSH FRAMEWORK — THE BEHAVIORAL STANDARD

<p>L1</p> <h3>Morality</h3> <p>Universal ethical foundations</p>	<p>L2</p> <h3>Virtue</h3> <p>Character and disposition</p>	<p>L3</p> <h3>Ethics</h3> <p>Applied principles in context</p>	<p>L4</p> <h3>Operational Excellence</h3> <p>Customer-customizable, domain-specific</p>
---	---	---	--