OFFICIAL-SENSITIVE / TECHNICAL ASSURANCE & UNDERWRITING BRIEF

EAVECORE: MITIGATING AI LIABILITY

The Senior Managers & Certification Regime (SM&CR) Liability Framework

Document Ref: #AXM-BRIEF-SMCR-2026.1

TARGET AUDIENCE:	Chief Risk Officers | General Counsels | Cyber-Liability Underwriters
REGULATORY ANCHOR:	FCA Senior Managers and Certification Regime (SM&CR)
PRIMARY MECHANISM:	Cryptographic Human-in-the-Loop (C-HITL) & Modular Micro-Cartridges
STATUTORY STANDING:	Section 7 Electronic Communications Act 2000 (ASAR Receipts)

1 1. The Insurability Crisis: Infinite Risk vs. Statutory Accountability

The fundamental barrier to enterprise AI adoption is not technological capability; it is legal liability. Under the Financial Conduct Authority's Senior Managers and Certification Regime (SM&CR), human executives bear ultimate, non-delegable responsibility for regulatory breaches in the areas of the firm for which they are responsible.

Currently, deploying a Large Language Model (LLM) into a regulated financial environment introduces an uninsurable, infinite risk surface. Because foundational models are inherently probabilistic, their operational pathways cannot be mathematically bounded or guaranteed. When "automation bias" sets in, human operators rubber-stamp non-compliant algorithmic decisions, effectively outsourcing SM&CR liability to an opaque black box.

The Axiom Solution: EAVEcore does not attempt to make the LLM itself safe. Instead, it acts as a physical execution turnstile at the API perimeter, transforming uninsurable probabilistic workflows into mathematically bounded, deterministic assets.

2 2. Bounded Risk Surfaces via Modular Cartridges

Underwriters cannot price a policy against an entire LLM because a failure in one area (e.g., a data hallucination) can cascade into another (e.g., executing a prohibited trade or commingling funds). EAVEcore eliminates this structural exposure through the **Modular Micro-Cartridge Registry**:

- **Sterile Context Windows:** Statutory rules are never bundled into a vulnerable, non-deterministic "mega-prompt." They execute as isolated, hot-swappable micro-cartridges (e.g., FCA CASS 7.13, UK-GDPR).
- **Sub-Millisecond Interception:** Live empirical testing confirms that when an AI attempts to violate a hardcoded policy—such as commingling retail client funds with corporate treasury—EAVEcore intercepts and drops the API payload in **0.31 milliseconds**.
- **Zero Business Interruption:** If a specific model exhibits semantic drift, the firm does not need to shut down its entire AI infrastructure. EAVEcore simply locks the specific failing cartridge while the rest of the enterprise pipeline remains fully operational.

3 3. Curing Automation Bias: The C-HITL Protocol

To maintain SM&CR integrity, the human operator must remain the ultimate arbiter of risk. EAVEcore enforces this through the **Cryptographic Human-in-the-Loop (C-HITL)** gateway.

When EAVEcore detects a divergence from statutory maxims, it physically halts the execution pipeline. The system systematically rejects vague, rubber-stamp approvals (such as `"looks fine"`). To unlock the gateway, the operator must provide a **Maxim-Anchored Covenant**—a legally binding justification explicitly citing the relevant statutory rule.

THE C-HITL LIABILITY TRANSFER GATEWAY

1. The Intercept: The non-compliant command is blocked at the border in 0.31 ms.
2. The Rejection: Rubber-stamp bypass attempts fail automatically without SM&CR attachment.
3. The Covenant: The operator signs a Maxim-Anchored Covenant referencing specific statutory guidelines.
4. The Result: Liability is deterministically re-absorbed by the certified human manager, generating an un-alterable cryptographic receipt before ledger settlement can proceed.

4 4. Rapid Rectification & Ombudsman Reporting Protocol

When an incident occurs or a systemic flaw is identified, Senior Managers must demonstrate to the FCA or the Financial Ombudsman Service (FOS) that immediate, reasonable steps were taken to halt the breach and rectify the system. EAVEcore delivers a mathematically proven incident response pipeline:

1. **Instant Granular Isolation:** The compromised API route is severed in fractions of a millisecond without affecting adjacent compliance turnstiles.
2. **Surgical Hot-Swapping:** The isolated EAVEcore rule cartridge is updated and hot-swapped into the Dynamic Routing Bus with zero system downtime.
3. **Cryptographic Proof:** Every successful EAVEcore intercept and human override generates an immutable Sovereign Audit Receipt (ASAR) logged via TallySticks UK CIC.

Structured under **Section 7 of the Electronic Communications Act 2000**, these cryptographic receipts are legally admissible in UK courts. EAVEcore allows Senior Managers to hand regulators an indisputable, millisecond-accurate timeline of the cure, proving exactly when the system was secured and mathematically guaranteeing that no further violations passed the perimeter after the patch.

5 5. Strategic Conclusion for Underwriters & Boards

THE UNDERWRITING AXIOM

By wrapping agentic workflows in the EAVEcore turnstile, FCA-supervised entities eliminate the uninsurable liability of non-deterministic AI. We provide the deterministic border that allows insurers to underwrite AI operations based on hard, mathematical interception metrics rather than probabilistic guesswork.