CYTOMATE

AEVA

Artificial Intelligence Evidence, Validation & Assurance

A CONCEPT PAPER

Defining the Open Standard for Continuous AI Assurance



EVIDENCE
Capturing AI system behavior, artifacts and decisions

VALIDATION
Verifying integrity, performance and compliance

ASSURANCE
Delivering trust through transparency and accountability

AI FORENSICS
Reconstruct, analyze and explain AI incidents

AI THREAT HUNTING
Detect anomalies, malicious behavior and abuse

MODEL MRI
Deep inspection of AI models and their lineage

DOHA, QATAR
2026

"AI adoption is accelerating.
Assurance must scale with it."
Countries cannot safely deploy AI at national scale without infrastructure that makes assurance **CONTINUOUS, MEASUREABLE AND PROVABLE.**

AI ASSURANCE | AI FORENSICS | AI THREAT HUNTING | MODEL MRI

www.cytomate.net



AEVA

by **Cytomate**

Artificial Intelligence Evidence, Validation & Assurance

A Concept Paper

Defining the Open Standard for Continuous AI Assurance

> *AI adoption is accelerating. Assurance must scale with it. Countries cannot safely deploy AI at national scale without infrastructure that makes assurance continuous, measurable and provable.*

I. The Idea

Every transformative technology eventually demands its own assurance infrastructure. For aviation, it was airworthiness certification. For finance, it was audit. For AI, that infrastructure does not yet exist — and its absence is becoming the defining risk of the decade.

AEVA — Artificial Intelligence Evidence, Validation & Assurance — is Cytomate's answer to this absence. It begins with a simple but radical observation: an AI system is not a product you test once before shipping. It is a living ecosystem of models, agents, memory, knowledge sources, tools and defenses, operating in an environment that evolves continuously. Assurance over such a system cannot be periodic. It must be perpetual.

AEVA is built around this conviction. It is not a compliance checklist. It is not a vulnerability scanner dressed in AI clothing. It is a Command & Control Center for the AI estate — a platform that discovers every AI asset in an organization's environment, tests it under adversarial conditions, validates the defenses surrounding it, collects structured evidence and produces a living, signed record of each system's assurance state. From the boardroom to the regulator's desk, that record is what accountable AI deployment requires.

> *From control claims to measurable, reproducible assurance — this is the transformation AEVA makes possible.*

Conceived and built in Doha, Qatar, AEVA carries a regional ambition and a global aspiration. The GCC sits at the epicentre of sovereign AI investment — with national programs in Saudi Arabia, the UAE and Qatar committing to AI as a pillar of economic transformation. Yet the governance infrastructure to underpin that commitment has not kept pace. AEVA exists to close that gap, first for the region, and in time for the world.

II. The Problem

Across government and enterprise, AI is being deployed faster than it is being governed. The result is not merely a technology risk — it is a sovereignty risk, a liability risk and, in sectors where AI makes consequential decisions, a public trust risk.

The structural failure is not one of intent. Organisations want to govern their AI. Regulators want to set standards. Governments want to protect citizens. The failure is one of infrastructure: the tools, frameworks and evidence mechanisms required to make AI governance operational simply do not exist at the scale or maturity that the moment demands.

Consider what most organisations face today. AI systems are introduced across departments without central inventory — meaning no-one can enumerate what AI assets exist, who owns them or what they do. Where testing occurs, it is manual, isolated and conducted once, against a system that will change many times before the next assessment. There is no centralised visibility, no continuous monitoring, no mechanism for collecting evidence that a regulator could inspect with confidence.

01	AI systems are deployed without complete inventory. Shadow AI is endemic. No organisation can claim full visibility of its AI estate.
02	Testing is periodic and manual. A model assessed six months ago may have drifted in behaviour. A RAG pipeline updated last week may carry poisoned knowledge.
03	Governance exists on paper. Policies and frameworks are written but rarely operationalised into continuous evidence collection and accountability mechanisms.
04	Defences are assumed, not validated. Guardrails, firewalls and human oversight workflows are claimed as controls but seldom tested under adversarial conditions.
05	Evidence is absent. When a regulator, a board or a court asks whether an AI system was safe at the time of a consequential decision, there is no answer.

AI introduces attack vectors that existing security frameworks were never designed to address: prompt injection, jailbreaks, RAG poisoning, memory manipulation, agent hijacking, hallucination under adversarial pressure, data leakage through carefully crafted queries. These are not theoretical risks. They are the operational reality of any organisation that deploys AI in a contested environment — which, in 2026, means every organisation.

> *AI systems require validation of the intelligence, the defences, the operational response and the evidence — not just the infrastructure.*

III. The Vision

AEVA's ambition is not to build a better compliance tool. It is to build the infrastructure layer that makes accountable AI deployment possible — open by design, sovereign by architecture, trusted by evidence.

The vision is structured around a single, clarifying concept: closed-loop assurance. Not a point in time. Not a snapshot. A continuous cycle — Discover, Assess, Attack, Validate, Observe, Collect Evidence, Find, Remediate, Retest — that runs in perpetual motion against every AI system an organisation operates. Every turn of the loop produces structured evidence. Every piece of evidence feeds a living AI Passport — a signed, time-bound record of a specific system's assurance state, its components, its findings, its residual risk and its next reassessment trigger.

The AI Passport is AEVA's central innovation. In the way that a financial audit produces a signed opinion on the state of an organisation's accounts, the AI Passport produces a signed assurance record on the state of an AI system. It travels with the system. It exposes its lineage — which model version, which prompts, which data sources, which defences were in place and tested on what date. And when any of those components change, the Passport's assurance state is automatically invalidated, triggering the next validation cycle.

> *The AI Passport makes the implicit explicit — transforming assurance from a claim into a credential.*

Beyond the enterprise, AEVA's vision extends to the national scale. The National AI Assurance Fabric — AEVA's federated architecture for sovereign AI governance — enables country-level visibility over the AI systems operating in government and critical national sectors, without requiring organisations to expose their proprietary models or sensitive data to a central authority. Local Assurance Nodes keep evidence local. Signed attestations travel upward to a national coordination layer, where regulators and AI authorities receive the situational awareness they need to govern AI at national scale while preserving organisational autonomy.

This is the North Star: not a product, but infrastructure. Not a service, but a standard. An open reference architecture that governments, enterprises, universities and international standards bodies can build on, contribute to and ultimately trust — the way they trust the foundations of financial audit, aviation safety or pharmaceutical approval.

IV. The Platform

AEVA is already operational. The vision described above is not a roadmap aspiration — it is a working platform, with core assurance capabilities validated and ready for enterprise deployment today.

What AEVA Does

AEVA functions as an AI Command & Control Center — the single pane of glass through which an organisation governs its complete AI estate. It maps every component of every AI system: foundation models, fine-tuning layers, system prompts, embedding models, vector databases, knowledge sources, agents, memory stores, tool connectors, MCP servers, AI firewalls, gateways, runtimes and human approval workflows. It then validates each of these components — and the interactions between them — through multi-engine adversarial testing, integrating best-in-class open-source AI security frameworks alongside Cytomate's own native validation capabilities.

Four Domains of Assurance

AI System Assurance	Continuous validation of models, prompts, RAG pipelines, agents, memory and tools — testing behaviour under adversarial conditions and tracking drift over time.
AI Defence Assurance	Testing of the controls that protect AI systems: guardrails, firewalls, identity and access management, data loss prevention, sandboxes and human approval workflows.
Operational Assurance	Validation of the organisation's ability to detect and respond to AI-specific incidents — through AI SOC capabilities, SIEM integration, alerting and containment telemetry.
Governance & Evidence	Automated generation of structured evidence packages, AI Passports and compliance mappings — the artefacts that make assurance demonstrable to boards, regulators and courts.

Current Capabilities

The AEVA MVP is operational, with the following capabilities validated and deployment-ready: AI system and model registration; multi-engine adversarial testing; RAG security assessment and model vulnerability evaluation; automated evidence collection; unified findings dashboard; assurance profiles and sector-specific templates; continuous compliance mapping; and exportable audit-ready reports. Enterprise automation, CI/CD pipeline integration and advanced governance capabilities are in active development.

The Road Ahead

The platform's development trajectory moves in four phases: first, engineering maturity and production hardening; second, advanced closed-loop assurance with the Evidence Graph and AI Assurance Passports; third, multi-tenant enterprise deployment with high-availability architecture; and fourth, sovereign and national infrastructure — the National AI Assurance Fabric, government deployment, federated assurance and the national AI registry. At each phase, capability compounds: what begins as an enterprise validation platform evolves into the AI Command & Control Centre that national AI governance requires.

V. The Ecosystem

AEVA is both a platform and a founding contribution to something larger: the AI Assurance Commons — an open ecosystem of specifications, scenarios, evidence schemas and research that any organisation can use, contribute to and build upon.

The ecosystem is structured in three layers. The AI Assurance Commons is the foundation: open standards, open specifications, community research, shared assurance scenarios and the AI-BOM and Passport schemas that define what structured AI evidence looks like. AEVA is the reference implementation of those commons: the Console, API, CLI, Evidence Test Engine and Assurance Compiler that operationalise the standards into a working platform. Cytomate Enterprise is the commercial scale layer: sovereign deployments, air-gapped environments, managed assurance services, national fabric support and advanced validation for regulated and mission-critical environments.

> *The Commons defines the ecosystem. AEVA operationalises it. Cytomate scales it for enterprise and sovereign environments.*

Open source is a strategic choice, not a concession. When government AI authorities, international standards bodies and university research programs can inspect the specifications underlying assurance claims, they become contributors to and advocates for the standard. The network effect this creates — more scenarios, more regulatory mappings, more evidence packs, more real-world validation — is what transforms a platform into infrastructure. No proprietary product can replicate it.

Universities are a particular lever. Bachelor, master and doctoral research programs can continuously expand AEVA's attack scenario coverage, evaluate emerging AI architectures, develop new evidence methodologies and contribute to the Commons at a scale that no commercial R&D program alone could match. This is how AEVA becomes comprehensive — not by building everything, but by creating the infrastructure that draws the world's research capacity into a single, coherent assurance ecosystem.

VI. Why Cytomate

The AI assurance market is nascent. Most participants are traditional cybersecurity vendors extending existing tooling, or AI testing startups without the security validation pedigree to build credible assurance. Cytomate stands at the intersection of both — with a track record and a regional position that no competitor can replicate.

Cytomate is the GCC's specialist in cyber validation. Its product portfolio — Rasid for external attack surface management, Battle Twin for OT breach and attack simulation, Breach+ for IT breach simulation and Sarab for cyber deception — demonstrates a consistent methodology: simulate real attacks, validate real controls, produce real evidence. AEVA applies that same methodology to AI. The parallel is direct and deliberate. Where traditional cyber validation runs Attack, Control, Detection, Response, Mitigation — AI assurance runs AI Attack, Behaviour, Controls, AI SOC, Evidence, Governance, Remediation. Cytomate's tooling, methodology and evidence culture extend cleanly into the AI domain.

Cytomate operates across the GCC and MENA region with established relationships in Saudi Arabia, the UAE and Qatar, and deep familiarity with the regulatory frameworks that govern AI and cybersecurity in these markets — NCA ECC, SAMA CSF, SDAIA AI Ethics, UAE Information Assurance. AEVA's compliance mappings are not generic — they are built for the actual governance environment its customers operate in, by a team that has spent years inside that environment.

And then there is geography. Doha is not an accident. Qatar's position as a regional hub for technology, finance and government innovation, combined with its active engagement in international AI governance discussions, makes it the right place to originate a platform with global ambitions. Born in Qatar, designed for regional leadership, built for global AI assurance — this is not a slogan. It is a strategy.

VII. The Call

The convergence of sovereign AI ambition and AI security risk creates a defining moment. The infrastructure does not yet exist at the scale required. The window to define it is open — but not indefinitely.

AEVA is ready. The platform is operational. The methodology is proven. The regulatory environment in the GCC is aligning rapidly, with AI governance mandates being formalised across every major market in the region. The question is not whether a sovereign AI assurance standard will emerge. It is who will define it, and from where.

Cytomate's proposition is that the standard should be open, that it should be built from the region's own regulatory reality, and that it should be anchored in the most rigorous validation methodology available. AEVA is the vehicle for that proposition. The AI Assurance Commons is its open foundation. And the partnerships required to make it the global reference standard are the ones we are seeking today.

Governments & National AI Authorities	We are seeking partnerships with national AI authorities, cybersecurity agencies and digital economy ministries ready to deploy sovereign AI assurance infrastructure — including SDAIA, NCSA, MCIT and equivalent bodies across the GCC.
Regulated Enterprises	We are inviting early adopters from financial services, healthcare, critical infrastructure and government services who are ready to deploy AEVA as their AI assurance platform and contribute real-world scenarios to the Commons.
University Research Programs	We are building a research network — bachelor, master and doctoral programs in AI, cybersecurity and governance — to continuously expand AEVA's testing coverage and contribute to the open Commons.
Strategic Investors	We are engaging investors who understand that AI assurance is infrastructure — long-cycle, high-moat, regulatory-mandated — and who want to own a position in the platform that defines the standard.

> *Traditional AI testing evaluates the intelligence. Traditional security validation evaluates the defences. AEVA validates the intelligence, the defences, the operational response and the evidence connecting them.Born in Qatar. Designed for Regional Leadership. Built for Global AI Assurance.*

CYTOMATE · AI ASSURANCE · OPEN INFRASTRUCTURE · DOHA