



WHITE PAPER DIGEST



Three-Path Feedback Dual-Loop Information Processing Architecture

for Pre-Execution Governance

Not to stop AI, but to enable it to act responsibly.



Executive Focus

A concise corporate digest of the EVA3 architecture for pre-execution governance, explainable decision-making, and three-path feedback learning.

Version 1.1 Executive Edition

July 2026

Amuletplus G.K.

What EVAE Solves

Designing the decision structure immediately before AI moves into action.

Core Proposition

AI safety and explainability cannot be secured through post-output auditing alone. Before execution, the origin, options, decision boundary, and observation record must be structured, and action outcomes must be returned to the next decision cycle.

EVAE is a pre-execution decision architecture that connects a consciousness loop and an action loop asymmetrically. It can be connected outside or inside existing AI systems, agents, and business systems to establish a decision boundary immediately before execution.

- | | | |
|---|--------------------------------|---|
| 1 | E (Origin) | Clarifies the starting point of a decision, such as a request, intent, danger signal, or anomaly. |
| 2 | V (Possibility) | Generates multiple options and future impacts instead of jumping to a single answer. |
| 3 | Λ (Decision Boundary) | Determines the conditions and authority for execution, hold, stop, or reconsideration. |
| 4 | Ξ (Observation / Trace) | Records the execution result and decision process and returns them to the next cycle. |

Operational Positioning

EVAE does not replace laws or standards. It connects them to operational decisions and treats speed and control, rapid response and deliberation, execution and learning as one continuous cycle.

AI Has Become Smarter.

But the structure of execution responsibility has not kept pace.

Output-Centric

Many AI safeguards focus on content evaluation and post-event audits. Real harm begins when output is connected to authority, APIs, equipment, or customer touchpoints.

Single Path

Conventional workflows tend to follow 'generate → execute → log' and lack structural routes for hold, reconsideration, or stop.

Responsibility Gap

Organizations often cannot later explain who authorized an action, on what basis, or where it should have been stopped.

As AI agents perform work more autonomously, the central question shifts from “Can the system produce the right answer?” to “Under what conditions may it change reality?”

Five Questions Required Before Execution

What triggered the decision? / What alternatives were available? / Where was the stop point? / Who decided? / How was the trace preserved?


- Hallucinations become critical when unverified output is connected to execution authority.
- Black boxes also arise when the chain of judgment, authority, approval, and retry is not recorded.
- Safety requires not only prohibitions, but structured paths to act, wait, stop, or reconsider according to the situation.

The EVA3 Answer

Answer these five questions through architecture, not merely through prose.

What Is EVAE?

Decomposing pre-execution AI decision-making into four reusable stages.



Stage	Role	Primary Question
E – Origin / Impulse	Captures requests, intent, anomalies, constraints, and danger signals.	What started this decision?
V – Possibility	Generates alternatives, future candidates, impacts, and exceptions.	What other options exist?
Λ – Decision Boundary	Determines execution conditions, prohibitions, authority, and thresholds.	Who may authorize what, and how far?
∃ – Observation / Trace	Preserves outcomes, differences, rationale, approvals, and reevaluation conditions.	What actually happened, and what changes next?

EVAE Is Not an AI Model

It is used with existing technologies such as LLMs, rule engines, RPA, robots, and business APIs. Its value lies not in model capability, but in making the decision structure before and after execution explicit.

Because each stage can be treated as an independent module, each business process can define its own inputs and outputs, responsible parties, thresholds, and log format. Explainability is therefore implemented not as generated prose, but as the ability to reconstruct the actual decision process.

Four-Part Logic

Origin captures the trigger. Possibility expands alternatives. Decision Boundary defines authority. Observation/Trace closes the loop through evidence.

E · V · Λ · ∃: Four Stages That Bring the Future Back to Reality

Expand possibilities from the origin, define execution conditions at the boundary, and return observations to the next cycle.



E Origin	V Possibility	Λ Decision Boundary	∃ Observation / Trace
Captures purpose, context, danger signals, and external constraints to clarify the starting point.	Generates multiple options, non-execution options, and staged-execution options.	Integrates expected value, risk, cost, alignment, and authority to determine the executable range.	Structures results, differences, approver, data used, exception handling, and improvement conditions.

Candidate Scoring

$$\text{score} = \alpha \times \text{expected value} - \beta \times \text{risk} - \gamma \times \text{cost} + \delta \times \text{alignment}$$

+ Expected Value future value	– Risk danger / uncertainty	– Cost time / burden	+ Alignment policy / authority
---	---------------------------------------	--------------------------------	--


Key Point

The formula itself is not the central value. The essential requirement is to make evaluation axes, weights, and approval authority explicit and to preserve why they were changed.


The Asymmetric Dual Loop

Separating the consciousness loop that deliberates from the action loop that changes reality.

Consciousness Loop



Action Loop




Why Asymmetric?

Thinking and changing reality do not carry the same weight. Action involves authority, cost, and potential irreversibility. Therefore, action outcomes must always return to the consciousness loop for reevaluation.

- In urgent situations, the action loop can be shortened to reduce latency.
- When uncertainty is high, the consciousness loop can be expanded to increase alternatives and evidence.
- Results do more than become the next input: Va transforms them into alternatives and new future candidates.

Three-Path Feedback

Returning action outcomes to the consciousness loop at different depths depending on the situation.



Decision Indicator	Typical Condition	Primary Return Point
Risk level	Collision, critical failure, authority deviation, security anomaly	Ec: immediate restart
Satisfaction / achievement	Below expectations, willingness to retry, formation of next objective	Va → Ec: new origin
Prediction error	Error exceeds a threshold for a specified number of consecutive cycles	Va → Vc: reconstruct option space

Point

This is not merely “continue on success, stop on failure.” The system chooses whether the meaning of the result should return to immediate response, the next intention, or assumption revision.

Pre-Execution Governance

Turning policies into actual decision gates.



Control Element	Example Implementation in EVA3
Authority boundary	Define roles, approvers, executable scope, spending limit, and target environment in Λ .
Risk thresholds	Branch to execute, hold, or stop according to risk, confidence, impact scope, and reversibility.
Human oversight	Escalate critical operations for human approval and record the reason in Ξ .
Audit trail	Store origin, options, selection rationale, model/data, and execution outcome as a decision chain.
Rollback	Include reversible execution, staged execution, test environments, and recovery conditions as options in V .



Value of Pre-Execution Control

Instead of only investigating why a failure occurred, organizations can define machine-readable boundaries that specify when the system must not act.

- Translate laws, standards, and policies from abstract rules into operational decision gates.
- Design execution authority, stop conditions, and audit evidence as one integrated structure.
- Return outcomes through the three paths and incorporate them into the next decision.

Enterprise Implementation: EVA3 ARMOR

Adding a layer of decision, authority, and traceability between existing AI and business systems.



ARMOR is not a product dedicated to a single model. It is positioned as an enterprise PoC and implementation package that can connect to cloud LLMs, internal models, RPA, DevOps, and robotics.

- Introduce it gradually at execution points without replacing the entire existing system.
- Design logs as auditable structured events rather than narrative text.
- Change Λ decision boundaries by department, use case, and risk level.

PoC Focus

The question is not simply whether AI produced an answer, but what was compared before execution, where the system stopped, and what evidence was preserved.

Application Areas

Attaching the same structure to decision-making across different industries.

Domain	E: Origin	Δ: Decision Boundary	Σ: Record
IT / DevOps	Change request, incident, vulnerability alert	Production authority, impact scope, approval, rollback	Change diff, executor, recovery, incident impact
Education	Learning objective, wrong answer, motivation decline	Material difficulty, next unit, need for intervention	Understanding, emotional change, continuity, error
Finance	Transaction request, market movement, fraud signal	Limit, suitability, approval, hold	Rationale, price, risk, customer impact
Healthcare Support	Symptoms, test values, clinical purpose	Scope of use, human confirmation, urgency	Recommendation basis, physician judgment, outcome difference
Robotics	Sensors, instruction, environmental change	Safety distance, speed, stop conditions	Trajectory, anomaly, avoidance, re-control

Shared Design Principle

The structure of origin, options, decision boundary, and observation record remains the same across industries. What changes are the metrics, authority, thresholds, and log fields.

Execution Risk

Evaluate authority deviation, incorrect execution, and excessive automation before action.

Auditability

Include not only the outcome, but also the options that were not selected.

Adaptive Learning

Separate improvement into emergency response, next intention, and assumption revision.

Implementation Value

Improving safety, explainability, implementation speed, and organizational learning at the same time.

Transparency

Decompose decisions into modules and reconstruct which options were excluded and who authorized execution.

Adaptability

Return outcomes through three paths and distinguish rapid response, short-term learning, and assumption revision.

Ease of Implementation

Connect outside existing AI without depending on a particular model or industry.

Conventional Approach	EVA3
Detect problems after output	Determine permission and stop conditions before execution
Single success/failure log	Record origin, options, selection, authority, and result as a chain
Rebuild controls when the model changes	Reuse model-independent decision boundaries
Human approval easily becomes procedural	Structure approval rationale, authority, and exceptions
Learning is a simple reflection into the next input	Separate immediate response, new intention, and option-space revision

Redefining Explainability

In EVA3, explainability does not mean asking AI to write a plausible explanation. It means preserving the actual decision chain in a form that can be verified later.

Alignment with Regulation and Standards

EVΛΞ does not replace laws or standards; it connects them to operational execution decisions.

External Framework	Primary Focus	Connection to EVΛΞ
EU AI Act	Risk-based regulation, transparency, human oversight, recordkeeping, GPAI obligations	Define execution conditions, oversight, and prohibitions in Λ; preserve evidence in Ξ.
GPAI Code of Practice	Practical support for transparency, copyright, safety, and security compliance	Track rights reservations and use conditions through interpretation, decision, and downstream use as a decision chain.
ISO/IEC 42001:2023	Establish, operate, maintain, and continually improve an AI management system	Translate organizational policy into E, V, Λ, and Ξ operating processes.
NIST AI RMF 1.0	Risk management through Govern / Map / Measure / Manage	Cycle risk context, evaluation, response, and observation through the dual loop.

2026 Operating Environment

The EU AI Act entered into force in August 2024, with general full applicability on 2 August 2026. GPAI rules have applied since August 2025, and the European Commission states that enforcement powers apply from August 2026.

Sources: European Commission AI Act overview; GPAI Code of Practice; ISO/IEC 42001:2023; NIST AI RMF 1.0. Accessed July 2026.

Intellectual Property

Protecting the core architecture by patent and the name and deployment domains by trademarks.

Patent Application

Application No.: 2025-160873
Title: Three-Path Feedback Dual-Loop Information Processing System Connecting a Consciousness Loop and an Action Loop

The core of the patent is not the symbols themselves. It is the information-processing structure that asymmetrically connects a consciousness loop and an action loop and returns action outcomes to the consciousness side through three paths: Emergency Response, Learning A, and Learning B.

Class	Trademark	Application / Registration No.	Registration Date / Status
Class 25	EVΛΞ	No. 70000549	24 Dec 2025
Class 16	EVΛΞ	No. 7039353	24 Apr 2026
Class 42	EVΛΞ	Trademark Application No. 2025-73597	Pending (amendment/correction completed)

IP Strategy

The patent protects the core architecture, while trademarks protect market distinctiveness and service deployment domains. The Class 42 filing expands protection to AI software, decision-support services, and applications, linking research papers, white papers, PoCs, and licensing.

External Engagement and Validation

Updating the architecture through continuous movement between research, policy dialogue, and enterprise implementation.

Area	Primary Activity
EU AI governance	Registered in the EU Transparency Register. Participated in a workshop and submitted a Position Paper concerning GPAI Code of Practice Measure 1.3.
Public communication	Published articles through Civil Society News Network on AI governance, A MA I, and responsibility-by-design for children and AI.
Enterprise / PoC	Developing ARMOR implementation designs and demonstrations for IT / DevOps, education, manufacturing / FA, finance, and other areas.
Research publication	Publishing analyses of the EVA3 stages and framework to articulate the architecture and invite third-party review.
Human development	Training human thinking, judgment, action, and observation as a dual loop through EVA3 MIND.

Positioning

These activities do not imply regulatory authority. Their purpose is to test EVA3 at the intersection of technology, society, and institutions and refine it into an implementable form.

Validation Logic

Research informs structure. Policy dialogue tests alignment. Enterprise PoCs test practicality. Human development extends the framework into training and use.

EU Transparency Register: 1702237103721-40. Activity descriptions are based on Amuletplus G.K. records as of July 2026.

The EVAE Ecosystem and Expansion

Bringing philosophy, architecture, human development, enterprise implementation, and applications into one cycle.

Layer	Role
A MA I	A philosophical foundation for restoring human 'Ma' (space / interval) and 'Ai' (love) in the AI era.
EVAE Dual Loop	The core architecture consisting of a consciousness loop and an action loop.
EVAE MIND	Training and a thinking method for updating human cognitive pathways.
EVAE MIRROR	Applications that support reflection and review.
EVAE NEXUS	A decision application for option generation, future simulation, and decision recording.
EVAE ARMOR	A PoC and implementation package that introduces pre-execution control into enterprise AI and agents.



Next Stage

Through accumulation of reference PoCs, standardization of audit-event formats, industry-specific profiles, and partner licensing, EVAE will evolve into a general-purpose foundation for pre-execution governance.

Not to Stop AI, but to Enable It to Act Responsibly.

In an era when AI is connected to the real world, capability competition alone is not enough. Before execution, organizations must structure the purpose, the possibilities considered, who may authorize what, and what must be recorded.

The EVAE Promise

Capture the origin. Expand the possibilities. Establish the decision boundary. Observe the result and carry it forward into the next decision cycle.

EVAE is not a technology confined to a single industry or model. It can be connected wherever decisions are made, actions change reality, and learning from outcomes is required.

Contact / Framework Scope

Amuletplus G.K.
EVAE Framework / EVAE ARMOR / EVAE MIND
Website: amuletplus.co.jp
EU Transparency Register: 1702237103721-40
Patent Application: 2025-160873

Disclaimer: This document is provided for technical and research information only and does not constitute legal advice or a guarantee of regulatory compliance. The content is based on information available as of July 2026 and information provided by Amuletplus G.K.

Reference Links and Demos

Public reference links for reviewing the EVAE framework, its patent-pending architecture, and working demonstrations. These links are provided for private review and professional feedback.

Primary English References

ENGLISH

Official Website

<https://amuletplus.co.jp>

Japanese-first corporate site; English reference materials are available.

ENGLISH

EVAE Framework / Patent Overview - English

<https://amuletplus.co.jp/evae-framework-patent-en/>

English overview page for the patent-pending EVAE architecture.

ENGLISH

Conscious Loop Demo - English

<https://evae-conscious-loop.vercel.app>

Public demo illustrating EVAE as a decision architecture and loop structure.

ENGLISH

IT / DevOps Governance Demo

<https://evae-financial-demo.vercel.app/it-devops-governance-demo>

Pre-execution governance flow for IT / DevOps decisions.

Supplemental Demo

JAPANESE-FIRST

Financial Governance Demo - Japanese-first

<https://evae-financial-demo.vercel.app/dashboard/financial>

Supplemental demo. Browser translation is recommended. It shows separation of AI output from execution decisions, a policy gate, and decision trace before action.

Note: EVAE is not proposed as a replacement for existing AI models, laws, or standards. It is a model-agnostic pre-execution governance architecture for decision boundaries, human review, and structured audit traces.

Financial Governance Demo Guide

This appendix adds the Japanese-first Financial Governance Demo guide to support readers who want to understand what the demo is intended to demonstrate. Browser translation is recommended where needed.

Why this Japanese-first appendix is included

The Financial Governance Demo is a concrete feasibility example of pre-execution governance. It demonstrates how EVAE separates AI-generated considerations from operational execution decisions, applies a policy gate, and preserves a decision trace before action.

The demo is not a financial advice tool and does not approve or reject loans. It is a conceptual demonstration for governance design.

<https://evae-financial-demo.vercel.app/dashboard/financial>

How to read the following guide pages

- Open the Financial Governance Demo link above.
- Use browser translation if needed.
- Compare the screen flow with the guide pages that follow.
- Focus on the separation between AI output and execution governance.

Financial Demo 操作ガイド

AIに融資判断を委ねるデモではありません。

AIが提示した内容を、**どの条件で次の手続きへ進めるのか**、そして判断の経路をどう記録するのかを体験するデモです。

このデモが示すこと

AIの「答え」と、業務上の「実行判断」を分離し、実行前に判断条件・責任・記録を構造化する。

対象シナリオ：仮想銀行A / 住宅購入資金の事前検討

本デモは実運用ではなく、融資の承認・否決を決定しません。

<https://evae-financial-demo.vercel.app/dashboard/financial>



DEMOを開く

操作方法

1

基本情報を入力する

「E - Origin | 人間の判断の起点」に、検討の前提となる情報を入力します。

- 年齢
- 年収
- 自己資金
- 申込金額
- 職業
- 家族構成
- 他の借入

2

確認サインを入れる

次の2項目を確認し、チェックを入れます。

- 入力内容が事実であることを確認しました
- 本デモは融資可否を決定しないことを理解しています

3

検討プロセスを表示する

入力と2つの確認が完了したら、「**検討プロセスを表示する →**」を押します。V・A・3の各領域に、検討論点・ゲート結果・トレースが表示されます。

4

Action Loopへ進む

表示内容を確認した後、「**Continue to Action Loop →**」を押します。プレゼンで各段階を確認できるよう、画面遷移は手動になっています。

確認情報は、判断経路を再現するための「3 - Trace」に含める想定です。

画面の見方

E

Origin | 人間の判断の起点

人間が、目的と前提情報を入力します。本デモでは、人は意思の起点のみを与えます。

V

Vision | AIが提示する論点

AIは融資可否を決めず、人間が検討すべき論点や可能性を整理して提示します。

A

Choice | 判断方法の確定

定義されたルールに基づき、Policy Gateの結果が表示されます。これは融資の承認・否決ではなく、次の手続きへ進める状態かどうかの確認です。

PASS：条件を満たす

HOLD：確認・情報が必要

E

Trace | 判断理由と責任の記録

入力、確認、検討論点、適用ルール、Gate結果を、再生可能な判断記録として表示します。後から生成する「説明」ではなく、判断に至った経路を残す考え方です。

→

Action Loop | 実行・記録・適応

実行、保留、停止、記録、結果確認、次の改善へ接続します。EVA3は出力で終わらず、実行後の結果を次の判断へ戻します。

このデモで確認してほしい3つの点

01 AIの回答と実行判断が分離されている

AIが論点を提示しても、それだけで業務上の実行へ進みません。生成と実行の間に判断構造を置きます。

02 条件不足をHOLDできる

必要な情報・証拠・権限が揃っていない場合、曖昧なまま次の手続きへ進めず、確認のために止められます。

03 判断の経路をTraceとして残せる

誰が、何を入力し、どのルールが適用され、なぜその状態になったのかを後から確認できる形にします。

一般的なAIは「**答えを生成する**」。

EVA3は「**その答えを、実行可能な決定として通してよいか**」を構造化する。

Amuletplus G.K.

EVA3 Creator / Hiro Yokoki

Financial Demo - Design-by-Transparency



DEMOを開く