**PUBLIC ABSTRACT**
**THE ADMISSIBILITY GAP IN AI DECISION SYSTEMS**

AI is now embedded across **enterprise decision surfaces**: executive briefings, board materials, finance and risk workflows, productivity platforms, and increasingly, agent-driven systems that recommend or initiate actions. Even as model quality, grounding, and transparency improve, organizations still lack a deterministic mechanism to decide **when AI-derived information is legitimate to rely on** in decision-grade contexts, and when downstream workflows must be prevented from proceeding.

This is not primarily a model capability problem.
It is an internal control and liability boundary problem.

**THE PROBLEM**
Modern AI governance has advanced across:

- Safety controls and evaluation
- Explainability and transparency
- Data provenance and monitoring
- Human review checkpoints
- Post-hoc audit and logging

These are necessary, but they do not address a specific failure mode that emerges at decision altitude:

**Intelligence silently becoming authority.**

In practice, AI outputs are often treated as "helpful information," but under time pressure they become decision justifications. In agentic workflows, they can become execution triggers. Human review alone does not reliably prevent this; review often becomes rubber-stamping, and auditability after impact does not prevent illegitimate reliance before impact.

The missing control is not "better answers."
The missing control is permission.

Specifically, organizations typically lack a deterministic boundary that answers these two questions in an enforceable way:

1. Is this AI-derived information legitimate to rely on for action in this declared context, at this authority level, within this time window?

2. If it is not legitimate to rely on, what prevents downstream workflows—human or automated—from proceeding anyway?

This gap increases exposure to:

- Illegitimate reliance under time pressure
- Authority laundering across organizational layers
- Board-level accountability without decision-time control
- Inconsistent delegation in agentic systems
- Post-incident defensibility failures ("why was this relied upon at all?")

**THE SOLUTION CATEGORY**

This work defines a **governance control primitive** that introduces an explicit phase boundary between:

AI-derived intelligence → human/institutional authority → execution.

It consists of two separable layers:

1. **Knowledge Accessibility Layer (KAL)**
   A decision-time admissibility system that determines whether AI-derived information is legitimate to rely on for a declared context, scope, and authority level. KAL is refusal-first, produces finite outcomes, and issues audit-grade governance artifacts. It does not advise, recommend, explain, or authorize action.
2. **Liminal Control Plane (LCP)**
   An execution-time enforcement layer that ensures no workflow, agent, or automation proceeds without an explicit legitimacy signal from an upstream authority. LCP is authority-agnostic and enforces allow/block semantics deterministically. It does not interpret, optimize, or decide.

Together, these layers prevent silent escalation from "suggestion" to "action" without formal authorization.

**STATUS**

A complete reference architecture, governance specification, conformance regime, and deterministic reference execution exist. The system is intentionally non-productized and designed to be implementable within existing enterprise infrastructure by a competent engineering organization.

## INTENDED USE

This work is intended for legal, risk, audit, and governance leadership; platform owners operating AI-driven decision systems; and organizations deploying agentic workflows where reliance and actuation carry material risk.

It is not a consumer product, advisory service, or AI model.

## ACCESS

Detailed specifications and the reference execution are available for controlled diligence and governance evaluation.

BB-AGPA-2026
Issued: February 2026
Prepared as an independent analysis. Sources available upon request.