RCA RELIANCE CONTROL ARCHITECTURE

The Admissibility Gap

The missing boundary between AI output and institutional action

THE QUESTION THAT MATTERS

Is this claim legitimate to rely on for this decision, at this authority level, in this context, right now?

Output is not reliance

AI output can be useful long before it is legitimate to rely on. The risk begins when useful information quietly becomes institutional justification.

RCA

The missing boundary - Reliance Control Architecture (RCA)

AI is entering institutional decision environments.

AI is no longer confined to drafts, summaries, and low-risk suggestions. It is moving into executive briefings, board materials, legal analysis, financial review, risk workflows, operational planning, product decisions, customer-facing systems, and agentic workflows.

As that happens, a governance question becomes harder to avoid.

CENTRAL QUESTION

When is AI-derived information legitimate to rely on?

In this brief, admissibility means fitness for institutional reliance. It does not refer to legal evidentiary admissibility and does not state legal conclusions.

THE DECISIVE QUESTION IS NARROWER THAN MODEL QUALITY

Fluency, accuracy, citation, human review, policy compliance, and audit logging all matter. They are necessary parts of an institutional AI stack. But they do not fully answer whether a particular claim may legitimately justify a particular decision at a particular moment.

RCA

Output is not reliance

AI output can be useful long before it is legitimate to rely on.

A model may summarize a document accurately, identify a plausible risk, compare alternatives, draft a memo, or surface a pattern that deserves attention. None of that automatically makes the output decision-grade.

The risk appears when useful information quietly becomes justification. Under time pressure, a generated claim starts to anchor a recommendation. Human review becomes implied endorsement. A meeting consensus starts to look like ownership. A workflow advances because no one stopped it.

Reliance is not something a model can assume on behalf of an institution. Reliance is an institutional act.

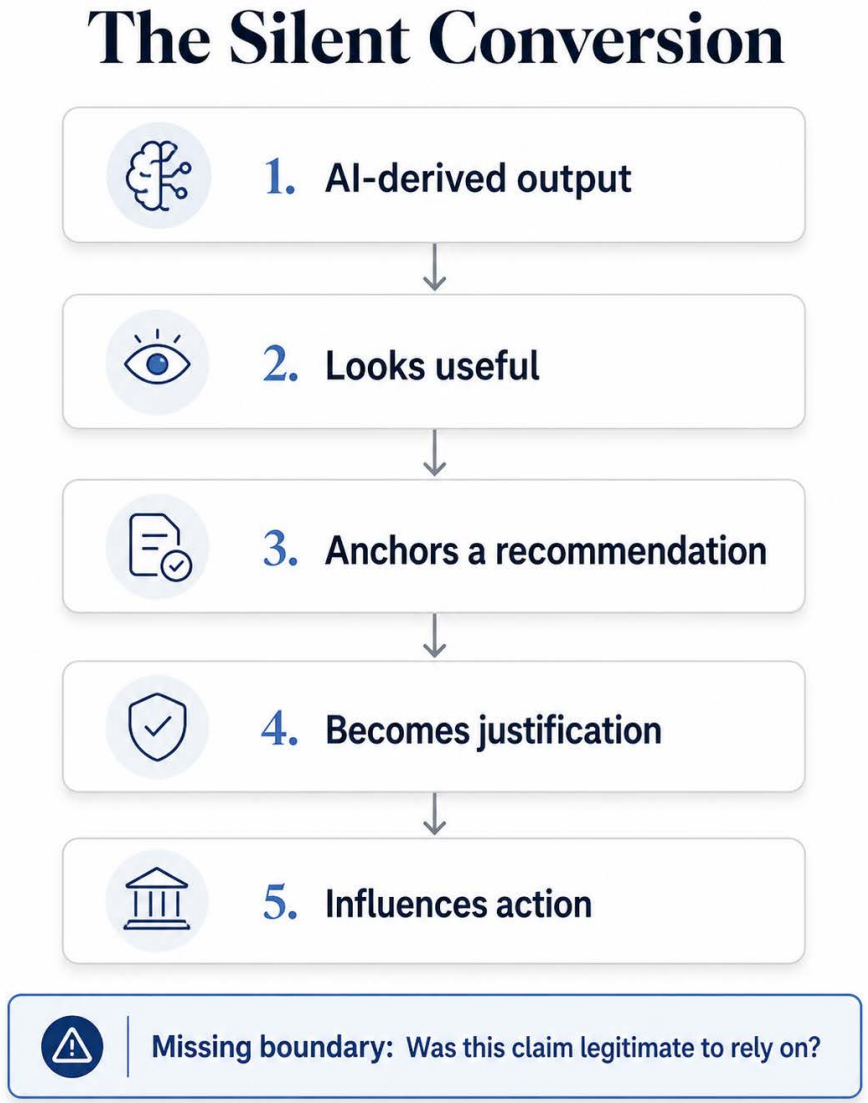

RCA

Why adjacent controls still stop short

The AI governance stack is improving quickly. Institutions are investing in model evaluation, provenance, observability, policy enforcement, security controls, agent governance, workflow approvals, audit logging, and human review. Each layer matters. None should be dismissed.

Provenance

Shows where information came from; it does not decide whether a claim may be relied on for a specific institutional use.

Model evaluation

Improves confidence in system behavior; it does not determine whether a particular output is legitimate to rely on for a particular decision.

Human review

Adds judgment and process; it does not automatically establish that a named authority owns reliance.

Runtime governance

Controls what systems, workflows, or agents may do; it does not necessarily decide whether the claim driving the action was legitimate to rely on.

Audit logging

Can reconstruct events after impact; it does not necessarily prevent illegitimate reliance before action.

Reliance boundary

Asks the missing question: may this claim legitimately justify this decision?

Adjacent controls can create the appearance that reliance has already been governed while the central question remains unanswered.

RCA

The failure is not always the error.

Sometimes the deeper failure is that no one can reconstruct why the output was allowed to matter.

Better models may reduce some errors, but they also increase the risk that fluent outputs become trusted too quickly.

The reconstruction problem

- Who relied on the claim?
- What authority did they have?
- What was the decision context?
- Was the claim bounded in time?
- Was the evidence sufficient for that use?
- Was the claim allowed to support action, or only discussion?

A governance system that only asks whether an output was plausible or traceable may still miss the institutional question.

RCA

Reliance needs its own boundary

A mature institutional AI stack needs a boundary between information and authority.

That boundary should be explicit, context-bound, authority-bound, time-bound, reviewable, and capable of refusal. It should not depend on vague trust language, confidence theater, or the assumption that a human somewhere in the process absorbed responsibility.

The question is not simply whether the system produced a better answer. The question is whether the institution created a legitimate basis for relying on that answer.

NOT THE SAME STATE

Explore
A claim may be considered.

Discuss
A claim may inform discussion.

Analyze
A claim may support analysis.

Justify
A claim may justify a decision.

Act
A claim may support downstream action.

RCA

The Reliance Gradient

Each step requires a different reliance posture.

Each step requires a different reliance posture.

RCA

What changes when admissibility is explicit

When admissibility is explicit, the organization gains a cleaner separation between thinking, relying, and acting.

Assistance remains useful

AI can still assist analysis, drafting, comparison, simulation, search, and planning.

Human authority is protected

Humans can still exercise judgment, but reliance is no longer silently absorbed by process momentum.

Refusal becomes governance

Refusal is treated as a valid outcome, not an inconvenience to route around.

Execution does not inherit legitimacy

Downstream action does not become legitimate merely because a claim sounded convincing.

Audit gets a clearer object

The audit question becomes not only what happened, but whether the claim was admissible before it influenced action.

Momentum slows at the right moment

The goal is not friction for its own sake. It is visibility before information becomes action.

RCA

Why this matters now

Institutions are already building controls for what AI systems may generate, what agents may do, what policies must be followed, what data may be accessed, and what logs must be retained.

Those controls will continue to improve. But as AI systems become more fluent, more integrated, and more action-capable, the reliance boundary becomes more important, not less.

The risk is not only that an AI system produces a bad answer. The risk is that a plausible answer becomes institutional justification before anyone has established that it was legitimate to rely on.

THE GOVERNANCE FRONTIER

The next governance frontier is deciding when AI-derived claims are allowed to matter.

Closing the gap is not a rejection of AI.

It is one condition for serious institutional use.

The Admissibility Gap: the missing boundary between useful AI output and legitimate institutional reliance.

RCA

Reliance Control Architecture

Reliance Control Architecture is a governance architecture for this boundary. It separates the question "May this AI-derived claim be relied on?" from the question "May downstream action be executed?"

That distinction matters because permission to rely is not permission to execute. Before AI-derived information influences decisions, workflows, agents, or real-world action, institutions need a way to make reliance explicit, bounded, reviewable, and capable of refusal.

CONTROLLED REVIEW

Controlled review inquiries may be directed through the RCA founder review process.

brian@bosworthcapitalgroup.com

The question is not whether AI will continue to improve.
It is whether institutions will govern when AI-derived claims are allowed to matter.

RCA

Serious AI use requires a visible boundary between useful information and institutional authority.

Reliance Control Architecture

Public Distribution Artifact v0.4

This public artifact introduces the Admissibility Gap at a conceptual level. It does not disclose RCA implementation mechanics, legal conclusions, regulatory determinations, or buyer-specific materials.