

The Market Timing Is Uncertain. The Control Requirement Is Not.

Why effect-bearing AI creates a structural requirement for pre-execution authorization.

An investor brief on runtime authorization as an enterprise control category, and on FERZ, Inc.'s formalization of Deterministic AI Governance as the engineering discipline behind it.

Governance is enforced at the authorization boundary: the point where an output becomes an effect-bearing action, and where evidence, not explanations, must exist.

PUBLISHER

FERZ, Inc.

VERSION

1.3 · July 2026

CONTACT

info@ferz.ai

FOR

Investors · Analysts · Strategic Partners

THE PROBLEM

Effect-Bearing AI Changes the Governance Problem

AI systems are moving beyond generating text, scores, and recommendations. They are beginning to initiate transactions, modify records, invoke infrastructure, communicate with external parties, and control operational processes.

That transition changes the governance problem. When an AI system only recommends an action, a human or conventional application may remain the authorized actor. When the AI system can cause the action directly, the organization must determine whether that action is permitted before it reaches the operational environment.

Monitoring does not answer that question. Monitoring can show that an action occurred. It may record the request, the model output, the tool call, the transaction, or the resulting system state. That evidence can support investigation and remediation. It cannot establish that the action was authorized before execution.

A post-execution record cannot satisfy a pre-execution authorization requirement.

THE STRUCTURAL RESULT

The Runtime Authorization Boundary

Assume that an AI system proposes an effect-bearing action, that organizational policy requires authorization before the action executes, and that the organization must later demonstrate the basis for the authorization decision. Under those conditions, the system requires a control point in the execution path that:

1. evaluates the proposed action before release;
2. renders a determinate authorization verdict;
3. blocks execution when authorization cannot be established; and
4. produces an authorization artifact sufficient for an independent third party to reconstruct the verdict.

This is the runtime authorization boundary. It is not an observability overlay placed beside the system. It is a non-bypassable enforcement boundary placed between the AI system and the operational environment.

THE CONTROL PATH

AI proposes action → authorization verdict → action released or blocked.

If the boundary is unavailable, the applicable policy state cannot be established, or the authorization artifact cannot be completed, the action does not proceed. That is fail-closed governance.

THE CHOICE

The Autonomy Trilemma

An organization deploying effect-bearing AI has three basic choices.

1. Restrict autonomy

Keep the AI system advisory. Require an authorized human or conventional application to initiate every consequential action.

2. Authorize autonomy

Allow the AI system to propose effect-bearing actions, but require every action to pass through a non-bypassable runtime authorization boundary before execution.

3. Accept unauthorizable execution

Permit the AI system to act without an independently reconstructable pre-execution verdict, relying on monitoring, investigation, and remediation after the event.

The third position may remain commercially acceptable for low-consequence uses. It becomes progressively harder to defend as action volume, autonomy, irreversibility, regulatory exposure, and operational impact increase.

THE ARCHITECTURAL PRESSURE IS CONDITIONAL BUT STRUCTURAL

An organization does not have to deploy autonomous AI. But once it permits autonomous, effect-bearing action subject to an ex-ante authorization requirement, observability alone cannot satisfy that requirement.

THE COMPANY

FERZ Is Building the Authorization Boundary

Governance is the market category. Authorization is the enforcement mechanism. FERZ builds runtime authorization infrastructure, the enforcement layer of AI governance, for AI systems operating in regulated and high-consequence environments.

The governed AI system may remain probabilistic. FERZ does not attempt to make model cognition deterministic. The determinate requirement exists at the point of execution.

Before a consequential action is released, the runtime authorization boundary must produce one of three verdicts:

ALLOW	The action is authorized and may proceed.
DENY	The action is not authorized and remains blocked.
ABSTAIN	Authorization cannot be resolved automatically. Execution remains blocked pending authorized human override.

For each verdict, the system is designed to produce a tamper-evident authorization artifact containing the evidence required to reconstruct how the verdict was reached.

The intended result is not a better dashboard. It is an independently reconstructable verdict rendered before execution.

Observability Versus Authorization

OBSERVABILITY	AUTHORIZATION
Records what occurred	Determines what may occur
Usually retrospective	Pre-execution
Supports investigation	Controls release
Can exist outside the execution path	Controls release at the execution boundary
May detect a violation	Blocks actions that lack authorization
Produces activity records	Produces an authorization artifact

Observability remains necessary. It is not sufficient for authorization. FERZ does not replace monitoring, logging, model evaluation, cybersecurity, or human governance. It supplies the enforcement boundary those functions cannot provide on their own.

TECHNICAL FOUNDATION

The Five Tests Standard

FERZ's published Five Tests Standard (5TS) provides five vendor-neutral tests for whether an AI governance control reaches the authorization threshold:

Stop	The action can be stopped before it executes when policy does not permit it.
Ownership	Authority for the action is established.
Replay	The verdict can be independently reconstructed.
Escalation	The action escalates at a policy boundary.
Provenance	The inputs grounding the verdict have an established origin. Origin, not truth.

The standard separates governance claims from governance evidence. A dashboard, policy document, model card, or audit log may be valuable without satisfying these tests.

PUBLISHED STANDARD

Five Tests Standard (5TS) · DOI: 10.5281/zenodo.21040295 · Foundational paper:
10.5281/zenodo.21048661

THE OPPORTUNITY

The Investment Thesis

The formal proposition is limited: observability cannot, by itself, supply pre-execution authorization.

The market proposition is broader: as AI systems receive authority to take consequential actions, regulated enterprises will require controls that can determine and prove whether those actions were authorized before execution.

The FERZ investment thesis is that:

- effect-bearing AI autonomy will continue to expand;
- the authorization deficit will become a recognized enterprise control category;
- regulated buyers will require independently reconstructable authorization evidence, not only monitoring evidence;
- runtime authorization will become a distinct infrastructure layer; and
- FERZ's position in that layer is already on the public record: the published Five Tests Standard, a dated research corpus including the formal impossibility result for observability-based authorization, priority-dated patent filings, defensive publications creating dated public disclosures, and formal submissions into federal AI governance proceedings.

These are market assumptions, not formal results.

THE DISTINCTION

What Is Inevitable, and What Remains a Market Question

Structurally required under the stated conditions

- a verdict before execution;
- a boundary that controls release;
- fail-closed behavior when authorization cannot be established;
- evidence sufficient to reconstruct the verdict.

Market timing remains uncertain

- when buyers recognize the requirement;
- which regulated sectors adopt first;
- how quickly procurement language catches up with the engineering requirement;
- how rapidly FERZ converts its category lead into commercial adoption.

THE POSITION

The Wager

FERZ is not betting that probabilistic AI will disappear. FERZ is betting that enterprises will eventually distinguish between observing an AI action and authorizing it.

FERZ has already formalized that distinction as Deterministic AI Governance: a defined engineering category grounded in pre-execution authorization, fail-closed governance, and independently reconstructable verdicts.

Nor is FERZ waiting for that recognition to arrive on its own. It has carried the distinction into federal AI governance and procurement processes through feedback on the NIST AI RMF Trustworthy AI in Critical Infrastructure Profile, responses to requests for information from the FDA and the Department of Transportation, and a public comment on GSA's proposed safeguarding clause for LLM systems. FERZ is contributing to the record where that language is being shaped.

When that distinction becomes a procurement requirement, FERZ will not be entering a category defined by others. It will be commercializing the category it already formalized through its published standard, dated research corpus, public doctrine, and priority-dated filings. Those dates and that sequence cannot be recreated retroactively.

The market timing remains uncertain. The control requirement does not.

NEXT STEP

FERZ, Inc. is raising a seed round to convert this position into deployed enforcement infrastructure. The investor deck and terms are available on request. Technical and architectural detail is available under NDA. Contact: info@ferz.ai.