

THE AGENTIC AI GOVERNANCE PLAYBOOK

GARTNER · CFO GUIDANCE

Agentic AI Governance Before ROI: *Runtime Controls for Enterprise AI Agents*

Governance readiness is becoming production readiness.

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

GARTNER · CFO GUIDANCE

Agentic AI Governance Before ROI: *Runtime Controls for Enterprise AI Agents*

Gartner's governance-first advice to CFOs points to a broader enterprise shift: autonomous systems need an authority boundary before they become an efficiency engine.

MAUREEN DOYLE-SPARE

Governance readiness is becoming production readiness. For AI agents, traceability and human oversight matter, but enterprise AI governance also needs a runtime authority boundary over the meaning that drives execution.

Enterprise AI governance is arriving before enterprise AI returns. Gartner's August 20 guidance[1] says the first finance AI agent should be treated as a governance pilot rather than an ROI pilot. The point is not that return on investment no longer matters. It is that agentic AI introduces a different control problem: an autonomous agent can interpret an objective, plan a sequence, access multiple systems and execute steps without a person approving each intermediate decision. Governance of autonomous decision systems therefore has to be proven before autonomy is scaled.

For finance leaders, that is a practical change in sequence. The institution first defines the AI controls, decision boundaries, human oversight, traceability and failure evidence that make autonomous execution governable. ROI then becomes a question about activity the institution is prepared to authorize, not simply how much work an AI agent can perform.

Governance readiness is becoming production readiness for AI agents

Alex Levine, Director Analyst in the Gartner Finance practice, frames the test plainly: "Success should be measured by governance readiness, not just autonomy or ROI." Gartner recommends starting with a low-risk, contained workflow where errors are visible and reversible, defining data access and allowed actions before development, assigning ownership across finance, technology and audit or risk, sandboxing the pilot, and preserving complete traceability. The finance team should be able to reconstruct what the agent planned, accessed and produced for every run, while maintaining a comprehensive failure log.

Those controls are important because AI workflow

© 2026 Maureen Doyle-Spare.

governance now extends beyond a model and its output. Agentic systems operate through AI orchestration: they retrieve context, use tools, coordinate steps and can hand state into downstream systems. In enterprise AI architecture, the unit of risk increasingly becomes the runtime decision and the trajectory it initiates.

Finance makes this visible because authority is already formalized through approval limits, maker-checker controls, reconciliations, treasury controls, financial reporting and audit evidence. The same pattern extends to healthcare, life sciences, insurance, manufacturing, critical infrastructure and government. Autonomous agents move consequential interpretation into runtime, so runtime governance has to move there with it.

Why traceability alone does not answer the reasoning-layer governance question

Being able to reconstruct what an agent planned, accessed and produced is necessary. It does not establish that the meaning under which the agent acted was institutionally authorized. An agent can use approved data, call approved tools and produce a technically valid result while resolving a business term under an Operational Interpretation[2] the institution never sanctioned.

This is the reasoning-layer governance problem. Terms such as approved, eligible, material, cleared, complete and within tolerance often draw their operational meaning from several policies and enterprise systems. Much of that institutional knowledge is distributed rather than centrally held, so an autonomous workflow may reconcile individually valid sources into a working interpretation that differs from institutional meaning without creating a conventional model-risk, access-control or output-quality signal. That condition is Invisible Failure[2] in autonomous workflows: every adjacent control reports correctly while the institution is still wrong.

That is why responsible AI architecture for autonomous systems needs an explicit runtime control surface. The institution maintains a Reasoning Baseline[2] for the decision class, exposes the Operational Interpretation that will drive the pending action, measures semantic divergence and applies an enforceable policy before execution. This is semantic integrity as an enterprise control problem rather than a documentation exercise, and it is the substance of a semantic-layer control framework.

The control boundary belongs before execution: deterministic gating in practice

The Semantic Control Plane[2] is the runtime governance architecture for that boundary, and it functions as a control-plane architecture for autonomous agents rather than a monitoring layer. The Ontology provides the authoritative resolution backbone for the Reasoning Baseline. The Semantic Layer extracts the Operational Interpretation. The Semantic Deviation Index[3] provides semantic deviation quantification. The Deterministic Gate[2] converts the measured state and institutional threshold into an enforceable disposition: Permit, Permit-with-Flag, Hold or Mandatory Hold. The RSSR then preserves the per-execution governance record.

This sequence is pre-execution assurance, or semantic-layer validation before execution. It is not another dashboard and it is not a replacement for cybersecurity, identity controls, model evaluation, explainable AI, AI compliance or human review. Those controls govern adjacent surfaces. Runtime oversight and deterministic gating govern whether this interpretation, for this decision, carries institutional authority to proceed before execution authority is emitted.

© 2026 Maureen Doyle-Spare.



Figure 1. *The Pre-Execution Assurance Protocol. Source: Doyle-Spare (2026).*

How runtime governance changes the economics of agentic AI

A governance-first pilot is not separate from the ROI discussion. It determines which reasoning and execution the enterprise is prepared to admit. Agentic AI can retrieve, plan, call tools, verify, retry and delegate before a business action is complete. Lower model prices can make each step cheaper while total autonomous activity continues to grow.

The Token Economics research, The Economics of Agentic AI: Runtime Governance as a Distinct Determinant of Enterprise AI Cost[4] (Doyle-Spare, 2026), defines Preventable Computation as execution-dependent reasoning, tool use and downstream organizational activity that would not occur if an institutionally available runtime control held an unauthorized or materially divergent trajectory at the earliest governable checkpoint. Token Liability is the inference-cost component of that larger exposure. The economic point is not that every long workflow is wasteful. Authorized reasoning is the cost of doing useful work. Governance changes whether additional work should be admitted at all.

This distinction is important for CFOs because AI risk management and AI economics converge at the runtime authority boundary. Model routing, caching and smaller models reduce the price of computation that proceeds. Runtime governance primitives determine whether execution-dependent computation should proceed in the first place. Governance does not guarantee savings, but it changes the cost surface the institution is exposed to.

Finance is the visible test case for enterprise AI governance, not the boundary

Finance is a useful proving ground because authority, evidence and accountability are already explicit. A close process, reconciliation workflow, forecasting process or internal reporting agent sits inside established control

© 2026 Maureen Doyle-Spare.

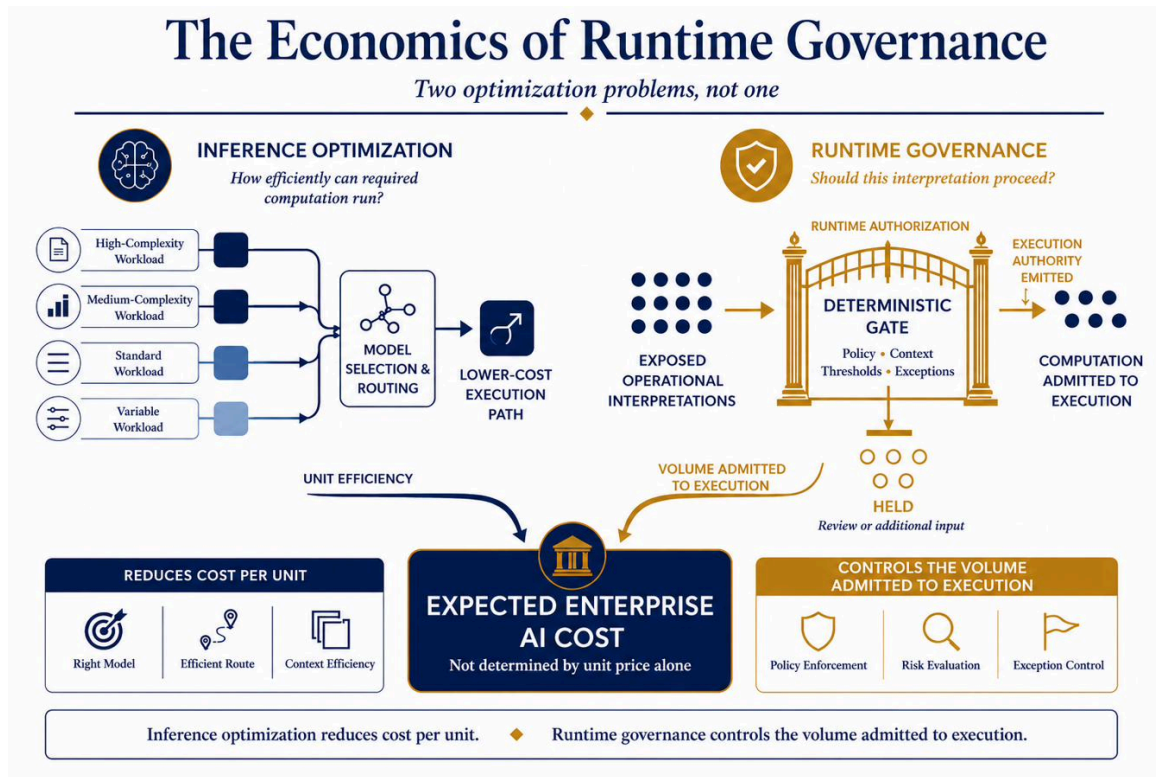


Figure 2. *The Economics of Runtime Governance. Source: Doyle-Spare (2026).*

ownership. That makes governance gaps easier to see and makes a contained first pilot sensible.

The same enterprise AI governance principle applies when a clinical agent interprets eligibility, an insurer interprets coverage, a manufacturer interprets a release condition, a cyber agent interprets evidence, or a multi-agent system carries upstream state into a downstream decision. That last case is why governance architecture for multi-agent systems matters: authority created in one environment can become executable state in another. AI agents can be technically authorized while the semantic layer has moved away from institutional meaning. The governing question is whether the autonomous system is still operating under the authority the institution actually intended to give it.

The risk is not only accidental. The Agentic Workflow Drift and Agentic Workflow Subversion taxonomy[5] distinguishes divergence that emerges through ordinary coordination from divergence that is deliberately induced. Separately, the Semantic Layer Integrity Attack[6] identifies a cyber threat class in which an adversary targets the reasoning-layer attack surface to alter the semantic conditions under which an agent acts. Reasoning-layer drift measurement is relevant across both accidental and adversarial divergence.

The first pilot should leave a governance asset behind

A useful agentic AI governance pilot should produce more than confidence in one agent. It should leave reusable decision-class boundaries, ownership rules, Reasoning Baselines, traceability requirements, review triggers, failure knowledge and evidence standards that can be carried into the next deployment. That is how governance readiness becomes scalable enterprise AI architecture rather than one-off oversight.

© 2026 Maureen Doyle-Spare.

The broader shift is clear. Agentic AI moves the unit of governance from the model and the output toward the runtime decision. An agentic-runtime governance architecture therefore has to govern whether an Operational Interpretation carries institutional authority to proceed before execution authority is emitted. The scalable autonomous agent is not simply the one that performs well. It is the one the institution can permit, hold, reconstruct and govern before its actions become the institution's actions. ◆

References

1 *Gartner. Gartner Says CFOs Must Pilot Governance First Before Scaling AI Agents. August 20, 2026.* gartner.com/en/newsroom/new/pr-q-a-new-vis-template/2026-08-20-gartner-says-cfos-mus-pilot-governance-first-before-scaling-ai-agents

2 *Doyle-Spare, Maureen. Agentic AI Systems Governance: A Runtime Reference Architecture for the Reasoning Layer and the Semantic Control Plane in Regulated Financial Institutions. Zenodo, 2026.* doi.org/10.5281/zenodo.20749051

3 *Doyle-Spare, Maureen. Agentic AI Drift Measurement and the Semantic Deviation Index. SSRN No. 6531238.* doi.org/10.2139/ssrn.6531238

4 *Doyle-Spare, Maureen. The Economics of Agentic AI: Runtime Governance as a Distinct Determinant of Enterprise AI Cost. July 16, 2026. SSRN No. 7129939.* doi.org/10.2139/ssrn.7129939

5 *Doyle-Spare, Maureen. Agentic Workflow Drift and Agentic Workflow Subversion: A New Risk Taxonomy for the Governance of Agentic AI in Financial Services. SSRN No. 6459612.* doi.org/10.2139/ssrn.6459612

6 *Doyle-Spare, Maureen. Agentic AI Cyber Subversion: The Semantic Layer Integrity Attack as a New Threat Class Against the Reasoning Layer. SSRN No. 6926219.* doi.org/10.2139/ssrn.6926219

Originally published in *The Agentic AI Governance Playbook* on Substack, August 21, 2026. maureendoylespare.substack.com

© 2026 Maureen Doyle-Spare.



The scalable autonomous agent is the one the institution can permit, hold, reconstruct and govern before its actions become the institution's actions.

THE AGENTIC AI GOVERNANCE PLAYBOOK

Download the research

Papers, briefs and reference architecture at maureendoylespare.com

Visit the Playbook

WEEKLY NEWSLETTER

Follow the newsletter

The Agentic AI Governance Playbook Weekly Newsletter

Follow Newsletter

INTELLECTUAL PROPERTY: © 2026 Maureen Doyle-Spare. All rights reserved. No part of this article, including its text, figures, tables, and images, may be reproduced or distributed without the author's prior written permission.