GLOBAL FINTECH FEST · MUMBAI 2026

India's Agentic AI Warning: Authorization Is Now *Too Slow*

Autonomous systems can compress processes from days into seconds.

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

GLOBAL FINTECH FEST · MUMBAI 2026

India’s Agentic AI Warning: Authorization Is Now *Too Slow*

India’s finance minister is warning that autonomous systems can compress processes from days into seconds. Institutional authorization has to operate at the same speed.

MAUREEN DOYLE-SPARE

Agentic AI[1] is forcing institutions to confront a control problem that conventional automation rarely created at the same scale: the interval between interpretation and action is disappearing. At the Global Fintech Fest in Mumbai on September 11, India’s Finance Minister Nirmala Sitharaman warned that agentic systems can move beyond recommendations to direct action, compressing processes that once took days into seconds. She paired that efficiency with a systemic-risk warning: greater speed can also accelerate errors, fraud and market shocks.

The point is especially visible in financial services because payments, fraud responses, customer servicing, trading support and operational decisions already move through tightly controlled infrastructures. Yet the principle reaches much further. A clinical agent can escalate care, an industrial agent can alter a production process, a cyber agent can launch remediation, and a logistics agent can reroute physical activity. In each case, faster autonomous execution reduces the time available for a person or a downstream control to discover that the system resolved the wrong operational meaning.

The governance challenge is no longer simply whether an AI system can produce an accurate recommendation. Institutions need a control architecture that operates inside the compressed interval between what an autonomous system has interpreted and what it is about to do.

Speed changes the control problem

Sitharaman’s comments were delivered at a conference whose 2026 theme explicitly describes agentic AI as enabling financial systems to sense, decide and act in real time. The efficiency case is clear. Long sequences of retrieval, analysis, routing, approval preparation and execution can be compressed into a far shorter autonomous workflow.

The risk is created by the same compression.

© 2026 Maureen Doyle-Spare.

Traditional controls often assume that there is enough time between stages for review, reconciliation or escalation. Agentic systems can collapse those stages. If the system forms an incorrect interpretation early in the sequence, the error can propagate into action before a conventional after-the-fact control has an opportunity to intervene.

That is why speed should be treated as a governance variable, not only a productivity metric. The faster an agent can convert meaning into action, the more important it becomes to place an enforceable boundary immediately before execution.

The missing control object is the Operational Interpretation

The Agentic AI Governance Playbook frames this boundary around the Operational Interpretation: the operational meaning an autonomous system resolves from policy, context, instructions and other authoritative inputs before it acts. The runtime governance architecture[4] treats the institutionally authorized Reasoning Baseline and the runtime Operational Interpretation as distinct governance states.

That distinction matters because an agent can be technically authorized and still reach the wrong operational conclusion. It can use the approved model, access permitted data, invoke an allowed tool and remain inside its identity and permission envelope while misapplying a policy, reconciling conflicting instructions incorrectly or resolving an exception outside the institution's intended boundary.

When execution occurs in seconds, an institution cannot rely on a human reviewer appearing somewhere downstream to repair that semantic error. The interpretation itself has to become a governed event.

Semantic deviation gives the institution a measurable boundary

The Semantic Deviation Index[5] provides a way to evaluate material divergence between the authorized Reasoning Baseline and the Operational Interpretation resolved at runtime. SDI does not attempt to make probabilistic reasoning deterministic. Its role is to make divergence governable before the system is permitted to cross the execution boundary.

That distinction is important in fast-moving environments. Institutions may accept considerable variation in how an agent reasons through low-risk administrative work while imposing a far tighter threshold around high-consequence decisions. The control policy can differentiate those decision classes without requiring every autonomous action to stop for manual approval.

The result is a scalable form of human oversight. Human judgment is concentrated where semantic divergence, decision criticality or institutional policy requires it, instead of being used as a generic label attached to every workflow.

The execution boundary has to operate at machine speed

A useful runtime sequence separates interpretation, resolution, validation, authorization and execution. That separation prevents technical permission from being mistaken for institutional approval. The agent can prepare an action, but execution authority is formed only after the institution has evaluated whether the meaning driving that action remains within the authorized boundary.

This is where the Deterministic Gate becomes important. The Gate applies an institutionally defined

© 2026 Maureen Doyle-Spare.

threshold policy after semantic divergence has been evaluated and before Authorization Emission creates downstream Execution Authority. It is a deliberately narrow use of determinism at the moment when probabilistic reasoning is about to create a real-world consequence.

The architecture does not depend on access to hidden chain-of-thought. It governs observable institutional states: which Baseline applied, which Operational Interpretation was resolved, what divergence was measured, which authorization outcome was issued and what action followed.

Pre-execution assurance is the answer to compressed decision time

The five-step Pre-Execution Assurance Protocol turns that boundary into an operating sequence: establish the applicable Reasoning Baseline, extract the Operational Interpretation, compute semantic divergence, enforce the Deterministic Gate and generate the runtime governance record.

That sequence becomes more valuable as autonomous systems accelerate. A workflow that once allowed hours for reconciliation may offer seconds. A process that once involved several human handoffs may now proceed end to end through software. The control cannot remain slower than the autonomy it is supposed to govern.

This is the deeper implication of the speed-versus-systemic-risk tension highlighted in Mumbai. Institutions do not need to slow every agentic workflow back to human tempo. They need a governance layer capable of evaluating authorization at machine tempo before consequential execution occurs.

© 2026 Maureen Doyle-Spare.

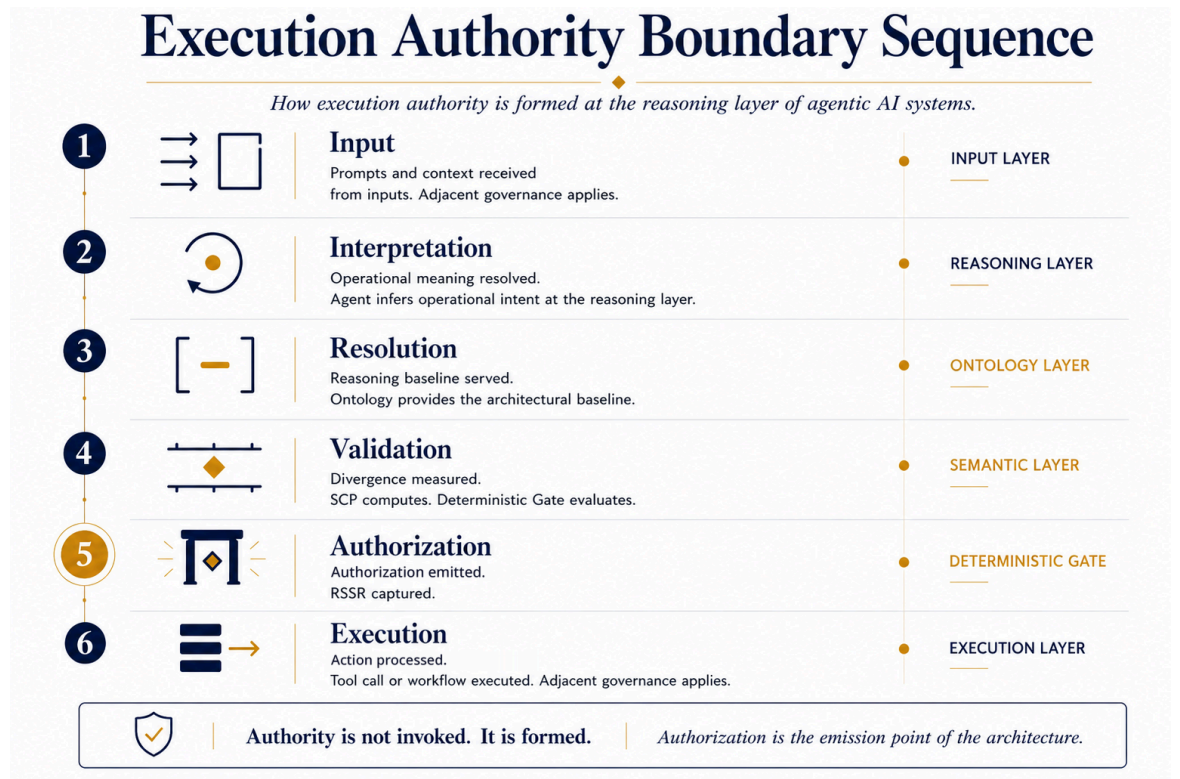


Figure 1. *Execution Authority Boundary Sequence. Source: Doyle-Spare (2026).*

The principle travels well beyond banking

Regulated industries make the problem easy to see because money movement, market activity and customer decisions already have explicit authority boundaries. The same pattern applies across healthcare, life sciences, industrial operations, critical infrastructure, cybersecurity and government. As autonomous systems gain authority to act, each institution needs an authoritative Baseline, an observable runtime interpretation and an enforceable point at which execution can be allowed, held or escalated.

This does not mean every agentic failure is a Semantic Layer Integrity Attack. SLIA is a specific threat class involving deliberate manipulation of the semantic conditions from which an Operational Interpretation is formed. Ordinary error, drift, misconfiguration and adversarial semantic manipulation need to remain distinct because their controls and evidence requirements differ.

The broader agentic AI governance research is built around that separation of failure modes and control objects. Executive summaries and implementation-oriented views are also available through the Agentic AI Governance Executive Briefs.

Autonomy does not remove institutional responsibility

The most important part of Sitharaman's warning is not that agentic AI is too fast. Speed is part of its value. The governance issue is whether institutional control can keep pace when autonomous systems begin resolving and executing decisions without the natural friction that once separated analysis from action.

Strong governance should preserve that efficiency while making the execution boundary explicit. The institution defines the authorized meaning. The agent

© 2026 Maureen Doyle-Spare.



Figure 2. *The Pre-Execution Assurance Protocol. Source: Doyle-Spare (2026).*

resolves an operational interpretation. Divergence is evaluated. Authority is emitted only when the decision remains inside the permitted boundary, with escalation where it does not.

Agentic AI can compress days into seconds. The control architecture now has to compress with it. ◆

References

1 *Global Fintech Fest. "Global Fintech Fest 2026 theme: Agentic AI, Tokenisation and Quantum."* globalfintechfest.com/theme

2 *Financial Express. "Need stronger guardrails against Agentic AI risks: FM." September 11, 2026.* financialexpress.com/life/technology-need-stronger-guardrailsagainst-agentic-ai-risks-fm-4337279

3 *Economic Times. GFF 2026 remarks on Agentic AI compressing processes from days into seconds. September 11, 2026.* economictimes.indiatimes.com/news/economy/policy/gff-2026-fm-nirmala-sitharaman-calls-indias-demographic-dividend-an-active-unstoppable-reality/articleshow/134058104.cms

4 *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

5 *Doyle-Spare, Maureen. Semantic Deviation Index. SSRN No. 6531238, April 2026.* doi.org/10.2139/ssrn.6531238

Originally published in *The Agentic AI Governance Playbook* on Substack, September 13, 2026. maureendoylespare.substack.com

© 2026 Maureen Doyle-Spare.



Agentic AI can compress days into seconds. The control architecture now has to compress with it.

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.