


MORTGAGE LENDING · FANNIE MAE LL-2026-04

# Agentic AI Governance Comes to *Mortgage Lending*

The governance obligation is arriving. Closing the gap means runtime governance for the Reasoning Layer

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

# Agentic AI Governance Comes to *Mortgage Lending*

The governance obligation is arriving. Closing the gap means runtime governance for the Reasoning Layer

MAUREEN DOYLE-SPARE

Fannie Mae's new AI requirements move governance into the operating obligations of lenders and servicers. The harder question begins at runtime.

On August 6, a new AI governance requirement became operational across an important part of the U.S. mortgage market.

Fannie Mae's Lender Letter LL-2026-04[1], published April 8 and effective 120 days later, establishes a governance framework for sellers and servicers using artificial intelligence or machine learning in the origination of loans sold to or guaranteed by Fannie Mae, or in servicing loans on its behalf.

On their face, the requirements are conventional. Firms using AI or machine learning must maintain policies and procedures governing development, implementation, use, maintenance, measurement, and risk management. Those policies must reflect applicable legal requirements and institutional risk tolerance, assign ownership, incorporate trustworthy and ethical AI characteristics, and be reviewed at least annually.

Coverage also extends beyond technology developed inside the institution. Sellers and servicers must govern subcontractor and vendor use of AI or machine learning to a standard no less protective than their own. Fannie Mae[1] can also request information about the types of AI being used, their purpose and manner of use, and the safeguards implemented around them.

That is an important development for AI governance in financial services because it moves the discussion out of general principles and into an operating relationship. For affected mortgage firms, AI governance is no longer simply a developing policy practice. It is part of the conditions under which AI can be used in origination and servicing.

Fannie Mae has since carried the same direction into multifamily lending. Its July Lender Letter 26-05[2] establishes a parallel AI and machine-learning governance framework for multifamily lenders across origination, servicing, and asset-management practices.

Taken together, these developments make the direction fairly clear. Mortgage AI is moving toward explicit institutional accountability.

The more difficult governance problem begins one layer deeper, at the Reasoning Layer.

There, an agentic system interprets what an institutional rule means and commits to that Operational

© 2026 Maureen Doyle-Spare.

Interpretation before it acts. Governing that layer at runtime, before execution, is the problem a runtime reference architecture for the Reasoning Layer[3] sets out to solve.

## Policies govern the system. The decision still happens at runtime.

Fannie Mae's framework appropriately requires governance, risk management, ownership, information security, business resiliency, and third-party controls. None of those requirements should be treated as optional or secondary.

But agentic AI introduces a different operating condition than the one traditional model governance was built around.

A conventional model generally produces an output that enters a downstream decision process. An agentic system can instead reason across multiple sources, resolve how institutional terms apply to a particular case, invoke tools, and move directly toward execution. In that environment, operational authority begins to form during the reasoning process itself.

In mortgage lending, the distinction has direct consequences.

Terms such as income, capacity, collateral value, eligibility, risk rating, exception, approval, hardship, modification, and loss mitigation are not merely data fields. They carry institutionally governed meaning. An autonomous system may have access to valid data, operate through approved tools, satisfy its access permissions, and still resolve one of those terms in a way the institution never authorized.

In the underlying Runtime Reference Architecture, this is the governance surface of Operational Interpretation. Here the problem is the meaning itself: the system formed a runtime interpretation that departed from the institution's authorized Reasoning Baseline before action was taken, even though the model performed correctly and the security controls held. The architecture distinguishes that single-decision condition as Runtime Semantic Divergence, which can accumulate into Agentic Workflow Drift when unresolved interpretations repeatedly execute.

This is particularly relevant to credit underwriting. That architecture already treats consumer, commercial, mortgage, trade finance, and specialized credit as Reasoning Layer environments in which terms such as income, capacity, collateral value, exposure, and approval authority must resolve against institutionally authorized definitions, not an implicit framework the agent assembles.

Fannie Mae's new requirements do not prescribe a solution to that problem. They should not be represented as doing so. But they make the underlying governance obligation much more concrete.

If an institution must be able to explain how AI is being used and what safeguards surround it, the quality of that answer becomes increasingly important as AI moves from assisting a mortgage workflow to coordinating it.

## This is where runtime governance becomes distinct from model governance

The revised U.S. interagency model-risk guidance issued earlier this year is important to this boundary as well. SR 26-2 places generative and agentic AI models outside the direct scope of the model-risk guidance while stating that institutions' broader risk-management and governance practices should guide the control of systems and processes not covered by it.

© 2026 Maureen Doyle-Spare.

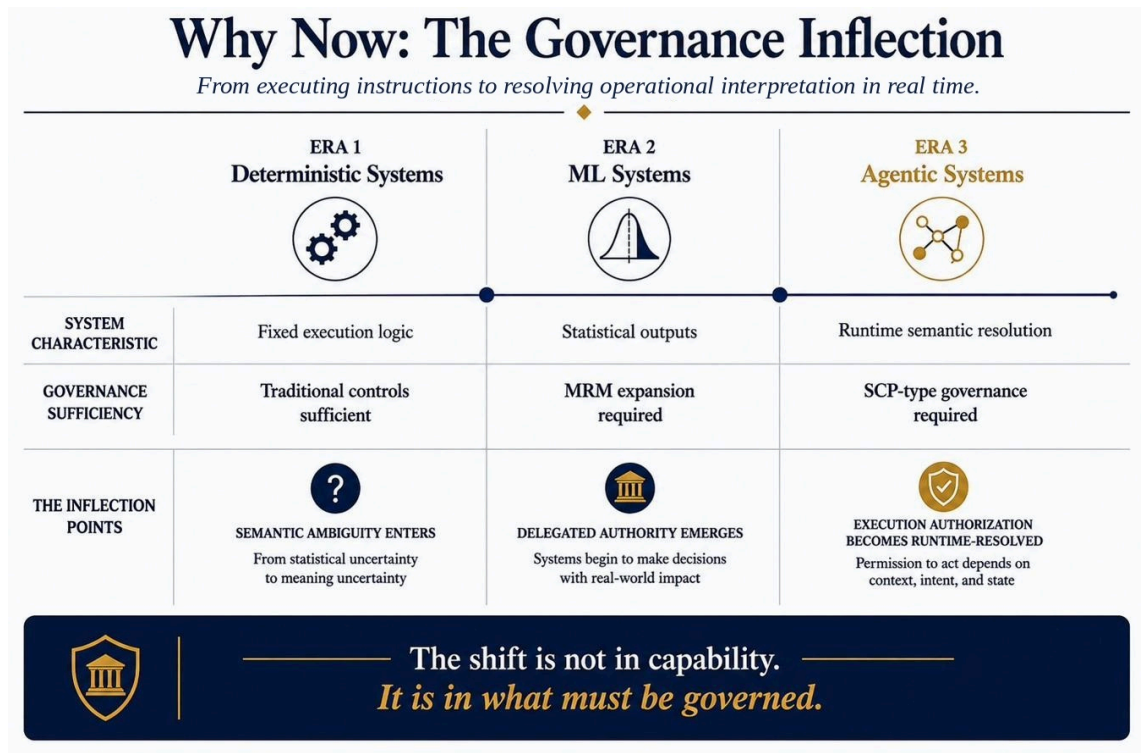



***Figure 1.*** *The Governance Inflection: From Output Governance to Runtime Semantic Governance. Source: Doyle-Spare (2026)[3], Agentic AI Systems Governance (Zenodo).*

A Semantic Control Plane[3] architecture treats that distinction explicitly. It does not replace model risk management, information security, workflow orchestration, or existing lending controls. Those disciplines remain authoritative at their own surfaces. Instead, the SCP operates as an Adjacent Governance Surface at the Reasoning Layer, where an agent resolves the Operational Interpretation under which execution authority will be formed.

It begins with an institutionally authorized Reasoning Baseline. At runtime, the Semantic Layer derives a Runtime Semantic State from the governable representations exposed by the workflow. From there, the Semantic Deviation Index, or SDI, measures divergence between that runtime state and the authorized baseline. A Deterministic Gate then converts the institution's calibrated tolerance into a defined authorization decision before execution authority is emitted.

*This is pre-execution assurance, not post-event monitoring.*

For an underwriting workflow, that difference can be quite practical. An agent may evaluate a borrower through a valid credit model, retrieve the right documents, apply authorized tools, and still substitute a generalized capacity heuristic for the institution's authorized capacity methodology. Those surrounding systems remain valid. The meaning governing the decision does not.

In the architecture's credit-underwriting reference case, that deviation becomes measurable in the Runtime Semantic State, is evaluated through SDI, and is intercepted by the Deterministic Gate before the approval proceeds. That decision state is then preserved through the RSSR and Semantic Audit Trail.

That is the distinction between proving that an AI

© 2026 Maureen Doyle-Spare.

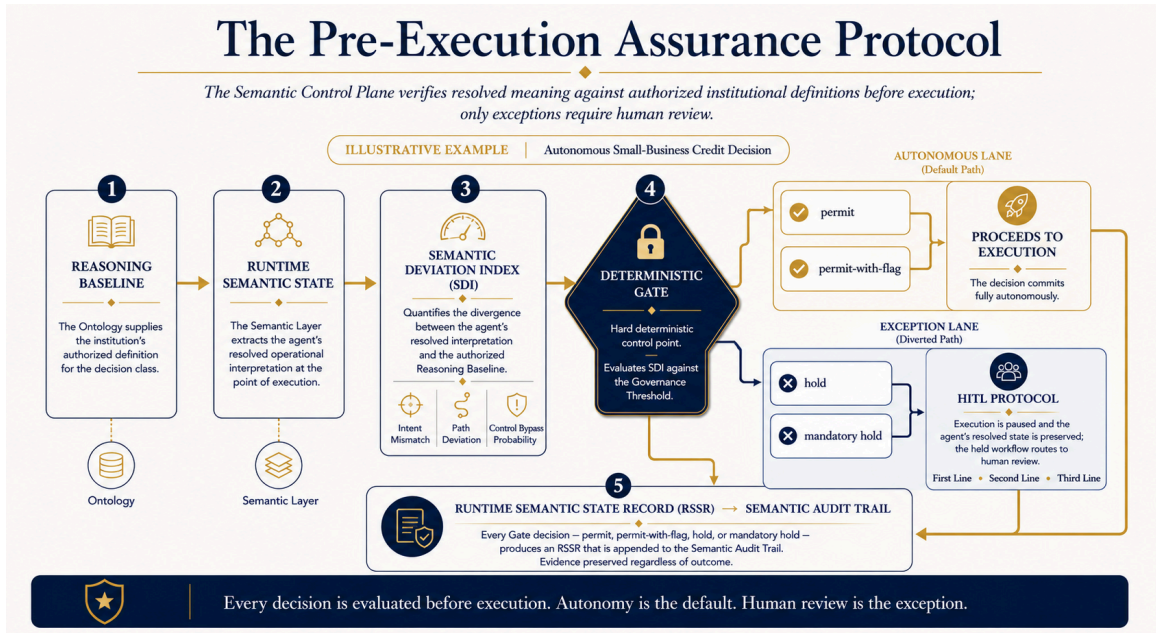



***Figure 2.*** *The Pre-Execution Assurance Protocol, traced through an autonomous credit underwriting decision. Source: Doyle-Spare (2026)[3], Agentic AI Systems Governance (Zenodo).*

system was governed and proving that a particular autonomous decision operated under meaning the institution authorized.

## Fannie Mae's disclosure requirement makes the evidence problem more interesting

There is another part of LL-2026-04 that deserves attention. Fannie Mae can require a seller or servicer to disclose the types of AI or machine learning it uses, the purpose and manner of use, the safeguards surrounding it, and other information Fannie Mae considers necessary.

At a basic level, firms can answer that through inventories, policies, model records, vendor documentation, and control descriptions.

As agentic systems become more autonomous, however, a second evidence class becomes relevant. A supervisor, auditor, risk function, or contractual counterparty may eventually need to reconstruct not just which system was used, but what happened during the particular decision.

Under the Runtime Reference Architecture, that evidence separates into three tiers: V-RSR is the verification activity at the point of decision, the RSSR the per-execution governance reconstruction artifact, and the Semantic Audit Trail the institutional longitudinal record. Together they preserve the Reasoning Baseline in force, the Runtime Semantic State, the SDI result, the Gate decision, and the authorization condition under which the workflow proceeded.

That is a substantially different evidentiary proposition from an AI inventory.

It lets an institution establish that safeguards existed and that they operated against the specific meaning under which a consequential decision was made.

## Third-party AI makes the boundary even more important

Fannie Mae's requirement that vendor and

© 2026 Maureen Doyle-Spare.

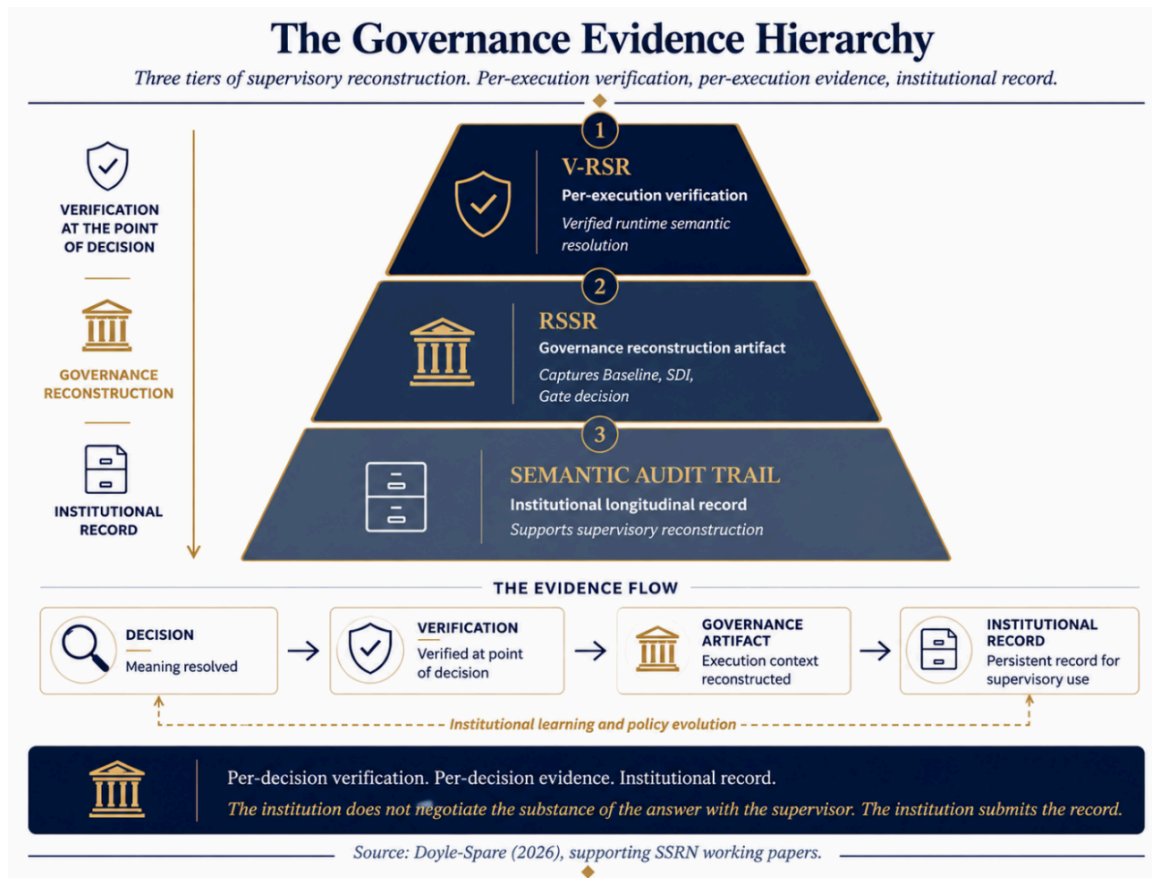



***Figure 3.*** *The Governance Evidence Hierarchy. Source: Doyle-Spare (2026)[3], Agentic AI Systems Governance (Zenodo).*

subcontractor AI governance be no less protective than the institution's own controls may ultimately become one of the more important elements of the framework.

Financial institutions increasingly consume intelligence, models, services, data, and agentic capabilities from outside their own control boundary. A vendor can be approved, authenticated, contractually governed, and technically secure while still supplying context or an interpretation that changes the meaning used by a downstream autonomous workflow.

This is where vendor governance and semantic governance intersect.

An institution remains accountable for the action it authorizes. It needs a control point at which external reasoning components are reconciled against institutionally authorized meaning before their outputs become execution authority.

The architecture does not require the institution to own every model. It requires the institution to retain authority over the Operational Interpretation under which its own systems act.

That distinction will matter far beyond mortgages.

## The direction of travel is becoming clearer

Fannie Mae has not endorsed a Semantic Control Plane, SDI, or any particular runtime architecture. LL-2026-04 is also not a federal banking regulation. It is a governance requirement imposed through Fannie Mae's seller and servicer relationship, and it should be described precisely that way.

Its significance is still considerable.

© 2026 Maureen Doyle-Spare.

AI governance is steadily moving from statements of principle toward expectations that institutions can identify where AI is being used, assign accountability, govern vendors, manage risk, and produce evidence of the safeguards surrounding deployment.

Agentic AI pushes the next part of that evolution into the Reasoning Layer.

An institution will increasingly need to know more than whether the AI system was approved, the model validated, the tool call authorized, and the workflow executed correctly. It will need to establish whether the operational interpretation under which execution authority was formed was itself institutionally authorized.

That is the control boundary that begins to matter when AI stops merely informing the mortgage process and starts acting inside it. ◆

## References


1 *Fannie Mae. "Lender Letter LL-2026-04, Governance Framework on Use of Artificial Intelligence and Machine Learning." April 8, 2026.* singlefamily.fanniemae.com/news-events/lender-letter-ll-2026-04-governance-framework-use-artificial-intelligence-and-machine-learning

2 *Fannie Mae Multifamily. "Lender Letters."* multifamily.fanniemae.com/doing-business-us/lender-letters

3 *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



Originally published in *The Agentic AI Governance Playbook* on Substack, August 7, 2026. maureendoylespare.substack.com



© 2026 Maureen Doyle-Spare.



*That is the control boundary that begins to matter when AI stops merely informing the mortgage process and starts acting inside 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.