

Governing Agentic AI Beyond Model Risk Guidance

Bank Control Coverage Under SR 26-2

Phil Hatch | Akholi

October 23, 2026

US bank leaders should require a complete control coverage decision for every outsourced agentic service. Exclusion from model risk guidance leaves the bank with a governance task that a provider model review cannot settle.

Require a Complete Coverage Decision

The bank's chief risk officer should identify who governs every material risk in each outsourced agentic service before management approves consequential action. A model review addresses only part of that decision. The relevant unit is the service: what the agent can read, decide, change, send, and delegate through systems controlled by the bank and its provider.

SR 26-2, issued April 17, 2026, replaces SR 11-7 and SR 21-8 and is expected to be most relevant to banking organizations above \$30 billion in assets. Its accompanying guidance excludes generative and agentic AI while directing banks to determine appropriate governance and controls for tools outside its scope.¹

Directors should therefore require an explicit coverage conclusion rather than accept the statement that an agent is outside model risk management. The conclusion should identify applicable obligations, responsible functions, unresolved gaps, operating restrictions, and the

¹ Board of Governors of the Federal Reserve System, *Revised Guidance on Model Risk Management* (Board of Governors of the Federal Reserve System, April 17, 2026), SR 26-2, applicability and supersession, <https://www.federalreserve.gov/supervisionreg/srletters/sr2602.htm>; Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency, *Supervisory Guidance on Model Risk Management* (Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency, April 17, 2026), sec. II and notes 1–3, <https://www.federalreserve.gov/frs/guidance/supervisory-guidance-on-model-risk-management.htm>.

executive authorized to accept residual exposure. Internal policy may deliberately apply model validation disciplines more broadly than the guidance's scope.

Separate the Model from the Service

An outsourced credit-document service illustrates the distinction. A conventional statistical score may sit inside the model inventory, while a language model extracts information and an agent writes to the case system. Validation of the score cannot establish that the agent used the correct file, respected permissions, preserved an adverse-action reason, or stopped after its authority expired.

The study's reported identity and observability problems make this separation operationally important. Respondents described shared credentials and tool reports that indicated completion without a corresponding result. These are reported delivery mechanisms requiring investigation, not evidence that every bank has the same condition.²

Management should map the complete path from source information to business consequence. Every handoff needs a named owner and evidence of successful operation. If a provider excludes its orchestration, tools, or subcontractors from assurance, the bank should record precisely which exposures remain untested and constrain the service until that gap is resolved.

Assign Risks to Functions with Authority

The proposed coverage map assigns each material action to a governing function. Information security covers identity and access; operational risk covers failed execution and continuity; compliance connects actions to obligations; privacy governs processing purposes and rights. Model risk specialists assess relevant model components and can contribute evaluation expertise without becoming the default owner of every agentic failure.

For each assignment, require a responsible executive, a control owner, an assurance owner, and a decision deadline. A committee label alone is insufficient. The owner needs access to provider evidence and authority to restrict deployment. Gaps between functions should go to the chief risk officer for resolution before they become disputes over incident ownership.

Figure 1 shows the questions the board should expect management's coverage map to answer. A row is complete only when the bank can connect an obligation or exposure to an

² Phil Hatch, *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*, version 1.8.2 (Akholi, forthcoming manuscript dated November 1, 2026), Part 3, Financial Institution Readiness, sec. 4.1, "Ability to Provision Access and Identities for Agents," indicator C108, <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 3, AI Technology, sec. 4.1, "Accuracy of the Tools' Own Success and Failure Reports," indicator C172.

operating control and a current test result. Several functions may contribute, but one executive should remain accountable for the service as a whole.

Figure 1. The coverage map the bank should require

Exposure	Bank decision	Evidence across both firms
Model component	Determine scope and evaluation owner	Inventory and component assessment
Agent action	Assign authority and operating limits	Permissions and denied action test
Legal obligation	Identify applicable entity and duty	Local applicability decision
Service failure	Assign containment and continuity	Joint reconstruction and rehearsal

Author’s proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

Source: Phil Hatch, proposed framework developed for this paper.

Keep Jurisdictional Conclusions Separate

A multinational bank should maintain local applicability annexes to the same service map. FINMA’s AI guidance describes Swiss supervisory observations on governance, testing, documentation, and independent review. Those observations provide useful comparison with the US scope decision; they do not extend SR 26-2 into Switzerland or convert every observed measure into a universal legal requirement.³

EU and UK entities need their own analysis of outsourcing, operational resilience, privacy, and sector duties. The bank should distinguish a rule applicable to the entity from a voluntary group standard. Where the group adopts a common control, local counsel should confirm whether local access, secrecy, or record-retention requirements change how the evidence can be shared.

The practical outcome is one service description with several legal applicability decisions, rather than a different technical story in every country. A cross-border provider should receive consistent operational instructions while each bank entity retains a defensible explanation of the requirements governing its own activities.

³ Swiss Financial Market Supervisory Authority, *Governance and Risk Management When Using Artificial Intelligence* (Swiss Financial Market Supervisory Authority, December 18, 2024), FINMA Guidance 08/2024, secs. 2.1–2.7, https://www.finma.ch/~media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmittelungen/20241218-finma-aufsichtsmittelung-08-2024.pdf?hash=60E2745D7AC1D7F4E06CE43474539F68&sc_lang=en.

Test Coverage Before Expanding Authority

Management should select a material outsourced workflow and reconstruct one completed case, one exception, and one stopped run. These are proposed starting samples, not sufficient statistical assurance. The review should follow the original input, policy and model versions, delegated permissions, actual tool actions, human corrections, and final bank acceptance.

A credible objection is that another governance process could duplicate existing review and delay useful deployment. The map should prevent that duplication by recognizing controls already supported by evidence and assigning only uncovered work. The chief risk officer should request one coordinated service decision, with specialist assessments attached, rather than require every function to repeat the same examination. Proportionality concerns should change the depth of review; they should not leave a material action without an accountable owner.

The assurance function should then test whether the mapped controls operated where the service crossed organizational boundaries. Provider certification or a clean model evaluation should remain evidence within its stated scope. Neither should substitute for the bank's test of its own approval, access, disclosure, or continuity duties.

Mayer Brown recommends adapting existing AI governance to autonomous action, including checkpoints and controls for consequential decisions. The additional outsourcing requirement proposed here is to assign and test those controls across both firms, so an apparent division of responsibility cannot leave the action ungoverned.⁴

Approve the Map and Restrict Uncovered Actions

The board risk committee should request a portfolio view of agentic services showing covered exposure, unresolved ownership, and restricted actions. Management should complete high-exposure coverage decisions first and assign funded remediation. An unresolved control owner should be treated as an unresolved deployment condition, rather than absorbed into a general maturity score.

The conclusion is specific: SR 26-2's scope exclusion requires a deliberate bank decision about the remaining governance. Continue only within authority supported by evidence. Expansion should depend on closing the relevant coverage gaps and verifying the resulting controls in the joint delivery model, with the chief risk officer reporting exceptions to the committee.

⁴ Arsen Kourinian, "Governance of Agentic Artificial Intelligence Systems," Mayer Brown, February 24, 2026, "Mitigation Measures," <https://www.mayerbrown.com/en/insights/publications/2026/02/governance-of-agentic-artificial-intelligence-systems>.

Evidence and Source Note

The study is a nonrepresentative cross-section of 850 vendor delivery-team respondents on 150 engagements serving Tier 1 financial institutions, collected during January–June 2026. Neither banks nor vendors corroborated the accounts. It establishes reported experience, not causation, legal breach, or market prevalence. The author-supplied version 1.8.2 is dated November 1, 2026 and is cited as a forthcoming manuscript consulted October 23, 2026.⁵

This is an analytical decision paper based on the supplied study and selected primary public sources checked through October 23, 2026. The original contribution is a service-level control coverage map separating model scope from bank accountability. The proposed method has not been independently validated. Legal requirements depend on the entity, activity, jurisdiction, and effective provisions; supervisory observations, research, and professional commentary have distinct authority.

Bibliography

Board of Governors of the Federal Reserve System. *Revised Guidance on Model Risk Management*. Board of Governors of the Federal Reserve System, April 17, 2026.

<https://www.federalreserve.gov/supervisionreg/srletters/sr2602.htm>.

Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency. *Supervisory Guidance on Model Risk Management*. Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency, April 17, 2026. <https://www.federalreserve.gov/frs/guidance/supervisory-guidance-on-model-risk-management.htm>.

Hatch, Phil. *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*. Version 1.8.2. Akholi. Forthcoming manuscript dated November 1, 2026. Author-supplied manuscript consulted October 23, 2026. <https://doi.org/10.2139/ssrn.7164559>.

Kourinian, Arsen. “Governance of Agentic Artificial Intelligence Systems.” Mayer Brown. February 24, 2026. <https://www.mayerbrown.com/en/insights/publications/2026/02/governance-of-agentic-artificial-intelligence-systems>.

Swiss Financial Market Supervisory Authority. *Governance and Risk Management When Using Artificial Intelligence*. Swiss Financial Market Supervisory Authority, December 18, 2024. https://www.finma.ch/~media/finma/dokumente/dokumentencenter/myfinma/4dokumentation/finma-aufsichtsmittelungen/20241218-finma-aufsichtsmittelung-08-2024.pdf?hash=60E2745D7AC1D7F4E06CE43474539F68&sc_lang=en.

⁵ Hatch, *Agentic AI Outsourcing*, Introduction, “Executive Summary” and “Evidentiary Basis and Its Limits.”