

Decoding the New Supervisory Guidance for Model Risk Management



On April 17, 2026, U.S. regulatory agencies issued SR 26-2, a landmark update to the foundational Model Risk Management (MRM) guidance. As the industry grapples with increasing complexity, SR 26-2 emphasizes a more nuanced, risk-sensitive approach that prioritizes "effective challenge" where it matters most. Superseding and replacing SR 11-7 and SR 21-8, the new framework largely preserves the underlying principles but introduces several important changes and adopts a higher-level approach with less detailed guidance.

Why Now? The Rationale for Evolution

The shift from SR 11-7 was driven by a decade of rapid technological change and industry friction. The agencies identified several catalysts for this update:

- **Modeling Complexity:** The rapid evolution of high-efficiency, innovative modeling techniques since 2011—including the acceleration in AI adoption and the emergence of Generative AI and Agentic AI systems—necessitated an updated definition of the guidance’s scope.
- **Tailored Supervision:** Regulators recognized that a "one-size-fits-all" approach was inefficient. SR 26-2 focuses primarily on institutions with \$30 billion or more in assets, allowing smaller firms to scale their practices to their specific risk profiles.
- **Definition Refinement:** For the definition of "model," consistent with industry practice, regulators have explicitly excluded simple arithmetic and deterministic rules (e.g., basic spreadsheets), focusing scrutiny on complex quantitative methods.

Comparison: SR 11-7 vs. SR 26-2

The transition represents important changes in how firms should manage model risk.

Feature	Old Guidance (SR 11-7)	New Guidance (SR 26-2)
Philosophical Approach	More prescriptive with some detailed guidance.	Principle-based with reduced detailed guidance.
Model Definition	Included most quantitative methods with no explicit exclusions.	Excludes simple arithmetic calculations and deterministic tools.
AI & Automation	Addressed generically as standard quantitative risk without explicit treatment.	Differentiates GenAI and Agentic AI from more traditional AI.
Scope	All supervised banking organizations.	Primarily organizations with >\$30B in assets, though still relevant to smaller institutions with significant model risk.
Enforcement & Expectations	Relatively high, often prescriptive expectation for validation and documentation.	Explicitly avoids "enforceable standards"; instead emphasizes tailored and commensurate risk management.

Guidance on Vendor Models

SR 11-7 and SR 26-2 are directionally consistent in affirming that vendor and other third-party models are fully subject to model risk management expectations, but they differ in depth and operational emphasis. SR 11-7 provides detailed, prescriptive guidance on how banks should select, validate, customize, and monitor vendor models, including expectations for vendor due diligence, assessment of development data and assumptions, use of alternative validation techniques when code access is limited, and bank-specific outcomes testing. It also emphasizes the need for in-house understanding, documentation of customization choices, and contingency planning in case a vendor relationship ends.

SR 26-2, by contrast, presents a more concise, principles-based articulation of these same concepts. It reinforces that validation, conceptual soundness, ongoing monitoring, and fitness-for-purpose remain essential for vendor models, including customized implementations, but without the procedural detail found in SR 11-7. In effect, SR 26-2 modernizes and reaffirms the SR 11-7 framework at a higher level, focusing on supervisory expectations and outcomes rather than step-by-step implementation guidance.

Setting the Stage for AI and Agentic Systems

Perhaps the most forward-looking aspect of SR 26-2 is how it handles emerging technology. It creates a "bridge" for future regulation by drawing a line in the sand:

- **Current Exclusions:** GenAI and Agentic AI are explicitly excluded from the scope of SR 26-2, labeled as "novel and rapidly evolving."
- **The "Standard AI" Foundation:** While GenAI and Agentic AI are out, the principles of SR 26-2 do apply to traditional statistical models and non-generative, non-agentic AI models (e.g., machine learning models used for fraud detection). This establishes a baseline for how AI should be governed before regulators tackle autonomous systems.
- **Principle over Prescription:** By focusing on outcomes and "effective challenge" rather than static technical checklists, the agencies have built a framework that can adapt as agentic systems mature.

While SR 26-2 expressly excludes Generative AI and Agentic AI from its scope, AI-related expectations continue to evolve across jurisdictions.

SR 26-2 Contrasts with the EU AI Act

Formally adopted in 2024 with phased implementation beginning in 2025, the EU Artificial Intelligence Act (EU AI Act) addresses model risk through a risk-based regulatory framework that classifies AI systems according to their potential impact on fundamental rights, safety, and societal outcomes. For high-risk AI systems, the Act requires a formal risk management system across the model lifecycle, including controls over training and testing data quality, documentation of model design and limitations, human-in-the-loop oversight, robustness and accuracy thresholds, and post-market monitoring to detect performance drift or unintended harm. Rather than focusing narrowly on model validation techniques, the EU AI Act embeds model risk within a broader governance regime that emphasizes preventive controls, transparency, accountability, and continuous monitoring, ensuring that model risks are identified, mitigated, and reassessed as AI systems evolve in real-world use.

Although the EU AI Act differs in scope and structure from SR 26-2, it provides a useful point of comparison in showing how model and AI-related risks are increasingly being addressed through broader governance, oversight, and lifecycle controls.

Anticipated Impact on Model Risk Supervision

SR 26-2 is likely the largest indicator of the shifting supervisory posture, signaling a move away from prescriptive, process-oriented reviews in favor of a tailored, principles-based supervisory focus. In upcoming horizontal exams, regulators may place less emphasis on comprehensive process-driven annual reviews and instead test the frameworks for "Model Materiality," specifically how firms justify quantitative exposure and qualitative purpose tiering. Examiners will likely focus on evidence of continuous lifecycle monitoring and the use of interim overlays, prioritizing dynamic risk mitigation over sequenced validation cycles.

For smaller institutions, this should not be read as an absence of supervisory expectation. Rather, the revised guidance reinforces that model risk management should be aligned to the size, complexity, and model risk profile of the institution. While smaller institutions may not require the same level of formalized infrastructure as larger firms, they would still be expected to demonstrate appropriate governance.

Most critically, because the agencies deliberately carved generative and agentic AI out of this guidance, banks should anticipate heightened scrutiny on model boundaries. Supervisors will likely zero in on the governance gaps where traditional machine learning

intersects with novel AI pipelines, particularly in high-stakes areas like transaction monitoring and compliance, meaning exam readiness will hinge on demonstrating governance of interconnected systems, data feeds, and configuration changes, rather than isolated mathematical weights.

Overall, the supervisory lens is not expected to become less onerous, but rather more focused on risk-tiered aspects, and will expect firms to bring that level of maturity to their frameworks.


Industry Takeaways

For MRM consultants and practitioners, SR 26-2 reiterates the fundamental principles of the prior guidance while explicitly clarifying key changes to scope and definitions. By narrowing definitions and embracing materiality, SR 26-2 alleviates the administrative burden of low-risk tools, allowing institutions to redirect their most skilled resources toward high-impact, complex models. While it offers a "stable harbor" for now by carving out Generative AI, it signals that the unique risks of autonomous agents will eventually require a specialized—and likely more rigorous—governance overlay. More than ever, banks are encouraged to prioritize holistic risk management over the management of individual models.

IBM Promontory can support banks in adopting SR 26-2 by helping translate its principle-based guidance into clear, operationally embedded model risk management practices. We can help banks integrate SR 26-2 requirements into existing model risk management processes, including model inventories and risk tiering, model documentation standards, validation and model review protocols, and ongoing monitoring routines.

Authors

For more information, please contact **Shweta Jain** at shweta.jain@ibm.com or **Miles Ravitz** at miles.ravitz1@ibm.com.



Shweta Jain

Partner, Risk and
Compliance Practice
Leader

IBM Consulting
shweta.jain@ibm.com



Miles Ravitz

Associate Partner,
Promontory Director

IBM Consulting
miles.ravitz1@ibm.com



Bo Qian

Senior Principal


IBM Consulting
bo.qian@ibm.com



Yuliya Ivanova

Senior Principal


IBM Consulting
yuliya.ivanova@ibm.com



Yiwei Ma

Senior Principal

IBM Consulting
yiwei.ma@ibm.com



Yidong Wang

Principal

IBM Consulting
yidong.wang@ibm.com

© Copyright IBM Corporation 2026

Produced in the
United States of America
April 2026

IBM, the IBM logo, and IBM Trademarks List, are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

