AI RISK CONSULTING

TURNING AI RISK INTO ADVANTAGE

REGULATORY ANALYSIS · AUGUST 2026

The Material Financial Risk Pivot

The OCC and FDIC have redrawn the perimeter of bank supervision around material financial harm. Combined with SR 26-2, the change moves most model, process and AI governance outside the supervisory enforcement net — and squarely onto the board.

What changed, what it means, and a tailored pathway for

community, regional and large banks — and for fintech and the wider BFSI sector

Satyaki Ghosh Dastidar · AI Risk Consulting LLC

airisk-consulting.com

CONTENTS

1 Executive summary

On 27 August 2026 the OCC and the FDIC finalised a rule that, for the first time, defines "unsafe or unsound practice" by regulation and sets a binding standard for when an examiner may issue a Matter Requiring Attention. On the same day the OCC reissued its bank enforcement manual, published its MRA manual for the first time, and proposed a further rule splitting violations of law into *substantive* and *technical* categories.

Read together with SR 26-2 — the revised model risk management guidance issued four months earlier — these documents describe a single, coherent supervisory philosophy: examiners should spend their attention on things that are likely to hurt the balance sheet, and should stop spending it on policies, procedures and documentation that do not.

That philosophy has four consequences that matter to anyone running a risk or governance function.

- **Harm is now defined narrowly and financially.** "Harm to financial condition" means losses or negative impacts to capital, asset quality, earnings, liquidity, or sensitivity to market risk. Reputational risk unrelated to financial condition is expressly out.
- **Process findings lose their teeth.** Weak policies and procedures now generally land as a *supervisory observation* — which carries no corrective-action requirement, need not go to the board, cannot be tracked by the examiner, and cannot be escalated into an MRA merely because the bank declined to act on it.
- **Tailoring cuts both ways, and the market is misreading it.** The rule states plainly that as an institution's risk factors rise, *the materiality threshold falls, the assessment becomes more granular, and remediation expectations increase*. For a large or complex bank this package is not deregulation. The OCC says outright that it may escalate against a large bank on practices that would not trigger a response at a community bank, and may escalate faster.

- **Model and AI governance now sit almost entirely outside the enforcement net.** SR 26-2 narrowed the definition of a model, set a $30 billion relevance threshold, and placed generative and agentic AI outside its scope altogether. Non-conformance with the 12 CFR 30 safety-and-soundness guidelines is expressly *not* a violation of law. The result is a genuine expectation vacuum around AI — one that boards, counterparties and litigants will fill even where examiners now will not.

THE THESIS IN ONE LINE

Supervision has stopped asking institutions to prove they followed a process and started asking them to demonstrate they were prudent. Those are different evidentiary burdens, and most governance functions are currently built for the first one.

The practical implication is not "do less governance." It is that the *output* of governance has to change. An MRA-driven control inventory produces artefacts. The new standard rewards a defensible, contemporaneous record that the institution's choices were consistent with generally accepted standards of prudent operation — assessed on objective facts and sound reasoning, at a granularity that scales with the institution's size and complexity.

2 What actually changed on 27 August 2026

Four instruments were released together. They are easy to conflate and they do different things.

THE 27 AUGUST 2026 PACKAGE

Instrument	Status	Effect
Interagency final rule OCC 12 CFR 4.92 · FDIC 12 CFR 305.1	Final; effective 60 days after Federal Register publication	Defines "unsafe or unsound practice"; sets the binding MRA standard; defines harm to financial condition; codifies tailoring, the objective-facts standard, supervisory observations, and "other violations."
PPM 5310-3 Bank Enforcement Actions	Revised internal policy (replaces the 25 May 2023 version)	Escalation, tailoring and focus as guiding principles; when MOUs, IMCRs, C&Ds and capital directives are used; 180-day timeliness target; mandatory substantial reliance on satisfactory internal audit for validation.
PPM 5400-11 Matters Requiring Attention	Published publicly for the first time	Operationalises the MRA standard for examiners and makes the OCC's internal MRA methodology visible to banks for the first time.
Violations NPRM Docket OCC-2026-0529	Proposed; 30-day comment period from FR publication	Would replace "actual violation" with "substantive violation" in the MRA standard, define five substantive categories, and create a new "technical violation" channel with sharply limited examiner powers.

2.1 The two definitions that carry the weight

An **unsafe or unsound practice** is now a practice, act, or failure to act that (1) is contrary to generally accepted standards of prudent operation *and* (2) if continued is *likely* to materially harm the institution's financial condition or present a material risk of loss to the Deposit Insurance Fund — or that already materially harmed financial condition.

An **MRA** may be issued on a lower probability: the same imprudence element, plus harm that *could reasonably be expected* under current or reasonably foreseeable conditions. Or, separately, for an actual violation of a banking or banking-related law or regulation. The agencies were explicit that this second, lower rung is what allows proactive supervision — they argue the MRA standard would have reached Silicon Valley Bank's interest-rate vulnerability during 2022.

> **THE IMPRUDENCE GATE MATTERS MORE THAN THE HARM TEST**
>
> Both standards require a finding that the institution acted contrary to generally accepted standards of prudent operation *first*. Likely material harm alone is not enough — the agencies have said they will train examiners accordingly. Equally, the agencies declined to publish a list of those standards and confirmed that practices identified in horizontal reviews of peers are *not* automatically the standard. Both facts are usable by an institution defending a supervisory finding.

2.2 Five choices in the final rule worth knowing by heart

- **Institution-affiliated parties were dropped.** The proposal would have applied the new definition to IAPs; the agencies pulled that back, because a large bank's balance sheet is hard to move and an individual could otherwise misbehave at a large institution with no recourse, then move to a small one. Enforcement against IAPs continues under prior standards and case law.
- **"Likely" was not quantified.** Commenters pushed for 51% or a time horizon. The agencies refused both, and refused to codify a burden of proof, resting instead on "objective facts and sound reasoning."
- **Materiality was not quantified either** — but it is explicitly relative to the institution, and it moves with the tailoring dial.
- **Consumer law above $10 billion is CFPB territory.** Consistent with 12 U.S.C. 5515–5516, the agencies will not issue MRAs for violations of Federal consumer financial laws at insured depository institutions above $10 billion in assets. That does not make the exposure disappear; it relocates it.
- **The Federal Reserve is not a party.** The final rule amends OCC and FDIC regulations only. The Board joined SR 26-2 in April but did not join this rule. State member banks and holding companies therefore continue under the prior, judicially developed standard — a real divergence for any multi-charter organisation.

2.3 The violations proposal

The NPRM would deem a violation **substantive** if its nature, duration, frequency or severity could meaningfully impact the institution or its customers, *and* it falls into at least one of five categories:

PROPOSED SUBSTANTIVE VIOLATION CATEGORIES · 12 CFR 4.92(D)

Category	Illustrations given by the OCC
1 · Systemic, or a pattern	BSA compliance-programme or pillar violations would generally qualify; repeated appraisal failures on higher-risk mortgages; ongoing OFAC reporting failures. Isolated violations are excluded.
2 · Direct, clear, predictable and more-than-minimal impact on financial condition	Purchase of low-quality assets from an affiliate in violation of Regulation W. Correlation or attenuated connection is not enough; a civil money penalty is the regulatory *response*, not the impact.
3 · More-than-minimal impact on books and records	Call Report inaccuracy — not per se, but generally where it would be a "material misstatement or omission" under FFIEC instructions.
4 · More-than-minimal restitution or adverse impact to customers	Both a few customers owed a lot and many owed a little. Non-financial impacts count: a failed FCRA identity-theft programme qualifies whether or not losses materialise.
5 · Insider misconduct or self-dealing	Regulation O; in some cases section 22(e) of the Federal Reserve Act. Qualifies *regardless* of size or prevalence.

Everything else becomes a **technical violation**. For those, the OCC could direct correction but could not specify how, could not require notification that the fix was made, could not require an action plan, and generally would not track remediation. Non-conformance with the 12 CFR part 30 guidelines would be neither substantive nor technical — it sits outside the violation framework altogether.

3 The new four-rung supervisory ladder

The practical output of the package is a clearer, steeper ladder. Knowing which rung a finding sits on now determines whether it consumes management attention at all.

Supervisory observation

Informal weakness in policies, practices, condition or operations. **No corrective action required. No board presentation required. Examiner may not track your response, may not require an action plan, and may not escalate solely because you declined the suggestion across multiple cycles.** But the underlying information can still support your ratings.

Technical violation (proposed) **/ "other violation"** (final rule)

An actual violation that does not meet the substantive threshold. **You must remediate; the agency may not prescribe how, and generally will not track it.** At the FDIC, an uncorrected other violation can be cited as an MRA at the next examination.

Matter Requiring Attention

Imprudence plus reasonably expected material financial harm or DIF risk — or a substantive violation. **Goes to the board; corrective action is mandatory.** An MRA is not a binding order, and mere failure to remediate one is not itself an unsafe or unsound practice.

Enforcement action

Informal (MOU, IMCR, notice of deficiency) escalating to formal (C&D or consent order, capital directive, safety and soundness order). **MOUs are generally unavailable where there has been an unsafe or unsound practice or a substantive violation** — which pushes those matters straight to formal action.

THE QUIET ESCAPE HATCH: 12 CFR 30 NOTICES OF DEFICIENCY

PPM 5310-3 states that a notice of deficiency may be appropriate where a bank shows *serious, systemic or repeat information technology or information security deficiencies* that must be addressed but are not likely to cause material financial harm — and, for a large or complex bank, where *serious internal control or risk management deficiencies exist but do not constitute unsafe or unsound practices or meet the MRA standard.*

Read that twice. The very findings the final rule appears to demote — control weakness, risk management weakness, IT and information security weakness — retain a dedicated enforcement channel that bypasses both the MRA standard and the unsafe-or-unsound definition. Anyone reading this package as an amnesty on control deficiencies has read only half of it.

3.1 Changes to how findings are worked, not just issued

- **Corrective actions must be minimal and targeted.** Enforcement actions must contain only what is necessary to remediate, must not elevate process and procedure over substance, and must not include requirements unrelated to the core purpose of the action.
- **Internal audit becomes load-bearing.** When validating corrective actions, examiners *must* substantially rely on the bank's internal audit work where that function is rated satisfactory — and may not require the bank to add validation to the audit scope purely to enable OCC reliance. The quality and rating of internal audit now directly determines examination burden.
- **MRAs should close when remediation is complete.** The agencies acknowledged they have held MRAs open past full remediation to observe sustainability, and identified that

practice as a distortion — though they deferred the fix to a later process.

- **Clocks are explicit.** Enforcement actions should be presented within 180 days of the start of the supervisory activity that indicated one was warranted; the first compliance assessment falls within 90 days of no-supervisory-objection on an action plan, or 120 days of execution where no plan is required.
- **Termination on substantial compliance** — the essential requirements satisfied, even if minor or isolated technical requirements are not.

4 The SR 26-2 convergence — and the expectation vacuum

SR 26-2, issued jointly by the Board, OCC and FDIC on 17 April 2026, superseded SR 11-7 after fifteen years. It runs on the same logic as the August package, and the two must be read together.

SR 11-7 → SR 26-2: WHAT MOVED

Dimension	SR 11-7 (2011)	SR 26-2 (2026)
Definition of a model	Quantitative method applying statistical, economic, financial or mathematical theories	Adds a **complexity** element; expressly excludes simple arithmetic and spreadsheet calculations and deterministic rule-based processes
Applicability	No uniform threshold across agencies	Most relevant to organisations **over $30 billion**; smaller institutions in scope where model prevalence or complexity, or non-traditional activities, warrant it
Generative and agentic AI	Not contemplated	**Expressly out of scope** as novel and rapidly evolving; institutions expected to apply broader risk management practices independently
Validation	Prescriptive cadence expectations in practice	Timing, nature and frequency vary with purpose, methodology and change; use before validation is permissible for urgent business need with enhanced attention to limitations and controls
Tiering	Implicit	Explicit materiality framework: immaterial models may be limited to monitoring and performance assessment
Third-party models	Detailed validation protocol	Understand the model; validate internally or externally; ongoing monitoring and outcome analysis for fitness and reliability

4.1 The vacuum

Stack the two reforms and a gap opens with unusual clarity. Consider a large bank deploying a generative AI system into a customer-facing process.

- It is **outside SR 26-2**, which places generative and agentic AI beyond its scope.
- Weaknesses in the policies, controls and documentation around it are, on their face, **the classic profile of a supervisory observation** — no corrective action, no board presentation, no tracking.
- Non-conformance with the 12 CFR 30 safety-and-soundness guidelines is **expressly not a violation** of law or regulation, so it cannot found an MRA on the violation prong.
- And it will rarely be *likely* to materially harm the financial condition of a large institution's balance sheet, so it will rarely be an unsafe or unsound practice.

Four independent routes to supervisory attention, and a well-run AI deployment misses all four. That is not an accident of drafting — it is the intended reallocation of supervisory resource. But an institution should not confuse the absence of a supervisory finding with the absence of risk. Three things fill the vacuum immediately:

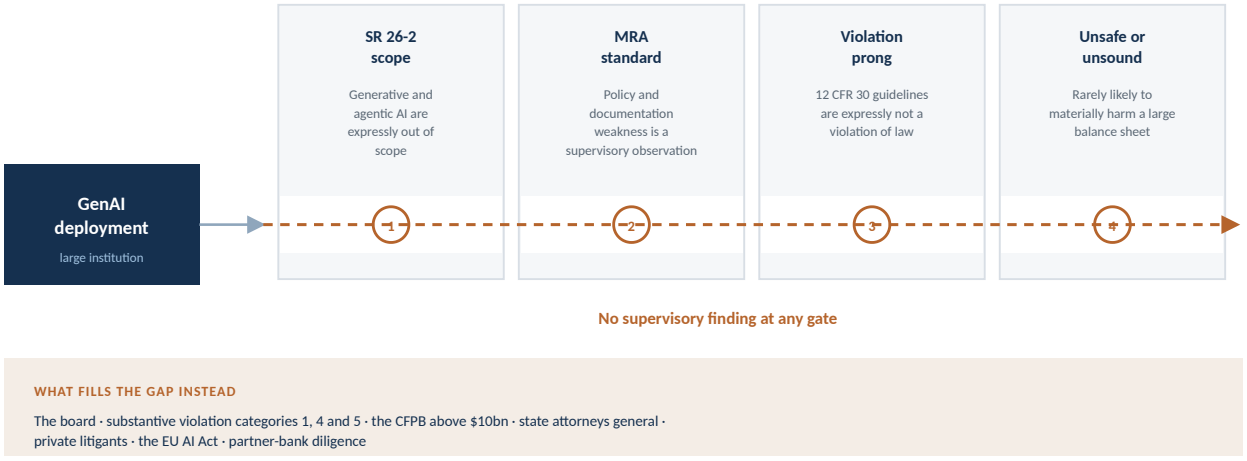

FIGURE 1 Four independent routes to supervisory attention, and a well-run generative AI deployment misses all four. The reallocation of supervisory resource is deliberate; the accountability gap it leaves is not filled by nothing.

WHAT FILLS THE SPACE EXAMINERS HAVE VACATED

- **The board.** The rule repeatedly frames its purpose as returning discretion to the board and management to exercise business judgement. Discretion returned is accountability returned. Where an examiner previously validated a control environment, the board is now the primary and often only reviewer.
- **Substantive violation categories 1, 4 and 5.** An AI system that produces systemic adverse-action or disclosure failures, or more-than-minimal customer impact, is squarely substantive — even with no balance-sheet effect and no financial loss to any customer. The FCRA identity-theft example in the NPRM makes that explicit. Fair lending, ECOA, EFTA, BSA/AML and OFAC remain banking-related laws.
- **Everyone else.** The CFPB above $10 billion, state attorneys general, private litigants, the EU AI Act for institutions with EU exposure, and — for anyone selling into banks — the diligence and contractual demands of partner institutions.

5 What did *not* change: five traps

TRAP 1 · READING TAILORING AS UNIVERSAL RELIEF

The rule is symmetric and the market has largely reported only one half. As risk factors rise, the materiality threshold *falls*, assessment becomes *more* granular — down to specific business lines, products or services — and remediation requirements and prudent-operation expectations *increase*. PPM 5310-3 says the OCC may escalate against a large or complex bank on practices that would not trigger a similar response at a community bank, and may escalate faster. For the largest institutions this package tightens as much as it loosens.

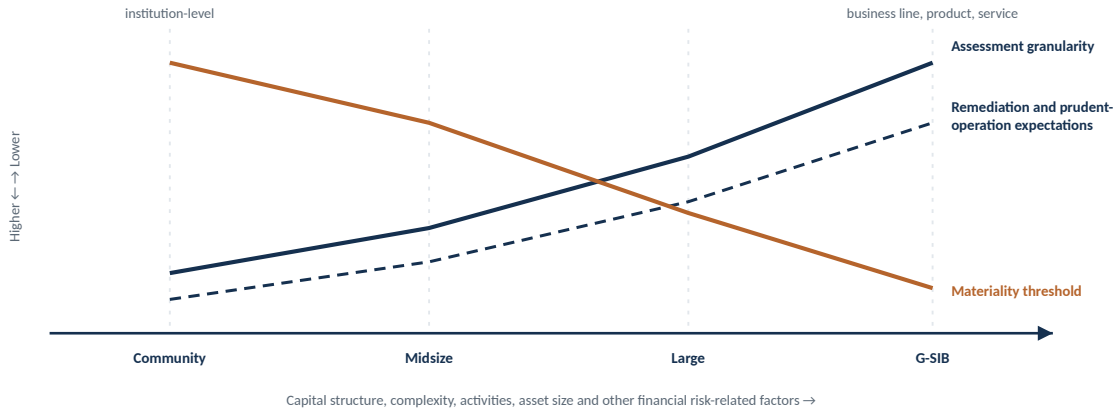

FIGURE 2 Tailoring is symmetric. As risk factors rise, the materiality threshold falls, harm is assessed more granularly, and remediation and prudent-operation expectations increase — codified at 12 CFR 4.92(e)(3) and 12 CFR 305.1(e)(3). The inverse applies as risk factors fall.

TRAP 2 · ASSUMING SUPERVISORY OBSERVATIONS ARE FREE

They require nothing. But the agencies confirmed that the information underlying supervisory observations can still support assigned ratings. A pattern of unaddressed observations that never becomes an MRA can still become a composite downgrade — with all the licensing, capital-distribution and M&A consequences that follow.

TRAP 3 · TREATING SR 26-2 AS PERMISSION TO SHRINK THE INVENTORY

The narrowed model definition removes deterministic rule-based processes and simple calculations from the guidance. It does not remove them from your risk. A deterministic pricing or eligibility rule that systematically produces a prohibited outcome is not a model risk problem — it is a substantive violation under category 1 or 4. Scoping tools out of the model inventory without scoping them into something else creates an uncontrolled population, and that population is exactly where systemic and pattern violations breed.

TRAP 4 · FORGETTING THE FEDERAL RESERVE DID NOT SIGN

A holding company with a national bank and a state member bank now faces two different unsafe-or-unsound standards. Group-level policy written to the OCC/FDIC definition will not automatically satisfy Board supervision, and internal issue taxonomies that assume a single standard will misclassify findings at the margin.

TRAP 5 · LOSING THE APPEAL ADVANTAGE

Codification cuts for the institution as well. Examiners must now justify findings on objective facts and sound reasoning; the agencies expect that justification to appear in the report of examination or supervisory letter precisely so that institutions can decide whether to appeal. Both agencies have live appeals reforms — the FDIC published revised guidelines for appeals of material supervisory determinations in January 2026 and the OCC proposed a Bank Appeals Process rule in February 2026. An institution that does not track the stated basis for each finding forfeits leverage it has just been handed.

6 Tailored pathways by segment

The rule's tailoring provision is a sliding scale, not a set of tiers — the agencies expressly declined to publish asset bands or a matrix. What follows translates the scale into four operating postures.

Community & small banks

Under ~$10B

WHAT CHANGES FOR YOU

The materiality threshold rises. Harm is assessed at the institution level rather than by business line or product. Expectations regarding prudent operation are lower and escalation is slower. You are below the SR 26-2 $30 billion relevance threshold, and the agencies have confirmed that practices identified in horizontal reviews of larger peers are not automatically generally accepted standards of prudent operation.

WHAT TO DO IN 90 DAYS

- Re-baseline the open issue inventory against the new MRA standard. Findings that were process-only are candidates for closure or downgrade — document the reasoning and raise it at the next exam.
- Strip out the controls you copied from bigger banks. Model validation cadences, committee structures and documentation standards borrowed from a $100B peer are now affirmatively hard to justify as *your* standard of prudent operation.
- Build the DIF-risk case file for liquidity. The DIF prong applies regardless of size, and the agencies' worked example is inadequate contingency funding — which can present material DIF risk without materially harming your financial condition. A small institution can present a material *risk* of loss even if the potential loss amount is small.
- Separate the Regulation O and insider file from everything else. Category 5 substantive violations qualify regardless of size or prevalence.

WHAT TO STOP

Building controls whose only justification is that a larger institution has them. Tracking supervisory observations as if they were MRAs.

THE TRAP Pattern violations. Tailoring means a *higher number* of violations is needed to establish a pattern at a large institution — but a *lower error rate* threshold applies there. Inverted for you: fewer absolute violations establish a pattern at a small bank, because the count is assessed relative to how often you undertake the activity. Ten flood-insurance failures against a small origination volume is a pattern.

Regional & midsize banks

~$10B – $100B

WHAT CHANGES FOR YOU

This is where the dial moves most, and where two independent thresholds cross your segment. Above $10 billion, the agencies will not issue MRAs for Federal consumer financial law violations — that supervision sits with the CFPB. Around $30 billion, SR 26-2 becomes squarely relevant. Between those lines the same institution can be simultaneously relieved of one supervisory channel and newly exposed to another.

WHAT TO DO IN 90 DAYS

- Map your issue population against *both* supervising authorities. A consumer-compliance finding that would previously have arrived as an OCC or FDIC MRA now arrives — if at all — from the CFPB, on a different timetable, with different remedies and different publicity. Your issue-management workflow probably assumes one prudential channel.
- Decide your SR 26-2 posture deliberately rather than by drift. If you are approaching $30 billion, the question is not whether the guidance applies but what your model risk profile requires. Write that assessment down; it is the artefact that demonstrates prudent operation.
- Rebuild the issue taxonomy around the four rungs. Most midsize banks run two states — MRA or not. You need four: supervisory observation, technical violation, MRA, enforcement action — each with its own routing, ownership and board treatment. The difference between rungs is now legally consequential.
- Raise internal audit's standing. Substantial examiner reliance on satisfactory internal audit converts audit quality into direct examination relief. For a bank at your scale that is the highest-leverage investment in this package.

WHAT TO STOP

Treating the CFPB threshold as a compliance-department fact. It is now a supervisory-architecture fact.

THE TRAP Escalation speed. You are large enough that the OCC will escalate faster than it would at a community bank, but you likely lack the dedicated remediation infrastructure of a $100B+ institution. The 180-day presentation target and the 90/120-day first-assessment clocks are tight if your action-plan process runs through quarterly committees.

Large & systemically important banks

$100B+

WHAT CHANGES FOR YOU

Less than the headlines suggest, and in some respects the opposite. Your materiality threshold *decreases*. Harm is assessed at the level of specific business lines, products or services. Remediation requirements and expectations regarding prudent operation *increase*. Escalation is faster, and the OCC may act on practices that would not draw a response at a smaller bank.

What you gain is not leniency but the removal of process-only findings and, potentially, faster MRA closure once remediation is complete.

WHAT TO DO IN 90 DAYS

- Re-triage the open MRA population into three buckets: findings that survive the new standard, findings that should now close as remediated, and findings that were always process-only and should be renegotiated to observations. Bring the analysis to the exam team rather than waiting for it.
- Build the granular harm model. If materiality is assessed at business-line and product level for you, your issue-materiality methodology has to operate at that level too. Enterprise-level materiality thresholds will systematically understate what your examiner will consider material.
- Treat the 12 CFR 30 notice-of-deficiency channel as live. It is named explicitly for large or complex banks with serious internal control or risk management deficiencies that fall outside both the MRA standard and the unsafe-or-unsound definition. Your operational resilience, IT and information security programmes should be governed to that standard, not to the MRA standard.
- Instrument the appeal record. Capture, per finding, the examiner's stated objective facts and reasoning, whether the imprudence element was independently established, and whether the harm assessment was granular or enterprise-level. That file is your appeal, and it is also your evidence of prudent operation.
- Close the AI governance gap deliberately. Generative and agentic AI sits outside SR 26-2 and largely outside the MRA net. Your board, your auditors and your counterparties will not accept that as an answer. Govern it under an internal standard that you set, publish and can evidence.

WHAT TO STOP

Describing this package internally as deregulation. It is a reallocation, and for you the net direction on material matters is tighter.

THE TRAP Internal audit dependency. Mandatory substantial reliance on satisfactory internal audit is relief only while the audit rating holds. An audit downgrade now costs you twice — once directly, and again in the examination burden that returns.

Fintechs, BaaS partners & the wider BFSI sector

Indirect exposure

WHAT CHANGES FOR YOU

Nothing directly — the final rule is expressly limited to institutions the agencies supervise. But three transmission channels carry it to you. First, PPM 5310-3 notes that its principles may be considered in taking enforcement action against a *third-party service provider*. Second, your partner banks' diligence questions will change shape: they now need evidence of prudent operation and of material-risk containment, not evidence of process. Third, the substantive-violation categories map almost perfectly onto how BaaS and embedded-finance programmes actually fail — systemic or patterned failures, more-than-minimal customer impact, restitution, and books-and-records accuracy.

WHAT TO DO IN 90 DAYS

- Re-write your bank-partner evidence pack around the five substantive categories. Show, per category, how your programme prevents the failure mode and how you would detect it. That is the artefact your partner's second line now needs.
- Instrument for pattern detection. Category 1 turns on systemic-or-pattern findings. Programmes that report errors as individual incidents cannot demonstrate the absence of a pattern; error-rate telemetry against activity volume can.
- Own the reconciliation and reporting chain. Books-and-records accuracy is a standalone substantive category, and in a BaaS structure the ledger accuracy that feeds a partner's Call Report often originates with you.
- For insurers, asset managers and non-bank BFSI: resist reading SR 26-2's narrowed model definition across to your own regime. Your supervisors have not adopted it, and the complexity carve-out was calibrated to bank supervisory resource allocation, not to model risk itself.

WHAT TO STOP

Selling into banks on the promise of MRA remediation. That market is contracting by the OCC's own estimate. Sell on evidence of prudent operation instead.

THE TRAP A false read of relief. If your partner bank concludes that its process findings no longer matter, its oversight of you gets thinner — right up until a customer-impact or pattern finding arrives as a substantive violation, at which point the bank's exposure is immediate and formal, and the contractual consequences land on you.

7 Six moves for the next two quarters

1. **Re-baseline the issue inventory against the four rungs** — supervisory observation, technical violation, MRA, enforcement action. Every open finding should carry an explicit classification and a stated rationale. This is the single highest-value exercise in the package, and it is time-limited — the argument is strongest before the first examination cycle under the new standard sets precedent.
2. **Rebuild the evidence model around prudence, not process.** The question has changed from "did you follow your policy" to "was this consistent with generally accepted standards of prudent operation, on objective facts." Contemporaneous decision records, documented rationale for risk acceptance, and evidence of challenge answer the second question. Control-testing evidence answers only the first.
3. **Recalibrate materiality to your tailoring position.** Large institutions need business-line and product-level materiality. Community institutions should raise enterprise thresholds and say so. Both need the calibration written down.
4. **Close the AI governance gap with an internal standard.** Generative and agentic AI is outside SR 26-2 and largely outside the MRA net. Adopt a proportionate internal standard — risk-tiered, with defined controls at each tier and runtime evidence for the highest tiers — and be able to show the board why it is sufficient. The absence of a supervisory expectation is not the absence of a duty.
5. **Elevate internal audit.** Its rating now directly determines examination burden through the mandatory substantial-reliance provision. Fund it accordingly.
6. **Comment on the violations NPRM.** The comment window is 30 days from Federal Register publication and the OCC asked fourteen open questions — including whether to drop "banking or banking-related" from the MRA standard entirely, and whether to extend substantive violations to a "more than minimal adverse impact to the public." Both would materially widen the net that has just been narrowed.

WHERE THIS FITS IN A READINESS ROADMAP

Institutions at the *Ad hoc* and *Defined* levels of maturity will experience this package as relief, and will be tempted to bank it. Institutions at *Operationalized* and above will recognise that the burden has shifted rather than lifted — from demonstrating conformance to an external standard to demonstrating the sufficiency of an internal one. That is a harder problem, and it is the one that separates a governance function that accelerates the business from one that merely documents it.

8 What to watch

Item	Timing	Why it matters
Violations NPRM comment deadline Docket OCC-2026-0529	30 days after Federal Register publication	Fourteen open questions, several of which would materially change the perimeter — including removal of the "banking or banking-related" limit and a possible "adverse impact to the public" category.
Final rule effective date	60 days after Federal Register publication	Applies prospectively to supervisory and enforcement activity only — not to the agencies' rulemaking authority.
FFIEC CAMELS proposal	Published 19 May 2026; pending	Would refocus component and composite ratings on factors materially affecting financial condition. The agencies deferred a proposal requiring every composite 3-or-below to carry an MRA or enforcement action until this concludes.
Appeals reform	FDIC guidelines Jan 2026; OCC proposed rule Feb 2026	Codified standards plus reformed appeals is the combination that makes challenging a supervisory determination realistic.
Federal Reserve alignment	Open	The Board joined SR 26-2 but not this rule. Whether it follows determines how long the two-standard problem persists for multi-charter groups.

SOURCES

OCC News Release 2026-72, *OCC Acts to Improve Transparency and Consistency to Bank Enforcement and Supervisory Standards* (27 August 2026), and OCC/FDIC News Release IA-2026-71, *Agencies Issue Final Rule to Prioritize Material Financial Risks* (27 August 2026). Interagency final rule, *Unsafe or Unsound Practices*, *Matters Requiring Attention*, OCC Docket ID OCC-2026-0174 / RIN 1557-AF35 and FDIC RIN 3064-AG16, adding 12 CFR 4.92 and 12 CFR part 305; proposal at 90 FR 48835 (30 October 2025). OCC PPM 5310-3, *Bank Enforcement Actions and Related Matters* (27 August 2026), superseding the 25 May 2023 issuance; OCC PPM 5400-11, *Matters Requiring Attention* (27 August 2026); OCC Bulletins 2026-40, 2026-41 and 2026-42. OCC notice of proposed rulemaking, *Violations of Laws or Regulations*, Docket ID OCC-2026-0529 / RIN 1557-AF56. Board of Governors of the Federal Reserve System, OCC and FDIC, SR 26-2, *Revised Guidance on Model Risk Management* (17 April 2026), superseding SR 11-7. FDIC, *Guidelines for Appeals of Material Supervisory Determinations*, 91 FR 3184 (26 January 2026); OCC, *Bank Appeals Process* proposed rule, 91 FR 7163 (17 February 2026); FFIEC CAMELS proposal (19 May 2026).

DISCLAIMER

About AI Risk Consulting

AI Risk Consulting helps banks, fintechs and regulated organisations build AI, GenAI and model risk governance that accelerates the business rather than blocking it — governance as a rail, not a roadblock.

Our work is grounded in twenty years inside tier-1 financial institutions building model risk and AI governance programmes from the ground up, and follows a four-stage method: Assess, Design, Implement, Sustain.

Satyaki Ghosh Dastidar
AI Risk Consulting LLC
airisk-consulting.com · info@airisk-consulting.com