ALEXANDER BARRETT

Proposal and Comment Information

Title: Proposed Third-Party Risk Management Guidance, OP-1881

Comment ID: FR-2026-0031-01-C63

Submitter Information

Name: Alexander Barrett

Submitted Date: 11/16/2026

Please see the attached comment letter.

Alexander D. Barrett
Senior Fellow, Verida Charter Foundation (affiliation for identification only)
ORCID 0009-0003-0359-3163
adb@veridafoundation.org

November 16, 2026

Chief Counsel's Office, Office of the Comptroller of the Currency (Docket ID OCC-2026-0793)
Benjamin W. McDonough, Secretary, Board of Governors of the Federal Reserve System (Docket No. OP-1881)
Jennifer M. Jones, Deputy Executive Secretary, Federal Deposit Insurance Corporation (ZRIN 3064-ZA58)
Melane Conyers-Ausbrooks, Secretary of the Board, National Credit Union Administration (Docket No. NCUA-2026-1684)

Re: Proposed Third-Party Risk Management Guidance, 91 FR 58536 (September 15, 2026)

Dear Sir or Madam,

Thank you for the opportunity to comment on the proposed guidance. I write as an independent researcher. For the past 15 months, I have studied what happens when one regulated organization relies on another's decision. Since September 2025, the research has comprised 85 semi-structured interviews with senior practitioners in six regulated industries across five countries, weighted toward insurance in the United Kingdom and the United States. It has also produced a series of working papers, listed at the end of this letter, and forms part of a longer working thesis, The Aetiology of Boundary Blindness.

I support the proposal's direction. A risk-based, tailored approach that lets banking organizations rely on consortia, certifications and credible summaries, instead of repeating every assessment themselves, is the right one. My comment concerns one word that this reliance now depends on: credible.

Summary

The proposed guidance allows a banking organization to rely on "credible summaries of technical data or proprietary information, as well as certifications or results of assessments." Still, it does not say what makes them credible. It also does not address artificial intelligence or the use of models. However, third parties use models, including generative and agentic AI, to produce the assessments and decisions that banking organizations rely on. I suggest four changes, all consistent with the proposal's non-prescriptive approach:

1. Illustrate what can make a summary, certification, or assessment credible, as the proposal already does for credible indemnification.
2. If the agencies list high-risk characteristics in answer to footnote 15, include third-party decisions that could themselves cause a violation of law or material harm, where the basis of those decisions does not reach the banking organization.

3. Note, for higher-risk relationships, the value of recording what a banking organization relied on when it accepted a third party's assessment, or accepted residual risk on limited information.
4. Note that any common format for carrying the basis of a certification should not itself become a single point of dependence.

1. The outcome travels, the basis does not

A decision has two parts: the outcome and the basis for it. When a decision passes from one organization to the next, the outcome travels, but the basis does not. The receiving party gets the outcome, not who decided, under what authority, against which rules and regulations, on what evidence, or whether it still holds. And where a model was involved, no record of which model, or of what the human did with its output: accepted it, changed it, overrode it, or never read it. So the receiving party starts again. I call this boundary blindness.

A growing part of third-party risk management is managing reliance on other organizations' decisions. A certification is a decision that a third party meets a standard. A credible summary is a decision about what technical data shows. An audit opinion is a decision about whether controls work. A fintech partner's onboarding approval, fraud alert, or credit decision is a decision the banking organization will answer for. As the proposal puts it, a banking organization's use of third parties "does not diminish its responsibility to meet these requirements to the same extent as if the activities were performed by the banking organization internally." The proposal also notes that "the degree of discretion granted to a third party can vary significantly in different types of relationships." Discretion is the power to decide. Where a third party decides, the banking organization receives the outcome. Whether it also receives the basis is the question this comment addresses.

2. What the proposal already recognizes

The proposal already declines to let an external assessment stand in for a banking organization's own judgment. Footnote 14 says banking organizations "should not rely on the existence of any agency supervision or regulation as a substitute for managing third-party risk or as a proxy for whether such third parties provide services safely and soundly." Footnote 17 says the agencies' own reports of examination of certain service providers "are not intended as a proxy or substitute for any banking organization's responsibilities to undertake effective risk management." It gives the reasons: such reports "are not tailored to any individual banking organization's circumstances or risk management needs and may be based on targeted rather than comprehensive examinations." Both reasons are often true of certifications and attestation reports as well.

The cross-cutting section then permits reliance on outputs from consortia, standard-setting and certification organizations, consultants, auditors and law firms. It says that "review of credible summaries of technical data or proprietary information, as well as certifications or results of assessments provided by such arrangements, may be adequate for a banking organization's due diligence needs, depending on facts and circumstances." It adds that "reliance on third parties for risk management purposes can itself involve risks."

That is a sensible position, but the guidance leaves "credible" undefined. The proposal's concern with the 2023 Guidance is that it "frequently has been interpreted in an overly broad manner," and it aims to "promote consistency." An undefined standard at the center of reduced-cost reliance invites the same problem, read differently by different examiners and banking organizations. It also risks letting the existence of a certificate become the proxy that footnotes 14 and 17 warn against.

Certifications and attestation reports already carry part of their basis. An ISO/IEC 27001 certificate states its scope, the standard version, and a validity period. A SOC 2 Type II report describes the system, the criteria, the tests performed, any exceptions, and the controls the user must operate itself. But a certificate does not carry the audit evidence, and an attestation report covers a past period, usually bridged to the present only by the provider's own unaudited letter. Neither speaks to the individual decisions a third party then makes for the banking organization. And where a model did part of the work, neither says which model was used or what a person did with its output.

3. The model gap

The proposed guidance does not address artificial intelligence or the use of models. Yet third parties use models, including generative and agentic AI, to produce the exact outputs banking organizations rely on, such as compliance screening, fraud detection, identity verification, and credit decisions. Models can also sit inside the "third-party technologies or processes" that the proposal says may "supplement or assist the banking organization's own risk oversight activities."

Adjacent guidance leaves a gap here. The interagency Revised Guidance on Model Risk Management, issued with SR 26-2 on April 17, 2026, superseded SR 11-7. It places generative and agentic AI models outside its scope, leaving them to each banking organization's own risk management and governance practices. It is expected to be most relevant to banking organizations with over $30 billion in total assets, and the NCUA is not a party to it. For most banking organizations, then, a third party's use of generative or agentic AI will be assessed largely through third-party risk management, and the proposed guidance gives them nothing specific to ask about it. The banking organization nonetheless remains responsible for the resulting decision as if it had made it internally.

Two features of models make the basis especially important. The same named model can behave differently from one release to the next, and providers of hosted models retire them on notice, after which requests to them fail. A record that does not say which model and which version produced an output cannot later be used to reconstruct what happened. The proposal lists "reviews of updates to information utilized in due diligence" among monitoring activities. A third party's model change is such an update, but a banking organization will rarely learn of it unless someone records which model was used. The human role is also often invisible. A person may have accepted a model's output, changed it, overridden it, or never read it, and the receiving organization usually cannot tell which.

Recording this is feasible. European banking guidelines on internal ratings provide that institutions should document the scale and rationale of each override of the inputs or outputs of the rating assignment process (EBA/GL/2017/16, paragraph 204). They should also store the date of the override and the person who performed and approved it. That provision is narrow. It covers overrides only, in one model class, and the record stays inside the firm. But it shows that a per-decision record of what a person did with a model's output can be produced in ordinary practice.

The same problem compounds along chains. The proposal lets a banking organization oversee subcontractors by "assessing, confirming, and monitoring the adequacy of its third party's own third-party risk management programs." Each step, whether taken by a person or by an AI agent acting for another, passes on an outcome. Unless the basis travels with it, each step leaves the banking organization knowing less about what it relies on.

4. Recommendations

Each recommendation below is illustrative and non-prescriptive. None would create a checklist or a new requirement, and the proposal's tailoring principles would govern how far each applies.

Recommendation 1: Illustrate "credible" as the proposal already illustrates credible indemnification. The only other use of "credible" in the proposal, "credible indemnification provisions," comes with illustrative factors: "the ability of the relevant party to meet any obligations," terms "relevant to the specific risks of the third-party relationship," and "confirming its ongoing sufficiency." The agencies could do the same for summaries and certifications. After the sentence on credible summaries, they could add language along these lines: "In assessing whether a summary, certification, or assessment result is credible, a banking organization may consider who produced it, under what authority or accreditation, and whether they are independent of the third party. It may also consider which standard and version were applied, whether the scope fits the banking organization's use, what evidence the result rests on, and whether it remains current. Where a model contributed to the result, a banking organization may consider whether it identifies which model and version was used, and what the responsible person did with the model's output."

Recommendation 2: answer footnote 15 narrowly. Footnote 15 asks to what extent a list of characteristics indicating a high-risk relationship would help. If the agencies include one, I suggest a single characteristic, tied to the proposal's own test of a non-trivial violation of law or material harm. The characteristic has two parts. The third party makes decisions for the banking organization that could themselves cause such a violation or harm, such as credit, account-opening, sanctions, or fraud decisions. And the basis of those decisions does not reach the banking organization. It weighs more where those decisions are produced or materially shaped by generative or agentic AI. This does not call for AI-specific guidance. It asks only that the guidance name the question.

Recommendation 3: record the reliance as well as the decision. Under governance, item (5) suggests documenting "key elements of risk management for third-party relationships." The agencies could note that, for higher-risk relationships, this may include recording what the

banking organization relied on, and for what purpose, when it accepted a third party's assessment or decision. Where it accepted residual risk "based on limited information," it may also include what information was lacking. The proposal says examiners will give "due consideration" to a banking organization's reasonable judgment. That consideration is easier to give, and to justify later, when the basis of the judgment was recorded at the time.

Recommendation 4: avoid a new single point of dependence. The proposal recognizes consortia and "standard-setting and certification organizations" as ways to share the cost of due diligence. Where they make the basis of their conclusions checkable without disclosing proprietary detail, that can support reliance on them, including by community banks that cannot repeat the assessment. If common formats for this emerge, one controlled by a single provider would itself become the kind of concentration the guidance asks banking organizations to manage.

5. Why this fits the proposal's aims

These changes are framed as considerations, not expectations. They can lower the cost of reliance, because a banking organization that can see the basis of an assessment need not repeat it. And they support innovation in the sense of Executive Order 14405. A newer third party with a short operating history may be unable to supply all the due diligence information a banking organization requests, as the proposal recognizes. It can still show the basis of the decisions it makes for the banking organization. That gives banking organizations a way to rely on new entrants with confidence.

6. Limits of this comment

The interviews behind my research are testimony, not checked against case files, and were not randomly sampled. Most were with insurance practitioners in the United Kingdom and the United States; those in banking were exploratory. This comment about models goes beyond what practitioners described. Nothing in this comment rests on the interviews. The points above stand on the text of the proposal and on the published sources cited.

Disclosure

I submit this comment in my personal capacity, and the views are my own. The research is self-funded and has received no external funding. I founded the Verida Charter Foundation, a nonprofit standards body, and serve as its Founding Trustee and Senior Fellow. Until its anchor members are seated, I approve changes to its specification. Two of the papers listed below propose a record of the kind Recommendations 1, 3 and 4 describe, and name the Foundation as one possible steward of it. By the tests those papers set, it does not yet qualify. Its charter bars license fees, usage royalties, and certification fees, and I receive no royalty or fee from any record format those papers describe. I also have an interest in the infrastructure this comment points toward. Readers should weigh the recommendations with that in view. This letter is not a submission of the Foundation.

Respectfully submitted,

Alexander D. Barrett

Sources and related working papers

Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency and Federal Deposit Insurance Corporation, Revised Guidance on Model Risk Management, issued with SR 26-2, April 17, 2026 (superseding SR 11-7 and SR 21-8). Cover letter and attachment, footnote 3.

European Banking Authority, Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposures, EBA/GL/2017/16, paragraphs 201 to 204.

Barrett, A. D., Boundary Blindness in Insurance: Early Findings from Eight Months of Senior-Practitioner Research, and an Invitation to Phase Four, June 2026. SSRN 6966099.

Barrett, A. D., The Decision Economy: How Organizations Rely on Decisions They Cannot Check, July 2026. SSRN 7201998.

Barrett, A. D., Boundary Blindness Under Artificial Intelligence: Early Cross-Industry Findings on the Missing Decision-Evidence Layer, July 2026. SSRN 7210798.

Barrett, A. D., Signed, Not Seen: The Evidence Nobody Is Required to Carry, August 2026. Selected as 2026 Featured Research, DC Fintech Week. SSRN 7556438.

Barrett, A. D., The Machine-Action Boundary: AI Control at the Point Where the World Says Yes, October 2026. SSRN 7548730.

Barrett, A. D., The Missing Evidence Layer of the UN 2030 Agenda: How Decision Basis Records Could Make Institutional Accountability Portable, October 2026. SSRN 7585738.