STEPHEN BISHOP, ET. AL.

Proposal and Comment Information

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

Comment ID: FR-2026-0031-01-C31

Subject

Docket No. OP-1881: Comment on Proposed Third-Party Risk Management Guidance

Submitter Information

Name: Stephen Bishop, et. al.

Submitted Date: 11/01/2026

NONCONFIDENTIAL // EXTERNAL

To the Board of Governors of the Federal Reserve System, Docket No. OP-1881

Below is a comment letter from Stephen Bishop and Tony del Fierro on the Proposed Third-Party Risk Management Guidance, 91 FR 58536 (September 15, 2026). Recommendation 5 also responds to the Board's companion Proposed Third-Party Risk Management Guide for Traditional Community Banking Organizations (Docket No. OP-1880). The same letter is being filed with the Office of the Comptroller of the Currency, the Federal Deposit Insurance Corporation and the National Credit Union Administration. Questions may be directed to Stephen Bishop at steve@ambaassador.com<mailto:steve@ambaassador.com> or

---

November 1, 2026

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

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

Ladies and Gentlemen:

We write as the co-authors of Know Your Agent: The Governance Standard for AI in Financial Services (Examiner Press, July 2026), a governance framework for artificial intelligence (AI) agents in banks, credit unions and their partners. Stephen Bishop has spent his career in banking and financial technology, including sponsor banking and embedded finance programs, and built a 115-item enhanced due diligence package for fintech partners. Tony del Fierro is a payments and infrastructure executive who served as senior vice president and chief technology officer of Sound Credit Union and now works in connectivity strategy at a large bank. We submit this letter in our individual capacities.

We support the direction of the proposal. Putting risk identification and assessment first, scaling oversight to the risk a relationship presents, recognizing that residual risk is accepted rather than eliminated, and giving due consideration to a banking organization's reasonable judgment are the practices operators have asked for. The four-component structure is easier to run than the five-stage life cycle it replaces, and the joint statement on core service providers names problems community banks have described for years.

Our comments concern one subject the proposal does not mention. The proposed guidance text uses the word "model" once, in the phrase "business model," and does not use the terms artificial intelligence, machine learning, automated decisioning or agent at all. Meanwhile, the agencies' revised model risk management guidance of April 17, 2026 states that generative AI and agentic AI models "are not within the scope of this guidance." As a result, an AI system that takes actions on a banking organization's behalf through a vendor's product sits in neither document. In our conversations with bank and credit union leaders this year, that is where the hardest examiner questions land: what is this system, what is it allowed

to do, who owns it, and what evidence supports trusting it. We offer six recommendations, each tied to a question the agencies asked or to text in the proposal.

1. Name AI-enabled action as a characteristic of higher risk (footnote 15). The agencies ask whether a list of characteristics that generally indicate a high-risk relationship would help. It would, provided the list is illustrative. We recommend one characteristic that no existing third-party risk list captures: the third party's product includes an AI system that takes actions or makes decisions within the banking organization's operations or customer relationships with limited human review before the action takes effect. Examples include initiating or releasing payments, changing account status or limits, communicating with customers, resolving disputes, and shaping a credit or compliance decision. The characteristic fits the proposal's two-part test. On magnitude, an acting system can produce a non-trivial violation or customer harm at the speed and scale of software. On likelihood, the behavior of these systems changes when the underlying model changes, often with no change to the contract. We call this the system's practical authority: what it can do in practice, whatever the product is called. A relationship that was lower risk when the product displayed information may become higher risk when the same product begins to act. The reassessment factors the proposal already lists (changes in reliance, new concerns, a change in the services provided) should expressly include a change in the practical authority of an AI-enabled feature.

2. Keep the written-agreement default, and say that AI systems acting inside the bank are in scope whichever contract carries them (footnote 9). A written-agreement test clarifies scope for most relationships, and we see its advantages. Its disadvantage is specific to AI. Three common arrangements would fall outside a strict reading: an AI feature a business unit provisions under a platform's click-through terms; an AI feature in a fintech partner's product that reaches the bank's customers under a contract between the bank and the fintech, with the model provider two steps removed; and a model embedded in a core provider's product that the bank never contracted for separately. In each case the bank bears the risk of the system's actions while the agreement sits elsewhere. We recommend the agencies retain the written-agreement observation and add one sentence: an AI system that acts within a banking organization's operations or customer relationships is within the scope of this guidance regardless of the contractual path by which it arrived, and the banking organization may oversee it through whichever party holds the contract.

3. Treat model changes as the subcontractor question of this decade. The subcontractor paragraph observes that a third party's use of subcontractors "may lessen the banking organization's control of activities" and points to contract terms and the third party's own risk management program as oversight tools. For AI-enabled products, the subcontractor that matters is the model provider, and the event that matters is a model change. A vendor can change the model underneath a feature, and the feature's behavior with it, without any change in the services the contract describes. We recommend the paragraph name this case and offer two illustrative practices: a contractual notice obligation when the model behind an AI-enabled feature changes materially, and treatment of such a change as a reassessment factor under the risk assessment section. Both are consistent with the proposal's statement that a banking organization "may be able to exercise effective risk oversight" through contract terms, and neither asks a community bank to supervise a large model provider directly, a point on which we agree with the May 2026 joint report of the Consumer Bankers Association, the American Fintech Council, the Coalition for Financial Ecosystem Standards and the Independent Community Bankers of America.

4. Give banks and examiners a shared, illustrative evidence set for AI-enabled relationships (governance practice 5). The governance section encourages banks to "document key elements of risk management for third-party relationships," and the proposal is careful that its examples are "illustrative only." We recommend one illustrative example for AI-enabled relationships, so that banks and examiners use the same vocabulary. The evidence set we use with institutions has five parts: an inventory of AI-enabled

features in third-party products and what each is permitted to do; a named owner for each; a bounded statement of authority, including limits and the point at which a person reviews; monitoring proportionate to that authority; and triggers for re-review. A community bank can keep this on one page per relationship; a large bank can keep it in a system. The Conference of State Bank Supervisors published an AI Supervisory Framework for state examiners on September 16, 2026, with an examiner guide and work program. Since the agencies and the states examine many of the same institutions, we urge the agencies to align any illustrative evidence language with that framework, so that a bank is not asked for two different files by two different examiners.

5. Add AI disclosure to the core provider statement's transparency factor, and treat AI as a feature inside the Board guide's existing vendor categories (Board guide, question on additional categories). The Joint Statement on Community Banks' Engagement with Core Service Providers lists transparency factors that will shape supervisory allocation decisions, including a provider's "willingness to provide reasonably relevant and timely due diligence information" and its "transparency and timely disclosure regarding operational issues and security incidents." We recommend adding a provider's disclosure of which products contain AI-enabled features, what those features are permitted to do, which model providers stand behind them, and when those models change. A community bank cannot inventory what its core provider will not disclose, and today many cannot find out; one of us has used a short inventory letter to core providers for this purpose, and the answers vary widely by provider. On the Board's proposed community bank guide, we do not recommend a separate AI vendor category. AI arrives as a feature inside the categories the guide already lists: automated credit decisioning in loan management systems, alert logic in financial crime platforms, calibrated models in fraud tools, and assistants in digital banking. We recommend a short cross-cutting consideration in Section IV.B that applies to every category: identify the AI-enabled features, confirm the bank can set or review the parameters that matter to its risk appetite, and confirm it will be told when the model changes. The guide's existing language on "automated credit decisioning" and on a provider's ability to "calibrate its model based on the TCBO's direction" is the right starting point.

6. Say plainly that customer-brought agents are outside this guidance, and where they will be addressed. A growing share of the AI agents reaching banks belong to customers, who hand an assistant a login and a task. Under the proposal's definition no business arrangement exists between the bank and that agent, so third-party risk management does not apply, and it should not; the right disciplines are authentication, authorization and the customer agreement. We recommend the agencies say so in the final guidance, because in its absence some institutions will force customer agents into vendor programs and others will conclude the agencies have no view. The Chairman of the Federal Deposit Insurance Corporation said in October that the agency is considering soliciting information from the public on AI. The customer-brought agent is the case most in need of it, and we would welcome the chance to contribute.

A further note on embedded finance. The proposal would replace the July 25, 2024 joint statement on banks' arrangements with third parties to deliver deposit products without successor text. Whatever the agencies decide about that statement's form, its substance on account records, ledgering, reconciliation and the bank's access to customer-level data remains the practical foundation of a safe embedded finance program. It is also the foundation on which any oversight of AI-enabled features in partner products rests: a bank that cannot see its own customers' records cannot see what an agent did to them. We encourage the agencies to preserve that substance.

We appreciate the agencies' work and the clarity of the proposal's structure. We would welcome the opportunity to discuss these recommendations with staff, and we can provide the evidence set described in recommendation 4 on request.

Respectfully submitted,

Stephen Bishop
Co-author, Know Your Agent: The Governance Standard for AI in Financial Services
steve@ambaassador.com<mailto:steve@ambaassador.com> |

Tony del Fierro
Co-author, Know Your Agent: The Governance Standard for AI in Financial Services