

ARROW INTELLIGENT SYSTEMS, INC., KARTHIKL KRISHNAN

Proposal and Comment Information

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

Comment ID: FR-2026-0031-01-C14

Submitter Information

Organization Name: Arrow Intelligent Systems, Inc.

Organization Type: Company

Name: Karthikl Krishnan

Submitted Date: 10/07/2026

We appreciate the opportunity to comment on the Proposed Third-Party Risk Management Guidance published September 15, 2026. Please see attached document for our detailed comments.

ARROW INTELLIGENT SYSTEMS INC.

arrowintsys.com · contact@arrowintsys.com

October 7, 2026

Benjamin W. McDonough, Secretary
Board of Governors of the Federal Reserve System
20th Street and Constitution Avenue NW, Washington, DC 20551

Re: Third-Party Risk Management Guidance — Board Docket No. OP-1881

(Interagency proposal: OCC Docket ID OCC-2026-0793, RIN 1557-ZA18; FDIC ZRIN 3064-ZA58; NCUA Docket No. NCUA-2026-1684, RIN 3133-ZA54)

To the Board of Governors of the Federal Reserve System:

Arrow Intelligent Systems Inc. ("Arrow") appreciates the opportunity to comment on the Proposed Third-Party Risk Management Guidance published September 15, 2026 (91 FR 58536) by the OCC, the Board of Governors of the Federal Reserve System, the Federal Deposit Insurance Corporation, and the National Credit Union Administration (together, the "agencies") (the "Proposed Guidance"). Because the Proposed Guidance is interagency, we are providing copies of this letter to the OCC, the FDIC, and the NCUA.

Arrow builds workflows and agentic systems for compliance operations, and we are creating new tools to test whether AI systems used in regulated financial services apply their business rules accurately and consistently. In the interest of transparency, independent validation of AI systems is part of our business. Our comments address what evidence supports sound oversight of AI third parties, not any particular product.

Summary

As proposed, no federal guidance squarely addresses a generative or agentic AI system that a banking organization obtains from a vendor. The agencies' April 2026 revised model risk management guidance excludes generative and agentic AI models as "novel and rapidly evolving" (SR 26-2, n.3; OCC Bulletin 2026-13), and the Proposed Guidance does not mention artificial intelligence. Banking organizations, and the fintech companies that partner with them, are considering acquiring AI capabilities through third parties, in credit decisioning, servicing, complaint handling, fraud, and compliance. For those systems, third-party risk management is now the main framework a banking organization has. More importantly, in our discussions with some banks, the lack of direct guidance is making some reluctant to adopt these technologies until expectations for vendor assessment are clearer.

We support the Proposed Guidance's risk-based, principles-based approach, its recognition that newer and innovative third parties can "enhance access to financial products and services," and its goal of replacing "check-the-box exercises" with oversight focused on material risk. We do not ask the agencies to add prescriptive requirements or checklists, and we understand why the agencies removed detailed examples from the 2023 Guidance. Our concern is the reverse problem: for AI third parties, silence also creates ambiguity, and ambiguity is costly.

Today, a banking organization evaluating an AI vendor has no clear answer to a basic question: what evidence is enough? Each banking organization sends its own questionnaire. Vendors may be unwilling to share code or AI model weights to protect their IP. The banking organization then either accepts answers it cannot evaluate or treats every AI vendor as high risk, the outcome the Proposed Guidance seeks to avoid. Fintechs

and vendors that serve several banking organizations answer the same questions many times, in different forms.

The evidence that resolves this is direct and inexpensive to produce: whether the system produces correct outputs, whether it produces the same output when given the same input, and whether its behavior changes when the vendor updates its model. Generative and agentic systems can pass every individual test and still decide the same case differently from one run to the next. Testing of this kind uses synthetic test cases rather than customer data, which simplifies testing and protects customer privacy. It measures how the system behaves without requiring the vendor to disclose its code, model weights, or training data, which protects the vendor's intellectual property. And a vendor can test once and share the results with every banking organization it serves.

Brief, permissive recognition of that evidence would reduce burden for everyone involved. Banking organizations would know what they may rely on and could right-size their oversight of AI vendors that provide it. Vendors and fintechs could produce one body of evidence instead of answering many questionnaires. Examiners would have a common reference point when giving "due consideration" to a banking organization's judgment. The result is a clear and consistent approach that banking organizations, fintechs, and vendors can all follow, without a checklist. This would allow banking organizations and regulated fintechs to extract productivity gains from AI and agentic systems at a reasonable and manageable compliance cost.

These recommendations are not new to the agencies. For the vendor models that SR 26-2 does cover, the agencies already describe validation of vendor products "either by internal or outside parties," and ongoing monitoring and outcome analysis to assess whether vendor models "are accurate, remain fit for purpose, and continue to be reliable," even where a banking organization does not receive "the underlying code, data, or methodology" (SR 26-2, Section VII). SR 26-2 also calls for evaluating adjustments made to customize vendor models, and for review by experts with "sufficient independence to maintain objectivity" (SR 26-2, Sections III and VII). Our recommendations apply the same principles to the AI systems SR 26-2 excludes.

We recommend three short additions, each permissive and each placed where the Proposed Guidance already invites it:

1. Risk assessment (Section B.1.iii; footnote 15): identify AI systems that materially influence customer outcomes as a characteristic that may indicate higher risk, while recognizing that documented testing can support a lower assessment of a particular relationship; and recognize vendor model updates as a changing circumstance that may warrant reassessment.
2. Due diligence (Section B.2.i): recognize documented accuracy and consistency testing, including testing on synthetic data, as evidence that may be adequate for due diligence where a vendor does not share proprietary information.
3. Ongoing monitoring (Section B.2.iii): illustrate "periodic testing of the third party's services" for AI systems as repeated testing of the system as configured, repeated after vendor updates.

We also comment on the scope question in footnote 9 and the rescission question in footnote 5, and identify three points the agencies could address in the preamble, a frequently-asked-questions document, or forthcoming AI guidance rather than in the text of the guidance.

1. Risk assessment

Proposed Guidance, Section B.1.iii (Risk Assessment) and footnote 15.

Footnote 15 asks whether the guidance should include characteristics that generally indicate a third-party relationship is high risk. We suggest one: the third party provides an automated or AI system whose outputs materially influence decisions about customers, such as credit, pricing, servicing, or complaint handling or resolution; or interactions with customers, such as customer service chatbots. Under the Proposed Guidance's own test, these are relationships in which non-business-as-usual performance "could cause the banking organization to incur an actual non-trivial violation of law or regulation," such as fair lending and adverse action requirements under the Equal Credit Opportunity Act and Regulation B, or prohibitions on unfair, deceptive, or abusive acts or practices. State requirements governing automated decision-making in consumer decisions are also developing.

Identifying this characteristic need not increase burden. Footnote 15 contemplates characteristics that may indicate higher risk, not conclusions. A banking organization that obtains credible testing evidence showing that a system performs accurately and consistently for its intended use could reasonably assess that particular relationship as lower risk and tailor its oversight accordingly. The characteristic tells a banking organization where to look; the evidence tells it how much oversight is warranted.

The Proposed Guidance also notes that "changing circumstances" may indicate that a relationship's risk has changed. For AI third parties, a common changing circumstance is invisible to the banking organization: the vendor updates its system architecture, model, prompts, or settings, and the system's behavior shifts without any change to the contract or services.

SUGGESTED LANGUAGE

Characteristics that may indicate a higher-risk relationship include a third party's provision of an automated or artificial intelligence system whose outputs materially influence decisions affecting customers or interactions with customers. Documented testing of such a system's accuracy and consistency may support a banking organization's assessment that a particular relationship presents lower risk.

For third parties that provide such systems, changes to the underlying system, model, its configuration, or its instructions may be changing circumstances that warrant reassessment, even where the contract and the services provided are otherwise unchanged.

2. Due diligence

Proposed Guidance, Section B.2.i (Due Diligence and Third Party Selection); Section B.2.v ("Use of Co-ventures, Consortia, ...").

Due diligence assesses "a third party's ability to perform the activity as expected." For an AI system, the most direct evidence of that ability is documented testing: the business rules or decisions the system was tested against, how test cases were built, the measured rate of correct outputs, and the measured rate at which the system returns the same output for the same input across repeated runs. Due diligence also considers the third party's "legal and regulatory compliance"; for an AI system that makes or informs decisions about customers, the same testing can help show whether its outputs are consistent with the laws that apply to those decisions.

The Proposed Guidance recognizes that a third party "may not share (or is not permitted to share) information that the banking organization requests," and that "credible summaries of technical data or proprietary information, as well as certifications or results of assessments," may be adequate. AI vendors

commonly decline to share model weights, training data, or source code. Testing results do not require that information, so they bridge the gap: a banking organization can learn how a system behaves without seeing how it is built.

Two features make such testing credible and practical. First, it should be documented well enough to be reproduced, and it should be performed or reviewed by a person or system independent of the developers. Second, it can use synthetic test cases rather than customer or production data. Synthetic test cases are simpler to build and share, can cover rare or sensitive scenarios that production data may not contain, and keep customer data out of the testing process altogether. Because the testing measures the system's outputs rather than its internals, it also does not require the vendor to expose its code, model weights, or training data. Customer privacy and vendor intellectual property, two common obstacles to AI due diligence, are both addressed by the method itself.

SUGGESTED LANGUAGE

For third parties that provide automated or artificial intelligence systems, documented testing of the system's accuracy and of the consistency of its outputs when given the same inputs may help a banking organization assess the third party's ability to perform the activity as expected, and may also help it assess whether the system's outputs are consistent with applicable laws and regulations, including consumer protection and fair lending laws. Such testing may use synthetic data rather than customer data, and may be performed or reviewed by the third party, the banking organization, or an independent party. Where the third party does not share proprietary technical information, the results of such testing may be adequate for the banking organization's due diligence needs, depending on facts and circumstances.

3. Ongoing monitoring

Proposed Guidance, Section B.2.iii (Ongoing Monitoring); Section B.2.v ("Subcontractors"; "Insurance, Indemnification, and Other Limitations on Liability").

Among its examples of monitoring, the Proposed Guidance lists "periodic testing of the third party's services to the banking organization to assess the third party's performance and effectiveness." For AI systems, two refinements make that testing meaningful:

- **Repeat identical inputs.** A single test run cannot detect inconsistency. Running the same test cases several times shows whether the system gives the same answer each time, and running them again after a vendor update shows whether its behavior changed.
- **Test the system as configured.** Most deployed AI systems combine the vendor's model with the banking organization's own policies, rules, or instructions. A vendor's generic results do not show how the system behaves with those policies. Testing the configured system does, and it also shows whether a problem originates in the vendor's model or in the banking organization's configuration. That distinction informs the allocation of responsibility under contract and the value of any indemnification.

The same approach works when AI reaches a banking organization indirectly. Fintech partners and other third parties increasingly rely on their own AI vendors, which are subcontractors from the banking organization's perspective. The Proposed Guidance notes that a banking organization may oversee subcontractors by "negotiating and monitoring compliance with contractual terms governing the use and oversight of

subcontractors" or by "assessing, confirming, and monitoring the adequacy of its third party's own third-party risk management programs." For AI subcontractors, asking the third party to provide test results for the AI systems it relies on is a practical way to do both, and avoids the cost of direct oversight of each subcontractor.

SUGGESTED LANGUAGE

For third parties that provide artificial intelligence systems, periodic testing may include running the same test cases more than once to assess the consistency of outputs, testing the system as configured with the banking organization's own policies or instructions, and repeating testing after material changes to the system.

4. Scope and rescission (footnotes 9 and 5)

Proposed Guidance, Section A (Third-Party Relationships), footnote 9; Section I, footnote 5.

Footnote 9 asks whether the guidance should apply only to third parties subject to a written agreement. AI often reaches a banking organization without a separate agreement: as a feature added to an existing core, servicing, or compliance product, or through a fintech partner's own AI vendor. A written agreement limit can work if the agencies make clear that it does not exclude AI capabilities added to products or services already provided under an existing agreement. Otherwise, the scope test could leave a fast-growing source of automated decisions outside third-party oversight altogether.

The same point bears on the agencies' companion statement on community banks' engagement with core service providers. Core providers increasingly add AI features to products community banks already use. Testing evidence that a core provider produces once, for all of its client banks, is an efficient way for community banks to understand those features without negotiating for it individually.

Footnote 5 asks whether other guidance should also be rescinded. We take no position on specific documents, but suggest that the agencies not rescind guidance bearing on vendor-provided models or technology without confirming that the final guidance, or forthcoming AI guidance, covers the same ground.

5. Additional considerations

Consistent with the agencies' preference for fewer examples in the text, we suggest addressing the following outside the guidance itself:

- **Contract terms.** We agree that there are "no generally applicable expected contract terms," and that a service level agreement "can define performance measures." For AI systems, two measures are especially useful and inexpensive to provide: accuracy and consistency rates on an agreed set of test cases, and notice of material model changes accompanied by updated test results. Noting these as options would help banking organizations, particularly those with limited negotiating power, know what to ask for without implying an expected contract term.
- **Shared and independent testing.** The Proposed Guidance recognizes that consortia, certifications, and other third parties "may allow community banks, in particular, to benefit from advanced technology and industry expertise otherwise unavailable." AI testing is a clear case. In our experience, few community banks employ machine-learning specialists, and many use the same

vendors (both with and without AI capabilities). Shared or independent testing lets them evaluate those vendors at a fraction of the cost of building the capability themselves. The need extends beyond AI: because SR 26-2 is "expected to be most relevant to banking organizations with over \$30 billion in total assets," community banks are likely to rely on third-party risk management for most vendor-provided models. As the Proposed Guidance notes, reliance on third parties for risk management carries its own risks; for testing, the best safeguard is transparency of method, so a banking organization can see what was tested, how, and with what results.

- **Coordination with forthcoming AI guidance.** Because SR 26-2 excludes generative and agentic AI, a vendor-provided AI system may currently fall outside both model risk guidance and, as proposed, any AI-specific discussion in third-party risk guidance. We encourage the agencies to coordinate any forthcoming AI guidance with the final Third-Party Risk Management Guidance, so the same evidence of accuracy and consistency can serve both purposes and banking organizations, fintechs, and vendors are not asked to produce it twice.

Conclusion

The Proposed Guidance rightly shifts third-party oversight from process to assessed risk. For AI third parties, assessed risk depends on evidence of how the system behaves across repeated decisions and over time. Recognizing that evidence in a few permissive sentences would give banking organizations, fintechs, vendors, and examiners a clear and common reference point; let one vendor's testing serve many banking organizations at once, using synthetic data that protects customer privacy and vendor intellectual property; and let banking organizations adopt AI from innovative third parties with confidence while right-sizing their oversight, rather than treating every AI vendor as high risk.

Arrow would welcome the opportunity to meet with agency staff, provide technical detail on testing methods, or demonstrate how consistency is measured in practice.

Respectfully submitted,

/s/ Karthik Krishnan

Karthik Krishnan
Founder and President
Arrow Intelligent Systems Inc.
contact@arrowintsys.com

cc:

Chief Counsel's Office, Office of the Comptroller of the Currency (Docket ID OCC-2026-0793)
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)