POLSINELLI

Publications

November 17, 2026 • Updates

OCC's New AI Supervisory Framework: Cybersecurity, Privacy and Vendor Management Takeaways for Banks

Key Updates

- On Nov. 4, 2026, the Office of the Comptroller of the Currency (OCC) issued its dedicated supervisory framework for examiners evaluating banks' use of generative and agentic artificial intelligence (AI).
- Although the framework does not impose new regulatory requirements or supervisory expectations, it provides insight into how OCC examiners may evaluate AI-related risks, including cybersecurity, customer data privacy, third-party risk management and operational resilience.

Why It Matters

- The framework addresses increasingly autonomous AI technologies that can access sensitive customer information, interact with banking systems and execute transactions. These capabilities create significant opportunities for banks but also raise security and privacy concerns that traditional risk management programs may not fully address.

Next Steps

- Banks and other financial institutions should consider reviewing their existing AI uses, governance, security controls and vendor arrangements, and developing AI-specific contractual protections and incident response procedures where appropriate.

On Nov. 4, 2026, the OCC issued *Bulletin 2026-54, Super Intelligence: Supervisory Framework*, introducing a dedicated examiner framework for evaluating banks' use of "super intelligence," including generative and agentic artificial intelligence (collectively, "AI" for purposes of this article). The development is notable because the federal banking agencies' April 2026 revised model risk management guidance expressly excluded generative and agentic AI from its scope. The new OCC framework addresses those technologies through 16 components organized around six pillars: purpose, risk, control, oversight, resilience and assurance.

Related People

- Alexander D. Boyd

Related Capabilities

- Privacy & Cybersecurity
- Banking & Financial Institutions

The OCC expressly states that the framework is intended for examiners, does not establish new regulatory requirements or supervisory expectations, and is not intended as a required self-assessment tool for banks. Nevertheless, it identifies important considerations for evaluating the risks associated with increasingly autonomous AI systems.

Although directed at OCC examiners supervising national banks, federal savings associations and federal branches and agencies of foreign banks, the framework's considerations may also be useful to other financial institutions, including state-chartered banks and credit unions.

1. What Are the Six Pillars of the OCC's AI Supervisory Framework?

The OCC's framework organizes its supervisory considerations into six pillars: **purpose, risk, control, oversight, resilience and assurance.** Together, these pillars address how a bank uses an AI system, the risks associated with that use, the controls in place to prevent harmful activity and the bank's ability to respond when those controls fail.

In practical terms, the framework identifies several important risk management considerations:

- **Define permitted uses and assess risks.** Banks should understand what each AI system is authorized to do, what information and systems it can access, and the potential consequences if it operates incorrectly or exceeds its intended purpose. Higher-risk applications, particularly those capable of accessing sensitive customer information or taking autonomous actions, may warrant more extensive assessments and safeguards.
- **Establish controls and meaningful oversight.** Banks should consider controls designed to prevent unauthorized, inappropriate or unsafe AI actions, including restrictions on system permissions, human approval for consequential decisions, ongoing monitoring and clear accountability for AI activities.
- **Build resilience and verify that controls work.** Banks should consider how they would detect, contain and recover from AI failures or unauthorized activity while maintaining critical operations. This includes testing controls, maintaining records sufficient to reconstruct significant AI actions, and establishing procedures to suspend AI systems or reverse harmful actions when necessary.

2. What Cybersecurity and Data Privacy Risks Does the OCC Identify?

The OCC recognizes that AI introduces privacy risks beyond conventional data handling concerns. AI tools may use sensitive customer information and then retain or reproduce it, infer information that was not expressly provided or use data in ways that are inconsistent with existing privacy policies.

The framework also identifies cybersecurity threats specific to AI, including prompt injection, data and model poisoning, model extraction and sensitive information leakage. These threats may allow attackers to manipulate AI systems or obtain sensitive information without a conventional network intrusion.

Agentic AI presents additional concerns because these systems may use credentials, access internal applications, modify records or execute transactions. Excessive permissions or insufficient human oversight could allow a compromised or manipulated AI agent to cause significant financial, security or operational harm.

Financial institutions should consider whether their existing privacy assessments, access

controls, security testing and monitoring practices adequately address these risks.

3. How Should Banks Address AI Risks in Vendor Due Diligence, Oversight and Contracts?

Rather than build their own tools, banks often obtain AI capabilities through third-party technology providers. The OCC recognizes that conventional vendor due diligence and information security reviews may not fully address AI-related risks. For example, existing vendors may introduce new AI functionality, modify underlying models or deploy AI agents that interact with bank systems without corresponding changes to the institution's existing contractual protections.

The framework specifically addresses visibility into vendor-operated AI agents, notification of AI-related incidents, changes to AI functionality and contractual responsibility for model-related errors or customer harm.

Banks should consider whether existing vendor management practices sufficiently address these issues. Depending on the risks involved, an AI-specific contractual addendum may provide additional protections concerning approved AI uses, restrictions on customer data use and model training, downstream AI providers, material changes to AI functionality, incident reporting, investigation cooperation, audit rights and allocation of liability.

4. Do Banks Need an AI Incident Response Plan?

While some AI-related incidents may fall within the scope of existing data security incident response plans, AI use can result in a wider array of incidents that may not be adequately covered by existing response plans.

An AI agent could improperly disclose customer information, modify account records, generate inaccurate customer communications or execute unauthorized transactions without a traditional system compromise or outage. These events may require investigation, containment and customer remediation even when conventional cybersecurity monitoring does not identify an intrusion.

The framework identifies AI-specific considerations involving incident detection, containment, suspension of AI functionality, preservation of relevant records, business continuity and customer remediation.

Banks should consider developing AI incident reporting and response procedures that supplement existing cybersecurity incident response plans. These procedures should address not only conventional security incidents but also unauthorized AI use, harmful or anomalous outputs, privacy violations, compromised models and unintended autonomous actions.

5. What Should Banks Do in Response to the OCC's AI Framework?

While the OCC's framework does not impose new compliance requirements, it identifies risks that may warrant closer attention under existing privacy, cybersecurity and third-party risk management programs. Banks and other financial institutions should consider the following practical steps.

- **Inventory existing and proposed AI uses.** Identify AI systems currently deployed or under consideration, including AI functionality embedded in third-party products. Document permitted uses, data access, system permissions, autonomous capabilities and human oversight requirements. Establish approval procedures for new AI

applications and material changes to existing uses.

- **Reevaluate vendor vetting and develop AI-specific contractual protections.** Update third-party due diligence procedures to address AI capabilities, underlying model providers, data handling practices and security controls. Consider an AI-specific vendor contract addendum addressing approved uses, data protection, restrictions on model training, material changes, incident notification, investigation cooperation, audit rights and liability for unauthorized or harmful AI actions.
- **Develop an AI incident reporting and response policy.** Establish procedures for identifying, internally reporting, escalating, investigating and remediating AI-related incidents. Address incidents involving unauthorized AI use, privacy violations, security vulnerabilities, harmful outputs and unintended autonomous actions.
- **Establish AI-specific containment and recovery procedures.** Evaluate the institution's ability to suspend AI agents, revoke access permissions, preserve prompts and activity logs, reverse unauthorized actions and maintain critical operations through alternative processes. Incorporate AI-related scenarios into incident response tabletop exercises and business continuity testing.
- **Evaluate AI governance and risk management practices.** Assess AI applications in light of the risk factors and supervisory considerations identified in the OCC framework. Consider benefits, potential customer harm, access to sensitive information, system autonomy and human oversight.
- **Conduct AI-specific privacy and cybersecurity assessments.** Evaluate how sensitive customer information is accessed, processed, retained and potentially disclosed through AI systems. Assess whether existing controls adequately address AI-specific threats, including prompt injection, model manipulation, data leakage, unauthorized system access and excessive agent permissions.

The central issue for banks is whether existing controls remain effective when AI systems can access customer information, interact with critical systems and act without direct human approval. The OCC's framework provides a useful starting point for evaluating those risks.

For more information about the OCC's AI supervisory framework, financial institution cybersecurity and privacy requirements, AI vendor contractual protections, or incident response preparedness, contact Polsinelli's Privacy and Cybersecurity team.