Artificial Intelligence Supervisory Framework

Core Examiner Guide

Conference of State Bank Supervisors

Version 1.0

Approval Date: August 13, 2026

Guide Contents

Considerations and Background

This examiner guide helps examiners understand how institutions use artificial intelligence and assess the governance and risk management of that use. It is a practical tool for examiners of bank and nonbank institutions.

The examiner guide should be applied in a manner proportionate to an institution's size, complexity, risk profile, and use of artificial intelligence. Not all sections will be relevant in every exam, and examiner judgment remains important in determining the depth and focus of review.

The examiner guide is based on existing public materials, including supervisory guidance, regulator-issued materials, cross-sector risk management frameworks, and implementation resources. It does not create new legal obligations or supervisory requirements.

Overview

This examiner guide helps examiners identify whether and how an institution uses artificial intelligence and focuses review on the areas most relevant to that use. It includes the following sections:

1. Initial Scoping: Whether the institution uses artificial intelligence, the nature of that use, and which sections of the framework are most relevant.
2. Document Request List: Core documents to support examiner understanding of the institution's artificial intelligence use, governance, inventory, and related controls.
3. Governance and Oversight: Accountability, governance structures, reporting, and oversight of artificial intelligence use.
4. AI Inventory and Use Cases: Where and how artificial intelligence is used, the purpose of that use, and how use cases are evaluated based on risk.
5. Generative AI and Emerging Use: Generative artificial intelligence, large language models, emerging uses, and associated risks and controls.
6. Use of Existing Supervisory Resources: Use of existing supervisory resources for third-party and vendor review, model risk review, consumer protection review, and operational risk review.

Initial Scoping

Use these questions at the beginning of the supervisory event to identify whether the institution uses artificial intelligence, the nature of that use, and the most relevant sections of this framework. A negative or unclear response to IS-1 may warrant confirmation against the institution's vendor inventory, software inventory, approved tools, and recent product or platform changes before concluding that the framework does not apply.

Scoping Question	Yes/ No	If Yes, Consider Prioritizing
IS-1. Does the institution use artificial intelligence in any products, services, operations, compliance activities, or internal support functions?		• Governance and Oversight • AI Inventory and Use Cases
IS-2. Has the institution identified the artificial intelligence systems, tools, models, or use cases it currently uses?		• AI Inventory and Use Cases
IS-3. Is artificial intelligence used in customer-facing activities or in activities that support or influence decisions?		• Use of Existing Supervisory Resources (Consumer Protection Review; Model Risk Review)
IS-4. Does the institution rely on third parties, vendors, or external platforms for artificial intelligence capabilities?		• Use of Existing Supervisory Resources (Third-Party and Vendor Review)
IS-5. Has the institution sought to identify where artificial intelligence is embedded in third-party products and services it uses, and what information has been provided by vendors or service providers?		• AI Inventory and Use Cases • Use of Existing Supervisory Resources (Third-Party and Vendor Review) If embedded vendor AI is unknown, unclear, or not disclosed, consider Third-Party and Vendor Review.
IS-6. Does the institution use generative artificial intelligence, large language models, or similar AI tools?		• Generative AI and Emerging Use
IS-7. Does the institution differentiate among its AI use cases based on risk, impact, or required level of review?		• Governance and Oversight • AI Inventory and Use Cases • Generative AI and Emerging Use, as applicable • Use of Existing Supervisory Resources
IS-8. Does the institution enter, process, or transmit customer, consumer, confidential, or sensitive information through artificial intelligence systems, tools, or third-party platforms?		• Governance and Oversight • AI Inventory and Use Cases • Generative AI and Emerging Use, as applicable • Use of Existing Supervisory Resources

Examiner Next Steps: If no AI use is identified, no further review under this framework may be necessary.

If AI use is identified, use the Document Request List to request supporting materials. Then proceed to the relevant sections based on the institution's artificial intelligence use.

Document Request List

Based on the institution's use of artificial intelligence, examiners may request the following materials, as applicable:

Governance and Oversight

- AI-related policies, standards, or procedures
- Governance, committee, or oversight materials
- Management or board reporting on AI use and related risks
- Management information, dashboards, reports, or other materials used to monitor AI use, performance, issues, or related risks, where applicable.
- AI-related employee guidance or training materials, if applicable

AI Inventory and Use Cases

- Inventory or list of AI systems, tools, models, or use cases
- Documentation describing business purposes, ownership, and use
- Materials describing AI use cases, including any available risk assessments, risk ratings, or tiering information
- List of vendor products or services with known or potential embedded AI capabilities, where available
- Change management documentation for AI systems, models, tools, or use cases, where applicable

Generative AI and Emerging Use

- Policies, guidance, or restrictions related to generative AI or large language model use
- List of approved or known generative AI tools in use
- Documentation describing business purposes, data controls, and output review

Additional Materials, as applicable

- Vendor or third-party oversight materials related to AI capabilities
- Model-related validation, testing, monitoring, or control materials
- Consumer-facing compliance, notice, complaint, or remediation materials, where applicable
- Samples of recent AI-generated or AI-assisted consumer-facing outputs, where applicable, such as notices, chatbot transcripts, generated communications, or other customer-facing content
- Contract terms, service terms, or vendor documentation addressing the use, retention, sharing, or training of AI systems on customer, consumer, confidential, or institutional data, where applicable
- Internal audit reports, compliance reviews, control testing, tracking logs, or other independent assessments related to AI use, where available

Governance and Oversight

Objective

Enable examiners to determine whether the institution has established a governance structure for artificial intelligence use, including clear accountability, oversight, and processes for identifying and managing AI-related risks.

Areas for Examiner Review

- Roles and responsibilities for oversight of AI use
- Policies, standards, committees, or other governance structures related to AI
- Processes for evaluating and approving AI use cases based on risk
- Alignment of governance to the size, complexity, risk profile, and scope of the institution's AI use
- Reporting to management and, where appropriate, the board or equivalent oversight body
- Identification and documentation of risks associated with AI use
- Controls, monitoring, quality assurance, change management, and accountability across the AI lifecycle
- Coverage of both internally developed tools and externally sourced or embedded AI tools
- Internal audit, compliance review, control testing, or other independent review of AI governance, controls, or use cases, where applicable.

Procedures

GO-1. Identify whether the institution uses artificial intelligence and who is responsible for oversight of that use.

GO-2. Review policies, standards, committees, or other governance structures related to artificial intelligence, including generative AI or emerging capabilities.

GO-3. Review how AI use cases are evaluated for risk and how those evaluations are reviewed and approved.

GO-4. Review information provided to management and, where appropriate, the board or equivalent oversight body regarding artificial intelligence use and related risks.

GO-5. Review how AI-related risks are identified, documented, and updated over time.

GO-6. Review whether the institution has processes to test or monitor AI use, performance, limitations, and related risks over time, as applicable, including when AI systems, data, vendors, or operating conditions change.

GO-7. Determine whether governance and oversight apply to internally developed AI tools, externally sourced AI tools, and AI capabilities embedded in third-party products and services.

AI Inventory and Use Cases

Objective

Enable examiners to understand where and how the institution is using artificial intelligence, the purpose of that use, and the scope of AI use across the institution.

Areas for Examiner Review

- Inventory of AI systems, tools, models, or other uses across the institution and whether the inventory identifies key factors across use cases, such as internal/external, customer-facing, or decision-making/decision-support.
- Business purpose and business owner for each identified AI use case
- How AI use cases are evaluated and tiered based on risk
- Use of AI in core business lines, support functions, operations, compliance, or customer interactions
- Documentation describing how AI is used and where it is embedded in processes or products
- Processes for keeping the institution's AI inventory current as use evolves
- Awareness of AI capabilities embedded in third-party products, platforms, or services
- Connections between inventory items and related controls, monitoring, testing, validation, or other lifecycle review activities, where applicable.

Procedures

AI-1. Identify whether the institution maintains an inventory of artificial intelligence systems, tools, models, or functionalities in use.

AI-2. Review documentation describing the business purpose and role of AI for identified AI use cases.

AI-3. Assess whether artificial intelligence use is internal, customer-facing, decision-support, or decision-making.

AI-4. Review the institution's process for evaluating and tiering AI use cases based on risk (see risk tiering worksheet).

AI-5. Review documentation describing how and where artificial intelligence is used in institutional processes or products.

AI-6. Determine whether the institution is aware of artificial intelligence capabilities embedded in third-party products, platforms, or services.

Examiner Next Steps: If generative AI is identified, continue to Generative AI and Emerging Use. Utilize Use of Existing Supervisory Resources where artificial intelligence use involves third-party support, model-related use, or consumer-facing activities.

Generative AI and Emerging Use

Objective

Enable examiners to understand whether the institution is using generative artificial intelligence, large language models, agentic tools, or similar AI capabilities, how those tools are being used, and whether their tool type, deployment model, data controls, output oversight, and evolving risks may warrant additional review.

Areas for Examiner Review

- Use of generative AI tools, large language models, or similar AI capabilities across the institution
- Distinction between public, private, internally deployed, and vendor-provided generative AI tools
- Business purposes for generative AI use, including internal productivity, customer interaction, content generation, analysis, or decision support
- Controls over data entered into or processed through generative AI tools
- Review and oversight of outputs produced by generative AI tools
- Awareness of risks such as inaccuracy, hallucination, prompt injection, and data exposure
- Processes for monitoring, updating, or restricting emerging AI use as risks and capabilities evolve
- Permission boundaries, human checkpoints, logging, reversibility, and the ability to restrict or halt AI systems that can take actions with limited human direction

Procedures

GE-1. Identify whether the institution uses generative artificial intelligence tools, large language models, or similar AI capabilities.

GE-2. Review the business purposes for generative artificial intelligence use, including internal productivity, customer interaction, content generation, analysis, or decision support.

GE-3. Determine whether generative artificial intelligence tools are public, private, internally deployed, or vendor-provided.

GE-4. Assess controls over data entered into or processed through generative artificial intelligence tools.

GE-5. Review oversight of outputs produced by generative artificial intelligence tools.

GE-6. Assess whether the institution has identified risks associated with generative AI use, including inaccuracy, hallucination, prompt injection, and data exposure.

GE-7. Review processes for monitoring, updating, or restricting generative artificial intelligence use as risks and capabilities evolve.

GE-8. Where AI systems can take actions with limited human direction, review how the institution defines the actions the system is permitted to take, human checkpoints, logging, reversibility, and the ability to restrict or halt the system.

Use of Existing Supervisory Resources

Use this section when the institution's artificial intelligence use suggests the need for additional review under existing supervisory resources. Some AI-related issues may fall within supervisory areas for which examination resources already exist, including third-party and vendor review, model risk review, and consumer protection review. Examiners may use those resources, as applicable, together with this framework.

Third-Party and Vendor Review

Where artificial intelligence capabilities are provided, supported, or embedded by third parties, vendors, service providers, or external platforms, examiners may use existing third-party and vendor oversight resources, as appropriate. This may include review of due diligence, contractual arrangements, performance monitoring, controls, change management, and contingency planning related to those relationships.

Model Risk Review

Where artificial intelligence systems or outputs inform significant decisions, analysis, or operations, examiners may use existing model risk management resources, as appropriate. This may include review of development, use, validation, monitoring, documentation, assumptions, limitations, controls, and oversight applicable to the institution's use of artificial intelligence. Because some existing model risk resources may not fully address generative AI, agentic AI, or similar capabilities, examiners may use this framework and work program to identify AI-specific considerations that may warrant additional review.

Consumer Protection Review

Where artificial intelligence affects consumer-facing activities, communications, decisions, or outcomes, examiners may use existing consumer protection and compliance resources, as appropriate. This may include review of consumer impact, accuracy, explainability, notices, disclosures, adverse outcomes, consumer data handling, privacy or data-use concerns, fair lending, UDAP/UDAAP, and other applicable state or federal consumer protection requirements.

Operational Risk Review

Where artificial intelligence affects the institution's operations, internal controls, risk management, reporting, compliance, or ability to manage key activities, examiners may use applicable safety and soundness or operational risk resources, as appropriate. This may include review of governance, controls, risk assessment, reporting, change management, contingency planning, and compliance with applicable laws and regulations.