AI Roadmap

for Regulated Industries

A Practical Playbook for Secure, Compliant, Risk-Aligned AI Adoption

SBS CyberSecurity

The Safe Path Forward for AI
in Regulated Organizations

AI is reshaping how financial institutions and other regulated organizations operate. But for security leaders, the pressure is two-fold: adopt AI quickly enough to stay competitive while managing risk with enough rigor to satisfy regulators and auditors.

Regulators have already shifted their stance from broad awareness to expectations around governance, documentation, and oversight. They are no longer asking if your institution plans to use AI. They are asking:

Where is AI already in your environment?

Who is accountable for oversight?

What controls and policies are in place today?

Only **18%** of organizations have formally adopted AI, but **78%** of employees are already exposed to it through workplace tools.

Federal Reserve, 2026

For many institutions, AI is already present — often introduced quietly through third-party platforms, cloud tools, embedded features, or staff experimenting with generative AI — while the gap between AI use and organizational visibility continues to widen.

To help bring structure to this challenge, the following approach provides a practical, risk-aligned path forward.

Assess AI Governance Readiness

Build Your AI Implementation Roadmap

Identify Safe, High-Value Use Cases

Understand the Role of a vCAIO

Take Practical Next Steps

This roadmap is designed to support institutions at any stage of AI adoption — from early exploration to advanced use.

Understanding AI Risk and Opportunity in Regulated Environments

The 2 Pressures

Regulated organizations are balancing business demand with increasing regulatory expectations.

1. Business Demand for Efficiency and Innovation

Organizations want AI to:

- Reduce manual work
- Accelerate decision-making
- Improve customer/member experience
- Modernize internal processes

Boards are increasingly asking leadership how AI fits into the strategic plan.

2. Regulatory Demand for Explainability, Governance, and Controls

Regulators expect:

- Documented governance
- Accountability
- Policies
- Explainability
- Risk management

They expect leadership to understand and explain how AI influences decisions, even when those capabilities come from a vendor. For example, AI-driven outputs from fraud detection tools or underwriting support systems may influence decisions that institutions must be able to explain during an exam.

AI Governance Readiness

A Checklist for Governance, Accountability, and Oversight

Understanding where your institution stands today is essential before expanding AI use. Many organizations already have AI present across vendors, internal workflows, or employee use, often without full visibility or formal governance. This checklist helps assess whether foundational elements are in place to support secure, accountable, and exam-ready AI adoption.

1. Visibility and Inventory

Know where AI exists before risks surface.

- [ ] Do you know where AI exists across your organization, including vendor tools and embedded features?
- [ ] Do you maintain and update that inventory as capabilities evolve?

2. Executive Ownership and Accountability

Clear ownership ensures risks are monitored and decisions are documented.

- [ ] Is a clearly defined role accountable for AI risk and oversight?
- [ ] Are responsibilities, reporting, and escalation paths documented?

3. Governance Policies and Guardrails

Policies set boundaries for responsible AI use.

- [ ] Are policies in place for acceptable AI use, including generative tools?
- [ ] Are policies reviewed as capabilities and regulations evolve?

4. Data Protection and Security

AI introduces new risks to sensitive data.

- [ ] Have you assessed how AI interacts with sensitive or regulated data?
- [ ] Are controls in place to prevent data leakage or unauthorized sharing?

5. Vendor and Third-Party Risk Management

You remain accountable for vendor-driven AI.

- [ ] Do you understand how vendors use AI on your behalf?
- [ ] Are AI risks addressed in due diligence, contracts, and reviews?

6. Regulatory and Compliance Alignment

Examiners expect clear oversight and accountability.

- [ ] Can you demonstrate AI governance and oversight during an exam?
- [ ] Are compliance, risk, and legal teams involved early in AI initiatives?

7. Strategic Alignment and Roadmap

AI delivers value when aligned to strategy.

- [ ] Is AI aligned with your strategic plan?
- [ ] Do you have a prioritized, risk-aligned roadmap?

Bridge Governance Gaps with a vCAIO

Managing AI responsibly requires executive-level oversight, but full-time roles aren't always practical. SBS's Virtual Chief AI Officer (vCAIO) provides the structure and expertise to build and manage a responsible AI program.

With a vCAIO, organizations move forward with confidence.

link.sbscyber.com/vcaio

Interpreting Your Results

✔ **CHECKED MOST BOXES**
Strong foundation for scaling AI responsibly

✔ **CHECKED SOME BOXES**
Awareness present, but structure needs strengthening

✔ **CHECKED FEW BOXES**
Insufficient oversight and elevated risk

10 Steps to Securely Implement AI in Your Organization

These steps are designed to be sequential. Each builds on the governance, visibility, and accountability established earlier, reducing risk as AI adoption expands.

Phase 1: Establish the Foundation

Focus on visibility, ownership, and alignment before introducing AI.

1. Define AI Team and Ownership

Form a cross-functional team and designate an AI champion to lead planning and coordination of all AI-related efforts.

2. Inventory Existing AI Use

Take stock of the AI tools and capabilities currently in use. Identify what's working, what's missing, and where gaps exist.

3. Define Your AI Roadmap

Work with your team and other experts to build a clear AI roadmap. This should include short-term and long-term goals and align with your overall business strategy.

Phase 2:
Build Governance and Enablement

Establish governance, define expectations, and prepare for controlled AI adoption.

4. Set Risk and Governance Controls

Set up a framework to manage AI risks and ensure compliance with ethical and legal standards. Ensure data is secure, governed, and appropriate for AI use.

5. Define Acceptable AI Use

Write an acceptable use policy that clearly outlines how AI should be used responsibly in your organization. Review and update it regularly.

6. Select Tools and Partners

Evaluate AI tools and vendors for alignment with governance requirements, anticipated use cases, and key criteria such as functionality, security, scalability, and support.

Talk to SBS about building your AI roadmap.
link.sbscyber.com/vcaio

Phase 3:
Implement and Scale

Introduce AI use cases and expand adoption within a governed environment.

7. Train Your People

Develop staff training aligned with approved tools, governance expectations, and anticipated use cases to teach everyone how to use AI effectively.

8. Select High-Value, Low-Risk Use Cases

Identify and prioritize use cases that align with your governance framework, available tools, and trained staff.

9. Embed Cross-Functional Coordination

Host regular department meetings to explore how AI can improve workflows. Promote collaboration and share success stories across teams.

10. Continuously Monitor and Improve

Use feedback and lessons learned from current projects to update your AI strategy. Stay flexible so you can adapt to new technologies and changing needs.

AI Use Cases
in Regulated Institutions

These practical, low-risk starting points are well-suited for leadership teams exploring AI adoption.

Safe, High-Value Uses

These are most effective once foundational governance and controls are in place.

Security and Risk

- Pattern detection in threat intelligence
- Log summarization
- Fraud pattern detection
- Accelerating vendor AI risk reviews

Operations and Efficiency

- Inventory cleanup using summarization models
- Ticket triage
- Document summarization
- Reconciliation support
- Training content creation and reporting

Business and Decision Support

- Campaign content development and refinement
- Audience segmentation
- Portfolio analytics
- Decision support models

Risks to Monitor

As AI use expands, these risks require ongoing monitoring, clear controls, and defined accountability.

Data Risks

- Data privacy
- Data boundary enforcement
- Inadvertent exposure of sensitive data

Model Risks

- Model bias
- Source transparency
- Reliance on unverified outputs in security decisions

Governance and Operational Risks

- Shadow AI use
- Integration controls
- Explainability requirements for examiners

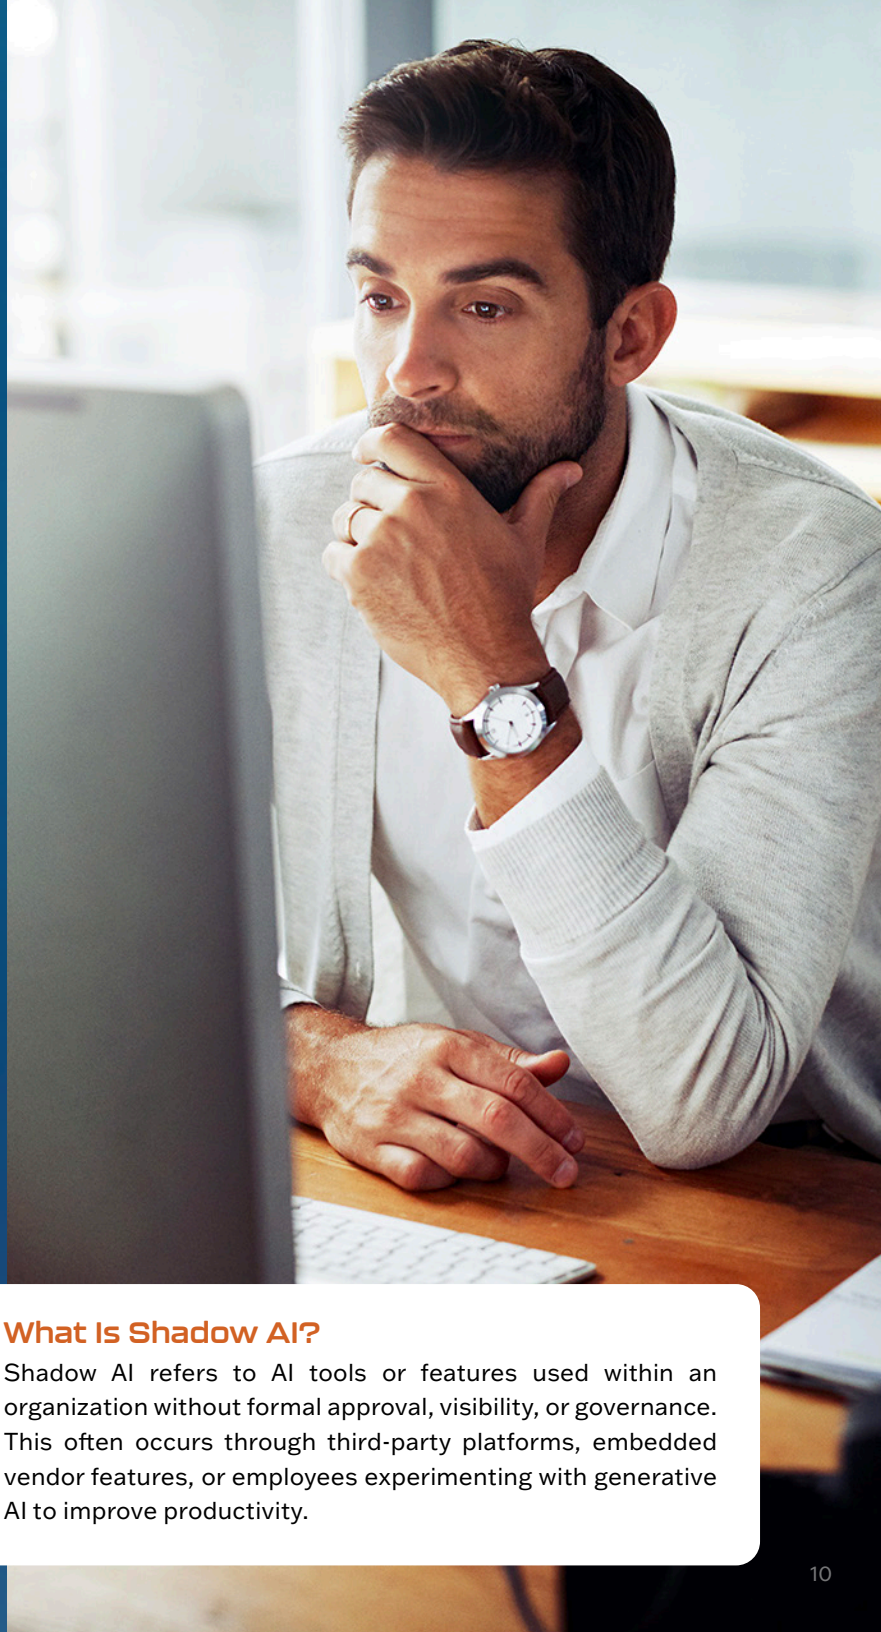

What Is Shadow AI?

Shadow AI refers to AI tools or features used within an organization without formal approval, visibility, or governance. This often occurs through third-party platforms, embedded vendor features, or employees experimenting with generative AI to improve productivity.

How a vCAIO Works
with Your Team

SBS CyberSecurity's Virtual Chief AI Officer (vCAIO) program provides executive-level oversight to ensure AI adoption is governed, accountable, and aligned with institutional risk and strategy.

Rather than operating as a full-time internal role, a vCAIO engages through structured touchpoints, defined responsibilities, and ongoing advisory support tailored to your organization's needs.

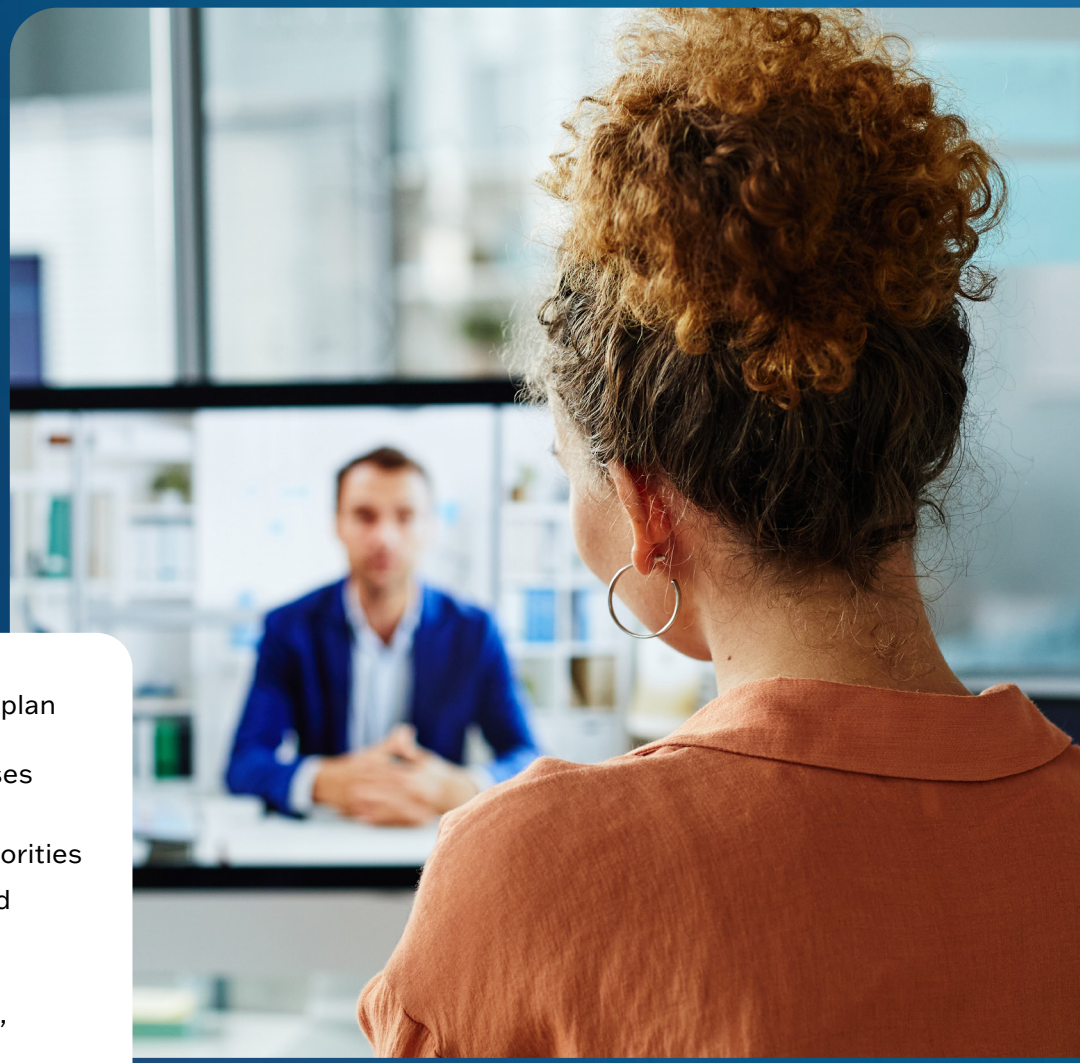

Cadence of Engagement

- Ongoing engagement with defined workplan milestones across pre-deployment, deployment, and post-deployment phases
- Quarterly status reporting summarizing progress, open items, and upcoming priorities
- Regular check-ins with management and AI/IT leadership, with escalation to committees or the board as needed
- Event-driven support when new AI tools, vendors, or risks emerge

Division of Responsibilities

vCAIO
responsibilities include:

- AI governance structure and operating model
- AI strategy and roadmap development
- AI policy, standards, and procedure guidance
- AI risk assessments and ongoing oversight
- Training strategy and phased rollout planning
- Management and board reporting

Client
responsibilities include:

- Day-to-day system administration
- Business ownership of AI use cases
- Final approval authority
- Operational control execution
- Internal change management and adoption

Where the vCAIO Provides Targeted Support

In addition to ongoing responsibilities, a vCAIO provides focused support in key areas such as:

- AI policy development and periodic updates
- Use case definition and risk evaluation for new initiatives
- Roadmap alignment with business strategy
- Monitoring and oversight design
- Board and committee reporting support

How Governance and Oversight Are Supported

- Ongoing governance as a program, not a one-time project
- Formal workplan with defined scope, status, and approvals
- Clear distinction between SBS-led and client-led responsibilities
- Regular reporting to support informed decision-making by management and the board
- Alignment with regulatory expectations and internal risk appetite over time

Talk to an SBS expert about how a vCAIO can transform your AI strategy.
link.sbscyber.com/vcaio

Next Steps
for Your AI Program

As AI programs mature, these considerations help reinforce strong governance and risk-aware decision-making.

1 AI Vendor Due Diligence: Critical Questions to Ask

- What type of AI models does your product use?
- How are training data and outputs handled?
- How do you prevent sensitive data retention?
- What controls reduce hallucinations or incorrect outputs?
- What is your incident response plan for model failures?

Limited transparency into data handling or model behavior should be treated as a potential risk indicator.

2 Structure of an AI Governance Policy: A Framework, Not a Template

- Purpose and scope
- Definitions
- Acceptable/prohibited use
- Data handling
- Security expectations
- Vendor responsibilities
- Review and escalation process

Specific requirements may vary based on regulatory expectations and institutional risk.

3 What to Track in an AI Inventory: Fields to Include

- Tool/system name
- Owner
- AI function
- Data inputs/outputs
- Vendor involvement
- Risk rating
- Controls
- Review schedule

Maintaining this inventory supports ongoing oversight, audit readiness, and informed decision-making as AI use evolves.

A Clear
Path Forward

AI adoption in regulated industries works best when it is deliberate, governed, and aligned with risk and oversight expectations. This roadmap is designed to help leaders understand where AI already exists, identify governance gaps, and take practical steps toward responsible adoption.

By focusing on visibility, accountability, and alignment early, organizations can reduce uncertainty, strengthen oversight, and move forward with clarity and control. Organizations that move step by step, starting with visibility and governance, are better positioned to scale AI safely.

Looking for guidance on turning this roadmap into action? Connect with an SBS AI governance specialist.

link.sbscyber.com/contact

sbscyber.com

About SBS CyberSecurity

SBS CyberSecurity empowers organizations to make informed cybersecurity decisions with clarity. By putting people at the center, SBS helps teams navigate complex challenges and turn them into real-world impact. SBS is committed to helping organizations stay secure, resilient, and ready for what's ahead.