Allied Solutions®

The AI Modernization Paradox

Balancing Aggressive Innovation with Practical Risk Governance



Table of Contents

3 Executive Overview

5 The Risk of Saying No

7 A New Model for AI Governance

- 8 Pillar 1 Classify AI Use by Risk
- 11 Pillar 2 Set Contractual Boundaries
- 13 Pillar 3 Give Teams a Secure Path to Experiment

14 Using AI to Strengthen Cyber Defense

15 The Monday Morning Playbook

18 Final Thought

19 About Allied

Executive Overview

AI is moving from experimentation to infrastructure. For financial institutions, the question is no longer whether AI will be used, but how quickly organizations can establish the visibility, governance and controls needed to use it responsibly.

The adoption curve is already underway in U.S. banking. Bank Director's 2026 Technology Survey found that 72% of responding bank executives and directors said their institutions had implemented generative AI. The survey also found that 88% had drafted an acceptable-use policy for AI technologies. These findings suggest that AI is becoming part of banking operations, reinforcing the need to establish governance alongside innovation.

That gap creates the AI modernization paradox: organizations are being asked to accelerate innovation while becoming more cautious about security, privacy, compliance, operational resilience and reputation.

Effective governance makes AI risk visible, measurable and manageable so leaders can make informed decisions.

The AI Governance Gap

72% *of responding bank execs say they've implemented AI*

87% *of banks could significantaly improve AI governance maturity*

49% *of organizations expect a Shadow AI incident*

30% *of breaches involve third parties*

AI adoption is accelerating. Governance, visibility and defense have to accelerate with it.



A risk-based approach starts with four questions: Where is AI being used? What data does it touch? What decisions does it influence? Which controls fit the use case? Governance must also cover vendor agreements, data protections, identity controls and ongoing monitoring beyond the institution's network.

This paper presents a practical model built around three principles:

1. **Classify AI use by risk.** Match controls to the data involved, the decisions affected and the potential consequences.
2. **Set contractual boundaries for vendor AI.** Require data isolation, prohibit external model training with institutional data and assign accountability.
3. **Give teams an approved path to experiment.** Provide vetted tools, controlled environments and clear review thresholds.

The same principle applies to defense. AI is accelerating the threat environment, from social engineering to vulnerability discovery. Financial institutions will increasingly need defensive capabilities that can operate at machine speed, using AI to process telemetry, correlate signals and automate repetitive security tasks while keeping human expertise at the center of critical decisions.

The Risk of Saying "No"

For financial institutions, the safest response to an emerging technology can sometimes appear to be the simplest: restrict it until the risks are better understood. With AI, that approach carries its own risks.

A blanket prohibition may reduce approved use, but it does not eliminate the underlying demand for faster analysis, automation, content creation, software development or decision support. The pressure to improve efficiency does not disappear because an institution has restricted the tools that can enable it.

In fact, saying "no" without providing a viable alternative can create a different set of vulnerabilities.

Competitive Disadvantage

AI is rapidly becoming embedded in how organizations operate, from customer service and marketing to software development, fraud detection and internal productivity.

Institutions that delay adoption indefinitely may find themselves competing against organizations that can deliver similar work faster, automate more processes and respond more quickly to changing customer and member expectations.

Delaying adoption can widen the operational gap between an institution and competitors that use AI effectively.

An institution does not need to adopt every new AI capability. It does need a deliberate way to determine which capabilities create meaningful value, which risks are acceptable and where waiting carries a greater cost than moving forward.

Lost Operational Efficiency

AI can reduce the time required for tasks that consume significant employee capacity without necessarily creating corresponding value, including summarization, research, drafting, coding assistance, data analysis and administrative workflows.

When access to appropriate AI tools is overly restricted, employees continue performing these tasks manually or rely on fragmented workarounds.

That creates an uncomfortable tradeoff: the organization may reduce one category of technology risk while preserving the operational inefficiency that prompted employees to seek AI in the first place.

Leaders should compare the risk of each use case with the operational cost of avoiding it.

Falling Behind Customer and Member Expectations

The pressure is not coming only from inside the organization. Customers are becoming comfortable using AI to understand financial products, compare options and seek advice. McKinsey found that nearly **one in four** consumers surveyed were already using GenAI for financial tasks at least monthly.

That creates a second modernization risk. An institution can decide to move slowly internally while its customers, competitors and technology providers continue moving forward.

Customer and member expectations now belong in the modernization decision, even when an institution chooses not to deploy a customer-facing AI application.

The Shadow AI Problem

The more consequential risk may be what happens outside the approved workflow. Proofpoint found that **49%** of organizations expected a Shadow AI incident within the following 12 months, making unapproved AI use one of the leading AI-related security concerns identified in its research.

A prohibition may reduce sanctioned use. It does not necessarily eliminate demand. When the approved path is too restrictive, difficult or slow, employees may find another path. This creates what is commonly referred to as **Shadow AI:** AI use that occurs without the organization's knowledge, oversight or formal approval.

Unauthorized tools remove visibility. The institution may not know which tools employees use, what information they enter, where providers process it, whether providers retain it or how providers use it.

The Governance Vacuum

This is where the risk of saying "no" becomes most consequential. When an institution prohibits AI without creating practical pathways for appropriate use, it can leave employees and business teams with a simple choice: wait, work manually or find another way.

The organization may have a policy, but lack the visibility, education, approved tools, risk classifications and technical controls needed to make that policy effective in practice.

Institutions should keep some uses off-limits when the potential exposure outweighs the benefit. They also need clear boundaries for appropriate use.

Organizations need a practical way to distinguish between low-risk experimentation, controlled enterprise use and applications that require a hard stop or formal review. They need controls that correspond to the level of risk, rather than applying the same process to every AI use case.

A workable policy tells employees where AI is appropriate, where it is prohibited and which safeguards approved use requires.

A New Model for AI Governance

Effective AI governance creates enough visibility and control to decide where AI can move quickly, where additional safeguards are necessary and where use should stop.

A practical operating model has three steps:

See, Classify and Enable

SEE

What is happening?

Identify where AI is being used, what systems and vendors are involved, what data is flowing into those systems and where employees may already be using AI outside formally approved channels.

CLASSIFY

What level of risk does this represent?

Evaluate AI use based on the data involved, the decisions being influenced, the degree of human oversight and the potential consequences if something goes wrong.

ENABLE

What controls allow us to move forward safely?

Apply controls that correspond to the level of risk, creating clear pathways for approved use while restricting or stopping activities where the potential exposure is unacceptable.

Together, these steps create informed control.

AI now appears across software, cloud platforms, productivity tools, customer experiences and third-party services. Employees can also access powerful public tools with little more than a browser and an account, so institutions can no longer evaluate AI as a small set of isolated applications.

A boundary drawn only around infrastructure misses much of this activity. The model begins by identifying the data, identities, decisions, vendors and workflows that surround AI. Three principles put it into practice.

Pillar 1: Classify AI Use by Risk

AI use cases pose different levels of risk. A tool that drafts an internal meeting summary differs materially from a system that analyzes sensitive member information, influences a credit decision or interacts directly with customers.

Yet organizations can easily fall into the trap of creating a single policy for "AI" as though every use case presents the same level of exposure. That approach creates problems in both directions.

AI use also varies in scope and potential impact. Research published by the American Bankers Association in March 2026 found that banks were prioritizing lower-risk applications that support human judgment, including document review, research, policy analysis and internal communications. This approach offers a practical starting point for governance: evaluate each use case and match oversight to its data, purpose and potential consequences.

Treating those use cases the same way creates unnecessary friction for low-risk applications while potentially understating the controls required for high-impact ones.

The solution is risk-based classification.

Rather than asking, *"Is this AI?"* organizations should ask:

- **What data does it access?**
- **Where does that data go?**
- **What decision or process does it influence?**
- **How much human oversight exists?**
- **What happens if the system produces an incorrect, biased or compromised output?**
- **How difficult would the impact be to detect and reverse?**



A practical classification model might include three categories:

1 Low-Risk Productivity



AI uses involving public information, generic content or routine productivity tasks with limited consequences if the output is incorrect.

Examples: Drafting public-facing marketing copy, creating generic training templates or summarizing non-sensitive information.

Governance response: *Approved tools, basic usage guidelines and employee education.*

2 Controlled Business Use



AI applications that provide meaningful operational value but operate within enterprise-controlled environments and established safeguards.

Examples: Internal code generation, meeting summaries, document analysis or workflow automation using approved, enterprise-bound AI environments.

Governance response: *Approved platforms, defined data boundaries, access controls, monitoring and appropriate human review.*

3 Sensitive or High-Impact Use



AI applications involving sensitive data, consequential decisions or processes where an incorrect or compromised output could materially affect an individual or the institution.

Examples: Systems involving member personally identifiable information, credit decisions, underwriting, fraud determinations or other high-impact financial processes.

Governance response: *Formal risk assessment, validation, documented controls, enhanced monitoring and defined human accountability.*

When High-Risk Use Should Stop

Some high-risk applications should not proceed. An institution should restrict or prohibit a use when it cannot protect the data, validate the output, maintain appropriate human oversight or establish clear accountability.

Examples include entering live customer or member information into a public AI tool, allowing an unvalidated model to make autonomous credit or underwriting decisions, or using a vendor that will not contractually guarantee tenant-level data isolation and prohibit the use of institutional data to train its foundational public models.

This three-level model allows institutions to move low-risk uses forward, place appropriate controls around business applications and stop high-risk uses when the necessary protections are unavailable.



Pillar 2: Set Contractual Boundaries

Financial institutions now operate across cloud platforms, software services, APIs and third-party providers. Security programs built mainly around a physical network boundary cannot govern every place data moves.

Organizational data moves through Microsoft 365, AWS, SaaS platforms, insurance distributors, claims platforms and other connected services. Each connection adds identities, vendors and contractual terms that the institution must govern.

A current risk review should answer:

- **Who can access it?**
- **What systems and vendors touch it?**
- **Where does it move?**
- **What happens to it along the way?**
- **What visibility do we have into that activity?**

These questions shift attention from the physical location of data to the people, systems and agreements that control it.

Institutions must govern the contracts, identities and technology relationships that touch their data, including cloud platforms, SaaS applications, APIs, embedded AI capabilities, public AI tools and third-party providers.

Verizon's 2025 Data Breach Investigations Report found that third-party involvement had doubled to 30% of breaches. As more technology and AI capabilities are delivered through third parties, those relationships become part of the organization's risk environment.

That means understanding how AI capabilities operate throughout the ecosystem, including:

- What data a third-party AI service can access
- Whether data is isolated from other customers
- Whether providers retain data or use it for model training
- Who has access to the system and underlying data
- How activity is monitored
- What happens when a vendor relationship ends
- Who remains accountable for the AI capability throughout its lifecycle

Vendor agreements should include a data isolation clause with two explicit guarantees:

- The provider will isolate institutional data and customer or member data within the institution's specific enterprise tenant.
- The provider will not use that data to train the vendor's foundational public models.

If a vendor will not make both commitments, the institution should not approve the service. The agreement should also define access, retention, deletion, audit rights, incident notification, subcontractor obligations and exit requirements.

These terms add contractual control when technical controls cannot keep data inside a physical boundary.

The goal is clear visibility into where AI operates, what it touches and who remains accountable for the outcome.



Pillar 3: Give Teams a Secure Path to Experiment

After an institution identifies AI use and classifies the risk, it can define the controls that allow work to proceed. Clear controls reduce review time and give employees and developers a known path forward.

Guardrails define what employees can do, where additional review is required and which uses remain off-limits. They prevent low-risk work from entering the same review process as high-impact applications.

Practical guardrails can include:

- Approved-use policies and AI tools
- Data classification and handling requirements
- Access controls
- Human oversight and review
- Testing and validation
- Monitoring and logging
- Documentation
- Escalation thresholds
- Clear ownership and accountability
- Secure environments for experimentation

The control structure should show employees where they can move quickly, where they need additional review and where they must stop.

When governance becomes too difficult to navigate, people have an incentive to route around it. Employees still need to summarize, research, draft, analyze and automate. Developers still need environments where they can build and test. If the approved path is too slow or restrictive, the demand does not disappear. It can move outside the organization's visibility.

Institutions can meet legitimate demand through secure pathways for experimentation.

Employees need approved tools. Developers need pre-approved, isolated environments with vetted infrastructure and APIs. Establishing controls up front lets teams test ideas without restarting the risk review for every experiment.

In this model, security provides infrastructure that teams can build on.

Organizations move faster when established controls let teams evaluate new use cases consistently.

Deloitte's research points to a similar distinction between AI pioneers and followers, with more advanced organizations demonstrating greater readiness across governance, talent and technology and greater ability to scale AI experimentation.

Using AI to Strengthen Cyber Defense

AI is increasing the speed and scale of cyberattacks. Financial institutions need defensive tools that can process large volumes of data and help security teams respond faster. For example, the IBM X-Force test

IBM's X-Force tested how quickly AI could generate a convincing phishing campaign and found that it could produce the campaign in approximately five minutes, compared with roughly 16 hours for a human social-engineering team. The comparison excluded infrastructure setup, but the implication is clear:

Threat actors use AI to accelerate reconnaissance, generate targeted social engineering content and scan for vulnerabilities. Security teams can use AI to process information, identify patterns and respond more quickly.

AI can help organizations:

- Ingest and correlate large volumes of security telemetry
- Prioritize potential threats and surface patterns that warrant investigation
- Automate repetitive analysis
- Give security professionals more capacity to focus on complex decisions

This matters particularly when security teams are operating with constrained budgets and limited access to specialized talent. AI does not eliminate the need for human expertise. It can extend the reach of the people who have it.

The purpose is to extend the security team's reach and speed while keeping people responsible for consequential decisions. But defensive AI still requires governance.

Organizations need to understand what data security tools can access, how AI capabilities process that data, where outputs are used and what decisions remain with human operators. The same principles established earlier in this paper still apply: **see what is happening, classify the risk and enable the appropriate response.**

Governance determines where defensive AI can operate and which decisions require human review.

Clear boundaries give employees approved tools, developers controlled environments, leaders defined review triggers and security teams more time for the threats with the greatest consequences.

Well-defined controls increase the institution's capacity to act. When teams know what is permitted, what requires additional controls and what must stop, they can scale useful capabilities without evaluating every opportunity from scratch.

The same governance model can support innovation and cyber defense.

The Monday Morning Playbook

AI governance can begin with a clear view of current activity and a repeatable way to decide what happens next.

The following actions create a practical starting point.

1 Inventory Current AI Use

Before creating new restrictions, understand where AI is already operating.

Look beyond formally approved tools. Include embedded AI capabilities, employee use, experimentation, workarounds and potential Shadow AI. The objective is visibility, not punishment. The organization needs to understand where demand already exists before it can determine how that demand should be governed.



Executive question:

Where is AI already operating across the organization?

2 Classify Use Cases by Risk

Not every AI use case requires the same level of scrutiny.

Use the four risk categories established earlier: low-risk productivity, controlled business use, sensitive or high-impact use, and prohibited or restricted use. Evaluate each use case based on the data involved, the decisions it influences, the level of human oversight and the potential consequences if something goes wrong.



Executive question:

Which uses can move quickly, which require additional controls and which should not proceed?

3 Set Vendor Contract Requirements

Identify what each vendor's AI can access, generate, influence and retain.

Map the data each third-party system processes. Require a data isolation clause that keeps institutional data and customer or member data within the institution's specific enterprise tenant and prohibits its use to train the vendor's foundational public models. Define access, retention, deletion, audit, incident notification and termination requirements.



Executive question:

Does every AI vendor isolate our data within our enterprise tenant and prohibit its use to train foundational public models?

4 Define Minimum Guardrails

Define the minimum controls required for each risk category.

These may include approved platforms, data restrictions, access controls, human review, testing and validation, monitoring, documentation and escalation requirements. Where appropriate, establish pre-approved environments for experimentation and development so teams have a secure place to build, test and learn.



Executive question:

What minimum controls allow us to move safely?

5 Assign Ownership and Escalation

Every AI use case needs clear accountability.

Determine who owns the risk, who approves the use, who monitors performance and who has authority to stop or escalate an issue. Depending on the use case, responsibility may span business leaders, data teams, risk, legal, security and technology.

This aligns with the core functions of the NIST AI Risk Management Framework: govern, map, measure and manage.



Executive question:

Who is accountable for the outcome if something goes wrong?

Five Questions to Keep in the Room

As new AI use cases emerge, organizations do not need to predict every future technology. They need a repeatable decision discipline.

Ask:

1. **What problem are we trying to solve?**
2. **What data does AI access or generate?**
3. **What decisions or outcomes could it influence?**
4. **What controls make the risk acceptable?**
5. **Who owns the outcome if something goes wrong?**

These questions create a practical governance habit. They can be applied to the next AI tool, the next vendor capability or the next use case without rebuilding the governance model from scratch.

Final Thought

AI modernization succeeds when institutions advance innovation and risk management together. Deloitte found an association between stronger AI governance maturity and stronger revenue growth. In its analysis, a 10-point increase in governance maturity corresponded with an average 10% increase in revenue growth. The finding shows association, not causation. Read Deloitte's Trustworthy AI Governance Index

Responsible AI adoption requires visibility, controls matched to risk, approved pathways and clear accountability. A repeatable governance model makes responsible AI adoption easier to execute and safer to scale.



> ***Strong governance helps institutions scale AI use without losing visibility or accountability.***

About Allied Solutions

Allied Solutions is one of the largest providers of insurance, lending, risk management, and data-driven solutions to financial institutions in North America. Allied Solutions uses technology-based solutions customized to meet the needs of over 6,000 banks, credit unions, finance companies, mortgage servicers, and auto dealers, along with a portfolio of innovative products and services from a wide variety of providers. Allied Solutions is headquartered in Carmel, Indiana and maintains several offices strategically located across the country. Allied Solutions is a wholly owned and independently operated subsidiary of Securian Financial Group.

Learn More

- Follow Allied on LinkedIn
- Read Allied Insights Thought Leadership
- Subscribe to Allied Insights Newsletters
- Listen to The Allied Angle on Podbean

Allied Solutions®

GROW, PROTECT AND EVOLVE YOUR BUSINESS.®

alliedsolutions.net

© 2026 Allied Solutions, LLC.

The information presented in this document is intended for informational purposes only and should not be construed as legal advice or a legal opinion and it may not reflect the most current legal developments. You should seek the advice of legal counsel of your choice before acting on any information provided in this document.

2369-R2-11.26