EXECUTIVE INSIGHTS · AI GOVERNANCE SERIES

Building an AI Governance Program

A Practical Guide for Financial Institutions

Introducing the Stokes AI Governance Framework™

By Alison Stokes, CRCM

Contents

Executive Summary

Governance Must Evolve as Artificial Intelligence Evolves

Artificial Intelligence is no longer an emerging technology reserved for innovation labs and pilot programs. It is becoming embedded in the way financial institutions serve customers, make decisions, detect fraud, assess risk, manage operations, and improve efficiency. As adoption accelerates, so do expectations from regulators, boards of directors, executive leadership, and consumers.

Much of the conversation surrounding Artificial Intelligence focuses on what the technology can do. Organizations are investing significant time and resources evaluating generative AI, predictive analytics, intelligent automation, and machine learning to create competitive advantages and improve customer experiences. While those conversations are important, they often overlook a more fundamental question:

How will Artificial Intelligence be governed?

For decades, financial institutions have successfully navigated transformational change by relying on disciplined governance. Whether implementing online banking, mobile applications, digital lending, cloud technologies, or enterprise data programs, organizations that succeeded shared one common characteristic: they integrated innovation into established governance frameworks rather than treating it as a separate initiative.

Artificial Intelligence should be approached no differently.

Effective AI governance is not about slowing innovation. It is about creating the governance structures that allow innovation to occur responsibly. Organizations that establish clear accountability, integrate AI into existing risk management processes, monitor performance continuously, and provide meaningful executive oversight are better positioned to realize the benefits of AI while managing regulatory, operational, reputational, and consumer risks.

This white paper introduces the **Stokes AI Governance Framework™**, a practical implementation model designed to help financial institutions incorporate Artificial Intelligence into existing Enterprise Risk Management (ERM) and Compliance Management Systems (CMS). Rather than replacing established governance practices or regulatory guidance, the framework demonstrates how organizations can leverage the governance structures they already trust to oversee AI responsibly.

The framework is intentionally practical. It reflects years of executive leadership in enterprise risk, regulatory compliance, governance, examination readiness, compliance testing, issue management, and operational transformation across financial services. It also aligns with widely recognized guidance, including the National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework, ISO/IEC 42001, Federal Reserve SR 11-7, Consumer Financial Protection Bureau guidance, Office of the Comptroller of the Currency supervisory expectations, and the OECD AI Principles.

Throughout this publication, we will explore how financial institutions can:

- Establish governance structures that clearly define accountability for AI initiatives.
- Integrate AI into existing Enterprise Risk Management and Compliance Management Systems.
- Develop policies, inventories, and risk assessments that support responsible AI adoption.
- Strengthen testing, monitoring, complaint management, and issue management processes.
- Enhance board and executive reporting to improve oversight and strategic decision-making.
- Build a practical roadmap for implementing and maturing an AI governance program.

One theme remains consistent throughout this paper:

> ***Artificial Intelligence does not replace good governance. It requires even better governance.***

Organizations that treat AI governance as a strategic business capability, not merely a technology initiative, will be better equipped to innovate responsibly, adapt to evolving regulatory expectations, and maintain the trust of customers, regulators, shareholders, and employees.

The future of Artificial Intelligence will undoubtedly continue to evolve. Governance must evolve with it.

This white paper is intended to help leaders take that next step.

Chapter 1: The AI Governance Imperative

Purpose

Artificial Intelligence is transforming financial services at a pace few technologies have matched. From fraud detection and underwriting to customer service, marketing, operations, and regulatory compliance, AI is reshaping how financial institutions make decisions and deliver products and services.

While the technology continues to evolve rapidly, one principle remains constant:

Innovation without governance creates risk.

The purpose of this chapter is to establish why AI governance has become an executive priority and why effective governance is essential to enabling responsible innovation.

Artificial Intelligence Is Changing Financial Services

Financial institutions have always embraced technology to improve efficiency, reduce costs, and enhance the customer experience. Artificial Intelligence represents the next stage of that evolution.

Today's AI capabilities can:

- Analyze vast amounts of information in seconds.
- Detect fraud patterns that humans may never identify.
- Improve customer interactions through intelligent automation.
- Strengthen regulatory monitoring and surveillance.
- Enhance operational efficiency across the enterprise.
- Support faster and more informed business decisions.

These capabilities offer tremendous opportunities. They also introduce new governance challenges.

Unlike traditional technology, AI systems can learn, adapt, and influence decisions that affect consumers, employees, and business operations. As organizations increase their reliance on AI, governance must evolve to ensure these technologies are deployed responsibly, ethically, and in alignment with regulatory expectations.

> ***Executive Insight:*** *Artificial Intelligence changes how organizations make decisions. It should never change the principles by which those decisions are governed.*

Governance Is the Foundation of Responsible Innovation

When organizations begin exploring AI, discussions often focus on technology, data, models, and implementation. Those conversations are important.

However, executive leadership should begin with a different question: **How will we govern Artificial Intelligence?**

Governance establishes accountability, defines decision-making authority, sets expectations for risk management, and ensures that innovation aligns with the organization's values, strategic objectives, and regulatory obligations.

Without governance, organizations may adopt AI quickly but struggle to demonstrate that decisions are transparent, explainable, consistent, and appropriately monitored.

Strong governance creates confidence. It enables organizations to innovate while maintaining the trust of customers, regulators, shareholders, employees, and the communities they serve.

> ***From Experience:*** *Throughout my career, I've found that successful organizations don't wait until a new technology is fully implemented before discussing governance. They establish governance early, creating clear accountability, defined responsibilities, and disciplined oversight that support innovation rather than slow it down.*

AI Governance Is a Leadership Responsibility

AI governance should not be viewed as a technology initiative. It is an enterprise responsibility.

Executive leadership establishes the organization's vision. The Board provides oversight. Business leaders remain accountable for outcomes. Risk, Compliance, Legal, Information Security, Privacy, Internal Audit, and Technology each contribute essential expertise.

Responsible AI governance requires these functions to work together within an integrated governance structure. Organizations that recognize AI as an enterprise issue are better positioned to adapt as technology, regulations, and stakeholder expectations continue to evolve.

Why This Guide Matters

Most organizations understand that AI governance is important. Many are less certain about where to begin.

Existing standards and regulatory guidance establish valuable principles and expectations. However, executive teams often need practical guidance for integrating those principles into existing governance, risk management, and compliance programs.

That is the purpose of this guide. The chapters that follow examine how governance must evolve to support Artificial Intelligence and present a practical approach for building an AI governance program that complements, not replaces, existing governance structures.

Executive Action Checklist

Before moving to the next chapter, consider whether your organization has established the following foundations:

- ☐ Executive leadership recognizes AI governance as an enterprise priority.
- ☐ AI governance objectives support the organization's strategic goals.
- ☐ Leadership understands that governance enables innovation rather than restricting it.
- ☐ Discussions regarding AI include governance, accountability, and risk management.
- ☐ The Board recognizes that AI governance requires ongoing oversight.

Executive Takeaway

Artificial Intelligence may transform the way financial institutions operate, but it does not change the principles of sound governance.

Organizations that establish governance as a strategic priority from the outset will be better positioned to innovate responsibly, adapt to regulatory change, and maintain the trust that is essential to long-term success.

In the next chapter, we examine why traditional governance approaches must evolve to address the unique characteristics and risks introduced by Artificial Intelligence.

Chapter 2: The Evolution of Governance

Purpose

Financial institutions have long relied on well-established governance frameworks to manage risk, maintain regulatory compliance, and support sound business practices. Enterprise Risk Management (ERM), Compliance Management Systems (CMS), Model Risk Management (MRM), Information Security, Privacy, Internal Audit, and Operational Risk have provided the foundation for responsible decision-making across the enterprise.

Artificial Intelligence does not eliminate the need for these governance disciplines. It expands them.

The purpose of this chapter is to examine why governance must evolve to address the unique characteristics of Artificial Intelligence while reinforcing the governance principles organizations already trust.

Strong Governance Is Already in Place

Many organizations begin their AI governance journey believing they must create an entirely new governance program. In most cases, that assumption is unnecessary.

Financial institutions already possess mature governance capabilities that have evolved over decades of regulatory oversight and industry experience. These programs establish accountability, define risk management practices, monitor regulatory compliance, oversee third-party relationships, and provide executive leadership and Boards with meaningful oversight.

Artificial Intelligence should build upon these existing governance capabilities, not replace them. The opportunity is not to start over. The opportunity is to evolve.

> ***Executive Insight:*** *The organizations best positioned to govern Artificial Intelligence are often those with the strongest governance foundations already in place.*

Why AI Changes the Governance Conversation

Although the principles of governance remain consistent, Artificial Intelligence introduces characteristics that require organizations to expand their governance practices.

Unlike traditional software, AI systems may learn from new information, generate probabilistic outcomes, or operate using complex models that are difficult to explain without specialized expertise.

Organizations must therefore consider questions that may not have existed in traditional governance programs. Examples include:

- How do we validate AI-generated outcomes?
- How do we identify and manage bias?
- How do we ensure appropriate human oversight?
- How do we govern generative AI used by employees?
- How do we monitor third-party AI providers?
- How do we explain AI-assisted decisions to regulators and consumers?
- How do we monitor AI performance after implementation?
- How do we respond when an AI solution behaves differently over time?

These questions expand the governance conversation beyond traditional technology oversight.

Governance Must Become More Dynamic

Traditional governance often focused on systems that changed infrequently. Artificial Intelligence changes continuously.

Models may evolve. Training data may change. Customer behavior may shift. Business strategies may adapt. Regulatory expectations continue to develop.

As a result, governance cannot rely solely on periodic reviews or annual policy updates. Organizations need governance processes that provide continuous visibility into emerging risks, changing performance, and evolving regulatory expectations.

Good AI governance is adaptive. It is designed to evolve alongside the technology it governs.

> ***From Experience:*** *One of the most effective governance practices I have observed is designing governance to evolve with the business. Organizations that regularly reassess their governance processes are far better prepared to respond to regulatory change, operational complexity, and emerging technologies than those that treat governance as static documentation.*

The Expanding Role of Executive Leadership

Artificial Intelligence introduces decisions that extend well beyond technology implementation. Executive leadership must determine:

- How much AI-related risk is acceptable?
- Which AI use cases align with the organization's strategy?
- What governance standards should apply to third-party AI solutions?
- How should AI-related issues be escalated?
- What information should be reported to executive leadership and the Board?
- How will responsible innovation be measured?

These are strategic governance decisions. They cannot be delegated solely to technology teams.

Preparing for What Comes Next

Organizations do not need to abandon existing governance frameworks to manage Artificial Intelligence. They need to strengthen them. They need governance that is:

- Integrated across business functions.
- Risk-based and scalable.
- Flexible enough to accommodate emerging technologies.
- Supported by executive leadership.
- Consistently monitored and continuously improved.

These characteristics provide the foundation for a mature AI governance program.

Executive Action Checklist

- ☐ Existing governance programs have been evaluated for AI readiness.
- ☐ Leadership understands how AI differs from traditional technology.
- ☐ AI-specific risks have been identified and discussed.
- ☐ Governance processes support continuous monitoring rather than periodic oversight alone.
- ☐ Executive leadership recognizes that AI governance is a cross-functional responsibility.
- ☐ Existing governance capabilities are viewed as the foundation for AI governance.

Executive Takeaway

Artificial Intelligence does not require organizations to replace their governance programs. It requires them to strengthen, expand, and integrate the governance capabilities they already possess.

Organizations that recognize this distinction will be better prepared to manage emerging risks while enabling responsible innovation.

In the next chapter, we introduce the **Stokes AI Governance Framework™**, a practical governance model designed to help financial institutions operationalize these principles through an integrated governance lifecycle.

Chapter 3: The Stokes AI Governance Framework™

Purpose

Financial institutions do not need another theoretical model for AI governance. They need a practical framework that helps them integrate Artificial Intelligence into the governance structures they already trust.

The Stokes AI Governance Framework™ was developed to provide that practical approach.

Rather than replacing Enterprise Risk Management, Compliance Management Systems, Model Risk Management, Information Security, Privacy, Internal Audit, or Operational Risk programs, the framework brings these disciplines together into a unified governance lifecycle for Artificial Intelligence.

Its purpose is simple: to help organizations innovate responsibly while maintaining strong governance, effective risk management, regulatory compliance, and executive accountability.

Why This Framework Was Developed

As organizations began adopting Artificial Intelligence, many governance discussions focused on individual risks such as model performance, data quality, privacy, bias, cybersecurity, or regulatory compliance.

While each of these areas is important, addressing them independently often creates fragmented governance: separate committees, separate reporting, separate assessments, separate accountability.

Effective governance should not operate in silos. It should function as a connected system.

The Stokes AI Governance Framework™ was designed to connect these governance activities into a single operating model that supports the entire AI lifecycle.

> ***Executive Insight:*** *Artificial Intelligence should not create separate governance programs. It should strengthen the governance ecosystem that already exists.*

The Framework

Figure 1. The Stokes AI Governance Framework™. A continuous governance lifecycle spanning Governance, Policies & Procedures, Risk Assessment, Regulatory Change Management, Testing & Monitoring, Complaint Management, Issue Management, and Board & Executive Reporting, centered on Responsible AI Governance.

The framework is built around one central objective: **Responsible AI Governance.**

Surrounding that objective are eight integrated governance domains that operate as a continuous lifecycle.

1. Governance
2. Policies & Procedures
3. Risk Assessment
4. Regulatory Change Management
5. Testing & Monitoring
6. Complaint Management
7. Issue Management
8. Board & Executive Reporting

Each domain performs a distinct function. Together, they create an integrated governance operating model that supports responsible innovation while strengthening enterprise oversight.

How the Framework Operates

Unlike traditional implementation models, the Stokes AI Governance Framework™ is not linear. Governance is continuous. Every governance activity informs the next.

Governance establishes accountability. Policies define expectations. Risk Assessments identify exposure. Regulatory Change Management monitors evolving requirements. Testing & Monitoring evaluates performance. Complaint Management identifies customer impacts. Issue Management drives corrective action. Board & Executive Reporting provides leadership with the information necessary to oversee the program and guide future decisions.

As governance decisions evolve, the lifecycle begins again. This continuous improvement model enables organizations to adapt as Artificial Intelligence, business strategies, and regulatory expectations continue to evolve.

Alignment with Existing Standards

The Stokes AI Governance Framework™ is intended to complement, not replace, existing governance standards and regulatory expectations.

Organizations should continue leveraging established guidance, including:

- NIST AI Risk Management Framework (AI RMF)
- ISO/IEC 42001 Artificial Intelligence Management Systems
- Federal Reserve SR 11-7
- OCC supervisory expectations
- CFPB guidance related to AI and automated decision-making
- OECD AI Principles

These resources describe important governance principles and regulatory expectations. The Stokes AI Governance Framework™ provides a practical operating model for implementing those principles within a financial institution.

> ***From Experience:*** *Throughout my career, I've learned that sustainable governance programs succeed because governance activities are connected. Policies inform risk assessments. Testing validates controls. Complaints reveal emerging issues. Executive reporting drives better decisions. When these activities function as an integrated lifecycle rather than independent processes, organizations become more resilient, more responsive, and better prepared for change.*

The Eight Governance Domains

Each governance domain represents a critical capability within a mature AI governance program.

Although every organization will implement the framework differently, each domain should operate as part of an integrated governance lifecycle.

The chapters that follow examine each domain in detail, providing practical guidance, leading practices, executive considerations, and implementation recommendations.

Together, these domains help organizations establish governance that is scalable, sustainable, and adaptable to future technological and regulatory change.

Executive Action Checklist

- ☐ Executive leadership supports the development of an AI governance program.
- ☐ Existing governance capabilities have been identified and assessed.
- ☐ Cross-functional stakeholders have been identified.
- ☐ Governance responsibilities will be integrated rather than siloed.
- ☐ Leadership understands the relationship between the eight governance domains.
- ☐ The organization is committed to continuous governance rather than one-time implementation.

Executive Takeaway

Artificial Intelligence introduces new opportunities, new risks, and new regulatory expectations.

Effective governance is not achieved by creating isolated controls. It is achieved by integrating governance activities into a coordinated operating model that evolves alongside the organization.

The Stokes AI Governance Framework™ provides that model.

The chapters that follow examine each governance domain individually, demonstrating how they work together to create a comprehensive and sustainable AI governance program.

Chapter 4: Governance

Purpose

Governance is the foundation of every successful AI governance program.

It establishes who is accountable, how decisions are made, how risks are escalated, and how executive leadership and the Board exercise oversight. Without governance, organizations cannot consistently manage AI risk, satisfy regulatory expectations, or ensure that AI initiatives align with business objectives.

Effective governance doesn't slow innovation. It enables responsible innovation.

Why Governance Matters

Every organization using Artificial Intelligence must answer four fundamental questions:

- Who owns it?
- Who approves it?
- Who oversees it?
- Who is accountable when something goes wrong?

Governance provides those answers. It creates the structure that allows organizations to innovate confidently while maintaining effective oversight, regulatory compliance, and sound risk management.

Strong governance also helps ensure that AI initiatives are aligned with the organization's strategy, risk appetite, and values.

> ***Executive Insight:*** *AI doesn't change who is accountable. It changes what leaders must oversee.*

Governance Principles

Every AI governance program should be built on five foundational principles.

Accountability. Every AI solution should have a clearly identified business owner. While technology teams may develop or support AI capabilities, accountability for business outcomes remains with the business.

Transparency. Leadership should know where AI is being used, why it is being used, and the level of risk associated with each implementation.

Collaboration. AI governance is a shared responsibility. Business leaders, Compliance, Enterprise Risk Management, Legal, Information Security, Privacy, Technology, Internal Audit, Human Resources, Procurement, and Third-Party Risk Management all have important roles to play.

Risk-Based Oversight. Not every AI use case presents the same level of risk. Governance activities should be proportional to the complexity, impact, and potential consequences of each implementation.

Continuous Improvement. Governance is not a one-time exercise. As AI capabilities evolve, governance processes should evolve with them.

Leading Practices

Organizations with mature AI governance programs often have several practices in common.

Appoint an Executive Sponsor. Executive sponsorship demonstrates organizational commitment and helps ensure AI governance remains aligned with strategic priorities.

Establish an AI Governance Committee. A cross-functional committee provides oversight, reviews significant AI initiatives, approves governance standards, and addresses emerging risks.

Maintain an Enterprise AI Inventory. Organizations should know where AI is being used across the enterprise. A centralized inventory should identify business owners, intended use, risk classification, third-party involvement, and governance status for every AI solution.

Classify AI by Risk. A chatbot answering routine customer questions should not receive the same level of governance as an AI model that influences lending, fraud detection, or customer eligibility decisions. Applying governance based on risk allows organizations to focus oversight where it matters most.

Document Governance Decisions. Committee decisions, approvals, policy exceptions, and significant governance discussions should be documented to support accountability, consistency, and regulatory examinations.

Report to Executive Leadership and the Board. Leadership reporting should provide meaningful insight into AI initiatives, emerging risks, governance activities, testing results, significant issues, and regulatory developments. Governance is effective only when leadership has the information needed to make informed decisions.

> ***From Experience:*** *One lesson has remained consistent throughout my career: governance works best when accountability is clear. During regulatory examinations, organizations rarely struggle because they lack policies. They struggle because ownership is unclear, responsibilities overlap, or decisions were never documented. Strong governance removes that uncertainty.*

Questions Every Executive Team Should Ask

- Do we know where AI is being used?
- Is every AI solution assigned to a business owner?
- Have we identified our highest-risk AI use cases?
- Are governance decisions documented?
- Does our governance committee have the authority to make decisions?
- Are significant AI risks consistently escalated?
- Does executive leadership receive meaningful reporting?
- Does the Board receive sufficient information to fulfill its oversight responsibilities?
- Are third-party AI providers subject to governance standards consistent with internally developed AI?

If leadership cannot confidently answer these questions, governance may need to be strengthened.

Common Pitfalls

- Treating AI governance as solely a Technology responsibility.
- Implementing AI before governance structures are established.
- Unclear ownership and accountability.
- Governance committees without decision-making authority.
- Limited executive involvement.
- Inconsistent Board reporting.
- Weak oversight of third-party AI providers.
- Governance processes that fail to evolve as AI capabilities mature.

Governance Readiness Checklist

- ☐ Executive sponsor appointed.
- ☐ AI governance committee established.
- ☐ Governance charter approved.
- ☐ Roles and responsibilities documented.
- ☐ Business owner assigned to every AI initiative.
- ☐ Enterprise AI inventory established.
- ☐ AI use cases classified by risk.
- ☐ Governance decisions documented.
- ☐ Risk escalation procedures defined.
- ☐ Executive reporting established.
- ☐ Board oversight responsibilities documented.
- ☐ Third-party AI governance expectations incorporated.
- ☐ Governance effectiveness reviewed periodically.

Executive Takeaway

Governance is more than an organizational chart or committee structure.

It is the operating model that enables responsible AI adoption across the enterprise.

Organizations that establish strong governance early create the accountability, oversight, and decision-making discipline necessary to innovate responsibly while maintaining the confidence of customers, regulators, executive leadership, and the Board.

In the next chapter, we'll examine how well-designed Policies & Procedures translate governance expectations into consistent business practices across the organization.

Chapter 5: Policies & Procedures

Purpose

Policies and procedures transform governance decisions into consistent business practices.

They establish organizational expectations, define responsibilities, support regulatory compliance, and provide employees with clear direction for the responsible use of Artificial Intelligence.

Without documented policies and procedures, governance becomes dependent on individual judgment, creating inconsistency and increasing operational and regulatory risk.

Well-designed documentation provides a common framework for how AI is developed, acquired, implemented, monitored, and retired across the organization.

Why Policies & Procedures Matter

Artificial Intelligence is evolving rapidly. Employees are experimenting with new tools. Business units are identifying new use cases. Third-party vendors are embedding AI into products and services. Regulatory expectations continue to evolve.

Policies and procedures help organizations respond consistently to these changes. They answer questions such as:

- What is considered AI within our organization?
- Who can approve the use of AI?
- What governance reviews are required?
- How should AI-related risks be assessed?
- When is executive approval required?
- What documentation must be maintained?
- How will AI solutions be monitored after implementation?

When expectations are documented, employees are better equipped to make informed decisions and leaders are better positioned to demonstrate effective governance.

> ***Executive Insight:*** *Strong governance begins with clear expectations. Policies explain what the organization expects; procedures explain how those expectations are carried out.*

Policy Principles

An effective AI documentation framework should be practical, scalable, and aligned with existing governance.

Integrate with Existing Governance. Organizations should build upon existing governance documents whenever possible rather than creating an entirely separate AI governance program. AI should become part of the organization's broader governance framework.

Clearly Define Roles and Responsibilities. Policies should identify accountable business owners, approval authorities, governance committees, and supporting functions. Clear ownership promotes consistency and accountability.

Establish Enterprise Standards. Organizations should define minimum expectations for AI governance, risk assessments, third-party AI solutions, human oversight, documentation, testing and monitoring, record retention, and escalation and reporting. Enterprise standards promote consistency across business units.

Keep Documentation Practical. Policies should establish principles. Procedures should describe the steps necessary to implement those principles. Documentation that is overly complex is difficult to maintain and less likely to be followed.

Review Documentation Regularly. Policies and procedures should evolve alongside changes in technology, business operations, and regulatory expectations. Periodic reviews help ensure documentation remains relevant and effective.

Leading Practices

Create an AI Governance Policy. The policy establishes the organization's overall governance philosophy, guiding principles, accountability structure, and governance expectations.

Develop Supporting Procedures. Procedures explain how governance activities will be performed, including approval processes, documentation requirements, monitoring activities, and escalation protocols.

Adopt Common Templates. Standardized templates promote consistency across business units, including an AI Risk Assessment, AI Use Case Intake Form, AI Inventory, Third-Party AI Review, Governance Committee Agenda, Exception Request, and AI Incident Report.

Establish a Document Hierarchy. Organizations often maintain a structured documentation framework: Policy, Standards, Procedures, Guidelines, Job Aids, and Templates. Each document serves a different purpose while supporting a common governance framework.

Align Documentation Across the Enterprise. AI governance documentation should complement existing policies related to enterprise risk management, compliance, privacy, information security, model risk management, records management, vendor management, and business continuity.

> ***From Experience:*** *One lesson I've learned is that organizations rarely struggle because they have too many policies. They struggle because policies don't connect to day-to-day operations. The most effective governance programs translate high-level expectations into practical procedures that employees understand and consistently follow.*

Questions Every Executive Team Should Ask

- Have we established an enterprise AI Governance Policy?
- Do employees understand when governance approval is required?
- Are AI responsibilities clearly documented?
- Are procedures consistent across business units?
- Have we standardized governance documentation?
- Are third-party AI providers subject to the same governance expectations?
- Are policies reviewed as technology and regulations evolve?
- Can employees easily locate and understand AI governance documentation?

Common Pitfalls

- Developing policies without supporting procedures.
- Creating documentation that is overly technical or difficult to understand.
- Allowing business units to create inconsistent governance practices.
- Failing to define ownership and approval responsibilities.
- Neglecting to update documentation as AI capabilities evolve.
- Overlooking third-party AI governance.
- Maintaining documentation that exists only to satisfy regulatory expectations rather than support business operations.

Policies & Procedures Readiness Checklist

- ☐ Enterprise AI Governance Policy approved.
- ☐ Supporting procedures documented.
- ☐ Roles and responsibilities clearly defined.
- ☐ AI approval process documented.
- ☐ Standard templates developed.
- ☐ AI documentation hierarchy established.
- ☐ Third-party AI governance incorporated.
- ☐ Policy review schedule established.
- ☐ Documentation integrated with existing governance programs.
- ☐ Employees trained on AI governance expectations.

Executive Takeaway

Policies and procedures provide the operational foundation for AI governance.

They transform governance principles into consistent organizational practices, establish clear expectations, and support responsible decision-making across the enterprise.

Organizations with practical, well-maintained documentation are better positioned to adapt to emerging technologies, respond to regulatory change, and scale AI responsibly.

In the next chapter, we'll examine Risk Assessment, the discipline that helps organizations identify, evaluate, and manage AI-related risks before they become enterprise issues.

Chapter 6: Risk Assessment

Purpose

Every AI initiative introduces risk.

The purpose of a risk assessment is to identify, evaluate, and understand those risks before AI is implemented and throughout its lifecycle.

A well-designed risk assessment enables organizations to make informed decisions, apply governance proportionate to risk, and allocate oversight where it is needed most.

Rather than preventing innovation, risk assessments help organizations understand potential impacts, evaluate whether risks fall within the organization's established risk tolerance, and determine whether additional governance or executive involvement is warranted.

Why Risk Assessments Matter

Artificial Intelligence is not inherently high risk. The level of risk depends on how the technology is used.

An internal productivity tool presents different considerations than an AI solution used to support lending decisions, fraud detection, customer service, marketing, or other activities that may directly affect consumers or business operations.

Organizations should avoid treating every AI solution the same. Instead, governance activities should be aligned with the level of inherent and residual risk associated with each AI use case.

A thoughtful risk assessment helps organizations identify higher-risk implementations, determine whether appropriate controls exist, and decide whether additional oversight or executive review is necessary before deployment.

> ***Executive Insight:*** *The objective of a risk assessment is not to eliminate risk. It is to understand risk well enough to make informed business decisions.*

Risk Assessment Principles

Risk Should Drive Governance. The level of governance should be appropriate for the level of risk. Higher-risk AI solutions generally require greater oversight, more rigorous testing, enhanced monitoring, increased executive visibility, and, when appropriate, Board awareness.

Assess Risk Before Implementation. Risk assessments should occur before an AI solution is implemented, not after it has been deployed. Understanding potential risks early allows organizations to incorporate appropriate controls before issues arise.

Consider the Entire AI Lifecycle. Risk assessments should not be limited to implementation. Organizations should periodically reassess AI solutions as technologies evolve, business objectives change, vendors introduce updates, or regulatory expectations develop.

Evaluate Multiple Risk Dimensions. Effective assessments consider a range of potential impacts, including consumer risk, regulatory compliance, operational resilience, model performance, data quality, privacy, information security, third-party risk, and reputational risk.

Reassess When Conditions Change. Risk assessments should be updated whenever there are significant changes to an AI solution, its intended use, supporting data, vendor relationships, or applicable regulatory requirements.

Establish a Risk Exception Process. Not every AI initiative will fall within an organization's established risk tolerance. When an AI initiative presents significant risk but there is a compelling business reason to proceed, organizations should establish a formal risk exception process. The process should require appropriate executive review, documentation of the decision, acknowledgement of the associated risks, and clearly defined accountability for accepting those risks.

Leading Practices

Organizations with mature AI governance programs often:

- Assess AI use cases before implementation.
- Apply governance based on the level of risk.
- Involve cross-functional stakeholders throughout the assessment process.
- Document significant assumptions and decisions.
- Reevaluate risks as AI capabilities evolve.
- Escalate higher-risk AI initiatives for executive review.
- Establish a formal risk exception process for AI initiatives that exceed the organization's defined risk tolerance.
- Evaluate third-party AI solutions using the same governance principles applied to internally developed AI.

> ***From Experience:*** *Throughout my career, I've found that the most effective risk assessments don't simply identify risk; they help leadership make informed decisions. The objective isn't to stop innovation. It's to understand the risks, determine what additional controls may be needed, and reduce those risks to a level the organization is comfortable accepting so the business can move forward responsibly.*

Questions Every Executive Team Should Ask

- Have we identified our highest-risk AI use cases?
- Do we understand why those use cases present elevated risk?
- Are governance activities aligned with the level of risk?
- Are risk assessments completed before implementation?
- Are assessments updated when significant changes occur?
- Are multiple business functions involved in the assessment process?
- Are third-party AI solutions evaluated using the same principles as internally developed AI?
- Have we established a formal process for reviewing AI initiatives that exceed our defined risk tolerance?
- Are executive decisions to proceed with higher-risk AI initiatives documented?
- Do executives receive appropriate visibility into the organization's highest-risk AI activities?

Common Pitfalls

- Treating every AI use case as equally risky.
- Completing risk assessments after implementation.
- Limiting assessments to technology risks while overlooking business, regulatory, and consumer impacts.
- Failing to reassess risks as AI solutions evolve.
- Conducting assessments without input from key stakeholders.
- Proceeding with higher-risk AI initiatives without executive visibility or documented risk acceptance.
- Allowing risk assessments to become a compliance exercise rather than a business decision-making tool.

Risk Assessment Readiness Checklist

- ☐ AI use cases are assessed before implementation.
- ☐ Governance activities are aligned with the level of risk.
- ☐ Multiple risk dimensions are considered.
- ☐ Cross-functional stakeholders participate in assessments.
- ☐ Significant assumptions and decisions are documented.
- ☐ Risk assessments are updated when material changes occur.
- ☐ Third-party AI solutions are assessed.
- ☐ A formal risk exception process exists for AI initiatives that exceed the organization's defined risk tolerance.
- ☐ Executive decisions involving risk exceptions are documented.
- ☐ Executive leadership receives appropriate visibility into higher-risk AI initiatives.

Executive Takeaway

Risk assessments are one of the most effective tools for enabling responsible AI adoption.

They help organizations understand potential impacts, apply governance proportionate to risk, and support informed decision-making throughout the AI lifecycle.

When AI initiatives exceed an organization's established risk tolerance, a formal risk exception process provides executive leadership with the visibility needed to evaluate the risks, determine whether additional controls are necessary, and make informed decisions about how the organization should proceed.

In the next chapter, we'll examine Regulatory Change Management and the role it plays in helping organizations adapt their AI governance programs to an evolving regulatory landscape.

Chapter 7: Regulatory Change Management

Purpose

The regulatory landscape for Artificial Intelligence continues to evolve at a rapid pace.

The purpose of Regulatory Change Management is to help organizations identify, assess, communicate, and implement changes resulting from new laws, regulations, supervisory guidance, enforcement actions, and industry expectations.

An effective regulatory change management process enables organizations to adapt their AI governance programs while maintaining compliance and supporting responsible innovation.

Rather than reacting to regulatory developments after they occur, organizations should establish a structured process for evaluating change and determining its potential impact on the business.

Why Regulatory Change Management Matters

Artificial Intelligence is transforming faster than most regulatory frameworks can adapt.

While many existing laws already apply to AI, regulators continue to issue new guidance, interpret existing requirements differently, and increase supervisory expectations as AI adoption grows.

Organizations that fail to monitor these developments may expose themselves to unnecessary regulatory, operational, legal, and reputational risk.

Regulatory change management helps organizations anticipate change, understand potential impacts, and make informed decisions before those changes affect business operations.

> ***Executive Insight:*** *Regulatory change should not be viewed as an interruption to the business. It should be viewed as an opportunity to strengthen governance, improve processes, and build greater organizational resilience.*

Core Principles

Regulatory Change Is a Continuous Process. Monitoring regulatory developments is not a one-time activity. Organizations should continuously evaluate changes affecting Artificial Intelligence and determine whether action is required.

Existing Regulations Still Apply. Organizations should not focus solely on AI-specific regulations. Many existing consumer protection, privacy, fair lending, model risk, third-party risk, information security, and compliance requirements continue to apply regardless of whether AI is involved.

Evaluate Business Impact. Organizations should determine whether a regulatory change applies to the organization, identify affected business areas, evaluate implementation requirements, assess associated risks, and establish appropriate priorities.

Risk-Based Prioritization. Organizations should prioritize regulatory changes based on their applicability, complexity, potential consumer impact, implementation effort, and overall risk to the organization.

Engage Cross-Functional Stakeholders. Legal, Risk Management, Information Security, Technology, Product, Model Risk Management, Internal Audit, and business leaders should collaborate to understand potential impacts and coordinate implementation activities.

Assign Clear Ownership. Every regulatory change should have an identified owner responsible for evaluating, implementing, and validating required actions.

Document Decisions. Organizations should document how regulatory developments were assessed, the decisions that were made, the rationale supporting those decisions, and any required actions.

Update Policies, Procedures, and Controls. Organizations should evaluate whether existing policies, procedures, internal controls, training, reporting, governance documentation, and related processes require updates to reflect new or revised regulatory expectations.

Validate Implementation. Organizations should validate that updated policies, procedures, controls, and operational processes are functioning as intended through testing, quality assurance reviews, monitoring activities, control validation, or other independent verification appropriate to the level of risk.

Leading Practices

Organizations with mature regulatory change management programs often:

- Continuously monitor regulatory developments affecting AI.
- Evaluate regulatory changes using a consistent governance process.
- Prioritize regulatory changes based on risk and business impact.
- Assess business impacts before implementation.
- Engage appropriate cross-functional stakeholders.
- Assign clear ownership for required actions.
- Document decisions and supporting rationale.
- Update policies, procedures, controls, training, and governance documentation following implementation.
- Validate that implemented changes are operating as intended.
- Communicate significant regulatory developments to executive leadership and governance committees, as appropriate.

> ***From Experience:*** *Throughout my career, regulatory change has been one of the few constants. The organizations that navigate change most effectively are not necessarily those with the largest compliance teams, but those with disciplined governance, clear accountability, and the ability to thoughtfully evaluate change before taking action.*

Questions Every Executive Team Should Ask

- Are we monitoring regulatory developments affecting AI?
- Do we have a consistent process for evaluating regulatory changes?
- Are regulatory changes prioritized based on risk and business impact?
- Do we understand how regulatory changes affect our existing AI governance program?
- Are responsibilities for implementation clearly assigned?
- Have policies, procedures, and controls been updated to reflect regulatory changes?
- Have implemented changes been validated to confirm they are operating as intended?
- Are significant regulatory developments communicated to executive leadership?
- Are regulatory decisions documented and supported?
- Are cross-functional stakeholders engaged throughout the evaluation and implementation process?

Common Pitfalls

- Reacting to regulatory changes instead of anticipating them.
- Monitoring regulations without evaluating business impact.
- Treating every regulatory development with the same level of urgency.
- Assigning responsibility without clear accountability.
- Failing to coordinate implementation across business functions.
- Overlooking existing regulations that already apply to AI.

- Failing to update policies, procedures, controls, or training following implementation.
- Assuming implementation is complete without validating that changes are operating effectively.
- Implementing changes without documenting decisions or supporting rationale.

Regulatory Change Management Readiness Checklist

- ☐ Regulatory developments affecting AI are monitored on an ongoing basis.
- ☐ Regulatory changes are evaluated using a consistent process.
- ☐ Regulatory changes are prioritized based on risk and business impact.
- ☐ Business impacts are assessed before implementation.
- ☐ Cross-functional stakeholders participate throughout the evaluation process.
- ☐ Responsibilities for implementation are clearly assigned.
- ☐ Regulatory decisions and supporting rationale are documented.
- ☐ Policies, procedures, controls, training, and governance documentation are updated following implementation.
- ☐ Implemented regulatory changes are validated to confirm they are operating as intended.
- ☐ Significant regulatory developments are communicated to executive leadership and governance committees, as appropriate.

Executive Takeaway

Regulatory change management is more than monitoring new laws and guidance.

It is a disciplined governance process that enables organizations to evaluate change, implement appropriate actions, update their governance framework, and validate that those changes are operating effectively.

Organizations that integrate regulatory change management into their broader AI governance program are better positioned to adapt to evolving regulatory expectations while continuing to support responsible innovation.

In the next chapter, we'll examine Testing & Monitoring and the important role it plays in providing ongoing assurance that AI governance processes and controls continue to operate effectively.

Chapter 8: Testing & Monitoring

Purpose

Testing and monitoring provide ongoing assurance that an organization's AI governance program is operating as intended.

While policies, procedures, and controls establish the governance framework, testing and monitoring help determine whether those controls continue to function effectively as business activities, technologies, regulatory expectations, and organizational risks evolve.

Together, they provide leadership with objective information about the effectiveness of the organization's AI governance program, identify opportunities for improvement, and support continuous enhancement of governance processes and controls.

Why Testing & Monitoring Matter

Designing effective controls is only the beginning.

Organizations should periodically evaluate whether governance processes are functioning as intended and whether established controls continue to effectively manage identified risks.

Artificial Intelligence systems are dynamic. Models evolve, data changes, business processes mature, third-party providers introduce updates, and regulatory expectations continue to develop.

Testing and monitoring help organizations identify control weaknesses, implementation gaps, emerging risks, and opportunities for improvement before they become significant issues.

> ***Executive Insight:*** *Organizations should not assume controls are effective simply because they have been implemented. Confidence is achieved through ongoing testing, monitoring, and continuous improvement.*

Core Principles

Testing Should Be Risk-Based. Organizations should prioritize testing activities based on the level of inherent and residual risk, regulatory requirements, consumer impact, business significance, and the complexity of the AI solution.

Testing Should Be Performed with Appropriate Independence. Testing should be performed objectively to provide leadership with an unbiased assessment of whether controls are designed appropriately and operating effectively. The level of independence will vary based on the size, complexity, and governance structure of the organization.

Evaluate Both Design and Operating Effectiveness. Effective testing evaluates whether a control is designed appropriately to address the identified risk, and whether it is operating consistently as intended.

Monitoring Should Be Ongoing. Monitoring provides continuous oversight between formal testing activities. Organizations should establish monitoring processes appropriate to their size and risk profile to identify operational changes, emerging risks, control breakdowns, performance trends, and evolving regulatory expectations.

Report Results Transparently. Testing and monitoring results should be communicated to the appropriate levels of management. Significant findings, recurring issues, emerging trends, and material control weaknesses should receive executive visibility through established governance and reporting processes.

Remediate Issues Promptly. Organizations should establish processes for assigning ownership, implementing corrective actions, tracking progress, and reporting remediation status until issues have been resolved.

Validate Corrective Actions. Organizations should confirm that corrective actions have been implemented effectively and that the underlying issue has been successfully addressed.

Leading Practices

Organizations with mature testing and monitoring programs often:

- Develop risk-based testing plans.
- Perform testing with appropriate independence.
- Evaluate both control design and operating effectiveness.
- Conduct ongoing monitoring between formal testing cycles.
- Escalate significant findings through established governance processes.
- Assign ownership and due dates for corrective actions.
- Validate remediation before closing issues.
- Analyze recurring findings to identify systemic issues.
- Report testing results, trends, and emerging risks to executive leadership and governance committees, as appropriate.

> ***From Experience:*** *Throughout my career, I've found that the most effective testing programs are not designed to find fault; they are designed to provide confidence. Effective testing validates what is working, identifies opportunities for improvement, and provides leadership with objective insight into whether governance processes and controls continue to operate as intended.*

Questions Every Executive Team Should Ask

- Are we testing the areas that present the greatest AI risk?
- Is testing performed with appropriate independence?
- Are both control design and operating effectiveness being evaluated?
- Do our monitoring activities identify issues before they become significant risks?
- Are significant findings communicated to executive leadership?
- Are corrective actions assigned, tracked, and completed on a timely basis?
- Are remediation efforts validated before issues are closed?
- Are recurring issues analyzed to identify broader governance concerns?
- Are testing results being used to strengthen our AI governance program?

Common Pitfalls

- Testing low-risk activities while overlooking higher-risk areas.
- Failing to maintain sufficient objectivity within the testing process.
- Evaluating control design without confirming operating effectiveness.
- Conducting monitoring inconsistently or infrequently.
- Closing issues before corrective actions have been validated.
- Repeating the same findings because underlying causes were not addressed.
- Viewing testing solely as a regulatory obligation rather than a governance function.
- Failing to communicate significant trends and recurring issues to executive leadership.

Testing & Monitoring Readiness Checklist

- ☐ Testing activities are prioritized using a risk-based approach.
- ☐ Testing is performed with appropriate independence.
- ☐ Both control design and operating effectiveness are evaluated.
- ☐ Monitoring activities occur between formal testing cycles.
- ☐ Significant findings are communicated through established governance channels.
- ☐ Corrective actions are assigned, tracked, and monitored to completion.
- ☐ Remediation efforts are validated before issues are closed.
- ☐ Recurring findings are analyzed to identify systemic issues.
- ☐ Testing and monitoring results are used to strengthen governance processes and controls.

Executive Takeaway

Testing and monitoring provide ongoing assurance that an organization's AI governance program continues to operate effectively.

They help organizations validate controls, identify emerging risks, strengthen governance, and support continuous improvement.

Whether performed by dedicated compliance testing teams, operational management, internal audit, or a combination of governance functions appropriate to the organization's size and complexity, effective testing and monitoring provide leadership with confidence that AI governance processes remain effective as the organization, technology, and regulatory environment continue to evolve.

In the next chapter, we'll examine Complaint Management and how customer feedback, complaints, and issue trends can provide valuable insight into the effectiveness of an organization's AI governance program.

Chapter 9: Complaint Management

Purpose

Complaint management provides organizations with valuable insight into how Artificial Intelligence systems perform in real-world situations.

Customer complaints can identify emerging risks, reveal control weaknesses, highlight unintended outcomes, and provide early warning of issues that may not be detected through testing, monitoring, or other governance activities.

An effective complaint management program helps organizations identify trends, strengthen governance, improve customer outcomes, and support continuous improvement across the AI lifecycle.

Why Complaint Management Matters

Customer complaints are more than service issues. They provide valuable information about the effectiveness of an organization's AI governance program.

Complaints may identify concerns involving fairness, transparency, privacy, inaccurate decisions, operational failures, discrimination, model performance, or customer experience.

Organizations that actively analyze complaint data are often better positioned to identify emerging risks before they become regulatory, legal, operational, or reputational issues.

> ***Executive Insight:*** *Every complaint represents an opportunity to learn. Organizations that listen carefully to customer feedback often identify governance weaknesses long before they become larger business or regulatory concerns.*

Core Principles

Make It Easy for Customers to Raise Concerns. Organizations should provide customers with clear, accessible methods for submitting complaints or raising concerns regarding AI-supported products, services, or decisions.

Investigate Complaints Objectively. Every complaint should be evaluated fairly and consistently. Organizations should establish processes for investigating complaints, gathering relevant information, determining root causes, and documenting outcomes.

Identify Trends and Root Causes. Patterns across multiple complaints are often even more valuable than individual complaints. Organizations should analyze complaint data to identify recurring issues, emerging risks, operational weaknesses, and opportunities for improvement.

Escalate Significant Issues. Organizations should establish criteria for escalating complaints involving significant consumer harm, legal or regulatory concerns, operational failures, model performance issues, or reputational risk.

Use Complaints to Strengthen Governance. Organizations should evaluate whether complaints indicate a need to improve policies, procedures, controls, models, training, governance processes, or customer communications.

Monitor Complaint Trends Over Time. Organizations should periodically review complaint trends to identify changing risks, evaluate the effectiveness of corrective actions, and assess whether governance improvements are producing the intended results.

Leading Practices

Organizations with mature complaint management programs often:

- Maintain centralized complaint tracking.
- Categorize complaints consistently.

- Analyze complaint trends regularly.
- Perform root cause analysis.
- Escalate significant complaints through established governance processes.
- Coordinate investigations across business functions.
- Use complaint data to strengthen policies, controls, training, and governance processes.
- Periodically report complaint trends to executive leadership and governance committees.

> ***From Experience:*** *Throughout my career, I've found that complaint data often tells a story long before formal testing or regulatory examinations do. Organizations that treat complaints as a source of operational intelligence rather than simply a customer service function are better positioned to identify emerging risks, improve customer outcomes, and strengthen governance across the enterprise.*

Questions Every Executive Team Should Ask

- Are customers able to easily report concerns?
- Are complaints investigated consistently?
- Do we perform root cause analysis?
- Are complaint trends regularly analyzed?
- Are significant complaints escalated appropriately?
- Are complaints identifying weaknesses in our AI governance program?
- Are corrective actions implemented and monitored?
- Are complaint trends reported to executive leadership?

Common Pitfalls

- Treating complaints solely as customer service issues.
- Investigating complaints individually without analyzing trends.
- Failing to identify root causes.
- Closing complaints without implementing corrective actions.
- Inconsistent complaint categorization.
- Poor communication between business functions.
- Failing to escalate significant issues.
- Missing opportunities to strengthen governance based on customer feedback.

Complaint Management Readiness Checklist

- ☐ Customers have accessible methods for submitting complaints.
- ☐ Complaints are investigated consistently.
- ☐ Complaint data is categorized and tracked.
- ☐ Root cause analysis is performed when appropriate.
- ☐ Complaint trends are reviewed regularly.
- ☐ Significant complaints are escalated through established governance processes.
- ☐ Corrective actions are implemented and monitored.
- ☐ Complaint trends are reported to executive leadership.
- ☐ Complaint insights are used to improve governance processes and controls.

Executive Takeaway

Complaint management is an important component of an effective AI governance program.

Customer feedback provides valuable insight into how AI systems perform in practice and helps organizations identify emerging risks, strengthen controls, and improve customer outcomes.

Organizations that integrate complaint management into their broader governance framework are better positioned to detect issues early, support continuous improvement, and maintain trust as AI technologies continue to evolve.

In the next chapter, we'll examine Third-Party Risk Management and the important role vendors, service providers, and technology partners play in an organization's AI governance program.

Chapter 10: Third-Party Risk Management

Purpose

Third-party relationships play an increasingly important role in the development, deployment, and operation of Artificial Intelligence.

Organizations frequently rely on vendors, cloud providers, software platforms, data providers, consultants, and other external partners to support AI capabilities.

The purpose of Third-Party Risk Management is to help organizations identify, assess, manage, and monitor the risks associated with these relationships while maintaining appropriate governance and accountability.

Although organizations may outsource AI solutions, they cannot outsource responsibility.

Why Third-Party Risk Management Matters

Third parties often have access to sensitive data, influence business decisions, develop AI models, or provide technology that directly affects customers and business operations.

Weak governance over third-party relationships may expose organizations to operational, regulatory, legal, cybersecurity, privacy, financial, and reputational risk.

Effective third-party risk management helps organizations understand these risks, establish appropriate oversight, and promote responsible use of externally provided AI solutions.

> ***Executive Insight:*** *Organizations may outsource technology, but they cannot outsource accountability. Effective governance requires understanding and managing the risks introduced by third-party relationships.*

Core Principles

Maintain Accountability. Organizations remain accountable for AI solutions used within their business, regardless of whether those solutions are developed internally or provided by a third party.

Perform Risk-Based Due Diligence. Organizations should evaluate third parties based on the services provided, access to sensitive information, use of AI, regulatory impact, operational criticality, and potential customer impact.

Understand the AI Solution. Organizations should develop an appropriate understanding of how third-party AI solutions support business activities, including the intended purpose, key inputs and outputs, limitations, governance practices, and associated risks.

Establish Appropriate Contractual Protections. Contracts should clearly define roles, responsibilities, performance expectations, information security requirements, confidentiality obligations, audit rights, regulatory cooperation, and other governance expectations appropriate to the relationship.

Monitor Third-Party Performance. Organizations should periodically monitor vendor performance, risk indicators, regulatory developments, cybersecurity events, financial condition, service quality, and compliance with contractual obligations.

Coordinate Across Business Functions. Procurement, Compliance, Risk Management, Legal, Information Security, Privacy, Technology, business leadership, and vendor management teams should work together to evaluate and oversee third-party AI relationships.

Prepare for Change. Organizations should establish processes for managing material changes, contract renewals, service modifications, significant incidents, and relationship termination.

Leading Practices

Organizations with mature third-party risk management programs often:

- Maintain an inventory of AI-related third-party relationships.
- Perform risk-based due diligence before onboarding vendors.
- Evaluate AI capabilities and associated risks.
- Establish contractual governance expectations.
- Monitor vendor performance throughout the relationship.
- Coordinate oversight across business functions.
- Reassess vendor risk periodically.
- Report significant third-party risks to executive leadership and governance committees, as appropriate.

> ***From Experience:*** *Throughout my career, I've found that some of an organization's most significant risks originate outside the organization itself. Effective third-party risk management is not about eliminating vendor relationships; it is about understanding the risks they introduce, establishing appropriate oversight, and maintaining accountability for customer outcomes regardless of who provides the technology.*

Questions Every Executive Team Should Ask

- Do we maintain an inventory of AI-related third-party relationships?
- Are vendors evaluated using a risk-based approach?
- Do we understand how third-party AI solutions support our business?
- Have appropriate contractual protections been established?
- Are higher-risk vendors receiving enhanced oversight?
- Are vendors monitored throughout the relationship?
- Are material third-party risks communicated to executive leadership?
- Are we prepared to respond if a critical third-party relationship changes or fails?

Common Pitfalls

- Assuming vendors are responsible for regulatory compliance.
- Applying the same level of due diligence to every vendor regardless of risk.
- Failing to understand how third-party AI solutions operate.
- Weak contractual governance.
- Treating vendor oversight as a one-time onboarding activity.
- Poor coordination across business functions.
- Inadequate monitoring of changing vendor risks.
- Failing to plan for significant vendor disruptions or termination.

Third-Party Risk Management Readiness Checklist

- ☐ AI-related third-party relationships are identified and inventoried.
- ☐ Vendors are evaluated using a risk-based approach.
- ☐ Higher-risk vendors receive enhanced due diligence.
- ☐ AI solutions are understood sufficiently to support effective oversight.
- ☐ Appropriate contractual protections have been established.
- ☐ Vendor performance is monitored throughout the relationship.
- ☐ Material third-party risks are communicated through established governance processes.
- ☐ Vendor risks are reassessed periodically.
- ☐ Plans exist for significant vendor changes or relationship termination.

Executive Takeaway

Third-party AI solutions offer significant opportunities for innovation, but they also introduce unique risks that require thoughtful governance.

Organizations remain accountable for AI-enabled products and services regardless of whether they are developed internally or provided by external partners.

A disciplined, risk-based third-party risk management program helps organizations strengthen oversight, support responsible innovation, and maintain confidence that third-party relationships continue to align with the organization's governance objectives.

In the next chapter, we'll examine Training & Awareness and the important role education plays in building a culture of responsible AI governance.

Chapter 11: Training & Awareness

Purpose

An effective AI governance program depends on more than policies, procedures, and controls. It also depends on the people responsible for designing, deploying, overseeing, and using Artificial Intelligence.

The purpose of Training & Awareness is to help employees understand their roles and responsibilities, promote responsible AI practices, and foster a culture of accountability throughout the organization.

Effective training supports informed decision-making, strengthens governance, and helps organizations adapt as AI technologies and regulatory expectations continue to evolve.

Why Training & Awareness Matter

Even the strongest governance framework can be undermined if employees do not understand how to apply it.

Employees make decisions every day that influence how AI systems are developed, implemented, monitored, and used. Without appropriate education, organizations increase the risk of inconsistent practices, control failures, regulatory noncompliance, and unintended consequences.

Training and awareness help create a shared understanding of governance expectations and reinforce responsible AI practices across the organization.

> ***Executive Insight:*** *AI governance is not achieved through documentation alone. It becomes effective when employees understand their responsibilities and consistently apply governance principles in their daily work.*

Core Principles

Establish a Culture of Responsible AI. Training should reinforce that responsible AI is everyone's responsibility. Employees should understand that governance is intended to support responsible innovation, sound decision-making, regulatory compliance, and customer trust.

Provide Role-Based Training. Training should be tailored to the responsibilities of executives, boards, business leaders, developers, data scientists, compliance professionals, risk managers, internal auditors, customer-facing employees, and others based on their roles.

Communicate Governance Expectations Clearly. Employees should understand the organization's AI governance framework, applicable policies, approval processes, risk management expectations, and reporting responsibilities.

Keep Training Current. Training should be reviewed and updated periodically to reflect changes in technology, business practices, organizational priorities, and regulatory expectations.

Reinforce Learning Through Ongoing Awareness. Organizations should reinforce governance expectations through ongoing communications, leadership messaging, awareness campaigns, job aids, and periodic reminders.

Measure Training Effectiveness. Participation rates, employee feedback, knowledge assessments, governance observations, and other appropriate measures can provide insight into the effectiveness of the training program.

Leading Practices

Organizations with mature training and awareness programs often:

- Develop role-based AI training.

- Provide governance education during employee onboarding.
- Update training as regulatory expectations evolve.
- Reinforce key governance principles through ongoing awareness activities.
- Maintain records demonstrating training completion.
- Evaluate training effectiveness periodically.
- Communicate governance expectations consistently across the organization.
- Report significant training initiatives and completion metrics to executive leadership, as appropriate.

> ***From Experience:*** *Throughout my career, I've found that governance is most effective when employees understand not only what is expected of them, but why it matters. Well-designed training programs build confidence, encourage consistent decision-making, and help organizations create a culture where governance becomes part of everyday business operations rather than a separate compliance activity.*

Questions Every Executive Team Should Ask

- Do employees understand their AI governance responsibilities?
- Is training tailored to different roles?
- Are governance expectations communicated clearly?
- Is training updated as technology and regulations evolve?
- Are awareness activities reinforcing key governance principles?
- Are we measuring whether training is effective?
- Does our culture encourage employees to raise concerns and ask questions?
- Are leaders visibly supporting responsible AI practices?

Common Pitfalls

- Treating training as a one-time event.
- Delivering the same training to every employee regardless of role.
- Allowing training materials to become outdated.
- Focusing on policies without explaining practical responsibilities.
- Measuring completion rather than effectiveness.
- Failing to reinforce governance expectations after formal training.
- Assuming employees understand AI risks without providing education.

Training & Awareness Readiness Checklist

- ☐ Employees understand their AI governance responsibilities.
- ☐ Training is tailored to different roles.
- ☐ Governance expectations are communicated clearly.
- ☐ Training materials are reviewed and updated periodically.
- ☐ Ongoing awareness activities reinforce responsible AI practices.
- ☐ Training effectiveness is evaluated.
- ☐ Training records are maintained appropriately.
- ☐ Leadership actively supports AI governance education.
- ☐ Training contributes to a culture of responsible AI.

Executive Takeaway

Training and awareness are essential components of an effective AI governance program.

They help employees understand their responsibilities, reinforce governance expectations, strengthen organizational culture, and support consistent decision-making across the enterprise.

Organizations that invest in ongoing education are better positioned to adapt to evolving technologies, regulatory expectations, and business priorities while fostering responsible innovation.

In the next chapter, we'll examine Board and Executive Reporting and the critical role governance reporting plays in providing leadership with meaningful oversight of the organization's AI governance program.

Chapter 12: Board and Executive Reporting

Purpose

Effective governance depends on informed leadership.

The purpose of Board and Executive Reporting is to provide leadership with meaningful, timely, and actionable information regarding the organization's AI governance program.

Reporting enables boards, executive management, and governance committees to understand the organization's AI activities, evaluate risks, monitor performance, oversee remediation efforts, and make informed strategic decisions.

Well-designed reporting promotes transparency, accountability, and sound governance.

Why Board and Executive Reporting Matter

Artificial Intelligence introduces opportunities as well as new and evolving risks.

Leadership is responsible for overseeing these risks, but effective oversight is only possible when decision-makers receive relevant, reliable, and understandable information.

Reporting should move beyond simply presenting data. It should provide insight into the effectiveness of the organization's AI governance program, highlight emerging risks, identify significant issues, and support informed decision-making.

> ***Executive Insight:*** *Effective governance reporting is not about providing more information. It is about providing the right information to the right people at the right time so leadership can make informed decisions.*

Core Principles

Report Information That Supports Decisions. Reporting should provide information that enables leadership to understand organizational risk, evaluate governance effectiveness, establish priorities, and make informed decisions.

Tailor Reporting to the Audience. Boards typically focus on strategic risks, governance effectiveness, emerging trends, and significant issues. Executive management may require greater operational detail to support implementation and decision-making.

Focus on Material Risks. Reporting should emphasize material risks, significant findings, emerging trends, unresolved issues, regulatory developments, and other information that could affect the organization's objectives or risk profile.

Report Trends, Not Just Individual Events. Long-term trends often provide greater insight into the effectiveness of the organization's AI governance program than individual issues alone.

Provide Balanced Reporting. Effective reporting should communicate both strengths and areas requiring improvement.

Support Accountability. Reporting should clearly identify responsible business areas, issue owners, implementation status, and significant decisions requiring executive attention.

Promote Continuous Improvement. Leadership should use reporting to identify opportunities for strengthening governance, improving controls, reducing risk, and supporting responsible innovation.

Leading Practices

Organizations with mature governance reporting programs often:

- Develop reporting tailored to different leadership audiences.
- Focus reporting on material risks and governance effectiveness.

- Include trend analysis rather than isolated metrics.
- Highlight significant regulatory developments and emerging risks.
- Report testing results, complaint trends, third-party risks, and remediation activities.
- Clearly identify ownership and implementation status for significant issues.
- Periodically review reporting to ensure it continues to meet leadership needs.

> ***From Experience:*** *Throughout my career, I've found that the best governance reports don't overwhelm leadership with information; they provide clarity. Executives and boards need meaningful insights that help them understand risk, evaluate governance effectiveness, and make informed decisions. Effective reporting builds confidence because it tells leadership not only what is happening, but why it matters.*

Questions Every Executive Team Should Ask

- Are we receiving the information necessary to oversee AI effectively?
- Do our reports focus on material risks and significant issues?
- Are reporting packages tailored to the needs of different audiences?
- Are we monitoring trends rather than isolated events?
- Do reports clearly identify ownership and accountability?
- Are significant issues being resolved in a timely manner?
- Are reporting results helping us strengthen governance?
- Does reporting support informed strategic decision-making?

Common Pitfalls

- Reporting excessive operational detail while overlooking strategic risks.
- Presenting data without meaningful analysis.
- Reporting individual issues without identifying trends.
- Delaying escalation of significant risks.
- Failing to identify clear ownership and accountability.
- Producing reports that are difficult for leadership to interpret.
- Measuring activity rather than governance effectiveness.

Board and Executive Reporting Readiness Checklist

- ☐ Reporting is tailored to different leadership audiences.
- ☐ Reports focus on material risks and governance effectiveness.
- ☐ Significant trends are identified and analyzed.
- ☐ Emerging risks are communicated promptly.
- ☐ Issue ownership and accountability are clearly identified.
- ☐ Reporting supports informed decision-making.
- ☐ Governance reports are reviewed periodically for continued effectiveness.
- ☐ Leadership uses reporting to drive continuous improvement.

Executive Takeaway

Board and executive reporting transforms governance information into meaningful oversight.

Effective reporting provides leadership with the insight needed to understand organizational risk, monitor governance effectiveness, evaluate emerging issues, and make informed strategic decisions.

Organizations that establish thoughtful, risk-based reporting practices strengthen accountability, improve governance, and position themselves to manage AI responsibly as technologies, business priorities, and regulatory expectations continue to evolve.

In the next chapter, we'll bring all of these governance components together by examining Continuous Improvement and how organizations can sustain and mature their AI governance program over time.

Chapter 13: Continuous Improvement

Purpose

An effective AI governance program is never static.

The purpose of Continuous Improvement is to help organizations evaluate the effectiveness of their AI governance program, respond to changing business and regulatory expectations, strengthen governance practices, and promote responsible innovation over time.

As Artificial Intelligence technologies continue to evolve, governance programs should evolve with them.

Why Continuous Improvement Matters

Organizations continuously adapt to changing business strategies, customer expectations, technologies, regulatory requirements, and emerging risks.

AI governance should evolve in the same manner. Governance programs that remain unchanged over time may become less effective as organizational priorities and risk profiles change.

Continuous improvement enables organizations to strengthen governance, improve operational effectiveness, enhance decision-making, and maintain confidence in their AI oversight framework.

> ***Executive Insight:*** *The objective of AI governance is not to build a perfect program. It is to build a program that continuously learns, adapts, and improves as the organization and its risks evolve.*

Core Principles

Evaluate Governance Periodically. Organizations should periodically assess the effectiveness of their AI governance program, considering governance structures, policies, procedures, controls, oversight activities, reporting, organizational responsibilities, and overall program performance.

Learn From Experience. Organizations generate valuable lessons through testing, monitoring, complaint management, incidents, audits, regulatory examinations, and day-to-day operations. These lessons should be incorporated into governance processes to strengthen future performance.

Respond to Change. Business priorities, technology, organizational structures, and regulatory expectations continue to evolve. Governance programs should adapt accordingly to remain effective and aligned with organizational objectives.

Encourage Innovation. Continuous improvement is not intended to slow innovation. Strong governance enables organizations to pursue innovation responsibly by providing clear expectations, effective oversight, and informed decision-making.

Measure Program Effectiveness. Organizations should periodically evaluate whether governance activities are achieving their intended objectives, helping leadership understand program effectiveness, identify trends, prioritize improvements, and allocate resources appropriately.

Promote Collaboration. Business leaders, technology teams, compliance professionals, risk managers, legal counsel, internal audit, and executive leadership should collaborate to identify opportunities to strengthen governance.

Foster a Culture of Improvement. Employees should feel encouraged to identify opportunities, raise concerns, recommend improvements, and participate in strengthening governance processes.

Leading Practices

Organizations with mature AI governance programs often:

- Periodically evaluate governance effectiveness.
- Review governance following significant organizational or regulatory changes.
- Incorporate lessons learned from testing, monitoring, complaints, incidents, audits, and examinations.
- Monitor governance trends over time.
- Encourage cross-functional collaboration.
- Measure governance effectiveness using meaningful performance indicators.
- Periodically review governance reporting to ensure it remains relevant.
- Continuously refine policies, procedures, controls, and governance practices.

> ***From Experience:*** *Throughout my career, I've found that the strongest governance programs are those that never stop improving. Organizations that regularly evaluate their governance framework, learn from experience, and adapt to change are better positioned to manage risk while continuing to support business growth and innovation. Continuous improvement transforms governance from a compliance obligation into a strategic business capability.*

Questions Every Executive Team Should Ask

- Is our AI governance program becoming more effective over time?
- Are we learning from testing, monitoring, complaints, incidents, audits, and examinations?
- Are governance processes evolving as technology and regulations change?
- Are we measuring the effectiveness of our governance program?
- Are improvements prioritized based on organizational risk?
- Do employees feel comfortable identifying opportunities for improvement?
- Are governance improvements communicated across the organization?
- Does our governance program continue to support responsible innovation?

Common Pitfalls

- Treating governance implementation as the finish line.
- Making improvements only after significant issues occur.
- Failing to incorporate lessons learned.
- Measuring activity rather than program effectiveness.
- Resisting governance changes as organizational needs evolve.
- Limiting improvement efforts to the compliance function.
- Allowing governance documentation to become outdated.

Continuous Improvement Readiness Checklist

- ☐ Governance effectiveness is evaluated periodically.
- ☐ Lessons learned are incorporated into governance improvements.
- ☐ Governance evolves as business, technology, and regulations change.
- ☐ Program effectiveness is measured using meaningful performance indicators.
- ☐ Cross-functional collaboration supports governance improvements.
- ☐ Employees are encouraged to identify opportunities for improvement.
- ☐ Governance documentation is reviewed and updated periodically.
- ☐ Leadership actively supports continuous improvement.
- ☐ Governance improvements are prioritized based on organizational risk.

Executive Takeaway

AI governance is not a one-time initiative.

It is an ongoing discipline that should continue to mature as the organization, its technologies, and its regulatory environment evolve.

Organizations that embrace continuous improvement strengthen governance, improve decision-making, enhance operational resilience, and create a sustainable foundation for responsible AI innovation.

In the next chapter, we'll examine AI Incident Management and how organizations can prepare for, respond to, and learn from AI-related incidents in a structured and disciplined manner.

Chapter 14: Issue Management

Purpose

AI-related issues are inevitable. No organization can eliminate every risk or prevent every unexpected outcome. What distinguishes mature AI governance programs is not the absence of issues, but how effectively organizations integrate AI-related issues into their existing governance framework. Rather than creating separate processes, organizations should leverage established issue management, business continuity, and operational resilience programs to ensure AI-related issues are addressed consistently and effectively.

Why Issue Management Matters

Organizations already have established processes for identifying, escalating, investigating, and resolving enterprise issues. AI-related issues should be incorporated into these existing governance processes rather than managed separately.

Depending on the nature and severity of the issue, organizations may leverage existing:

- Cybersecurity incident management
- Technology outage management
- Data breach response
- Operational disruption management
- Crisis management
- Business continuity and operational resilience programs

AI governance should complement these enterprise capabilities by ensuring AI-related risks, impacts, and lessons learned are incorporated into the organization's overall governance framework.

> ***Executive Insight:*** *AI governance should not create a separate issue management process. Instead, it should ensure AI-related issues are managed through the organization's existing governance framework, with appropriate oversight, transparency, and accountability.*

Core Principles

Integrate with Existing Issue Management. Organizations should manage AI-related issues through established enterprise issue management processes whenever practical, avoiding unnecessary duplication of governance activities.

Leverage Existing Business Continuity and Resilience Programs. When AI-related issues significantly impact operations, customers, or critical services, organizations should activate existing business continuity, operational resilience, crisis management, or incident response programs, as appropriate.

Recognize That AI Issues Can Originate from Multiple Sources. AI-related issues may be identified through testing and monitoring, customer complaints, internal audits, regulatory examinations, operational events, cybersecurity incidents, third-party oversight, or employee reporting.

Maintain Appropriate Executive Visibility. Material AI-related issues should be reported through existing governance and reporting structures so leadership can understand organizational impacts, make informed decisions, and provide appropriate oversight.

Learn from Issues. Organizations should evaluate significant AI-related issues to identify opportunities to strengthen governance, improve controls, enhance oversight, and reduce the likelihood of similar issues occurring in the future.

Leading Practices

Mature organizations recognize that AI governance does not operate in isolation. Instead, they integrate AI oversight with existing enterprise governance programs, allowing AI-related issues to be managed consistently alongside other operational, compliance, technology, and risk events.

> ***From Experience:*** *Throughout my career, I have found that organizations are most effective when they build upon existing governance processes rather than creating new ones for every emerging risk. Whether the issue involves a regulatory finding, a control deficiency, a technology disruption, or an AI-related concern, the objective remains the same: understand the issue, assess its impact, respond appropriately, and use the experience to strengthen the organization.*

Questions Every Executive Team Should Ask

- Are AI-related issues incorporated into our existing issue management framework?
- Would a significant AI event be addressed through our existing business continuity or crisis management processes?
- Do our governance committees receive appropriate visibility into significant AI-related issues?
- Are lessons learned from AI-related issues being used to improve governance across the organization?

Common Pitfalls

- Creating separate issue management processes exclusively for AI.
- Treating AI issues differently from comparable enterprise risks.
- Failing to coordinate AI governance with existing business continuity and operational resilience programs.
- Overlooking opportunities to strengthen governance through lessons learned.

Readiness Checklist

- ☐ AI-related issues are incorporated into existing issue management processes.
- ☐ Business continuity and operational resilience programs consider AI-related events, where appropriate.
- ☐ Executive reporting includes material AI-related issues.
- ☐ Lessons learned are incorporated into governance improvements.

Executive Takeaway

Organizations do not need a separate issue management program simply because AI is involved. The strongest governance programs integrate AI-related issues into existing enterprise processes, ensuring consistent oversight, effective decision-making, and continuous improvement across the organization.

Chapter 15: The Future of AI Governance

Purpose

Artificial intelligence will continue to evolve, bringing new opportunities, emerging risks, and changing regulatory expectations. While no organization can predict every technological advancement or regulatory change, organizations can build governance programs that are flexible, resilient, and capable of adapting to an evolving AI landscape.

Why the Future of AI Governance Matters

AI is transforming how financial institutions operate, serve customers, manage risk, and compete in the marketplace. As AI capabilities continue to advance, governance must evolve alongside them. Organizations that establish strong governance foundations today will be better positioned to embrace innovation while maintaining regulatory compliance, managing risk, and preserving customer trust.

> ***Executive Insight:*** *The future of AI governance is not about predicting what comes next. It is about building a governance program that can adapt to whatever comes next.*

Core Principles

Governance Must Evolve with Technology. AI capabilities will continue to mature, introducing new applications, risks, and business opportunities. Governance frameworks should be periodically evaluated to ensure they remain aligned with organizational objectives and emerging technologies.

Regulatory Expectations Will Continue to Change. Legislatures, regulators, and industry standards will continue to shape AI governance. Organizations should remain informed and be prepared to evaluate how new requirements may affect their governance framework.

Governance Should Enable Responsible Innovation. Effective governance should support innovation by providing appropriate oversight, clear accountability, and sound decision-making.

Collaboration Will Become Increasingly Important. The future of AI governance will require continued collaboration across business functions, including executive leadership, compliance, legal, risk management, information security, technology, data governance, internal audit, and business operations.

Trust Will Remain a Competitive Advantage. Customers, regulators, investors, employees, and business partners increasingly expect organizations to use AI responsibly. Organizations that demonstrate transparency, accountability, and sound governance will strengthen stakeholder confidence and differentiate themselves in the marketplace.

Leading Practices

Leading organizations recognize that AI governance is not a one-time initiative. They periodically assess their governance framework, evaluate emerging risks, monitor regulatory developments, and continuously strengthen governance capabilities as their business and technology evolve.

> ***From Experience:*** *Throughout my career, I have seen technology transform the financial services industry many times. Regulations change, business models evolve, and new risks emerge. The organizations that succeed are not necessarily those with the most sophisticated technology; they are the ones that build strong governance, remain adaptable, and make thoughtful decisions as circumstances change. AI is no different. Organizations that establish a solid governance foundation today will be better prepared for whatever the future brings.*

Questions Every Executive Team Should Ask

- Is our AI governance framework adaptable as technology and regulations evolve?
- Are we periodically evaluating whether our governance framework remains effective?
- Does our governance program support innovation while maintaining appropriate oversight?
- Are we prepared to respond to future regulatory expectations and emerging risks?
- Are we building stakeholder trust through responsible AI governance?

Common Pitfalls

- Viewing AI governance as a one-time implementation.
- Assuming today's governance framework will remain sufficient indefinitely.
- Allowing innovation to outpace governance.
- Treating governance solely as a compliance obligation rather than a strategic business capability.
- Waiting for regulatory mandates before strengthening governance.

Readiness Checklist

- ☐ Our AI governance framework is periodically reviewed and updated.
- ☐ We monitor emerging regulatory and industry developments.
- ☐ Governance supports responsible innovation across the organization.
- ☐ Executive leadership remains actively engaged in AI governance.
- ☐ Continuous improvement is embedded within our governance program.

Executive Takeaway

Artificial intelligence will continue to evolve, and so will the expectations surrounding its use. Organizations that build adaptable governance frameworks today will be better positioned to manage future risks, embrace new opportunities, and maintain the confidence of regulators, customers, and other stakeholders. AI governance is not a destination; it is an ongoing commitment to responsible innovation, sound oversight, and continuous improvement.

Conclusion: Leading with Confidence

Artificial intelligence represents one of the most significant technological advancements in modern financial services. It has the potential to improve operational efficiency, strengthen decision-making, enhance customer experiences, and create new opportunities for growth. Realizing these benefits, however, requires more than technology. It requires governance.

Throughout this guide, we have explored the principles that support an effective AI governance program, from establishing governance structures and assessing risk to regulatory change management, testing, oversight, third-party risk, executive reporting, and continuous improvement. While every organization will tailor its governance framework to its size, complexity, and risk profile, the underlying principles remain the same.

Successful AI governance is not about slowing innovation. It is about enabling innovation responsibly. It provides the structure, accountability, and oversight necessary to make informed decisions while managing risk and maintaining trust.

As AI continues to evolve, organizations should resist the temptation to create entirely new governance structures for every emerging technology. Instead, they should build upon the governance capabilities they already have, integrating AI into existing enterprise risk management, compliance, operational resilience, and governance programs wherever practical.

The future will undoubtedly bring new technologies, new regulatory expectations, and new challenges. Organizations that invest in strong governance today will be better equipped to adapt with confidence tomorrow.

> *AI will continue to evolve. Good governance will ensure your organization is ready when it does.*