Preparing your financial institution for AI starts now

Accelerate your data and AI use cases with AI governance

Table of contents

The problem with AI and fragmented governance

The latest IDC report estimates AI will add ~$20T to the global economy by 2030[1]

The hype is real. And smart, ambitious financial services organizations are beaming with visions of the productivity and profits they'll achieve with AI. The market forecasts are bullish. The latest IDC report estimates AI will add ~$20T to the global economy by 2030.[1]

But there's a critical reality emerging—a widening gap between what banks, wealth management, investment and other financial services organizations want to do with data and AI and what they can do. A fundamental challenge persists: fragmented governance. In most organizations, data exists in isolated pockets, creating blind spots in visibility, context and access control. Fragmentation isn't simply a technical obstacle. It stifles the capacity to scale data and AI initiatives. It can strangle your business.

The problems with fragmented governance:

- Inability to leverage data across silos

- Inconsistent policy enforcement

- Reduced trust in data quality

- Limited visibility into data lineage

- Barriers to scalable AI adoption

[1]Source: 'IDC: Artificial Intelligence Will Contribute $19.9 Trillion to the Global Economy through 2030 and Drive 3.5% of Global GDP in 2030.' September 17 2024.

The new AI reality: Unified governance is your strategic accelerator

The stakes couldn't be higher. Financial organizations are rushing to adopt AI without solving their governance gaps, causing them to build their futures on fractured foundations. The organizations that will thrive in this new era are those that recognize a fundamental truth: unified governance isn't just about control—it's about acceleration. By implementing unified governance for data and AI across every system, source and user, you can create a foundation for rapid innovation while ensuring compliance and trust.

This shift requires a new paradigm: Data Confidence.

Data Confidence™ | Explained

Accelerating all your data and AI use cases, safely and with well-understood data

Data Confidence is the way you and your colleagues feel when your organization can accelerate every data and AI use case — without compromising on safety or quality.

It happens when governance becomes an enabler rather than a bottleneck. Your people can find, understand and use trusted data across every system. Business context flows alongside technical metadata. And policies apply consistently everywhere data lives.

Bottom line: When your people can trust, comply and consume data confidently, innovation accelerates. That's Data Confidence.

How this ebook can help you and your organization

Despite traditional attitudes about governance, AI governance isn't technological redtape. In reality, AI governance is an accelerator that can speed up all your data and AI use cases and create positive outcomes for both customers and employees when done right.

Taking an initial measured approach to AI and implementing proper AI governance delivers two clear benefits:

- Allows you to deploy ethical AI that's in compliance with regulatory requirements

- Helps stimulate your AI program for faster use case creation, increased ROI, and a better understanding of where to invest

Implementing AI governance may sound challenging. However, it doesn't have to be. This ebook is designed to:

- Introduce and guide you through an easy-to-implement AI governance framework

- Help you make an effective, successful AI governance program a reality at your organization

All together now: Why collaboration is key

Don't do it alone. Before you get started, gathering the right set of stakeholders to form an "AI roundtable" is critical to your success. Because of the inherent opportunities and risks involved with AI, it's important to have a wide range of viewpoints to help make well-rounded decisions. While every organization will be different and the make-up of your AI roundtable will be unique, there are a few critical organizational members you need to ensure have a seat at the table.

- Legal, Compliance, and Privacy teams to help assess legal risks and understand AI and data regulations and usage

- Chief Data Officer / Data Office participation to help with all data-related matters, including accessibility of trusted data

- Business Unit leaders who may have an interest in AI usage that could impact their teams

Some organizations may even have an Ethics Chief (an increasingly important position) and even members of HR to assess impacts to employees. Regardless of your organizational structure, it's essential to ensure that you collaborate on AI initiatives from the start.
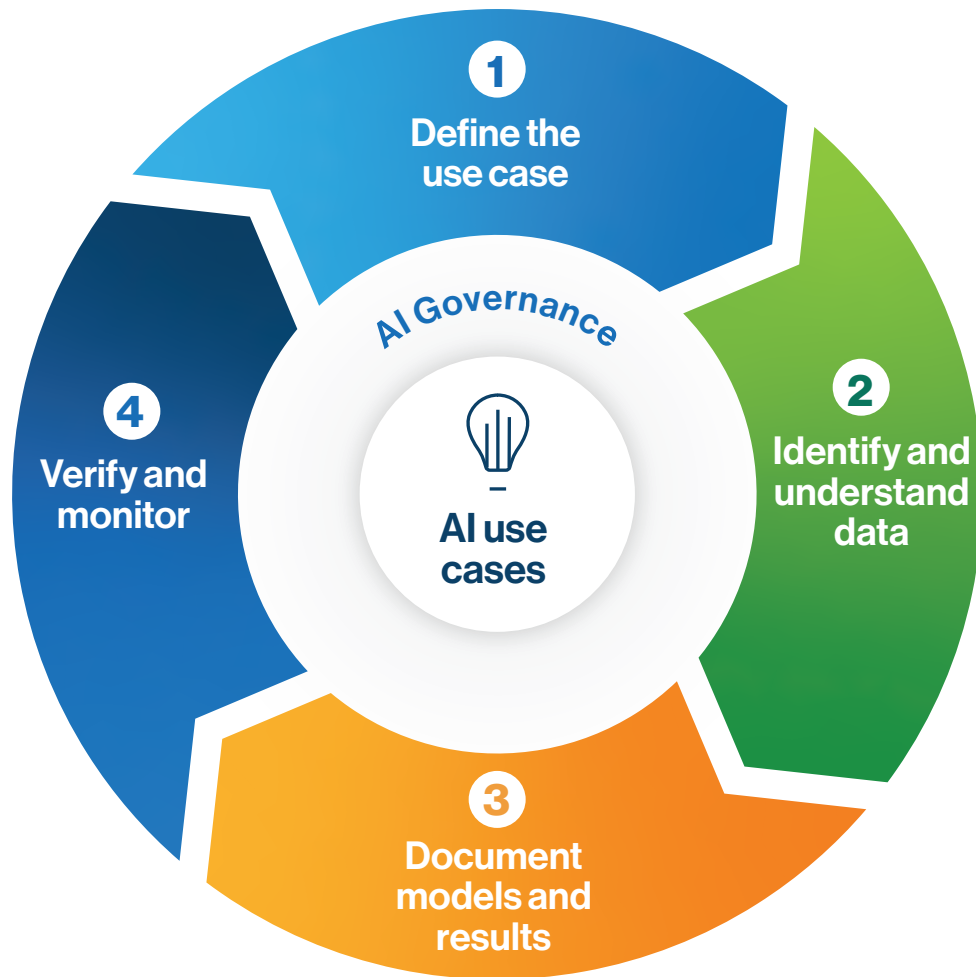
Our AI governance framework

Integrating AI into your roadmap requires a strategic approach. To navigate effectively, there are four essential steps your organization must take. These steps form the framework of a systematic, repeatable approach to AI.

1 Define the use case
Examples include credit score calculations or customer risk profiling

2 Identify and understand data
Will you use customer data that includes PII?

3 Document models and results
What do outputs look like?

4 Verify and monitor
Continuously check results to watch for drift or other irregularities

By following these steps, you can harness the full potential of AI, driving innovation and achieving significant competitive advantages.

1.
Define the use case

Before you start anything, you need to know what you're doing. The first step in your AI journey is to clearly define your use case. Knowing the intended purpose of your AI model — and where it will be deployed — should always be your first step.

A well-defined use case serves several purposes. It clarifies why the AI model is necessary, outlines the specific problems it aims to solve, and details the type of data it will utilize. Additionally, it provides a clear vision for the desired outcomes and the business value. More importantly, it helps your stakeholders and reviewers make informed decisions about how to move forward.

What goes into defining a use case?

Business context: Develop a well-documented use-case description that includes an analysis of the business value, the business policies the model may impact, and a list of business owners and their respective responsibilities.

Legal, ethics, and compliance: Assess whether the model will handle sensitive or private information, such as personal identifiable information (PII). Understand and document any specific regulations like the EU AI Act, GDPR, CCPA or other US state data and AI laws that may impact your AI model, along with risk assessments to ensure compliance and ethical considerations are addressed.

Data usage: Clearly outline the data required for the model, including what data will be used as input, how the model will be trained, and the nature of the output data.

By addressing these key areas, you'll ensure your AI roadmap is grounded in a thorough understanding of the broader context before you invest any resources into building an AI model.

Why Collibra AI Governance:

An AI governance solution that provides a single location to document AI use cases and collaborate with a wide range of stakeholders will help you keep track of AI use cases across the project lifecycle.

2.
Identify and understand data

Trust is everything in data and AI. It's why the old adage — "Garbage in, garbage out" — still holds true in our AI era. It explains why once you've defined your AI use cases, you need to take a close look at your data.

But how? Most financial institutions face an intractable challenge. Fragmented governance tethers control and visibility to specific data systems, sources or compute platforms and prevents companies from scaling their data and AI use cases safely. Data exists in pockets across apps, multiple public and private clouds, and on-prem, creating blind spots for what data exists, context for what it means and who has access.

The news doesn't get better the longer you wait. As data estates become more complex and LOBs spin up more use cases, this problem will only grow more complicated, while AI multiplies the risk of unreliable and noncompliant use. The disconnect extends to people as most systems offer no way to bring business users to access data and give it meaning.

It's why the cornerstone of any successful AI initiative is a deep understanding of the data your model will leverage. It includes understanding the nature of your data, as well as ensuring your compliance with all relevant laws and regulations. And it's why delivering trusted data for AI models starts with implementing the right data governance strategy. Rigorous guardrails will ensure you can operationalize AI workflows and processes to deliver trusted data.

To operationalize successfully, an enterprise data catalog is key. A game-changer for data scientists, a data catalog streamlines discovery and understanding of data across sources. The inclusion of a user-friendly data marketplace, in addition to a data catalog, helps data scientists find and access data in a fraction of the time compared to traditional methods of consulting stakeholders.

Data quality is another pivotal factor in the success of your AI initiatives. The active monitoring of data pipelines using advanced data quality and observability tools is crucial. These tools help quickly identify and resolve problems before they reach downstream outputs, such as your AI models. Finally, implementing clear data privacy policies ensures only authorized users can access specific datasets, reducing the risk of inappropriate data usage and reinforcing the integrity of your AI program.

3.

Document AI models and results

It's time to build. With a well-defined use case and high-quality data to feed your model, your focus shifts to building the AI model. It's crucial to document every detail during this process, including model outputs and challenges faced.

This step is where data scientists will focus most of their time. They'll document, trace and track the model, associated data products and usage. Comprehensive documentation is vital for model analysis and reporting. Data lineage is particularly essential in this phase; it ensures that you have clarity on the origin of the data, any transformations to it, and how and where outputs are used. This is especially useful for financial services institutions, which as you know, are highly regulated with reporting and compliance requirements on the minds of many stakeholders (you're probably thinking about it right now).

In this step, your primary goal is to get initial results.

Once you land on a model that passes scrutiny, you're ready for the hard step of moving into production.

4.

Verify and monitor

The final step isn't really a final step. It's important to remember that AI governance is not a one-time effort. Once your model is ready for production, it's vital you continually monitor results and revisit the legal and compliance requirements as new AI- and data-specific regulations are always coming into play.

The key aspects of this step include:

- Verifying model performance: Prior to full-scale deployment, it's critical to verify that the AI model acts as intended. Verification is a quality check, confirming that the model meets technical and business expectations

- Putting the model into production: Moving the AI model from a controlled testing environment into production is a big step. It involves integrating your model into your operational environment where it will start affecting real-world decisions. You'll also trace and document the flow of data through the AI system to understand how data is transformed and used in decision-making, which is crucial for troubleshooting and compliance

- Ongoing monitoring for data quality and compliance: Monitoring is vital to detect and address performance issues, data drifts or unexpected behavior. It involves tracking model output for accuracy, bias, and adherence to regulatory and ethical standards. You'll also vigilantly protect sensitive data, adhering to privacy regulations and ethical standards, especially as the model interacts with new datasets

- Retraining the model as needed: AI models are not set-and-forget tools. They require periodic retraining to incorporate new data, new regulations, and new technologies. Retraining is crucial to ensuring the model's accuracy and relevance

By following this step with an emphasis on data quality, data lineage, and data privacy, you can ensure your AI models remain relevant, robust, and compliant, as well as capable of adapting to new challenges and ensuring long-term effectiveness.

AI governance: Critical capabilities checklist for financial institutions

Researching AI governance solutions? A comprehensive AI governance platform that focuses on ensuring trusted data requires an array of advanced tools. Here are the key capabilities to look for:

☐ **Data Governance**

Automate data governance workflows to deliver trusted data to AI teams faster

☐ **Data Catalog**

Help AI teams discover and understand data they have for performant AI

☐ **Data Quality & Observability**

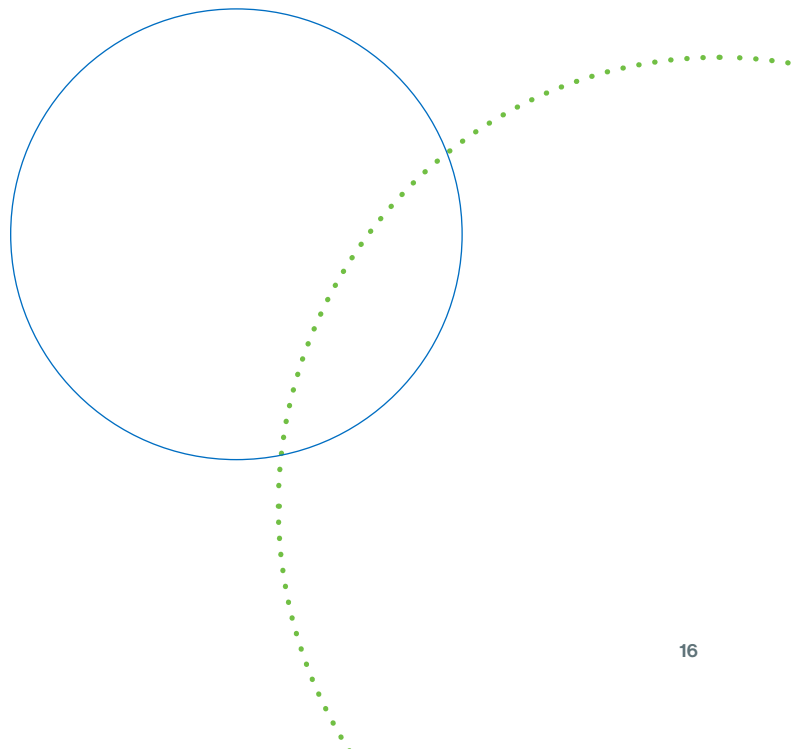Ensure high-quality data in every model built

☐ **Data Lineage**

Map input, training and output data for analysis and regulatory reporting

☐ **Data Privacy**

Automate and operationalize privacy and address both global AI and data regulations

☐ **Data Access Governance**

Classify and protect data across all data sources

From ambition to acceleration: Your four-step journey to AI governance

AI is widening a gap between what financial institutions want to do with AI and what they can do. The difference? Governance. Organizations stifled by fragmented governance will watch their AI initiatives stall and competitors leave them behind.

But there's a better way. When you unify governance across your entire data ecosystem — breaking down silos, connecting teams and creating a single source of truth—you unlock Data Confidence. You're confident that all your data and AI initiatives are fueled by trusted, high-quality data. And you're accelerating every use case.

Confidence in your data will help lead to positive AI outcomes. But building a scalable and trusted AI program isn't a one time event — it's a living entity within your organization. These four steps form your roadmap:

Step 1: Define the use case

Step 2: Identify and understand data

Step 3: Document models and results

Step 4: Verify and monitor

Unified governance and the benefits of Data Confidence are your goals.

Ready for AI governance?

Learn more about Collibra AI Governance

Collibra AI Governance allows you to define, measure and monitor AI models across your enterprise to maximize results while minimizing risks. Built with the tested platform of Collibra Data Governance, Collibra AI Governance serves as the backbone for your AI strategy.