

AI governance ↪ for the enterprise



Contents

01 →
Introduction

02 →
Challenges
of scaling AI

03 →
All models
need governance

04 →
Holistic AI governance

05 →
watsonx.governance for
responsible, transparent
and explainable AI

06 →
AI governance in action

07 →
Next steps



01

Introduction

Governance makes
AI practical at the
enterprise level

Are your colleagues pushing
to operationalize AI? They're
right to be excited.

The Harvard Business Review reports¹ that
“to call generative AI revolutionary is not
hyperbole. It has the potential to improve
productivity in any function that involves
cognitive tasks.”

Certainly, the promise of AI is undeniable.
And just as surely, the risks of AI are real.
A well-considered approach to governance
gives everyone permission to move ahead.

With governance as your safety net,
there's no reason to hold back from
the revolutionary aspects of AI.

Set your enterprise on a fast path.

Keep reading for the full story
or try [watsonx.governance](https://www.ibm.com/watsonx/governance)
at no charge.



The market size in the generative AI market is expected to show an annual growth rate of 24.40%.²

02

Challenges of scaling AI

The influence of AI is growing exponentially as organizational leaders deploy the technology in nearly every industry.

At the same time, employees and leaders at many of these organizations have difficulty with the following aspects of implementing AI.

It's hard to operationalize AI with confidence

A wide variety of tools exists for AI governance—but too often, models are built without proper clarity, monitoring or cataloging. Without end-to-end AI lifecycle tracking using automated processes, scalability and transparent processes are hindered. Explainable results are elusive.

You may have heard of “black box models,” which are a growing concern for AI stakeholders. AI models are built and deployed, but it isn’t always easy to trace how and why decisions were made, even for the data scientists who created them. These challenges lead to inefficiencies resulting in scope drift, models that are delayed or never placed into production, or that have inconsistent levels of quality and unperceived risks.



Read key takeaways from a poll of global IT senior decision-makers on the pace of AI adoption.

[IBM Global AI Adoption Index 2022 →](#)

It's difficult to manage risk and reputation

You've seen the headlines: unfair, unexplainable or biased AI models, in production. The resulting incorrect assumptions and decisions can affect customers and harm your brand.

Explainable processes and results help auditors and customers know how specific analytic results were reached. Such processes help ensure that results don't reflect bias around race, gender, age or other key factors. These processes are critical for patient diagnoses and treatment plans, reviewing transactions flagged as suspicious, and loan applications that are denied.

Take action to build AI systems that are transparent, explainable, fair and inclusive. You'll help preserve privacy, security, customer loyalty and trust.


AI regulations just keep changing

Successful AI requires adherence to laws and regulations—local, regional and national—which are proliferating at a rapid pace. Noncompliance could cost your organization tens of millions of dollars in fines, as demonstrated by some of the most stringent AI regulations currently debated globally, such as the proposed EU AI Act. The current draft of the EU AI act contemplates fines of up to €30 million, or 6% of a company's global revenue.

Model documentation is crucial—and it's an area with aspects that are easy to miss for a data scientist who's pressed for time and whose organization lacks clear requirements.


Don't disregard this step: new regulations will require model documentation for metadata and lineage.




Timeline of global
AI regulation

80% of business leaders see at least one of these ethical issues as a major concern³


■ Agree ■ Neutral ■ Disagree




Explainability
Believe decisions made by generative AI are not sufficiently explainable.



Ethics
Concerned about the safety and ethical aspects of generative AI.



Bias
Believe that generative AI will propagate established biases.



Trust
Believe generative AI cannot be trusted.

03

All models need governance

AI models are not created equally. But all models must be governed.

Most organizations as of 2023 employ traditional machine learning, and their leaders are beginning to adopt generative AI.

Machine learning models

ML models use predictive analytics to identify trends and patterns in data. They learn from their experience, so that they can improve skills and make more accurate analytic decisions. These models are created from algorithms that are trained using either classified, unclassified or mixed data. ML enables models to learn automatically, without human intervention.

Different machine learning algorithms are suited for different goals, such as classification or prediction modeling, so data scientists use different algorithms as the basis for different models. As data is introduced to a specific algorithm, it's modified to better manage a specific task, and it becomes a machine learning model.



Generative models

These AI models include both foundation models (FMs) and large language models (LLMs). They have the potential to unlock trillions in economic value, because they boost productivity with their remarkable performance, and because they're extensible to a wide range of tasks.

Such models are highly customizable, scalable and cost effective. They can query extremely large volumes of data—and they're learning all the while. “Off the shelf” generative applications require little expertise and have the potential to eliminate many tedious, time-consuming tasks.

In statistics, generative models have been used for years to analyze numerical data. Recently, deep learning has made it possible to extend these models to generate images, music, speech, video, text and even code. Use cases can include marketing, customer service, retail and education.

While generative models have pushed AI high on the agenda for most business leaders, their capabilities drive a new complexity which can pose risks for organizations and for society alike.

■
Learn how to
scale AI responsibly.

[Read the blog →](#)



04

Holistic AI governance

Like any other initiative, successful AI governance depends upon the intersection of people, process and technology.

To implement AI properly, you need a strong cross-functional team. AI is very much a strategic imperative for many leaders, and it can feel like the list of stakeholders grows longer by the day. Some of these people are new to the AI lifecycle concept, and others have new reasons to be involved in AI efforts.

Try to meet the needs of all these groups without overburdening your data scientists, who have little time to route or manage the approvals and requests for information.

Start by putting your stakeholders into alignment. Get buy-in from the proper interested parties and encourage them to participate in ideation, align on outcomes and adopt responsible AI. Then, take steps to ensure that the correct set of metrics, KPIs, and objectives are defined in accordance with your company's business controls and regulations. You'll also want to monitor the specific metrics that have been identified for AI models.




Learn how to build a holistic approach to AI governance

[Read the blog →](#)



Roles across the AI lifecycle



Encourage collaboration with key stakeholders and understand their top concerns:

- CFO, risks to profitability
- CMO, risks to brand
- CRO, risks to enterprise
- CDO, efficient data operations
- CHRO, potential talent impacts
- CEO, organizational accountability
- CPO, regulatory accountability

Process

AI governance traces and documents the origin of data, associated models and metadata, and overall data pipelines for audit. Your documentation should include the techniques that trained each model, the hyperparameters that were used, and the metrics from testing phases. This results in increased transparency visibility by the appropriate stakeholders into the model's behavior throughout the lifecycle, including the data that was influential in its development and the model's possible risks.

You'll first want to benchmark and evaluate your organization's current AI technology and processes. Some processes and stakeholders may already be aligned and can be extended, while others might need to be replaced. Then create a set of automated governance workflows in line with compliance requirements. New and existing AI models can adopt these workflows, which should be designed to avoid the process delays mentioned above. Finally, set up a framework to alert owners and users when a model's metrics exceed the acceptable threshold.

Technology

The establishment of well-planned, well-executed, and well-controlled AI requires specific technological building blocks. Look for a solution that governs the end-to-end AI lifecycle and has the following capabilities:

- Integrates data of many types and sources across diverse deployments
- Is open, flexible and works with your existing tools of choice
- Offers self-service access with privacy controls and a way to track lineage
- Automates model building, deployment, scaling, training and monitoring
- Connects multiple stakeholders through a customizable workflow
- Provides support to build customized workflows for different personas using governance metadata



A framework for responsible, governed AI

	Operationalize with confidence	Manage risk and reputation	Strengthen compliance	Meet stakeholder demands
Plan	Define measurable performance metrics for AI usage across your organization	Review existing processes that monitor fairness and explainability	Conduct gap analysis against current and potential AI regulations	Review existing skills and demand for responsible AI, and align with business objectives
Build	Establish traceability and auditability of current processes	Operationalize updated processes and checkpoints throughout the AI lifecycle	Make sure model documentation is accessible	Specify the new roles, skills and learning agendas required to implement responsible AI
Create	Create automatic documentation of model lineage and metadata	Enable AI models that are fair, explainable and high-quality, minimize drift and conduct regular policy reviews	Act to strengthen regulatory compliance for data science teams without overhead	Establish a repeatable, end-to-end workflow with built-in stakeholder approvals to lower risk and increase scale

watsonx.governance for responsible, transparent and explainable AI.

Meet the toolkit for AI governance. The IBM® watsonx.governance™ approach helps you to direct, manage and monitor your organization's AI activities.

Built with IBM® watsonx™ a portfolio of AI products, this toolkit employs software automation to strengthen your ability to meet regulatory requirements and address ethical concerns. You get comprehensive AI governance without the excessive costs of switching from your current data science platform.

Before a model is put into production, it's validated to assess business risks. Once the model goes live, it's continuously monitored for fairness, quality and drift. Regulators and auditors can get access to documentation that provides explanations of the model's behavior and predictions.

You can offer visibility into how the model works, and which processes and training the model received. Watsonx.governance spans the entire lifecycle, and your teams get help as they design, build, deploy, monitor, and centralize facts for AI explainability.

With this governance toolkit, audits can become easier. Trace and document the origin of data, the models and their associated metadata, and the pipelines.

The documentation will include the techniques that trained each model, the hyperparameters used, and the metrics from testing phases.

Expect increased transparency into each model's behavior throughout its lifecycle, knowledge of the data that was influential in its development, and the ability to determine possible risks.

IBM principles of responsible AI



The purpose of AI is to augment human intelligence




Data and insight belong to their creator



AI systems must be transparent and explainable

Consider these
components:




Regulatory compliance

Manage AI to meet upcoming safety and transparency regulations and policies worldwide—a “nutrition label” for AI.

- Translate external AI regulations into policies for automated enforcement
- Enhance adherence to regulations for audit and compliance
- Use dynamic dashboards for compliance across policies and regulations

Automatic metadata

Data transformation and lineage capture through Python notebooks.




Risk management

Proactively detect and mitigate risks monitoring fairness, bias, drift, and new LLM metrics.

- Automate facts and workflow for compliance to business standards
- Identify, manage, monitor and report on risk and compliance at scale
- Use dynamic dashboards for clear, concise, customizable results
- Enhance collaboration across multiple regions and geographies

Open

Support governance of models build and deployed in third party tools.



Lifecycle governance

Manage, monitor, and govern AI models from IBM, open source communities, and other model providers.

- Monitor, catalog and govern AI models from where they reside
- Automate the capture of model metadata
- Increase prediction accuracy, identifying how AI is used and where it lags

Comprehensive

Govern the end-to-end AI lifecycle.

06

AI governance in action

IBM Chief Privacy Officer ↻

Scaling automation to address AI regulatory requirements

Building on the company's AI framework to address AI regulatory requirements, IBM's Chief Privacy Office (CPO) has taken significant steps in putting into practice AI and data industry-leading capabilities built on a strong combination of privacy, security, AI governance, ethics, processes, technology and tooling.

The IBM CPO, supported by the IBM AI Ethics Board, developed a set of enhanced processes that enable more detailed tracking of compliance with existing standards and applicable legal requirements.

Using IBM's integrated governance framework and process to manage and monitor the development and use of AI across the company so teams can:

- Create a robust workflow using IBM tools to collect, consolidate, display and monitor the workflow
- Automate the capture and integration of facts from the AI lifecycle to accelerate the maintenance of the global AI inventory

[Learn more →](#)



07

Next steps

See how quickly you can create responsible, transparent and explainable AI workflows with the watsonx.governance toolkit—without the costs of switching from your current data science platform.


- Operationalize AI governance
- Manage risk and reputation
- Support regulatory compliance

Get started

[Request a demo →](#)

[Learn more about the governance toolkit →](#)

[Try it yourself at no charge →](#)





1. “How to capitalize on generative AI,” Harvard Business Review, 2023.
2. “Generative AI worldwide,” Statista, 2023.
3. “Generative AI: The state of the market,” IBM Institute for Business Value, 2023.

© Copyright IBM Corporation 2024
IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
December 2024

IBM, the IBM logo, IBM watsonx and IBM watsonx.governance are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on [ibm.com/trademark](https://www.ibm.com/trademark).

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

All client examples cited or described are presented as illustrations of the manner in which some clients have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions. Generally expected results cannot be provided as each client’s results will depend entirely on the client’s systems and services ordered. It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Statement of Good Security Practices: No IT system or product should be considered completely secure, and no single product, service or security measure can be completely effective in preventing improper use or access. IBM does not warrant that any systems, products or services are immune from, or will make your enterprise immune from, the malicious or illegal conduct of any party.

The client is responsible for ensuring compliance with all applicable laws and regulations. IBM does not provide legal advice nor represent or warrant that its services or products will ensure that the client is compliant with any law or regulation.