# Introducing IBM AI Governance

**Driving Trust, Transparency and AI Explainability**

**Priya Krishnan**
Director Product Management
IBM AI Governance

**Doug Stauber**
Principal Product Manager
IBM AI Governance

**Manish Bhide**
Distinguished Engineer & CTO,
AI Governance

# Agenda

Introduction

Industry Changes

Four Pillars of AI Governance

How IBM can help

Demo

Common Use Cases

Q&A

# Industry Changes in 4 key areas:

1. Regulation
2. Reputation
3. Risk
4. Stakeholders

# **Regulation:** AI Policies are accelerating over time

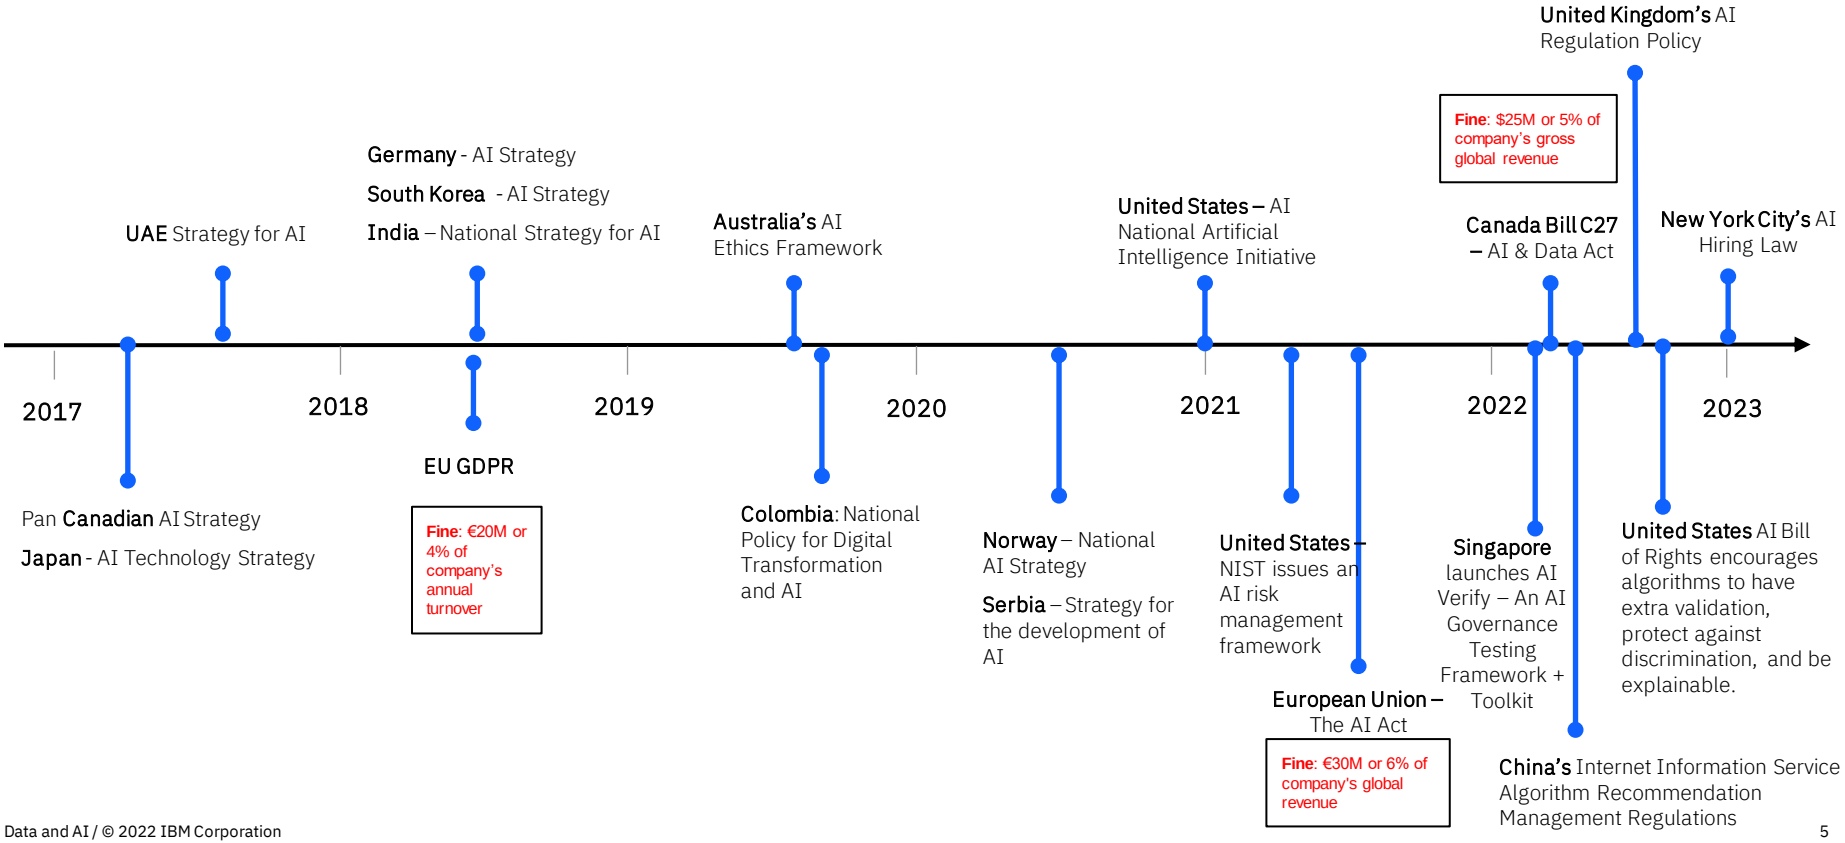

# **Reputation:** poor governance can damage consumer trust

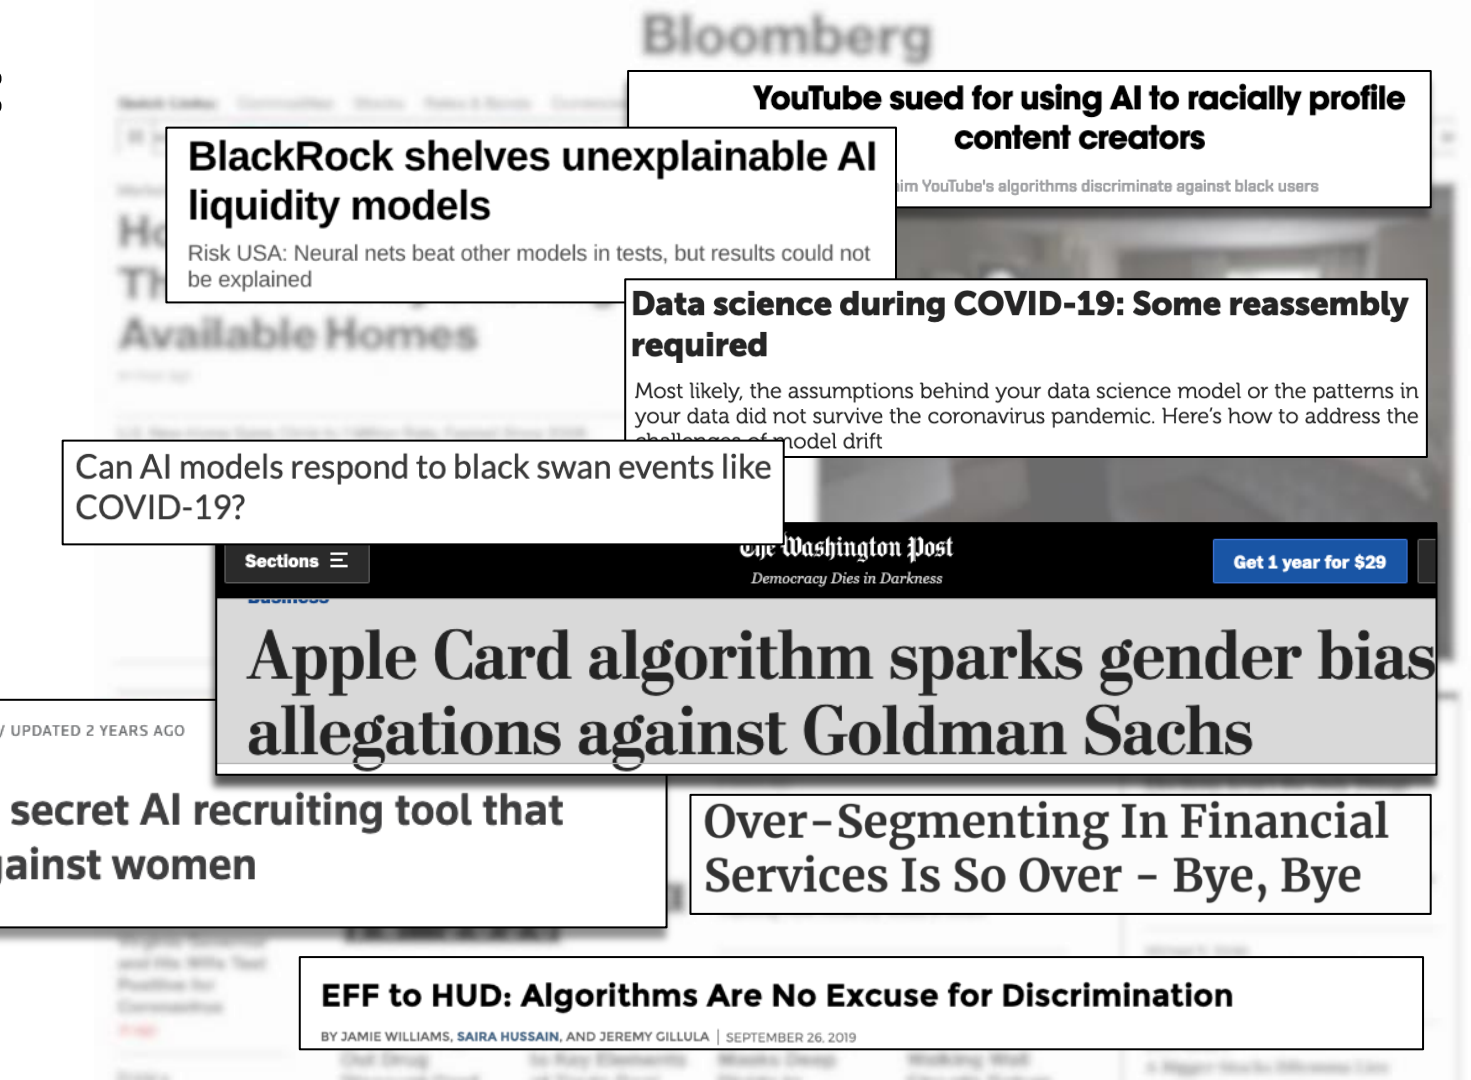

# **Risks** throughout entire AI workflow

Organizations need to mitigate risks of:

- Using personal information when needed and with user's consent
- Ensuring automated decisions are free from bias
- Customer confidence by providing explanations for business decisions
- Fraud to the organization and to customer's accounts
- Delays in putting models into production
- Inefficiency of AI lifecycle stakeholders

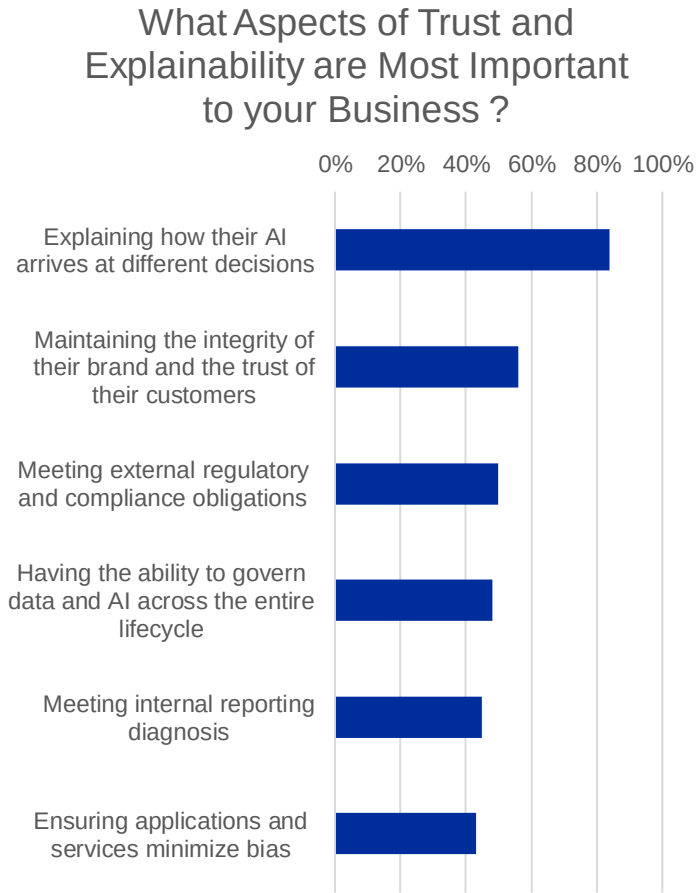


Source: IBM Global AI Adoption Index 2022

# **Stakeholders** increasing beyond traditional players

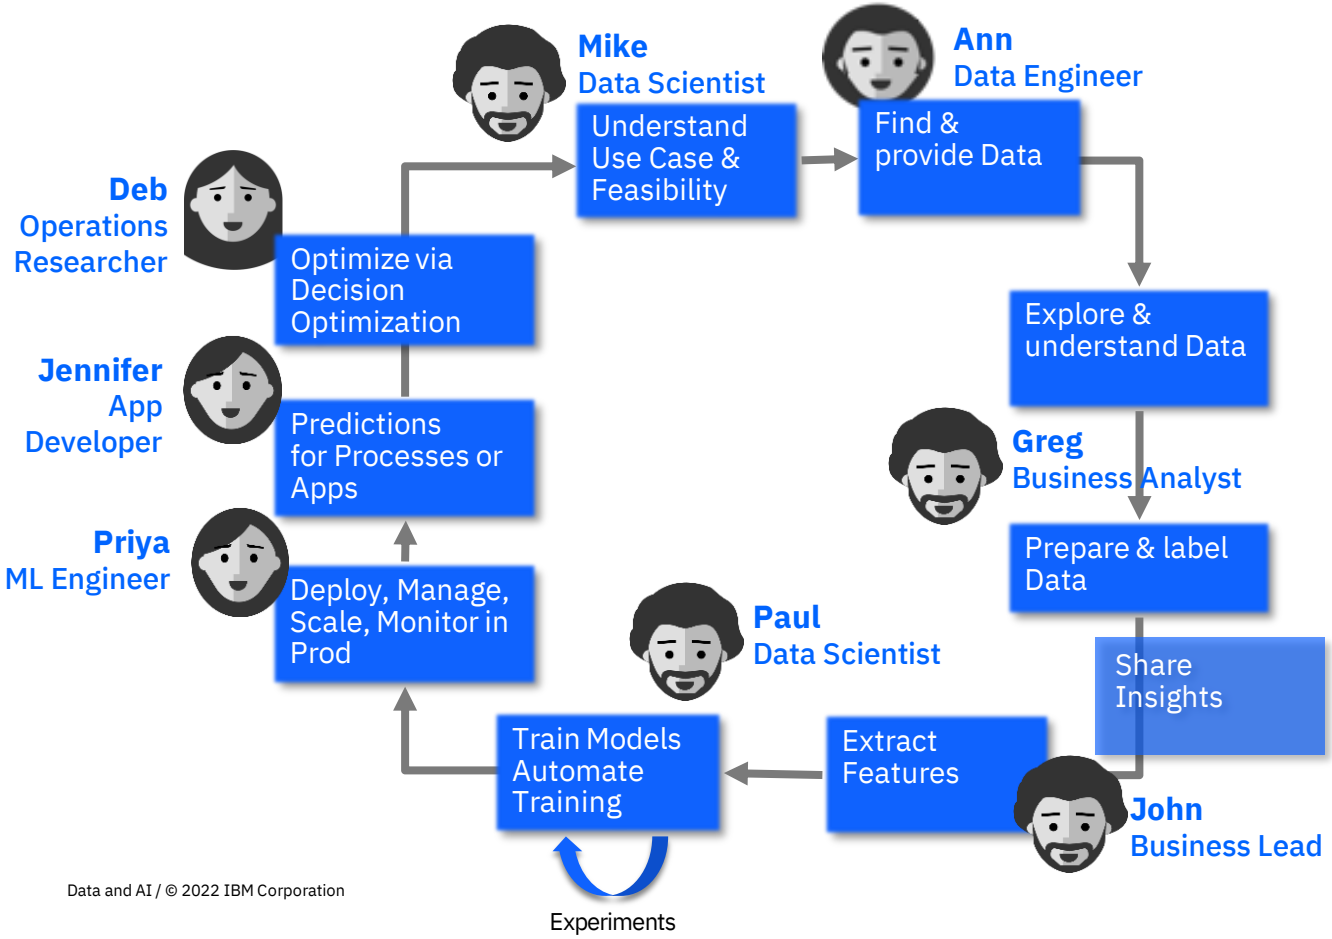


- **CFO** – Risks to profitability
- **CMO-** Risks to brand
- **CRO** – Risks to enterprise
- **CDO** – Efficient Data Operations
- **HR Lead** – Potential Job Impacts
- **CEO** – Organizational Accountability

New!

GRC Tools provide governance

**Maria**
Model Validator / Risk Officer

# Industry changes leading to a clear before and after

| AI Governance | Before | After | IBM AI Governance Solution |
|---|---|---|---|
| **Reputation** | • Accuracy and model performance are rated above all else | • **Responsible AI** to drive equal weight for a broader set of metrics like fairness, drift, explainability, quality, etc. | Lifecycle Governance |
| **Risk** | • Model trust established at the end of AI Lifecycle (if at all) | • Trust established throughout… starting from data collection… to ensure model robustness | Risk Management |
| **Stakeholders** | • Silo'd projects focused on collaboration between Data Science, Business leadership | • **Enterprise-wide** organization required driving C-suite discussions | |
| **Regulations** | • Limited regulations<br>• Data scientists time spent prepping, building, and deploying models | • Newly imposed regulatory requirements<br>• **Regulations require** them to document lineage and metadata | Regulatory Compliance |

# The IBM AI Governance Software Solution

IBM AI Governance, built on IBM Cloud Pak for Data, was designed to direct, manage and monitor the AI activities of an organization. It was designed to meet regulatory requirements, and ethical concerns through software automation.

The solution includes the following capabilities:

**Lifecycle Governance**

- Enable fair, explainable high-quality, drift-free AI models
- Monitor and automatically act upon a broad set of metrics like fairness, drift, quality, etc.
- Enable the businesses to operate and automate AI at scale with transparency and explainability
- Auto-building of fair and accurate models
- Increases accuracy of predictions by identifying how AI is used and where it is lagging.

**Risk Management**

- Establish a repeatable end-to-end workflow with approvals to lower risk and increase scale
- Align the new personas via customized dashboards to organize an enterprise-wide view
- Enhanced collaboration and drives business compliance across multiple regions and geographies.

**Regulatory Compliance**

- Automatic documentation of model lineage and metadata
- Translate external AI regulations into a set of policies for various stakeholders that can be automatically enforced to ensure compliance.
- Ensure regulation compliance for data science teams without overhead
- Manage through dynamic dashboard for up-to-date compliance status across all policies and regulations.

**Capture model metadata automatically**

**Manage AI risk across the organization**

**Monitor models for bias and explainability**

**Exceed AI regulations**

# Let's get ↷ **started** together

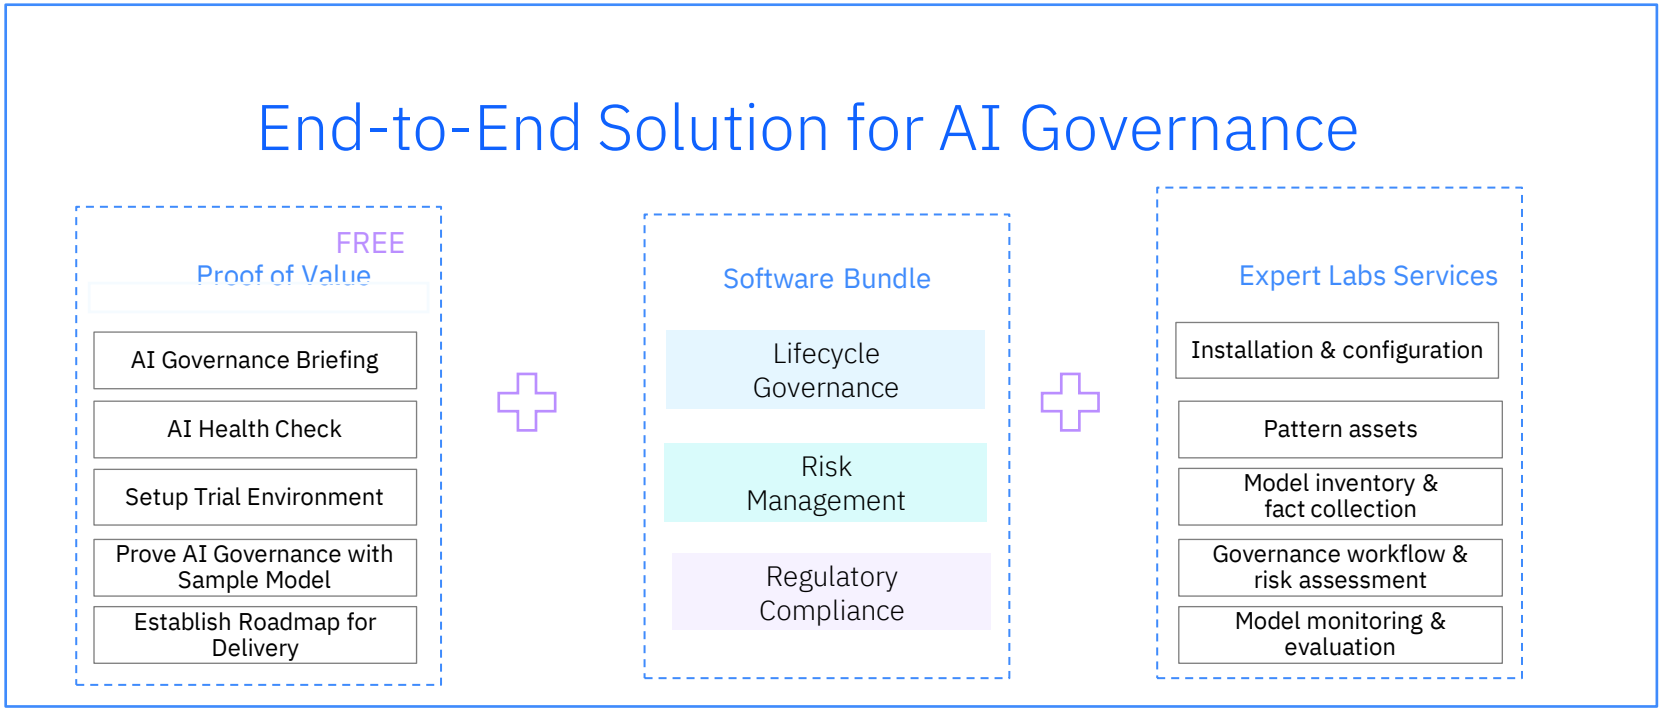

# What you will see in the demo...

**Persona**: Beth - CIO of CM Financial

**What does she need?**

- Understand the state of AI in her organization
- Govern the end-to-end lifecycle of AI models
- Adhere to internal best risk management best practices
- Understand the state of compliance to the new EU AI Ethics Regulations

**Demo**: How does Beth achieve all the above objectives?



IBM

# Demo of capabilities

# **Scenario 1:** Reducing model deployment time

## Business Challenge

Data Scientists build multiple AI models before selecting a model to be sent for validation. Model Risk Management teams need additional information from the data scientist about why a specific model was selected, what parameters were used, etc., before deciding. This back and forth leads to a delay in model approval and deployment to production.

## Solution

IBM AI Governance can accelerate the time to get models into production by adding:

1. Meta data collection of the models developed and deployed
2. Automatic metric data collection on models deployed
3. Providing always-up-to-date dashboards to organize views for all personas in the organization

# **Scenario 2:** Eliminate Bypassing best practices when deploying models

## Business Challenge

Monitoring of models in production is not just limited to monitoring if they are fair, have drift, etc. Models need to be used for the right use cases for which they have been approved. E.g., a model which has been approved for use for clients in CA will require a different sets of checks as compared to a model being used for clients in NY or Europe. How can organizations make sure that models undergo the right tests before being deployed and used in production?

## Solution

IBMs AI Governance can provide a framework to operationalize AI with confidence by:

1. Establishing a customized workflow with the right checks and thresholds to cater to the unique needs for each region
2. Reporting of all the information and results in a standardized and comprehensive governance tool for the risk advisory team to manage.

# **Scenario 3:** Continuous monitoring of models to ensure Regulatory Compliance

## Business Challenge

A big focus area of the new AI Regulations is that of fairness in hiring and promotion decisions. AI models which make these decisions are built without having knowledge of the gender or ethnicity of the person for whom the hiring decision is being made. However, this information can be leaked to the model from correlated features. How can organizations ensure that such models are fair?

## Solution

IBM AI Governance provides state of the art technology to solve such problems by:

1. Ensure the models are fair when they are built
2. Continuously monitor models for Indirect bias where gender/ethnicity information can be leaked to the model due to correlated features
3. Recommend features on which the model is likely to exhibit bias

# Call to Action:

**Learn more:**

Research: IBM Global AI adoption Index 2022

Blog: “AI Governance – break open the black box

Blog: “From principles to actions: building a holistic approach to AI Governance

Try it out: Free trial

Contact us: Talk to an expert: IBM Expert Labs team

# Questions and Answers

# Thank you

# Which of these Common Scenarios most apply to you?

**Scenario 1**: Reduce time to deploy models

**Scenario 2**: Model deployment by bypassin

**Scenario 3**: Continuously monitor models t
regulatory compliance

Poll question:
Which of these scenarios are you most interested in?
- Scenario 1: Reduce time to deploy models into production
- Scenario 2: Model deployment by bypassing best practices
- Scenario 3: Continuously monitor models to ensure regulatory compliance

 

# Strategic Vision for AI Governance –Potential addition

# AI Governance standardizes AI processes and manages risks related to AI/ML deployment

–Wont share slide as it is, but good summary

## AI Governance personas:

**C-suite** Including CEO (organizational accountability), CFO (risks to profitability), CMO (brand damage), and HR Lead (potential job losses).
**Chief Risk Officers & Chief Data Officers** the ability to manage risk pertaining to the AI/ML lifecycle and ensure AI models adhere to AI regulations
**Data Scientists & Model Validators** to automatically capture metadata about models and monitor them for fairness, drift, quality and explainability post deployment

## Key Use Cases:

–New use case

- **Data observability**:
  Before models are built, ensure underlying data is able to be used for modeling and is fair. Monitor data for changes, ensure it continues to flow into workflow process to simplify and strengthen the generation of fair models. (Using AutoAI fairness)
- **Capture model metadata automatically**:
  Monitor and catalog AI/ML models regardless of where they are built. Automatically capture metadata at development time, monitoring time and approval/validation time without consuming data scientist resources
- **Manage AI risk across your organization**:
  Establish an enterprise-wide model dashboard with an end-to-end workflow including approvals to lower risk of incorrect model lifecycle decisions and increase model throughput
- **Monitor models for bias and explainability**:
  Monitor model accuracy, fairness, drift, and explainability to protect brand reputation and practice Responsible AI
- **Exceed AI regulations**:
  Establish a repeatable method to govern the AI process, to exceed external regulations through automatic workflows and real-time model monitoring.

## Differentiators:

–Summarized differentiators

**Comprehensive**
We govern the **end-to-end AI lifecycle** from model development through to model monitoring. Metadata is captured at all stages through on Factsheets, for which we hold a patent

**Open**
We are **AI vendor neural** and support governance of models built and deployed in 3rd party tools such as Amazon Sagemaker, Azure Machine Learning etc
including any GRC stack (visionary).

**Automated**
**Automatic metadata** and data transformation/lineage capture though Python notebooks.

# Backup

*Gartner Forecasts Worldwide Artificial Intelligence Software Market to Reach*

# *$62 Billion*

*in 2022*[1]

---

*Gartner predicts that 30% of IT organizations that fail to adopt AI will no longer be operationally viable by 2022.*[3]

In a world where *trust, transparency and explainable AI* matters, every organization wants the comfort and compliance of understanding how analytic insights and decisions are being made.

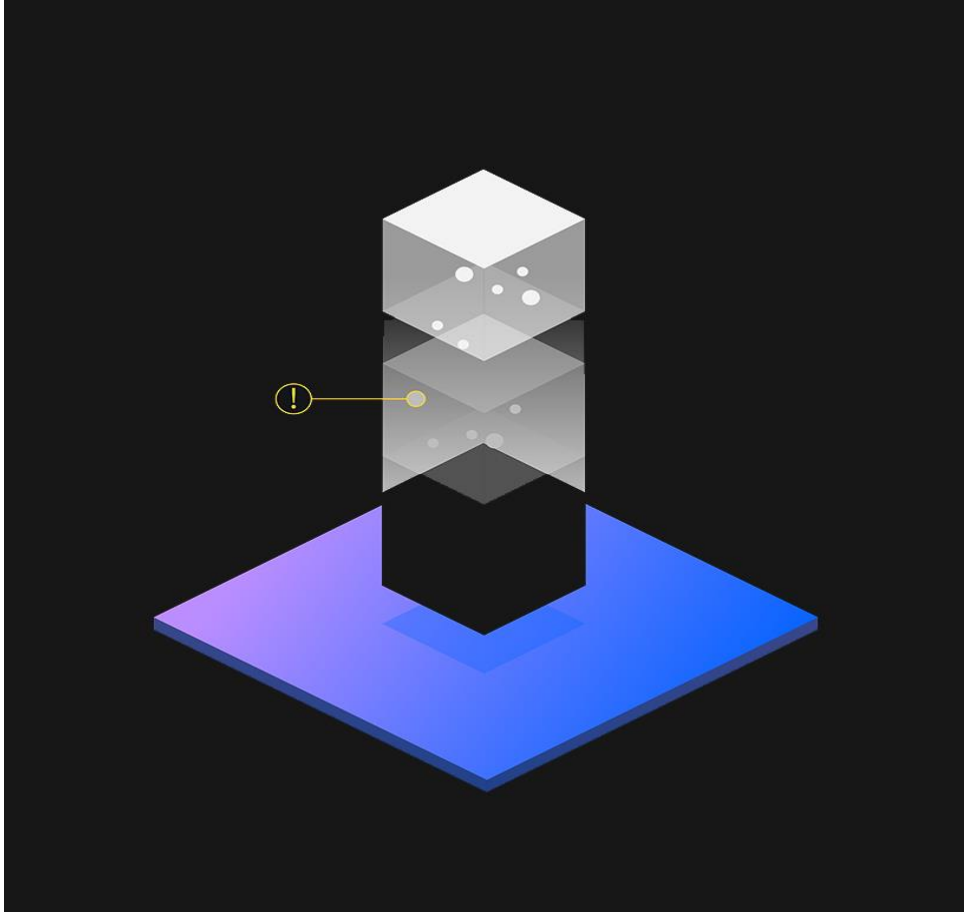

# Most organizations haven’t taken key steps towards governance and trusted AI

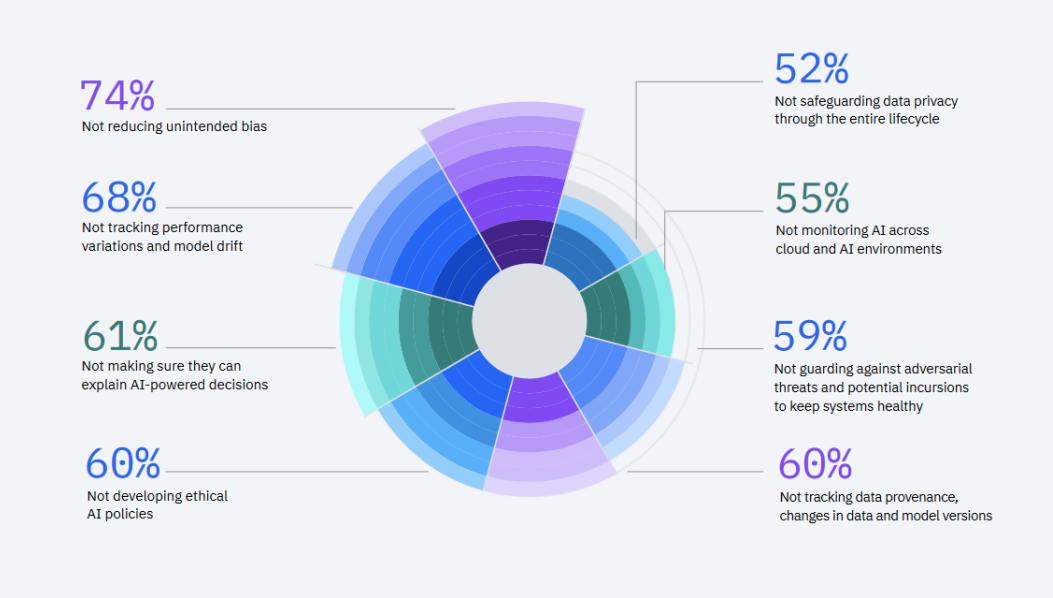


And 60% Lack the AI governance and management tools that don’t work across all data environments

# IBM’s strategic vision for AI Governance

# Industry changes leading to a clear before and after

| AI Governance | Before | After | Solution |
|---|---|---|---|
| **Regulations** | • Limited regulations<br>• Data scientists time spent prepping, building, and | • Newly imposed regulatory requirements<br>• **Regulations require** them to ge and | • Ensure regulation compliance for DS teams without overhead<br>• Automatic documentation of model lineage and metadata |
| **Reputation** | | **AI** to drive r a broader ike fairness, lity, quality, | • Enable fair, explainable high-quality, drift-free AI models<br>• Auto-building of fair models |
| **Risk** | all) | ed starting from data collection… to ensure model robustness | • Establish a repeatable end-to-end **workflow with approvals** to lower risk and **increase scale** |
| **Stakeholders** | • Silo'd projects focused on collaboration between Data Science, Business leadership | • **Enterprise-wide** organization required driving C-suite discussions | • Align the new personas via **customized dashboards** to organize an enterprise-wide view |

Poll question:
Which of these concerns your organization the most?
- Increased regulations
- Brand Reputational Damage
- Risks from not tracking entire AI Lifecycle
- Aligning new stakeholders