An IBM Chief Privacy Office Case Study

How IBM's Chief Privacy Office is building on the company's trustworthy AI framework and scaling automation to address AI regulatory requirements

A long-standing commitment to trust and transparency has been central to IBM's strategy for the use and development of technology systems.

IBM's Chief Privacy Office has taken significant steps in putting into practice industry-leading AI and data capabilities, building on a strong combination of privacy, security, AI governance, ethics, processes, and technology and tooling.

The Challenge

A growing focus on ethics and transparency for artificial intelligence (AI) and algorithmic systems is driving the need for enterprises to capture, integrate, and make transparent metadata throughout the AI system lifecycle, from design to deployment and in everyday use.

The metadata in scope typically includes:

- Purpose and intent at model creation
- Data and model descriptions, and other metrics which may be used in model development
- Bias and disparate impact analysis during model validation
- Performance and bias detection during model deployment and monitoring
- AI ethics metrics and practices used throughout the lifecycle

Working with an inventory of systems and collecting facts determined by regulations, ethics policies, and transparency requirements in an organization the size and maturity of IBM is a challenge of scale and complexity which needs to be met with technical innovation and automation.

IBM's mature and robust method of deploying solutions leveraging AI, and ethics adherence, has resulted in thousands of internal systems being deployed and, as such, monitored for compliance. Each system has a responsible owner, and these thousands of systems each have multiple data points being used to track compliance requirements.

The Response

Using IBM's integrated governance framework and processes to manage and monitor the development, deployment, and use of AI across the company, our 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.
- Accelerate model building at scale. Automate and consolidate multiple tools, applications, and platforms while documenting the origin of datasets, models, associated metadata, and pipelines.

People

Essential roles in our teams responding to the challenge:

The model owner is accountable for the overall compliance, development, delivery, and operation of the solutions within their Business Unit. The model owner should gain assurance, from their model lead, that the solutions for which they are accountable are compliant.

The legal role has due diligence responsibility to evaluate product-level compliance and ethical concerns. Engaged in the development and monitoring of AI tools that are impacted by AI laws and regulations. Can assist by explaining the steps IBM has taken to mitigate ethical concerns.

The data scientist evaluates data through selection, preparation, modeling and deployment, for ethical concerns. Has clear guidelines on ethical standards and best practices. Knows what to measure, what to communicate and to whom – and knows how and why the model was built, what data was used, and what the results mean.

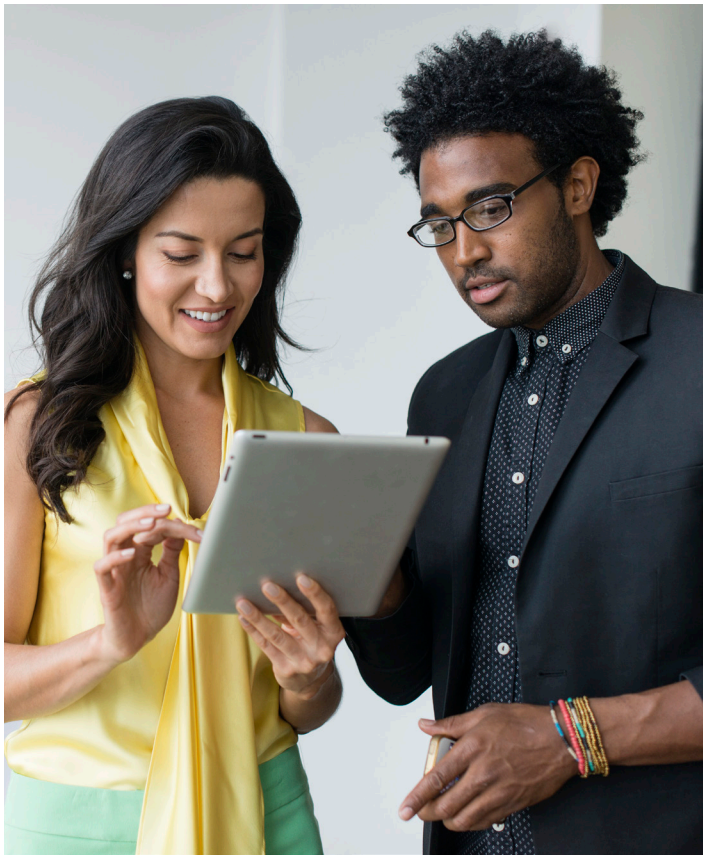

Ethics by Design

These AI processes are underpinned by IBM's Ethics by Design methodology.

Ethics by Design (EbD) is a framework created by IBM to integrate technical ethics solutions into the technology lifecycle and development pipeline, both for AI and other applicable algorithmic systems.

EbD is sponsored by the AI Ethics Board and is embedded in IBM's governance structure through IBM's Tech Ethics by Design Corporate Directive, enabling processes and an EbD Playbook. EbD complements existing IBM security and privacy practices.

EbD and the Playbook support the responsible development of AI and other technologies within IBM and its ecosystem, in a way that is aligned with IBM's multi-dimensional approach to tech ethics.

IBM system owners are trained and required to follow the EbD methodology which includes completing a comprehensive EbD Evidence Questionnaire.

EbD positions IBM to continue its commitment to Trust and Transparency principles, enabling the company to be a trusted partner for clients and grow as a Hybrid Cloud and AI company. Embedding these principles in IBM's products and services enables AI and other technologies as a force for good.

Ethics by Design Playbook

To support the implementation of the Ethics by Design framework, the Ethics by Design Playbook was developed to centralize IBM's recommended methodology for developers and data scientists. This Playbook outlines the tools, methods, and best practices that developers and data scientists should use to integrate tech ethics into their everyday jobs. It includes guidance on how to validate that AI systems are aligned with principles as well as information on metrics and potential mitigation strategies. Some of the tools referenced within the Playbook are AI-powered solutions developed by IBM that help enable compliance with the methodology outlined within the Playbook.

For example, some of these solutions provide functionality to help enable fairness assessments (AI Fairness 360 and Watson OpenScale), transparency documentation (AI FactSheets 360), explainability facilitation (AI Explainability 360), and privacy enablement (AI Privacy 360). When used together these AI-powered solutions can support a holistic approach to the development of responsible AI that ties into IBM's wider AI governance structure.

The Ethics by Design Playbook was developed through a collaborative effort across the company. This included input from the Chief Privacy Office, AI Ethics Project Office, IBM Research, Chief Information Security Office, Chief Data Office, Legal, Government and Regulatory Affairs, Human Resources, Software, and Services.

The Playbook is deployed across IBM through education, and continuous awareness efforts solicit feedback for improvement. Including a diverse set of stakeholders in the development and continuous improvement of playbook content helps to maintain its quality and effectiveness.

The Chief Privacy and Trust Officer and AI Ethics Global Leader, as co-chairs of the IBM AI Ethics Board, were accountable for overseeing and leading the process of developing the Playbook.

Compliance

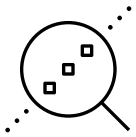

Collect

IBM OpenPages® is used to identify AI Systems (“AI Inventory”) and collect compliance data from AI systems and system owners.

This is integrated with IBM’s Privacy and AI Management System (PIMS).

Facts are collected as required via the IBM Ethics by Design Playbook and AI policies.

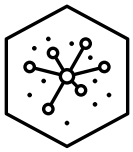

Consoldate

The IBM Knowledge Catalog-based metadata store is used to consolidate compliance facts through the Cognitive Enterprise Data Platform (CEDP) ML-Ops Pipeline.

The cloud-based enterprise metadata repository activates information for AI, machine learning (ML) and deep learning supported by active metadata.

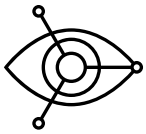

Display

AI Factsheets is the core of the AI Governance solution. It works as an automated process using IBM products to produce Factsheets that document the process of training the AI data model as well as its expected results and use. The completed Factsheets help to build trust in AI services through declarations of conformity. This in turn enables visibility and transparency for internal compliance and audit teams.

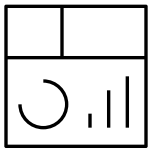

Monitor

Watson OpenScale is used to detect and monitor for bias and drift. It includes the capability to conduct internal bias assessments for disparate impact.

It can help to provide visibility into how the AI is built and used. Its open platform enables businesses to operate and automate AI at scale with transparent, explainable outcomes that are free from harmful bias and drift.

Putting this into practice

The IBM Chief Privacy Office, 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 - for example the NYC Local Law (or Local Law 144).

To simplify the process for owners of technology systems that might be impacted by multiple standards and legal obligations, these potential requirements were bundled into a single workflow. Integrated with existing privacy tooling, additional tracking capabilities were enabled for artificial intelligence and other applicable automated tools or algorithmic systems. Through this capability, it became possible to institute a more detailed review of how a system complies with applicable corporate directives, such as Privacy by Design and Ethics by Design, legal requirements, and industry standards and guidance, such as the Data & Trust Alliance Algorithmic Bias Safeguards for Workforce framework and NIST Risk Management Framework. This helps enable HR tools and processes that use workforce data to monitor transparency, robustness, and explainability, and adherence to privacy principles.

Key benefits of the enhanced process

- Simplified experience for IBMers that integrates related requirements
- Draws from expertise across multiple IBM functional areas
- Addresses IBM's regulatory readiness position through an inventory of technology systems and their operations
- Enables HR systems to provide access to opportunities to IBM employees and candidates

The specific steps taken within IBM HR help demonstrate IBM's continued commitment to building trustworthy technology and addressing regulatory readiness as new requirements are established.

Through IBM's leadership in this area, we are able to help clients and partners in their own journey towards these objectives.

Looking ahead with watsonx.governance

IBM's multidisciplinary, multidimensional approach is helping to advance responsible AI.

The watsonx.governance solution employs software automation to strengthen the ability to address regulatory requirements and ethical concerns. It spans the entire lifecycle from design, to building, deploying, monitoring, and centralizing facts for AI explainability.

Before a model is put into production, it can be assessed for risks to the business. Once the model is live, it can be continuously monitored for fairness, quality, and drift.

Regulators and auditors can be provided with access to the model's documentation, which provides explanations of the model's behavior and predictions. These explanations provide visibility into how the model works and what processes and training the model received.

Audits can become more manageable as does tracing and documenting the origin of data, the models and their associated metadata. Documentation can include the techniques that trained each model, the hyperparameters used, and the metrics from testing phases. The solution can bring about the advantages of increased transparency into the model's behavior throughout the lifecycle, a record of the data that was influential in its development, and the ability to determine possible risks.

Learn more

To learn more and discuss your enterprise's next steps on its AI ethics and automation journey visit:

Learn more about IBM ethics →

Explore the watsonx.governance solution →

IBM Corporation

New Orchard Road, Armonk, NY 10504

Note: all numbers are approximate and subject to change over time.