

Brought to you by:



Informatica™

AI Governance

for
dummies®

A Wiley Brand



Understand
AI governance

Reimagine governance
for the AI Era

See AI governance
in action

Informatica
Special Edition

Audrey O'Shea

About Informatica

Informatica from Salesforce is a leader in AI-powered enterprise cloud data management. Its Intelligent Data Management Cloud (IDMC) platform enables organizations to connect, manage, and unify AI-ready data across the enterprise. With capabilities spanning data cataloging, integration, governance, quality, privacy, metadata management, and master data management, Informatica supports a broad partner ecosystem and helps customers unlock the full value of their data and AI initiatives.

About Salesforce

Salesforce is the #1 AI CRM, empowering companies to connect with their customers in a whole new way through the power of artificial intelligence, data, and trust. For more information about Salesforce (NYSE: CRM), visit www.salesforce.com.



AI Governance

Informatica Special Edition

by Audrey O'Shea

for
dummies[®]
A Wiley Brand

AI Governance For Dummies®, Informatica Special Edition

Published by
John Wiley & Sons, Inc.
111 River St.
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2026 by John Wiley & Sons, Inc., Hoboken, New Jersey. All rights, including for text and data mining, AI training, and similar technologies, are reserved.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Informatica and the Informatica logo are registered trademarks of Informatica. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@dummies.biz, or visit www.dummies.com/custom-solutions. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN 978-1-394-39844-7 (pbk); ISBN 978-1-394-39845-4 (ebk);
ISBN 978-1-394-39846-1 (ePub)

Publisher's Acknowledgments

Project Editor and Manager:

Carrie Burchfield-Leighton

Acquisitions Editor: Traci Martin

Senior Managing Editor: Rev Mengle

Client Account Manager:

Jeremith Coward

Special Help from Informatica:

Vaibhav Suresh, Donal Dunne,

Suzanne Weller,

Melissa Whitaker

Production Editor:

Umeshkumar Rajasekhar

Introduction

If you're not already using artificial intelligence (AI) in your business, you're falling behind the curve. AI is now a core part of business strategy at the enterprise level. A recent survey conducted by Wakefield Research, reported in Informatica's CDO Insights 2025 report, involved 600 data leaders from different regions worldwide, each working at companies generating over 500 million dollars in revenue. The CDO Insights 2025 survey reported that 86 percent of respondents expect to increase their investments in data management, and nearly half of those identify AI readiness as the reason.

The CDO Insights 2025 survey also reported that 92 percent of respondents had concerns about moving projects forward even though known problems from previous projects haven't been resolved. According to the same report, the obstacles that kept AI projects from reaching production included data quality, privacy, and using AI responsibly. The report also showed that because of these concerns, many AI projects fail to make the transition from proof of concept to live systems despite increasing investment in AI projects. While 87 percent of companies that are already using or plan to use AI expect to increase their investment, 25 percent of those expect their AI investments to significantly increase.

As AI applications become ubiquitous across diverse business units and embedded into business processes, effective governance of AI systems is even more important. This is especially true at the enterprise level where data is large scale, and the number of AI use cases in operation across an enterprise is also significant. AI requires processes, tools, and automation to govern it and remain in compliance with emerging regulations concerning its use.

About This Book

This book provides you with a governance framework that's straightforward and simple but powerful in its execution. *AI Governance For Dummies*, Informatica Special Edition, is

your resource for developing an AI governance system. This book contains

- » An explanation of the role of AI governance and the challenges in developing AI and GenAI systems (Chapter 1)
- » A four-step approach to AI governance at the enterprise level that recommends a clear, scalable path from experimentation to production (Chapter 2)
- » An example of the implementation of an AI system following AI governance processes (see Chapter 3)
- » The ten best practices to follow to initiate AI governance (Chapter 4)

Icons Used in This Book

This book displays icons in the margin to alert you to key information. The two icons I use are



TIP

The Tip icon indicates ideas and helpful information that can save you time or avoid frustration.



REMEMBER

Remember icons highlight key information to store in your memory for future use.

Beyond the Book

This book provides an introduction to AI governance, but I can only cover so much in these short pages. For a deeper understanding of AI governance, take advantage of the following resources:

- » www.informatica.com
- » www.informatica.com/blogs/streamline-ai-governance-with-informatica.html
- » www.informatica.com/cdoinsights2025
- » www.informatica.com/dx-datamanagement
- » www.informatica.com/resources/articles/ai-governance-explained.html

IN THIS CHAPTER

- » Making sense of AI governance
- » Looking at why AI governance matters
- » Applying AI in AI governance
- » Identifying AI governance challenges

Chapter 1

Understanding AI Governance

With the advancement and widespread use of artificial intelligence (AI), organizations are challenged with a new kind of governance — AI governance. Simply put, AI governance is your organization's blueprint for using AI responsibly. It sets up the policies and frameworks that help you avoid common AI pitfalls, such as biased results, compliance issues, security vulnerabilities, and privacy violations. These safeguards become absolutely essential when you're dealing with AI and data at scale — because in the world of large-scale AI, small problems can quickly snowball into major headaches.

Governing AI

AI governance considers the organization as a whole with its scope focusing on practices and policies that control the development and use of AI in your organization, manage how those AI systems behave, and figure out their impact. A key tenet of AI governance is to ensure that AI is used in a reliable, responsible, and ethical manner, continuously monitored to operate within defined regulatory boundaries — which mitigates risk and safeguards enterprise reputation.

AI governance is far reaching and affects the design, development, delivery, and scope of AI systems. AI has become a tool that organizations now rely on. Its use can affect the success or failure of your business. From automating routine tasks to providing superior customer experiences, generating content, and analyzing data, AI is quickly becoming deeply embedded in organizations' daily operations, workflows, and business processes. Successful governance of these AI operations helps ensure that they continue to deliver reliable and responsible value, promoting positive business outcomes.

Regulations imposed by government entities are continuously evolving. AI governance frameworks must be dynamic and able to address these new and emerging mandates, while remaining consistent with organizational goals.



REMEMBER

AI governance encompasses the entire AI life cycle from design to development, deployment, implementation, and operation. In short, AI governance

- » Consists of frameworks, policies, and guidelines
- » Defines best practices and processes to ensure completeness, transparency, and accountability
- » Works in harmony with and supports corporate goals
- » Is designed to minimize risks associated with AI
- » Facilitates regulatory compliance
- » Enables the unbiased, accurate, and ethical use of data and AI in all applications through strong governance, ensuring transparency and compliance
- » Protects sensitive data such as personally identifiable information (PII) and sensitive company information, and its use in training and prompting AI systems

AI governance doesn't end when an AI system is deployed. The accuracy and performance of AI systems can deteriorate over time, known as *drift*. AI policies and frameworks ensure that the data sets used by AI models are reliable, transparent, and operate within prescribed boundaries. This helps ensure AI outcomes remain accurate, unbiased, and trustworthy. AI policies and frameworks must be consistent with regulatory compliance, ensuring data privacy and safety while minimizing risk.



TIP

Using intelligent, unified software tools is a necessity when working with large volumes of diverse data, and the explosion in the number of AI use cases in an enterprise. While discrete tools are available from various vendors for supporting different aspects of AI governance, an integrated system like Informatica Intelligent Data Management Cloud (IDMC) eliminates the need for and complexity of managing multiple software stacks.

Understanding Why AI Governance Matters

AI governance provides your organization with a pathway to develop and use AI in a manner that protects your data while fostering innovation and growth. Additionally, AI governance ensures trustworthy outcomes (for example, predictions, decisions, generated content, and recommendations) from the AI system.



TIP

The benefits of AI governance include

- » **Transparency and explainability:** Governance requires that the data source used, how the AI model is used, what data and tools the AI system has access to, and what controls are in place are all transparent and explainable, so there's no mystery surrounding how a particular outcome was derived.
- » **Improved reliability of output:** By putting in place systemic frameworks and documentation, AI governance helps improve the accuracy and trustworthiness of AI models.
- » **Improved AI fluency:** AI governance fosters visibility and a greater understanding of AI for both technical and non-technical users. Documented guidance on the use data and AI allows users to engage with AI and encourages them to understand it better.
- » **Fostering trust among stakeholders:** Trust is fostered by the presence of standardized definitions, transparent processes, and consistent governance frameworks available to various stakeholders across the enterprise.
- » **Promoting secure collaboration:** AI governance includes guidelines for handling and sharing AI and data.

ENSURING PROPER AI GOVERNANCE

The importance of *proper* AI governance can be best illustrated by a real-world example. Failure to ensure that proper AI governance is followed can be costly: In 2023, iTutor Group, a remote tutoring company, paid \$365,000 to settle an unlawful age discrimination suit because its AI-powered recruiting software automatically rejected female applicants ages 55 and older and male applicants ages 60 and older.

- » **Mitigated risk:** AI governance guidelines promote regulatory compliance and provide the means to demonstrate adherence, reducing liabilities and non-compliance risk. Other mitigated risks can include harm to society or individuals, lack of return on investment (ROI), or financial and reputation harm to the organization deploying AI.

Utilizing AI in AI Governance

Using AI to automate AI governance tasks such as discovering, cataloging, tagging data, and monitoring its use is a necessity because of the sheer volume and complexity of information being processed.



TIP

Using AI to improve and automate AI governance has several advantages:

- » Enhances auditing and reporting accuracy
- » Accelerates time to insights and value
- » Provides trustworthy and reliable outcomes
- » Helps modernize data and application safely
- » Delivers faster data discovery and onboarding
- » Supports lower data and AI risk exposure

- » Eases access to trusted customer data for elevating the customer experience
- » Increases employee productivity
- » Improves visibility into data sources
- » Simplifies collaboration



REMEMBER

Look for a software solution that uses AI for AI governance such as Informatica IDMC.

Navigating AI Governance Challenges

Informatica's CDO Insights 2025 report states that 67 percent of respondents say that they're unable to transition even half of their pilots to production. The challenges organizations face in developing systems include



REMEMBER

- » **Integrating data and managing high-volume, high-velocity data for AI:** Data integration takes data from multiple sources and consolidates it into a single view. Using AI to gather and process this data is essential in an enterprise where the volume and velocity of data are high.

When a company is developing and training AI models, it's particularly important to ensure that the data the AI models are trained on is comprehensive and diverse. This helps reduce inherent biases; however, as the number of sources increases, integrating and managing them becomes more of a challenge, so using AI for integrating data is paramount.

- » **Data quality:** Data, including unstructured data, needs to be cleaned and monitored for accuracy, completeness, and consistency across the AI life cycle. Using AI to perform those tasks greatly reduces cost in the form of manual labor and decreases human error.
- » **Data privacy, protection gaps, and regulatory and compliance scrutiny:** Privacy and security are also major concerns of regulatory compliance. Establishing data-use processes and using AI to classify and analyze data and apply policies can help keep your organization in compliance and help fill privacy and protection gaps.

» **Data and AI fluency:** Improving data and AI fluency usually requires investing in AI and data training for your employees. Another avenue to improving fluency is to empower users with tools that enable them to understand data and AI better. Tools that are easy to use and integrated, designed with user needs in mind can help them learn more about how data and AI are used and build trust with it.

As a part of project management, your leaders should anticipate employees' apprehension and culture shock when using new AI or GenAI systems. Explaining to your employees and other concerned parties how new AI systems benefit both the company and employees can foster employee and stakeholder buy-in, garner support, and facilitate an easier transition to using AI in company endeavors. Employee training should also emphasize the responsibility of employees for the output of AI systems.

» **Ensuring responsible AI:** This challenge is multifaceted. AI must deliver accurate output that's ethical, resilient, secure, unbiased, accountable, and transparent. It needs to ensure privacy of personal and other restricted data while providing valid and accurate output. Following the AI governance processes assists in ensuring that AI is developed and deployed in a responsible manner and provides responsible outcomes.

» **Fragmented technology and complexity:** Enterprise data is large, complex, and fast moving. With a broad ecosystem of tools to choose from, organizations could end up with fragmented, disconnected tools to develop and deploy AI systems. Using fragmented systems for the various components of AI governance only adds to the confusion and can result in problems with their interoperability, as well as challenges in applying governance consistently.



TIP

Using a single solution with multiple facets such as Informatica IDMC can alleviate some of the angst and facilitate smooth and cohesive AI design, development, deployment, implementation, and operation. Check out Chapter 3 for a hypothetical situation to illustrate how AI governance works to facilitate AI development.

IN THIS CHAPTER

- » Developing a comprehensive inventory
- » Ensuring an effective system of access control
- » Providing trustworthy, high-quality data
- » Observing how your data is used

Chapter 2

Reimagining Governance for the AI Era

If you're wondering how to manage the data and models used in artificial intelligence (AI) systems, you've come to the right place. In this chapter, you discover the best way to find and manage diverse data, enrich it, protect it, and provide excellent and trustworthy results from AI for all your stakeholders. You follow a reliable four-step approach that supports your AI governance framework. The approach is both simple and comprehensive in its execution. It's designed to address challenges in the emerging field of AI governance (which I cover in Chapter 1) and provide a catalyst for the adoption of reliable AI systems and promote business innovation. The four steps are

- » Inventory
- » Control
- » Deliver
- » Observe



REMEMBER

These steps provide a blueprint to ensure that your AI systems are accurate and trustworthy while keeping data and AI operations safe and secure.

Creating a Comprehensive Inventory of Data and AI Assets

A comprehensive inventory must include all your data and AI assets. Creating a data catalog is a great place to start.

Developing a centralized catalog and automating data discovery

A catalog, which is a centralized, comprehensive inventory of data and AI assets and their associated metadata, is one of your first and possibly most important AI governance steps toward taking advantage of the power that AI offers.



TIP

The advantages of a catalog include the following:

- » Consolidates data and provides a holistic view
- » Provides centralized oversight of all AI use cases, systems, and models
- » Enables easier data discoverability, tracking of data lineage, and enhances understanding
- » Provides easier compliance with regulatory requirements
- » Provides consistency and accuracy of data
- » Improves collaboration by making data sharing easier
- » Reduces manual efforts by streamlining processes and automation
- » Provides a repository of AI models, projects, regulations, policies, processes, and geographies
- » Encourages and enables data-driven decision making

The problem that most companies face is that their data landscape is large and diverse, and scanning it can be difficult, produce multiple records of the same data, and yield data taken out of context. For large organizations, managing data effectively can be a major stumbling block to moving ahead with AI projects. Before you panic, this section aims to help you make sense of your data and create a comprehensive, useful, and trustworthy inventory

for your AI projects. The quality of the data you start with largely determines how trustworthy your AI application's output will be.

Data can exist in many places and forms:

- »» Cloud storage and platforms
- »» On-premises storage
- »» Business intelligence (BI) tools
- »» Databases
- »» Multi-vendor extract, transform, load (ETL) tools
- »» Financial and business reports and spreadsheets
- »» API and IoT feeds
- »» Customer relationship management (CRM) files
- »» Multimedia, images, and PDF files

Identifying all relevant data can be complex and costly when performed manually, so automating the discovery of data, AI models, and systems across your organization is essential. Your AI-powered data discovery solution should be able to locate data wherever it exists, including in the cloud, in local systems, and in mainframes and legacy nodes.

When you're finding the most relevant data, the quality of that data is key to the quality and success of your AI application's output. When seeking relevant data, keep in mind the following:

- »» **Data boundaries:** Data boundaries should be clearly defined.
- »» **Regulatory compliance:** Regulations protect personally identifiable information (PII), intellectual property (IP) through acts like the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and the California Consumer Privacy Act (CCPA).
- »» **Classifying sensitive data:** Data discovery should capture information about the sensitivity of data to help stewards identify possible risk exposures.



TIP

Ensuring data boundaries keeps your system streamlined, efficient, and compliant.

Automating the identification of changes and updating information with AI-powered tools ensures that data is always reliable, accurate, and trustworthy.

Cataloging data and AI assets

After assets are discovered, they should be cataloged by using an automated solution to streamline the process. Without automation, creating a comprehensive catalog can be time-consuming and challenging. Your catalog should include metadata related to your AI assets such as model purpose, performance metrics, training data characteristics, and limitations.

Your cataloging solution should provide relationship discovery and create a knowledge graph mapping the relationships between datasets and AI models. The catalog should include a holistic view that can show the data's life cycle, beginning with its creation, ending with its final destination, and delineating all the transitions in between. AI models have a similar life cycle that shows the original foundation model they were based on and the data used to pre-train or fine-tune them. This life cycle information is known as *lineage*. Ensure that you've mapped data and AI models and applications to their owners to facilitate communication, problem solving, and stewardship of data.



REMEMBER

One solution for cataloging is the Cloud Data Governance and Catalog (CDGC) service of Informatica Intelligent Data Management Cloud (IDMC). Find out more about data catalog benefits at www.informatica.com/resources/articles/what-is-a-data-catalog-benefits-and-use-cases.html.

Organizing structured and unstructured data

Is your data AI-ready? It's simpler to pull data from structured sources such as databases than from unstructured sources like PDF files. Despite advances in AI models such as large language models (LLMs), making unstructured data suitable and available for AI use requires sophisticated preprocessing. This comes with many challenges:

- » Unstructured data can be distributed across multiple sources of different systems and may exist in data silos.

- » Unstructured data doesn't have a schema or metadata to aid in understanding its content and quality.
- » Unstructured data may contain personally identifiable information (PII) or sensitive company information.
- » Properly configuring access and privacy controls.

There's no schema with unstructured data, so identifying whether PII or sensitive data is included is even more difficult than with structured data.

To discover unstructured data for AI use cases or manage sensitive information, the data must first be scanned and categorized. This process helps identify the subject matter, sensitivity level, and context of the data. Doing this manually at scale is nearly impossible, so using AI-driven software to automate and simplify the task is essential.

Categorizing data in this way enables you to discover structured, semi-structured and unstructured data that is

- » Relevant to the AI use case
- » High quality — meaning it's complete, accurate, and up to date
- » Enriched with metadata describing the context so AI systems can better understand it



TIP

Look for a software solution like Informatica IDMC that's designed to scan large volumes of diverse data, parse and classify data objects, and identify their metadata attributes, representing the data by commonly used fields.

Enriching data and AI assets with context and keeping it current

After you organize your data, it's time to enrich it by adding contextual metadata information to improve AI system accuracy and output quality. Raw data is often fragmented, residing in varying sources. Enriching the data enables you to add missing bits of information, deduplicate records, and provide context to the data. For example, a customer record may have only a name, address, and phone number.



REMEMBER

Enriching structured data means adding information like

- » Order history
- » Customer projects
- » Tagging and gap-filling
- » Assumptions and audience profiles
- » Contextual metadata
- » Cultural information

Enriching unstructured data may include

- » Who created it
- » What part of the organization it's from
- » Tagging and classifications applied to the data
- » Its version
- » Whether it's linked to other data
- » How it's typically used

Including assumptions, audience profiles, and cultural information gives AI better direction in how to respond. Context provides deeper meaning to data and AI asset output. Improved context can

- » Help determine risk and risk mitigation
- » Improve AI accuracy and relevance
- » Provide explainability and auditability
- » Reduce hallucinations
- » Enable AI to provide more human-like responses

Enriching data with context is an iterative process. Keeping data and AI assets continuously updated means you'll be able to provide the most accurate, reliable, and trustworthy output from your AI applications.



REMEMBER

High-quality, relevant, current, accurate, and trustworthy data and metadata improve AI outputs. Because the quality of AI results depends directly on the quality of the input data, maintaining data integrity is essential for preventing common errors, such as hallucinations.

Viewing AI projects and operations

Make sure you have a unified view of all AI projects and their operations, including the data they're using. Ensure that you have documented the following for all your AI projects:

- » Use case for AI
- » Purpose of the AI
- » Stakeholders
- » Risk analysis
- » Expected ROI
- » Legal and ethical considerations

Ensuring asset discoverability

A system needs to be in place to ensure that users can easily find and understand data and AI models, including their lineage, context, sensitivity, and how they've changed over time. A centralized master data catalog helps establish this visibility into data and AI assets, including their relationships, regardless of where the data resides. Your software solution needs to maintain an updated inventory so users can quickly locate the latest and most relevant data.



TIP

Because you've done all this work to curate your data, it's a good idea to maintain a library of pre-approved models, data, and products for reuse. Remember to keep it updated, too. And to be truly useful, make sure your cataloging software includes a semantic search feature enabling users to search in natural language.

Controlling Access with Confidence

Controlling access to data, and the AI systems that use it, employs access control lists (ACLs), user roles, and similar tools to ensure that the right people have access to the right data at the right time. An effective system of access control can result in

- » Avoiding privacy and compliance breaches and the costs associated with them

- » Confidence that data is safe and fit for the designated use
- » Enhanced data mining for analysis needs

Proper access control can both avoid expenses incurred when unauthorized people access information and increase profits by readily providing appropriate information for detailed analysis that leads to better business decisions and innovation. Both directly affect your company's bottom line.



TIP

Products like Informatica IDMC provide robust features to automate access controls. It also supports customized AI projects and use case approval workflows and provides a unified view of AI projects and operations.

Automating controls to protect data

When managing the large volumes of data in your enterprise environment, you must leverage the advantages of automated controls to protect data. This includes the ability of a tool to provide access management for data by analyzing, and classifying the data and automatically applying data protection policies based on that analysis.

The ability to use granular controls for protecting specific columns or rows of data in structured sources enables those sources to be utilized where access may otherwise be blocked. This means that users have access to what they need without gaining access to prohibited or sensitive information.

Automatic analysis and classification of new resources and de-identification of data (which can be done by leveraging software features such as redaction, hashing, tokenization, generalization, and substitution measures applied as the data is used) further enhance the ability to provide safe-to-use data while controlling access, keeping data in the hands of those who should have it, and making sure it's out of the hands of those who shouldn't.

Although manual tagging is an option, automatic generation of metadata tags can be used to limit access based on how the data within the resource is classified instead of relying on user roles to provide access to the resource.

Enforcing policies

With the proliferation of endpoints in an enterprise environment, it's important that policies can be centralized and pushed down to those endpoints to ensure they're enforced regardless of where the data is being used. This includes diverse environments like data silos and public facing resources.

Customizing workflows

Workflows are an important pillar of an AI governance strategy from project conception to its decommission. Your AI project is likely to involve different parties with various roles. Your software solution for managing control should include a component that enables you to customize the approval workflow and collaboration between the stakeholders who request approval, and the people tasked with approving or rejecting the project.

AI governance requires multiple workflows to ensure that systems run smoothly and all necessary tasks are completed. Workflows include the following and many more:

- » Approving a project
- » Approving an AI model for use within an organization
- » Approving an AI system for production
- » Maintaining risk assessment and compliance
- » Documenting AI use cases
- » Monitoring and addressing performance drift
- » Managing fairness and bias
- » Decommissioning an AI use case

Assessing risk

As a part of the customized workflows, ensure that you have a way to assess risk and ensure compliance with any regulations involving the data and AI system in question. Customizable workflows enable you to bring together the stakeholders involved in AI governance, for example, legal, ethical, security officer, data protection officer, compliance officer, architecture, and anyone

else involved. Each will be able to access the system of record for the AI system complete with all the context needed to perform risk assessment and compliance checking.

Delivering Trustworthy, High-Quality Diverse Data at Scale

Even though the volume of data belonging to your enterprise can be massive and diverse, you must ensure that you're delivering data that's trustworthy and of the highest quality possible because people, the applications they use, and automated business processes will be relying on it for quality output.

Ensuring reliability and quality of data

The quality of output from AI applications is limited by the quality of the data that goes into it. One aspect of delivering high-quality data is adding context (see the earlier section, "Enriching data and AI assets with context and keeping it current"). Another is creating a data quality score for that information based on its completeness, accuracy, relevance, reliability, security, and freedom from bias. Determining data quality scores manually at scale isn't feasible.

Providing AI-ready data at scale requires automation to assess the data quality, whether there's schema drift, and adherence to prescribed data standards. Assessing unstructured data is a bit more difficult. Points to consider are bias, relevancy, and completeness.

Using tools with AI capabilities that automate data quality tasks like matching, cleansing, and anomaly detection can assist you in improving data quality.

Providing data protection and availability

Data protection in AI systems includes the proactive removal of PII or other sensitive information. It also requires that fine-grained access controls are in place to protect the integrity of the AI system as well as the data to ensure security.

Delivering quality drives responsible use of AI

Data quality scores drive responsible use of AI by assisting in identifying drifts, biases, and inaccuracies that can creep in and then by remediating them. Early identification of data anomalies in data pipelines and fixing those anomalies helps to ensure that AI systems operate on trustworthy data that meets the principles of responsible AI.

Bringing Monitoring and Observability to AI Systems

When an AI system goes live, the work is far from being done. You move from creating an excellent system to vigilantly observing the output from those AI systems, with the goal being continuous improvement. That process has four parts:

- » **Monitoring AI pipelines:** Monitoring is essential throughout the AI life cycle, from the data used to pre-train or fine-tune an AI model, to the data being fed into the AI system during its operational or prediction phase. Detecting anomalies early and correcting issues is imperative to building and running trustworthy AI systems.
- » **Ensuring data and AI observability:** For data or an AI system to be trusted, it must be observable. Using automated tools is essential to identify data quality issues and detecting drift and anomalies as well as solving those anomalies and providing insights into the system.
- » **Delivering continuous improvement:** One of the goals of observability is to deliver continuous improvement to enhance the trustworthiness of data. This requires an iterative process where AI systems performance is analyzed over time with routine testing and auditing, leveraging the power of AI to verify the integrity of information and adding context.

» **Ensuring transparency, accountability, and risk management:** Observing systems also includes providing transparency into where data comes from and how it's manipulated in the process. When problems arise, it's important to be able to view the accountability for those platforms such as what roles are assigned to whom and tracking policies and decisions to ensure that systems are free of bias and are compliant. Observing systems should also include a process for monitoring and mitigating risk.



REMEMBER

AI and an array of automation tools are required to achieve the lofty objectives of observability in an enterprise environment with its vast amounts of distributed and disparate data.

- » Facilitating easier AI governance
- » Identifying key roles in an AI governance project

Chapter 3

Examining AI Governance in Action

Artificial intelligence (AI) governance requires a union of people, processes, and tools that work together to ensure that AI projects result in reliable, trustworthy, safe, and meaningful outcomes, which provide confidence in the development and use of AI. Chapter 2 explains the components of AI governance, but it may be difficult to envision how all the moving parts work in real life. So, follow along in this chapter as a project goes through the initial AI governance process.

Making Governance Easier

AI projects in an enterprise are no longer optional. They're often an integral part of corporate strategy and add significant value to the company whether through enhanced customer experience or saving time and encouraging innovation.

The aim of AI governance is to promote confidence in the development and use of AI. It ensures that the data used is safe, relevant, reliable, trustworthy, enriched, and curated so the result is without bias. AI governance provides oversight of AI models to ensure they're created and used in a compliant and responsible

manner. It monitors the lifecycles of AI systems to ensure they solve use cases and deliver value. When AI governance is done properly, it drives forward your corporate image, innovation, and bottom line.

AI governance features and functions at each of the four AI governance steps (inventory, control, deliver, and observe) work together to achieve those goals. These steps provide a simple framework to streamline AI governance; however, they require execution focus to produce desired results. Minimizing governance isn't advisable because a misstep can result in lawsuits, expenses, and a damaged corporate reputation.

Successful AI governance programs rely on three key factors:

- » **People:** Successful AI governance requires input from different roles throughout your organization. As you develop your AI governance framework, bring together representatives from security, compliance, architecture, ethics, IT, and legal. These stakeholders bring diverse expertise to examine AI projects from multiple perspectives.
- » **Process:** Process refers to the steps and actions that people perform to implement AI governance. They include processes to approve AI use cases, allow the use of specific datasets or models, and perform checks and risk assessments before an AI system is deployed. Processes coordinate people with workflows through the use of powerful tools.
- » **Tools:** In an enterprise environment with large, rapidly changing and diverse data, performing the steps manually in a timely manner is impossible. Careful selection of the tools employed to assist in the process is paramount.



REMEMBER

AI governance builds on the solid foundations of data governance. Companies need to adopt approaches where data and AI governance is integrated. Reexamine the tools available to deliver governance. With the right tools, people can get on board with this paradigm shift, and processes can be streamlined. Tools for AI governance should therefore meet these key requirements:

- » **Provide governance tools at each stage of the AI life cycle.** Governance doesn't happen in one place, and it's not a one-and-done activity. AI life cycles are long, iterative, and dynamic.

- » **Integrate diverse platforms providing a single, comprehensive view.** AI system developers have many tools to choose from, and data is everywhere — from Internet of Things (IoT) devices to mainframes to various online providers. AI governance tooling needs to integrate diverse AI platforms and data wherever they exist to provide insights into a single cohesive view, facilitating overview and decision-making.
- » **Support both compliance and innovation.** Tools need to support governance teams by providing organization-wide oversight while also streamlining governance processes for AI teams. This includes helping them quickly find the data they need and reducing the time it takes to gather governance information and seek approval.
- » **Manage data effectively.** Data and AI are inextricably linked. Any AI governance system must be efficient and effective at managing data and AI models, systems, and agents. The system should provide lineage and relationship information so governance officers can understand and control the data used to train, test, ground, and operate AI systems. By nature, an AI governance system is dynamic, so the data and other assets involved must be monitored and observed.
- » **Encourage collaboration.** Work gets done in a more effective and efficient way when ideas are shared in a collaborative environment. AI systems are accessed from the perspective of many different stakeholders including, for example, compliance, security, legal, and data protection. Orchestration of approval and risk assessment processes is a major factor that slows development.

Playing the Roles

By following the activities of key players, this section illustrates the AI life cycle of a project through the four steps of AI governance: inventory, control, deliver, and observe. (Check out Chapter 2 to find out more about the four steps.) This section shows the responsibilities of people in key roles. For the purposes of this example, Emily and Chris are the AI product manager and AI governance steward, but they each represent a team (or teams) of people.

Introducing the AI product manager

The AI product manager gets the ball rolling when it comes to new AI projects. In this example, Emily, the product manager, wants to develop an AI agent for the sales department that examines customer records and provides suggestions for additional products related to customer purchases, enabling the salespeople to better direct sales efforts in providing solutions to those customers' pain points. Emily's focus is on developing a successful AI project that creates a return on investment (ROI). In Emily's role, she follows the project through its entire life cycle:

- » Define the use case and get approval.
- » Discover data and AI models.
- » Develop and test the AI system.
- » Deploy and monitor the AI system.
- » Iterate the process to drive improvement.
- » Decommission the AI system when it's no longer needed.

Introducing the AI governance steward

The AI governance steward's role is to ensure responsible and compliant use of AI throughout the project life cycle. Chris is responsible for

- » Risk assessment
- » Driving approval
- » Data governance
- » Compliance checking
- » Oversight in testing the model
- » Auditing and documentation

Working together

The teams represented by Emily and Chris work on each step of the process following the AI life cycle while working within the AI governance framework. In this section, you notice that data management is a critical component in every step of the AI life cycle.

Defining the use case and getting approval

The AI life cycle begins with the discovery of a use case where AI can add value — for example, optimizing a workflow or generating content. The use case example in this section is to increase sales by making targeted product recommendations. The life cycle is as follows:

1. **Emily defines the use case, the impact she's trying to achieve, and why AI is a potential solution.**
2. **Emily creates an AI project in the inventory, and triggers an approval workflow for the project.**
3. **Chris makes the first approval, triaging the project and ensuring that all the details about the project are documented.**
4. **Data protection, legal, security, architecture, and ethics officers review the project and give their stamp of approval.**
5. **Emily goes to the data catalog and searches for the data and AI models that are the best fit for her AI project.**
6. **After gathering data and AI models, Emily updates the inventory, registering the data and AI model she's using with her AI project.**

This step initiates a new cycle of the compliance workflow. Data is controlled through an access management system to ensure that sensitive or protected data isn't being exposed throughout the entire project.

Developing and testing the AI system

At this point, Emily and her team develop and test the AI system by performing evaluations to ensure that its output remains accurate and aligned with the use case under a range of conditions. Chris provides oversight to ensure that the evaluation and testing are comprehensive. Again, a workflow is triggered and the system is approved to go into production.

Deploying and monitoring the AI system

In this part of the AI life cycle, the system is put into production. Emily monitors the AI system and makes enhancements

as needed to improve its performance. Chris is alerted whenever changes occur, such as a drop in performance or when new biases or anomalies are detected. Meanwhile, Chris is publishing information to the data marketplace to ensure visibility of the assets Emily created so other teams can find and reuse them.



REMEMBER

For this process to work smoothly, Emily and Chris need to have access to the same tools and shared information.

Iterating the process to drive improvement

The process is iterative. Over time, model upgrades, new data, and controls can be applied to drive up the performance of the system. Going forward, Chris performs auditing and documentation to support governance and regulatory compliance.

Decommissioning the AI system when it's no longer needed

When the project is no longer useful, Emily decommissions it. This step is necessary to ensure that sensitive data is no longer available and to prevent legal risk and risks to the enterprise's reputation.



TIP

For Emily and Chris's work to be aligned, they need to be able to view the same information on a shared dashboard and shared marketplace. By using a solution like Informatica Intelligent Data Management Cloud (IDMC), the product manager and steward would have these capabilities and much more.

IN THIS CHAPTER

- » Establishing a clear vision
- » Using an integrated platform
- » Using AI for AI governance
- » Celebrating successes

Chapter 4

Ten Best Practices for Initiating AI Governance

Implementing artificial intelligence (AI) governance is a multi-faceted process, and following AI governance guidance is paramount to fostering the kind of growth and efficiency that only AI can bring to an organization. If you're wondering where to start, you've come to the right place. To initiate AI governance in your organization, follow the ten best practices in this chapter.

Establish Clear AI Governance Vision, Sponsorship, and Buy-In

Establishing your organization's AI governance vision provides expectations and boundaries for developing your AI governance frameworks. Finding sponsors who champion the project and provide leadership and financial support is key to driving the project forward.

Getting stakeholder buy-in is also crucial when a shift in culture happens like when you establish an AI governance system. To get that buy-in, make sure to engage and train stakeholders so they understand how establishing an AI governance system positively

impacts your organization as a whole and them as individuals. This is integral for buy-in, which positively impacts the success of your venture.

Establish Integrated and Connected Data Management

When choosing data management tools, look for an integrated platform like Informatica Intelligent Data Management Cloud (IDMC). An integrated platform streamlines the process by not relying on niche, single point solutions, enabling you to use a single pane for multiple data and AI governance capabilities. Data integration, data quality, cataloging, governance, and access management can be viewed from a single platform instead of hopping between separate platforms that are often difficult to connect and synchronize.

Maintain Updated Inventories

After you've integrated projects, AI systems, and AI models, you can more easily access a broader array of information for visibility. The inventory must be kept up to date to ensure that output derived from it is accurate and to ensure full visibility. Locate and use tools that scan the data resources automatically to ensure that they're constantly current.

Centralize Your Metadata

Metadata provides a contextual understanding of data that helps AI interpret data more effectively. A centralized collection of metadata helps with the following:

- » Enables more thorough searching of data for your particular purposes
- » Acts as a system of intelligence that can be utilized to establish a common understanding of data, its lineage and operations across the enterprise

- » Identifies discrepancies and missing information, prompting correction of records and avoiding misinterpretation of data

Capitalize on AI Power for AI Governance

At the enterprise level, using the power of AI to automate AI governance tasks such as data discovery, cataloging and tagging data, and identifying data issues is essential. The power of AI improves data quality and enhances security by applying policies to data based on its metadata, limiting access to only approved fields and information.

Improve Your AI Governance Literacy

Understanding the goals and processes involved in AI governance helps employees better understand and support compliance, bias mitigation, and valuable new uses for AI. A team that's highly literate in AI governance can more effectively explain AI governance processes and leverage AI tools for improved AI governance. As AI use and AI governance mature, the need for educating those involved in AI governance continues.

Implement Responsible AI

Use AI tooling and frameworks, and choose solutions that enable you to deploy responsible AI governance tooling and frameworks to ensure responsible AI practices. This results in increased trust and confidence with AI. AI can easily perform repetitive tasks relieving the burden on people of using AI safely.

Adapt Governance to Shifting Markets, Tech and Regulations

Technology changes, markets continuously shift, and new regulations follow closely behind. Foster an environment of continual learning to stay abreast of those changes and modify AI governance

to avoid falling behind the competition — and to ensure that your organization remains in compliance. This is possible when your AI governance infrastructure is flexible, agile, and can align with market, business, and regulation requirements dynamically.

Analyze Success by Tracking ROI and Celebrating Wins

Constantly analyze your successes. Calculating your return on investment (ROI) for AI governance can be measured in faster processing time, shorter time to implement AI projects, and costs saved by implementing measures resulting in improved compliance levels and enhanced accuracy of AI-generated insights. Being able to quantify your ROI and celebrating wins encourages further investment in AI and AI governance initiatives, and promotes stakeholder buy-in.



TIP

To see how other companies have gained value through the use of AI governance, take a look at the companies highlighted here: www.informatica.com/ai-datagovernance.

Iterate: Observe, Learn, and Refine

AI governance isn't a one-and-done. It's an iterative process. Ensure that the tools you choose for AI governance provide visibility and leverage the power of AI to analyze data and AI outputs for bias and other undesirable effects. Then make necessary changes and use AI to determine if the outputs are better. The iterative process continually improves data and AI systems.

USE KPIS TO FIND ROI

Use key performance indicators (KPIs) to track the changes that result in your ROI. A report by Nucleus Research used KPIs to study the impact of Informatica's integration platform. The study revealed a 413 percent ROI, four months payback period, and \$1.25 million dollar cost savings.



Build Trust in the Era of AI

Generative AI sparks excitement — and heightens worries about data integrity, privacy and new risks of AI misuse. Without clear governance, non-transparent and unreliable AI can lead to inconsistent results, eroded trust and jeopardized brand reputation.

Managing data and AI across numerous fragmented sources is complex — and growing. Success demands modern, intelligent and unified AI-driven governance that drives oversight over AI models from creation through deployment.

Informatica's Intelligent Data Management Cloud™ provides unique features supporting AI model governance across the stages of “inventory, control, deliver, and observe.”

- Connect seamlessly across diverse data and AI platforms with rich metadata
- Embed AI governance into a broader data governance framework, not separately
- Harness a centralized, AI-powered metadata system for easier, more robust AI governance
- Safeguard AI assets throughout their lifecycle with granular access management
- Monitor data and AI pipelines continuously for quality and performance

With Informatica, the market leader in governance platforms, you gain the confidence to scale AI responsibly — unlocking the value while upholding privacy, security and compliance.

Ready to turn AI uncertainties into trusted outcomes?

Send us a message at www.informatica.com/contact-us.

**Where data
& AI come to** The LIFE logo features the word "LIFE" in a bold, sans-serif font, where each letter is filled with a colorful, abstract pattern of lines and shapes.

Adopt a clear governance framework

AI applications are abundant across business units and processes, and effective governance of your AI systems is more important than ever. In this book, you discover a governance framework that's straightforward and simple but powerful in its execution. *AI Governance For Dummies*, Informatica Special Edition, is your go-to resource for developing an AI governance system.

Inside...

- The role of AI governance
- The challenges in developing AI and GenAI systems
- A four-step approach to AI governance
- Ten best practices to initiate AI governance



Informatica

Audrey O'Shea is a computer consultant, former educator, and IT author with numerous industry certifications. She has written four technology books leading to IT industry certifications and has served as technical editor on several others. She holds two degrees and a New York State teaching license.

Go to **Dummies.com**[™]
for videos, step-by-step photos,
how-to articles, or to shop!

ISBN: 978-1-394-39844-7

Not For Resale



for
dummies[®]
A Wiley Brand

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.