AI FOR BUSINESS AND NON-TECHNICAL AUDIENCES

AI Governance for Leaders

Level: Practitioner • 2 days (expandable to 3) • Virtual, In-person

Overview

Most organizations are already using AI faster than they are governing it: employees paste sensitive data into chatbots, teams pilot tools no one has reviewed, and vendors embed AI into products the organization already owns. The hard part of AI governance is not writing a policy document; it is building a system of roles, reviews, and controls that manages real risk without becoming the office that says no to everything.

This is a hands-on, practitioner course. It starts with the fundamentals of governance itself: what can go wrong, how to see your organization's actual AI exposure, and how the major frameworks (the NIST AI Risk Management Framework, ISO/IEC 42001, and the EU AI Act) think about risk, then builds up to the operational machinery of policies, roles, and controls. Rather than survey every framework clause, it goes deep on the working parts a leader must actually build. Every module includes a hands-on lab and builds on the one before.

Who Should Attend

- Leaders responsible for establishing or improving AI governance in their organization
- Risk, compliance, legal, and security professionals extending their remit to AI
- Program and technology leaders who must get AI initiatives through review and into production

Attendees new to AI itself should take *AI Fundamentals for Business Leaders* first.

Prerequisites

- Working familiarity with how your organization currently uses or plans to use AI
- Some experience with organizational policy, risk, or compliance work is helpful
- No technical AI background is required; take *AI Fundamentals for Business Leaders* first if AI itself is new to you

What You Will Learn

- Explain why AI needs governance distinct from existing IT and data governance
- Navigate the major frameworks: NIST AI RMF, ISO/IEC 42001, and the EU AI Act, and judge what fits your organization
- Build an AI use inventory and classify use cases by risk
- Draft workable AI policies, including acceptable use
- Design governance roles, an intake process, and a review workflow with clear accountability
- Establish controls, monitoring, and incident response for AI systems, including vendor AI

Course Outline

Day one: the foundations of AI governance

- Why AI Governance, and Why Now

  - What actually goes wrong: data leakage, bias, hallucinated decisions, and regulatory exposure
  - Why existing IT and data governance does not cover it
  - Governance as an enabler: making yes safe instead of making everything no
  - Lab: inventory the AI already in use in your organization and flag what is ungoverned
- The Governance Frameworks
  - The NIST AI RMF and its four functions: govern, map, measure, manage
  - ISO/IEC 42001 and management-system thinking; the EU AI Act and risk tiers
  - Choosing reference points that fit your organization instead of adopting everything
  - Lab: map one of your real AI use cases onto the NIST AI RMF functions
- Seeing and Classifying AI Risk
  - A working taxonomy of AI risk: data, model behavior, misuse, and third parties
  - Risk tiering: matching the weight of review to the weight of the risk
  - Lab: build a risk-tier classification for a portfolio of sample use cases

Day two: making governance operational

- Policies and Acceptable Use
  - What an AI policy must actually cover, and what to leave out
  - Acceptable-use rules employees can actually follow
  - Lab: draft an AI acceptable-use policy for your organization
- Roles, Accountability, and Review
  - Who decides: governance bodies, owners, and escalation paths
  - The intake and review workflow: proportionate, fast, and documented
  - Lab: design an AI intake and review workflow with named roles and decision points
- Controls, Monitoring, and Incidents
  - Controls that match the risk tier, from human review to access limits
  - Monitoring deployed AI and handling incidents when it misbehaves
  - Governing vendor and embedded AI you did not build
  - Lab: create a monitoring and incident-response checklist for one high-risk use case

Extended Version

The three-day version keeps the same gradient and adds depth on the areas organizations struggle with most:

- A deeper working session on the EU AI Act and preparing for regulatory obligations
- Vendor and third-party AI assessment in practice
- Audit readiness: evidence, documentation, and answering hard questions
- A capstone in which teams assemble a complete governance charter for their organization and defend it in review