AI Governance and Risk Management

GRC302

Govern an organization's AI systems on ISO/IEC 42001 and the NIST AI RMF: find and tier them, decide the organization's role, assess risk and impact on people, control models, agents and vendors, map laws onto the controls and show they work.

Modules	19
Lessons	114
Phases	4
Free preview	Course Orientation (5 lessons)

The method

Built on standards, with laws as dated overlays. ISO/IEC 42001 and the NIST AI RMF are the spine, with ISO/IEC 42005 for impact assessment and NIST AI 600-1 for generative AI. The EU AI Act and US state and local laws are worked examples of a method that works for any law, each dated to the day it was checked.

Many organizations, one method. Each lesson works an example from a different organization, a hospital group, a bank, a school district, a staffing company, so the method is seen in many settings rather than one. Practice in every lesson is a record for your own organization, or for a short scenario if you have none.

Practiced in records you keep. Eight templates, from the inventory entry to the audit test, are free downloads with a completed example each. Twelve practice-tool scenarios and two predict blocks per module check your judgment as you go. The course ends on a 40-question exam.

Full syllabus

Phase 1: Foundations

Course Orientation FREE PREVIEW

What AI governance decides and how this course teaches it: the five questions asked of every AI system, the standards the course uses, reading the course for your organization, how practice works, and a first inventory entry.

1. 0.1 What AI Governance Decides
2. 0.2 The Standards This Course Uses
3. 0.3 Reading the Course for Your Organization
4. 0.4 How Practice Works
5. 0.5 A First Inventory Entry

What AI Governance Governs

The three questions AI governance asks of every system: is it an AI system, what role does the organization hold, and who answers for it.

1. 1.1 What Counts as an AI System
2. 1.2 The AI System Lifecycle
3. 1.3 Developer, Provider, Deployer and User
4. 1.4 Deciding Your Role
5. 1.5 Where AI Risk Differs From IT Risk
6. 1.6 Accountability and the Governing Body
7. 1.7 The AI Governance Charter

Inventory and Classification

Building an AI inventory and tiering it: finding AI already in use, shadow AI and AI inside licensed software, use cases against models, tiering criteria and legal overlays, what each tier decides, and a worked inventory for a regional insurer.

1. 2.1 Why the Inventory Comes First
2. 2.2 Finding AI Already in Use
3. 2.3 Shadow AI and AI Inside SaaS
4. 2.4 Use Case Against Model
5. 2.5 Tiering Criteria
6. 2.6 What the Tier Decides
7. 2.7 Building the Inventory

Phase 2: The Frameworks

The NIST AI Risk Management Framework

The NIST AI RMF as a working method: how it is built, the Govern, Map, Measure and Manage functions outcome by outcome, the generative AI profile NIST AI 600-1, and a current and target profile for a software company.

1. 3.1 How the AI RMF Is Built
2. 3.2 Govern
3. 3.3 Map
4. 3.4 Measure
5. 3.5 Manage
6. 3.6 The Generative AI Profile
7. 3.7 Building a Profile

ISO/IEC 42001, the AI Management System

ISO/IEC 42001 clause by clause: context and scope, leadership and AI policy, planning, support and operation, Annex A by control area, evaluation and improvement, and a Statement of Applicability for a hospital group.

1. 4.1 What a Management System Standard Requires
2. 4.2 Context and Scope
3. 4.3 Leadership and AI Policy
4. 4.4 Planning
5. 4.5 Support and Operation
6. 4.6 Annex A by Control Area
7. 4.7 Evaluation and Improvement
8. 4.8 The Statement of Applicability

Phase 3: Risk and Controls

AI Risk and Impact Assessment

Assessing an AI system for risk and impact: harms and affected parties, likelihood from probabilistic outputs, severity and a grid tied to approvers, treatment and residual risk, the ISO/IEC 42005 record, and a full assessment of a city council's benefit fraud model.

1. 5.1 Two Assessments, One Process
2. 5.2 Harms and Affected Parties
3. 5.3 Likelihood When Outputs Are Probabilistic
4. 5.4 Severity and the Grid
5. 5.5 Treatment and Residual Risk
6. 5.6 The Impact Assessment Record
7. 5.7 Assessing One System

Data and Model Governance Across the Lifecycle

Governing AI across its life: the stages and their gates, data provenance and quality, bias testing, model documentation, change control, drift monitoring and retirement, and a motor insurer's pricing model from design to retirement.

1. 6.1 The Life Cycle and Its Gates
2. 6.2 Data Provenance and Quality
3. 6.3 Testing for Bias
4. 6.4 Model Documentation
5. 6.5 Change Control
6. 6.6 Monitoring for Drift
7. 6.7 One Model's Life

Generative AI and AI Agents

Governing generative AI and agents: prompt injection, data leakage, unreliable output, excessive agency, human oversight and an acceptable-use policy, with permission sheets for each assistant, and a law firm assessing a document assistant and an email agent.

1. 7.1 What Is Different About Generative AI and Agents
2. 7.2 Prompt Injection
3. 7.3 Data Leakage
4. 7.4 Unreliable Output
5. 7.5 Excessive Agency
6. 7.6 Oversight and an Acceptable-Use Policy
7. 7.7 A Law Firm Deploys Two Assistants

Third-Party AI

Governing AI from suppliers: the chain behind a product, due-diligence questions and ratings, contract terms, AI switched on inside existing software, ongoing assurance, and a school district assessing a tutoring platform and new learning system features.

1. 8.1 What Is Different When AI Comes From a Supplier
2. 8.2 Due-Diligence Questions
3. 8.3 Contract Terms
4. 8.4 AI Switched On Inside Existing Software
5. 8.5 Ongoing Assurance
6. 8.6 A School District Buys AI

Phase 4: Regulation and Assurance

Regulation as a Method

Turning AI law into obligation lines mapped to controls: reading a law, mapping and gaps, the EU AI Act and US state and local law as dated examples, overlaps and conflicts, and keeping registers current.

1. 9.1 Reading a Law as Obligations
2. 9.2 Mapping Obligations to Controls
3. 9.3 The EU AI Act, a Dated Example
4. 9.4 US State and Local Law
5. 9.5 Overlaps and Conflicts
6. 9.6 Staying Current

Assurance and Reporting

Assurance on an AI management system: the three lines, metrics, internal audit, writing an audit test, ISO/IEC 42001 certification readiness, management review and the board paper, and a finding followed from test to closure.

1. 10.1 What Assurance Is For
2. 10.2 Metrics That Mean Something
3. 10.3 Internal Audit of the Management System
4. 10.4 Writing an Audit Test
5. 10.5 Certification Readiness
6. 10.6 Management Review and the Board
7. 10.7 From Finding to Board Paper

Phase 0: Course Resources

AI Governance Cheatsheets

One sheet per area of AI governance work: the question, the instrument that settles it, how to use it and what it does not prove.

1. Scope, Roles and What Counts as AI
2. Inventory and Tiering
3. The NIST AI RMF
4. ISO/IEC 42001 and the Statement of Applicability
5. Risk and Impact Assessment
6. The Data and Model Life Cycle
7. Generative AI and Agents
8. Third-Party AI
9. Regulation and Assurance
10. Lookup Values

AI Governance Cookbooks

Nine runbooks for the recurring work of AI governance, from the inventory to the board paper, and the eight templates they fill.

1. Building the Inventory and Tiering
2. Running a Risk and Impact Assessment
3. Writing the Statement of Applicability
4. Gating a Model Release
5. Approving a Generative AI or Agent Deployment
6. Assessing a Vendor's AI
7. Mapping a Law to the Controls
8. Running the Internal Audit
9. Preparing Management Review and the Board Paper

Lab Setup

Build a governance workspace from the course templates, wire the registers together, seed it with a practice system and verify it.

1. The Governance Workspace
2. Wiring the Registers Together
3. A Seed System to Practice On
4. Verify the Workspace

Walkthroughs

Four governance cases worked end to end across the modules, including the wrong turns, each ending in a decision, a report card and the point that transfers.

1. A New System, From First Question to Release
2. A Complaint, to a Changed Gate
3. A Vendor Switches On a Feature
4. A Law Changes, and the Register Follows

Playbooks

Eight playbooks, one per thing that fires in AI governance, each a procedure to execute from the moment it lands to a recorded disposition.

1. Unregistered AI Found in Use
2. An AI Incident or Harm Is Reported
3. A Vendor Changes Its Model or Terms
4. A Monitor Crosses Its Threshold
5. A Release Asks to Skip the Gate
6. A Regulator or Customer Asks for Evidence
7. A Law or Ruling Changes an Obligation
8. Someone Asks You to Soften a Finding or Rating

Practice Playground

Every practice-tool scenario and predict block in the course, by module, and the exam.

1. Tools, Predict Blocks and the Exam

References and Reading

The sources the course opened, by area, with what each was used for, and the ones that could not be opened.

1. Standards and Frameworks
2. Law and Regulation
3. Security, Models and Assurance

Course Completion

Course Exam

AI Governance and Risk Management end-of-course exam: 40 questions in four parts, from scope and inventory to assurance and the board, ending on a case study.

1. Course Completion. AI Governance and Risk Management

Generated 2026-11-05 from the published course. The current syllabus is always at ridgelinecyber.com/training/courses/ai-governance-risk-management/