

Auditing Artificial Intelligence (AI): A Hands-On Course for Internal Auditors

OEM: Institute of Internal Auditors

COURSE DESCRIPTION

In this hands-on course, participants will explore the core concepts, risks, and governance implications of Artificial Intelligence (AI) and Generative AI—focusing on their impact on audit scope, risk assessment, and controls. Designed for internal and external auditors, the course covers how to assess AI initiatives using recognized frameworks and implement governance policies aligned with trust principles, regulations, and lifecycle oversight. Learners will apply the full AI audit lifecycle to a real-world case, evaluate governance maturity, and produce audit-ready documentation in line with ISO/IEC 42001 and the NIST AI Risk Management Framework.

AUDIENCE PROFILE

- This course is intended for anyone in an Internal Auditing role or Compliance role.

PREREQUISITES

- A fundamental understanding of AI
- Familiarity with the IIA AI Framework
- Review the NIST AI RMF
- Completion of the basic level course: Essentials for AI Auditing

COURSE OBJECTIVES

By the end of this course, participants will be able to:

- Define key AI concepts including Artificial Intelligence, Machine Learning, Generative AI, LLMs, and Small Language Models (SLMs).
- Explain how LLMs generate responses and identify the risks of hallucination, inconsistency, and lack of explainability.
- Differentiate between traditional automation, ML models, GenAI, and SLMs in terms of audit risk and control requirements.
- Evaluate core risks that apply across all AI systems, including data bias, model drift, and overfitting.
- Describe the auditor's responsibilities in reviewing early-stage GenAI adoption, with attention to role separation between internal and external audit functions.
- Compare major AI auditing and governance frameworks used by internal and external auditors (NIST AI RMF, ISO/IEC 42001, EU AI Act, GAIA, etc.).
- Apply risk and governance concepts in a hands-on LLM prompt lab and generate auditor-aligned reflections.
- Identify key points in the AI lifecycle where internal and external auditors should engage

- Differentiate responsibilities between audit, IT, and data science teams regarding AI risk
- Recognize the audit implications of model degradation, drift, and bias ownership
- Apply prompt engineering as a method to surface audit-relevant model behaviors
- Interpret how prompt-based audit observations support assurance, documentation, and escalation

COURSE OUTLINE

Module 1: Course Outline

Learning Path 1: Understanding AI Systems And Establishing Audit Scope

- Module 1: Exploring AI and Generative AI Services
 - Intro and Objectives
 - AI Fundamentals
 - GenAI and Language Models
 - Risk Awareness
 - Frameworks and Governance
 - Internal versus External Role Comparison
 - Optional Hands-On Lab
 - Summary and Takeaways
- Module 2: How to Audit the Intricate Components of AI Applications
 - What to Audit in AI Systems
 - How Audit Checklist Items Map to Frameworks
 - Metrics for Evaluating GenAI Outputs
 - Mapping Metrics to NIST AI RMF Functions
 - GenAI Output Review: Internal vs External Auditor Roles
 - Introduction to Auditing Tools
 - Auditing Tools by Role
 - Optional Lab: Microsoft Purview
 - Optional Lab: Aequis
 - Classifying Data and AI Models
 - Real-World Risk Case: \$1 Car Chatbot
 - Module Summary and Takeaways
- Module 3: Investigating Internal AI Usage – Governance
 - Introduction to Internal AI Governance
 - AI Activity Logging and Monitoring Practices
 - Using Microsoft Purview to Audit AI Usage
 - Policy Adherence and Risk Signal Evaluation
 - Internal vs. External Auditor Responsibilities
 - Hands-On Lab: Reviewing Copilot Activity and Prompt Trails
 - Summary and Key Takeaways
 - Knowledge Check

Learning Path 2: Structuring Risk-Base AI Engagements

- Module 4: Redefining Audit Engagement Across the AI Lifecycle
 - Framing the Auditor's Role in AI Governance
 - Auditor Engagement Across the AI Lifecycle
 - Who Owns AI Risk? Role Differentiation Matrix
 - Understanding Model Degradation, Drift, and Accountability
 - Prompt Engineering as an Audit Tool
 - Prompt Audit Patterns: Red Flag Prompts for Risk Discovery
 - Hands-On Lab (Optional): Conducting a Prompt-Based AI Audit
 - Reflection and Role Exercise: Who Should Respond to This Risk?

Learning Path 3: Executing Fieldwork Across the AI Lifecycle

- Module 5: Execute AI Project Management Efficiently
 - AI Project Governance: Scope and Oversight
 - Auditing the AI Vision, Strategy, and Roadmap
 - Evaluating Project Roles and Cross-Functional Accountability
 - Auditing Risk Registers, Use Case Alignment, and Business Impact
 - Hands-On Lab: Reviewing AI Project Governance Templates
 - Summary and Key Takeaways
 - Knowledge Check
- Module 6: Monitoring AI Systems and Governance in Action
 - Monitoring What Matters
 - What Should Be Audited Post-Deployment
 - Types of Audit Evidence: Logs, Outputs, Labeling, Risk
 - Frameworks in Action
 - Manual Audit Techniques: No-Tools? No Problem
 - Optional Lab: Investigating Copilot and Purview Logs (Demo)
 - Summary and Maturity Takeaways
 - Knowledge Check
 - AI Governance, Monitoring, and Capstone Execution

Learning Path 4: Assessing Maturity, Governance, and Strategic Closure

- Module 7: Designing Governance in New AI and GenAI Applications
 - AI Governance Principles and Standards Overview
 - Translating Governance Principles into Policy
 - Domain-Based Governance Structures
 - Auditing Governance Implementation
 - Governance Gaps and Red Flags
 - Knowledge Check and Reflection Questions
 - Optional Lab: Assigning AI and Data Governance Roles
 - Summary and Key Takeaways
- Module 8: Auditing AI Improvement Cycles and Profile-Driven Risk Tailoring
 - Understanding AI Improvement Obligations.
 - NIST AI RMF Profiles and Audit Customization.
 - Auditing the AI Feedback Loop: Are Controls Evolving?

- Internal versus External Auditor Roles in the Improvement Lifecycle.
- Evaluating Evidence of Corrective and Preventive Actions.
- Knowledge Checks.
- Lab: Auditing Evidence of AI Governance Improvement.
- Summary and Key Takeaways
- Module 9: Administering Trust and Accountability in Emerging AI Platforms
 - Introduction to Trust and AI Platform Governance
 - Auditing Platform-Level Trustworthiness Characteristics
 - Provisioning and Onboarding AI Services Securely
 - Controls for Post-Deployment Behavior and Drift
 - Auditing Multimodal AI and GenAI Capabilities
 - Lab: Privacy Trust Assessment
 - Summary and Key Takeaways
- Module 10: Finalizing the AI Audit – Synthesis, Reporting, and Strategic Readiness
 - Reviewing Multi-Domain AI Audit Findings
 - Mapping Risks to ISO Clauses and NIST Functions
 - Evaluating Governance Maturity and Improvement Signals
 - Final Audit Judgment: Certification, Readiness, or Escalation
 - Internal versus External Auditor Roles in Final Reporting
 - Knowledge Check: Risk Readiness versus Risk Documentation
 - Capstone Simulation
 - Summary and Key Takeaways

Capstone Final Event: Business Audit Simulation

- Capstone: Auditing a National AI Program – The Australian Taxation Office Case
- Learners will submit one of the following, based on their assigned role:
- Internal Auditor Role:
 - A completed AI Audit Closure Memo, including:
 - Summary of findings
 - Residual risk analysis
 - Clause 10.2 alignment
 - Closure determination or monitoring plan
- External Auditor Role:
 - A completed Readiness Opinion Letter, including:
 - Scope of review
 - Key observations
 - Maturity and risk assessment
 - Certification readiness opinion
 - Recommendations for improvement