



Generative AI and Agentic Analytics: Governance Guardrails for Data-Driven Leaders



Generative AI and Agentic Analytics: Governance Guardrails for Data-Driven Leaders

Technical depth: Conceptual · **Practical mode:** Case study

Introduction

Analytics teams are moving from dashboards to AI assistants and agents that write queries, summarise results and trigger actions on their own. Leaders are asked to fund and approve these tools before anyone can show that their figures are traceable, their data access is controlled or their actions stay within risk appetite. This Core Concept course gives executives the judgement to direct generative and agentic analytics against recognised AI governance standards and security guidance. Participants examine cases from several sectors and leave with an Agentic Analytics Governance Charter for their organisation.

Course Objectives

- Evaluate where generative and agentic analytics create value and where autonomy should be limited
- Judge proposed architectures for conversational and agent-based analytics, including semantic layers and retrieval, at decision level
- Direct an AI management system and risk approach for analytics use cases in line with ISO/IEC 42001 and the NIST AI RMF
- Set guardrails, approval points and accountabilities that keep agent actions within risk appetite
- Oversee assurance of AI-generated analysis through traceability testing, monitoring indicators and the Three Lines Model
- Approve an Agentic Analytics Governance Charter for the organisation

Target Audience

- Chief data, analytics and digital officers setting the direction for AI-enabled analytics
- Chief financial and operating officers whose functions will rely on AI-generated analysis
- Heads of risk, compliance and internal audit responsible for oversight of AI use
- Chief information and information security officers accountable for platforms and access control
- Directors of strategy and performance who commission management reporting
- Board committee members overseeing technology, data and risk

Course Outline

Day 1: The Generative and Agentic Analytics Landscape

- Evolution from Dashboards to Conversational and Agentic Analytics
- Large Language Model Capabilities and Limits in Analytical Work
- Agentic Analytics Value Cases in Finance, Operations and Public Services
- Autonomy Levels Scale: Assistant, Copilot, Supervised Agent and Autonomous Agent
- Analytics and AI Readiness Assessment Across Data, Skills and Controls

Day 2: Architecture and Governance Frameworks

- Reference Architecture: Semantic Layer, Retrieval-Augmented Generation and Tool Access
- ISO/IEC 42001 AI Management System Requirements Applied to Analytics
- NIST AI RMF Functions and the NIST AI 600-1 Generative AI Profile
- ISO/IEC 38507 Governance Implications of AI for Governing Bodies
- ISO/IEC 5259-5 Data Quality Governance for Analytics and Machine Learning

Day 3: Directing Agentic Analytics in Practice

- Use Case Intake and Prioritisation Scorecard
- Certified Metrics and Semantic Layer Ownership Model
- Guardrail Principles: Row-Level Security, Query Validation and Read-Only Tool Permissions
- Human-in-the-Loop Approval Points for Agent Actions
- AI System Impact Assessment Using ISO/IEC 42005

Day 4: Risks, Failure Modes and Assurance

- OWASP Top 10 for LLM Applications: Prompt Injection and Sensitive Information Disclosure
- OWASP Top 10 for Agentic Applications: Goal Hijack, Tool Misuse and Privilege Abuse
- Hallucinated Figures and Answer Traceability Testing
- AI Risk Register Aligned to ISO/IEC 23894
- Three Lines Model Assurance and Agent Monitoring Indicators

Day 5: Case Work and the Agentic Analytics Governance Charter

- Banking Finance Function Case: Conversational Reporting Rollout Review
 - Public Health Agency Case: Agent Access to Sensitive Data Decision
 - Guardrail and Autonomy Decision Matrix Build
 - Agentic Analytics Governance Charter Drafting
 - Executive Panel Challenge and Charter Defence
-

Skills You Will Gain

- Generative AI Literacy
- Agentic Risk Oversight
- AI Management Systems
- Decision Rights Allocation
- Semantic Data Governance
- AI Impact Assessment
- Assurance Planning
- Technology Investment Judgement

Why Attend This Course

- Return to work with an Agentic Analytics Governance Charter ready for executive committee or board approval
- Ask vendors and internal teams the questions that reveal whether an AI analytics product is safe to scale
- Distinguish a pilot that looks impressive from one whose figures can be traced and trusted
- Test decisions against peers from banking, government, energy and other sectors facing the same choices

Conclusion

Generative and agentic analytics can shorten the path from question to decision, but only if leaders decide in advance what agents may read, calculate and do, and how their answers will be checked. This course moves from the landscape and architecture, through the governance standards and guardrail choices, to the security failures and assurance that expose weak deployments. The final day turns that material into an Agentic Analytics Governance Charter that participants take back to their executive team as the basis for scaling AI in analytics.