



Responsible and Generative AI in Government: Deploying Trustworthy AI in Public Services

2 – 6 November 2026

London - Central London four-star



Responsible and Generative AI in Government: Deploying Trustworthy AI in Public Services

Ref: 1145_5963 **Date:** 2 – 6 November 2026 **Location:** London - Central London four-star **Fees:** 23000 SAR

Introduction

Deploying responsible generative AI in government requires balancing administrative efficiency with strict public duties of equity, transparency, and accountability. Public sector organisations face growing operational demands to summarize policy documentation, accelerate citizen correspondence, and automate routine departmental workflows. However, poorly governed language models risk generating inaccurate outputs, exposing confidential public records, and producing administrative determinations that cannot be explained or contested. Addressing these institutional risks demands robust governance frameworks, technical guardrails, and systematic impact evaluations aligned with international standards. This five-day practitioner course from Core Concept enables public sector managers to design, evaluate, and oversee generative artificial intelligence initiatives that protect public trust.

Course Objectives

- Assess public value generation across public services, policy drafting, and departmental communication.
- Classify and screen public sector generative model deployments against international trustworthy artificial intelligence frameworks.
- Design resilient generative architecture, including retrieval-augmented generation systems and citizen-facing assistants, incorporating human oversight and accessibility standards.
- Execute impact assessments and produce structured algorithmic transparency records for public disclosure.
- Establish technical security guardrails and procurement specifications to prevent prompt injection and data leakage.
- Assemble a complete deployment plan ready for institutional leadership review and operational validation.

Target Audience

- Service delivery and citizen experience directors managing public interaction channels
 - Digital transformation and innovation leaders directing artificial intelligence initiatives in government agencies
 - Policy analysts and legal counsel drafting official publications and reviewing regulatory instruments
 - Data governance and information security managers safeguarding public sector data assets
 - Procurement and commercial officers contracting enterprise model providers and cloud infrastructure
-

Course Outline

Day 1: Generative AI and Public Value

- Generative artificial intelligence operational use cases across policy formulation, administrative operations, and public interaction
- Technical capabilities and boundaries: large language models, retrieval-augmented generation architectures, and autonomous agents
- Public value evaluation methodologies for measuring efficiency gains and citizen service quality
- Frameworks for trustworthy artificial intelligence governance in public administration: foundational enablers, institutional guardrails, and stakeholder engagement
- Institutional readiness diagnosis across institutional data maturity, technical infrastructure, workforce skills, and internal policy

Day 2: Principles, Standards and Classification

- International ethical recommendations and principles governing trustworthy artificial intelligence systems
- Taxonomy and risk classification models for public administration automated systems
- ISO/IEC 42001 management system controls for public administration
- Generative risk profile categories based on NIST AI 600-1
- Government data classification protocols, cloud hosting models, and sovereign processing architectures

Day 3: Designing Responsible Generative AI Services

- Screening matrices and qualification criteria for government generative use cases
- Retrieval-augmented generation design patterns for administrative knowledge bases and legislative repositories
- Citizen-facing conversational agent workflows with mandatory human escalation pathways and WCAG 2.2 accessibility compliance
- Prompt engineering protocols and output verification standards for official government communication
- Procurement criteria and vendor evaluation standards for foundational model providers and software integrations

Day 4: Impact, Transparency and Security Risks

- Artificial intelligence impact assessment workflows for public institutions aligned with ISO/IEC 42005
 - Algorithmic transparency records and public accountability reporting templates
 - Evaluation frameworks for multilingual linguistic precision, factual grounding, and mitigation of demographic bias
 - Vulnerability countermeasures addressing prompt injection and confidential data leakage based on the OWASP Top 10 for LLM Applications
 - Procedural contestability mechanisms and human-in-the-loop controls for assisted administrative findings
-

Day 5: Case Work and the Deployment Plan

- Public assistance eligibility simulation: evaluating automated intake and procedural fairness
- Municipal citizen support centre simulation: deploying a knowledge-grounded assistant
- Drafting the institutional deployment plan for an agency operational service
- Red-team stress testing and scenario analysis for adversarial prompts and systemic failures
- Executive defense and review of the completed deployment plan

Skills You Will Gain

- Public value analysis
- Trustworthy artificial intelligence governance
- Use case risk screening
- Retrieval-augmented service design
- Algorithmic impact assessment
- Public transparency reporting
- Language model security controls
- Procurement specification drafting

Why Attend This Course

- Develop an operational deployment plan tailored to an institutional service, validated through stress testing against operational failure modes.
- Establish defensible human oversight controls that allow public bodies to explain system reasoning and preserve citizen contestability rights.
- Identify vulnerabilities in procurement contracts and technical system designs that risk data compromise or ungrounded outputs.
- Benchmark governance practices and deployment workflows against international institutional peers across various operational disciplines.

Conclusion

Deploying generative artificial intelligence within public institutions offers significant operational improvements in administrative speed and public engagement, provided that ethical safeguards, transparency, and accountability remain central to implementation. This programme guides participants from foundational public value assessment and governance principles through technical architecture, impact analysis, and operational security controls. Participants complete the course equipped with a tested deployment plan and practical oversight methodologies, preparing them to guide artificial intelligence adoption across their institutions with professional rigor and institutional integrity.
