



AI Risk and Accountability Course for Government Agencies



AI Risk and Accountability Course for Government Agencies

AI Risk and Accountability for Government Agencies is a five-day practitioner course for public servants who approve, buy, run or audit AI tools that affect citizens. You learn to inventory every AI use in your agency, rate its risk, assign oversight roles, publish a plain transparency record and watch live tools for harm. You leave with a completed Agency AI Risk Register. The course is delivered by Mercury Training Center.

About This Course

Agencies adopt chatbots, scoring models and document tools faster than anyone can explain who answers for them when they go wrong. You should already work in policy, risk, audit, procurement or public IT. You practice on agency scenarios: rating tools, writing records and defending oversight choices to a review panel.

Who It Is For

- Risk and internal audit staff who must show how AI tools are controlled.
- Policy and service owners who decide where AI touches citizens.
- Procurement teams who buy AI products and write supplier obligations.
- Public IT teams who deploy, configure and monitor AI tools.
- Legal and ethics advisers who review fairness, privacy and appeal routes.

This course is not for software engineers who want to build or tune models, who are better served by a technical machine learning course, or for senior officials who need only a short briefing on AI policy, who gain more from an executive session.

Competencies You Will Build

- AI use inventory: you find and list every AI tool in use, including those bought inside other software.
- Risk tiering: you rate each use by harm to citizens, scale and reversibility.
- Oversight design: you name who approves, who monitors and who can stop each tool.
- Transparency recording: you write public records that explain what a tool does in plain words.
- Supplier accountability: you put testing, logging and audit rights into AI contracts.
- Live monitoring: you track errors, complaints and drift and act on them.

What You Will Be Able to Do

By the end of the course you will be able to:

- From a list of agency services, identify where AI is used or planned and record each use.
- Given an AI use case, apply ISO/IEC 23894 risk guidance to rate likelihood and impact on citizens.
- Using a RACI template, assign approval, monitoring and stop authority for a deployed tool.
- Given a vendor proposal, draft contract clauses for testing, logs and audit access.
- From complaint and error logs, decide whether a live tool needs review, retraining or suspension.
- Using all findings, build an agency AI risk register with owners and review points.

Course Content

Day 1: Where AI Sits in Public Services

- Types of AI Use in Agencies: Chatbots, Scoring, Triage and Document Tools
- AI Use Inventory Template and Discovery Interviews
- Hidden AI Inside Purchased Software and Cloud Services
- Citizen Impact Mapping for Each Service Touchpoint
- Public Accountability Duties and the Right to an Explanation

Day 2: Rating AI Risk

- ISO/IEC 23894 Guidance on AI Risk Management
- Risk Tiering by Harm, Scale and Reversibility
- Bias and Fairness Testing on Eligibility and Scoring Tools
- Privacy and Security Threats Specific to AI Tools
- Recording Ratings and Rationale in the Risk Register

Day 3: Oversight Roles and Management Systems

- RACI for AI Approval, Monitoring and Stop Authority
- Human-in-the-Loop Design for Decisions About Citizens
- ISO/IEC 42001 AI Management System Elements for Agencies
- AI Review Boards: Terms of Reference and Escalation
- Appeal and Redress Routes for People Affected by AI Outputs

Day 4: Transparency, Suppliers and Live Monitoring

- Public Algorithm Registers and Transparency Record Templates
- Writing Plain-Language Explanations of AI Tools
- AI Procurement Clauses: Testing, Logging and Audit Rights
- Monitoring Dashboards for Errors, Complaints and Drift
- Incident Handling and Suspension Criteria for AI Tools

Day 5: Practice Review and Risk Register

- Exercise: Inventory and Tiering of a Sample Agency Portfolio
- Exercise: Transparency Record for a Benefits Triage Tool
- Exercise: Review Panel Challenge of an Oversight Design
- Exercise: Supplier Contract Gap Review
- Completing the Agency AI Risk Register

Case Studies and Exercises

The following are suggested activities the group may adapt to its own agencies.

- Case study: a social benefits office using a scoring tool to flag claims; you decide its risk tier and the human review it needs.
- Case study: a municipal service chatbot giving wrong permit advice; you trace accountability and draft the incident response.
- Exercise: a public hospital buying an AI triage product; you write the supplier testing and audit clauses.
- Exercise: a tax authority publishing its AI tools; you produce a plain-language transparency record.

What You Take Back

You take back the Agency AI Risk Register, built from your own inventory, risk ratings, oversight roles and transparency records. In your first month you use it to list the AI tools in one directorate, rate the highest-risk use and agree its owner with your manager.

- An AI use inventory with risk tiers.
- Oversight assignments and stop authority for each tool.
- Transparency record drafts for public-facing tools.
- A monitoring and review schedule.

Quick Answers FAQ

What should I know before an AI risk and accountability course for government agencies?

You need no coding skills. You should know how your agency approves new services, handles complaints and buys technology, and be comfortable reading a simple risk matrix.

How does AI risk and accountability training for agencies differ from an AI strategy course?

A strategy course decides where AI should create value. This one starts once tools exist or are being bought: rating their risk to citizens, naming who answers for them, publishing records and monitoring them in use.

What is an AI transparency record in the public sector?

It is a short public document describing an AI tool an agency uses: what it does, which decisions it supports, what data it uses, who is responsible and how people can question its results.

What do I take back from AI risk and accountability training for government agencies?

You take back the Agency AI Risk Register: an inventory with risk tiers, oversight assignments, transparency record drafts and a monitoring schedule, ready to apply to your own directorate.

For Your Manager

After this course your team member can list the AI tools your agency uses, rate their risk to citizens, assign clear oversight and write public transparency records. They return with the Agency AI Risk Register, first used to review one directorate's AI tools and to brief the body that approves new AI purchases.