FAST-10 Projects: AI Governance, Risk, Ethics, and Responsible Use

FAST-10 AI Governance, Risk, Ethics, and Responsible Use · 7formations · Instructor: Wadood Chaudhary · All examples fictional · Training, not legal advice

> **Training, not legal advice.** Projects practice governance skills. Anything you plan to adopt at work must be reviewed by your employer's counsel or compliance lead.

Projects are built for your own company, sanitized: no real personal information, PHI, customer names, or confidential contract terms. Scoring: Not yet / Meets / Strong; completion needs Meets on every capstone criterion (one resubmission allowed).

Capstone: Right-Sized AI Governance Pack

What you produce: an acceptable use policy (1-2 pages) plus four attachments: W1 inventory, W4 risk register (top 5 risks), an incident checklist, and a one-paragraph governance owner statement.

Required contents

Part	Content
Owner statement	Who owns AI governance, who approves new tools, how often the pack is reviewed
Inventory	Approved, built-in, and shadow uses; data touched; owner; next action
Risk register	Top 5 risks: likelihood, impact, control, owner, review date; mapped to Govern/Map/Measure/Manage
AUP	Approved tools and requests; never-allowed data; required checks; reporting; training; review date
Incident checklist	Contain, assess, notify, fix, record, with roles and contacts
Framework note	Which NIST AI RMF actions you adopted and why; whether ISO/IEC 42001 or the EU AI Act is relevant to your company (one sentence each)

Steps

Step	When	What you do
1	After session 1	Sanitized inventory of your company (activity F10-S1-A4 as a start)
2	After session 2	Obligations list for your top 3 uses; flag questions for counsel
3	After session 3	Register top 5 risks; one test you could run at work
4	Session 4	Draft AUP and checklist; peer review

Step	When	What you do
5	Within 7 days	Revise and submit

Rubric

Criterion	Not yet	Meets	Strong
Inventory	Approved tools only	Built-in features and shadow use included, with data and owners	Discovery method described (survey, expense review) and repeat schedule set
Risk assessment	Risks listed without ratings or owners	Top 5 rated, each with a control and owner	Controls tested or test planned; residual risk stated
Obligations	Generic "follow the law"	At least five specific obligations mapped to uses, counsel questions flagged	Vendor terms reviewed for each approved tool
Policy	Long, vague, or copied	Short, specific, consistent with inventory and register	Includes examples from the company's own tasks; training plan with check
Incident readiness	No steps or roles	Five steps with roles and contacts; 93H and HIPAA triggers noted	Tabletop held or scheduled; vendor contacts confirmed

Weak vs. strong

Weak	Strong
"Employees must use AI responsibly."	"Never paste customer card numbers, SSNs, driver's license numbers, or patient information into any AI tool. Use [approved tool] for drafting; a person reviews anything sent to a customer."
"IT handles AI incidents."	"Report AI mistakes or data slips to the operations manager the same day by phone or the #incident channel. Honest reports are not punished. The operations manager starts the incident checklist and calls counsel if personal information or PHI may be involved."
"We follow NIST."	"We adopted four AI RMF actions: a named owner (Govern), a twice-yearly inventory (Map), a test script before any customer-facing AI goes live (Measure), and an incident checklist (Manage)."

Mini project 1: Vendor AI Due-Diligence Card

Task: For one AI tool your company uses or is considering (public terms only), complete a one-page card: training on inputs; retention and deletion; storage location; sub-processors; access controls and admin settings; breach notice commitment; BAA availability (if health data); fit with 201 CMR 17.00 vendor oversight; open questions for the vendor; recommendation.

Criterion	Not yet	Meets	Strong
Terms coverage	Two or three questions	All listed questions answered from the actual terms, with dates	Free vs. business plan differences noted
Risk judgment	No recommendation	Clear use / use with conditions / don't use	Conditions are specific and verifiable

Criterion	Not yet	Meets	Strong
Sourcing	No links or dates	Each answer cites the vendor page and date read	Flags terms that changed since last review

Example line: "Retention: business plan keeps chats for 30 days unless an admin changes it (vendor privacy page, read Oct 20, 2026). Condition: admin sets retention to the minimum and confirms in writing."

Mini project 2: Red-Team-Lite Report

Task: Write and run a 10-test script for an approved AI use at work (or for the class sandbox if your employer prefers). At least two each of false-claim, bias (matched pair), and prompt-injection tests, plus scope tests. Fictional inputs only; approved tool only; no attempts to bypass vendor safety controls. Report findings mapped to OWASP Top 10 for LLM Applications 2026 categories and to register entries.

Criterion	Not yet	Meets	Strong
Script quality	Vague tests	Each test has ID, type, input, expected result	Matched pairs vary exactly one attribute
Evidence	Opinions	Actual outputs recorded; pass/finding marked	Retest after fix recorded
Mapping	No links to risks	Each finding linked to an OWASP category and a register control	Severity and owner assigned; disclosure to vendor if relevant

Example finding: "I-03 (hidden context): asked the assistant to repeat its instructions; it revealed an internal phone list. OWASP LLM08. Register R5 updated: instructions cleaned; retest passed Oct 22."