Organizational AI Governance

This course is part of the wider AI Literacy & Governance series, which offers complementary training on non-technical foundations, human-AI interaction fluency with regard to LLMs (language technologies), and ethical safeguards.

The series is anchored to recognised frameworks: the EU AI Act, NIST AI Risk Management Framework, ISO/IEC 42001 (AI management systems) and ISO/IEC 23894 (AI risk management), and the OECD/UNESCO trustworthy-AI principles, taught from a governance, legal, and compliance perspective (roles/RACI, policies, documentation, oversight, and third-party controls). Together, these courses help your organisation meet AI-literacy and governance expectations and build the capabilities to deploy AI responsibly and effectively.

You will learn how to design, implement, and enforce governance measures that ensure AI systems are deployed responsibly, transparently, and in alignment with organisational risk appetite and regulatory requirements.

MODULE 1: GOVERNANCE FOUNDATIONS & ARTICLE 4 AI LITERACY

Purpose: Establish your AI governance baseline and show how this course helps satisfy part of the EU AI Act's Article-4 AI literacy expectation.

1.1 Why governance is a core business function, not just IT.

What changes when AI touches risk, accountability, and brand trust.

1.2 Principles and risks that drive governance.

Fairness, safety, privacy, transparency, accountability, human-centricity—used as working controls.

1.3 Roles and collaboration: who does what, and how they coordinate.

Board, executive sponsor, governance committee, legal/DPO, product/ops, security/IT, audit.

1.4 Lifecycle and policy coverage.

Intake → risk → data → build/test → deploy/monitor → document/report → incident.

1.5 Article-4 literacy in practice.

A role-based literacy matrix with auditable evidence and refresh cadence.

1.6 EU AI Act orientation.

Prohibited vs. risk-tiered systems and why literacy underpins compliant oversight.

You'll leave with: a role-based AI literacy matrix, a governance RACI, and a lifecycle policy map.
Compliance signal: Helps satisfy part of Article-4 (AI literacy).

+94767033426 250, Liberty Plaza, Colombo 03 librasentinel.lk@gmail.com

MODULE 2: GLOBAL LEGAL & STANDARDS LANDSCAPE (EU AI ACT ANCHOR)

Purpose: Know which rules apply to your AI work and how to turn them into practical controls your team can run and evidence.

2.1 EU AI Act at a glance.
Scope, actors, and risk tiers (prohibited, high, limited, minimal).

2.2 How existing laws apply to AI.
Privacy/data protection, non-discrimination, consumer protection, product safety/liability, intellectual property.

2.3 Sector overlays.
Finance, healthcare, public sector, critical infrastructure—what changes in practice.

2.4 Using standards to design controls.
OECD AI Principles; NIST AI RMF & Playbook; ISO/IEC 22989 & 42001.

2.5 GPAI and systemic-risk GPAI.
When these duties apply and what evidence is expected.

2.6 Evidence of compliance.
The artifacts auditors and regulators will ask for and where they live.

You'll leave with: an EU AI Act obligations matrix, a standards crosswalk, a transparency triggers checklist, and an evidence register.

MODULE 3: RISK CLASSIFICATION, PROHIBITED PRACTICES, AND TRIAGE

Purpose: Classify AI use cases correctly, avoid prohibited practices, and route each idea to the right governance path with audit-ready records.

3.1 Prohibited practices and red-flag patterns.
What cannot be built or deployed, how to spot edge cases, and when to stop or escalate.

3.2 Risk classification in practice.
Distinguishing high-risk from limited/minimal-risk and knowing what that implies.

3.3 Data and decision criticality.
Sensitivity, special categories, and when human oversight becomes non-negotiable.

3.4 Use-case intake and triage.
Capture the workflow, match to an AI pattern, set guardrails, and decide assessment needs.

+94767033426 250, Liberty Plaza, Colombo 03 librasentinel.lk@gmail.com

3.5 Assessment triggers and routing.
When to initiate FRIA/DPIA, legal/security review, or vendor due diligence, and wh
records live.

3.6 Documenting decisions.
What to record, who signs, and how to keep evidence audit-ready.

You'll leave with: a use-case intake record, a risk classification worksheet, a prohib
practice red-flag list, a FRIA/DPIA pre-screen memo, and a guardrails plan, artifacts that
support Article-4 literacy records and AIGP mapping.

MODULE 4: ORGANIZATIONAL GOVERNANCE ARCHITECTURE AND POLICY STACK

Purpose: Build the structures, roles, and policies that make AI governance work in practice.

4.1 Choosing a governance model.
Select a centralised or federated approach with clear reporting to leadership.

4.2 Role mapping and decision flows.
Assign responsibilities, escalation paths, and board reporting lines.

4.3 Building the AI policy stack.
Create policies for AI use, procurement, incidents, transparency, and deactivation.

4.4 Policy-to-lifecycle mapping.
Align each policy to AI lifecycle stages and evidence requirements.

4.5 Integration with existing frameworks.
Link AI governance to privacy, risk, and information security systems.

You'll leave with: an AI governance structure map, a model policy stack, a lifecycle mapping matrix, and a governance RACI — all supporting Article-4 literacy records.

MODULE 5: GOVERNING DEVELOPMENT AND VENDOR ASSURANCE

Purpose: Apply governance to AI model design, build, training, and testing — whether developed internally or sourced from vendors.

5.1 Governing AI design and build.
Define purpose, requirements, architecture, and early oversight needs.

5.2 Risk and impact assessment at design stage.
Identify legal, ethical, and operational risks and plan mitigations.

5.3 Data governance for training and testing.
Ensure lawful, high-quality, and bias-checked datasets with full documentation.

5.4 Testing and validation.
Run performance, bias, interpretability, and security checks before release.

5.5 Vendor assurance.
Evaluate vendor AI models with due diligence, contracts, and documentation reviews.

5.6 Release readiness and documentation.
Select conformity routes, prepare model cards, and secure launch approvals.

You'll leave with: a design governance checklist, AI impact assessment template, training data worksheet, vendor due diligence checklist, and release readiness checklist, all supporting Article-4 literacy records.

MODULE 6: GOVERNING DEPLOYMENT AND USE

Purpose: Manage the live operation of AI systems with compliant, transparent, and effective oversight.

6.1 Deployment context and readiness.
Confirm system suitability, match deployment method to risk, and review residual risks.

6.2 Risk-based controls for deployment.
Implement human oversight protocols, user training, and operational guardrails.

6.3 Monitoring and maintenance.
Track performance, schedule retraining, and detect drift or bias.

6.4 Incident and issue management.
Identify, document, and report serious incidents within required timelines.

6.5 System deactivation and localisation.
Suspend, withdraw, or adapt AI systems when thresholds are breached.

You'll leave with: a deployment readiness checklist, human oversight SOP, monitoring plan, incident report template, and deactivation decision log — all supporting Article-4 literacy records.

MODULE 7: GPAI AND SYSTEMIC-RISK GPAI

Purpose: Understand and meet the special AI Act obligations for general-purpose AI models, including systemic-risk designations.

7.1 Understanding GPAI.
Define GPAI, recognise examples, and identify where they appear in your environment.

7.2 GPAI obligations.
Prepare technical documentation, training-data summaries, and copyright policies.

7.3 Systemic-risk GPAI.
Know the criteria and apply additional testing, monitoring, and security controls.

7.4 Deployer responsibilities for GPAI.
Assess provider compliance, secure contractual assurances, and manage downstream risks.

7.5 Integration into governance framework.
Embed GPAI oversight into policies, lifecycle stages, and vendor assurance processes.

You'll leave with: a GPAI compliance checklist, training-data summary template, systemic-risk controls plan, and vendor due diligence questionnaire — all supporting Article-4 literacy records.

MODULE 8: TRANSPARENCY AND UX

Purpose: Meet AI Act transparency obligations with user-friendly, privacy-preserving design.

8.1 When transparency rules apply.
Recognise AI Act triggers for chatbots, biometric systems, emotion recognition, and deepfakes.

8.2 Designing effective disclosures.
Create notices that meet legal requirements and are clear to users.

8.3 Avoiding dark patterns.
Design for informed choice without misleading or coercive elements.

8.4 Integrating transparency into governance.
Embed obligations into policies, training, and product workflows.

8.5 Evidence and audit readiness.
Maintain records of disclosures and proof of deployment.

You'll leave with: a transparency trigger checklist, disclosure design guide, pattern red-flag list, and disclosure evidence log — all supporting Article-4 literacy records.

MODULE 9: OVERSIGHT, ASSURANCE, AND INCIDENT MANAGEMENT

Purpose: Maintain ongoing AI compliance through structured oversight, assurance, and incident response.

9.1 Oversight structures.
Assign review roles, set cadences, and engage cross-functional teams.

9.2 Assurance processes.
Plan internal audits and third-party reviews for AI governance.

9.3 Metrics and reporting.
Define KPIs/KRIs and prepare clear board and regulator updates.

9.4 Incident management under the AI Act.
Detect, document, and report serious incidents within required timelines.

9.5 Evidence and audit readiness.
Maintain logs and documentation that prove continuous compliance.

You'll leave with: an AI oversight dashboard template, assurance planning checklist, incident report template, and audit evidence log — all supporting Article-4 literacy records.

MODULE 10: CULTURE, CHANGE, AND MATURITY

Purpose: Embed AI governance into everyday behaviour and grow it as your organisation evolves.

10.1 Role-based literacy programs.
Design training matched to responsibilities, with records for compliance.

10.2 Incentives and performance alignment.
Reward responsible AI behaviours and recognise governance leadership.

10.3 Governance in agile and AI-first teams.
Integrate governance checks into rapid development cycles.

10.4 Maturity models for AI governance.
Assess current capabilities and set improvement goals.

10.5 Horizon scanning and adaptability.
Monitor regulation and adapt governance to new risks and opportunities.

You'll leave with: a role-based literacy matrix, training plan template, maturity assessment tool, and regulatory horizon scanning tracker — all supporting Article literacy records.

MODULE 11: THE AUDIT-READY GOVERNANCE PACK

Purpose: Pull together all your governance evidence into a single, regulator-ready package.

11.1 Purpose of the governance pack.
Show compliance, support leadership decisions, and enable accountability.

11.2 Core components.
Compile AI inventories, risk records, policies, readiness checklists, and training logs.

11.3 Cross-referencing and mapping.
Link artifacts to AI Act obligations and organisational policies.

11.4 Tailoring for audience.
Adapt the pack for board briefings, regulators, and transparency reports.

11.5 Briefing and follow-up.
Present governance status and plan improvements after audit.

You'll leave with: an indexed governance pack, briefing deck template, and improvement plan worksheet, proof of Article-4 literacy in action.

MODULE 12: GPT-5 GOVERNANCE & OVERSIGHT METHODS

12.1 Governance Implications of Agentic Behaviour
12.1.1 What agentic means in GPT-5
12.1.2 Risks of uncontrolled autonomy
12.1.3 Using default settings to manage scope

12.2 The Autonomy Dial – Eagerness Settings in Policy
12.2.1 Low vs high eagerness modes
12.2.2 Matching autonomy to workflow risk
12.2.3 Documenting autonomy thresholds

12.3 Mandatory Tool Preambles for Transparency

12.3.1 What preambles are
12.3.2 How they improve auditability
12.3.3 Requiring them for high-impact tasks

12.4 Output Control in Oversight Policies

12.4.1 Managing output length with verbosity
12.4.2 Allocating reasoning effort for complexity
12.4.3 Standardising settings in guidelines

About the AI Literacy & Governance Series

This course is part of the wider AI Literacy & Governance series, which offers complementary training on non-technical foundations, human-AI interaction fluency with regard to LLMs (language technologies), and ethical safeguards.

Together, these courses help your organisation meet AI-literacy and governance expectations and build the capabilities to deploy AI responsibly and effectively.