

S201 SECURE & GOVERN · PRACTITIONER

AI Governance in Practice

Turning regulation and frameworks into a working governance system for AI.

FORMAT

Knowledge pack · 14 slides

FOR

Risk, security and IT leaders

LEVEL

Practitioner



THE PREMISE

AI governance as a continuous management system

Effective AI governance answers four questions continuously: what AI we use, what could go wrong, **who decided it was acceptable**, and how we know it still is.

Inventory

Every system, including features inside SaaS.

Risk

Harm to people and the business, by tier.

Accountability

Who decided, and on what basis.

Assurance

Monitoring, evaluation and periodic review.

REGULATION


The EU AI Act regulates by risk

CATEGORY	EXAMPLES	OBLIGATIONS
Prohibited	Social scoring, manipulation, emotion recognition at work	Banned since February 2025; two more bans from December 2026
High-risk	Annex III uses such as employment, credit and critical infrastructure	Risk management, data governance, logging, human oversight, conformity assessment
Transparency	Chatbots, and AI-generated or manipulated content	Disclosure and labeling
Minimal	Most other uses, such as spam filters	Voluntary codes; AI literacy measures still apply
General-purpose models	Broadly capable models, such as large language models	Documentation, copyright policy and training-content summary since August 2025; more for systemic-risk models

TIMELINE

Key dates after the 2026 amendments

The Digital Omnibus on AI, Regulation (EU) 2026/1744, in force since July 27, 2026, deferred the high-risk dates, added two prohibitions and eased the AI literacy duty.



ROLES

Your obligations depend on your role

Most enterprises are deployers. They become providers by building a system, by rebranding or substantially modifying a high-risk one, or by repurposing one into high-risk use (Art. 25).

ROLE	WHO IT IS	CORE OBLIGATIONS FOR HIGH-RISK SYSTEMS
Provider	Develops or commissions a system and places it on the market or puts it into service under its own name or trademark	Quality management system, risk management, data governance, technical documentation, logging, human-oversight design, conformity assessment, registration, post-market monitoring
Deployer	Uses a system under its own authority	Use it as instructed, assign competent oversight, check input data it controls, monitor and report incidents, keep logs at least six months, inform workers and affected people, assess fundamental-rights impact where Art. 27 applies

FRAMEWORKS

How ISO/IEC 42001, NIST AI RMF and the law fit together

Use ISO/IEC 42001 for the management system, NIST AI RMF for risk methods and the law for obligations. ISO/IEC 42005 (2025) guides AI system impact assessment.

ISO/IEC 42001

A certifiable AI management system: policy, roles, risk and impact assessment, controls, internal audit and improvement. Published in 2023.

NIST AI RMF


A voluntary framework in four functions: Govern, Map, Measure and Manage. Version 1.0 (2023), now under revision; Generative AI Profile (NIST AI 600-1), 2024.

Legal obligations

The EU AI Act, sector regulation and data protection law; in the United States, state laws such as Colorado's, challenged by the Justice Department and rewritten in 2026.

INTAKE

From use-case intake to periodic review



TIERING

Tier risk on factors you can assess

FACTOR	LOWER RISK	HIGHER RISK
Who is affected	Staff as users of assistive tools	Customers, applicants, the public
Role in the decision	Suggests; a person decides	Determines or materially shapes the outcome
Autonomy	A person reviews every output	Acts without review
Data	Public or internal	Personal or special-category data
Reversibility	Easily corrected	Hard to reverse once acted on
Domain	No sector rules	Employment, credit, health, critical services

AGENTS

Govern agents by what they are allowed to do

Singapore's Model AI Governance Framework for Agentic AI (January 2026) begins with two duties: assess and bound the risks upfront, and make humans meaningfully accountable.

- ✓ **Tier by action**
Set the risk tier by the most consequential action the agent can take.
- ✓ **Accountable person**
Each agent traceable to a named person who answers for what it does.
- ✓ **Checkpoints**
Human approval before significant or irreversible actions.
- ✓ **Bound in advance**
Approve tools, data and spending limits with the use case, at the minimum necessary.
- ✓ **Delegated authority**
Never beyond the rights of the person or role it acts for.
- ✓ **Stop**
A tested way to interrupt it, as Art. 14(4)(e) requires for high-risk systems.

INVENTORY AND SUPPLIERS

Govern the AI you buy as rigorously as the AI you build

Many AI features arrive through updates to software you already license.

RECORD FOR EVERY SYSTEM

- ◆ Purpose, owner and users
- ◆ Provider, model and version
- ◆ Data used, and where it is processed
- ◆ Legal category, role and risk tier
- ◆ Controls, evaluations, incidents and next review

REQUIRE FROM SUPPLIERS

- ◆ No training on your data without consent
- ◆ Data residency and sub-processors
- ◆ Notice of material model changes
- ◆ Incident notice, cooperation and access to logs
- ◆ Settings to disable features you have not approved

CONTROLS

Eight baseline controls

- ✓ **AI literacy**
Training matched to how each role uses AI.
- ✓ **Human oversight**
Designed into the workflow, with authority to stop.
- ✓ **Logging**
Inputs, outputs and decisions; high-risk logs kept at least six months.
- ✓ **Incident response**
Existing processes extended to AI harms; Art. 73 reports due within 2, 10 or 15 days.
- ✓ **Approved tools**
Sanctioned options that reduce use of unapproved AI tools.
- ✓ **Transparency**
People are told when they are interacting with AI.
- ✓ **Evaluation**
On representative cases, before release and on every material change.
- ✓ **Security testing**
Prompt injection and data leakage tested and controlled.

PITFALLS

How AI governance goes wrong

✗ Governing only what you build

✓ Include SaaS features and purchased AI

✗ Approval as a one-time event

✓ Review on change and on a fixed cycle

✗ Blocking every tool

✓ Offer approved alternatives

✗ Policy without an inventory

✓ Start by knowing what is in use

✗ Treating every use case as high risk

✓ Tier, and spend effort where risk is real

✗ Working apart from security and privacy

✓ One intake, shared controls and evidence

GETTING STARTED

A 90-day plan

DAYS 1–30

Foundations

- ◆ Name an accountable executive and a cross-functional council
- ◆ Publish acceptable use and approved tools
- ◆ Start the inventory, including SaaS AI features

DAYS 31–60

Process

- ◆ Launch intake and risk tiering
- ◆ Formalize role-based AI literacy (Art. 4)
- ◆ Add AI clauses to supplier contracts

DAYS 61–90

Assurance

- ◆ Assess the highest-risk use cases first
- ◆ Define monitoring and incident handling
- ◆ Map gaps against ISO/IEC 42001 and the EU AI Act

SUMMARY

Key takeaways

- 01 Governance is a continuous system: inventory, risk, accountability and assurance.
- 02 Under the amended EU AI Act, Annex III high-risk obligations apply from December 2027.
- 03 Know, system by system, whether you are a provider or a deployer.
- 04 Use ISO/IEC 42001 for the management system and NIST AI RMF for risk methods.
- 05 Hold purchased and embedded AI to the same intake, tiering and review.

NEXT COURSE · S202

Exposure Management

Secure & Govern · Practitioner