cycore

FREE DOWNLOAD

2026 EDITION

AI GOVERNANCE

AI Governance Starter Kit

Cycore's starter kit for governing the AI you build or buy. Compare ISO/IEC 42001, NIST AI RMF, and the EU AI Act, check the current EU application dates, and triage your systems in the inventory at the end.

PREPARED BY

Cycore

ONLINE

cycoresecure.com

If your team builds AI into a product or uses AI tools to make decisions, someone will ask how you govern it: a customer's security questionnaire, an auditor, or a regulator. This kit compares the three frameworks people mean when they say AI governance, ISO/IEC 42001, the NIST AI Risk Management Framework, and the EU AI Act (Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744), tells you which ones apply to you, and ends with an inventory you can fill in this week. We wrote it for technology and health care teams.

The three frameworks compared

The three differ in force. ISO/IEC 42001 is a certifiable standard, NIST AI RMF is a voluntary method, and the EU AI Act is law. They overlap in content, so the first thing to settle is which ones bind you and which ones you choose.

	ISO/IEC 42001:2023	NIST AI RMF 1.0	EU AI ACT
What it is	The international standard for an AI management system (AIMS). Certifiable.	A voluntary risk management framework from the US National Institute of Standards and Technology.	EU law. Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744.
Force	A voluntary standard. Customers and contracts ask for certification.	Voluntary. There's no certification.	Binding on the roles Article 2 covers: providers, deployers, importers, distributors, and others.
Who it's for	Any organization that develops, provides, or uses AI systems and wants certification.	Any organization that wants a structured way to manage AI risk.	Organizations in an Article 2 role, wherever they're established, once a system or its output reaches the EU.
Core focus	An AI management system: AI policy, risk assessment and treatment, lifecycle governance, continual improvement.	Four functions across the AI lifecycle: Govern, Map, Measure, Manage.	Risk-based classification of AI systems, prohibited practices, high-risk requirements, transparency, and protection of fundamental rights.
What it asks of you	Set up the AIMS. Define an AI policy and objectives. Assess and treat AI risks. Keep an AI system inventory and	Set up governance structures and culture. Map AI risks in context. Measure and track them.	Classify each system. Stop any prohibited practice. Meet the high-risk requirements and go through conformity assessment

	ISO/IEC 42001:2023	NIST AI RMF 1.0	EU AI ACT
	documentation. Implement the Annex A controls, with the guidance in Annex B. Monitor, measure, and keep improving.	Prioritize and act on them. Document and communicate the work. Tie it into your existing governance.	where they apply. Meet Article 50 transparency obligations where they apply. Register Annex III high-risk systems in the EU database. Run post-market monitoring.
Certification or assessment	Stage 1 and Stage 2 audits by an accredited certification body, annual surveillance audits, and a three-year recertification cycle.	None. Self-assess, or bring in a third party for a maturity review.	Conformity assessment for high-risk systems, by self-assessment or a notified body depending on the system, plus market surveillance by national authorities.
Best for	Companies that need a recognized certificate to show AI governance to customers and partners, especially B2B and SaaS companies selling to enterprises.	Organizations that want a flexible, risk-based starting point: US companies, federal contractors, and teams early in AI governance.	Any organization whose AI systems or their outputs reach the EU market. Confirm your role under Article 2.

ISO/IEC 42001:2023 is a management system standard, edition 1, published in December 2023. Certification shows that you run an AI management system. It doesn't prove compliance with the EU AI Act, which sets its own duties by role and by system.

NIST AI RMF 1.0 (NIST AI 100-1, January 2023) is voluntary. NIST says version 1.0 is being revised, and no revised version had been released when we checked on 2026-09-23, so confirm the current version before you build a program on it. For implementation, see Cycore's NIST AI RMF page.

The EU AI Act is law, and the 2024 text alone no longer gives you the current application dates. Regulation (EU) 2026/1744 moved several of them, and the schedule later in this kit reflects it. For the work involved, see Cycore's EU AI Act page.

Who the EU AI Act applies to

The Act applies by role and by facts, set out in Article 2. Start from what you do with the system and where.

- **Provider.** You develop an AI system or a general-purpose AI model, or have one developed, and you place it on the market or put it into service under your own name or trademark, paid or free (Article 3(3)). Providers are covered when they place a system on the EU market or put it into service in the EU, wherever they're established (Article 2(1)(a)).
- **Deployer.** You use an AI system under your authority, outside personal non-professional activity (Article 3(4)). Deployers established or located in the EU are covered (Article 2(1)(b)).
- **Providers and deployers outside the EU** are covered where the output the system produces is used in the EU (Article 2(1)(c)).
- **Importers and distributors** of AI systems, **product manufacturers** who place a product on the market with an AI system under their own name, and **authorized representatives** of non-EU providers are covered too (Article 2(1)(d) to (f)).
- Article 2 also carves out areas, for example systems placed on the market, put into service, or used exclusively for military, defense, or national security purposes (Article 2(3)). The exclusions have conditions, so read the text before you rely on one.

A US company with no EU office is in scope if it puts a system on the EU market or the system's output is used there. A third-party tool you use under your authority makes you the deployer, with a deployer's duties for that system. Close classifications, exclusions, and cross-border questions go to qualified legal review. General-purpose AI models carry their own obligations under Chapter V of the Act, which this kit doesn't cover.

Which frameworks you need

Answer these in order. Most teams end up with more than one.

1. **Do you develop, deploy, or use AI systems, including third-party tools with AI features?** If no: You may not need a dedicated framework yet. Keep a list of the third-party tools you use that have AI in them; there are usually more than people think. If yes: Continue.
2. **Do you place AI systems on the EU market, put them into service in the EU, use them in the EU, or have their outputs used in the EU?** If yes: The EU AI Act likely applies. Confirm your role under Article 2 in the section above, then keep going to see whether ISO/IEC 42001 or NIST AI RMF belong in your program too. If no: The Act is a lower priority. Revisit this question when your market changes.
3. **Do customers or partners ask for a certificate as proof of AI governance?** If yes: ISO/IEC 42001 certification is the recognized way to give it to them. If no: Certification still has value, and it can come later.
4. **Do you want a flexible, risk-based method to build your internal AI governance process?** If yes: NIST AI RMF is a good starting point. If no: ISO/IEC 42001 gives you a more structured, auditable approach.
5. **Do you work with US federal agencies or government contractors?** If yes: Read your contracts for a NIST AI RMF alignment clause. The framework itself is voluntary. If no: NIST AI RMF is still useful, and there's no pressure from this angle.
6. **Do you deploy AI in a high-risk area, such as health care, financial services, hiring and HR, law enforcement, or critical infrastructure?** If yes: Expect to need a combination: the EU AI Act if you have an EU nexus, ISO/IEC 42001 for certifiable governance, and NIST AI RMF for the risk method. If no: Start with one framework and expand as needed. NIST AI RMF for flexibility, ISO/IEC 42001 for certification value.

These are the combinations we'd start with. Your Article 2 role decides whether the Act applies.

SITUATION	WHERE WE'D START
US-based SaaS company selling to EU customers	EU AI Act plus ISO/IEC 42001
Health tech company, US market only	NIST AI RMF plus ISO/IEC 42001
Fintech with global customers	All three
US federal contractor building AI tools	NIST AI RMF first, ISO/IEC 42001 if the product is customer-facing
Early-stage startup exploring AI features	NIST AI RMF, as a lightweight start
Enterprise SaaS answering customer questionnaires	ISO/IEC 42001; a certificate carries weight in a questionnaire

EU AI Act risk tiers

The Act sets duties by category. It bans some practices, regulates high-risk systems closely, puts transparency obligations on certain systems, and leaves the rest without specific obligations. The labels "unacceptable risk", "limited risk", and "minimal risk" that you'll see in articles and vendor decks are editorial shorthand.

Prohibited practices (Article 5)

These can't be placed on the market, put into service, or used. Article 5(1) lists them, each with conditions and some with narrow exceptions, so read the text before you conclude a system is clear. In short:

- (a) subliminal, manipulative, or deceptive techniques that materially distort behavior and cause, or are likely to cause, significant harm
- (b) exploiting vulnerabilities of age, disability, or a social or economic situation to distort behavior in the same way
- (ba) and (bb), added by Regulation (EU) 2026/1744 and applying from 2 December 2026: generating or manipulating realistic intimate or sexually explicit material of an identifiable person without that person's explicit consent, and generating or manipulating material covered by the child sexual

abuse Directive (2011/93/EU). Article 5(1a) and (1b) set the conditions under which a provider or deployer falls under the ban.

- (c) social scoring that leads to detrimental treatment in unrelated contexts, or treatment that's unjustified or disproportionate
- (d) predicting the risk that a person commits a criminal offense based solely on profiling or personality traits
- (e) building or expanding facial recognition databases by untargeted scraping of facial images from the internet or CCTV
- (f) inferring emotions in the workplace or in education institutions, except for medical or safety reasons
- (g) biometric categorization to infer race, political opinions, trade union membership, religious or philosophical beliefs, sex life, or sexual orientation
- (h) real-time remote biometric identification in publicly accessible spaces for law enforcement, with narrow exceptions

The 2024 items have applied since 2 February 2025.

High-risk AI systems (Article 6, Annex I, Annex III)

A system is high-risk in one of two ways.

- **Article 6(1).** It's a safety component of a product, or is itself a product, covered by the EU product laws listed in Annex I, and that product needs a third-party conformity assessment under those laws. Annex I lists those product laws; medical devices are one example. The 2026 amendment adds that AI used only for non-safety aspects (user assistance, improving performance, service efficiency, automation or convenience, and quality control) doesn't count as a safety component, unless its failure would endanger health and safety (Article 6(1a) and (1b)).
- **Article 6(2).** It's an intended use listed in one of the eight Annex III areas:
  - biometrics
  - critical infrastructure
  - education and vocational training

- employment, workers' management, and access to self-employment
- access to essential private and public services and benefits, which includes creditworthiness and credit scoring and life and health insurance pricing
- law enforcement
- migration, asylum, and border control
- administration of justice and democratic processes

Each area lists specific intended uses, so read the entry for yours. Article 6(3) takes some Annex III systems out where they don't pose a significant risk, for example a system that performs a narrow procedural task. Document the reasoning either way.

Once a system is high-risk, its provider has to:

- run a risk management system across the lifecycle (Article 9)
- govern the training, validation, and testing data (Article 10)
- keep technical documentation and automatic logs (Articles 11 and 12)
- give deployers the information they need to use it correctly (Article 13)
- design it for human oversight (Article 14)
- meet accuracy, resilience, and cybersecurity requirements (Article 15)
- run a quality management system, go through conformity assessment, draw up the EU declaration of conformity, and apply the CE marking (Articles 17, 43, 47, and 48)
- register Annex III systems in the EU database before placing them on the market (Articles 49 and 71)
- monitor the system after it's on the market and report serious incidents (Articles 72 and 73)

Deployers of high-risk systems have their own duties under Article 26, starting with using the system in line with the provider's instructions, and some deployers must complete a fundamental rights impact assessment before use (Article 27).

Systems subject to transparency obligations (Article 50)

You'll hear these called "limited risk". That's shorthand; the Act's term is transparency obligations, and they attach to specific systems whatever their other classification.

- Providers of systems that interact directly with people must design them so people know they're dealing with an AI system, unless that's obvious in context (Article 50(1)).
- Providers of systems that generate synthetic audio, image, video, or text must mark the output in a machine-readable way as artificially generated or manipulated (Article 50(2)).
- Deployers of emotion recognition or biometric categorization systems must inform the people exposed to them (Article 50(3)).
- Deployers must disclose deepfakes, and AI-generated text published to inform the public on matters of public interest (Article 50(4)).

Each paragraph has exceptions, for example for evidently artistic or satirical work and for lawful crime detection. These obligations generally apply from 2 August 2026, with a transition for pre-existing synthetic-content systems in the schedule below.

Everything else

Systems outside Article 5, Article 6, and Article 50 have no specific obligations under the Act. People call this "minimal risk", and that's shorthand too. Article 95 encourages voluntary codes of conduct for them. Spam filters, inventory forecasting, and most recommendation engines sit here, and they still belong in your inventory, because a change of use can move a system into a regulated category.

How to classify a system

1. Start with Article 5. If the system does anything on that list, stop. Redesign it or don't deploy it.

2. Check Article 6. Is it a safety component of a product under the Annex I laws, or an intended use in Annex III? Read the specific Annex III entry for your area and the Article 6(3) conditions.
3. Check Article 50. Does it interact with people, generate synthetic content, recognize emotions, categorize people by biometrics, or produce deepfakes?
4. Classify the deployment. The same model can be unregulated in one use and high-risk in another.
5. Write down why you classified it the way you did, whatever the outcome. Regulators and auditors ask.

The current EU AI Act schedule

Regulation (EU) 2026/1744 rewrote the application dates for high-risk systems and added transitions. The table gives the dates as they stand in the consolidated text: Regulation (EU) 2024/1689, Articles 111 and 113, as amended by Regulation (EU) 2026/1744. Whether a specific system falls under a given provision remains a matter for qualified legal review.

DATE	WHAT APPLIES
2025-02-02	The original Chapter I and Chapter II application date remains relevant, subject to the specific amendments below.
2026-08-02	Article 50 transparency obligations generally apply.
2026-12-02	Providers of systems placed on the market before 2026-08-02 that generate synthetic audio, image, video, or text must comply with Article 50(2) by this date. Certain prohibited-practice amendments introduced by Regulation (EU) 2026/1744 also apply from this date.
2027-12-02	Chapter III, Sections 1 through 3 apply to high-risk systems covered by Article 6(2), which refers to Annex III use cases.
2028-08-02	Chapter III, Sections 1 through 3 apply to high-risk systems covered by Article 6(1), which concerns specified product and safety-component systems.
2030-08-02	Providers and deployers of high-risk AI systems intended for use by public authorities, placed on the market or put into service before the Chapter III application date in Article 113 that applies to them (2027-12-02 for Article 6(2) and Annex III systems, 2028-08-02 for Article 6(1) and Annex I systems), must comply by this date.

Two notes.

- **Existing high-risk systems.** Under Article 111(2) as amended, the Act applies to a high-risk system placed on the market or put into service before its Chapter III date only if the system's design changes significantly after that date. The Article 5 prohibitions and the public-authority rule in the last row apply regardless. Recital 39 of Regulation (EU) 2026/1744 says this runs by type and model, from the date the first unit of that type and model was placed on the market or put into service, and that a significant design change after the Chapter III date ends it.
- **Penalties.** Article 99 sets them. The amounts and how they scale depend on the infringement and the operator, so read the current text for your exposure.

AI policy outline

Ten parts. It's an outline of what the policy should cover; the specifics come from your own systems, risk profile, and obligations.

1. **Purpose and scope.** Your commitment to responsible AI development and use. Which AI systems are in scope: internally developed, third-party, and AI features embedded in other products. Which frameworks you align to (ISO/IEC 42001, NIST AI RMF, the EU AI Act). Roles: AI governance lead, risk owners, executive sponsor.
2. **AI system inventory.** How you find and catalog every AI system. How you classify each one, by risk tier, business function, and data sensitivity. The fields you capture: purpose, owner, data inputs, decision outputs, deployment date, vendor. Review cadence; we suggest quarterly.
3. **Risk assessment method.** How you assess AI-specific risks: bias, accuracy degradation, security weaknesses, privacy impact. Scoring criteria and thresholds. The link to your organizational risk register. Fundamental rights impact assessments where Article 27 applies.
4. **Data governance for AI.** Training data quality, representativeness, and provenance. Bias assessment and mitigation. Retention and deletion for

training and inference data. Data protection law (GDPR, HIPAA, and others) in the AI context.

5. **Model development and testing standards.** Documentation during development. Testing for accuracy, fairness, resilience, and security. Validation before production. Version control and change management for models. Standards for AI you buy.
6. **Monitoring and performance.** What you monitor in deployed systems. The metrics: accuracy, drift, fairness, error rates. Alert thresholds and escalation. A revalidation schedule.
7. **Human oversight.** Human-in-the-loop, human-on-the-loop, or human-in-command, decided per system. Minimum oversight by risk classification. Training for the people who oversee. Override and intervention procedures.
8. **Incident response for AI failures.** What counts as an AI incident: a wrong output, a bias event, a security breach, a safety issue. Reporting chain and deadlines. Investigation and root cause. Remediation and communication. Regulatory notification, including serious incident reporting under Article 73 where it applies.
9. **Transparency and explainability.** How you explain AI-driven decisions to the people they affect. Documentation of each system's logic and limits. What you disclose publicly about AI use. Article 50 obligations where they apply.
10. **Third-party AI vendor management.** Due diligence before you buy. Contract terms: transparency, audit rights, incident notification, data handling. Ongoing monitoring of the vendor. A responsibility matrix for what the vendor owns and what you own. Reassessment when the rules or the risk change.

Common mistakes

1. **Waiting for the rules to settle.** The prohibited practices have applied since February 2025 and the transparency obligations apply from August 2026. A program built after an obligation applies costs more and takes longer than governance built into development. Start the inventory now.

2. **Buying a framework before you know your inventory.** You can't govern what you haven't found. Before you pick a framework, hire a consultant, or buy a GRC tool, list the AI systems you have, including the ones nobody calls AI, such as the support chatbot, an automated fraud rule, and the tool marketing signed up for last quarter.
3. **Ignoring third-party AI tools.** Under the EU AI Act you have duties as a deployer of the AI you use, including tools someone else built. If you use a third-party tool to screen resumes, assess credit risk, or make clinical recommendations, the deployer duties for that system are yours.
4. **Over-scoping or under-scoping.** Over-scoping is trying to govern every AI system at once; the team burns out and progress stalls. Under-scoping is governing only the obvious systems and missing the ones that carry the most risk. Start from a complete inventory, prioritize by risk, and implement in phases.
5. **No human oversight process.** Every framework here expects human oversight, and Article 14 requires it for high-risk systems. AI systems degrade, hit edge cases, and produce unexpected outputs. If nobody is named to review, override, or intervene in an AI-driven decision, documentation won't close that gap. Write the process, train the people, and test it.

Your inventory and next steps

Start this week. A spreadsheet is enough.

1. **Inventory** every AI system you develop, deploy, or buy, including third-party tools with AI features and AI features inside products you already use. Capture the fields in the table below.
2. **Name an owner** for each system, and one person who owns the program.
3. **Triage** each system with the classification steps: prohibited, high-risk candidate, transparency obligations, or no specific obligation. Record the EU

nexus, meaning your Article 2 role if you have one, and whether the system sits in a high-risk area.

4. **Note the controls** in place against the policy outline: human oversight, monitoring, documentation, data governance, incident process, vendor terms. Mark the gaps, and give each gap an owner and a date.
5. **List the systems that need legal or conformity-assessment review:** anything close to Article 5, anything that may be high-risk under Article 6, any EU nexus where your role is unclear, any exclusion you're relying on, and cross-border cases.

FIELD	WHAT TO RECORD
System	Name and a one-line description
Owner	The person accountable for it
Built or bought	Internally developed, third-party, or an AI feature inside another product; the vendor's name if bought
Purpose and use	What it does and which decision or output it feeds
Data in, decisions out	Data inputs, including personal data; outputs and the people they affect
Deployment	Date; where it's placed on the market or used; EU nexus and your Article 2 role
Triage tag	Prohibited, high-risk candidate (Annex I or Annex III), transparency obligations (Article 50), or no specific obligation
Controls in place	Oversight, monitoring, documentation, data governance, incident process, vendor terms
Gaps	What's missing, with an owner and a date
Legal review	Yes or no, and why

Whatever the inventory shows, the next moves are the same. Work the high-risk candidates and the EU-nexus systems first, get the legal review list in front of counsel, and put the policy outline into your own words. If you already run a GRC platform, load the inventory and its triage tags into it; the platform tracks the controls, and the classification and the legal review list still come from you.

More Cycore guides and checklists

- SOC 2 Readiness Checklist
- ISO 27001 Gap Assessment
- HIPAA Compliance Checklist
- HITRUST Readiness Guide

Next step

Talk to Cycore about your AI governance program