

AI Law 2026 – Founder Checklist

1. First Pass: Do We Even Touch “High- Risk” AI?

- Do we use AI to decide or score:
 - Credit, lending, or insurance eligibility?
 - Hiring, promotion, or firing?
 - Access to education or exams?
 - Access to public benefits or services?
 - Critical infrastructure (energy, transport, utilities, healthcare systems)?
- Do we use biometrics (face recognition, voice, gait, emotion, etc.)?
- Do we analyze or try to influence emotions or vulnerabilities (kids, elderly, at- risk groups)?
- Are any of our systems deployed or sold into the EU?

If “yes” to any of the above, treat the product as potentially **high- risk** and move through the rest of this checklist more carefully.

2. “Unacceptable Risk” / Obviously Bad Ideas

- We do NOT do any kind of:
 - Social scoring of individuals for general “trustworthiness.”
 - Covert manipulation of vulnerable people.
 - Real- time remote biometric ID in public spaces (outside narrow, lawful exceptions).
- If we are unsure whether a use case is banned, we have written it down and flagged it for legal review before building or deploying.

If any box above would be checked “yes,” pause the project.

3. EU AI Act – Basic Readiness (If We Touch the EU)

- We know whether:

- We provide AI into the EU market.
- Our customers deploy our AI in the EU.
- We have roughly classified each AI feature:
 - High- risk (in sensitive domains listed above).
 - General- purpose model provider (foundation / base model).
 - Limited or minimal risk (e.g., content drafting, internal tools).

For **high- risk** or potentially high- risk systems:

- We have a short “AI risk memo” for each system that describes:
 - Purpose and use case.
 - Who is affected and how.
 - Main risks (errors, bias, security, misuse).
- We have some form of:
 - Basic risk management process (identify, mitigate, review).
 - Data governance (where data comes from, quality checks, retention).
 - Logging and monitoring (we can reconstruct what the system did).
 - Human oversight points (when people can override, audit, or stop the system).

4. 2026 EU Deadlines – Are We On Track?

If we are high- risk / EU- exposed:

- We understand that key obligations for high- risk systems become enforceable around **August 2, 2026**.
- Before that date we plan to:
 - Finish classifying our systems and documenting risks.
 - Implement logging and monitoring for those systems.
 - Define human- in- the- loop or human- on- the- loop controls.
 - Prepare technical documentation customers or authorities may request.

- We know who in the company “owns” EU AI compliance (even if it’s just one founder wearing the hat).

5. U.S. – Patchwork Reality + Federal Shift

- We know which US states our customers are in (or where we deploy AI).
- We’ve identified any **state laws** that clearly apply to:
 - Consumer AI services.
 - Data privacy.
 - Employment, credit, or insurance.
- We understand that a recent federal Executive Order:
 - Signals a move toward a unified national AI framework.
 - May eventually weaken or preempt some conflicting state AI laws.
- For now, we:
 - Treat the strictest relevant state requirement as our default.
 - **Keep a simple internal note of which state rules we’re relying on.**

6. Sector- Specific Rules (Beyond “AI Law”)

For our product, we have considered whether we fall under:

- Financial services / lending / payments
- Health / medical / patient data
- Employment / HR / recruiting
- Education / testing
- Critical infrastructure / security
- **Children’s services / apps aimed at minors**

For any “yes”:

- We have at least skimmed the main regulator’s AI or automation guidance.

- We understand any **extra** duties (e.g., explainability, record keeping, fairness tests).

7. Data & Transparency Basics

- We have a short, plain- language document (internal or public) that answers:
 - What data we collect for AI features.
 - How we use and store that data.
 - Whether we train or fine- tune models on user data.
 - How users can opt out where appropriate.
- We can show, if asked:
 - Where our training data came from (at a high level).
 - What steps we take to reduce bias and obvious errors.
 - How people can contest or override an important AI- assisted decision.

8. Enterprise / Buyer- Readiness

- We have a simple “AI FAQ” or “AI use and risk” one- pager for customers and investors.
- It covers:
 - What parts of our product are AI- driven.
 - Where we host and process data.
 - What third- party models or APIs we rely on.
 - Our basic security and privacy posture.
 - Our current stance on EU AI Act and U.S. compliance (even if rough).
- We’re prepared for security / legal questionnaires with a reusable base set of answers.

9. Internal Ownership & Review

- Someone in the company is clearly responsible for AI governance (even if part- time).
- We do at least one **annual** (ideally quarterly) AI use review:

- What AI we use and where.
- What has changed in law or guidance.
- **What new features might push us toward “high- risk.”**
- We have a simple process for:
 - Approving new AI features.
 - Retiring risky or non- compliant ones.
 - Responding to user or regulator complaints.

10. Red- Flag Situations (Stop and Get Help)

- We are rolling out AI that:
 - **Directly affects someone’s ability to get money, a job, education, housing, or public benefits.**
 - Uses live biometrics or emotion tracking in public or at scale.
 - Is sold to governments, law enforcement, or critical infrastructure operators.
- If any of the above is true, we have explicitly decided to:
 - Slow down and get specialist legal input.
 - Adjust the design to lower risk.
 - Or not ship at all until we understand the implications.