The Tech Founder's Guide to AI Compliance

A practical guide to navigating the evolving AI regulatory landscape

LAWSEL ADVISORY

www.thelawsel.com

Table of Contents

CHAPTER 1

The AI Regulatory Landscape in 2026

Artificial intelligence regulation has moved from policy discussion to enforceable law. Technology founders building AI-powered products now operate in an environment where non-compliance carries real financial and operational consequences. This chapter surveys the major regulatory frameworks you need to know.

The EU AI Act

The EU AI Act, which entered into force in August 2024 with phased compliance deadlines through 2027, is the most comprehensive AI-specific regulation globally. It introduces a risk-based classification system that determines the compliance obligations for AI systems based on their intended use.

Key deadlines for founders:

- February 2025: Prohibitions on unacceptable-risk AI systems take effect
- August 2025: Obligations for general-purpose AI (GPAI) models apply
- August 2026: Full compliance required for high-risk AI systems

Penalties under the EU AI Act are substantial — up to €35 million or 7% of global annual turnover for prohibited practices, and up to €15 million or 3% of global turnover for other violations.

United States: A Patchwork Approach

The US lacks a single federal AI law but has adopted a sector-specific and state-level approach. Key developments include:

- Executive Order 14110 on Safe, Secure, and Trustworthy AI (October 2023), which directs federal agencies to develop AI safety standards
- Colorado AI Act (SB 24-205), effective February 2026, requiring developers and deployers of high-risk AI to implement risk management practices
- New York City Local Law 144, regulating automated employment decision tools
- The NIST AI Risk Management Framework, which is voluntary but increasingly referenced in procurement and contract requirements

Founders serving US customers should expect growing state-level regulation and should build compliance programmes that can adapt to new requirements.

India: DPDPA and AI Implications

India's Digital Personal Data Protection Act (DPDPA) 2023, while primarily a data privacy law, has direct implications for AI systems that process personal data. The Act establishes:

- Consent-based processing requirements that affect training data collection
- Purpose limitation obligations that constrain how AI models can use personal data
- Data principal rights including the right to correction and erasure, which create challenges for trained models

- The Data Protection Board of India as enforcement authority

India has also signalled intent to develop AI-specific regulation. The Ministry of Electronics and Information Technology has issued advisories requiring government approval for deploying 'unreliable' AI models.

United Kingdom

The UK has adopted a pro-innovation, sector-specific approach through its AI Regulation White Paper (March 2023). Rather than creating new AI-specific legislation, the UK relies on existing regulators — the FCA, ICO, CMA, Ofcom, and others — to apply five cross-cutting principles:

- Safety, security, and robustness
- Appropriate transparency and explainability
- Fairness
- Accountability and governance
- Contestability and redress

While the UK approach is less prescriptive than the EU's, founders should not assume it is less demanding. Sector regulators have enforcement powers and are actively developing AI guidance.

China's AI Regulations

China has taken an early and active approach to AI regulation. Key regulations include:

- The Interim Measures for the Management of Generative AI Services (August 2023)
- The Algorithm Recommendation Regulations (March 2022)
- The Deep Synthesis Provisions governing deepfakes (January 2023)

These regulations impose requirements including algorithm registration, content moderation, training data transparency, and user notification obligations. Founders with any exposure to the Chinese market need specialist advice.

CHAPTER 2

Risk Classification: Where Does Your AI Product Fall?

The EU AI Act's risk-based framework is becoming the global reference point for AI classification. Understanding where your product falls determines your compliance obligations.

Unacceptable Risk (Prohibited)

The following AI practices are prohibited outright:

- Social scoring systems by public authorities
- Real-time remote biometric identification in public spaces for law enforcement (with narrow exceptions)
- AI systems that exploit vulnerabilities of specific groups (age, disability, social or economic situation)
- AI systems that infer emotions in workplace and educational settings
- Untargeted scraping of facial images from the internet or CCTV to build facial recognition databases
- AI-based manipulation techniques that cause harm

If your product falls into any of these categories, it cannot be deployed in the EU regardless of safeguards.

High-Risk AI Systems

High-risk AI systems face the most extensive compliance requirements. These include AI used in:

- Biometric identification and categorisation
- Critical infrastructure management (water, gas, electricity, transport)
- Education and vocational training (admissions, assessment, monitoring)
- Employment (recruitment, task allocation, performance monitoring, termination)
- Access to essential services (credit scoring, insurance pricing, emergency services)
- Law enforcement (risk assessment, polygraphs, evidence analysis)
- Migration and border control
- Administration of justice

High-risk systems must comply with requirements for risk management, data governance, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness, and cybersecurity.

Limited Risk

Limited-risk AI systems have transparency obligations only. This category primarily covers:

- Chatbots and conversational AI (must disclose that users are interacting with AI)
- Emotion recognition systems (must inform subjects)
- Deepfake generators (must label content as AI-generated)
- AI-generated content (must be machine-readable as AI-generated)

Many commercial SaaS products with AI features will fall into this category.

Minimal Risk

AI systems that pose minimal risk — such as AI-powered spam filters, inventory management systems, or video game AI — have no specific obligations under the EU AI Act. However, providers are encouraged to voluntarily adopt codes of conduct.

Even for minimal-risk systems, general obligations around data protection (GDPR), consumer protection, and product safety still apply.

Practical Classification Exercise

To classify your AI product, work through these questions:

- Does the system make or materially influence decisions about individuals?
- In which sector does the system operate?
- What data does the system process (personal, biometric, sensitive)?
- Who are the end users and affected persons?
- What is the potential impact if the system fails or produces biased outputs?
- Does the system interact directly with individuals?
- Does the system generate or manipulate content?

Document your classification reasoning — regulators will expect to see a defensible analysis.

CHAPTER 3

Documentation and Technical Requirements

For high-risk AI systems, the EU AI Act mandates detailed technical documentation. Even if your system is not classified as high-risk, maintaining thorough documentation is good practice and increasingly expected by enterprise customers and investors.

Technical Documentation Requirements

High-risk AI systems must maintain documentation covering:

- General description of the AI system (intended purpose, developer identity, version history)
- Detailed description of system elements (architecture, computational resources, third-party components)
- Description of the development process (design specifications, development methodology, training approach)
- Information on training, validation, and testing data (data sources, preparation methods, data characteristics, known biases)
- Performance metrics (accuracy, robustness, cybersecurity measures, discriminatory impact assessment)
- Description of the risk management system
- Description of post-market monitoring arrangements

This documentation must be kept up to date throughout the system's lifecycle.

Data Governance

Training, validation, and testing datasets must meet specific quality criteria:

- Data must be relevant, representative, and as free of errors as possible
- Data must be appropriate for the system's intended geographical, behavioural, or functional setting
- Bias detection and mitigation measures must be documented
- Data provenance must be traceable
- Privacy and data protection requirements must be met

For founders using third-party or open-source training data, establishing clear data lineage is critical. Licence terms, consent bases, and data processing agreements should be documented for each data source.

Logging and Monitoring

High-risk AI systems must be designed to automatically log events relevant to:

- Identifying situations that may result in risks to health, safety, or fundamental rights
- Facilitating post-market monitoring
- Monitoring the operation of the system

Logs must be retained for a period appropriate to the system's intended purpose and applicable legal obligations. As a practical matter, build logging infrastructure from the start — retrofitting is significantly more expensive.

Human Oversight

High-risk AI systems must be designed to allow effective human oversight, including:

- Measures to enable human operators to understand the system's capabilities and limitations
- Awareness of automation bias risks
- Ability to correctly interpret outputs
- Ability to override or reverse the system's output
- Ability to interrupt or stop the system

This requirement has practical design implications. AI-assisted decisions in high-risk contexts should include human review workflows, not just an override capability.

CHAPTER 4

Board-Level AI Governance

Effective AI governance starts at the top. Regulators, investors, and enterprise customers increasingly expect companies to demonstrate board-level oversight of AI risks and opportunities.

Establishing an AI Governance Committee

An AI governance committee should include:

- A board member or C-suite sponsor with accountability for AI governance
- Legal/compliance representation with knowledge of applicable AI regulations
- Technical leadership who can assess AI risk and capability
- Product/business leadership who understand commercial context
- An ethics or responsible AI lead (internal or advisory)

The committee should meet at minimum quarterly, with additional sessions triggered by significant model deployments, regulatory changes, or incidents.

Governance Framework Components

A functional AI governance framework should address:

- AI strategy and risk appetite — what level of AI risk is the organisation willing to accept?
- AI inventory — a register of all AI systems in development and production
- Risk assessment methodology — how AI-specific risks are identified, assessed, and mitigated
- Approval workflows — who authorises the development and deployment of AI systems
- Monitoring and audit — how deployed systems are monitored for drift, bias, and performance
- Incident management — how AI-related incidents are detected, escalated, and resolved
- Training and awareness — ensuring relevant staff understand AI governance requirements
- Third-party AI management — due diligence and oversight of AI vendors and partners

Reporting Framework

The AI governance committee should receive regular reporting on:

- AI system inventory changes (new systems, retirements, significant updates)
- Risk assessment outcomes and mitigation actions
- Compliance status against applicable regulations
- Incidents, near-misses, and complaints
- Regulatory developments and their impact
- Third-party AI vendor assessments
- Training completion rates

Reporting should be concise and decision-oriented. Dashboards that track key risk indicators are more useful than lengthy narrative reports.

CHAPTER 5

Practical Compliance Roadmap

Building an AI compliance programme does not require a large team or unlimited budget. The following phased approach helps founders prioritise and build capability incrementally.

Phase 1: Foundation (Weeks 1–4)

- Map all AI systems in your product — identify what qualifies as AI under applicable definitions
- Classify each system by risk level using the EU AI Act framework
- Identify which jurisdictions you serve and which regulations apply
- Appoint an AI governance lead (this can be an existing role with expanded responsibilities)
- Conduct a gap assessment against applicable requirements
- Establish an AI system register

This phase is about understanding your current position. Do not attempt to fix everything at once.

Phase 2: Core Compliance (Weeks 5–12)

- Develop risk management procedures for each high-risk system
- Create or update technical documentation
- Implement logging and monitoring capabilities
- Establish human oversight mechanisms where required
- Review and update data governance practices for training data
- Develop an AI-specific incident response plan
- Update contracts with customers to address AI-specific terms (transparency, liability, data use)
- Brief the board or leadership team on AI governance obligations

Phase 3: Maturity (Weeks 13–24)

- Implement ongoing bias monitoring and fairness testing
- Establish a regular AI audit cycle (internal or third-party)
- Develop AI-specific procurement and vendor management procedures
- Build AI governance into product development processes (by design, not as an afterthought)
- Create staff training programmes on responsible AI
- Establish regulatory monitoring to track new requirements
- Prepare for conformity assessment requirements (EU AI Act)
- Document your compliance programme for regulator and customer inquiries

Quick Wins

Some actions deliver disproportionate value:

- Add AI disclosure notices where your product uses AI in customer-facing interactions — low effort, high trust signal
- Document your training data sources and licence terms — this is often the hardest gap to close later

- Add a 'human in the loop' option for any AI-assisted decisions that affect individuals — even if not legally required, it builds customer confidence
- Publish a responsible AI statement on your website — signals commitment to customers and regulators

CHAPTER 6

Key Contract Considerations for AI Products

AI products introduce contract considerations that traditional software agreements do not adequately address. Whether you are the provider or customer of AI-powered services, your contracts need to cover these areas.

Intellectual Property in Training Data

Key questions to address contractually:

- Who owns or has rights to the training data?
- Can customer data be used to train or improve the AI model?
- If customer data improves the model, who owns the resulting improvements?
- What happens to the model if the customer relationship ends?
- Are there restrictions on using the model with competing customers?

Best practice: be explicit about data use for training purposes. Ambiguity in this area creates legal risk and erodes customer trust.

Liability Allocation

AI-specific liability issues include:

- Who is liable when the AI produces incorrect, biased, or harmful outputs?
- How is liability allocated between the AI provider, the deployer, and the end user?
- Are there specific exclusions for AI-generated outputs?
- What standard of care applies — is the AI expected to perform at human expert level?

The EU AI Act imposes specific liability on both providers and deployers of high-risk AI, which must be reflected in commercial agreements.

Performance and SLA Considerations

Traditional SLAs focus on uptime and response time. AI products need additional performance metrics:

- Accuracy thresholds (and how accuracy is measured)
- Bias and fairness metrics
- Latency for real-time inference
- Model drift monitoring and retraining commitments
- Transparency reporting (e.g., confidence scores)

Define these metrics clearly in the contract and establish measurement methodologies that both parties agree on.

Regulatory Compliance Obligations

Contracts should address:

- Which party is responsible for regulatory compliance (and in which jurisdictions)?
- Cooperation obligations for regulatory inquiries or audits
- Notification requirements for regulatory changes that affect the service
- Conformity assessment obligations under the EU AI Act
- Record-keeping and documentation access

As AI regulation evolves, include mechanisms for contract adaptation — such as regulatory change clauses that trigger renegotiation.

Ready to discuss your compliance needs?

Book a free 30-minute consultation to explore how LawSel Advisory can support your business.

BOOK A CONSULTATION

calendly.com/rini-thelawsel/30min

GET IN TOUCH

rini@thelawsel.com
www.thelawsel.com