AI Development & Deployment Governance Framework

Navigate EU AI Act, state regulations, bias audits, and compliance requirements

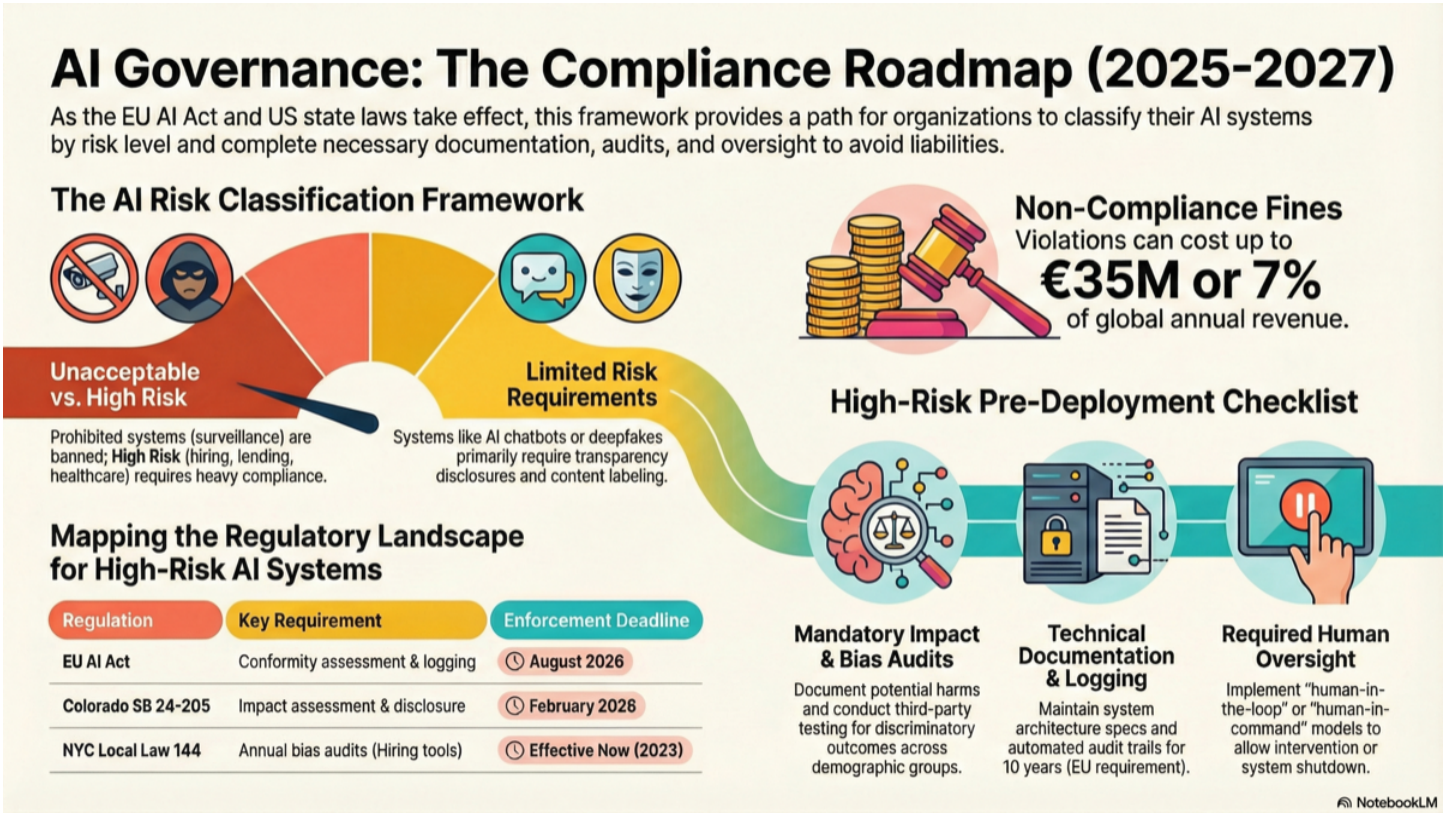

Scott Armbruster

AI Strategy & Systems Partner

For organizations building AI-powered products, features, and automated decision systems

scottarmbruster.com | linkedin.com/in/scottrarmbruster | info@scottarmbruster.com

The Regulatory Landscape (2025-2027)

AI governance isn't optional. Between the EU AI Act, state laws, and sector-specific regulations, ignorance is expensive. Here's what's actually enforceable.

EU AI Act

Adopted April 2024 | Full enforcement by 2027

• Risk classification (Unacceptable/High/Limited/Minimal) • Conformity assessments for high-risk AI • Technical documentation and audit trails • Fines up to €35M or 7% of global revenue

Applies to ANY company offering AI in EU market

Colorado SB 24-205

Effective February 2026

• Impact assessments for high-risk AI • Disclosure requirements • Risk management framework • Fines up to $20K per violation

Applies to high-risk AI deployed in Colorado

NYC Local Law 144

Effective July 2023 (Active Now)

• Annual bias audit for hiring AI • Published summary statistics • Candidate notification requirements • Daily fines for non-compliance

Applies to automated employment decision tools in NYC

GDPR Article 22

Active since 2018

• Right to explanation for automated decisions • Right to human review and contest • Data Protection Impact Assessments (DPIA) • Fines up to €20M or 4% of global revenue

Applies to automated decisions affecting EU residents

AI System Risk Classification

First step: classify your AI system. This determines what documentation, testing, and oversight you need.

UNACCEPTABLE RISK (Banned)

Examples: Social scoring by governments, mass surveillance, subliminal manipulation, exploiting vulnerabilities

Requirements: Do not deploy. Stop development immediately.

HIGH RISK (Heavy Compliance)

Examples: AI for hiring, lending, healthcare, law enforcement, education, insurance, housing

Requirements: Impact assessment + bias audit + technical docs + human oversight + logging required before deployment

LIMITED RISK (Disclosure Required)

Examples: Chatbots, AI-generated content, deepfakes, emotion recognition

Requirements: Disclose AI interaction to users. Label AI-generated content. Basic documentation.

MINIMAL RISK (Voluntary)

Examples: Spam filters, video games, inventory management, weather prediction

Requirements: No specific AI regulations apply. Follow general software best practices.

Quick Decision Tree:

1. Does your AI manipulate, exploit, or surveil people in prohibited ways? → UNACCEPTABLE
2. Does it make decisions about people in sensitive areas (hiring, lending, healthcare)? → HIGH RISK
3. Does it interact with users or generate content? → LIMITED RISK
4. Otherwise → MINIMAL RISK

High-Risk AI: 8-Step Pre-Deployment Checklist

All 8 must be complete before launching a high-risk AI system in regulated markets. One missing requirement = non-compliant.

1. Pre-Deployment Impact Assessment Completed

Document intended use, potential harms, mitigation strategies, affected populations

• System description and decision-making role • Data sources and training data • Reasonably foreseeable risks • Risk mitigation measures • Affected demographic groups • Human oversight mechanisms

Required by: Colorado SB 24-205 (2026), EU AI Act Article 27

2. Independent Bias Audit Conducted

Third-party or internal audit testing for discriminatory outcomes across protected classes

• Disparate impact analysis by race, gender, age • Selection rate ratios or impact ratios • Confusion matrix metrics by group • Statistical significance testing • Audit date, methodology, sample size

Required by: NYC Local Law 144 (annual), EU AI Act (before deployment + major changes)

3. Technical Documentation Package Created

Comprehensive documentation of system design, training, testing, performance

• System architecture and design specs • Training data description and methodology • Model performance metrics • Validation and testing results • Known limitations and failure modes • Cybersecurity measures • Change log and version control

Required by: EU AI Act Annex IV. Retention: 10 years (EU), 3-7 years (US)

4. Model Card Published

Standardized ML model documentation (Mitchell et al. framework)

• Model details (version, type, license) • Intended use and out-of-scope uses • Performance metrics by subgroup • Training and evaluation data • Ethical considerations • Caveats and recommendations

Best practice (not legally required but recommended)

5. Human Oversight Mechanism Implemented

GDPR Article 22 and EU AI Act require human review capability

• Human-in-the-loop (approves each decision) • Human-on-the-loop (monitors and can intervene) • Human-in-command (can override or shut down) • Document which model is used, who has authority, intervention triggers

Required by: GDPR Article 22, EU AI Act

6. Automated Logging and Audit Trails Active

Enable traceability through comprehensive logging

• Input data for each decision • Model output and confidence scores • Timestamp and system version • User who invoked the model • Human override events • Errors and exceptions

Required by: EU AI Act Article 12. Retention: 3-10 years, protected from tampering

7. User-Facing Transparency Disclosures

Inform affected individuals that AI is used and provide meaningful information

• Clear notice AI is making/assisting decisions • Purpose and general logic of the system • Right to contest or request human review • Contact for questions or complaints

Required by: GDPR, NYC Local Law 144, California AB 2930

8. Post-Deployment Monitoring Plan Established

Continuous monitoring required for high-risk AI in operation

• Performance metrics tracked • Drift detection (data and concept drift) • Bias metrics tracked over time • Incident reporting procedures • Review cadence and responsible team • Triggers for re-audit

Required by: EU AI Act Article 61

AI System Inventory Template

You can't govern what you can't see. Maintain a living inventory of all AI systems in development or production.

Track These Fields for Each AI System:

- **System Name:** Internal identifier
- **Owner/Product Team:** Team responsible for system
- **Risk Classification:** Unacceptable / High / Limited / Minimal
- **Use Case:** What the system does
- **Deployment Status:** Research / Dev / Staging / Production
- **Deployment Date:** When system went live
- **Affected Geographies:** EU / US (which states) / Other
- **Impact Assessment:** Link to document or 'N/A'
- **Last Bias Audit:** Date of most recent audit
- **Technical Docs:** Link to documentation package
- **Human Oversight Model:** In-loop / On-loop / In-command / None
- **Logging Enabled:** Yes / No
- **Next Review Date:** Scheduled re-audit or monitoring review

Governance tip: Store this inventory in a shared system (Airtable, Notion, internal wiki) accessible to Legal, Compliance, Product, and Engineering. Review quarterly. Add new systems as they move from dev to staging.

AI Incident Response: 5-Step Protocol

When your AI system fails, discriminates, or causes harm, response speed matters. EU AI Act Article 62 requires reporting serious incidents to regulators within 15 days.

1. Detect & Triage (Hour 0-2)

• Log incident with timestamp • Assess severity (Critical / High / Medium / Low) • Notify AI system owner • Preliminary impact assessment (how many users affected?) • Critical severity: immediate halt/rollback if ongoing user harm

2. Contain & Mitigate (Hour 2-24)

• Decide system status (continue / modify / suspend / shut down) • If suspended: notify affected users • Preserve logs and evidence (do not delete) • Activate human oversight or fallback • Brief Legal and Compliance teams

3. Investigate Root Cause (Day 1-7)

• Review logs and system behavior • Reproduce issue in test environment • Identify root cause (model, data, code, process) • Document findings in incident report • Assess if similar issues exist in other AI systems

4. Remediate & Test (Day 7-30)

• Implement fix (retrain, update data, change logic) • Re-run bias audit and impact assessment if High Risk • Test fix in staging • Validate no new issues introduced • Update technical documentation

5. Redeploy & Monitor (Day 30+)

• Gradual rollout with enhanced monitoring • Communicate fix to affected users • Report to regulators if required (EU: serious incidents within 15 days) • Post-incident review with team • Update incident response procedures

EU AI Act Reporting: Providers of high-risk AI must report serious incidents (death, serious health issues, fundamental rights violations) to national authorities within 15 days. Establish your reporting chain now.

Regulatory Compliance Quick Reference

Map your AI system to applicable laws.

Regulation	Applies To	Key Requirement	Deadline
EU AI Act	High-risk AI in EU market	Conformity assessment, docs, logging	Aug 2026 (high-risk) May 2027 (all)
Colorado SB 24-205	High-risk AI in Colorado	Impact assessments, disclosure	Feb 2026
NYC Local Law 144	Automated hiring tools (NYC)	Annual bias audit, publish stats	Effective July 2023
GDPR Article 22	Automated decisions (EU residents)	Right to review, explanation	Active since 2018
ECOA	AI credit decisions (US)	Adverse action notices, non-discrimination	Active
FDA AI/ML	AI medical devices (US)	Premarket review, algorithm change protocol	Active (evolving)

Not sure which applies? Start with geography (where are your users?) and use case (what decisions does your AI make?). If you're making decisions about people in sensitive areas, assume multiple regulations apply.

Building AI Systems? Don't Navigate Compliance Alone.

The regulatory landscape is complex and penalties are severe. Whether you need help classifying your AI systems, running bias audits, building impact assessments, or preparing for EU AI Act compliance, I can help you get it right before launch.

Scott Armbruster

AI Strategy & Systems Partner

Website: scottarmbruster.com

LinkedIn: linkedin.com/in/scottrarmbruster

Email: info@scottarmbruster.com

Services: scottarmbruster.com/agency