

THE COMPLIANCE PRACTITIONER'S
**FIELD GUIDE TO AI,
PRIVACY, AND
SECURITY REGULATIONS**

AN ENCYCLOPEDIC REFERENCE - MAPPED TO THE AI GOVERNANCE
STACK

A companion to

GOVERNING INTELLIGENCE

Law, Privacy, Security, and Compliance in the Age of Artificial Intelligence

Noah M. Kenney

Founder & Principal Consultant, Digital 520

President, Ethical Tech Forum • President & Chief Scientist, Disruptive AI Lab

FIRST EDITION • 2026

Copyright

© 2026 Noah M. Kenney. All rights reserved. Published by Digital 520.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the author, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Legal Disclaimer

This publication is provided for informational and educational purposes only and does not constitute legal, regulatory, compliance, or professional advice. The content reflects the author's analysis and interpretation of artificial intelligence governance, law, privacy, security, and compliance frameworks as of the date of publication. It is not intended to serve as a substitute for advice from qualified legal counsel, regulatory authorities, or other licensed professionals.

Artificial intelligence governance is a rapidly evolving field. Laws, regulations, regulatory guidance, and enforcement practices may change over time and may vary significantly across jurisdictions. Readers are responsible for ensuring that their use of any concepts, frameworks, or recommendations in this publication complies with applicable laws, regulations, and organizational requirements in their specific context.

The frameworks, models, thresholds, and decision rules presented in this text are intended as reference implementations and starting points for analysis. They must be adapted, validated, and approved based on the specific risk profile, industry requirements, and legal obligations of each organization.

The author and publisher disclaim all liability for any loss, damage, or legal consequences arising directly or indirectly from the use, application, or reliance on the information contained in this publication.

Nothing in this publication creates any attorney-client, advisory, fiduciary, or other professional relationship between the author and the reader. By using this publication, the reader acknowledges and agrees to these terms.

Permissions and Contact

For permission requests, speaking engagements, or consulting inquiries, contact:

Noah M. Kenney • Info@NoahKenney.com

How to Use This Field Guide

Artificial intelligence law in 2026 is not one body of law. It is a layered, jurisdictionally fragmented, sector-conditioned, and rapidly evolving network of statutes, regulations, agency guidance, common-law doctrines, voluntary standards, and contractual norms. A practitioner advising on AI governance is, in practice, advising simultaneously on privacy law, consumer protection, civil rights, sectoral regulation, cybersecurity, intellectual property, and emerging horizontal AI regimes - each with its own enforcement authority, evidentiary expectations, and remedial structure.

This Field Guide is the practitioner-facing companion to "Governing Intelligence: Law, Privacy, Security, and Compliance in the Age of Artificial Intelligence." Where the textbook builds the conceptual foundation, this volume is an operational translation layer that maps legal obligations into enforceable system controls. It is not a catalog. It is the operating manual practitioners use to convert regulatory text into auditable, testable, monitor-able infrastructure.

The AI Governance Stack as Operational Translation Layer

Each entry concludes with a Stack Lens callout mapping the law's obligations onto the five-layer AI Governance Stack:

- **Layer 1 - Data Governance:** inventory, classification, quality, bias assessment, provenance, consent.
- **Layer 2 - Model Governance:** architecture, training, fairness/robustness testing, interpretability, documentation.
- **Layer 3 - System Integration:** integration architecture, data-pipeline security, cascading failure analysis, human-AI interaction, boundary testing.
- **Layer 4 - Control & Monitoring:** access control, real-time monitoring, anomaly detection, incident response, deployment gates.
- **Layer 5 - Audit & Evidence:** documentation standards, evidence preservation, audit mechanisms, regulatory reporting, stakeholder communication.

The Stack is not interpretive. It is executable. It enables organizations to convert regulatory ambiguity into system-level controls that can be tested, monitored, and audited. It reduces audit failure risk by eliminating the gap between policy commitments and operational implementation. It compresses compliance implementation time by giving every new regulation a known landing zone. And it standardizes cross-jurisdiction execution: a control built once at the right Stack layer satisfies obligations across multiple regimes without reimplementing.

A regulator's expectation that a high-risk AI system be "subject to appropriate human oversight" is a Layer 4 deployment-gate problem long before it is a courtroom problem. A regulator's demand for "training data documentation" is a Layer 1 lineage problem long before it is a Layer 5 disclosure

problem. The Stack Lens callout in each entry identifies where, in operational terms, the law actually bites - and where the system control belongs.

Coverage

This volume covers more than one hundred legal regimes, regulatory frameworks, and voluntary standards across:

- **United States - Federal:** the FTC Act, HIPAA/HITECH, GLBA, FERPA, COPPA, FCRA, ECOA, TCPA, CFAA, ECPA, the Privacy Act, FISMA/FedRAMP, EO 14179, VPPA, CAN-SPAM, DPPA, Section 230, the CLOUD Act, FISA Section 702, Title VII, ADEA, GINA, OFCCP, the Fair Housing Act, ADA Title III, the NLRA, CISA/CIRCA, NERC CIP, TSA cybersecurity directives, NRC, FAA, FERC, and DOD Responsible AI.
- **United States - State:** twenty-plus comprehensive state privacy laws (CCPA/CPRA, VCDPA, CPA, CTDPA, UCPA, ICDPA, INCDPA, TIPA, MCDPA, NJDPA, DPDPA, NHDPA, KCDPA, MODPA, MNCDPA, RIDTPPA, NDPA, FDBR, OCPA, TDPSA, and more), all major biometric statutes (BIPA, CUBI, RCW 19.375), the leading state AI statutes (Colorado AI Act, TRAIGA, California AB 2013/SB 942/AB 3030/SB 53, NYC Local Law 144, Illinois HB 3773, Tennessee ELVIS Act, Utah AIP), New York SHIELD/Part 500, Massachusetts 201 CMR 17, Washington MHMDA, Colorado biometric/neural data, and the multistate landscape for breach notification, election deepfakes, NCII, student privacy, and workforce surveillance.
- **Sector-specific:** FDA AI/ML SaMD, Federal Reserve SR 11-7, CFPB Circulars, EEOC AI guidance, Section 1557 of the ACA, the NAIC Model Bulletin, SEC predictive analytics, and the wave of state healthcare AI laws.
- **Intellectual Property:** a comprehensive chapter on copyright (US Copyright Office AI guidance, NYT v. OpenAI, Thaler v. Perlmutter), USPTO AI inventorship guidance, the EU CDSM Article 4 TDM exception, UK and Japan TDM regimes, trade secret protection of AI models, open source AI licensing (OSI definition, OpenRAIL, Llama Community License), and right of publicity / NO FAKES Act.
- **European Union:** the EU AI Act, GDPR, Digital Services Act, Digital Markets Act, NIS2 (with implementing acts), DORA, Data Act, Data Governance Act, ePrivacy Directive, Cyber Resilience Act, Product Liability Directive update, AI Liability Directive, eIDAS 2.0, Cybersecurity Act, EU Health Data Space, and the GPAI Code of Practice.
- **United Kingdom and the Americas:** UK GDPR/DUAA, UK sectoral AI, UK Online Safety Act, Canada PIPEDA + Quebec Law 25, Brazil LGPD, Mexico LFPDPPP, Argentina, Chile (Law 21.719), Colombia, Peru, Uruguay, Costa Rica, Dominican Republic, Panama, Ecuador.
- **Asia-Pacific:** China PIPL/CSL/DSL plus Generative AI Measures, Deep Synthesis, and Algorithmic Recommendation Provisions; Japan APPI; Korea PIPA + AI Basic Act; Singapore PDPA + Model AI Governance Framework; India DPDP Act; Australia Privacy Act; New Zealand Privacy Act; Vietnam PDPD; Indonesia PDP Law; Thailand PDPA; Philippines DPA; Malaysia PDPA; Hong Kong PDPO; Taiwan PDPA.

- **Middle East and Africa:** Israel PPL Amendment 13, UAE Federal PDPL + DIFC + ADGM, Saudi Arabia PDPL, Bahrain PDPL, Qatar PDPPL, Oman PDPL, Kuwait DPPR, Egypt PDPL, Nigeria NDPA, Kenya DPA, Ghana DPA, Morocco 09-08, Tunisia, Turkey KVKK, South Africa POPIA.
- **International soft law and standards:** the Council of Europe Framework Convention on AI, OECD AI Principles, UNESCO Recommendation on AI Ethics, NIST AI RMF + Generative AI Profile, ISO/IEC 42001 / 23894 / 27001 / 27701, MITRE ATLAS, OWASP LLM/ML Top 10, NIST SP 800-53, SOC 2, CSA Cloud Controls Matrix, IEEE 7000 series, CIS Controls, ENISA AI threat materials, and PCI DSS as applied to AI.

How Each Entry Is Structured

Every entry follows the same template: title, citation, jurisdiction/effective/regulator/scope metadata, applicability, core obligations, penalties and enforcement, recent developments, a Stack Lens callout, and (on higher-stakes entries) a Practitioner Note and a Common Failure Pattern callout. The structure permits direct comparison across jurisdictions for any dimension of interest: penalty exposure, breach timing, automated-decisioning rights, sensitive-data treatment, fairness obligations, or vendor flow-down requirements.

Practitioner Notes compress advisory experience into a single operational instruction. **Common Failure Pattern** boxes call out the specific implementation mistakes that produce most of the enforcement actions in a given regime: the audit findings to design out before they appear.

Currency and Use

Coverage is current through April 2026. AI regulation is a moving target; the 2024–2026 period has seen the most sustained legislative and regulatory activity in the field's history, and the trajectory is accelerating. Practitioners must verify the operative text and most recent regulatory guidance before relying on any entry. Citations to primary regulations are provided to enable that verification. Soft-law instruments and voluntary standards are included where they materially shape compliance expectations.

This Guide is designed as a baseline implementation model, not a theoretical construct. Use it in three modes: (1) jurisdictional or sectoral lookup via the Table of Contents and alphabetical index; (2) cross-jurisdictional comparison of equivalent obligations using the Stack Lens callouts; and (3) audit and program preparation, mapping internal controls onto the regulations and standards in scope. The companion textbook provides the underlying methodology in full.

Table of Contents

When opening this document in Microsoft Word, the Table of Contents will populate automatically. If page numbers do not appear, right-click anywhere in the table and select "Update Field" (or press F9 with the cursor placed in the table) to refresh.

How to Use This Field Guide.....	3
The AI Governance Stack as Operational Translation Layer.....	3
Coverage.....	4
How Each Entry Is Structured.....	5
Currency and Use	5
P A R T I	14
Privacy, Consumer Protection, and Cross-Cutting Statutes	15
The Federal Trade Commission Act, Section 5 (UDAP).....	15
The Health Insurance Portability and Accountability Act (HIPAA) and the HITECH Act	17
The Gramm-Leach-Bliley Act (GLBA) and the FTC Safeguards Rule.....	19
The Family Educational Rights and Privacy Act (FERPA).....	21
The Children's Online Privacy Protection Act (COPPA).....	22
The Fair Credit Reporting Act (FCRA)	24
The Equal Credit Opportunity Act (ECOA) and Regulation B.....	25
The Telephone Consumer Protection Act (TCPA).....	27
The Computer Fraud and Abuse Act (CFAA).....	28
The Stored Communications Act (SCA) and the Wiretap Act.....	30
The Privacy Act of 1974	31
The Federal Information Security Modernization Act (FISMA) and FedRAMP.....	33
Executive Order 14179 and the Federal AI Action Plan.....	34
Communications, Intermediary, and Surveillance Laws.....	35
Video Privacy Protection Act (VPPA)	36
Communications Decency Act Section 230	37
CLOUD Act and Cross-Border Government Access	38
EU–U.S. Data Privacy Framework (DPF) and Successor Mechanisms	39
Foreign Intelligence Surveillance Act (FISA) Section 702 and U.S. Surveillance Reform	40
Electronic Communications Privacy Act (ECPA) - Comprehensive	41
CAN-SPAM Act and Commercial Electronic Communications.....	42
Driver's Privacy Protection Act (DPPA) and Other Federal Sectoral Privacy Statutes	43
Civil Rights and Employment Laws Applied to AI	43

Title VII of the Civil Rights Act and AI in Employment44

Age Discrimination in Employment Act (ADEA) and AI45

Genetic Information Nondiscrimination Act (GINA).....46

Office of Federal Contract Compliance Programs (OFCCP) and AI in Federal Contractor Employment47

Fair Housing Act (FHA) and HUD AI Guidance48

Americans with Disabilities Act (ADA) Title III and AI Accessibility.....49

National Labor Relations Act (NLRA) and AI in the Workplace.....50

Workforce Surveillance Statutes (State Survey).....51

Critical Infrastructure and Sector-Specific Federal Frameworks51

 Cybersecurity and Infrastructure Security Agency (CISA) Act and Critical Infrastructure AI Guidance52

 NERC Critical Infrastructure Protection (CIP) Standards53

 TSA Pipeline and Rail Cybersecurity Directives54

 Nuclear Regulatory Commission (NRC) Cybersecurity and AI Frameworks55

 FAA AI in Aviation and Aviation Cybersecurity.....56

 Federal Energy Regulatory Commission (FERC) Order 901 and Critical Infrastructure56

 Department of Defense Responsible AI Strategy and Procurement57

Sector-Specific Federal Regulators on AI.....58

 FDA AI/ML-Based Software as a Medical Device (SaMD) Framework58

 Federal Reserve SR 11-7 - Model Risk Management.....60

 CFPB Adverse Action and AI - Circulars 2022-03 and 2023-0361

 EEOC AI in Employment Guidance and ADA Title I62

 Section 1557 of the Affordable Care Act and AI Discrimination.....63

 NAIC Model Bulletin on Use of Artificial Intelligence Systems64

 SEC Predictive Data Analytics Proposal and Existing Disclosure Obligations65

P A R T I I67

Comprehensive State Privacy Laws - California, Virginia, and the First Wave68

 California Consumer Privacy Act (CCPA) and California Privacy Rights Act (CPRA)68

 Virginia Consumer Data Protection Act (VCDPA)70

 Colorado Privacy Act (CPA).....71

 Connecticut Data Privacy Act (CTDPA)73

 Utah Consumer Privacy Act (UCPA).....74

 Texas Data Privacy and Security Act (TDPSA)75

 Oregon Consumer Privacy Act (OCPA)76

Texas Capture or Use of Biometric Identifier Act (CUBI)..... 77

Illinois Biometric Information Privacy Act (BIPA) 78

Washington My Health My Data Act (MHMDA)..... 79

New York SHIELD Act and NY DFS 23 NYCRR Part 500 80

Comprehensive State Privacy Laws - The Second and Third Waves..... 82

 Iowa Consumer Data Protection Act (ICDPA) 83

 Indiana Consumer Data Protection Act (INCDPA) 83

 Tennessee Information Protection Act (TIPA) 84

 Montana Consumer Data Privacy Act (MCDPA) 85

 New Jersey Data Privacy Act 86

 Delaware Personal Data Privacy Act..... 87

 New Hampshire Data Privacy Act..... 88

 Kentucky Consumer Data Protection Act 89

 Maryland Online Data Privacy Act (MODPA)..... 90

 Minnesota Consumer Data Privacy Act (MNCDDPA)..... 91

 Rhode Island Data Transparency and Privacy Protection Act 92

 Nebraska Data Privacy Act 93

 Florida Digital Bill of Rights..... 94

State Cybersecurity, Breach Notification, and Sector Privacy Regimes 95

 U.S. State Data Breach Notification Laws (Survey)..... 95

 Washington Biometric Identifier Statute (RCW 19.375) 96

 New York Biometric and Student Data Laws..... 97

 Colorado Student Data Transparency and Security Act 98

 California Student Online Personal Information Protection Act (SOPIPA) 99

 New York Stop Hacks and Improve Electronic Data Security (SHIELD) - Detailed 100

 Massachusetts 201 CMR 17 - Standards for the Protection of Personal Information..... 101

P A R T I I I 103

Leading State AI Statutes..... 104

 Colorado AI Act (CAIA, SB 24-205)..... 104

 Texas Responsible Artificial Intelligence Governance Act (TRAIGA) 106

 California Generative AI Training Data Transparency Act (AB 2013) 107

 California AI Provenance Act (SB 942) and California AI Transparency Act..... 108

 California AI in Healthcare Disclosure Act (AB 3030) 109

 New York City Local Law 144 - Automated Employment Decision Tools (AEDT) 109

 Illinois Artificial Intelligence Video Interview Act..... 111

Illinois HB 3773 - AI in Employment Decisions 111

Tennessee Ensuring Likeness, Voice, and Image Security Act (ELVIS Act)..... 112

State Deepfake and AI-Generated Content Laws (Survey)..... 113

Additional State AI Statutes, Frontier Regulation, and Sectoral AI Laws 114

Utah Artificial Intelligence Policy Act (UAIP) 114

California SB 1047 - Safe and Secure Innovation for Frontier AI Models Act (Vetoed) 115

California SB 53 - Transparency in Frontier AI Act..... 116

California Generative AI Accountability Act (AB 2885) and Related California AI Statutes..... 117

Massachusetts Information Privacy and Security Act (Pending) and Massachusetts AI Initiatives.. 118

Connecticut SB 2 - AI Consequential Decision Bill (Pending) and Other Connecticut AI Initiatives . 119

New York Local Law 144 - AEDT Detailed Requirements and NY State AI Bills 119

Oklahoma AI in Healthcare Act and Related State Sectoral AI Laws 120

State Election AI Deepfake Laws (Survey) 121

State Non-Consensual Intimate Imagery (NCII) and Synthetic Sexual Content Laws 122

P A R T I V 124

IP Frameworks for the AI Lifecycle 125

U.S. Copyright Office Guidance on AI-Generated Works 125

NYT v. OpenAI and the Generative AI Training Litigation Landscape 126

Thaler v. Perlmutter and AI Inventorship/Authorship Doctrine 127

USPTO Inventorship Guidance for AI-Assisted Inventions 128

EU Copyright Directive Article 4 (Text and Data Mining Exception)..... 129

UK Text and Data Mining and AI Copyright Reform 130

Japan Copyright Act Article 30-4 - TDM Exception 131

Trade Secret Protection of AI Models and Training Data 132

Open Source AI Licensing - OSI Definition, OpenRAIL, MIT/Apache, and Custom Licenses 133

Right of Publicity, NO FAKES Act, and AI Voice/Likeness Protection..... 134

P A R T V 136

Core EU AI, Privacy, and Digital Regulation 137

EU AI Act (Regulation (EU) 2024/1689) 137

General Data Protection Regulation (GDPR, Regulation (EU) 2016/679)..... 139

EU Digital Services Act (DSA) 141

EU NIS2 Directive..... 143

EU Digital Operational Resilience Act (DORA) 144

EU Data Act (Regulation (EU) 2023/2854) 145

ePrivacy Directive (Cookie Law) 146

EU Digital Markets Act (DMA) 147

Additional EU Regulation - CRA, PLD, eIDAS 2.0, EHDS, and the GPAI Code of Practice 148

EU Cyber Resilience Act (CRA) 148

EU AI Liability Directive (Pending) and Revised Product Liability Directive 149

eIDAS 2.0 and the European Digital Identity Wallet 150

EU Cybersecurity Act and ENISA 151

EU Data Governance Act 152

EU Network and Information Security (NIS2) Sectoral Implementing Acts 153

EU AI Act Code of Practice for General-Purpose AI 154

EU Health Data Space (EHDS) Regulation 155

P A R T V I 157

United Kingdom and Major American Regimes 158

UK GDPR and the Data Protection Act 2018 158

UK Sectoral AI Regulation and AI Safety Institute 159

UK Online Safety Act 160

Canada - PIPEDA and the Pending Consumer Privacy Protection Act 161

Brazil - Lei Geral de Proteção de Dados (LGPD) and Marco Legal da Inteligência Artificial 162

Mexico - Federal Law on Protection of Personal Data Held by Private Parties 163

Argentina - Personal Data Protection Law and Pending Reform 164

Additional Latin American Regimes 165

Chile - Personal Data Protection Law 165

Colombia - Personal Data Protection Law 166

Peru - Personal Data Protection Law 167

Uruguay - Personal Data Protection Law 167

Costa Rica - Protection of the Person against the Treatment of Personal Data 168

Dominican Republic - Personal Data Protection Law 169

Panama - Personal Data Protection Law 170

Ecuador - Organic Law on Protection of Personal Data 170

P A R T V I I 172

Leading Asia-Pacific and MENA Regimes 173

China - Personal Information Protection Law (PIPL) 173

China - Generative AI Measures, Deep Synthesis Provisions, and Algorithmic Recommendation Provisions 174

Japan - Act on the Protection of Personal Information (APPI) 176

South Korea - Personal Information Protection Act (PIPA) and AI Basic Act 177

Singapore - Personal Data Protection Act (PDPA) and Model AI Governance Framework 178

India - Digital Personal Data Protection Act (DPDP Act)..... 179

Australia - Privacy Act 1988 and AI Ethics Framework 180

Israel - Privacy Protection Law (PPL) and AI Policy 181

United Arab Emirates - Federal Personal Data Protection Law and AI Strategy 182

Saudi Arabia - Personal Data Protection Law (PDPL)..... 183

South Africa - Protection of Personal Information Act (POPIA) 184

Additional Asia-Pacific Regimes 185

 Vietnam - Personal Data Protection Decree and Cybersecurity Law 185

 Indonesia - Personal Data Protection Law (PDP Law) 186

 Thailand - Personal Data Protection Act (PDPA) 187

 Philippines - Data Privacy Act 188

 Malaysia - Personal Data Protection Act (PDPA)..... 188

 Hong Kong - Personal Data (Privacy) Ordinance (PDPO) 189

 Taiwan - Personal Data Protection Act..... 190

 New Zealand - Privacy Act 2020 191

Deeper MENA and African Regimes 192

 Bahrain - Personal Data Protection Law..... 192

 Qatar - Personal Data Privacy Protection Law..... 193

 Oman - Personal Data Protection Law 194

 Kuwait - Data Privacy Protection Regulation..... 194

 Egypt - Personal Data Protection Law 195

 Nigeria - Nigeria Data Protection Act (NDPA)..... 196

 Kenya - Data Protection Act 197

 Ghana - Data Protection Act..... 198

 Morocco - Data Protection Law..... 198

 Tunisia - Personal Data Protection Law 199

 Turkey - Personal Data Protection Law (KVKK)..... 200

P A R T V I I I 202

AI-Specific International Frameworks 203

 Council of Europe Framework Convention on Artificial Intelligence..... 203

 OECD AI Principles and OECD Recommendation on AI 204

 UNESCO Recommendation on the Ethics of Artificial Intelligence..... 204

 NIST AI Risk Management Framework (AI RMF 1.0) and Generative AI Profile 205

 ISO/IEC 42001 - AI Management Systems..... 206

ISO/IEC 23894 - AI Risk Management Guidance 207

MITRE ATLAS - Adversarial Threat Landscape for AI Systems 208

OWASP LLM Top 10 and ML Top 10 209

Foundational Security and Management System Standards 210

ISO/IEC 27001 - Information Security Management Systems 210

ISO/IEC 27701 - Privacy Information Management Systems 211

SOC 2 - Service Organization Controls..... 212

NIST SP 800-53 - Security and Privacy Controls..... 213

Cloud Security Alliance (CSA) Cloud Controls Matrix and AI Controls..... 214

IEEE 7000 Series - Ethics-Driven Standards 214

CIS Controls and CIS Benchmarks 215

ENISA AI Threat Landscape and Multilateral AI Cybersecurity Frameworks 216

PCI DSS and AI in Payment Card Processing..... 217

P A R T I X 219

Maximum Civil Penalties - Comparative View..... 220

Breach Notification Timelines 221

Automated Decision-Making - Cross-Jurisdictional Comparison..... 221

Stack Layer Cross-Reference - Where to Find Each Layer's Heaviest Hitters 222

P A R T X 224

Closing Note 230

P A R T I

United States - Federal Law

Federal statutes, agency rules, executive instruments, civil rights frameworks, and critical-infrastructure regimes that shape how AI systems handling personal information, making consequential decisions, or operating in regulated markets are developed, deployed, and supervised across the United States.

Privacy, Consumer Protection, and Cross-Cutting Statutes

Even without comprehensive federal privacy legislation, the United States has assembled a powerful, layered federal regime governing AI through consumer protection, civil rights, sector-specific privacy, communications, computer-misuse, and procurement law. The statutes in this chapter touch every meaningful U.S. AI deployment and are the de facto federal AI regulatory baseline.

The Federal Trade Commission Act, Section 5 (UDAP)

15 U.S.C. § 45

Jurisdiction	United States - Federal
Effective	1914 (continuously amended)
Regulator	Federal Trade Commission (FTC)
Scope	All persons, partnerships, and corporations engaged in commerce, with limited carve-outs (banks, common carriers, certain nonprofits)

Applicability

Section 5 prohibits "unfair or deceptive acts or practices in or affecting commerce." It is the FTC's most powerful and most flexible tool, and the de facto federal AI and privacy enforcement statute in the absence of comprehensive federal privacy legislation. The Commission has used Section 5 to police algorithmic discrimination, biased model outputs, deceptive AI marketing claims, undisclosed automated decisioning, dark patterns, faulty data security practices, and the use of training data obtained through deception or in violation of representations made to consumers.

Core Obligations

- Avoid deceptive claims about AI capabilities, accuracy, fairness, training data, or human review (the FTC has explicitly warned against "AI-washing").
- Substantiate marketing claims about AI performance with competent and reliable evidence before disseminating them.
- Implement reasonable data security commensurate with the sensitivity of data processed and the risk of harm.
- Honor representations made in privacy notices, terms of service, and consumer-facing statements about model behavior.
- Be prepared for "algorithmic disgorgement" - the destruction of models, algorithms, and derivative work products built using improperly obtained data (e.g., Everalbum, Cambridge Analytica, WW International).

Penalties & Enforcement

No general civil penalty for first-time Section 5 violations, but consent orders impose 20-year compliance regimes, monetary redress, and structural relief. Civil penalties up to \$53,088 per violation (2026 adjustment) for violations of consent orders or trade-regulation rules. Each affected consumer or each non-compliant transaction can constitute a separate violation, creating exponential liability exposure in large-scale AI deployments where a single deceptive claim may touch millions of users. State AG concurrent enforcement under "little FTC Acts" and parallel class action litigation under unfair competition statutes routinely follow Commission action, multiplying exposure beyond the federal penalty itself. The FTC's 2024 "Operation AI Comply" sweep and 2025–2026 enforcement priorities place AI deception, biometric misuse, and children's privacy at the top of the docket.

Recent Developments (through 2026)

In 2024–2025, the FTC issued the Negative Option Rule (click-to-cancel), expanded the Health Breach Notification Rule to non-HIPAA health apps and AI wellness tools, and finalized the COPPA Rule update tightening training-data uses involving children. The 2026 enforcement agenda emphasizes generative AI deception, AI-enabled impersonation, and unfair biometric data practices.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training-data provenance and consent representations are directly enforceable as Section 5 obligations; deceptive collection or scraping triggers algorithmic disgorgement.
Layer 2 - Model Governance	Performance, fairness, and accuracy claims about models must be substantiated; un-tested fairness assertions are deceptive.
Layer 4 - Control & Monitoring	Inadequate monitoring that allows known model drift or biased outputs to persist can be both unfair and deceptive.
Layer 5 - Audit & Evidence	Consent orders mandate detailed recordkeeping, third-party assessments, and 20-year reporting cadences - Layer 5 evidence machinery becomes a regulatory artifact.

PRACTITIONER NOTE
 Treat every public statement about an AI system as a Section 5 representation. The FTC reads marketing pages, model cards, and sales decks as evidentiary exhibits. Misalignment between external claims and internal evaluation results is one of the fastest paths to enforcement action and the precondition for algorithmic disgorgement orders.

COMMON FAILURE PATTERN
 Marketing teams publish AI accuracy or fairness claims drafted before the model is finalized; product teams quietly retrain or scope-narrow the system without updating the public statements. The FTC then aligns the original representation against the deployed system in discovery, and the resulting deviation supports both

deception and unfairness theories simultaneously. Build a representation-review gate that requires sign-off by the engineering owner of the underlying evaluation before any external claim ships.

The Health Insurance Portability and Accountability Act (HIPAA) and the HITECH Act

42 U.S.C. §§ 1320d et seq.; 45 C.F.R. Parts 160, 162, 164

Jurisdiction	United States - Federal
Effective	HIPAA 1996; Privacy Rule 2003; Security Rule 2005; HITECH 2009; Omnibus Rule 2013; proposed Security Rule update 2025
Regulator	HHS Office for Civil Rights (OCR); CMS; state attorneys general (concurrent under HITECH)
Scope	Covered Entities (health plans, healthcare clearinghouses, healthcare providers conducting covered transactions) and their Business Associates and Subcontractors

Applicability

HIPAA governs the use and disclosure of Protected Health Information (PHI), defined as individually identifiable health information held or transmitted by a Covered Entity or Business Associate. The Privacy Rule restricts uses and disclosures; the Security Rule mandates administrative, physical, and technical safeguards for electronic PHI (ePHI); the Breach Notification Rule requires notification to individuals, HHS, and (for breaches affecting 500+ individuals) the media. AI systems processing PHI for clinical decision support, ambient documentation, predictive analytics, claims adjudication, or population health are all in scope.

Core Obligations

- Execute Business Associate Agreements (BAAs) before any AI vendor receives PHI; the BAA must flow down to Subcontractors, including model-hosting providers and inference-as-a-service platforms.
- Apply the minimum necessary standard to AI training datasets - PHI used for model development should be limited to what is reasonably required.
- Conduct risk analyses under § 164.308(a)(1)(ii)(A) for each AI system processing ePHI; the proposed 2025 Security Rule update would mandate annual risk analyses, asset inventories including AI components, and explicit encryption of ePHI at rest and in transit (eliminating the current "addressable" flexibility).
- Honor individual rights of access, amendment, accounting of disclosures, and (where applicable) restriction; AI-generated notes and predictions in the designated record set are subject to access requests.

- Notify affected individuals within 60 days of breach discovery; Section 1557 of the ACA layers a non-discrimination obligation on patient-care decision-support tools (effective 2024).

Penalties & Enforcement

Civil monetary penalties tiered by culpability: \$137 to \$2,067,813 per violation per calendar year (2026 inflation-adjusted). Criminal penalties up to 10 years imprisonment for knowing wrongful disclosures for personal gain or malicious harm. State AGs may pursue civil suits under HITECH § 13410(e). OCR resolution agreements increasingly require AI vendor inventories and BAA audits.

Recent Developments (through 2026)

The HHS Notice of Proposed Rulemaking (December 2024) substantially modernizes the Security Rule, removing the addressable/required distinction, mandating multi-factor authentication, requiring written documentation of all security policies, and explicitly anticipating AI workloads. Final rule expected in 2026. Section 1557 final rule (2024) prohibits algorithmic discrimination in patient-care decision-support tools and requires Covered Entities to make reasonable efforts to identify and mitigate the risk of discrimination from such tools.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	PHI provenance, minimum-necessary scoping, and de-identification (Safe Harbor or Expert Determination under § 164.514) are core Layer 1 controls. Training-set lineage must trace each PHI element back to a permissible use.
Layer 2 - Model Governance	Model documentation supports the "reasonable efforts" standard under Section 1557; fairness and subgroup testing on protected characteristics is the practical compliance vehicle.
Layer 3 - System Integration	BAAs and Subcontractor agreements govern Layer 3 integration boundaries; data-flow diagrams must mirror the BAA chain.
Layer 4 - Control & Monitoring	Access controls, audit logs, and integrity monitoring under § 164.312 apply unchanged to AI inference services and prompt logs.
Layer 5 - Audit & Evidence	Six-year retention of HIPAA documentation aligns naturally with Layer 5 evidence preservation; OCR investigations rely on this trail.

PRACTITIONER NOTE
 Audit the full prompt path before any clinical AI deployment: retrieval-augmented context, system prompts, telemetry, and logging endpoints must each have documented BAA coverage. A single uncovered subprocessor in the chain converts a clinical workflow into a reportable breach.

COMMON FAILURE PATTERN
 Vendors market "HIPAA-compliant" AI APIs while routing prompts and outputs through model providers, fine-tuning platforms, or evaluation services that lack BAAs. The Covered Entity inherits liability for every prompt

that contained PHI, every output that reached the wrong recipient, and every log entry retained beyond the BAA chain. OCR resolution agreements increasingly itemize each AI subprocessor; what was sold as a single vendor relationship becomes a multi-party breach disclosure.

The Gramm-Leach-Bliley Act (GLBA) and the FTC Safeguards Rule

15 U.S.C. §§ 6801–6809; 16 C.F.R. Part 314

Jurisdiction	United States - Federal
Effective	GLBA 1999; Safeguards Rule 2003; revised Safeguards Rule 2021 (effective 2023); breach notification amendment effective May 13, 2024
Regulator	FTC (non-bank financial institutions); federal banking agencies (banks); SEC (broker-dealers, investment advisers); state insurance regulators
Scope	Financial institutions as defined under § 509(3) - includes lenders, mortgage brokers, debt collectors, tax preparers, check cashers, payday lenders, certain finders, and a growing list of fintech actors

Applicability

GLBA imposes three principal regimes: the Privacy Rule (notices and opt-outs for sharing nonpublic personal information with non-affiliated third parties), the Safeguards Rule (administrative, technical, and physical safeguards), and the Pretexting Rule. AI systems used for credit underwriting, fraud detection, KYC/AML, robo-advisory, claims processing, and customer-service automation operate squarely within Safeguards Rule scope.

Core Obligations

- Designate a Qualified Individual responsible for the information security program.
- Conduct a written risk assessment that explicitly addresses AI/ML systems and is updated periodically.
- Implement and maintain a written information security program with eight enumerated elements, including access controls, encryption of customer information at rest and in transit, MFA for any individual accessing customer information, secure disposal, change management, and continuous monitoring or annual penetration testing plus biannual vulnerability assessments.
- Oversee service providers - selection due diligence, contractually required safeguards, and periodic reassessment.
- Notify the FTC within 30 days of discovery of any security event involving the unauthorized acquisition of unencrypted customer information of 500 or more consumers (effective May 2024).
- Provide initial and annual privacy notices and respect opt-out elections under the Privacy Rule.

Penalties & Enforcement

No statutory cap; FTC may seek injunctive relief and consumer redress. Banking-agency enforcement includes cease-and-desist orders, civil money penalties, and management removal. Recent FTC orders against mortgage and fintech firms have included multimillion-dollar penalties and 20-year compliance obligations.

Recent Developments (through 2026)

FTC enforcement against TaxSlayer, Drizly, and Cafepress established that the Safeguards Rule reaches AI vendors and that personal liability for executives is on the table where the security program is materially deficient. The 2024 breach-notification amendment created a public-facing breach database that competitors and plaintiffs' counsel monitor in real time.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Customer information used to train fraud, credit, or KYC models is "customer information" under § 314.2 and inherits encryption, access, and disposal obligations.
Layer 2 - Model Governance	Model risk management practices must align with prudential expectations (SR 11-7 for banks, NAIC Model Bulletin for insurers); GLBA does not prescribe model governance but the risk assessment must consider model-specific risks.
Layer 3 - System Integration	Service-provider oversight reaches AI hosting and inference platforms; flow-down provisions, audit rights, and right-to-terminate-on-breach are standard contract terms.
Layer 4 - Control & Monitoring	MFA and continuous monitoring requirements apply to AI development environments, not just production data stores - an under-appreciated audit finding.
Layer 5 - Audit & Evidence	30-day FTC breach notification requires Layer 5 incident-detection telemetry capable of confirming "unauthorized acquisition" - not merely access - within tight windows.

PRACTITIONER NOTE
 Map every AI system that touches customer information against the eight Safeguards Rule elements. Any element marked "addressed at the platform level" must be supported by a documented review of the platform's controls and the contractual basis for relying on them. Unverified platform reliance is the most-cited finding in 2024-2025 FTC Safeguards Rule consent orders.

COMMON FAILURE PATTERN
 AI development environments are excluded from the formal information security program because they sit outside the production data plane. The Qualified Individual signs the annual certification, MFA is on production, and adversarial testing is run on the inference endpoint, but the model training environment, the experiment tracking platform, and the prompt-evaluation tooling all hold customer information without parallel controls. The 30-day breach notification clock starts when any of those side environments is compromised, not just production.

The Family Educational Rights and Privacy Act (FERPA)

20 U.S.C. § 1232g; 34 C.F.R. Part 99

Jurisdiction	United States - Federal
Effective	1974 (continuously amended)
Regulator	U.S. Department of Education, Student Privacy Policy Office (SPPO)
Scope	Educational agencies and institutions receiving funds from a program administered by the U.S. Department of Education

Applicability

FERPA protects the privacy of education records and personally identifiable information (PII) derived from them. Schools may not disclose education records to third parties without parental consent (or student consent for those 18+ or in postsecondary institutions), subject to enumerated exceptions including the "school official" exception for vendors performing institutional services. EdTech vendors deploying AI tutoring, predictive enrollment, plagiarism detection, or behavioral analytics generally rely on this exception and inherit its conditions.

Core Obligations

- Limit disclosures of education records to parties with legitimate educational interests under direct institutional control.
- Designate "directory information" carefully and provide opt-out - directory information is the only category that may be disclosed without consent.
- For EdTech vendors operating under the school official exception: act under direct control, use education records only for authorized purposes, and refrain from re-disclosure or secondary use.
- Provide parents and eligible students rights to inspect, review, and seek amendment of education records - including AI-generated risk scores and predictive labels in the educational record.
- Maintain records of disclosures and provide them on request.

Penalties & Enforcement

FERPA contains no private right of action and no statutory civil penalties for institutions; the ultimate sanction is loss of federal education funding. In practice, ED resolves matters through Compliance Determinations and corrective action plans. State student-privacy statutes (notably California SOPIPA, New York Education Law § 2-d, and Colorado HB 16-1423) provide independent enforcement and private remedies.

Recent Developments (through 2026)

SPPO's 2023–2025 guidance addresses generative AI tools, prompting the proposition that prompt logs and chat transcripts containing PII may themselves constitute education records. The proposed FERPA

modernization rulemaking (anticipated 2026) is expected to clarify AI vendor obligations and tighten the school-official exception.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Student PII used for AI training is generally impermissible under FERPA absent consent; the school-official exception does not authorize commercial model improvement.
Layer 2 - Model Governance	AI-generated predictions about student performance, behavior, or risk inserted into the educational record are subject to access and amendment rights.
Layer 3 - System Integration	The "direct control" requirement constrains vendor architecture: prompt routing through subprocessors, telemetry, and logging endpoints all require institutional oversight.
Layer 5 - Audit & Evidence	Disclosure recordkeeping under § 99.32 is a Layer 5 obligation; AI inference calls disclosing PII to vendors should be logged at the same fidelity as traditional record disclosures.

PRACTITIONER NOTE
 When negotiating EdTech contracts, insist on contractual prohibitions against using student data to train models that benefit other customers. The "improve service" exception is routinely abused; FERPA requires authorized purpose use only.

The Children's Online Privacy Protection Act (COPPA)

15 U.S.C. §§ 6501–6506; 16 C.F.R. Part 312

Jurisdiction	United States - Federal
Effective	1998; Rule revised 2013; major Rule revision finalized April 2025 (effective June 23, 2025; full compliance April 22, 2026)
Regulator	Federal Trade Commission; state attorneys general (concurrent)
Scope	Operators of websites or online services directed to children under 13, and operators with actual knowledge of collecting personal information from children under 13

Applicability

COPPA conditions the collection, use, and disclosure of personal information from children under 13 on verifiable parental consent. The 2025 Rule revision substantially expands the definition of personal information, restricts secondary use of children's data for AI training, requires written information security programs, mandates separate consent for certain disclosures, and adds biometric identifiers and inferences derived therefrom to the protected categories.

Core Obligations

- Provide direct notice to parents and obtain verifiable parental consent before any collection, use, or disclosure of personal information from children under 13.
- Limit data collection to what is reasonably necessary for the activity; do not condition participation on the disclosure of more information than necessary.
- Obtain separate, opt-in consent for disclosures of personal information to third parties unrelated to the operator's internal operations (2025 Rule).
- Obtain separate, opt-in consent before using personal information for behavioral advertising or to develop or train AI/ML systems for purposes beyond providing the service requested by the parent (2025 Rule).
- Establish, implement, and maintain a written information security program with reasonable safeguards.
- Honor data deletion rights and limit data retention to what is reasonably necessary.

Penalties & Enforcement

Civil penalties up to \$53,088 per violation (2026 adjusted); each affected child may constitute a separate violation. Major settlements include TikTok (\$5.7M, 2019; \$92M class action 2021; pending 2024 DOJ action), YouTube/Google (\$170M, 2019), Epic Games (\$275M, 2022), and Microsoft Xbox (\$20M, 2023).

Recent Developments (through 2026)

The 2025 COPPA Rule revision is the most significant overhaul since 2013, explicitly addressing AI training, biometric data, and dark patterns in consent flows. Compliance deadline is April 22, 2026 - most operators are in active remediation.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Use of children's data for model training requires separate parental consent - an explicit Layer 1 data-collection gate.
Layer 3 - System Integration	Third-party SDK and ad-tech integrations transmitting children's data are direct-disclosure events triggering consent and recordkeeping.
Layer 4 - Control & Monitoring	Dark-pattern prohibitions in the 2025 Rule reach Layer 4 deployment of consent UX; A/B testing of consent flows requires governance review.
Layer 5 - Audit & Evidence	Operators must maintain consent records for as long as the personal information is retained; AI systems with long-lived training datasets create indefinite Layer 5 obligations.

PRACTITIONER NOTE

Run a "general audience" assessment under the multi-factor test: subject matter, visual content, language, age of models, presence of child celebrities, music, and the use of animated characters. The FTC has rejected "we are not directed to children" defenses where these factors point the other way.

The Fair Credit Reporting Act (FCRA)

15 U.S.C. § 1681 *et seq.*

Jurisdiction	United States - Federal
Effective	1970 (continuously amended)
Regulator	CFPB (primary); FTC; federal banking agencies; state attorneys general
Scope	Consumer reporting agencies, furnishers of information, and users of consumer reports

Applicability

FCRA governs consumer reports - communications bearing on creditworthiness, credit standing, character, general reputation, personal characteristics, or mode of living that are used or expected to be used for an enumerated permissible purpose. The CFPB has taken the position (Circular 2022-03 and subsequent guidance) that AI-driven background screening, tenant screening, employment screening, and certain marketing-by-AI tools are subject to FCRA obligations including accuracy, dispute handling, and adverse action notices that explain the specific reasons for the decision - not merely "the algorithm decided."

Core Obligations

- Maintain reasonable procedures to assure maximum possible accuracy of information in consumer reports.
- Provide adverse action notices that include the specific principal reasons for adverse action - boilerplate is insufficient even for ML models (CFPB Circular 2022-03).
- Investigate and resolve consumer disputes within 30 days; furnishers must conduct reasonable investigations.
- Limit use of consumer reports to permissible purposes enumerated in § 1681b.
- Comply with the Disposal Rule for consumer report information (16 C.F.R. Part 682).
- For "investigative consumer reports," provide additional disclosures.

Penalties & Enforcement

Statutory damages of \$100–\$1,000 per willful violation, plus actual damages, punitive damages, and attorneys' fees. Negligent violations: actual damages plus fees. CFPB enforcement: penalties up to \$1,362,567 per day (2026 adjusted) for knowing violations. The FCRA private right of action drives high-volume class action litigation.

Recent Developments (through 2026)

CFPB's 2023–2024 enforcement against TransUnion, Equifax, and several tenant-screening firms emphasized that AI-driven scoring outputs without underlying explainability cannot satisfy adverse

action requirements. The CFPB's 2024 proposed rule on data brokers (under § 1681e) would expand FCRA coverage to data brokers selling header information and behavioral profiles - substantial AI-training implications. Status of finalization uncertain under post-2025 administrative changes.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training data for credit, employment, or tenant decisioning models that includes consumer-report-equivalent information triggers FCRA accuracy duties.
Layer 2 - Model Governance	Model interpretability is a regulatory requirement, not a best practice - adverse action explanations must derive from the model's actual decision logic.
Layer 4 - Control & Monitoring	Dispute-handling workflows require Layer 4 monitoring and operational controls; 30-day reinvestigation timelines are not aspirational.
Layer 5 - Audit & Evidence	Reinvestigation files, dispute records, and adverse action documentation are Layer 5 evidence repositories; private litigation routinely tests their adequacy.

PRACTITIONER NOTE

If the model uses any feature derived from consumer-report data (credit, eviction, criminal, employment, prior-residence), assume FCRA applies. The "we don't pull credit" defense fails where derived features carry the same predictive signal. FCRA's private right of action and statutory damages produce class action exposure that dwarfs the regulator-driven penalty.

COMMON FAILURE PATTERN

A modeling team builds a non-credit risk score using features that predict creditworthiness (utility payment patterns, e-commerce return behavior, social signals). Counsel concludes FCRA does not apply because the source is not a CRA. Plaintiffs' counsel then plead the model output as a "consumer report" assembled from third-party data and used for an FCRA-permissible purpose. Resolving that question takes years of discovery; in the meantime the company is litigating a class action it could have foreclosed by treating the system as FCRA-covered from the start.

The Equal Credit Opportunity Act (ECOA) and Regulation B

15 U.S.C. § 1691 et seq.; 12 C.F.R. Part 1002

Jurisdiction	United States - Federal
Effective	1974 (continuously amended)
Regulator	CFPB; federal banking agencies; DOJ; HUD (overlapping)
Scope	Creditors, including any entity that regularly extends, renews, or continues credit

Applicability

ECOA prohibits discrimination in credit transactions on the basis of race, color, religion, national origin, sex (including sexual orientation and gender identity per CFPB interpretation), marital status, age, receipt of public assistance, or exercise of rights under the CCPA. CFPB Circular 2022-03 and subsequent guidance establish that creditors using AI/ML must (a) provide statements of specific reasons for adverse action that accurately reflect the model's decision, and (b) conduct fair-lending testing of model outputs including disparate impact analysis.

Core Obligations

- Provide written adverse action notices within 30 days, including the specific principal reasons for the adverse action - black-box justifications fail.
- Conduct fair-lending testing including disparate impact analysis on model outputs across protected classes.
- Maintain records of credit applications for 25 months (12 months for business credit).
- Limit collection of demographic information to authorized purposes (HMDA, monitoring).
- For mortgage credit, comply with HMDA reporting and Regulation B monitoring requirements.

Penalties & Enforcement

Actual and punitive damages (up to \$10,000 individual; lesser of \$500,000 or 1% of net worth for class actions). CFPB and federal banking agencies may seek civil money penalties, restitution, and injunctive relief. DOJ may pursue pattern-or-practice litigation.

Recent Developments (through 2026)

CFPB Circular 2023-03 reaffirmed that complex AI models do not exempt creditors from specific-reason adverse action requirements. The 2024 settlement between DOJ and a major fintech imposed model-validation obligations and a fair-lending monitor. Post-2025, the regulatory posture may shift but the statutory obligations remain unchanged.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Proxy variables for protected characteristics (ZIP code, name, school) are well-known fair-lending traps - Layer 1 bias assessment is essentially a fair-lending exercise.
Layer 2 - Model Governance	Model fairness testing across demographic subgroups is a regulatory requirement; the choice of fairness metric (disparate impact ratio, marginal effects, demographic parity) must be defensible.
Layer 3 - System Integration	Where AI scores are inputs to human decisioning, the system-integration boundary determines whether the AI is the "decision" or an "input"; both views carry ECOA exposure.
Layer 5 - Audit & Evidence	25-month application file retention is a baseline; fair-lending defense files (disparate impact analyses, business necessity justifications, less discriminatory alternatives reviews) are essential Layer 5 artifacts.

PRACTITIONER NOTE

The "less discriminatory alternative" doctrine in disparate impact analysis is the single most important fair-lending consideration for AI underwriting. Document the search for and rejection of alternative models or features that achieve comparable predictive power with less disparate impact prospectively, not after litigation. The contemporaneous search record is what survives summary judgment; the post-hoc rationalization does not.

COMMON FAILURE PATTERN

A modeling team selects the highest-AUC model and ships it. Fair-lending review is performed downstream as a check, finding a disparate impact disparity within "tolerance." When the regulator or a private plaintiff later asks for the search for less discriminatory alternatives, no record exists, because no formal search was performed. The disparate impact prima facie case shifts the burden, business necessity is asserted, and the LDA analysis is reconstructed under litigation pressure. Make the LDA search a gated step before model selection, with a decision memo signed at the model-approval stage.

The Telephone Consumer Protection Act (TCPA)

47 U.S.C. § 227; 47 C.F.R. § 64.1200

Jurisdiction	United States - Federal
Effective	1991; substantial 2024–2025 amendments and FCC orders affecting AI-generated voice
Regulator	Federal Communications Commission (FCC); state attorneys general; private right of action
Scope	Persons making calls or sending texts to U.S. numbers using automated or prerecorded technology

Applicability

TCPA restricts the use of automatic telephone dialing systems (ATDS), prerecorded and artificial voice messages, and certain text messages. The FCC's February 2024 declaratory ruling clarified that AI-generated voice calls are "artificial voice" under § 227(b), requiring prior express written consent for marketing calls regardless of whether the call is initiated by a human operator. The FCC opened a parallel Notice of Proposed Rulemaking on AI-disclosure obligations and consumer-facing transparency.

Core Obligations

- Obtain prior express written consent before placing marketing calls or texts using an ATDS or artificial/prerecorded voice to any wireless number.

- Disclose at the start of any AI-generated voice marketing call that the call uses an artificial voice, the identity of the caller, and the nature of the call (per FCC 2024 ruling and pending rulemaking).
- Honor do-not-call requests within a reasonable time, not exceeding 30 days.
- Maintain a written do-not-call policy and train personnel.
- For political and informational AI voice calls, comply with disclosure obligations even where consent is not required.

Penalties & Enforcement

Statutory damages of \$500 per negligent violation and \$1,500 per willful violation, with no statutory cap on aggregate damages. Class actions routinely seek hundreds of millions in damages. FCC may impose civil forfeiture penalties up to \$25,000 per violation.

Recent Developments (through 2026)

The 2024 FCC ruling on AI voice was triggered by AI-generated robocalls impersonating a sitting president before a primary election. State legislatures (including Michigan, Washington, and Texas) have layered additional AI-voice disclosure requirements. The 2026 FCC AI transparency rulemaking is the most-watched item on the consumer-protection docket.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Consent records are the entire compliance basis; a deficient consent log is per se liability.
Layer 3 - System Integration	AI voice integrations with telephony platforms must enforce real-time disclosure overlays and consent verification at the call boundary.
Layer 4 - Control & Monitoring	Real-time monitoring for opt-out and STOP texts is required; AI agent calling at scale magnifies the cost of latency in honoring withdrawals of consent.
Layer 5 - Audit & Evidence	TCPA defense files - consent records, call logs, IVR recordings, complaint handling - are Layer 5 audit artifacts that determine litigation outcomes.

PRACTITIONER NOTE

AI voice agents in customer service workflows can drift into "marketing" if they cross-sell, even if the inbound call was initiated by the consumer. Build content classifiers and call-flow gates that prevent marketing content where prior express written consent has not been documented.

The Computer Fraud and Abuse Act (CFAA)

18 U.S.C. § 1030

Jurisdiction	United States - Federal
Effective	1986 (continuously amended)

Regulator	U.S. Department of Justice; private civil right of action
Scope	Anyone who accesses a "protected computer" without authorization or in excess of authorized access

Applicability

CFAA criminalizes unauthorized access and certain access-with-intent activities involving protected computers (any computer used in or affecting interstate commerce). For AI systems, CFAA implications arise in three principal contexts: (1) training-data scraping and the post-Van Buren / hiQ v. LinkedIn doctrine governing access to publicly available data; (2) adversarial attacks against deployed models, where prompt injection or model extraction may constitute exceeding authorized access; and (3) red-teaming activities, which require careful authorization scoping.

Core Obligations

- Obtain explicit, documented authorization before red-teaming production AI systems owned by third parties.
- Avoid scraping content where access is restricted by technical or contractual measures targeted at the scraper (the post-Van Buren analysis remains highly fact-specific).
- Document the authorization scope for internal AI security testing - "in scope" must be defined in writing.
- For bug bounty programs and AI red-team engagements, use Safe Harbor language modeled on DOJ's 2022 CFAA prosecution policy.

Penalties & Enforcement

Misdemeanor and felony criminal penalties scaling with intent and damage; private civil remedies including damages and injunctive relief where the loss exceeds \$5,000 in any one-year period.

Recent Developments (through 2026)

Van Buren v. United States (2021) narrowed "exceeds authorized access" but left the contours of "without authorization" - particularly for scraping - unsettled. The Ninth Circuit's hiQ v. LinkedIn line of cases generally protects scraping of public data but warns against bypassing technical access controls. AI training-data acquisition remains a high-risk area.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training-data acquisition pipelines must include legal review of source authorization; "publicly available" is not legally synonymous with "permissible to scrape at scale."
Layer 2 - Model Governance	Adversarial robustness testing programs need authorization documentation; cross-team red-teaming should rely on the same scoping rigor as third-party engagements.

STACK LENS - How this law maps to the AI Governance Stack

Layer 3 - System Integration	API integration with third-party AI services should map terms-of-service restrictions onto operational guardrails.
Layer 5 - Audit & Evidence	Authorization records, ToS reviews, and red-team rules-of-engagement are Layer 5 evidence in any CFAA inquiry.

PRACTITIONER NOTE

A common audit finding: data-engineering teams scraping competitor content, social media posts, or public records to "augment" training datasets without legal review. Maintain a "data acquisition gate" - no new training source enters the pipeline without a written authorization assessment.

The Stored Communications Act (SCA) and the Wiretap Act

18 U.S.C. §§ 2510–2523 (*Wiretap*); §§ 2701–2713 (*SCA*)

Jurisdiction	United States - Federal
Effective	1986 (Electronic Communications Privacy Act)
Regulator	DOJ; private civil right of action
Scope	Providers of electronic communication services and remote computing services; persons intercepting electronic communications

Applicability

The Wiretap Act prohibits the interception of electronic communications in transit; the SCA restricts disclosure of stored communications by service providers. AI implications include: ambient and call-center recording; AI meeting assistants that capture, transcribe, and summarize conversations; SaaS providers that host customer communications and use them for model improvement; and the disclosure of customer communications to AI subprocessors.

Core Obligations

- Obtain consent (one-party or two-party depending on jurisdiction; many states require all-party consent) before recording or transcribing communications using AI tools.
- For service providers, restrict disclosure of customer communications to authorized recipients; do not use customer communications for AI training without explicit consent.
- Comply with judicial process for compelled disclosure (warrant, court order, or subpoena depending on data type).
- Implement technical controls preventing inadvertent capture by AI agents not authorized to monitor the conversation.

Penalties & Enforcement

Wiretap: criminal penalties and civil damages of the greater of actual damages, \$100/day per violation, or \$10,000 per violation, plus punitive damages and attorneys' fees. SCA: criminal penalties; civil damages of actual damages, \$1,000 per violation minimum, plus fees.

Recent Developments (through 2026)

A wave of class actions (2023–2025) against companies using AI chat-replay, session-recording, and meeting-assistant tools has pushed Wiretap Act analysis to the front of vendor diligence. Several federal courts have allowed Wiretap claims to proceed against companies whose customer-service chat tools route content to third-party analytics or AI providers without consent.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Communications captured by AI assistants enter Layer 1 datasets; consent and minimum-necessary scoping are upstream gates.
Layer 3 - System Integration	AI subprocessor flows for transcription, summarization, and analytics expose third-party disclosure surfaces; data-flow diagrams should explicitly map every recipient.
Layer 4 - Control & Monitoring	Operational controls preventing unauthorized AI-agent participation in protected communications (e.g., privileged legal calls) are Layer 4 deployment gates.
Layer 5 - Audit & Evidence	Consent records and disclosure logs are central Layer 5 artifacts in any SCA or Wiretap claim.

PRACTITIONER NOTE

Map every AI tool that processes communications - meeting assistants, call summarization, chat analytics, session replay - against state two-party consent regimes. The single most common failure is using a vendor's default consent banner that does not name the AI subprocessor.

The Privacy Act of 1974

5 U.S.C. § 552a

Jurisdiction	United States - Federal (federal agencies)
Effective	1974 (continuously amended)
Regulator	Office of Management and Budget; private civil right of action
Scope	Federal agencies maintaining systems of records on individuals

Applicability

The Privacy Act governs federal agency collection, maintenance, use, and dissemination of records about U.S. citizens and lawful permanent residents. The Act's provisions concerning matching programs, system-of-records notices (SORNs), Privacy Impact Assessments (PIAs) under the E-Government Act, and the prohibition on solely automated decisioning that produces adverse effects without human review apply directly to federal agency AI systems. OMB Memoranda M-24-10 and M-25-21 substantially extended these obligations.

Core Obligations

- Publish a SORN for any system of records about individuals; AI systems that retrieve records by personal identifier are covered.
- Conduct a PIA before procuring or developing AI systems that handle personally identifiable information (E-Government Act § 208).
- Provide individuals access to their records and the ability to seek correction.
- Honor the Privacy Act's no-disclosure principle except under enumerated exceptions.
- For federal AI under M-24-10 / M-25-21: designate a Chief AI Officer, maintain an AI use-case inventory, conduct impact assessments for rights- and safety-impacting AI, and implement minimum risk management practices.

Penalties & Enforcement

Civil damages of actual damages or \$1,000 minimum, attorneys' fees, and equitable relief; criminal misdemeanor for willful disclosures by agency employees.

Recent Developments (through 2026)

OMB M-24-10 (March 2024, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence") and its 2025 successor establish a federal AI governance baseline that aligns closely with the AI Governance Stack. The 2025 administrative changes have produced revised guidance with lighter pre-deployment assessment requirements but maintained the core inventory and impact-assessment obligations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	SORN coverage and routine-use authorization are Layer 1 gates for federal AI training data.
Layer 2 - Model Governance	PIA requirements drive Layer 2 documentation including model purpose, data flows, and risk mitigations.
Layer 3 - System Integration	Inter-agency data-matching programs require Computer Matching and Privacy Protection Act compliance - a Layer 3 integration constraint.
Layer 5 - Audit & Evidence	AI use-case inventories under M-24-10 are foundational Layer 5 artifacts for federal agencies and their contractors.

PRACTITIONER NOTE

Federal contractors building AI for agency clients should treat SORN and PIA artifacts as deliverables, not afterthoughts. Procurement officials increasingly expect contractors to draft the privacy documentation as part of the proposal package.

The Federal Information Security Modernization Act (FISMA) and FedRAMP

44 U.S.C. § 3551 et seq.; FedRAMP authorities under OMB Memoranda

Jurisdiction	United States - Federal
Effective	FISMA 2002 (modernized 2014); FedRAMP 2011
Regulator	OMB; CISA; NIST; agency CIOs
Scope	Federal agencies and federal contractors operating information systems on behalf of the government; cloud service providers serving federal customers (FedRAMP)

Applicability

FISMA establishes the framework for protecting federal information systems through risk-based controls implementing NIST SP 800-53. FedRAMP applies these controls to commercial cloud services used by federal agencies and is the principal compliance gateway for AI vendors selling to government. FedRAMP's 2024 modernization (OMB M-24-15) and the 2025 FedRAMP Authorization Act streamline reciprocity but tighten supply-chain and AI-specific control expectations.

Core Obligations

- Categorize information systems per FIPS 199 and apply NIST SP 800-53 controls at the appropriate baseline.
- For federal AI systems and AI-as-a-service offerings, comply with the NIST AI Risk Management Framework and CISA's 2024 secure-by-design AI guidance as overlays to the standard control catalog.
- Achieve and maintain FedRAMP authorization (Low, Moderate, or High baseline) for cloud services hosting federal data.
- Submit continuous monitoring reports and material-change notifications to authorizing officials.
- For high-impact systems and critical software (per EO 14028), satisfy SBOM, vulnerability disclosure, and provenance requirements.

Penalties & Enforcement

Loss of authorization to operate (ATO) and contractual remedies. False Claims Act exposure where FISMA or FedRAMP attestations are materially false (as in the SAIC, Aerojet, and Penn State cases).

Recent Developments (through 2026)

The 2025 FedRAMP modernization permits agency-led authorizations and accelerated reciprocity. The 2024 NIST IR 8470 and the 2025 CISA "Secure AI System Development" guidance establish federal baseline expectations for AI components. Generative AI providers serving federal customers have been working through extensive technical questionnaires beyond the standard FedRAMP package.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Federal data classification (CUI categories) constrains training-data eligibility and inference-time disclosures.
Layer 2 - Model Governance	Model documentation requirements piggyback on the NIST 800-53 baseline's system documentation controls (CM, SA families).
Layer 3 - System Integration	Boundary diagrams, interconnection security agreements, and supply-chain risk management (SR family) directly govern Layer 3 architecture.
Layer 4 - Control & Monitoring	Continuous monitoring (CA-7) is Layer 4 by another name; AI-specific KPIs map onto the existing CONMON program.
Layer 5 - Audit & Evidence	Authorization packages are Layer 5 evidence at industrial scale; AI overlays add model cards, impact assessments, and red-team reports.

PRACTITIONER NOTE

Treat FedRAMP as a strategic floor, not a ceiling. The marginal cost of meeting commercial frameworks (SOC 2, ISO 27001, ISO 42001) on top of a FedRAMP Moderate authorization is small, and customers across both federal and commercial segments increasingly expect both.

Executive Order 14179 and the Federal AI Action Plan

Executive Order 14179 (January 2025); subsequent agency guidance

Jurisdiction	United States - Federal Executive Branch
Effective	January 2025; ongoing implementation through 2026
Regulator	White House; OMB; agency Chief AI Officers; NIST
Scope	Federal agencies; federal procurement of AI systems; voluntary frameworks influencing private sector

Applicability

EO 14179 ("Removing Barriers to American Leadership in Artificial Intelligence") rescinded EO 14110 (2023) and reframed federal AI policy around competitiveness and deregulation. The companion America's AI Action Plan and updated OMB Memoranda M-25-21 and M-25-22 modify, but do not eliminate, the federal AI governance baseline. Practitioners should track this regime alongside, not in

place of, the underlying statutory authorities (Privacy Act, FISMA, sector-specific laws) which remain in force.

Core Obligations

- Federal agencies must maintain AI use-case inventories and designate Chief AI Officers (carried over from prior administration).
- Procurement of AI systems classified as rights- or safety-impacting requires risk management practices including testing for performance, fairness considerations, and security.
- Federal contractors should anticipate evolving solicitation requirements referencing NIST AI RMF and CISA secure-AI guidance.
- NIST is directed to revise the AI Risk Management Framework to reflect updated federal priorities; the 2026 RMF revision is in progress.

Penalties & Enforcement

Executive Orders themselves do not establish private rights of action; enforcement runs through procurement, agency performance management, and the underlying statutory frameworks.

Recent Developments (through 2026)

The pace of federal AI policy revision has accelerated; practitioners should monitor OMB and NIST publications quarterly. State legislatures have responded to perceived federal pullback by accelerating their own AI regulation (notably California, Colorado, Texas, and New York City).

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

Federal AI governance documentation requirements primarily live at Layer 5; the underlying technical practices (Layers 1–4) remain governed by NIST and sector-specific regulators.

PRACTITIONER NOTE

Build governance to the underlying statutory floor, not to executive policy ceilings. Statutes do not change with administrations; executive priorities do.

Communications, Intermediary, and Surveillance Laws

AI systems that capture, store, transmit, or analyze communications or device telemetry implicate a substantial body of federal communications and surveillance law. Section 230, ECPA, the VPPA, the CLOUD Act, FISA Section 702, and the EU–U.S. Data Privacy Framework together shape the cross-border and domestic data flows on which AI operations depend.

Video Privacy Protection Act (VPPA)

18 U.S.C. § 2710

Jurisdiction	United States - Federal
Effective	1988; substantial 2012 amendments
Regulator	Private right of action (no agency enforcement)
Scope	Video tape service providers - interpreted broadly to include streaming services, video-embedded websites, and (under recent case law) any business knowingly disclosing personally identifiable video viewing information

Applicability

Originally enacted in response to disclosure of a Supreme Court nominee's video rental records, the VPPA prohibits "video tape service providers" from knowingly disclosing personally identifiable information of "consumers" to third parties without consent. Modern litigation has applied the VPPA broadly to streaming services, news websites with embedded video, and (most consequentially for AI) websites using analytics pixels (Meta Pixel, Google Analytics) that transmit video viewing information together with personally identifying signals. AI-driven personalization, recommendation, and analytics features routinely create VPPA exposure.

Core Obligations

- Obtain informed, written consent before disclosing personally identifiable information identifying a person as having requested or obtained specific video materials.
- Consent must be in a form distinct and separate from any form setting forth other legal or financial obligations.
- Destroy personally identifiable information no later than one year from the date the information is no longer necessary for the purpose for which collected (subject to exceptions).

Penalties & Enforcement

Statutory damages of \$2,500 per violation, plus actual damages, punitive damages, and attorneys' fees. Class action exposure has driven the largest VPPA settlements (Hulu \$7M, NBCUniversal \$30M, multiple settlements in the \$5M–\$30M range).

Recent Developments (through 2026)

A wave of post-2022 VPPA litigation has applied the statute to website analytics and pixel tracking, particularly Meta Pixel transmissions. Several federal courts of appeals have wrestled with the "consumer" definition; the Second and Eleventh Circuits have applied the statute to free-content visitors who interacted with video features. AI personalization features that condition viewing recommendations on cross-site identifiers create significant exposure.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Viewing-history training data must include explicit VPPA consent - many AI recommender systems sit on top of unconsented data flows.
Layer 3 - System Integration	Third-party SDK and analytics integrations transmitting video viewing data are direct disclosure events; data flow diagrams must explicitly map every recipient.
Layer 4 - Control & Monitoring	Operational controls preventing inadvertent transmission via pixel-based tracking are Layer 4 deployment gates.

PRACTITIONER NOTE
 A common audit finding: marketing teams add Meta Pixel or similar tracking to video content pages without privacy-team review. Treat any analytics or advertising integration on pages with video content as a potential VPPA disclosure and require consent or pixel suppression.

Communications Decency Act Section 230

47 U.S.C. § 230

Jurisdiction	United States - Federal
Effective	1996; SESTA/FOSTA carve-out 2018
Regulator	N/A (immunity provision; courts apply on motion)
Scope	Providers and users of interactive computer services

Applicability

Section 230 provides federal immunity to interactive computer service providers for content provided by third parties (§ 230(c)(1)) and for good-faith content moderation actions (§ 230(c)(2)). The application of Section 230 to AI-generated content is the most actively contested intermediary-liability question of the post-ChatGPT era. Provider liability for generative AI outputs depends on whether the AI provider is the "information content provider" (no Section 230 protection for content the provider materially contributes to creating) or a passive intermediary. Courts have begun to address this question; the trajectory generally limits Section 230 protection for AI-generated content the provider produces in response to user prompts.

Core Obligations

- No affirmative obligations; Section 230 is an immunity defense.
- For good-faith moderation immunity: actions must be voluntarily taken in good faith to restrict access to or availability of material the provider considers objectionable.
- Carve-outs: federal criminal law, intellectual property, ECPA, sex trafficking (FOSTA-SESTA).

Penalties & Enforcement

N/A (immunity defense).

Recent Developments (through 2026)

Multiple 2024–2026 cases have addressed Section 230's application to AI: courts have generally treated chatbot-generated defamation claims as actionable against the AI provider where the provider materially contributes to the offending content; recommendation algorithms remain partly protected; product-liability and product-defect theories are increasingly being used to circumvent Section 230. The Section 230 reform debate continues but no significant statutory amendment has passed.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Generative AI output liability depends on Layer 2 design choices - output filtering, system prompts, and the degree of provider-induced content shape liability analysis.
Layer 5 - Audit & Evidence	Documentation of moderation policies, output filtering, and provider-induced content shaping is Layer 5 evidence in Section 230 motion practice.

CLOUD Act and Cross-Border Government Access

Clarifying Lawful Overseas Use of Data Act, 18 U.S.C. § 2713 (2018); ECPA § 2703

Jurisdiction	United States - Federal (extraterritorial)
Effective	March 23, 2018
Regulator	U.S. Department of Justice; foreign governments through executive agreements
Scope	Electronic communication service providers and remote computing service providers subject to U.S. jurisdiction; foreign government access via executive agreement

Applicability

The CLOUD Act amended ECPA to clarify that U.S. providers must produce data within their possession, custody, or control regardless of where the data is stored. Separately, it authorized executive agreements with qualifying foreign governments enabling those governments to compel data from U.S. providers (with reciprocity for U.S. authorities). The U.S.–UK CLOUD Act Agreement is the leading example. AI providers operating cross-border face complex jurisdictional analysis under CLOUD Act, GDPR Chapter V, and corresponding regimes.

Core Obligations

- Comply with legal process for stored content and non-content data, regardless of data location, subject to comity-based motion-to-quash standards.

- For executive agreement orders: comply with foreign government compulsion as authorized by the agreement.
- Maintain capability to produce data in response to lawful process; consider impact of architectural decisions (e.g., zero-knowledge designs) on compliance posture.

Penalties & Enforcement

Contempt and ECPA civil and criminal penalties for non-compliance.

Recent Developments (through 2026)

Multiple bilateral CLOUD Act executive agreements have been concluded or negotiated through 2026 (UK, Australia, EU framework discussions). The interaction with GDPR Chapter V (particularly Article 48) remains contested; multinational AI providers face dual-compliance regimes that occasionally conflict.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	AI architecture decisions - including data localization, encryption, and key custody - substantially shape CLOUD Act exposure; jurisdictional analysis is a Layer 3 design constraint.
Layer 5 - Audit & Evidence	Government request transparency reporting is increasingly expected; Layer 5 documentation should support both internal review and external transparency disclosures.

EU–U.S. Data Privacy Framework (DPF) and Successor Mechanisms

European Commission adequacy decision (July 10, 2023); Executive Order 14086 (2022); 28 CFR Part 201

Jurisdiction	European Union ↔ United States
Effective	July 10, 2023
Regulator	U.S. Department of Commerce (DPF self-certification administration); EU DPAs; Data Protection Review Court (DPRC)
Scope	U.S. organizations self-certifying to the Data Privacy Framework; EU controllers transferring personal data to certified U.S. organizations

Applicability

The DPF is the third generation of EU–U.S. data transfer mechanism (after Safe Harbor, struck down 2015, and Privacy Shield, struck down 2020 in Schrems II). The framework relies on the Commission's adequacy decision in conjunction with Executive Order 14086's signals intelligence safeguards and the Data Protection Review Court mechanism. AI training data and inference data flows from EU to U.S. routinely rely on the DPF as the transfer basis. The framework is the subject of pending challenges (Schrems III) and remains operative as of 2026.

Core Obligations

- For U.S. organizations: self-certify to Department of Commerce; comply with DPF principles including notice, choice, accountability for onward transfer, security, data integrity and purpose limitation, access, recourse and enforcement, and liability.
- Annual recertification.
- Subject to FTC or DOT enforcement of DPF commitments.

Penalties & Enforcement

FTC enforcement under Section 5; loss of DPF certification.

Recent Developments (through 2026)

Schrems III litigation pending before the CJEU as of early 2026; trajectory uncertain. Practitioners should maintain SCC fallback positions even where DPF coverage exists. EDPB statement on the DPF's first-year operation noted ongoing concerns about U.S. surveillance practices.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	DPF self-certification is a Layer 3 transfer mechanism; SCC plus supplementary measures remain the fallback architecture for entities not certified or for transfers outside DPF scope.
Layer 5 - Audit & Evidence	DPF compliance documentation, annual recertification artifacts, and complaint-handling records are Layer 5 deliverables subject to FTC inquiry.

Foreign Intelligence Surveillance Act (FISA) Section 702 and U.S. Surveillance Reform

50 U.S.C. § 1881a; *Reforming Intelligence and Securing America Act (April 2024)*

Jurisdiction	United States - Federal
Effective	Section 702 enacted 2008; reauthorized 2024 through 2026
Regulator	Foreign Intelligence Surveillance Court; Office of the Director of National Intelligence
Scope	U.S. providers compelled to assist with surveillance of non-U.S. persons reasonably believed to be located outside the U.S.

Applicability

Section 702 is the principal U.S. statutory authority for compelled provider assistance with foreign intelligence surveillance. The 2024 reauthorization expanded the definition of "electronic communication service provider" (ECSP) - a change significant for AI providers, datacenter operators, and providers of cloud services with foreign-located components. The CJEU's Schrems II concerns about

Section 702 underlie ongoing transfer-mechanism litigation. AI providers should anticipate compliance scenarios and maintain transparency reporting.

Core Obligations

- Compelled assistance with directives issued under Section 702.
- Confidentiality obligations limiting disclosure of specific compulsion.
- Aggregate transparency reporting permitted within statutory ranges.

Penalties & Enforcement

Contempt and other compulsion-enforcement remedies.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Aggregate transparency reporting is a Layer 5 disclosure regime increasingly material to enterprise AI customer due diligence.

Electronic Communications Privacy Act (ECPA) - Comprehensive

18 U.S.C. §§ 2510–2523 (Wiretap), §§ 2701–2713 (SCA), §§ 3121–3127 (Pen Register)

Jurisdiction	United States - Federal
Effective	1986; continuously amended
Regulator	Department of Justice; private civil right of action
Scope	Three statutes: Wiretap Act, Stored Communications Act, Pen Register/Trap and Trace

Applicability

ECPA's three components together govern interception, stored communications, and metadata collection. Beyond the AI-relevant provisions previously discussed, the Pen Register/Trap and Trace statute governs collection of non-content metadata (dialed numbers, addressing information). AI-driven tools that collect device identifiers, behavioral patterns, or addressing information may implicate Pen Register obligations in some contexts. The interaction with the Fourth Amendment's third-party doctrine post-Carpenter v. United States (2018) continues to develop.

Core Obligations

- Wiretap Act: criminal prohibition on interception of electronic communications absent consent; one- or all-party consent depending on jurisdiction.
- Stored Communications Act: restrictions on disclosure of stored communications by service providers; required process for compelled disclosure (warrant for content, subpoena/court order for non-content depending on age).

- Pen Register Act: court order required for installation of pen register or trap and trace device.

Penalties & Enforcement

Criminal penalties; civil damages of greater of actual or statutory minimum, plus attorneys' fees.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	AI tools collecting device telemetry, behavioral signals, or addressing information must consider ECPA implications at Layer 1 collection design.

CAN-SPAM Act and Commercial Electronic Communications

Controlling the Assault of Non-Solicited Pornography and Marketing Act, 15 U.S.C. §§ 7701–7713

Jurisdiction	United States - Federal
Effective	January 1, 2004
Regulator	Federal Trade Commission; FCC; state attorneys general; ISPs
Scope	Senders of commercial electronic mail messages

Applicability

CAN-SPAM regulates commercial email including AI-personalized marketing campaigns. The act preempts most state spam laws but several state laws remain operative for state-resident protections. AI-driven email targeting and content generation must comply with CAN-SPAM's sender identification, unsubscribe, and content requirements.

Core Obligations

- Do not use false or misleading header information.
- Do not use deceptive subject lines.
- Identify the message as an advertisement.
- Tell recipients where you are located (valid physical postal address).
- Tell recipients how to opt out of receiving future email; honor opt-out requests within 10 business days.
- Monitor what others are doing on your behalf (vicarious liability for vendors and partners).

Penalties & Enforcement

Civil penalties up to \$53,088 per violation per email (2026 adjusted); each non-compliant email is a separate violation.

STACK LENS - How this law maps to the AI Governance Stack

Layer 4 - Control & Monitoring	AI email campaign infrastructure must implement opt-out propagation across all systems; 10-day clock requires Layer 4 operational reliability.
---	--

Driver's Privacy Protection Act (DPPA) and Other Federal Sectoral Privacy Statutes

18 U.S.C. § 2721 (DPPA); various sectoral statutes

Jurisdiction	United States - Federal
Effective	DPPA 1994
Regulator	DOJ; private right of action
Scope	State motor vehicle records and personal information derived therefrom

Applicability

The DPPA prohibits state DMVs from disclosing personal information from motor vehicle records except under enumerated exceptions, and prohibits permitted recipients from re-disclosing. AI tools using motor vehicle records (insurance underwriting, fraud detection, identity verification) must comply with DPPA permissible-purpose requirements. Several other federal sectoral privacy statutes (Cable Communications Privacy Act, Right to Financial Privacy Act, Telecommunications Act customer proprietary network information) impose analogous restrictions.

Core Obligations

- Use motor vehicle record information only for enumerated permissible purposes.
- Maintain records of permitted disclosures.
- Restrict re-disclosure consistent with original permissible purpose.

Penalties & Enforcement

Civil damages of greater of actual damages or \$2,500, plus punitive damages and fees; criminal penalties.

STACK LENS - How this law maps to the AI Governance Stack

Layer 1 - Data Governance	Motor vehicle record use in AI underwriting and fraud detection requires Layer 1 permissible-purpose mapping; downstream model uses may exceed original authorization.
----------------------------------	--

Civil Rights and Employment Laws Applied to AI

Federal civil rights and employment statutes - Title VII, the ADEA, GINA, the Fair Housing Act, the ADA, OFCCP authorities, and the NLRA - apply with full force to AI systems used in employment, housing, education, and credit. The doctrinal frameworks of disparate treatment, disparate impact, and reasonable accommodation translate directly into Layer 2 model fairness practice.

Title VII of the Civil Rights Act and AI in Employment

42 U.S.C. § 2000e et seq.; EEOC Strategic Enforcement Plan 2024–2028

Jurisdiction	United States - Federal
Effective	1964 (continuously interpreted)
Regulator	Equal Employment Opportunity Commission (EEOC); private right of action
Scope	Employers with 15+ employees; employment agencies; labor organizations; certain federal contractors

Applicability

Title VII prohibits employment discrimination on the basis of race, color, religion, sex (including pregnancy, sexual orientation, and gender identity per *Bostock v. Clayton County*), or national origin. AI hiring, promotion, and termination tools are subject to Title VII's disparate treatment and disparate impact frameworks. The Uniform Guidelines on Employee Selection Procedures (29 CFR Part 1607) provide the operational framework for adverse impact analysis. The EEOC's 2024–2028 Strategic Enforcement Plan explicitly prioritizes algorithmic discrimination across all protected characteristics.

Core Obligations

- Avoid employment practices having disparate impact on protected classes that are not job-related and consistent with business necessity (the disparate impact framework).
- For AI selection tools: conduct validation studies under Uniform Guidelines; demonstrate job-relatedness and business necessity; identify less discriminatory alternatives.
- For all employment AI: avoid disparate treatment based on protected characteristics including via proxy variables.
- Maintain records sufficient to support adverse impact analysis (selection rates by demographic group).
- Provide reasonable accommodations for religious practices including in AI assessment processes.

Penalties & Enforcement

EEOC charge process leading to conciliation, litigation, or right-to-sue; private remedies including back pay, front pay, compensatory and punitive damages (capped by employer size, \$50,000–\$300,000), and attorneys' fees.

Recent Developments (through 2026)

EEOC v. iTutorGroup (2023) was the first publicly resolved EEOC AI discrimination case (\$365,000 settlement) and established the agency's direct enforcement appetite. Multiple agency guidance documents and the 2024 EEOC AI Joint Statement (with DOJ, CFPB, FTC) reinforce algorithmic disparate impact theory across federal regulators.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training data bias is a foundational disparate impact concern; demographic representation in training data is a Layer 1 audit element.
Layer 2 - Model Governance	Adverse impact testing (four-fifths rule) and validation studies are core Layer 2 fairness requirements; less discriminatory alternative analysis is the model-development discipline.
Layer 5 - Audit & Evidence	Validation studies and selection rate documentation are Layer 5 evidence in EEOC charges and litigation.

PRACTITIONER NOTE

For any AI tool used in selection, document the four-fifths rule analysis at deployment and at each material model update. Where adverse impact is observed, document either the validation supporting business necessity or the search for less discriminatory alternatives. Failing to perform the analysis is itself a litigation vulnerability.

Age Discrimination in Employment Act (ADEA) and AI

29 U.S.C. § 621 et seq.

Jurisdiction	United States - Federal
Effective	1967
Regulator	EEOC; private right of action
Scope	Employers with 20+ employees; protects individuals 40 years of age and older

Applicability

The ADEA prohibits employment discrimination on the basis of age (40+). AI tools that use age-correlated features (years of experience, graduation year, technology familiarity) may produce disparate impact on protected workers. The 2019 Facebook ad-targeting settlement established that age-based ad targeting in employment contexts can violate the ADEA; AI-driven hiring funnels face analogous exposure.

Core Obligations

- Do not discriminate against employees or applicants 40+ on the basis of age.

- Avoid AI-driven exclusion of older workers from recruitment, hiring, promotion, training, or benefits.
- For AI hiring: document the relationship between any age-correlated features and bona fide occupational qualifications.

Penalties & Enforcement

Liquidated damages for willful violations; back pay; attorneys' fees; injunctive relief.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Age-correlated feature analysis is a Layer 2 fairness practice often missed when teams focus on race and sex; Layer 2 testing should include age-banded subgroup analysis.

Genetic Information Nondiscrimination Act (GINA)

42 U.S.C. § 2000ff (Title II - employment); 42 U.S.C. § 300gg-91 (Title I - health insurance)

Jurisdiction	United States - Federal
Effective	November 21, 2009
Regulator	EEOC (employment); HHS / Departments of Treasury and Labor (insurance)
Scope	Employers (Title II); group health plans and health insurance issuers (Title I)

Applicability

GINA prohibits discrimination based on genetic information in employment and health insurance. AI tools that infer genetic information (from family medical history, surname patterns, demographic signals, or biometric patterns) may implicate GINA. The act also restricts collection, use, and disclosure of genetic information.

Core Obligations

- Do not request, require, or purchase genetic information of employees or family members (with limited exceptions).
- Maintain confidentiality of genetic information; restrict access.
- For health plans: do not use genetic information for underwriting; do not request or require genetic testing.

Penalties & Enforcement

EEOC enforcement under Title VII framework; private right of action; HHS enforcement for health plans.

STACK LENS - How this law maps to the AI Governance Stack

Layer 1 - Data Governance	AI inferences that implicitly capture genetic information (family medical history fields, photographic phenotype analysis) require Layer 1 GINA analysis.
----------------------------------	---

Office of Federal Contract Compliance Programs (OFCCP) and AI in Federal Contractor Employment

Executive Order 11246 (revoked 2025); Section 503 of Rehabilitation Act; VEVRAA; OFCCP guidance

Jurisdiction	United States - Federal (federal contractors)
Effective	Executive Order 11246 revoked January 2025; Section 503 and VEVRAA remain operative
Regulator	Office of Federal Contract Compliance Programs (OFCCP)
Scope	Federal contractors and subcontractors; coverage scope subject to ongoing changes following EO 11246 revocation

Applicability

OFCCP enforces affirmative action and non-discrimination obligations applicable to federal contractors. Following the January 2025 revocation of EO 11246, the affirmative action regime for race and sex in federal contracting has been substantially restructured; Section 503 (disability) and VEVRAA (protected veterans) obligations remain. AI hiring tools used by federal contractors face OFCCP scrutiny including audit obligations under remaining authorities.

Core Obligations

- Section 503: 7% utilization goal for individuals with disabilities; affirmative action plan including assessment of personnel processes including AI selection tools.
- VEVRAA: protected veteran utilization goal; affirmative action plan.
- For AI selection tools: maintain selection-rate data sufficient to support OFCCP audit; respond to data requests.

Penalties & Enforcement

OFCCP debarment; loss of federal contracting privileges; back pay and other relief through administrative or court proceedings.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	OFCCP audit data requests routinely include AI vendor information, model documentation, and selection-rate data; Layer 5 documentation should be assembled in anticipation rather than at audit notice.
---------------------------------------	---

Fair Housing Act (FHA) and HUD AI Guidance

42 U.S.C. §§ 3601–3619; HUD Office of General Counsel Guidance on Application of FHA Standards to AI (2024)

Jurisdiction	United States - Federal
Effective	1968; HUD AI guidance 2024
Regulator	Department of Housing and Urban Development (HUD); DOJ; state and local fair housing agencies
Scope	Housing-related transactions including sales, rentals, lending, advertising, and provision of services

Applicability

The FHA prohibits discrimination on the basis of race, color, national origin, religion, sex (including sexual orientation and gender identity), familial status, and disability in housing. AI tools used in tenant screening, mortgage underwriting, housing-related advertising, property valuation, and eviction prediction face FHA exposure under both disparate treatment and disparate impact theories. HUD's 2024 guidance addresses AI-specific applications including tenant screening and advertising algorithms.

Core Obligations

- Avoid housing-related practices having disparate impact on protected classes that are not necessary to achieve a substantial, legitimate, nondiscriminatory interest, where less discriminatory alternatives are available.
- For tenant screening AI: validate models for accuracy and fairness; provide adverse action notices consistent with FCRA where applicable; consider less discriminatory alternatives.
- For housing advertising: do not use targeting features that exclude protected classes (post-Facebook Settlement framework).
- For mortgage underwriting: comply with ECOA and FHA frameworks; maintain disparate impact analysis.
- For property valuation AI: address bias risk in automated valuation models.

Penalties & Enforcement

HUD administrative penalties; DOJ pattern-or-practice litigation; private FHA suits with statutory and punitive damages plus fees; civil penalties up to \$25,000 first violation, \$65,000 per subsequent (2026 adjusted).

Recent Developments (through 2026)

HUD's 2024 guidance and the 2023 Settlement with SafeRent (\$2.275M) over algorithmic tenant screening established federal enforcement against AI tenant screening. The Facebook Settlement (2022)

addressed algorithmic ad targeting; subsequent compliance is the operating model for housing-related ad delivery.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training data for tenant screening, mortgage underwriting, and valuation models often inherits historical bias; Layer 1 representativeness analysis is essential.
Layer 2 - Model Governance	Disparate impact testing for housing AI parallels ECOA fair-lending analysis; Layer 2 fairness testing must include disparate impact ratio across protected classes.
Layer 4 - Control & Monitoring	Ad-delivery integration with housing advertisers requires Layer 4 controls preventing prohibited targeting categories.

Americans with Disabilities Act (ADA) Title III and AI Accessibility

42 U.S.C. § 12182 (Title III); DOJ ANPRM on Web Accessibility (2024)

Jurisdiction	United States - Federal
Effective	ADA 1990; Title III interpretation continuing
Regulator	Department of Justice; private right of action
Scope	Public accommodations and commercial facilities

Applicability

ADA Title III prohibits discrimination against individuals with disabilities by public accommodations. AI-driven customer-facing systems including chatbots, voice interfaces, recommendation engines, and consumer-decision tools must be accessible. Web accessibility under Title III has produced extensive litigation; AI-specific accessibility (alternative text generation, voice agent compatibility with assistive technology, captioning quality) is the emerging frontier. DOJ's 2024 web accessibility ANPRM may produce technical standards for digital accessibility under Title III in 2026–2027.

Core Obligations

- Provide effective accessibility for individuals with disabilities; for digital interfaces, common practice references WCAG 2.1 AA.
- For AI-generated content: provide accessible alternatives (alt text, captions, audio descriptions where applicable).
- For AI customer service: ensure compatibility with assistive technologies; provide accessible alternative channels.
- Document accessibility practices and respond to accommodation requests.

Penalties & Enforcement

DOJ enforcement and consent decrees; private rights of action with attorneys' fees and injunctive relief; substantial settlement costs in private litigation.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	AI integration with customer-facing channels must include accessibility validation at the integration boundary.
Layer 4 - Control & Monitoring	Operational monitoring for accessibility regressions in AI outputs (e.g., generated images without alt text) is a Layer 4 deployment gate.

National Labor Relations Act (NLRA) and AI in the Workplace

29 U.S.C. § 151 et seq.; NLRB General Counsel Memo GC 23-02 (2022)

Jurisdiction	United States - Federal
Effective	1935; NLRB AI guidance 2022 onward
Regulator	National Labor Relations Board (NLRB)
Scope	Most private-sector employers and employees engaged in protected concerted activity

Applicability

NLRA Section 7 protects employees' right to engage in protected concerted activity. NLRB has issued guidance applying Section 7 to AI workplace surveillance, monitoring, and management tools - employer use of AI to surveil workers may unlawfully chill protected concerted activity. AI tools that monitor employee productivity, communications, location, or sentiment face NLRA scrutiny in addition to ECPA, state surveillance laws, and emerging worker protection statutes.

Core Obligations

- Disclose to employees the existence and general nature of AI workplace monitoring.
- Avoid AI-driven discipline based on protected concerted activity (e.g., union organizing, joint complaints about working conditions).
- Refrain from AI surveillance designed to identify or chill protected concerted activity.

Penalties & Enforcement

NLRB unfair labor practice complaints; remedies include back pay, reinstatement, and posting requirements.

STACK LENS - How this law maps to the AI Governance Stack

Layer 4 - Control & Monitoring	AI workplace monitoring deployments require Layer 4 governance review including Section 7 analysis; legal review at the procurement stage avoids costly retrofits.
---	--

Workforce Surveillance Statutes (State Survey)

NY Labor Law § 52-c (electronic monitoring notice); CA AB 2188 (off-duty cannabis); various pending state worker AI bills

Jurisdiction	Multiple U.S. states
Effective	NY 2022; varying other states
Regulator	State labor agencies; private rights of action
Scope	Employers using AI/electronic monitoring of workers

Applicability

A growing wave of state laws requires notice or consent for electronic worker monitoring, restricts use of AI in worker discipline, and protects against AI bias in employment decisions. New York requires written notice of electronic monitoring at hiring. Several states have considered or enacted notice obligations for AI worker management tools, including limits on continuous monitoring and prohibitions on AI-driven termination without human review.

Core Obligations

- NY: written notice at hire of electronic monitoring; prominent posting; acknowledgment.
- Several states: consent for biometric employee data, including time-clock biometrics.
- Pending state bills: notice and limits on AI-driven discipline; meaningful human review of AI termination recommendations.

Penalties & Enforcement

State labor agency enforcement; per-violation civil penalties; potential private rights.

STACK LENS - How this law maps to the AI Governance Stack

Layer 4 - Control & Monitoring	Workforce monitoring AI requires Layer 4 disclosure and operational governance; deployment without notice is a per se violation in growing list of states.
---	--

Critical Infrastructure and Sector-Specific Federal Frameworks

AI systems supporting energy, transportation, nuclear, aviation, financial, and defense infrastructure face the most stringent U.S. cybersecurity and operational expectations. CISA, NERC CIP, TSA directives, NRC frameworks, FAA software qualification, FERC INSM orders, and the DOD Responsible AI Strategy together constitute the federal critical-infrastructure overlay.

Cybersecurity and Infrastructure Security Agency (CISA) Act and Critical Infrastructure AI Guidance

6 U.S.C. § 651 et seq. (CISA Act); CISA Secure AI System Development guidance (April 2024); CIRCIA (Cyber Incident Reporting for Critical Infrastructure Act, 6 U.S.C. § 681 et seq.)

Jurisdiction	United States - Federal
Effective	CISA Act 2018; CIRCIA 2022 (rules pending); CISA AI guidance 2024
Regulator	CISA; sector-specific risk management agencies (SRMAs)
Scope	Critical infrastructure entities across 16 designated sectors; federal civilian executive branch agencies

Applicability

CISA is the federal coordinator for critical infrastructure cybersecurity. The 2024 CISA Secure AI System Development guidance (issued jointly with international partners including NCSC-UK) establishes baseline expectations for AI systems supporting critical infrastructure. CIRCIA (when rules are finalized, anticipated 2026–2027) will require covered critical infrastructure entities to report covered cyber incidents within 72 hours and ransomware payments within 24 hours. AI systems supporting critical infrastructure operations face cumulative federal expectations.

Core Obligations

- For critical infrastructure operators using AI: apply secure-by-design principles to AI systems supporting operations including model integrity, supply chain risk management, monitoring, and incident response.
- Under CIRCIA (when finalized): report covered cyber incidents within 72 hours and ransomware payments within 24 hours.
- Comply with Sector Risk Management Agency-specific requirements (e.g., Energy/DOE, Transportation/TSA, Financial/Treasury).
- For federal agencies: comply with CISA binding operational directives including emergency directives addressing AI-relevant vulnerabilities.

Penalties & Enforcement

CIRCIA: civil penalties for non-reporting (specifics in pending rule); SRMA-specific penalties under sector authorities.

Recent Developments (through 2026)

The 2024 joint guidance with international cybersecurity authorities established the most comprehensive multilateral baseline for AI cybersecurity. CISA has actively engaged with frontier AI developers on risk to critical infrastructure including through the 2024 voluntary commitments and ongoing red-team coordination.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	AI system integrity and supply chain risk management map to Layer 2 model governance for critical infrastructure deployments.
Layer 3 - System Integration	AI integration with operational technology (OT) systems requires Layer 3 security architecture distinct from IT-system practice.
Layer 4 - Control & Monitoring	Continuous monitoring and incident detection for AI systems in critical infrastructure are intensified Layer 4 obligations.
Layer 5 - Audit & Evidence	CIRCA reporting and SRMA documentation are Layer 5 evidence repositories with regulatory consequences.

NERC Critical Infrastructure Protection (CIP) Standards

NERC CIP-002 through CIP-014, plus CIP-002-016 series

Jurisdiction	United States - Federal (FERC-jurisdictional electric utilities); Canada (some provinces)
Effective	Continuous; CIP version 5/6 effective; subsequent versions in deployment
Regulator	North American Electric Reliability Corporation (NERC); Federal Energy Regulatory Commission (FERC) oversight
Scope	Bulk electric system (BES) cyber assets and associated facilities

Applicability

NERC CIP is the most prescriptive U.S. cybersecurity regulatory regime. AI systems supporting BES operations - operational technology AI, predictive maintenance AI for BES assets, AI-augmented control room decision support - fall within NERC CIP scope where they meet criticality criteria. CIP-013 (supply chain) and emerging AI-specific CIP guidance establish OT AI expectations.

Core Obligations

- Categorize BES Cyber Systems by impact level (high, medium, low).
- Implement CIP controls scaled to impact level: cyber security policies, personnel and training, electronic security perimeters, physical security, system security management, incident reporting and response, recovery planning, configuration change management, vulnerability assessments, information protection.

- CIP-013 supply chain risk management: vendor risk assessment, contractual security requirements, vulnerability disclosure.
- Audit and reporting to Regional Entity.

Penalties & Enforcement

FERC enforcement: civil penalties up to \$1.6M per day per violation; reputational and operational consequences are typically more material.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	CIP-013 supply chain obligations cover AI vendors providing OT-relevant tools; vendor inventories and contractual security flow-down are Layer 3 obligations.
Layer 4 - Control & Monitoring	OT AI deployment requires NERC CIP-aligned change management and monitoring; deployment gates are Layer 4 obligations with regulatory teeth.
Layer 5 - Audit & Evidence	NERC audits demand extensive documentation; AI-specific evidence packages should be assembled in anticipation.

TSA Pipeline and Rail Cybersecurity Directives

TSA Security Directive Pipeline-2021-01 through 2021-02C; SD 1580/82-2022-01; SD 1580/82-2024-01 (rail); Pipeline performance-based directive 2024

Jurisdiction	United States - Federal
Effective	May 2021 onward; performance-based directive 2024
Regulator	Transportation Security Administration (TSA)
Scope	Owners and operators of pipelines, freight rail, passenger rail, and certain transportation systems

Applicability

TSA cybersecurity directives establish mandatory cybersecurity requirements for designated transportation systems. The 2024 transition to performance-based directives provides flexibility but maintains baseline obligations including cybersecurity risk assessment, network segmentation between IT and OT, vulnerability management, incident reporting (within 24 hours to CISA), and cybersecurity coordination with TSA. AI systems supporting pipeline and rail operations are within scope.

Core Obligations

- Implement cybersecurity assessment plan and submit to TSA.
- Implement specified or equivalent technical controls including network segmentation, access control, continuous monitoring.
- Report cybersecurity incidents to CISA within 24 hours of identification.

- Designate Cybersecurity Coordinator available 24/7.

Penalties & Enforcement

TSA civil penalties; potential criminal referral for willful violations; operational consequences.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	IT/OT segmentation requirements affect AI deployment architecture in transportation environments.
Layer 4 - Control & Monitoring	Continuous monitoring and 24-hour incident reporting are Layer 4 operational obligations.

Nuclear Regulatory Commission (NRC) Cybersecurity and AI Frameworks

10 C.F.R. § 73.54 (cybersecurity); NRC Regulatory Issue Summary on AI (2024)

Jurisdiction	United States - Federal
Effective	Cybersecurity regulation 2009; AI guidance ongoing
Regulator	Nuclear Regulatory Commission
Scope	NRC-licensed nuclear power plants and certain other nuclear facilities

Applicability

NRC § 73.54 requires nuclear plant licensees to provide high-assurance protection of digital computers, systems, and networks associated with safety, security, and emergency preparedness functions. The NRC has signaled AI-specific guidance development through 2026, including for AI applications in plant operations, predictive maintenance, and decision support. Use of AI in safety-significant functions remains tightly constrained.

Core Obligations

- High-assurance cybersecurity protection of safety, security, and emergency preparedness digital systems.
- For AI in safety-significant functions: extensive review and qualification.
- For AI in non-safety functions: defense-in-depth and verification requirements.

Penalties & Enforcement

NRC enforcement: civil penalties; license modifications, suspensions, or revocations.

STACK LENS - How this law maps to the AI Governance Stack

Layer 2 - Model Governance	AI in safety-significant nuclear functions requires extensive Layer 2 model verification and validation distinct from commercial practice.
-----------------------------------	--

FAA AI in Aviation and Aviation Cybersecurity

14 C.F.R. (various parts); FAA Order 8110.49C (software approval); RTCA DO-178C; DO-326A; FAA AI Roadmap (2024)

Jurisdiction	United States - Federal
Effective	Continuous; AI Roadmap 2024
Regulator	Federal Aviation Administration; international counterparts (EASA, etc.)
Scope	Civil aviation aircraft, components, operations, and air traffic management

Applicability

FAA-certified aviation systems use stringent software qualification frameworks (DO-178C, DO-254). AI systems in aviation safety-critical functions face the most stringent qualification requirements in commercial deployment. The FAA AI Roadmap (2024) establishes the agency's phased approach to AI certification, anticipating 2030+ for fully autonomous systems in passenger air operations. Cybersecurity for aviation systems is governed by DO-326A and related standards.

Core Obligations

- For software in safety-critical aviation systems: DO-178C compliance with assurance level commensurate with criticality.
- For ML-based systems: emerging EUROCAE WG-114 / SAE G-34 framework guidance (Statement of Concerns, ARP 6983).
- Cybersecurity per DO-326A and related airworthiness security standards.

Penalties & Enforcement

FAA civil penalties; airworthiness directive consequences; certification suspension or revocation.

STACK LENS - How this law maps to the AI Governance Stack

Layer 2 - Model Governance	AI/ML model qualification for aviation requires extensive Layer 2 verification; emerging standards (ARP 6983) are reshaping practice.
-----------------------------------	---

Federal Energy Regulatory Commission (FERC) Order 901 and Critical Infrastructure

FERC Order 901 (2023) - Internal Network Security Monitoring; subsequent FERC orders

Jurisdiction	United States - Federal
Effective	Order 901 effective phases through 2027
Regulator	FERC; NERC implementation
Scope	Bulk electric system high and medium impact BES Cyber Systems

Applicability

FERC Order 901 directs NERC to develop standards requiring internal network security monitoring (INSM) for high and medium impact BES Cyber Systems and certain Electronic Access Control or Monitoring Systems (EACMS) and Physical Access Control Systems (PACS). AI-driven anomaly detection is a permissible technical approach to INSM compliance.

Core Obligations

- Implement INSM consistent with NERC standards developed under Order 901.
- Detect and respond to anomalous network activity within the electronic security perimeter.
- Maintain documentation supporting INSM implementation.

Penalties & Enforcement

FERC penalties via NERC enforcement.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	AI-driven anomaly detection deployments in BES contexts must support FERC INSM evidence requirements.

Department of Defense Responsible AI Strategy and Procurement

DOD Directive 3000.09 (Autonomy in Weapon Systems); DOD Responsible AI Strategy and Implementation Pathway; DOD AI Adoption Strategy (2023)

Jurisdiction	United States - Federal (Department of Defense)
Effective	DODD 3000.09 originally 2012, updated January 2023; RAI Strategy 2022
Regulator	DOD CIO; Chief Digital and Artificial Intelligence Office (CDAO)
Scope	DOD use of AI; DOD contractors and suppliers

Applicability

DOD has established the most extensive AI governance framework among federal agencies. DODD 3000.09 governs autonomous weapon systems with explicit human-judgment requirements. The DOD Ethical Principles for AI (2020) and Responsible AI Strategy (2022) operationalize five ethical principles:

responsible, equitable, traceable, reliable, governable. DOD contractors face cascading obligations through DFARS and program-specific requirements.

Core Obligations

- For autonomous weapon systems: appropriate levels of human judgment; senior review and approval; rigorous T&E.
- For all DOD AI: align with five Ethical Principles; implement RAI risk management; maintain documentation.
- For DOD contractors: comply with applicable DFARS cybersecurity (CMMC), program-specific AI requirements, and contractual flow-down.

Penalties & Enforcement

Procurement and contracting consequences; criminal referrals for serious violations; suspension and debarment.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	DOD T&E (test and evaluation) requirements for AI systems are intensified Layer 2 obligations including operational test in representative conditions.
Layer 3 - System Integration	Human judgment requirements shape Layer 3 integration architecture; meaningful human control is operationalized at the integration boundary.
Layer 5 - Audit & Evidence	DOD AI documentation expectations exceed commercial practice; contracting flow-down means DOD-grade Layer 5 documentation cascades to suppliers.

Sector-Specific Federal Regulators on AI

Federal sectoral regulators have moved aggressively to apply existing statutory authorities to AI. The FDA framework for AI/ML medical devices, Federal Reserve SR 11-7 model risk management, CFPB adverse action guidance, EEOC AI in employment, Section 1557 algorithmic discrimination, the NAIC Model Bulletin, and SEC predictive analytics together establish the U.S. sector-specific AI regulatory baseline.

FDA AI/ML-Based Software as a Medical Device (SaMD) Framework

FD&C Act § 520; 21 CFR Parts 803, 806, 820, 830; FDA Guidance: Predetermined Change Control Plans for AI/ML-Enabled Devices (December 2024)

Jurisdiction	United States - Federal
Effective	Continuous; PCCP guidance December 4, 2024

Regulator	FDA Center for Devices and Radiological Health (CDRH)
Scope	Software functions intended to treat, diagnose, cure, mitigate, or prevent disease or other conditions, where the software is a medical device

Applicability

AI/ML-enabled medical devices fall within the existing FDA medical device framework (510(k), De Novo, PMA pathways). The 2024 Predetermined Change Control Plan (PCCP) guidance permits manufacturers to specify modifications to AI/ML algorithms in advance, allowing post-market updates without new submission. The 2025 final guidance and 2024 Total Product Lifecycle approach establish FDA's operating model for adaptive AI in healthcare.

Core Obligations

- Premarket pathway selection (510(k), De Novo, PMA) based on classification and predicate device analysis.
- Quality System Regulation (21 CFR 820) including design controls, risk management (ISO 14971), software lifecycle (IEC 62304), and clinical evaluation.
- Post-market surveillance, MDR reporting (21 CFR 803), and corrective action.
- For PCCP-enabled devices: documented modification protocol describing types of permitted changes, methods to develop and validate modifications, and impact assessment.
- Real-world performance monitoring; post-market data collection.
- Cybersecurity per FDA Premarket Cybersecurity Guidance (2023) and the SBOM requirement under FDORA § 524B.

Penalties & Enforcement

Civil penalties, seizure, injunction, criminal prosecution under the FD&C Act; misbranded or adulterated devices subject to recall and import refusal.

Recent Developments (through 2026)

FDA approved hundreds of AI/ML-enabled devices through 2025; the public AI/ML device list serves as a benchmark database. The 2024 PCCP guidance is operationally significant for adaptive AI products. The 2025 NIH-FDA collaboration on real-world performance monitoring infrastructure is the leading post-market initiative.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Clinical training data must be representative; FDA expects subgroup performance documentation as part of Substantial Equivalence and De Novo determinations.
Layer 2 - Model Governance	Algorithm documentation, validation, and PCCP modification protocols are core Layer 2 deliverables.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	EHR integration architecture and clinical workflow integration are evaluated as part of intended use; software architecture is reviewed in 510(k) and PMA submissions.
Layer 4 - Control & Monitoring	Post-market surveillance and real-world performance monitoring align with Layer 4; MDR thresholds drive incident escalation.
Layer 5 - Audit & Evidence	Design History File and Quality System documentation are the most extensive Layer 5 requirements in U.S. healthcare AI; FDA inspection readiness is operational.

PRACTITIONER NOTE

For continuously learning AI, the PCCP is the operational mechanism. Build the modification protocol with sufficient specificity to permit meaningful updates while satisfying FDA's expectations for predetermined boundaries; vague PCCPs are routinely subject to additional information requests.

Federal Reserve SR 11-7 - Model Risk Management

SR Letter 11-7 (Federal Reserve); OCC Bulletin 2011-12; SR 21-8 (third-party model risk management)

Jurisdiction	United States - Federal (banking supervision)
Effective	April 4, 2011 (continuously updated)
Regulator	Federal Reserve System; Office of the Comptroller of the Currency; FDIC (concurrent)
Scope	Banking organizations supervised by the Federal Reserve, OCC, and FDIC; models used in business decisions

Applicability

SR 11-7 is the foundational U.S. model risk management framework. It applies to "models" defined as quantitative methods, systems, or approaches that apply statistical, economic, financial, or mathematical theories, techniques, and assumptions to process input data into quantitative estimates. ML and AI models are squarely in scope. The 2024 OCC and Federal Reserve guidance reaffirmed the application of SR 11-7 principles to AI, with attention to interpretability, bias, and model risk for emerging techniques.

Core Obligations

- Robust model development including documentation, evaluation of conceptual soundness, ongoing performance monitoring, and outcomes analysis.
- Independent validation by parties with appropriate competence, independence, and authority - including evaluation of conceptual soundness, ongoing monitoring, outcomes analysis, and benchmarking.

- Governance, policies, and controls including model inventory, model risk policies, roles and responsibilities, internal audit function review, and management oversight.
- For third-party models (SR 21-8): vendor due diligence, contractual provisions, ongoing monitoring, and exit strategies.
- Documentation of all phases including model development, implementation, use, validation, and ongoing monitoring.

Penalties & Enforcement

Supervisory consequences include matter-requiring-attention designations, supervisory letters, formal enforcement actions, civil money penalties, capital surcharges, and management remediation orders.

Recent Developments (through 2026)

Federal Reserve and OCC have issued multiple AI-related supervisory letters and have included AI-specific findings in examinations. SR 11-7 remains the operative framework; no AI-specific replacement has been issued.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	SR 11-7 is essentially a Layer 2 model governance specification adapted for banking; conceptual soundness, validation, and ongoing monitoring map directly.
Layer 3 - System Integration	Third-party model risk under SR 21-8 governs Layer 3 AI vendor relationships; due diligence packages should anticipate examiner review.
Layer 5 - Audit & Evidence	Model inventory, validation reports, and outcomes analyses are core Layer 5 examination artifacts; model documentation is the principal supervisory deliverable.

CFPB Adverse Action and AI - Circulars 2022-03 and 2023-03

CFPB Circulars 2022-03 (May 26, 2022) and 2023-03 (September 19, 2023); ECOA Reg. B § 1002.9

Jurisdiction	United States - Federal
Effective	2022; 2023
Regulator	Consumer Financial Protection Bureau
Scope	Creditors using AI/ML for credit decisions

Applicability

The two circulars together establish the CFPB's position that the use of complex AI/ML models does not exempt creditors from the requirement to provide specific principal reasons for adverse action under ECOA and Regulation B. Black-box explanations and post-hoc rationalizations are insufficient; the reasons must accurately reflect the principal factors that led to the adverse decision.

Core Obligations

- Provide specific principal reasons that accurately describe the factors actually considered or scored by the model that led to the adverse action.
- Avoid checklist-of-30-or-more-reasons approaches that fail to meaningfully convey what mattered.
- For complex models, ensure that explanation methods can support specific, accurate, and actionable adverse action notices.

Penalties & Enforcement

CFPB enforcement under ECOA; penalties up to \$1,362,567 per day (2026 adjusted) for knowing violations; private right of action under ECOA.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Model interpretability is a regulatory requirement, not a best practice; Layer 2 explainability infrastructure must produce specific principal reasons for each adverse decision.
Layer 5 - Audit & Evidence	Adverse action notice records and the supporting explanation documentation are core Layer 5 artifacts in fair-lending examinations and litigation.

EEOC AI in Employment Guidance and ADA Title I

EEOC Technical Assistance: The ADA and the Use of Software, Algorithms, and AI to Assess Job Applicants and Employees (May 2022); EEOC v. iTutorGroup (2023)

Jurisdiction	United States - Federal
Effective	EEOC ADA Guidance May 12, 2022
Regulator	Equal Employment Opportunity Commission
Scope	Employers covered by Title I of the ADA (15+ employees) and other federal employment discrimination laws

Applicability

EEOC guidance establishes that employers may be liable under the ADA where AI hiring tools screen out individuals with disabilities, including where the tool measures characteristics that are not job-related and consistent with business necessity. Reasonable accommodation obligations apply to AI assessment processes. The 2023 iTutorGroup settlement established direct EEOC enforcement against algorithmic hiring discrimination.

Core Obligations

- Ensure that AI hiring tools do not screen out individuals with disabilities in a way that violates Title I.
- Provide reasonable accommodations during AI assessment processes including alternative evaluation methods.
- Avoid reliance on AI tools that measure characteristics correlated with disabilities without demonstration of job-relatedness and business necessity.
- Validate AI hiring tools for adverse impact across protected classes (consistent with Uniform Guidelines on Employee Selection Procedures).
- Provide notice to applicants of AI assessment use; permit reasonable accommodation requests.

Penalties & Enforcement

EEOC charge process leading to conciliation, litigation, or right-to-sue letters; private remedies under ADA include injunctive relief, back pay, compensatory and punitive damages, and attorneys' fees.

Recent Developments (through 2026)

EEOC AI Strategic Enforcement Plan (2024–2028) prioritizes algorithmic discrimination; the agency has pursued multiple AI-related charges and provided AAA-supported guidance.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Adverse impact testing and validation of AI hiring tools are Layer 2 fairness obligations; the Uniform Guidelines four-fifths rule remains the practical screen.
Layer 4 - Control & Monitoring	Reasonable accommodation infrastructure during AI assessment is a Layer 4 deployment obligation.

Section 1557 of the Affordable Care Act and AI Discrimination

42 U.S.C. § 18116; 45 C.F.R. § 92.210 (2024 final rule)

Jurisdiction	United States - Federal
Effective	Section 1557 March 23, 2010; 2024 final rule effective July 5, 2024 (AI provisions effective May 1, 2025)
Regulator	HHS Office for Civil Rights
Scope	Health programs and activities receiving federal financial assistance, including healthcare entities receiving Medicare or Medicaid payments

Applicability

The 2024 Section 1557 final rule (45 C.F.R. § 92.210) prohibits discrimination on the basis of race, color, national origin, sex, age, or disability through the use of patient-care decision-support tools. Covered entities must make reasonable efforts to identify uses of patient-care decision-support tools that

employ input variables or factors that measure race, color, national origin, sex, age, or disability, and to mitigate the risk of discrimination from such tools.

Core Obligations

- Make reasonable efforts to identify uses of patient-care decision-support tools that employ inputs measuring or correlating with protected characteristics.
- Make reasonable efforts to mitigate the risk of discrimination from identified tools.
- Document the identification and mitigation efforts.

Penalties & Enforcement

OCR enforcement under Title VI, Title IX, Section 504, and Age Discrimination Act frameworks; loss of federal funding; OCR resolution agreements with corrective action.

Recent Developments (through 2026)

The Section 1557 AI provisions are the most explicit U.S. federal regulation of clinical decision-support fairness. Implementation is in early stages; OCR has issued FAQs and is conducting compliance reviews.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Identification of input variables that measure or correlate with protected characteristics requires Layer 1 data lineage analysis.
Layer 2 - Model Governance	Mitigation of discrimination risk requires Layer 2 fairness testing and remediation; the documentation of identification and mitigation is the principal compliance artifact.
Layer 5 - Audit & Evidence	Documentation of "reasonable efforts" must be assessment-ready; OCR will request the analysis and mitigation records.

NAIC Model Bulletin on Use of Artificial Intelligence Systems

NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers (December 4, 2023)

Jurisdiction	United States - State (insurance, varies by state adoption)
Effective	2024–2026 (state-by-state adoption)
Regulator	State insurance commissioners
Scope	Insurers using AI in any insurance practice (underwriting, rating, claims, marketing, fraud detection)

Applicability

The NAIC Model Bulletin establishes expectations for insurer AI governance including written AI program, governance and risk management, third-party AI risk management, and consumer protection. As of 2026, more than half of U.S. states have adopted or implemented the bulletin.

Core Obligations

- Maintain written AI program documenting governance, risk management, and use of AI systems.
- Implement testing for fairness, bias, and accuracy.
- Conduct third-party AI vendor due diligence.
- Maintain documentation supporting AI use in regulated activities including underwriting, rating, and claims.
- Comply with applicable state unfair trade practices laws including unfair discrimination prohibitions.

Penalties & Enforcement

State insurance commissioner enforcement under unfair trade practices and market conduct frameworks; civil penalties varying by state.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Fairness and accuracy testing are Layer 2 obligations; insurance commissioner examinations increasingly request testing documentation.
Layer 3 - System Integration	Third-party AI vendor due diligence is a Layer 3 governance requirement; vendor inventories must include AI components.
Layer 5 - Audit & Evidence	Written AI program documentation is the central Layer 5 deliverable; market conduct examinations evaluate this artifact.

SEC Predictive Data Analytics Proposal and Existing Disclosure Obligations

Proposed Rule 17 CFR §§ 240, 275 (Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers, 2023; pending)

Jurisdiction	United States - Federal
Effective	Proposed; existing disclosure and fiduciary obligations apply
Regulator	Securities and Exchange Commission
Scope	Broker-dealers and investment advisers using predictive data analytics or similar technology in investor interactions

Applicability

The SEC's 2023 proposed rule on predictive data analytics has been substantially revised in subsequent action and remains pending in modified form. Existing fiduciary, disclosure, and anti-fraud obligations

apply to AI use in investment advice and broker-dealer activities; the 2024 SEC enforcement actions against AI-marketing firms ("AI-washing") established the SEC's enforcement appetite. Investment advisers and broker-dealers using AI must consider Form ADV disclosures, marketing rule obligations, suitability and best-interest standards, and the use of AI in customer-facing communications.

Core Obligations

- Disclose material AI use in client communications, Form ADV, and marketing materials in a manner that is not false or misleading.
- Apply existing fiduciary and best-interest obligations to AI-driven recommendations and advice.
- For broker-dealers, comply with Reg BI and FINRA rules including suitability when AI generates or supports recommendations.
- Manage conflicts of interest associated with AI use including incentives that may misalign the AI with client interests.

Penalties & Enforcement

SEC civil enforcement including monetary penalties, disgorgement, censure, and injunctive relief; FINRA enforcement for broker-dealers; private securities litigation.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Conflict-of-interest analysis for AI models requires Layer 2 documentation of model objectives and incentive alignment.
Layer 5 - Audit & Evidence	AI-related disclosures in Form ADV, marketing materials, and client communications are Layer 5 transparency artifacts under continuing SEC scrutiny.

P A R T I I

United States - State Privacy & Cybersecurity Laws

Comprehensive state consumer privacy statutes from California to Maryland, the leading state biometric and cybersecurity regimes, and the multistate breach-notification patchwork - together the most rapidly evolving regulatory surface for U.S. AI practitioners.

Comprehensive State Privacy Laws - California, Virginia, and the First Wave

As of 2026, more than 20 U.S. states have enacted comprehensive consumer privacy laws. The leading regimes follow either the California pattern (CCPA/CPRA) or the Virginia pattern (VCDPA), with Colorado introducing the third-generation features that influence subsequent laws nationally. Practitioners with multistate operations build to the strictest applicable obligation by element - typically California for ADMT and SPI, Colorado for biometrics and profiling, Texas and Nebraska for breadth of coverage, and Washington for inferred health data.

California Consumer Privacy Act (CCPA) and California Privacy Rights Act (CPRA)

Cal. Civ. Code §§ 1798.100–1798.199.100; CPPA Regs. (Cal. Code Regs. tit. 11, §§ 7000 et seq.)

Jurisdiction	California
Effective	CCPA January 1, 2020; CPRA amendments January 1, 2023; ADMT and risk assessment regulations finalized 2025, phased compliance through 2027
Regulator	California Privacy Protection Agency (CPPA); California Attorney General (concurrent)
Scope	For-profit businesses doing business in California meeting one of three thresholds: (1) annual gross revenues exceeding \$26.625M (2026 adjusted); (2) buying, selling, or sharing personal information of 100,000 or more consumers/households; or (3) deriving 50% or more of annual revenues from selling or sharing personal information

Applicability

The CCPA, as amended by the CPRA, is the most comprehensive U.S. state privacy law and serves as a template for many subsequent state regimes. The 2025 CPPA regulations on Automated Decisionmaking Technology (ADMT), risk assessments, and cybersecurity audits substantially expand AI-specific obligations. Sensitive personal information (SPI) - including biometric, geolocation, racial/ethnic origin, religious beliefs, contents of communications, and certain inferences - receives heightened protection.

Core Obligations

- Provide notices at collection identifying categories of personal information collected, purposes of use, and retention periods.
- Honor consumer rights of access, deletion, correction, opt-out of sale/sharing, opt-out of cross-context behavioral advertising, limitation of use of SPI, and opt-out of automated decisionmaking technology (ADMT) for significant decisions.
- For ADMT used for "significant decisions" (employment, housing, healthcare, financial services, education, essential goods/services): provide pre-use notice describing the technology and its purpose, honor opt-out requests, and provide an alternative process where opt-out is honored.

- Conduct documented risk assessments before processing that presents a significant risk to consumers (including most uses of ADMT for significant decisions, training of ADMT models on personal information, and processing of SPI for inference).
- Conduct annual cybersecurity audits where a business processes personal information that "presents significant risk to consumers' security" (thresholds based on revenue and data volume).
- Implement and respect Global Privacy Control (GPC) signals as opt-out of sale/sharing.
- Limit retention to what is reasonably necessary; disclose retention period or criteria.
- Maintain service-provider, contractor, and third-party contracts with prescribed terms (§ 1798.140(ag), (j), (ai)).

Penalties & Enforcement

Civil penalties of \$2,500 per violation (\$7,500 for intentional violations or violations involving the personal information of consumers under 16), administered by the CPPA and AG. Limited private right of action for data breaches involving certain categories of personal information (\$100–\$750 per consumer per incident or actual damages, whichever is greater). Multiple settlements above \$1M; AG and CPPA enforcement against Sephora, DoorDash, and others established that data sharing for advertising constitutes a sale.

Recent Developments (through 2026)

CPPA finalized ADMT, risk assessment, and cybersecurity audit regulations in 2025 with phased effective dates running into 2027. Sephora settlement (\$1.2M, 2022); DoorDash settlement (\$375K, 2024); Honda settlement (\$632K, 2025). California has emerged as the de facto national privacy regulator, with multistate coordinated enforcement actions involving CT, CO, and OR AGs.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Notice-at-collection, retention limits, and SPI scoping are direct Layer 1 controls; the data inventory required for response to access requests is functionally a Layer 1 catalog.
Layer 2 - Model Governance	ADMT regulations require pre-use evaluations of accuracy and non-discrimination - Layer 2 fairness testing as a regulated obligation.
Layer 3 - System Integration	Service-provider, contractor, and third-party contracts with prescribed flow-down provisions govern Layer 3 boundaries; the CCPA contract regime is the most prescriptive in U.S. law.
Layer 4 - Control & Monitoring	GPC signal honoring requires real-time opt-out enforcement at every advertising integration point - a Layer 4 monitoring obligation.
Layer 5 - Audit & Evidence	Risk assessments and cybersecurity audits are Layer 5 deliverables, with CPPA explicitly empowered to demand them on inquiry; recordkeeping requirements run 24 months minimum.

PRACTITIONER NOTE

Build the ADMT pre-use notice and opt-out flow as a single, reusable component. Many businesses have multiple "significant decisions" use cases (hiring, credit, insurance) that share infrastructure; designing ADMT compliance once at the platform layer is dramatically cheaper than per-product retrofits. The CPPA's explicit risk-assessment-on-demand authority makes pre-built artifacts a strategic necessity.

COMMON FAILURE PATTERN

Cookie consent and Global Privacy Control honoring are treated as a CMP-vendor problem rather than a real-time data-flow enforcement problem. A "do not sell" toggle exists on the website footer, but the advertising and analytics SDKs continue to receive identifiers because the suppression is implemented client-side after the request fires. The CPPA reads the network capture, not the consent banner. Sephora, DoorDash, and Honda settlements all turned on this exact failure mode.

Virginia Consumer Data Protection Act (VCDPA)

Va. Code §§ 59.1-575 to 59.1-585

Jurisdiction	Virginia
Effective	January 1, 2023
Regulator	Virginia Attorney General
Scope	Persons conducting business in Virginia or producing products/services targeted to Virginia residents that (a) control or process personal data of 100,000+ consumers, or (b) control or process data of 25,000+ consumers and derive 50%+ of revenue from sale of personal data

Applicability

The VCDPA is the original "second-generation" U.S. privacy law and the template for the Virginia/Connecticut/Colorado-pattern statutes. It distinguishes between controllers and processors, requires opt-out for sale and targeted advertising, opt-in for sensitive data, and mandates data protection assessments for high-risk processing including profiling that presents reasonably foreseeable risks of unfair or deceptive treatment, financial/physical/reputational injury, or intrusion upon seclusion.

Core Obligations

- Provide privacy notices identifying categories of data, purposes, third-party recipients, and consumer rights mechanisms.
- Honor consumer rights of access, correction, deletion, portability, and opt-out of sale, targeted advertising, and certain profiling.
- Obtain opt-in consent before processing sensitive data (race, ethnicity, religious beliefs, mental/physical health, sexual orientation, immigration status, genetic/biometric data, precise geolocation, children's data).

- Conduct data protection assessments for high-risk processing including profiling presenting reasonably foreseeable risks.
- Maintain processor contracts with prescribed terms (§ 59.1-579).
- Respond to consumer rights requests within 45 days (extendable to 90 with notice).

Penalties & Enforcement

Civil penalties up to \$7,500 per violation; 30-day right to cure (sunsetting in many subsequent state laws but retained in Virginia).

Recent Developments (through 2026)

Virginia AG has prioritized enforcement of opt-out mechanisms and targeted advertising disclosures; the 2024–2025 docket includes ad-tech and data-broker investigations. Several state laws subsequently modeled on the VCDPA have removed the cure period - practitioners in multistate compliance should not rely on it as a structural assumption.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Sensitive data opt-in shapes Layer 1 collection design; biometric and inferred-attribute pipelines need explicit consent infrastructure.
Layer 2 - Model Governance	Profiling assessments map directly onto Layer 2 model documentation and fairness testing.
Layer 3 - System Integration	Processor contracts establish Layer 3 boundaries; the VCDPA flow-down requirements are now standard contract drafting.
Layer 5 - Audit & Evidence	Data protection assessments are confidential to the AG but discoverable on inquiry - Layer 5 evidence under privilege protection.

Colorado Privacy Act (CPA)

Colo. Rev. Stat. §§ 6-1-1301 to 6-1-1313; Colo. Code Regs. § 904-3

Jurisdiction	Colorado
Effective	July 1, 2023; AG biometric and ADMT rules effective 2025–2026
Regulator	Colorado Attorney General
Scope	Controllers conducting business in Colorado or producing commercial products/services intentionally targeted to Colorado residents that (a) control or process personal data of 100,000+ consumers, or (b) derive revenue from sale of personal data and process data of 25,000+ consumers

Applicability

The CPA follows the Virginia template but adds (1) explicit Universal Opt-Out Mechanism (UOOM) recognition (since July 2024), (2) the most detailed AG implementing regulations of any state, (3) substantial 2024 amendments creating standalone biometric and neural data protections, and (4) significant 2024–2025 amendments addressing children's data and AI profiling. The 2024 Colorado AI Act (separate statute) layers consequential decision obligations on top of the CPA framework.

Core Obligations

- Honor universal opt-out mechanisms (e.g., Global Privacy Control) for sale and targeted advertising.
- Conduct and document data protection assessments for high-risk processing.
- For biometric data (effective July 2024): notice, consent before collection, retention limits (24 months unless extended for permissible purpose), written biometric policies, and security obligations independent of general consumer status.
- For neural data (effective August 2024): treated as sensitive data under the CPA, requiring opt-in consent.
- For children's data (effective October 2025): enhanced consent and prohibition on certain dark patterns affecting minors.
- Comply with detailed AG regulations on consent, opt-out signals, profiling, and data protection assessments.

Penalties & Enforcement

Civil penalties under the Colorado Consumer Protection Act framework; per-violation penalties up to \$20,000. The cure period sunset on January 1, 2025; AG no longer offers automatic right to cure.

Recent Developments (through 2026)

Colorado was the first U.S. state to extend privacy protections to neural data (2024), and the Colorado AI Act (SB 24-205) is the most ambitious state-level AI consequential-decision regime. The AG's active rulemaking docket through 2026 makes Colorado the leading state regulator for substantive AI privacy practice.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Biometric and neural data carve-outs require segmented Layer 1 controls; combining these data with other personal data without consent triggers liability.
Layer 2 - Model Governance	Profiling assessments under § 6-1-1309 require Layer 2 documentation of accuracy, fairness, and consumer impact.
Layer 4 - Control & Monitoring	UOOM honoring at advertising integration points is a Layer 4 enforcement obligation; many businesses have failed audits because their cookie consent stack cannot honor GPC consistently.
Layer 5 - Audit & Evidence	AG's expansive rulemaking creates evolving Layer 5 documentation expectations; assessments must be refreshed when material changes occur.

PRACTITIONER NOTE

Colorado's biometric and neural data carve-outs are the leading edge of U.S. consumer privacy law. Even businesses without Colorado-specific operations should treat these as proxies for the next generation of multistate requirements.

Connecticut Data Privacy Act (CTDPA)

Conn. Gen. Stat. §§ 42-515 to 42-525

Jurisdiction	Connecticut
Effective	July 1, 2023; substantial 2024 amendments effective October 1, 2024
Regulator	Connecticut Attorney General
Scope	Persons conducting business in Connecticut or producing products/services targeted to Connecticut residents that (a) control or process personal data of 100,000+ consumers excluding payment processing, or (b) control or process data of 25,000+ consumers and derive 25%+ of revenue from sale of personal data

Applicability

The CTDPA tracks the Virginia template with several enhancements: opt-out mechanism honoring (effective January 1, 2025), heightened protections for minors (effective October 2024), explicit data minimization standards, and a sunseting cure period (sunset July 1, 2025). The 2024 amendments added robust protections for consumers under 18, including default privacy settings, restrictions on targeted advertising, and prohibitions on dark patterns directed at minors.

Core Obligations

- Honor universal opt-out mechanisms for sale and targeted advertising.
- For minors (under 18): obtain consent before collection of geolocation or processing for targeted advertising or sale; do not use system design features that prolong engagement among minors when reasonable to do otherwise.
- Conduct data protection assessments for high-risk processing.
- Provide notice and honor consumer rights consistent with Virginia/Colorado patterns.
- Maintain processor contracts with prescribed terms.

Penalties & Enforcement

Civil penalties of \$5,000 per willful violation under the Connecticut Unfair Trade Practices Act; cure period sunset July 1, 2025.

Recent Developments (through 2026)

Connecticut AG Tong has been an active multistate participant; coordinated enforcement actions with California and other states are increasingly the operating model. The 2024 minors' amendment is among the most protective in the U.S. and influences pending federal proposals.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Minor-specific data minimization shapes Layer 1 collection defaults; "off by default" is the operating principle.
Layer 2 - Model Governance	Engagement-prolongation prohibitions reach Layer 2 model design where recommender systems are optimized for time-on-platform.
Layer 4 - Control & Monitoring	Opt-out signal honoring is Layer 4 by definition; the same compliance patterns developed for California and Colorado generally satisfy Connecticut.

Utah Consumer Privacy Act (UCPA)

Utah Code §§ 13-61-101 to 13-61-404

Jurisdiction	Utah
Effective	December 31, 2023
Regulator	Utah Attorney General; Department of Commerce, Division of Consumer Protection
Scope	Controllers/processors with \$25M+ annual revenue meeting one of two data-volume thresholds (100,000+ Utah consumer records or 25,000+ records and 50%+ revenue from sale)

Applicability

The UCPA is the most business-friendly of the major state privacy laws - narrower scope, opt-out (not opt-in) for sensitive data, no profiling assessment requirement, and no UOOM honoring. AI applicability is correspondingly narrower but the law's presence affects multistate compliance program structure.

Core Obligations

- Provide privacy notices.
- Honor consumer rights of access, deletion, portability, and opt-out of sale and targeted advertising.
- Provide opt-out (not opt-in) for sensitive data processing.
- Maintain processor contracts.

Penalties & Enforcement

Civil penalties up to \$7,500 per violation; 30-day cure period.

STACK LENS - How this law maps to the AI Governance Stack**Layer 5 - Audit & Evidence**

Even where UCPA imposes minimal substantive obligations, multistate compliance documentation should reflect the UCPA exemption posture for clarity.

Texas Data Privacy and Security Act (TDPSA)

Tex. Bus. & Com. Code §§ 541.001 to 541.205

Jurisdiction	Texas
Effective	July 1, 2024
Regulator	Texas Attorney General
Scope	Persons conducting business in Texas or producing products/services consumed by Texas residents, that process or engage in the sale of personal data, and that are not small businesses (per SBA definition); no consumer-volume threshold

Applicability

The TDPSA reaches further than other state laws because it lacks consumer-volume thresholds - small-business status (per SBA) is the principal carve-out. Sensitive data requires opt-in consent. Data protection assessments are required for high-risk processing including profiling presenting reasonably foreseeable risks. The 2025 Texas Responsible AI Governance Act (TRAIGA, separate statute) layers AI-specific obligations.

Core Obligations

- Provide privacy notices including a clear and conspicuous disclosure if selling sensitive or biometric personal data.
- Honor consumer rights of access, correction, deletion, portability, and opt-out of sale, targeted advertising, and profiling.
- Obtain opt-in consent for sensitive data processing.
- Conduct data protection assessments for high-risk processing.
- Maintain processor contracts.
- Respond to consumer requests within 45 days.

Penalties & Enforcement

Civil penalties up to \$7,500 per violation; 30-day cure period preserved (in contrast to most newer state laws).

Recent Developments (through 2026)

Texas AG has aggressively enforced both the TDPSA and the parallel Capture or Use of Biometric Identifier Act (CUBI). The June 2024 settlement with Meta (\$1.4B over CUBI) was the largest privacy settlement against a single company in U.S. history at the time.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	No consumer-volume threshold means smaller AI vendors are routinely in scope; data inventories must include all Texas personal data, not just high-volume datasets.
Layer 5 - Audit & Evidence	Profiling assessments require Layer 5 documentation comparable to Virginia/Colorado.

Oregon Consumer Privacy Act (OCPA)

Or. Rev. Stat. §§ 646A.570 to 646A.589

Jurisdiction	Oregon
Effective	July 1, 2024 (controllers); July 1, 2025 (nonprofits)
Regulator	Oregon Attorney General
Scope	Persons conducting business in Oregon or providing products/services to Oregon residents and processing data of 100,000+ consumers (or 25,000+ where 25%+ revenue is from sale of personal data)

Applicability

Oregon's law is notable for (a) extending coverage to nonprofits (effective 2025) - the only major state law to do so; (b) requiring controllers to provide consumers with the specific identities of third parties to whom personal data has been disclosed (a transparency obligation more demanding than other state laws); and (c) extending sensitive data definitions to include national origin, status as transgender or nonbinary, and victim-of-crime status.

Core Obligations

- Provide privacy notices and respond to consumer rights requests, including the right to a list of specific third parties to whom personal data has been disclosed.
- Obtain opt-in consent for sensitive data processing.
- Conduct data protection assessments for high-risk processing.
- Maintain processor contracts.
- For nonprofits (from July 2025), the same obligations apply.

Penalties & Enforcement

Civil penalties up to \$7,500 per violation under the Unlawful Trade Practices Act; cure period sunset January 1, 2026.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	The third-party-identity disclosure requirement requires Layer 5 logging at the recipient level - many ad-tech and analytics integrations cannot satisfy this without architectural changes.

Texas Capture or Use of Biometric Identifier Act (CUBI)

Tex. Bus. & Com. Code § 503.001

Jurisdiction	Texas
Effective	2009
Regulator	Texas Attorney General (sole enforcement; no private right of action)
Scope	Persons capturing biometric identifiers (retina/iris scans, fingerprints, voiceprints, hand/face geometry) for commercial purposes

Applicability

CUBI prohibits capture of a biometric identifier for commercial purpose without prior informed consent, restricts sale and disclosure, requires reasonable security, and limits retention to one year after the purpose for collection ends. The 2024 Meta settlement (\$1.4B) and 2025 Google settlement (\$1.375B) established Texas as the most consequential biometric privacy enforcer in the U.S.

Core Obligations

- Obtain prior informed consent before capturing biometric identifiers for commercial purposes.
- Do not sell, lease, or otherwise disclose biometric identifiers except under enumerated exceptions.
- Implement reasonable security measures and destroy biometric identifiers within a reasonable time, in any event no later than one year after the purpose for collecting them expires.

Penalties & Enforcement

Civil penalties up to \$25,000 per violation. Texas AG has interpreted "per violation" to allow per-individual or per-image counting, producing the unprecedented Meta and Google settlements.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Biometric capture pipelines require explicit upstream consent gating; AI features that derive biometric identifiers from images (face geometry from photos) are within scope even when the original collection purpose was different.
Layer 4 - Control & Monitoring	One-year retention requires Layer 4 deletion automation tied to purpose-end events.

Illinois Biometric Information Privacy Act (BIPA)

740 ILCS 14/

Jurisdiction	Illinois
Effective	2008; major 2024 amendment limiting per-scan damages effective August 2, 2024
Regulator	Private right of action (no agency enforcement)
Scope	Private entities collecting, capturing, purchasing, receiving through trade, or otherwise obtaining biometric identifiers or biometric information

Applicability

BIPA is the most consequential privacy statute in the U.S. by litigation volume. It requires written notice, written consent, retention/destruction policies, and prohibitions on selling biometric data. The 2019 Illinois Supreme Court decision in *Rosenbach v. Six Flags* established that statutory violation alone confers standing - no actual injury required. The 2024 amendment (SB 2979) limits damages to one violation per individual per type of action (collection, sale, disclosure), substantially reducing aggregate damages but not eliminating litigation incentive.

Core Obligations

- Provide written notice that biometric identifiers/information are being collected or stored, and the purpose and length of collection, storage, and use.
- Obtain written executed consent before collection.
- Maintain a publicly available written policy establishing a retention schedule and destruction guidelines.
- Do not sell, lease, trade, or otherwise profit from biometric data.
- Implement reasonable care in storing, transmitting, and protecting biometric data.

Penalties & Enforcement

Statutory damages of \$1,000 per negligent violation and \$5,000 per intentional or reckless violation, plus attorneys' fees and equitable relief; 2024 amendment limits to one violation per individual per type. Major settlements include Facebook (\$650M, 2021), TikTok (\$92M, 2021), Google (\$100M, 2022), and BNSF (\$75M, 2023).

Recent Developments (through 2026)

The 2024 BIPA amendment was a partial business win, but BIPA litigation remains the leading privacy class action vehicle in the U.S. AI features that compute biometric measurements (face embeddings, voiceprints, gait analysis) routinely face BIPA suits even when the underlying images were public.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Biometric training-data collection in Illinois requires written consent infrastructure - a Layer 1 gate that many ML pipelines have not historically supported.
Layer 2 - Model Governance	Models that compute biometric features from non-biometric inputs (e.g., face geometry from photos) are within scope; Layer 2 model documentation should disclose biometric outputs.
Layer 3 - System Integration	Cross-state architecture should geofence biometric processing or apply the strictest standard universally - Illinois has effectively set the U.S. floor.

COMMON FAILURE PATTERN

A computer-vision feature is built and trained on a dataset of public images and ships globally. The face-geometry computation runs in production for every Illinois user. There is no written consent infrastructure because the team did not classify the system as a biometric pipeline. BIPA's private right of action and per-individual statutory damages convert this into a class action with eight or nine-figure settlement exposure. The 2024 amendment caps per-type aggregation but does not cure the underlying violation. Build the biometric classification check at the model-card stage, before any image-derived embedding reaches Illinois traffic.

Washington My Health My Data Act (MHMDA)

Wash. Rev. Code § 19.373

Jurisdiction	Washington
Effective	March 31, 2024 (regulated entities); June 30, 2024 (small businesses)
Regulator	Washington Attorney General; private right of action under Consumer Protection Act
Scope	Legal entities that conduct business in Washington or produce or provide products or services targeted to consumers in Washington and that, alone or jointly with others, determine the purpose and means of collecting, processing, sharing, or selling consumer health data

Applicability

The MHMDA was the first-mover state "consumer health data" law and reaches far beyond HIPAA-covered entities - fitness apps, wellness platforms, mental health applications, fertility tracking, and

(notably) any AI system that infers health-related information from non-health data. Several other states have followed (Nevada's SB 370, Connecticut's amendment, New York's pending bill).

Core Obligations

- Provide a Consumer Health Data Privacy Policy distinct from the general privacy policy.
- Obtain affirmative consent before collecting or sharing consumer health data; obtain valid authorization (a heightened standard) before selling consumer health data.
- Maintain a list of categories of consumer health data shared and the third parties or affiliates with which it is shared.
- Honor consumer rights of access, deletion, and withdrawal of consent.
- Implement reasonable security; restrict access to consumer health data to employees, processors, and contractors with a need to know.
- Do not implement geofences around in-person healthcare services to identify, track, or send notifications to consumers.

Penalties & Enforcement

Enforced under the Washington Consumer Protection Act; penalties up to \$7,500 per violation. The private right of action has produced significant litigation activity since 2024.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Inference scoping is critical: AI systems that derive health attributes from purchase history, search behavior, or location data create consumer health data within the meaning of the act.
Layer 3 - System Integration	Vendor and affiliate sharing requires distinct consent infrastructure; HIPAA BAAs do not satisfy MHMDA requirements where the data is not PHI.
Layer 4 - Control & Monitoring	The geofence prohibition is enforceable against ad-tech and customer-engagement integrations - Layer 4 deployment gates should preclude geofences around healthcare facilities.

PRACTITIONER NOTE

The MHMDA's reach is broader than its name suggests. If your AI system infers any health-related attribute - fertility, mood, addiction risk, sleep quality, chronic condition - assume MHMDA scope and build a separate consent flow.

New York SHIELD Act and NY DFS 23 NYCRR Part 500

N.Y. Gen. Bus. Law § 899-bb (SHIELD); 23 NYCRR Part 500

Jurisdiction	New York
---------------------	----------

Effective	SHIELD Act 2020; Part 500 originally 2017; substantial Part 500 amendments effective November 2023, with phased compliance through November 2025
Regulator	New York Attorney General (SHIELD); New York Department of Financial Services (Part 500)
Scope	SHIELD: any person or business that owns or licenses computerized data including private information of NY residents. Part 500: financial services entities licensed by NY DFS

Applicability

The SHIELD Act imposes data security obligations and breach notification on any entity holding NY residents' private information. NY DFS Part 500 is the most prescriptive U.S. cybersecurity regulation: explicit governance, risk assessment, MFA, encryption, third-party risk management, incident response, and CISO certification requirements. The November 2023 amendments substantially tightened obligations including 72-hour notification for cybersecurity events, expanded multi-factor authentication, and explicit recognition of AI-related risks in risk assessment.

Core Obligations

- SHIELD: implement reasonable safeguards (administrative, technical, physical); notify affected NY residents of breaches involving private information.
- Part 500: maintain a written cybersecurity program, designate a CISO, conduct annual risk assessments, implement MFA for all individuals accessing internal networks from external networks, encrypt nonpublic information in transit and at rest (or equivalent compensating controls), maintain audit trails, restrict privileged access, deploy continuous monitoring or annual penetration testing plus biannual vulnerability assessments.
- Notify DFS within 72 hours of determination that a "cybersecurity event" has occurred (broadened in 2023 amendments).
- Annual CISO certification of compliance to the board (in lieu of, or in addition to, CEO certification per 2023 amendments).
- Vendor risk management with prescribed elements; annual reassessment.

Penalties & Enforcement

SHIELD: civil penalties up to \$5,000 per violation. Part 500: penalties of up to \$1,000 per violation under DFS general penalty authority; settlements have ranged from \$1.5M (FNTG, 2021) to \$40M (PayPal, 2025) for systemic failures.

Recent Developments (through 2026)

NY DFS issued AI-specific guidance in 2024 addressing the risks AI poses to covered entities and the application of Part 500 to AI systems. The 2023 Part 500 amendments and 2024 AI guidance together constitute the most explicit AI cybersecurity regime among U.S. financial regulators.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Nonpublic information used in AI training inherits Part 500 encryption and access requirements.
Layer 2 - Model Governance	AI risk assessment is now an explicit Part 500 obligation; model risk and information security risk converge in the assessment artifact.
Layer 3 - System Integration	Third-party AI providers fall under § 500.11 (third-party service provider security policy); the 2024 AI guidance specifically calls out AI inference-as-a-service as a high-attention third-party category.
Layer 4 - Control & Monitoring	Continuous monitoring and access management apply to AI development environments equally with production data stores.
Layer 5 - Audit & Evidence	CISO certification creates personal accountability; the certification artifact and the supporting evidence base are the principal Layer 5 deliverables for NY DFS-regulated entities.

PRACTITIONER NOTE

NY DFS Part 500 is the de facto cybersecurity floor for U.S. financial services AI. If your governance program satisfies Part 500, it likely satisfies most state and sector-specific cybersecurity requirements; the inverse is rarely true. The CISO certification creates personal accountability that survives executive turnover.

COMMON FAILURE PATTERN

Third-party AI providers are mapped under § 500.11 only at procurement, then refreshed annually as a paperwork exercise. Part 500's 2024 AI guidance and the 72-hour cybersecurity-event clock require continuous monitoring of upstream AI providers, not periodic attestation. When a model-hosting provider is breached, the covered entity must determine within 72 hours whether nonpublic information was implicated and certify the determination to DFS. Build the third-party AI inventory as a live system fed by procurement, security questionnaires, and contract-renewal events, not a spreadsheet refreshed quarterly.

Comprehensive State Privacy Laws - The Second and Third Waves

The 2024–2026 wave of state privacy enactments - Iowa, Indiana, Tennessee (with NIST PF affirmative defense), Montana, New Jersey, Delaware, New Hampshire, Kentucky, Maryland (strict minimization), Minnesota (mandatory data inventory), Rhode Island, Nebraska, and Florida - completes the multistate privacy landscape and establishes operational expectations that exceed what California originally established. Practitioners should expect the substantive convergence to continue: data minimization, profiling assessments, UOOM honoring, and minor protections are the recurring themes.

Iowa Consumer Data Protection Act (ICDPA)

Iowa Code §§ 715D.1 to 715D.9

Jurisdiction	Iowa
Effective	January 1, 2025
Regulator	Iowa Attorney General
Scope	Persons conducting business in Iowa or producing products/services targeted to Iowa residents that during a calendar year (a) control or process personal data of 100,000+ Iowa consumers, or (b) control or process data of 25,000+ Iowa consumers and derive 50%+ of gross revenue from sale of personal data

Applicability

The ICDPA is among the most business-friendly of the comprehensive state privacy laws - it lacks several rights present elsewhere (no opt-out of profiling, no data correction right) and provides a 90-day cure period that does not sunset. Sensitive data requires opt-out (rather than opt-in as in most peer states). For AI practitioners, Iowa's structural posture means a multi-state baseline built to Virginia or Colorado will satisfy Iowa, but Iowa-specific compliance assessments should reflect the narrower obligations.

Core Obligations

- Provide privacy notices identifying categories of personal data, purposes, third-party recipients, and consumer rights mechanisms.
- Honor consumer rights of access, deletion, portability, and opt-out of sale and targeted advertising.
- Provide opt-out (not opt-in) for sensitive data processing.
- Maintain processor contracts with prescribed terms.
- Respond to consumer requests within 90 days.

Penalties & Enforcement

Civil penalties up to \$7,500 per violation; permanent 90-day cure period.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

Iowa's narrower obligation set means Layer 5 documentation can be simpler - but inventories should explicitly note where Iowa carve-outs apply for audit clarity.

Indiana Consumer Data Protection Act (INCDPA)

Ind. Code §§ 24-15-1 et seq.

Jurisdiction	Indiana
---------------------	---------

Effective	January 1, 2026
Regulator	Indiana Attorney General
Scope	Persons conducting business in Indiana or producing products/services targeted to Indiana residents that during a calendar year (a) control or process personal data of 100,000+ Indiana consumers, or (b) control or process data of 25,000+ Indiana consumers and derive 50%+ of gross revenue from sale of personal data

Applicability

The INCDPA closely tracks the Virginia template with a delayed 2026 effective date and a 30-day cure period that does not sunset. Sensitive data requires opt-in consent. Data protection assessments are required for high-risk processing including profiling presenting reasonably foreseeable risks of unfair or deceptive treatment, financial/physical/reputational injury, or intrusion upon seclusion.

Core Obligations

- Provide privacy notices.
- Honor consumer rights of access, correction, deletion, portability, and opt-out of sale, targeted advertising, and certain profiling.
- Obtain opt-in consent for sensitive data processing.
- Conduct data protection assessments for high-risk processing.
- Maintain processor contracts.
- Respond to consumer requests within 45 days.

Penalties & Enforcement

Civil penalties up to \$7,500 per violation; 30-day cure period.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	DPA's for profiling track Virginia/Colorado; multi-state assessment programs satisfy Indiana with documentation tagged to Indiana scope.

Tennessee Information Protection Act (TIPA)

Tenn. Code Ann. §§ 47-18-3201 to 47-18-3213

Jurisdiction	Tennessee
Effective	July 1, 2025
Regulator	Tennessee Attorney General
Scope	Persons conducting business in Tennessee or producing products/services targeted to Tennessee residents that exceed \$25M in revenue and (a) control or

	process personal data of 175,000+ Tennessee consumers, or (b) control or process data of 25,000+ consumers and derive 50%+ of revenue from sale of personal data
--	--

Applicability

TIPA is distinguished by its NIST Privacy Framework affirmative defense - controllers and processors that "create, maintain, and comply with" a written privacy program reasonably conforming to the NIST Privacy Framework (or comparable standards) are entitled to an affirmative defense to causes of action under the act. This is the first U.S. state privacy law with an explicit standards-based safe harbor and a model practitioners are watching for replication.

Core Obligations

- Standard Virginia-pattern obligations: notices, consumer rights, opt-in for sensitive data, DPAs for high-risk processing, processor contracts.
- For affirmative defense: maintain a documented privacy program reasonably conforming to NIST Privacy Framework, ISO 27701, or comparable; scale program to organization size, personal data volume/variety/sensitivity, and consumer protection cost.
- Respond to consumer requests within 45 days.

Penalties & Enforcement

Treble damages up to \$15,000 per violation for willful violations; 60-day cure period.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	NIST Privacy Framework conformance becomes a Layer 5 affirmative-defense artifact - the most explicit U.S. statutory recognition of a standards-based safe harbor.

PRACTITIONER NOTE
 Even outside Tennessee, the TIPA structure illustrates the strategic value of standards-based privacy programs. Documented NIST PF or ISO 27701 alignment supports defensibility in inquiries from other AGs even where no statutory safe harbor exists.

Montana Consumer Data Privacy Act (MCDPA)

Mont. Code Ann. §§ 30-14-2801 to 30-14-2817

Jurisdiction	Montana
Effective	October 1, 2024

Regulator	Montana Attorney General
Scope	Persons conducting business in Montana or producing products/services targeted to Montana residents that (a) control or process personal data of 50,000+ Montana consumers excluding payment transactions, or (b) control or process data of 25,000+ consumers and derive 25%+ of gross revenue from sale of personal data

Applicability

The MCDPA has the lowest consumer-volume threshold of any state privacy law (50,000 vs. 100,000 in most peers), capturing more mid-market businesses. It otherwise tracks the Virginia/Connecticut pattern. The cure period sunsets April 1, 2026.

Core Obligations

- Standard Virginia-pattern obligations.
- Honor universal opt-out mechanisms (effective January 1, 2025).
- Conduct DPAs for high-risk processing.
- Respond to consumer requests within 45 days.

Penalties & Enforcement

Penalties under Montana Unfair Trade Practices Act framework; cure period sunsets April 1, 2026.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Lower threshold extends Layer 5 documentation obligations to many mid-market operations not previously subject to comprehensive state privacy law.

New Jersey Data Privacy Act

N.J. Stat. Ann. §§ 56:8-166.4 to 56:8-166.18

Jurisdiction	New Jersey
Effective	January 15, 2025
Regulator	New Jersey Division of Consumer Affairs (Attorney General)
Scope	Controllers conducting business in New Jersey or producing products/services targeted to New Jersey residents that during a calendar year (a) control or process personal data of 100,000+ NJ consumers, or (b) control or process data of 25,000+ NJ consumers and derive revenue or receive a discount on the price of any goods/services from the sale of personal data

Applicability

NJDPA is notable for its expansive sensitive data definition (including financial information, account credentials, racial/ethnic origin, religious beliefs, mental/physical health condition, sex life/sexual orientation, citizenship/immigration, status as transgender or non-binary, status as victim of crime, and biometric/genetic), profiling-related opt-out rights, and a particularly aggressive AG. The act adopts the universal opt-out mechanism (UOOM) framework with an effective date of July 15, 2025 for honoring such signals.

Core Obligations

- Standard Virginia-pattern obligations with expanded sensitive data list.
- Opt-in consent for sensitive data processing.
- DPAs for processing presenting heightened risk including targeted advertising, sale, profiling for unfair or deceptive treatment risk, processing sensitive data, and processing of children's data.
- Honor UOOM signals (effective July 15, 2025).
- Honor consumer rights including correction.
- Children (13–17): obtain consent before targeted advertising, sale, or profiling for decisions producing legal/similarly significant effects.

Penalties & Enforcement

Penalties under New Jersey Consumer Fraud Act (CFA) framework - potentially substantial including treble damages; 18-month cure period for first violations (sunsets July 1, 2026).

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Expanded sensitive data list captures more AI training inputs as opt-in protected; data classification schemes need NJ-specific tagging.
Layer 4 - Control & Monitoring	UOOM honoring infrastructure must respect NJ's July 2025 effective date as part of multi-state opt-out enforcement.

Delaware Personal Data Privacy Act

Del. Code Ann. tit. 6, §§ 12D-101 to 12D-114

Jurisdiction	Delaware
Effective	January 1, 2025
Regulator	Delaware Department of Justice (Consumer Protection Unit)
Scope	Persons conducting business in Delaware or producing products/services targeted to Delaware residents that (a) control or process personal data of 35,000+ Delaware consumers excluding payment transactions, or (b) control or process data of 10,000+ Delaware consumers and derive 20%+ of gross revenue from sale of personal data

Applicability

Delaware's law is distinguished by its low thresholds (35,000 / 10,000 consumers - the lowest in the U.S.), its inclusion of nonprofits (alongside Oregon and Colorado), and its expansive children's data protections. The act otherwise tracks the Connecticut pattern.

Core Obligations

- Standard Connecticut-pattern obligations.
- Lower thresholds capture small and mid-market businesses.
- Includes nonprofits (with limited exceptions).
- Children: opt-in consent for processing of personal data of consumers known to be 13–17 for targeted advertising, sale, or profiling.
- Honor UOOM signals from January 1, 2026.
- Provide list of categories of third parties to which personal data has been disclosed.

Penalties & Enforcement

Penalties under Delaware Consumer Fraud Act; up to \$10,000 per willful violation under analogous frameworks. Cure period sunset 60 days from law's effective date.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

Third-party-category disclosure obligations require Layer 5 disclosure logging; combined with Oregon's identity-level disclosure, Delaware's category-level disclosure is the second tier.

New Hampshire Data Privacy Act

N.H. Rev. Stat. Ann. §§ 359-NN:1 to 359-NN:11

Jurisdiction	New Hampshire
Effective	January 1, 2025
Regulator	New Hampshire Attorney General
Scope	Persons conducting business in New Hampshire or producing products/services targeted to NH residents that during a calendar year (a) control or process personal data of 35,000+ NH consumers, or (b) control or process data of 10,000+ consumers and derive 25%+ of gross revenue from sale of personal data

Applicability

NHDPA tracks the Connecticut pattern with low thresholds matching Delaware. The Secretary of State is empowered to issue rules; cure period sunsets January 1, 2026.

Core Obligations

- Standard Connecticut-pattern obligations.
- Honor UOOM signals from January 1, 2025.
- Opt-in consent for sensitive data; opt-in for children (13–17) targeted advertising, sale, and profiling for decisions producing legal/similarly significant effects.

Penalties & Enforcement

Civil penalties under New Hampshire Consumer Protection Act framework.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Standard New England multi-state compliance posture; documentation can largely mirror Connecticut with NH-specific scope tagging.

Kentucky Consumer Data Protection Act

Ky. Rev. Stat. §§ 367.3611 to 367.3637

Jurisdiction	Kentucky
Effective	January 1, 2026
Regulator	Kentucky Attorney General
Scope	Persons conducting business in Kentucky or producing products/services targeted to Kentucky residents that during a calendar year (a) control or process personal data of 100,000+ KY consumers, or (b) control or process data of 25,000+ consumers and derive 50%+ of gross revenue from sale of personal data

Applicability

Kentucky's 2024 act tracks the Virginia template and includes a permanent 30-day cure period. It does not include UOOM honoring or correction rights.

Core Obligations

- Standard Virginia-pattern obligations including DPAs for high-risk processing.
- Opt-in consent for sensitive data.
- Respond to consumer requests within 45 days.

Penalties & Enforcement

Civil penalties up to \$7,500 per violation; permanent 30-day cure period.

STACK LENS - How this law maps to the AI Governance Stack**Layer 5 - Audit & Evidence**

Permanent cure period reduces immediate enforcement exposure but does not affect underlying compliance obligation; Layer 5 documentation should remain assessment-ready.

Maryland Online Data Privacy Act (MODPA)

Md. Code Ann., Com. Law §§ 14-4601 to 14-4619

Jurisdiction	Maryland
Effective	October 1, 2025
Regulator	Maryland Attorney General (Consumer Protection Division)
Scope	Persons conducting business in Maryland or producing products/services targeted to MD residents that during a calendar year (a) control or process personal data of 35,000+ MD consumers, or (b) control or process data of 10,000+ MD consumers and derive 20%+ of gross revenue from sale of personal data

Applicability

Maryland's MODPA is among the most consumer-protective state privacy laws. It includes a strict data minimization requirement (controllers may collect personal data only as is reasonably necessary and proportionate to provide or maintain a specific product or service requested by the consumer), a per-se prohibition on selling sensitive data, prohibitions on targeted advertising or sale of consumer data of consumers under 18, and unique constraints on "consumer health data."

Core Obligations

- Strict data minimization: collection limited to what is reasonably necessary and proportionate to provide the specific product/service requested.
- Per-se prohibition on the sale of sensitive personal data.
- Prohibition on targeted advertising or sale of personal data of consumers known to be under 18.
- Heightened restrictions on consumer health data including consent and prohibition on geofencing around healthcare facilities.
- Standard rights including access, correction, deletion, portability, opt-out of sale, targeted advertising, and profiling for decisions producing legal/similarly significant effects.
- Honor UOOM signals.
- DPAs for high-risk processing.

Penalties & Enforcement

Penalties under Maryland Consumer Protection Act; up to \$10,000 per violation, \$25,000 per repeat violation; 60-day cure period sunset April 1, 2027.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Strict data minimization is the most aggressive Layer 1 collection constraint in U.S. state law; AI training data acquisition must be tied to specific product/service necessity.
Layer 2 - Model Governance	Per-se sensitive-data sale prohibition forecloses certain Layer 2 monetization patterns common in adtech-funded AI.
Layer 3 - System Integration	Health-data and minor-targeting restrictions reach Layer 3 advertising integration architecture.

Minnesota Consumer Data Privacy Act (MNCDPA)

Minn. Stat. §§ 3250.01 to 3250.16

Jurisdiction	Minnesota
Effective	July 31, 2025
Regulator	Minnesota Attorney General
Scope	Persons conducting business in Minnesota or producing products/services targeted to MN residents that during a calendar year (a) control or process personal data of 100,000+ MN consumers, or (b) control or process data of 25,000+ MN consumers and derive 25%+ of gross revenue from sale of personal data

Applicability

MNCDPA tracks the Connecticut pattern with several Minnesota-specific additions: a right to question results of profiling decisions and review the data used; a data inventory requirement for controllers; and a chief privacy officer or equivalent oversight requirement (though not formally titled DPO). The act is among the most prescriptive in operational terms.

Core Obligations

- Maintain a documented inventory of personal data and a written data security and management policy.
- Designate a person responsible for compliance (de facto DPO obligation).
- Honor consumer right to question profiling decisions producing legal or similarly significant effects, review the data used, and request reassessment.
- Honor UOOM signals.
- DPAs for high-risk processing.
- Standard Connecticut-pattern obligations otherwise.

Penalties & Enforcement

Civil penalties under Minnesota Consumer Fraud Act; cure period available at AG discretion through January 31, 2026 then sunset.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Mandatory data inventory is the most explicit Layer 1 documentation requirement in U.S. state law.
Layer 2 - Model Governance	Right to question profiling and review data creates Layer 2 explainability obligation parallel to GDPR Article 22.
Layer 5 - Audit & Evidence	Designated compliance person and documented inventory become Layer 5 audit anchors.

PRACTITIONER NOTE
 Minnesota's data inventory and named-person obligations push U.S. practice toward GDPR-style accountability infrastructure. Organizations with multistate exposure benefit from building these artifacts uniformly rather than treating Minnesota as a unique exception.

Rhode Island Data Transparency and Privacy Protection Act

R.I. Gen. Laws §§ 6-48.1-1 et seq.

Jurisdiction	Rhode Island
Effective	January 1, 2026
Regulator	Rhode Island Attorney General
Scope	Commercial websites or internet services that conduct business in Rhode Island, customers in Rhode Island, or are controlled by an entity that does

Applicability

Rhode Island's law is unique in its hybrid structure - it combines comprehensive privacy obligations applying to entities meeting consumer-volume thresholds with a separate disclosure-only obligation applying to all commercial websites and internet services (regardless of size) requiring identification of categories of personal information collected and third parties to whom personal information is sold or disclosed. The disclosure obligation is broader in scope than substantive obligations.

Core Obligations

- For all in-scope commercial websites/internet services: disclose specific categories of personal information collected and sold/disclosed and the identity of third parties to whom personal data is sold.

- For larger entities meeting consumer thresholds: standard Virginia-pattern obligations including consumer rights, opt-in for sensitive data, and DPAs.
- Sale of biometric data prohibited without consent.

Penalties & Enforcement

Civil penalties up to \$10,000 per knowing or willful violation under Deceptive Trade Practices Act framework.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Universal disclosure obligation creates broader Layer 5 transparency baseline than peer state laws.

Nebraska Data Privacy Act

Neb. Rev. Stat. §§ 87-1101 to 87-1115

Jurisdiction	Nebraska
Effective	January 1, 2025
Regulator	Nebraska Attorney General
Scope	Persons conducting business in Nebraska or producing products/services consumed by Nebraska residents that process or engage in the sale of personal data, and that are not small businesses (per SBA definition); no consumer-volume threshold

Applicability

NDPA mirrors the Texas TDPSA structure - no consumer-volume threshold, small-business carve-out only, opt-in consent for sensitive data. AI practitioners with even small Nebraska footprints are routinely in scope.

Core Obligations

- Standard Virginia-pattern obligations.
- Opt-in consent for sensitive data.
- DPAs for high-risk processing.
- Respond to consumer requests within 45 days.

Penalties & Enforcement

Penalties under Nebraska Consumer Protection Act framework; 30-day cure period.

STACK LENS - How this law maps to the AI Governance Stack

Layer 1 - Data Governance	No-threshold scope captures small AI vendors that escape California and Virginia thresholds; data inventory must include all Nebraska consumer data.
----------------------------------	--

Florida Digital Bill of Rights

Fla. Stat. §§ 501.701 to 501.722

Jurisdiction	Florida
Effective	July 1, 2024
Regulator	Florida Department of Legal Affairs (Attorney General)
Scope	Controllers that (a) make in excess of \$1B in global gross annual revenues and (b) satisfy at least one of: derive 50%+ of revenue from online ad sale, operate a consumer smart speaker/voice command service, or operate an app store / digital distribution platform offering at least 250,000 different software applications

Applicability

Florida's law is targeted at very large platforms - the threshold limits substantive obligations to a small set of major technology companies. The act includes notable provisions including restrictions on targeted advertising to minors, government-controlled-account moderation prohibitions, and search engine result transparency. AI implications are concentrated on app-store and platform operators.

Core Obligations

- For in-scope controllers: standard Virginia-pattern privacy rights including opt-out of sale, targeted advertising, and profiling.
- Opt-in consent for sensitive data.
- Restrictions on targeted advertising to consumers under 18.
- Search engine result transparency obligations.

Penalties & Enforcement

Civil penalties up to \$50,000 per violation; treble damages for violations involving children, dishonoring opt-out signals, or knowingly violating with respect to a known minor.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	High-revenue threshold concentrates obligations on platform operators; for in-scope entities, Florida-specific compliance documentation is required even where multistate programs satisfy other states.
---------------------------------------	--

State Cybersecurity, Breach Notification, and Sector Privacy Regimes

Beyond comprehensive privacy law, U.S. practitioners navigate the leading state cybersecurity regulations (NY DFS Part 500, Massachusetts 201 CMR 17), the all-50-state breach notification patchwork, sector-specific student and biometric privacy statutes, and the SHIELD Act's data security obligations. These regimes operate independently of the comprehensive privacy laws and apply broadly across industry sectors.

U.S. State Data Breach Notification Laws (Survey)

All 50 states + D.C., Puerto Rico, Guam, U.S. Virgin Islands have breach notification statutes

Jurisdiction	All U.S. states and territories
Effective	California 2003 (first); all states by 2018; continuous amendment
Regulator	State attorneys general; some states (e.g., Massachusetts, New York) require regulator notification independent of consumer notice
Scope	Persons or businesses owning, licensing, or maintaining computerized data containing personal information of state residents

Applicability

The U.S. lacks a comprehensive federal breach notification law (sectoral laws aside); the operative regime is the patchwork of 54 jurisdictions. AI systems holding training data or inference logs containing personal information are subject to all jurisdictions where affected individuals reside. Common elements include: trigger events (unauthorized acquisition or access; some states require materiality); covered data definitions (typically name + SSN/DL/financial account; expanding to include biometric, health, login credentials, medical, and increasingly genetic and student data); notification timing (varying from "without unreasonable delay" to specific 30-, 45-, 60-, or 90-day clocks); regulator notice (required in approximately 30 jurisdictions); and notice content prescriptions.

Core Obligations

- Determine notification obligations based on each affected resident's state of residence - not the entity's headquarters or data location.
- Apply the state's definition of "personal information" - substantially divergent across states (e.g., Illinois includes biometric data; Maryland includes passport numbers; Connecticut includes biometric data and email-plus-password combinations).
- For larger breaches, comply with state regulator notification (e.g., NY DFS 72h, Indiana AG, California AG/web posting for 500+ residents).
- Use specific notification content elements where prescribed (e.g., Maryland and Iowa specify required content).

- Comply with credit monitoring obligations where required (e.g., California, Connecticut for SSN breaches).
- For HIPAA-covered information, notify HHS and (for 500+) media in addition to state notification.

Penalties & Enforcement

Varies by state; per-violation civil penalties (often per resident not notified) typical; private rights of action in several states (CA, AK, IL, MA among others). Massachusetts and New York actively pursue breach matters under SHIELD Act and 201 CMR 17 frameworks. Multistate AG coordinated enforcement is increasingly common (Equifax, Marriott, etc.).

Recent Developments (through 2026)

NY SHIELD Act, Maryland's 2025 expansion (genetic data), Connecticut's 2024 amendments (biometric data, expanded credentials), and California's ongoing CCPA cybersecurity audit regulations have continued the expansion trajectory. AI-relevant breach scenarios - exposure of training data, model artifacts containing memorized PII, prompt logs - present novel notification analyses under several state regimes.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	Breach detection and triage workflows require Layer 4 monitoring capable of identifying not just access events but unauthorized acquisition; AI inference-time disclosure events may trigger notification obligations distinct from training-data breaches.
Layer 5 - Audit & Evidence	Multistate notification readiness requires Layer 5 templates pre-tailored by jurisdiction; the patchwork makes 60-day clocks operationally tight.

PRACTITIONER NOTE
 AI training-data and prompt-log breaches present novel analyses: training data may be derived rather than directly collected; prompt logs may contain personal information not previously inventoried. Build the breach playbook with explicit AI-incident classifications mapped to each state's notification trigger language.

Washington Biometric Identifier Statute (RCW 19.375)

Wash. Rev. Code §§ 19.375.010 to 19.375.900

Jurisdiction	Washington
Effective	July 23, 2017
Regulator	Washington Attorney General (no private right of action under RCW 19.375; private action under Consumer Protection Act may be available)
Scope	Persons enrolling biometric identifiers in databases for commercial purposes

Applicability

Washington's biometric law applies a notice-and-consent regime to enrollment of biometric identifiers (fingerprints, voiceprints, eye retinas/irises, biometric data derived from face geometry) in commercial databases. Distinct from BIPA, the law lacks a private right of action and the AG has not historically enforced it aggressively, but the statute remains operative and its consent requirements affect AI biometric pipelines targeting Washington residents.

Core Obligations

- Provide notice and obtain consent (or comply with another permissible mechanism) before enrolling a biometric identifier in a commercial database.
- Take reasonable care to guard against unauthorized access.
- Limit retention to that necessary to provide the service for which the biometric identifier was obtained.

Penalties & Enforcement

AG enforcement under Consumer Protection Act framework; per-violation penalties up to \$7,500.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Notice/consent for enrollment is a Layer 1 collection gate; "enrollment in a commercial database" is the trigger, not mere capture.

New York Biometric and Student Data Laws

N.Y. Educ. Law § 2-d (Student Privacy); pending NYS 1144 / A 2434 (Biometric Privacy Act, modeled on BIPA)

Jurisdiction	New York
Effective	NY Ed. Law § 2-d effective 2014; biometric reform pending
Regulator	New York State Education Department (student privacy); NYS Attorney General; private right under proposed biometric law
Scope	Student privacy: educational agencies and third-party contractors of educational agencies; biometric (if enacted): private entities collecting biometric identifiers

Applicability

New York Education Law § 2-d is among the strictest student data privacy regimes in the U.S., applying directly to EdTech vendors in New York and requiring contractual obligations regarding data security, parental notification, breach response, and limitations on data use. The pending NY biometric privacy

bill - modeled on BIPA - would create a private right of action for biometric data violations and substantially expand New York's privacy enforcement landscape.

Core Obligations

- NY Ed. Law § 2-d: educational agencies must require third-party contractors to comply with privacy obligations including the agency's data security and privacy policy; contractors must adopt the NIST Cybersecurity Framework; parental rights to inspect and challenge student data; mandatory breach notification within prescribed periods.
- Pending biometric bill: written notice and consent before capture; written retention/destruction policy; prohibition on sale; reasonable care; private right of action.

Penalties & Enforcement

NY Ed. Law § 2-d: AG enforcement; loss of state contract eligibility. Pending biometric bill: BIPA-style statutory damages.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	NIST CSF adoption requirement for ed-data contractors creates Layer 1/4 framework alignment as a contract term.
Layer 5 - Audit & Evidence	NY Ed. Law contractor obligations require privacy and security documentation tailored to district contracts.

Colorado Student Data Transparency and Security Act

Colo. Rev. Stat. §§ 22-16-101 et seq.

Jurisdiction	Colorado
Effective	August 10, 2016
Regulator	Colorado Department of Education
Scope	School service providers and school service on-demand providers handling personally identifiable information from K-12 students

Applicability

Colorado's student privacy act prohibits use of student data for targeted advertising, restricts secondary uses, requires data security, and requires Colorado Department of Education to maintain a public list of school service contractors. The act's on-demand provider framework was the first U.S. state regulation of free EdTech tools used in classrooms.

Core Obligations

- Do not use student data for targeted advertising or to develop a profile for non-school purposes.

- Maintain reasonable security.
- Honor deletion requests and provide data to parents/students on request.
- School service contract providers: provide annual lists; comply with contract requirements.

Penalties & Enforcement

Department of Education enforcement; loss of approved-vendor status; potential civil action.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Targeted advertising prohibition forecloses certain Layer 1 monetization patterns in EdTech AI tools.

California Student Online Personal Information Protection Act (SOPIPA)

Cal. Bus. & Prof. Code §§ 22584 to 22585

Jurisdiction	California
Effective	January 1, 2016
Regulator	California Attorney General
Scope	Operators of websites, online services, online applications, and mobile applications used primarily for K-12 school purposes and designed and marketed for K-12 school purposes

Applicability

SOPIPA prohibits operators from using K-12 student data for targeted advertising, building profiles for non-school purposes, selling student data, or disclosing student data except under enumerated exceptions. AI-driven EdTech tools are squarely in scope. SOPIPA was the U.S. template for many subsequent state student privacy laws.

Core Obligations

- Do not engage in targeted advertising directed at students or based on covered information.
- Do not amass a profile about a K-12 student except in furtherance of K-12 school purposes.
- Do not sell student data.
- Do not disclose student data except for enumerated school-supporting purposes.
- Implement and maintain reasonable security; delete student data upon request of school or district.

Penalties & Enforcement

AG enforcement under unfair business practices framework; civil penalties up to \$2,500 per violation under California Unfair Competition Law.

STACK LENS - How this law maps to the AI Governance Stack

Layer 1 - Data Governance	Profile-building prohibition restricts Layer 1 derived-data practices in EdTech AI.
Layer 3 - System Integration	Disclosure restrictions reach Layer 3 EdTech vendor architecture.

New York Stop Hacks and Improve Electronic Data Security (SHIELD) - Detailed*N.Y. Gen. Bus. Law §§ 899-aa, 899-bb*

Jurisdiction	New York
Effective	March 21, 2020 (security obligations); July 25, 2019 (notification expansion)
Regulator	New York Attorney General
Scope	Any person or business that owns or licenses computerized data including private information of NY residents

Applicability

The SHIELD Act extends NY breach notification obligations to incidents involving "access" to private information (not just acquisition), expands the definition of private information (account credentials, biometric data, certain credit account combinations), and requires reasonable safeguards (administrative, technical, physical). The "reasonable safeguards" obligation reaches all entities holding NY private information regardless of size or industry.

Core Obligations

- Notify affected NY residents of breaches involving private information without unreasonable delay; notify NY AG, Consumer Protection Bureau, and NY State Police for breaches affecting 5,000+ residents.
- Implement reasonable safeguards: designated employee, identification of risks, training, vendor risk management, written program (with safe harbors for HIPAA-covered, GLBA-covered, NY DFS Part 500-regulated, and Title V GLBA-aligned entities).
- Small businesses (fewer than 50 employees, less than \$3M annual revenue, less than \$5M in year-end total assets): scaled safeguards.

Penalties & Enforcement

Civil penalties up to \$5,000 per violation under § 899-aa; injunctive relief and damages under § 899-bb.

STACK LENS - How this law maps to the AI Governance Stack

Layer 4 - Control & Monitoring	Reasonable safeguards include vendor risk management touching AI subprocessors; "access" trigger requires Layer 4 access monitoring beyond exfiltration detection.
---	--

Massachusetts 201 CMR 17 - Standards for the Protection of Personal Information

201 CMR 17.00; M.G.L. c. 93H; M.G.L. c. 93I

Jurisdiction	Massachusetts
Effective	201 CMR 17 effective March 1, 2010
Regulator	Office of the Attorney General; Office of Consumer Affairs and Business Regulation
Scope	Any person that owns or licenses personal information about a Massachusetts resident

Applicability

Massachusetts has the most prescriptive U.S. state data security regulation outside the financial sector. 201 CMR 17 mandates a written information security program (WISP), encryption of personal information transmitted across public networks or stored on portable devices, secure user authentication, monitoring, training, and oversight of third-party service providers. The WISP requirement reaches AI vendors handling Massachusetts personal information.

Core Obligations

- Develop, implement, and maintain a comprehensive WISP including: designated employee responsibility, risk identification, employee training, third-party service provider oversight (contractual security obligations and selection due diligence), reasonable monitoring, post-incident review, and updating.
- Encrypt personal information transmitted across public networks and on portable/wireless devices.
- Implement secure user authentication and access controls.
- Restrict physical access to records containing personal information.
- Reasonably monitor systems for unauthorized use or access.

Penalties & Enforcement

Civil penalties up to \$5,000 per violation under M.G.L. c. 93A; private right of action under c. 93A; AG enforcement.

STACK LENS - How this law maps to the AI Governance Stack

Layer 4 - Control & Monitoring	WISP and 201 CMR 17 controls map directly onto Layer 4 operational security; AI-specific monitoring requirements integrate naturally with the framework.
Layer 5 - Audit & Evidence	Documented WISP is a foundational Layer 5 artifact; many breach defenses turn on WISP adequacy.

P A R T I I I

United States - State AI-Specific Laws

State statutes specifically targeting AI systems - consequential decisions, employment, generative AI transparency, frontier model regulation, deepfakes, biometric likenesses, and sectoral AI regimes.

Leading State AI Statutes

State legislatures have moved aggressively into AI-specific regulation in the absence of comprehensive federal AI legislation. Colorado, Texas, California, Illinois, and New York City have produced the leading enacted frameworks. The cross-jurisdictional pattern is converging on (1) consequential-decision impact assessments, (2) employer transparency obligations, and (3) developer documentation duties - even where the specific statutory text varies.

Colorado AI Act (CAIA, SB 24-205)

Colo. Rev. Stat. §§ 6-1-1701 to 6-1-1707

Jurisdiction	Colorado
Effective	February 1, 2026 (delayed from original; phased)
Regulator	Colorado Attorney General
Scope	Developers and deployers of high-risk AI systems doing business in Colorado

Applicability

The Colorado AI Act is the first comprehensive U.S. state AI consequential-decision law. It applies to "high-risk AI systems" - those that, when deployed, make or are a substantial factor in making a consequential decision (employment, education, financial/lending, essential government services, healthcare, housing, insurance, legal services). Both developers and deployers have distinct obligations; the act adopts a risk-based framework reminiscent of the EU AI Act but with U.S. private-litigation exposure.

Core Obligations

- **Developers:** provide deployers with documentation describing system's intended uses, known/foreseeable harms, training data summaries, evaluation results, mitigation measures, and information sufficient for deployers to comply with their obligations.
- **Developers:** publicly publish a high-level summary of types of high-risk AI systems developed or substantially modified.
- **Developers:** notify the AG within 90 days of discovering that a high-risk AI system has caused or is reasonably likely to have caused algorithmic discrimination.
- **Deployers:** implement risk management policy and program covering the use of high-risk AI systems including governance, data, transparency, monitoring, and incident response.
- **Deployers:** complete an impact assessment for each high-risk AI system at least annually and after any intentional and substantial modification.

- Deployers: notify consumers when high-risk AI systems make or are a substantial factor in making consequential decisions, provide explanation rights, and offer opportunity to correct erroneous personal data.
- Deployers: provide consumers the opportunity to appeal adverse consequential decisions, including human review where technically feasible.
- Both: use reasonable care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination.

Penalties & Enforcement

Enforcement exclusively by the AG (no private right of action). Violations are deceptive trade practices under Colo. Rev. Stat. § 6-1-105; remedies include injunctive relief and civil penalties up to \$20,000 per violation. The act includes a rebuttable presumption that compliance with NIST AI RMF and similar standards constitutes use of reasonable care.

Recent Developments (through 2026)

The act was the subject of substantial 2024–2025 amendment efforts; the legislature delayed the effective date and the AG has been active in stakeholder rulemaking. The framework is closely watched as the U.S. precedent for state-level AI consequential-decision regulation; California, Connecticut, New York, and others have introduced similar bills.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Developer documentation requirements force Layer 1 transparency - training data summaries are now externally disclosed.
Layer 2 - Model Governance	Algorithmic discrimination prevention obligations require Layer 2 fairness testing across consequential-decision use cases.
Layer 3 - System Integration	The developer-deployer split allocates Layer 3 integration responsibilities; deployers must understand the developer's documentation to satisfy their own obligations.
Layer 4 - Control & Monitoring	Annual deployer impact assessments and consumer appeal rights are Layer 4 monitoring and operational deliverables.
Layer 5 - Audit & Evidence	AG notification within 90 days of discovering algorithmic discrimination is a Layer 5 incident-response obligation; documentation of the discovery process and remediation must be assessment-ready.

PRACTITIONER NOTE

The CAIA's "substantial factor" test for whether AI is making a consequential decision is fact-specific. Build the impact-assessment process to document precisely how AI outputs are used in the decision pipeline; this documentation drives both compliance scope and litigation defense.

Texas Responsible Artificial Intelligence Governance Act (TRAIGA)

Tex. Bus. & Com. Code § 552 (codification pending)

Jurisdiction	Texas
Effective	January 1, 2026 (substantive provisions); state agency provisions earlier
Regulator	Texas Attorney General; Texas AI Council
Scope	Persons that develop or deploy AI systems in Texas; state agencies; certain federal contractors

Applicability

TRAIGA combines a consequential-decision framework with prohibited-practice provisions, state-agency requirements, and a regulatory sandbox. It is narrower than initially drafted (the original 2024 version was substantially amended in 2025) but includes meaningful prohibitions on certain government and commercial AI uses, transparency obligations, and a developer-deployer accountability split.

Core Obligations

- Prohibitions: certain social scoring, manipulation that causes harm, and biometric categorization based on sensitive attributes (mirroring select EU AI Act prohibitions).
- For deployers using AI in consequential ways: provide notice that AI is in use, basic information about the system, and rights of explanation/appeal.
- State agencies must inventory AI systems, conduct impact assessments, and report to the Texas AI Council.
- Regulatory sandbox program permitting controlled testing of innovative AI uses with limited liability protection.
- Healthcare and biometric AI carve-outs with their own specific obligations.

Penalties & Enforcement

AG enforcement; civil penalties scaling with violation type, generally \$10,000 to \$200,000 per violation depending on whether the conduct is curable, incurable, or involves a prohibited practice. Cure periods available for first-time violations of certain provisions.

Recent Developments (through 2026)

TRAIGA is the leading U.S. state AI law of 2025 and represents a Republican-led approach distinct from Colorado's. The Act's sandbox provisions and federal-contractor preemption clauses are the subject of ongoing implementation work.

STACK LENS - How this law maps to the AI Governance Stack

Layer 2 - Model Governance

Prohibited-practice provisions reach Layer 2 model-design choices; biometric categorization features must be evaluated against the prohibitions.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	State-agency vendor obligations cascade to Layer 3 contracts for AI suppliers serving Texas government.
Layer 5 - Audit & Evidence	Sandbox participation requires extensive Layer 5 documentation; the trade-off is liability protection for compliant innovation.

California Generative AI Training Data Transparency Act (AB 2013)

Cal. Bus. & Prof. Code § 22757

Jurisdiction	California
Effective	January 1, 2026
Regulator	California Attorney General
Scope	Developers of generative AI systems or services made available to California residents

Applicability

AB 2013 requires generative AI developers to publicly post documentation describing the data used to train the system, including high-level summaries of dataset sources, time period of data collection, whether the data includes copyrighted, trademarked, or personal information, and whether the data was purchased, licensed, or scraped.

Core Obligations

- Post on the developer's website, prior to making the system publicly available, a documentation describing data used to train the GenAI system.
- Describe the sources or owners of the datasets, the number of data points, and whether the datasets include personal information or aggregate consumer information.
- Describe whether the data was purchased or licensed, modified by the developer, and the time period during which data in the datasets was collected.
- Update the documentation when training data changes materially.

Penalties & Enforcement

Enforcement under California Unfair Competition Law and Consumer Protection statutes; civil penalties and equitable relief available.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Direct Layer 1 transparency: training data provenance becomes a public artifact, constraining acquisition practices upstream.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	Public documentation creates a permanent evidence record subject to enforcement scrutiny; revisions are themselves disclosable artifacts.
---------------------------------------	---

California AI Provenance Act (SB 942) and California AI Transparency Act

Cal. Bus. & Prof. Code §§ 22757.1, 22949.90 et seq.

Jurisdiction	California
Effective	January 1, 2026
Regulator	California Attorney General
Scope	Covered providers of generative AI systems with at least 1,000,000 monthly visitors or users in California

Applicability

SB 942 requires covered GenAI providers to make available an AI detection tool, include both visible and invisible (cryptographic provenance) disclosures in AI-generated content, and provide content authenticity capabilities consistent with C2PA or similar standards.

Core Obligations

- Make freely available an AI detection tool that allows users to assess whether content was created or modified by the provider's GenAI system.
- Include in AI-generated content (where technically feasible) a clear and conspicuous disclosure that content was generated or substantially altered by AI.
- Embed in AI-generated content latent disclosures (digital watermarks or provenance metadata) consistent with widely adopted standards.
- Maintain provenance information for a reasonable period to support detection and authenticity verification.

Penalties & Enforcement

Civil penalties of \$5,000 per violation per day; AG and CPPA may seek injunctive and equitable relief.

STACK LENS - How this law maps to the AI Governance Stack

Layer 2 - Model Governance	Watermarking and provenance generation are Layer 2 model-output obligations; output filters and post-processing pipelines must implement disclosure infrastructure.
Layer 4 - Control & Monitoring	Detection-tool maintenance and provenance retention are Layer 4 operational responsibilities, not one-time engineering tasks.

California AI in Healthcare Disclosure Act (AB 3030)

Cal. Health & Safety Code § 1339.75

Jurisdiction	California
Effective	January 1, 2025
Regulator	Medical Board of California; Department of Consumer Affairs
Scope	Health facilities, clinics, physicians, and groups using GenAI to generate written or verbal communications to patients

Applicability

AB 3030 requires healthcare providers to disclose to patients when GenAI is used to generate clinical information communicated to the patient (including patient messaging, after-visit summaries, certain instructions). Communications that are reviewed by a human licensed or certified provider before transmission are exempt.

Core Obligations

- Provide a disclosure to the patient that the communication was generated by GenAI.
- Provide instructions for how the patient may contact a human healthcare provider, employee, or other appropriate person.
- Document the use of GenAI in patient-facing communications.

Penalties & Enforcement

Disciplinary action by the Medical Board of California for licensee violations; civil penalties available under unfair business practices framework.

STACK LENS - How this law maps to the AI Governance Stack

Layer 3 - System Integration	Disclosure is a Layer 3 integration boundary requirement: patient-communication systems must surface AI provenance to the recipient.
Layer 4 - Control & Monitoring	Operational workflows must distinguish reviewed-by-human from unreviewed AI outputs; Layer 4 monitoring should track the disclosure-trigger threshold.

New York City Local Law 144 - Automated Employment Decision Tools (AEDT)

N.Y.C. Admin. Code § 20-870 et seq.; 6 RCNY § 5-300 et seq.

Jurisdiction	New York City
Effective	January 1, 2023; enforcement July 5, 2023

Regulator	New York City Department of Consumer and Worker Protection (DCWP)
Scope	Employers and employment agencies using "automated employment decision tools" to make or substantially assist hiring or promotion decisions for positions located in NYC or candidates living in NYC

Applicability

Local Law 144 was the first U.S. AI-specific employment regulation. It requires (1) an annual independent bias audit of the AEDT, (2) public posting of the bias audit summary, and (3) candidate notice describing the tool's use and the qualifications/characteristics it assesses.

Core Obligations

- Conduct an annual independent bias audit of the AEDT calculating selection rate and impact ratio across categories of sex and race/ethnicity (including intersectional categories).
- Publish a summary of the most recent bias audit on the employer's public website.
- Provide candidates with notice at least 10 business days before use of the AEDT, including the tool's use, the job qualifications/characteristics assessed, and (if applicable) the source and types of data collected and the data retention policy.
- Allow candidates to request an alternative selection process or accommodation.

Penalties & Enforcement

Civil penalties of \$500 for first violation, \$500–\$1,500 per subsequent violation. Each day of continued non-compliance and each affected candidate may constitute a separate violation.

Recent Developments (through 2026)

DCWP enforcement has been sporadic but the bias-audit framework has been replicated in many subsequent state and local proposals (Illinois, New Jersey, Washington D.C.). The "substantially assist" test is the key compliance scope question.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Annual independent bias audits are an externalized Layer 2 obligation: third-party verification of model fairness becomes a regulatory deliverable.
Layer 5 - Audit & Evidence	Public bias audit summaries and candidate notices are Layer 5 transparency artifacts subject to ongoing scrutiny.

PRACTITIONER NOTE

The "substantially assist" test is interpreted broadly: AEDT analysis applies wherever the AI score is one of the principal factors weighed in the decision. Build the bias-audit cadence into AI procurement so renewals do not lapse.

Illinois Artificial Intelligence Video Interview Act

820 ILCS 42/

Jurisdiction	Illinois
Effective	2020; amended 2022
Regulator	Illinois Department of Labor; private right of action under similar civil rights frameworks
Scope	Employers using AI to analyze applicant video interviews for positions based in Illinois

Applicability

The Act requires notice to and consent from applicants subject to AI analysis of video interviews, restricts who may view the video, requires destruction within 30 days of request, and (for employers relying solely on AI to determine in-person interview eligibility) requires annual demographic data reporting to the Illinois Department of Commerce and Economic Opportunity.

Core Obligations

- Provide written notice before the interview that AI may be used and the general types of characteristics the AI will use to evaluate the applicant.
- Obtain consent from the applicant before the interview.
- Limit access to the video and AI evaluation to authorized persons.
- Destroy the video and all copies within 30 days of an applicant's request.
- Report demographic data annually if AI is solely relied upon for in-person interview decisions.

Penalties & Enforcement

Enforcement under generally applicable employment law remedies; private civil right of action; potential class action exposure.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Video data lifecycle management - consent collection, access controls, deletion on request - is a Layer 1 / Layer 4 obligation.
Layer 5 - Audit & Evidence	Annual demographic reporting requires Layer 5 statistical recordkeeping by demographic category, with data integrity controls.

Illinois HB 3773 - AI in Employment Decisions

775 ILCS 5/2-101 et seq. (amendments)

Jurisdiction	Illinois
Effective	January 1, 2026
Regulator	Illinois Department of Human Rights; private right of action under Illinois Human Rights Act
Scope	Employers in Illinois using AI to make employment decisions

Applicability

HB 3773 amends the Illinois Human Rights Act to make it a civil rights violation for employers to use AI that has the effect of subjecting employees to discrimination on the basis of protected classes, or to use ZIP codes as a proxy for protected classes, in recruitment, hiring, promotion, renewal of employment, selection for training or apprenticeship, discharge, discipline, tenure, or terms/conditions of employment. Notice to employees of AI use is required.

Core Obligations

- Do not use AI in employment decisions in a manner that has a disparate impact on protected classes.
- Do not use ZIP codes as a proxy for protected classes.
- Provide notice to employees and applicants when AI is used in employment decisions.

Penalties & Enforcement

Available remedies under the Illinois Human Rights Act, including damages, attorney fees, and injunctive relief; private right of action through the Illinois Department of Human Rights or directly.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Disparate impact analysis is a Layer 2 fairness-testing requirement; the explicit ZIP-code prohibition codifies what fair-lending practice already considers a leading proxy variable.
Layer 3 - System Integration	Notice obligations require integration into HR communication workflows.

Tennessee Ensuring Likeness, Voice, and Image Security Act (ELVIS Act)

Tenn. Code Ann. § 47-25-1101 et seq.

Jurisdiction	Tennessee
Effective	July 1, 2024
Regulator	Private right of action; criminal enforcement by district attorneys

Scope	Persons publishing, performing, distributing, transmitting, or otherwise making available the voice or likeness of an individual or knowingly providing tools whose primary purpose is the unauthorized use of voice or likeness
--------------	--

Applicability

The ELVIS Act extends Tennessee's right of publicity to expressly include voice and likeness in the AI context. It is the leading U.S. statute on AI-generated voice cloning and deepfakes, with civil and criminal penalties for unauthorized use and for distributing tools designed for unauthorized voice/likeness generation.

Core Obligations

- Do not publish, perform, distribute, transmit, or otherwise make available the voice or likeness of an individual without authorization.
- Do not knowingly distribute, transmit, or otherwise make available an algorithm, software, tool, or other technology, service, or device the primary purpose or function of which is the production of an individual's voice or likeness without authorization.

Penalties & Enforcement

Civil damages including actual damages, statutory damages of \$25,000 per violation, attorneys' fees, and equitable relief; misdemeanor criminal penalties for willful violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Voice and likeness training data must include rights documentation; biometric/likeness data acquisition pipelines need consent and licensing infrastructure.
Layer 2 - Model Governance	Voice cloning model deployment requires upstream rights verification at the model-input boundary.
Layer 3 - System Integration	Distribution of voice/likeness tools is itself actionable; product distribution channels are Layer 3 risk surfaces.

State Deepfake and AI-Generated Content Laws (Survey)

Jurisdiction	Multiple U.S. states
Effective	2019–2026 (varies by state)
Regulator	State attorneys general; private right of action (varies)
Scope	Varies - generally addresses (1) deepfakes used in elections, (2) non-consensual intimate imagery (NCII), and (3) deepfake fraud

Applicability

A patchwork of state laws addresses deepfake content. Election-deepfake laws (California, Texas, Michigan, Minnesota, Wisconsin, and others) prohibit or require disclosure of AI-generated political content within proximity to elections. NCII laws (45+ states) have been amended in 2024–2025 to address synthetic intimate imagery. Deepfake-fraud laws apply existing criminal statutes to AI-enabled impersonation.

Core Obligations

- For political/election content: disclose AI generation within prescribed periods before elections (specifics vary).
- For NCII: do not create, distribute, or solicit non-consensual synthetic intimate imagery.
- For deepfake fraud: existing wire-fraud, identity-theft, and impersonation statutes apply to AI-enabled offenses.

Penalties & Enforcement

Varies by state; civil and criminal remedies; private rights of action in many jurisdictions.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Generative model output filters that detect and refuse NCII or election-impersonation content are Layer 2 model-deployment controls.
Layer 3 - System Integration	Distribution-platform terms of service and AUPs operationalize state deepfake laws at the integration boundary.

Additional State AI Statutes, Frontier Regulation, and Sectoral AI Laws

A second wave of state AI legislation addresses generative AI consumer protection (Utah), frontier AI model transparency (California SB 53), state agency AI governance (Connecticut PA 23-16), pending consequential-decision frameworks (Massachusetts, New York, Connecticut SB 2), state healthcare AI laws (Oklahoma, Florida, Georgia patterns), and the state election deepfake / NCII patchwork. The aggregate trajectory is a multistate AI regulatory environment more stringent than the federal baseline.

Utah Artificial Intelligence Policy Act (UAIP)

Utah Code §§ 13-72-101 to 13-72-401

Jurisdiction	Utah
Effective	May 1, 2024 (initial provisions); subsequent amendments through 2026
Regulator	Utah Division of Consumer Protection; Utah Office of Artificial Intelligence Policy

Scope	Persons using or providing generative AI in interactions with consumers in Utah; regulated occupations using AI
--------------	---

Applicability

Utah's UAIP was the first U.S. state generative AI consumer protection law. It requires disclosure when consumers interact with generative AI in regulated occupations and certain other contexts, holds entities liable for generative AI consumer protection violations as if the entity made the statements directly, and establishes the Utah Office of Artificial Intelligence Policy with regulatory sandbox authority. The 2025 amendments narrowed the disclosure trigger and clarified scope.

Core Obligations

- For regulated occupations (e.g., healthcare, legal): clearly and conspicuously disclose that the consumer is interacting with generative AI when such interaction occurs.
- For other consumer interactions: disclose AI use upon request by the consumer.
- No defense based on AI involvement: liability for consumer protection violations attaches as if the entity made the statements directly.
- Sandbox: regulated participants may obtain limited regulatory mitigation in exchange for participation conditions.

Penalties & Enforcement

Standard Utah Consumer Sales Practices Act penalties; per-violation civil penalties up to \$2,500.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	Disclosure obligations are Layer 3 integration boundary requirements; consumer-facing AI interfaces must surface AI provenance.
Layer 5 - Audit & Evidence	Sandbox participation requires Layer 5 documentation comparable to traditional regulatory submissions.

California SB 1047 - Safe and Secure Innovation for Frontier AI Models Act (Vetoed)

Cal. Bus. & Prof. Code (proposed Section 22602 et seq., vetoed September 29, 2024)

Jurisdiction	California (proposed)
Effective	Vetoed; relevant as precedent
Regulator	Would have been California Attorney General
Scope	Would have applied to developers of frontier AI models above defined compute and capability thresholds

Applicability

SB 1047 was the most prominent U.S. frontier-AI regulation proposal of 2024. It would have required pre-deployment safety evaluations, kill-switch capability, third-party auditing, and whistleblower protections for developers of models exceeding certain compute thresholds. Vetoed by Governor Newsom on grounds that it focused on model size rather than risk profile. The 2025 California Working Group on AI Frontier Models report and SB 53 (2025) revived elements of the framework. Practitioners should treat SB 1047's structure as a precedent for likely state and federal frontier AI legislation.

Core Obligations

- As enacted in successor legislation: transparency reporting for frontier developers, critical safety incident reporting, and whistleblower protections (SB 53, 2025, narrower than SB 1047).

Penalties & Enforcement

SB 53 penalties scaled to frontier developer revenue.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Frontier model safety evaluations and red-teaming are anticipated Layer 2 obligations across emerging frontier AI regulation.
Layer 4 - Control & Monitoring	Critical incident reporting and rapid-response capability are Layer 4 deliverables that frontier developers should preposition.

California SB 53 - Transparency in Frontier AI Act

Cal. Bus. & Prof. Code §§ 22757.20 et seq.

Jurisdiction	California
Effective	September 2025 (signed); operative provisions phased into 2026
Regulator	California Attorney General; California Operations Agency for AI Reporting
Scope	Frontier AI model developers - defined by training compute thresholds and revenue floor

Applicability

SB 53 is the operative California frontier AI law (the successor to vetoed SB 1047). It requires frontier developers to publish a Frontier AI Framework describing safety practices, file critical safety incident reports with the AG, provide whistleblower protections for AI safety personnel, and (subject to limited transparency) describe training characteristics. SB 53 is among the first operative state frontier AI laws and is closely watched as a U.S. precedent.

Core Obligations

- Publish a Frontier AI Framework describing how the developer evaluates and mitigates catastrophic risks.
- Report critical safety incidents to AG within prescribed periods.
- Provide whistleblower protections for personnel reporting concerns about catastrophic safety risks; restrict NDAs that would suppress such reporting.
- Maintain records of safety evaluations sufficient to support AG inquiry.

Penalties & Enforcement

AG enforcement under unfair competition framework; per-violation civil penalties; injunctive relief.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Frontier AI Framework documentation is the most explicit Layer 2 model-safety disclosure requirement in U.S. state law.
Layer 5 - Audit & Evidence	Critical incident reporting and underlying evaluation records are Layer 5 obligations with regulator inquiry hooks.

California Generative AI Accountability Act (AB 2885) and Related California AI Statutes

Cal. Gov. Code § 11546.45.5 (AB 2885, AI definition); other related California AI bills enacted 2024–2025

Jurisdiction	California
Effective	Various 2025–2026
Regulator	California Office of Information Security; California Department of Technology; sectoral regulators
Scope	State agency AI procurement and use; certain commercial AI uses

Applicability

California enacted a substantial package of AI-specific legislation in 2024–2025 addressing state government AI use, election deepfakes, performer voice and likeness, watermarking, AI in healthcare communications, and AI training data transparency. The legislative cadence positions California as the most prolific U.S. AI legislature.

Core Obligations

- AB 2885: standardize AI definition across California code sections, supporting consistent regulatory application.
- AB 2655 / AB 2839: prohibit certain election-related deepfakes and require platform action.
- AB 1836: extend digital replica rights for deceased personalities.

- AB 2602: require informed consent in employment contracts for digital replica use.
- SB 942: AI provenance and watermarking (separately addressed).
- AB 2013: training data documentation (separately addressed).
- AB 3030: healthcare GenAI disclosure (separately addressed).

Penalties & Enforcement

Vary by statute; AG and private rights of action available in several.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	AB 2602 informed consent for digital replicas reaches Layer 1 talent and likeness data acquisition.
Layer 2 - Model Governance	AB 2655 / 2839 election deepfake prohibitions affect Layer 2 model output policy and content moderation.

Massachusetts Information Privacy and Security Act (Pending) and Massachusetts AI Initiatives

Massachusetts S 25 / H 80 (pending comprehensive privacy); Massachusetts Executive Order 25-1 (Strategic AI Plan)

Jurisdiction	Massachusetts
Effective	Pending; AI Executive Order effective 2024
Regulator	Massachusetts Office of the Attorney General; Executive Office of Technology Services and Security
Scope	Pending comprehensive privacy law would apply to most large processors; AI Strategic Plan applies to state agencies

Applicability

Massachusetts has been working toward comprehensive privacy legislation since 2022. The 2025 versions of the Massachusetts privacy bill are among the most consumer-protective proposals - strict data minimization, sensitive data prohibitions, and direct regulator enforcement. The AI Strategic Plan establishes state agency AI governance. Practitioners should monitor Massachusetts as a likely 2026–2027 enacting state.

Core Obligations

- Pending comprehensive privacy law: strict data minimization; opt-in consent for sensitive data; no sale of sensitive data; AI/profiling impact assessments; consumer rights including opt-out of profiling.
- AI Strategic Plan: state agency AI inventory, impact assessment, and procurement obligations.

Penalties & Enforcement

Pending.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Anticipated strict minimization parallels Maryland; readiness work should mirror MODPA preparations.

Connecticut SB 2 - AI Consequential Decision Bill (Pending) and Other Connecticut AI Initiatives

Connecticut SB 2 (pending); Connecticut Public Act 23-16 (state agency AI)

Jurisdiction	Connecticut
Effective	PA 23-16 effective 2023; SB 2 framework pending
Regulator	Connecticut Attorney General; Connecticut Office of Policy and Management
Scope	SB 2 (if enacted) would apply broadly to AI developers and deployers; PA 23-16 applies to state agencies

Applicability

Connecticut has been a leading state in AI legislative effort. SB 2 (in successive sessions 2023–2026) would establish a Colorado-style consequential-decision framework. PA 23-16 establishes state agency AI inventory, impact assessment, and procurement obligations.

Core Obligations

- PA 23-16: state agency AI inventory and impact assessment.
- SB 2 (if enacted): developer and deployer obligations for high-risk AI similar to Colorado AI Act.

Penalties & Enforcement

Pending SB 2 framework.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	PA 23-16 inventory and assessment requirements provide Connecticut state contractor AI documentation expectations.

New York Local Law 144 - AEDT Detailed Requirements and NY State AI Bills

NY State S 7623 (LOADInG Act); A 7634 (algorithmic discrimination); NYC Local Law 144 (separately addressed)

Jurisdiction	New York State and New York City
Effective	Local Law 144 enforceable July 5, 2023; state bills pending
Regulator	NYC DCWP; NYS AG (pending bills)
Scope	Local Law 144 (NYC employment AEDTs); state proposals broader

Applicability

New York State has multiple pending AI bills addressing automated decision-making in employment, housing, lending, healthcare, and education. The LOADInG Act (Legislative Oversight of Automated Decision-Making in Government) would establish state agency AI governance. Practitioners should track New York as a likely 2026 enacting state for at least one AI consequential-decision framework.

Core Obligations

- Local Law 144 obligations (separately addressed): bias audit, public posting, notice.
- Pending state bills would extend bias audit and impact assessment requirements to additional contexts.

Penalties & Enforcement

Local Law 144 penalties (separately addressed); pending state bills variable.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	NY momentum suggests bias-audit infrastructure built for Local Law 144 will see expanded use under future state requirements.

Oklahoma AI in Healthcare Act and Related State Sectoral AI Laws

Okla. Stat. tit. 63 (pending healthcare AI provisions); analogous bills in Florida (HB 1207), Georgia (HB 887)

Jurisdiction	Multiple states (Oklahoma example)
Effective	Various 2025–2026
Regulator	State health regulators
Scope	AI use in healthcare delivery in covered states

Applicability

A growing wave of state sectoral AI laws addresses healthcare specifically - generally requiring physician oversight of AI clinical decisions, disclosure to patients, and prohibitions on using AI to deny coverage without clinician review. California AB 3030 is the leading example; Oklahoma, Florida, Georgia, Pennsylvania, Texas, and others have considered or enacted similar provisions through 2026.

Core Obligations

- Physician oversight of AI clinical decisions (varying scope).
- Disclosure to patients of AI use in clinical care.
- For coverage decisions: clinician review of AI denials.
- Documentation of AI use in patient records.

Penalties & Enforcement

State medical board enforcement; civil and administrative penalties varying by state.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	Clinician-review requirements shape Layer 3 integration architecture: AI cannot be the terminal decision-maker for adverse coverage decisions.

State Election AI Deepfake Laws (Survey)

Tex. Elec. Code § 255.004; Cal. Elec. Code § 20012; Mich. Comp. Laws § 169.273; Minn. Stat. § 609.771; multiple other states

Jurisdiction	Multiple U.S. states
Effective	2019–2026 (varies)
Regulator	State attorneys general; state election authorities
Scope	Persons creating, distributing, or making available materially deceptive synthetic media of candidates or election officials within prescribed periods of elections

Applicability

Election-deepfake laws prohibit or require disclosure of AI-generated political content within proximity to elections. Most operative laws cover the period within 30, 60, or 90 days of an election; some apply year-round to materially deceptive content. The 2024 election cycle saw the first significant state-level deepfake enforcement; the 2026 cycle is expected to see more.

Core Obligations

- Do not distribute materially deceptive synthetic media of candidates within prescribed periods unless disclosed.
- Specific exemptions for satire, news media (often), and clearly disclosed parody.
- Some states require platforms to remove or label content on notice.

Penalties & Enforcement

Varies by state; civil and criminal remedies; private rights of action in several jurisdictions.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Generative model output filters and policies must address election-period content; Layer 2 model deployment must include geofenced or temporal restrictions.
Layer 4 - Control & Monitoring	Platform monitoring and takedown infrastructure are Layer 4 operational obligations under several state laws.

State Non-Consensual Intimate Imagery (NCII) and Synthetic Sexual Content Laws

47+ states have NCII statutes; 2024–2025 amendments cover synthetic content in most jurisdictions

Jurisdiction	Multiple U.S. states; federal Take It Down Act (2025)
Effective	Varies by state; federal law effective 2025
Regulator	State attorneys general; private rights of action; federal criminal authorities (Take It Down Act)
Scope	Persons creating, distributing, or threatening to distribute non-consensual intimate images including AI-generated synthetic content

Applicability

State NCII statutes generally prohibit creation, distribution, or threatened distribution of non-consensual intimate imagery. The 2024–2025 wave of amendments expanded coverage to AI-generated synthetic content. The federal Take It Down Act (2025) creates parallel federal criminal and civil remedies and platform takedown obligations.

Core Obligations

- Do not create, distribute, or threaten to distribute non-consensual intimate imagery including synthetic.
- For platforms (under federal Take It Down Act): respond to victim takedown requests within 48 hours.

Penalties & Enforcement

State criminal and civil; federal criminal under Take It Down Act; private rights of action in most jurisdictions.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Generative model safety filters preventing NCII generation are Layer 2 deployment controls.
Layer 4 - Control & Monitoring	Take-down infrastructure for hosting platforms is a Layer 4 operational obligation.

P A R T I V

Intellectual Property & AI

Copyright, patent, trade secret, open source licensing, and right-of-publicity frameworks as they apply to AI - the most actively contested doctrinal frontier in AI law.

IP Frameworks for the AI Lifecycle

The intellectual property questions raised by AI development span the entire lifecycle: training data acquisition (copyright, contract, scraping doctrine); model artifact protection (trade secret, patent, copyright); output authorship and ownership (the Thaler line of cases); and downstream user obligations (open source licenses, right of publicity). Pending litigation - most prominently *NYT v. OpenAI* - is reshaping the doctrinal landscape rapidly, but documented diligence remains a consistent partial defense across all theories.

U.S. Copyright Office Guidance on AI-Generated Works

U.S. Copyright Office, Copyright and Artificial Intelligence Reports (Part 1: Digital Replicas, July 2024; Part 2: Copyrightability, January 2025; Part 3: Generative AI Training, anticipated 2026); Compendium of U.S. Copyright Office Practices (Third Edition)

Jurisdiction	United States - Federal
Effective	Copyright Act 1976 (continuously interpreted); Copyright Office AI Reports 2024–2026
Regulator	U.S. Copyright Office (registration); federal courts (infringement and validity)
Scope	Copyright registration of AI-involved works; copyright infringement claims involving AI training and outputs

Applicability

The U.S. Copyright Office has taken the position that copyright protection requires human authorship - purely AI-generated outputs are not copyrightable, but human-authored elements within AI-assisted works may be. The 2024–2026 AI Reports establish the operating framework for registration and enforcement. The Office's pending Part 3 (generative AI training) is the most-anticipated guidance, addressing fair use analysis for AI training and the licensing market's development. Practitioners should distinguish between (a) copyrightability of AI outputs, (b) copyright in AI-assisted human works, (c) infringement of training inputs, (d) infringement of model outputs, and (e) DMCA Section 1201 / 1202 obligations for content provenance.

Core Obligations

- Disclose AI involvement in copyright registration applications when AI generated more than de minimis content; specify human-authored portions claimed.
- For digital replicas of identifiable persons: comply with emerging federal and state right-of-publicity frameworks (NO FAKES Act, ELVIS Act, others).
- For DMCA § 1202 (Copyright Management Information): do not knowingly remove or alter CMI; AI training that strips copyright notices may create § 1202 exposure.

- For training data acquisition: maintain records of source authorization, licensing, and (where applicable) permitted-use justification.

Penalties & Enforcement

Copyright Act statutory damages: \$750–\$30,000 per work for non-willful infringement; up to \$150,000 per work for willful infringement; plus actual damages, profits, and attorneys' fees. Class-action exposure in pending AI training litigation is in the multi-billion-dollar range.

Recent Developments (through 2026)

Major pending litigation: *New York Times v. OpenAI/Microsoft* (2023, ongoing); *Authors Guild v. OpenAI* (2023, ongoing); *Andersen v. Stability AI* (2023, ongoing); *UMG v. Anthropic* (2023, partial settlement 2024); *Concord Music Group v. Anthropic*; multiple visual artist cases. Several outcomes have favored AI providers on transformative-use grounds (e.g., partial *Bartz v. Anthropic* ruling 2025), others have allowed claims to proceed. The Open Source Initiative's 2024 Open Source AI Definition and the 2025 industry licensing initiatives are reshaping the practical landscape.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training data acquisition is the central IP risk surface; Layer 1 controls including source authorization, opt-out respect, and licensing infrastructure are foundational.
Layer 2 - Model Governance	Output filtering for verbatim reproduction of training content is a Layer 2 obligation in litigation defense; "memorization" mitigations are increasingly expected.
Layer 5 - Audit & Evidence	Documentation of training data sources, licensing, and permission infrastructure is Layer 5 evidence in pending and future copyright litigation.

PRACTITIONER NOTE
 For commercial generative AI products, build a training-data provenance ledger that traces every source corpus to a documented authorization basis. Pending litigation is reshaping the doctrinal landscape, but documented diligence is consistently a partial defense.

NYT v. OpenAI and the Generative AI Training Litigation Landscape

The New York Times Company v. Microsoft Corporation, OpenAI, Inc., et al., No. 1:23-cv-11195 (S.D.N.Y. filed Dec. 27, 2023); In re OpenAI Inc. Copyright Infringement Litig., MDL pending

Jurisdiction	United States - Federal courts
Effective	Filed December 2023; ongoing through 2026
Regulator	Federal courts

Scope	Copyright infringement claims against generative AI developers and deployers
--------------	--

Applicability

NYT v. OpenAI is the most consequential pending generative AI training litigation. The plaintiffs allege both training-stage infringement (use of copyrighted articles in training datasets) and output-stage infringement (verbatim or near-verbatim reproduction of articles in ChatGPT outputs). The case will likely shape the doctrinal framework for fair use analysis of AI training and output liability for verbatim reproduction. As of 2026, the litigation has survived motions to dismiss, with discovery underway. Class-action consolidation in MDL proceedings affects related visual-art and music cases.

Core Obligations

- Maintain training data provenance records sufficient to support source attribution and authorization defenses.
- Implement output filters preventing or detecting verbatim reproduction of training content.
- For commercial generative AI deployments: consider indemnification structures and customer-facing IP protection commitments (now standard in major commercial generative AI offerings).
- For publishers and content owners: establish opt-out infrastructure (robots.txt, ai.txt, do-not-train signals) and licensing programs.

Penalties & Enforcement

Copyright Act remedies including statutory damages and attorney's fees; injunctive relief that could materially affect deployed products.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Opt-out signal respect (robots.txt, do-not-train metadata) is becoming a Layer 1 acquisition gate.
Layer 2 - Model Governance	Output filtering for verbatim reproduction is a Layer 2 model-deployment control; near-duplicate detection is the operational implementation.

Thaler v. Perlmutter and AI Inventorship/Authorship Doctrine

Thaler v. Perlmutter, No. 23-5233 (D.C. Cir. 2025); Thaler v. Vidal, 43 F.4th 1207 (Fed. Cir. 2022)

Jurisdiction	United States - Federal
Effective	Federal Circuit 2022 (patent inventorship); D.C. Circuit 2025 (copyright authorship)
Regulator	USPTO; U.S. Copyright Office; federal courts
Scope	Patent inventorship and copyright authorship of AI-generated works

Applicability

Stephen Thaler's parallel patent and copyright cases established that under U.S. patent law, an "inventor" must be a natural person (*Thaler v. Vidal*, Federal Circuit 2022), and under U.S. copyright law, an "author" must be human (*Thaler v. Perlmutter*, D.C. Circuit 2025). These holdings establish the U.S. doctrinal foundation: AI cannot be an inventor or author; human contribution is required for both patent and copyright protection. The threshold of human contribution sufficient to support protection is the operative practical question.

Core Obligations

- For patent applications: name human inventors who made significant contributions to the conception of the claimed invention; AI-assisted inventions can be patentable where humans materially contributed.
- For copyright registration: claim only human-authored portions; describe AI involvement; do not claim authorship of purely AI-generated material.
- Maintain documentation of human contribution sufficient to support inventorship/authorship determinations.

Penalties & Enforcement

Loss of patent or copyright protection; allegations of inequitable conduct in patent prosecution; copyright cancellation.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Documentation of human contribution to AI-assisted creative or inventive work is Layer 5 evidence supporting IP protection.

USPTO Inventorship Guidance for AI-Assisted Inventions

USPTO Inventorship Guidance for AI-Assisted Inventions, 89 Fed. Reg. 10,043 (February 13, 2024); USPTO Subject Matter Eligibility Examples 47–49 (AI-related)

Jurisdiction	United States - Federal
Effective	February 13, 2024
Regulator	United States Patent and Trademark Office
Scope	Patent applications involving AI-assisted invention

Applicability

The USPTO's 2024 inventorship guidance establishes the framework for evaluating inventorship in AI-assisted inventions. Each named inventor must have made a significant contribution to the conception of the claimed invention. The guidance enumerates factors and provides examples. Patent applicants must analyze human contribution at the claim level - different claims may have different inventorship analyses. The guidance is the most operational federal AI patent doctrine in 2026.

Core Obligations

- Analyze human contribution at the claim level; name inventors whose contribution is significant.
- Document the human-contribution analysis in inventorship records.
- Comply with duty of disclosure including AI involvement where material to patentability.
- For continuation/divisional applications: re-evaluate inventorship as claims change.

Penalties & Enforcement

Patent invalidity or unenforceability; inequitable conduct allegations; loss of patent term where errors require correction.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Inventorship contribution records are Layer 5 evidence in patent prosecution and litigation; documentation should be contemporaneous with invention.

EU Copyright Directive Article 4 (Text and Data Mining Exception)

Directive (EU) 2019/790 (CDSM Directive), Articles 3 and 4

Jurisdiction	European Union
Effective	June 7, 2021 (Member State implementation deadline)
Regulator	National copyright authorities; EU courts
Scope	Reproductions and extractions of lawfully accessible works for text and data mining

Applicability

CDSM Article 3 provides a mandatory text-and-data-mining (TDM) exception for research organizations and cultural heritage institutions. Article 4 provides a broader TDM exception for any beneficiary, but rights holders may opt out by an "express reservation in an appropriate manner, such as machine-readable means in the case of content made publicly available online." The Article 4 opt-out has emerged as the central European mechanism for AI training data control. The EU AI Act (Article 53(1)(c)) requires GPAI providers to comply with Article 4 opt-outs in training. Practitioners must implement opt-out signal recognition and respect.

Core Obligations

- For AI training of works available in the EU: respect machine-readable opt-out signals (e.g., TDM reservation in robots.txt, ai.txt, HTTP headers, content-level metadata).
- For AI providers: maintain documentation of opt-out signal recognition and respect.
- For research-purpose AI training: Article 3 exception may apply with narrower beneficiary scope.

Penalties & Enforcement

Copyright remedies under Member State implementations; EU AI Act penalties for GPAI non-compliance with copyright policy obligations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	TDM opt-out signal recognition is a Layer 1 acquisition gate for EU operations; the technical standard for opt-out signals is evolving and practitioners should monitor.
Layer 5 - Audit & Evidence	EU AI Act Article 53(1)(c) requires written copyright policy and training data summary - Layer 5 deliverables enforceable by the AI Office.

UK Text and Data Mining and AI Copyright Reform

Copyright, Designs and Patents Act 1988, s.29A (TDM exception, currently research-only); UK government Code of Practice on Copyright and AI (development withdrawn 2023; reform consultations 2023–2026)

Jurisdiction	United Kingdom
Effective	CDPA TDM exception 2014; reform pending
Regulator	UK Intellectual Property Office (IPO)
Scope	Copyright in works lawfully accessed for TDM

Applicability

UK TDM exception currently applies only to non-commercial research. The 2022 government proposal for an unrestricted TDM exception was abandoned in 2023 after creative-industry opposition; the 2023–2026 consultations have considered an opt-out model similar to EU Article 4, but no statutory reform has been enacted as of 2026. UK AI training practice in the absence of clear statutory authority creates significant litigation risk; the Getty Images v. Stability AI litigation (UK High Court) is the leading test case.

Core Obligations

- For commercial AI training using UK-accessible content: obtain license or rely on judicial fair-dealing (limited and uncertain).

- Monitor pending reform; the most-anticipated UK creative-industries / AI compromise remains unresolved.

Penalties & Enforcement

Copyright remedies including damages, account of profits, and injunctive relief.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	UK acquisition without clear authority is high-risk; Layer 1 acquisition policy should default to opt-out respect even where statutory backstop is uncertain.

Japan Copyright Act Article 30-4 - TDM Exception

Japan Copyright Act Article 30-4 (2018 revision)

Jurisdiction	Japan
Effective	January 1, 2019
Regulator	Agency for Cultural Affairs
Scope	Reproductions of works for purposes other than enjoying their thoughts or feelings

Applicability

Japan's Article 30-4 provides one of the world's broadest text-and-data-mining exceptions, permitting reproduction of copyrighted works for purposes other than enjoying the thoughts or feelings expressed (including AI training) without rightsholder authorization, subject to limitations including not unreasonably prejudicing the interests of the copyright holder. The exception has positioned Japan as a favorable jurisdiction for AI training. Japanese courts have not yet substantially tested the boundaries; the Article 30-4 regime is widely cited but practitioners should monitor judicial development.

Core Obligations

- Reproduction for AI training purposes is generally permitted within Article 30-4 scope.
- Avoid uses that would unreasonably prejudice rights holders' interests; do not commercially distribute reproductions in ways that compete with authorized uses.

Penalties & Enforcement

Copyright remedies for uses outside the exception.

STACK LENS - How this law maps to the AI Governance Stack

Layer 1 - Data Governance	Japan is currently the most-permissive major-economy jurisdiction for AI training data acquisition; many AI providers structure training operations to leverage Japan-based facilities.
----------------------------------	---

Trade Secret Protection of AI Models and Training Data

Defend Trade Secrets Act, 18 U.S.C. § 1836; state Uniform Trade Secrets Act adoptions; EU Trade Secrets Directive (2016/943)

Jurisdiction	United States and European Union
Effective	DTSA 2016; EU Directive 2018
Regulator	Federal courts; state courts; EU national courts
Scope	Owners of information that derives independent economic value from not being generally known and is subject to reasonable secrecy efforts

Applicability

Trade secret protection is the principal IP regime protecting AI models, training data, weights, prompts, and evaluation methodologies. Unlike patent, trade secret protection requires no disclosure but is lost upon disclosure or failure to maintain secrecy. AI model weights, fine-tuning datasets, prompt-engineering know-how, and evaluation infrastructure typically qualify. The Defend Trade Secrets Act provides a federal civil cause of action with potential whistleblower immunity provisions.

Misappropriation through reverse engineering of deployed models, model extraction attacks, employee mobility, and supply chain leakage are the principal risk vectors.

Core Obligations

- Implement reasonable secrecy measures: access controls, NDAs (employee, contractor, vendor), classification, watermarking where appropriate, monitoring.
- For AI deployment: consider model extraction defenses (rate limiting, output noise, watermarking) where weights themselves are claimed as trade secrets.
- Comply with the Defend Trade Secrets Act's whistleblower immunity notice requirements in employee/contractor agreements.
- Maintain documentation of trade secret status, value, and protection measures.

Penalties & Enforcement

DTSA: damages including unjust enrichment, exemplary damages up to 2x compensatory for willful misappropriation, attorneys' fees, and injunctive relief; criminal penalties under Economic Espionage Act for willful misappropriation benefiting foreign government.

Recent Developments (through 2026)

AI-related trade secret litigation has accelerated through 2024–2026 including disputes between AI competitors over model weights, alleged misappropriation by departing engineers, and supply chain leakage. The 2024 SDNY decision in *Doe v. GitHub* (Copilot litigation) addressed but did not resolve trade secret claims related to training data.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training data classification and access controls are foundational Layer 1 trade secret protections.
Layer 2 - Model Governance	Model weights and fine-tuning artifacts as trade secrets require Layer 2 protection including model extraction defenses.
Layer 4 - Control & Monitoring	Monitoring for unauthorized model access, extraction attempts, and exfiltration is a Layer 4 detection obligation.
Layer 5 - Audit & Evidence	Trade secret protection programs must be documented; misappropriation litigation routinely tests the "reasonable measures" element.

PRACTITIONER NOTE
 For commercial AI products, the most under-protected trade secrets are typically prompt-engineering libraries and evaluation infrastructure - these have substantial value but are often less rigorously protected than model weights. Build trade secret protection consistently across all AI artifacts.

Open Source AI Licensing - OSI Definition, OpenRAIL, MIT/Apache, and Custom Licenses

Open Source Initiative, Open Source AI Definition v1.0 (October 28, 2024); OpenRAIL family of licenses; Llama Community License; OpenMDW License

Jurisdiction	International (license-based)
Effective	OSI OSAID v1.0 October 2024; license-specific dates
Regulator	License terms enforced by license owners; OSI is non-enforcing standards body
Scope	AI models, weights, code, data, and documentation distributed under open source or open-weights terms

Applicability

The 2024 OSI Open Source AI Definition is the first formal definition of open source AI, requiring open licensing of model parameters, source code, and "data information" sufficient to recreate the model. Most "open weights" model releases (including Llama, Mistral, and similar) do not meet the OSI definition due to use restrictions or training-data limitations. Practitioners should distinguish between (a) OSI-compliant open source AI, (b) Responsible AI Licenses (OpenRAIL family) imposing use

restrictions, (c) source-available licenses with commercial restrictions, and (d) custom community licenses. Each carries distinct compliance obligations and IP-strategy implications.

Core Obligations

- For consumers of open-source AI: review license terms before commercial use; OpenRAIL-family licenses impose substantive use restrictions including prohibitions on certain uses (military, surveillance, manipulation) that bind downstream users.
- For Llama Community License: comply with monthly active user thresholds, attribution requirements, and use restrictions including prohibition on using outputs to improve other LLMs.
- For modifications and derivative works: comply with copyleft or attribution obligations; document compliance.
- For distribution: include required licensing notices and (for OpenRAIL) compliant use restrictions.

Penalties & Enforcement

Contractual breach of license terms; potential copyright infringement where license is breached; reputational and ecosystem consequences.

Recent Developments (through 2026)

OSI OSAID v1.0 has prompted active community debate about which existing model releases meet the definition. The 2025 OpenMDW (Open Model Distribution and Weights) License attempts to provide an OSI-compatible model license. Industry licensing practice continues to evolve rapidly.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Open source AI licensing creates downstream Layer 1 obligations for both inputs (training data licensing) and outputs (use restrictions on derived models).
Layer 3 - System Integration	License compliance for open AI components in product architectures is a Layer 3 governance task; incompatible license combinations create exposure.
Layer 5 - Audit & Evidence	License inventory and compliance documentation are Layer 5 deliverables; SBOM-equivalent AI bill of materials (AIBOM) practice is emerging.

PRACTITIONER NOTE

Build an AI Bill of Materials (AIBOM) tracking each AI component, its license, its training data origin (where disclosed), and its compliance obligations. The AIBOM is the AI-equivalent of the SBOM and is rapidly becoming a procurement expectation.

Right of Publicity, NO FAKES Act, and AI Voice/Likeness Protection

NO FAKES Act (S.2691, pending); state right-of-publicity statutes (CA Civ. Code § 3344, NY Civ. Rights Law § 50, etc.); ELVIS Act (Tennessee, separately addressed)

Jurisdiction	United States - Federal pending; state law operative
Effective	State laws operative; NO FAKES Act pending
Regulator	State courts; private rights of action
Scope	Use of name, voice, likeness, signature, or other indicia of identity for commercial purposes

Applicability

AI voice cloning, deepfake generation, and synthetic likeness creation implicate the patchwork of state right-of-publicity laws and the pending federal NO FAKES Act. The federal proposal would create a federal property right in voice and likeness with civil and criminal penalties for unauthorized AI replicas. AI providers offering generative voice or image services must implement consent verification and provide notice-and-takedown infrastructure.

Core Obligations

- Obtain rights authorization before training or generating identifiable voice or likeness.
- Maintain consent records for licensed digital replicas.
- For platforms: implement notice-and-takedown processes for unauthorized replicas.
- For creative industries: comply with collective bargaining provisions (e.g., 2023 SAG-AFTRA agreement digital replica provisions).

Penalties & Enforcement

State right-of-publicity damages including disgorgement of profits; pending federal NO FAKES Act would add federal civil and criminal remedies including statutory damages.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Voice and likeness training data acquisition requires Layer 1 rights documentation; biometric/likeness acquisition pipelines need consent and licensing infrastructure.
Layer 2 - Model Governance	Generative voice and image models require Layer 2 prompt-time identity verification or restriction.

P A R T V

European Union

The most comprehensive horizontal AI regulatory regime in the world, layered on the foundational GDPR, the Digital Services Act, NIS2, DORA, the Data Act, the Cyber Resilience Act, and a deepening corpus of Union digital regulation.

Core EU AI, Privacy, and Digital Regulation

The European Union has produced both the world's leading data protection regime (GDPR) and the world's first comprehensive horizontal AI law (the AI Act), supported by a dense ecosystem of digital regulation addressing platforms, cybersecurity, financial sector resilience, data sharing, and intermediaries. EU obligations apply extraterritorially in many respects and shape AI compliance globally - a phenomenon often described as the "Brussels Effect."

EU AI Act (Regulation (EU) 2024/1689)

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024

Jurisdiction	European Union (extraterritorial - applies to providers placing AI systems on the EU market and to use of outputs in the EU)
Effective	Entered into force August 1, 2024; phased application - prohibited practices February 2, 2025; GPAI obligations August 2, 2025; high-risk Annex III obligations August 2, 2026; Annex I high-risk and product-safety obligations August 2, 2027
Regulator	European AI Office (within DG CNECT); national market surveillance authorities; European Data Protection Supervisor (Union institutions); European Artificial Intelligence Board
Scope	Providers, deployers, importers, distributors, and product manufacturers of AI systems where the system is placed on the EU market, put into service in the EU, or where the output is used in the EU

Applicability

The EU AI Act is the world's first comprehensive horizontal AI regulation, structured around a risk pyramid: (1) prohibited practices (Article 5), (2) high-risk AI systems (Article 6 and Annexes I and III), (3) limited-risk transparency obligations (Article 50), and (4) minimal-risk systems. General-Purpose AI Models (GPAI) face a separate regulatory regime under Articles 51–55 with additional obligations for models with "systemic risk" (training compute > 10²⁵ FLOPs threshold). Penalties scale with violation severity, reaching €35 million or 7% of worldwide annual turnover for prohibited practices.

Core Obligations

- Prohibited practices (Article 5): no deployment of subliminal manipulation causing significant harm; exploiting vulnerabilities of specific groups; social scoring by public authorities; real-time remote biometric identification in publicly accessible spaces (with narrow law-enforcement exceptions); biometric categorization by sensitive attributes; emotion recognition in workplace and education; certain predictive policing; untargeted scraping for facial recognition databases.
- High-risk AI systems (Annexes I and III): risk management system; data governance for training/validation/testing; technical documentation; record-keeping; transparency to deployers; human oversight; accuracy, robustness, and cybersecurity; quality management system;

conformity assessment (self-assessment for most Annex III; notified body for biometric systems and Annex I product-safety systems); CE marking; EU declaration of conformity; registration in EU database.

- Deployers of high-risk AI: use according to instructions; assign human oversight; ensure input data is relevant and representative; monitor operation and report serious incidents; comply with information obligations to natural persons; conduct fundamental rights impact assessment for certain deployers (public authorities and Article 6 high-risk systems used for credit/insurance assessment).
- GPAI providers: technical documentation; information for downstream providers; copyright compliance policy and training data summary; (for systemic-risk models) model evaluation; systemic risk assessment and mitigation; serious incident reporting; cybersecurity protections.
- Limited-risk transparency: AI systems interacting with humans, generating synthetic content, performing emotion recognition or biometric categorization, or generating deep fakes must inform users (with exceptions for content used for criminal investigation, satire, or where authorized by law).

Penalties & Enforcement

Article 99 penalties: up to €35M or 7% of worldwide annual turnover (whichever higher) for Article 5 prohibited practices; up to €15M or 3% for most other violations; up to €7.5M or 1% for supplying incorrect, incomplete, or misleading information. Lower caps for SMEs and start-ups. National enforcement supplemented by EU AI Office for GPAI matters.

Recent Developments (through 2026)

European Commission published the GPAI Code of Practice in 2025 to facilitate compliance with Articles 51–55. The European AI Office has begun building cases against frontier model providers. Standardization mandates have been issued to CEN-CENELEC for harmonized standards on Annex III high-risk AI systems; published standards will trigger presumption of conformity. The 2026 high-risk effective date is the central operational milestone for most enterprise practitioners.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Article 10 data governance imposes the most prescriptive Layer 1 requirements in any global regulation: relevant, representative, free of errors, complete training/validation/testing datasets; bias examination; appropriate statistical properties; data lifecycle documentation.
Layer 2 - Model Governance	Article 9 risk management, Article 13 transparency, Article 14 human oversight, Article 15 accuracy/robustness/cybersecurity - the entire Layer 2 model governance specification has been transposed into EU regulation.
Layer 3 - System Integration	Quality management system (Article 17) and the developer-deployer-distributor responsibility chain shape Layer 3 integration governance; deployer obligations specifically address operational integration.

STACK LENS - How this law maps to the AI Governance Stack

Layer 4 - Control & Monitoring	Post-market monitoring (Article 72), serious incident reporting (Article 73), and human oversight requirements are Layer 4 obligations with regulatory teeth.
Layer 5 - Audit & Evidence	Technical documentation (Annex IV), record-keeping (Article 12), conformity assessment (Article 43), CE marking, and registration in the EU database constitute the most extensive Layer 5 audit-trail requirements in any jurisdiction.

PRACTITIONER NOTE

Build compliance to the Annex IV technical documentation template even for systems not (yet) classified as high-risk. The template captures the full lifecycle of governance evidence and serves as a cross-jurisdictional defense file. For GPAI providers, the Code of Practice is the operational playbook: adopt it as your baseline.

COMMON FAILURE PATTERN

Risk classification is treated as a one-time triage at procurement, then never revisited as the system's deployment context evolves. A general-purpose model placed downstream of a hiring shortlist tool, a credit-scoring augmentation, or a public-service eligibility filter becomes a high-risk Annex III deployment without the provider's knowledge. Article 25 places the high-risk obligations on whoever puts the high-risk product on the market, regardless of whether the original developer designated the system as high-risk. Build the deployment-context review as a continuous process tied to use-case onboarding, not a single intake form.

General Data Protection Regulation (GDPR, Regulation (EU) 2016/679)

Regulation (EU) 2016/679

Jurisdiction	European Union (extraterritorial under Article 3)
Effective	May 25, 2018
Regulator	National Data Protection Authorities (DPAs); European Data Protection Board (EDPB); European Data Protection Supervisor (Union institutions)
Scope	Controllers and processors of personal data established in the EU; controllers and processors outside the EU offering goods/services to data subjects in the EU or monitoring their behavior in the EU

Applicability

The GDPR is the foundational EU privacy regulation and the model adopted (with variations) by jurisdictions worldwide. AI systems processing personal data - for training, inference, or downstream use - must comply with the GDPR's lawful-basis, transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity/confidentiality, and accountability principles. Article 22 governs

solely automated decision-making with legal or similarly significant effects. The interplay between the GDPR and the EU AI Act is the central compliance complexity for EU operations.

Core Obligations

- Identify and document a lawful basis for each processing activity (consent, contract, legal obligation, vital interests, public task, legitimate interests).
- Provide transparency information under Articles 13–14, including for profiling and automated decision-making - the existence of the logic involved, significance, and envisaged consequences.
- Implement data minimization, purpose limitation, and storage limitation throughout the AI development lifecycle.
- Honor data subject rights: access, rectification, erasure, restriction, data portability, objection, and rights related to automated decisions (Article 22) - including the right not to be subject to a solely automated decision producing legal or similarly significant effects (with enumerated exceptions).
- Conduct a Data Protection Impact Assessment (DPIA, Article 35) for processing likely to result in a high risk, including most AI systems engaged in profiling or large-scale processing of special-category data.
- For special-category data (Article 9) including biometric, health, racial/ethnic, religious, political, sexual orientation, genetic - obtain explicit consent or rely on a narrow Article 9(2) condition; many AI use cases require both an Article 6 lawful basis and an Article 9 condition.
- For cross-border transfers (Chapter V): use adequacy decisions, Standard Contractual Clauses with supplementary measures (post-Schrems II), Binding Corporate Rules, or derogations.
- Designate a Data Protection Officer where required (Article 37).
- Maintain Records of Processing Activities (Article 30).
- Report personal data breaches to the supervisory authority within 72 hours (Article 33) and to data subjects without undue delay where high risk to rights and freedoms (Article 34).

Penalties & Enforcement

Tier 1: up to €10M or 2% of worldwide annual turnover (whichever higher) for violations of certain articles. Tier 2: up to €20M or 4% for violations of basic processing principles, data subject rights, transfer rules, and certain DPA orders. Major AI-relevant fines include Meta (€1.2B, 2023, transfers); Clearview AI (multiple DPA fines totaling >€90M cumulative); Replika (€5M, 2025, Italian Garante); ChatGPT (Italian Garante temporary block 2023, settled with €15M penalty 2024).

Recent Developments (through 2026)

EDPB Opinion 28/2024 on AI models and personal data established the European regulatory position on three key questions: (1) when AI models can be considered to process personal data in inference; (2) the legitimate interest analysis for training on publicly available data; and (3) the consequences for downstream deployers of upstream training using unlawfully processed data. Italian Garante and other DPAs have actively enforced GDPR against AI providers. The interaction with the EU AI Act creates a dual-regime compliance architecture.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Lawful basis, purpose limitation, and data minimization are foundational Layer 1 controls; AI training data acquisition is the most legally exposed phase of the AI lifecycle in the EU.
Layer 2 - Model Governance	Article 22 obligations on automated decisions, accuracy principle, and DPIA requirements drive Layer 2 model documentation, fairness analysis, and explainability infrastructure.
Layer 3 - System Integration	Controller-processor structure and Standard Contractual Clauses define Layer 3 integration boundaries; AI subprocessor flows and joint-controller relationships require explicit allocation.
Layer 4 - Control & Monitoring	Right-to-object and right-to-erasure mechanics demand Layer 4 operational infrastructure capable of acting on inference data and (where feasible) training data.
Layer 5 - Audit & Evidence	Records of Processing Activities (Article 30) and DPIA documentation are Layer 5 audit artifacts; DPAs increasingly request these in early-stage inquiries.

PRACTITIONER NOTE
 For AI systems trained on personal data, the EDPB Opinion 28/2024 analysis (does the model itself contain personal data; what is the lawful basis for training; what are the consequences for downstream use) must be a documented decision artifact maintained alongside the DPIA. DPAs increasingly request the EDPB Opinion analysis explicitly in early-stage inquiries; absence of the analysis is itself a compliance signal.

COMMON FAILURE PATTERN
 Legitimate interest is asserted as the lawful basis for AI training without performing the three-part test (purpose, necessity, balancing) in writing. The Italian Garante, French CNIL, and Dutch AP have all opened inquiries on AI training questions where the controller could not produce a contemporaneous balancing test. Reconstructed analyses do not survive regulator scrutiny and convert what should have been a procedural finding into a substantive lawfulness violation with Tier-2 penalty exposure.

EU Digital Services Act (DSA)

Regulation (EU) 2022/2065

Jurisdiction	European Union (extraterritorial)
Effective	February 17, 2024 (general); August 25, 2023 (very large online platforms and search engines)
Regulator	European Commission (VLOPs/VLOSEs); national Digital Services Coordinators

Scope	Providers of intermediary services (mere conduit, caching, hosting), with tiered obligations for online platforms, very large online platforms (VLOPs, 45M+ monthly EU users), and very large online search engines (VLOSEs)
--------------	--

Applicability

The DSA imposes content-governance, transparency, risk-assessment, and due-process obligations on online intermediaries, with the most demanding obligations falling on VLOPs and VLOSEs. AI-implicated obligations include recommender system transparency, restrictions on dark patterns, advertising transparency, restrictions on profiling minors, and (for VLOPs/VLOSEs) annual systemic risk assessments and audits - including risks arising from generative AI and recommender systems.

Core Obligations

- Notice-and-action mechanisms; statement of reasons; internal complaint-handling.
- For platforms: trader traceability for marketplaces; suspension of users frequently providing illegal content; transparency reporting.
- For VLOPs/VLOSEs: annual systemic risk assessments (covering dissemination of illegal content, fundamental rights, electoral processes, gender-based violence, public health, minors' protection); risk mitigation; independent audits; data access for vetted researchers; crisis response mechanism.
- Recommender systems transparency including non-profiling option for VLOPs/VLOSEs.
- Restrictions on advertising based on profiling using sensitive data; prohibition on profiling-based advertising directed at minors.
- Code of Conduct on Disinformation incorporated by reference for VLOPs.

Penalties & Enforcement

Up to 6% of worldwide annual turnover; periodic penalty payments up to 5% of average daily worldwide turnover for non-compliance with orders. Commission has opened formal proceedings against multiple VLOPs/VLOSEs since 2024.

Recent Developments (through 2026)

Commission proceedings against X, Meta, AliExpress, Temu, and others have established enforcement patterns. The 2025 codes of conduct on AI-generated content and the integration of generative AI features into search and recommender systems have prompted formal information requests.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Recommender system models (Layer 2 design) are now externally auditable and subject to systemic risk assessment.
Layer 4 - Control & Monitoring	Real-time monitoring for illegal content and operational risk-mitigation measures are Layer 4 obligations.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	Annual systemic risk assessments and independent audits are the most demanding Layer 5 transparency exercise outside the financial sector.
---------------------------------------	--

EU NIS2 Directive

Directive (EU) 2022/2555

Jurisdiction	European Union (transposed by Member States)
Effective	October 17, 2024 (transposition deadline; many Member State delays)
Regulator	National competent authorities and CSIRTs; ENISA; Cooperation Group
Scope	Essential and important entities across 18 sectors including energy, transport, banking, financial market infrastructure, healthcare, drinking water, wastewater, digital infrastructure, ICT service management, public administration, space, postal/courier, waste management, manufacturing, food, manufacturing of certain critical products, digital providers, and research

Applicability

NIS2 substantially expands the scope and stringency of EU cybersecurity obligations beyond NIS1, with explicit risk management measures, incident reporting, supply chain security, and management body accountability. AI systems and AI providers are not specifically named but are functionally in scope where they support essential or important entities or where the entity itself qualifies (digital infrastructure, ICT service management, manufacturing of critical products).

Core Obligations

- Implement risk management measures including: incident handling; business continuity; supply chain security including direct supplier and service provider risk; security in network and information systems acquisition, development, and maintenance; policies and procedures for assessing effectiveness; basic cyber hygiene practices and cybersecurity training; cryptography and encryption policies; HR security, access control, asset management; multi-factor authentication; secure communications.
- Report significant incidents within 24 hours (early warning), 72 hours (incident notification), and one month (final report).
- Management body approval and oversight of cybersecurity risk management measures; management body training obligations.
- Supply chain security assessment for critical service providers.

Penalties & Enforcement

Essential entities: up to €10M or 2% of worldwide turnover. Important entities: up to €7M or 1.4% of worldwide turnover. Personal liability for management bodies for non-compliance with risk-management measure requirements.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	Supply chain security obligations cover AI providers serving essential/important entities; vendor inventories must include AI components and inference services.
Layer 4 - Control & Monitoring	Incident reporting timelines (24h/72h/1mo) are aggressive Layer 4 obligations; AI-specific incident detection and classification logic must be pre-built.
Layer 5 - Audit & Evidence	Management body accountability creates board-level Layer 5 reporting requirements; AI-related risk assessments feed into the management approval cycle.

EU Digital Operational Resilience Act (DORA)

Regulation (EU) 2022/2554

Jurisdiction	European Union (financial sector)
Effective	January 17, 2025
Regulator	European Supervisory Authorities (EBA, ESMA, EIOPA); national competent authorities; oversight framework for critical ICT third-party providers
Scope	Financial entities (credit institutions, payment institutions, electronic money institutions, investment firms, crypto-asset service providers, central securities depositories, central counterparties, trading venues, trade repositories, managers of alternative investment funds, management companies, insurance and reinsurance undertakings, intermediaries, ancillary service providers, credit rating agencies, administrators of critical benchmarks, crowdfunding service providers, securitization repositories) and ICT third-party service providers

Applicability

DORA establishes uniform requirements for financial sector ICT risk management, incident reporting, operational resilience testing, and oversight of ICT third-party providers - including AI providers serving the financial sector. Critical ICT third-party providers (designated by the ESAs) face direct EU-level oversight including on-site inspections.

Core Obligations

- ICT risk management framework approved by management body, with periodic review.
- Major ICT-related incident classification and reporting (initial, intermediate, final reports).
- Operational resilience testing program including threat-led penetration testing for significant entities every three years.

- Contractual provisions with ICT third-party providers including service descriptions, locations, security and SLA, exit strategies, audit rights, and notification of material changes.
- Concentration risk monitoring; pre-contractual assessment of critical ICT services.
- For ICT third-party providers designated as critical: oversight framework, recommendations, and penalties up to 1% of average daily worldwide turnover.

Penalties & Enforcement

National competent authority sanctions (varies by Member State); for critical ICT third-party providers: ESA penalties up to 1% of average daily worldwide turnover, applied daily until compliance.

Recent Developments (through 2026)

The first wave of critical ICT third-party provider designations is anticipated in 2026; major cloud, SaaS, and AI providers are expected. Financial entities have spent 2024–2026 retrofitting third-party contracts and exit strategies.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	DORA contractual provisions are the most prescriptive AI-vendor contract regime in the world for financial services; legacy AI contracts must be substantially renegotiated.
Layer 4 - Control & Monitoring	Operational resilience testing including threat-led penetration testing reaches AI components; Layer 4 testing programs must include adversarial ML.
Layer 5 - Audit & Evidence	ICT-related incident reporting and ICT risk management framework documentation are core Layer 5 deliverables; ESAs may require evidence on inquiry.

EU Data Act (Regulation (EU) 2023/2854)

Regulation (EU) 2023/2854

Jurisdiction	European Union
Effective	September 12, 2025
Regulator	National competent authorities (varies by Member State)
Scope	Manufacturers of connected products and providers of related services placing them on the EU market; data holders, data recipients, and data processing services

Applicability

The Data Act establishes user rights to access and port data generated by connected products (IoT) and related services, restricts contractual lock-in by data processing services (cloud, edge, AI inference), enables business-to-government data sharing in emergencies, and harmonizes contractual rules for data

sharing. AI implications include training-data sourcing from connected products, AI-as-a-service portability, and switching obligations for AI inference providers.

Core Obligations

- Connected product manufacturers must design products to make data accessible to users; provide data to users on request; permit users to share data with third parties (subject to certain restrictions).
- Data processing service providers (including AI inference services) must facilitate switching: contractual maximum 30-day transition period; gradual reduction of switching charges through 2027 with elimination thereafter; equivalent functionality requirements.
- B2G data sharing in case of exceptional need (public emergencies).
- Restrictions on unfair contractual terms imposed on SMEs.

Penalties & Enforcement

Member State penalties (varies); specific penalties for personal data violations align with the GDPR framework.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	IoT data sourcing for AI training implicates user data-access rights; data must be made available with metadata sufficient for use.
Layer 3 - System Integration	AI-as-a-service portability obligations reshape Layer 3 architecture; lock-in through proprietary embeddings and inference formats becomes a contractual liability.

ePrivacy Directive (Cookie Law)

Directive 2002/58/EC as amended (transposed by Member States)

Jurisdiction	European Union (Member State implementation varies)
Effective	2002 (continuously amended; pending replacement by ePrivacy Regulation)
Regulator	National DPAs and telecommunications regulators
Scope	Providers of publicly available electronic communications services and operators of websites/apps storing or accessing information on user terminal equipment

Applicability

The ePrivacy Directive governs cookies, similar tracking technologies, electronic marketing communications, and confidentiality of communications. Article 5(3) requires consent for storing or accessing information on terminal equipment except where strictly necessary for service requested. The pending ePrivacy Regulation has been delayed but remains on the agenda. AI implications include

consent for analytics, AI-driven profiling cookies, and the use of communications metadata in model training.

Core Obligations

- Obtain consent before storing or accessing non-essential cookies and similar technologies (Article 5(3)).
- For unsolicited electronic marketing communications: prior opt-in consent (with limited exceptions for existing customer relationships).
- Maintain confidentiality of communications and traffic data; restrict use of metadata.

Penalties & Enforcement

Member State penalties; the GDPR penalty framework applies where cookie consent failures also constitute unlawful processing of personal data.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Cookie-derived behavioral data feeds many AI models; consent failures upstream undermine downstream training-data lawfulness.

EU Digital Markets Act (DMA)

Regulation (EU) 2022/1925

Jurisdiction	European Union (extraterritorial)
Effective	May 2, 2023 (designated gatekeepers obligations from March 7, 2024)
Regulator	European Commission
Scope	Designated "gatekeepers" providing core platform services

Applicability

The DMA imposes ex ante obligations on designated gatekeepers including data-portability obligations, restrictions on combining personal data across services, and requirements for fair access to services. AI-relevant obligations include restrictions on using business user data to compete; requirements for interoperability; and consent requirements for combining personal data across services that condition access to AI-powered features.

Core Obligations

- For designated gatekeepers: refrain from combining personal data across services without explicit consent; permit business users to promote competing offers; provide interoperability obligations for messaging core platform services; provide effective data portability.

Penalties & Enforcement

Up to 10% of worldwide annual turnover (20% for repeated infringement); periodic penalty payments up to 5% of average daily worldwide turnover.

STACK LENS - How this law maps to the AI Governance Stack

Layer 3 - System Integration	AI feature delivery that depends on cross-service personal data combination requires Layer 3 architectural separation; gatekeepers must operationally enforce consent at the boundary.
-------------------------------------	--

Additional EU Regulation - CRA, PLD, eIDAS 2.0, EHDS, and the GPAI Code of Practice

Beyond the AI Act and GDPR, the EU has assembled a layered regulatory ecosystem addressing product cybersecurity (Cyber Resilience Act), product liability (the revised Product Liability Directive and pending AI Liability Directive), digital identity (eIDAS 2.0), cybersecurity certification (the Cybersecurity Act and ENISA), data sharing (Data Governance Act), and sector-specific frameworks (European Health Data Space). The GPAI Code of Practice operationalizes the AI Act's general-purpose AI obligations.

EU Cyber Resilience Act (CRA)

Regulation (EU) 2024/2847

Jurisdiction	European Union (extraterritorial)
Effective	December 10, 2024 (entry into force); main obligations apply December 11, 2027 (vulnerability and incident reporting from September 11, 2026)
Regulator	Market surveillance authorities; ENISA (single reporting platform); European Commission
Scope	Manufacturers, importers, and distributors of products with digital elements (PDEs) placed on the EU market

Applicability

The Cyber Resilience Act establishes mandatory cybersecurity requirements for products with digital elements throughout their lifecycle. AI models, AI-enabled software products, and connected devices incorporating AI are within scope. Critical and important products face conformity assessment by notified bodies; standard products use self-assessment. The CRA interlocks with the EU AI Act for AI-enabled products and with NIS2 for product cybersecurity.

Core Obligations

- Design products with digital elements in accordance with cybersecurity essential requirements (Annex I) including no known exploitable vulnerabilities at release, secure default configurations, vulnerability handling, security updates, encryption of data at rest and in transit, etc.
- Conduct cybersecurity risk assessment as part of planning, design, development, production, delivery, and maintenance.
- For "important" products (Annex III): self-assessment with documentation requirements. For "critical" products (Annex IV): notified body conformity assessment.
- Vulnerability handling: report actively exploited vulnerabilities to ENISA single reporting platform within 24 hours; severe incidents within 72 hours; final reports within 14 days.
- Provide security updates for the product's expected lifecycle (default 5 years where appropriate).
- CE marking, declaration of conformity, technical documentation.

Penalties & Enforcement

Up to €15M or 2.5% of worldwide turnover for essential requirement violations; up to €10M or 2% for other obligations; up to €5M or 1% for misleading information.

Recent Developments (through 2026)

CRA harmonized standards are being developed by CEN-CENELEC; the AI-specific overlay is in active development. The vulnerability and incident reporting effective date (September 2026) is the principal practitioner milestone in 2026.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	AI models embedded in products must satisfy CRA essential requirements; Layer 2 model security (adversarial robustness, prompt injection defense) is a regulatory expectation.
Layer 3 - System Integration	AI component supply chain integration falls under CRA vulnerability handling; product manufacturers inherit upstream AI vulnerability obligations.
Layer 4 - Control & Monitoring	24-hour exploited-vulnerability reporting and 72-hour incident reporting are aggressive Layer 4 obligations.
Layer 5 - Audit & Evidence	Technical documentation requirements parallel AI Act Annex IV; combining documentation packages is the efficient practice.

EU AI Liability Directive (Pending) and Revised Product Liability Directive

Proposed AI Liability Directive (COM/2022/496); Revised Product Liability Directive 2024/2853

Jurisdiction	European Union
Effective	PLD effective December 9, 2024 (Member State transposition by December 9, 2026); AILD pending - withdrawn from priority list 2025
Regulator	EU national courts; Member State product safety authorities

Scope	PLD: economic operators placing products on the EU market; AILD: providers and users of AI systems (if enacted)
--------------	---

Applicability

The revised Product Liability Directive (PLD) explicitly extends product liability to software, including AI systems; updates definitions of "product," "defect," "damage," and "economic operator"; eases the plaintiff's burden in complex AI cases through disclosure of evidence and rebuttable presumptions of defect or causation; and recognizes new categories of compensable damage (data corruption, psychological harm). The AI Liability Directive (AILD) was proposed to harmonize non-contractual civil liability rules for AI but has been delayed; the European Commission removed it from the 2025 priority list, and revival in 2026–2027 is uncertain.

Core Obligations

- PLD: economic operators face strict liability for defective products causing damage; AI providers should anticipate evidence-disclosure obligations and rebuttable defect/causation presumptions.
- Maintain technical documentation (paralleling AI Act Annex IV) supporting defect/causation analysis.
- AILD (if enacted): rebuttable causal presumptions for AI-related fault claims; disclosure obligations.

Penalties & Enforcement

PLD: damages including material harm, personal injury, property damage, data corruption, and psychological harm.

Recent Developments (through 2026)

PLD transposition deadline December 2026 is the central 2026 milestone. AILD's prospects are uncertain; the European Parliament has urged revival, but Commission strategy remains in flux.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	AI Act Annex IV documentation supports PLD defect-and-causation defense; Layer 5 evidence packages are the principal liability defense infrastructure.

eIDAS 2.0 and the European Digital Identity Wallet

Regulation (EU) 2024/1183 amending Regulation (EU) No 910/2014

Jurisdiction	European Union
Effective	May 20, 2024; Member State wallet provision by 2026
Regulator	National supervisory authorities; European Commission

Scope	Trust service providers; European Digital Identity Wallet (EUDIW) providers; relying parties
--------------	--

Applicability

eIDAS 2.0 establishes the European Digital Identity Wallet - a state-issued digital identity wallet enabling EU citizens to identify themselves and share verified attributes across the EU. AI implications include: (1) using EUDIW for AI-driven identity verification; (2) integrating EUDIW with KYC/AML AI tools; (3) restrictions on relying parties' use of identity data; and (4) the EUDIW's zero-knowledge attribute disclosure capabilities affecting AI training data acquisition.

Core Obligations

- For Member States: provide EUDIW by 2026.
- For relying parties (including AI services using EUDIW for identity): register, comply with attribute-use restrictions, support privacy-preserving disclosure.
- For trust service providers: comply with qualified trust service requirements.

Penalties & Enforcement

National supervisory authority penalties.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	EUDIW supports privacy-preserving identity verification - a Layer 1 alternative to bulk PII collection for AI services.
Layer 3 - System Integration	EUDIW integration is a Layer 3 architectural option for KYC and identity flows in EU operations.

EU Cybersecurity Act and ENISA

Regulation (EU) 2019/881

Jurisdiction	European Union
Effective	June 27, 2019
Regulator	ENISA (EU Agency for Cybersecurity); national cybersecurity certification authorities
Scope	EU cybersecurity certification framework; ICT products, processes, and services

Applicability

The EU Cybersecurity Act establishes the European cybersecurity certification framework, providing a unified mechanism for ICT product certification across the EU. The European Common Criteria-based

scheme (EUCC) was adopted in 2024; the European Cloud Services certification scheme (EUCS) and AI-specific certification schemes are in development. EU certification is increasingly referenced in NIS2, DORA, CRA, and procurement frameworks.

Core Obligations

- Voluntary certification under adopted schemes; in-scope manufacturers may seek certification at basic, substantial, or high assurance levels.
- For products in critical sectors: certification may become mandatory through implementing acts.
- ENISA maintains the Common Criteria-based EUCC scheme.

Penalties & Enforcement

Loss of certification; market access consequences in mandatory-certification contexts.

Recent Developments (through 2026)

The EUCS (cloud services) scheme has been delayed pending Member State agreement on sovereignty requirements. ENISA's 2024 AI Threat Landscape report and 2025 ML Security Guidance establish operational expectations for AI cybersecurity in EU contexts.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	EU certification provides a Layer 5 cross-jurisdictional compliance demonstration; alignment with EUCC or future AI schemes supports defensibility.

EU Data Governance Act

Regulation (EU) 2022/868

Jurisdiction	European Union
Effective	September 24, 2023
Regulator	National competent authorities; European Data Innovation Board
Scope	Public sector body re-use of protected data; data intermediation services; data altruism organizations

Applicability

The Data Governance Act facilitates re-use of certain categories of protected public sector data, regulates data intermediation services (including data marketplaces), and establishes a framework for data altruism. AI implications include access to public-sector training data, regulation of AI marketplaces and data exchanges, and the data altruism framework as a mechanism for ethically sourced AI training data.

Core Obligations

- For public sector bodies: facilitate re-use of certain protected data subject to conditions including anonymization, IP-respecting access, and processing in secure environments.
- For data intermediation service providers: notification to national authority; functional separation from data services using the data; neutrality requirements.
- For recognized data altruism organizations: registration; transparency; compliance with rulebook (when adopted).

Penalties & Enforcement

National competent authority penalties.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	DGA mechanisms expand Layer 1 lawful sourcing options for AI training including public-sector data with privacy-preserving processing requirements.

EU Network and Information Security (NIS2) Sectoral Implementing Acts

Commission Implementing Regulation (EU) 2024/2690 (technical and methodological requirements)

Jurisdiction	European Union
Effective	October 18, 2024 (technical requirements for digital service providers)
Regulator	National CSIRTs and competent authorities; ENISA
Scope	Digital service providers including DNS, TLD, cloud computing, data center, content delivery network, managed service providers, online marketplaces, online search engines, social networking, and trust service providers

Applicability

The NIS2 implementing regulation establishes specific technical and methodological cybersecurity risk-management measures for digital service providers, including detailed requirements applicable to cloud and AI-as-a-service offerings. The regulation transcribes generic NIS2 obligations into operational technical controls.

Core Obligations

- Implement specific controls including policy on cybersecurity risk management, risk-management framework, incident handling, business continuity, supply chain security, etc.
- Detailed requirements for human resources, asset management, access control, cryptography, secure development, network security, monitoring and logging.
- Compliance is mandatory for in-scope digital service providers regardless of national transposition.

Penalties & Enforcement

NIS2 penalties - up to €10M or 2% of worldwide turnover for essential entities.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	Implementing regulation provides specific Layer 4 control requirements for digital service providers; AI-as-a-service offerings should map controls explicitly.

EU AI Act Code of Practice for General-Purpose AI

GPAI Code of Practice (final text published 2025); EU AI Office

Jurisdiction	European Union (voluntary code, with regulatory weight)
Effective	Code of Practice published 2025; companies may adhere from publication
Regulator	EU AI Office; signatory companies; Code drafting facilitators
Scope	Providers of general-purpose AI models including those with systemic risk

Applicability

The GPAI Code of Practice operationalizes EU AI Act Articles 51–55 for general-purpose AI providers. Adherence creates a presumption of compliance with the Act's GPAI obligations; non-adherence requires alternative compliance demonstration. The Code addresses transparency, copyright compliance, training data summary publication, and (for systemic-risk models) model evaluation, systemic risk assessment, mitigation, incident reporting, and cybersecurity protections.

Core Obligations

- For all GPAI providers: technical documentation, information for downstream providers, copyright policy, training data summary.
- For systemic-risk GPAI providers (>10²⁵ FLOPs threshold): model evaluation, systemic risk assessment and mitigation, serious incident reporting, cybersecurity protections.
- Code adherents commit to specific operational practices; non-adherents must demonstrate equivalent compliance through other means.

Penalties & Enforcement

EU AI Act Article 99 penalties: up to €15M or 3% of worldwide turnover for GPAI obligation violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Systemic-risk GPAI evaluation obligations operationalize Layer 2 frontier model safety practice.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

Training data summary publication and technical documentation are extensive Layer 5 obligations; the Code provides operational templates.

EU Health Data Space (EHDS) Regulation

Regulation (EU) 2025/327 (European Health Data Space)

Jurisdiction	European Union
Effective	March 26, 2025 (entry into force); phased application through 2029
Regulator	National health data access bodies; European Health Data Space Board
Scope	Health data processing in primary use (healthcare delivery) and secondary use (research, public health, AI training, regulatory)

Applicability

EHDS establishes EU-wide rules for health data primary use (cross-border healthcare) and secondary use (research, AI development, public health, regulatory purposes). For AI in healthcare, EHDS's secondary-use framework is particularly significant: it permits use of health data for AI training under specified conditions including national health data access body authorization, data altruism principles, and specified processing environments. EHDS interacts with GDPR, EU AI Act, and Medical Device Regulation.

Core Obligations

- For healthcare providers: provide patient access to electronic health data and support cross-border exchange via MyHealth@EU infrastructure.
- For secondary use (e.g., AI training): apply via national health data access body; comply with permitted purposes, processing in secure environment, and rules on data minimization.
- For health data holders: provide data to health data access bodies on request for permitted secondary uses.
- For data users (including AI developers): comply with terms of authorized access including no re-identification attempts.

Penalties & Enforcement

EHDS-specific penalties under national implementations; alignment with GDPR and AI Act penalty frameworks for related violations.

STACK LENS - How this law maps to the AI Governance Stack

Layer 1 - Data Governance	EHDS secondary-use framework provides Layer 1 lawful basis for healthcare AI training within specified constraints; significant sourcing opportunity for EU operations.
Layer 3 - System Integration	Secure processing environment requirements shape Layer 3 architecture for EHDS-sourced data.

P A R T V I

United Kingdom & The Americas

Post-Brexit UK regulation, Canadian federal and provincial regimes, Brazil's LGPD and AI bill, and the leading Latin American privacy frameworks from Mexico to Chile.

United Kingdom and Major American Regimes

The UK pursues a sectoral, principles-based AI regulatory model alongside data protection reform via the Data (Use and Access) Act. Canada, Brazil, Mexico, and Argentina constitute the principal Americas regimes alongside the U.S.

UK GDPR and the Data Protection Act 2018

UK GDPR (retained EU law); Data Protection Act 2018; Data (Use and Access) Act 2025

Jurisdiction	United Kingdom
Effective	May 25, 2018; UK GDPR effective January 1, 2021 (post-Brexit); Data (Use and Access) Act 2025 progressively implemented
Regulator	Information Commissioner's Office (ICO)
Scope	Controllers and processors processing personal data of UK individuals; offering goods/services to UK individuals; or monitoring behavior in the UK

Applicability

The UK GDPR closely parallels the EU GDPR with several substantive divergences widening over time. The Data (Use and Access) Act 2025 (DUAA) introduces UK-specific reforms including changes to automated decision-making (Article 22 reform), legitimate interests (a new "recognised legitimate interests" framework), research provisions, smart data schemes, and digital verification services. The ICO's 2024 strategic AI plan and 2025 statutory codes establish UK-specific AI governance expectations.

Core Obligations

- All GDPR obligations as adapted for the UK context.
- DUAA-modified Article 22: solely automated decisioning may be permitted on a wider basis (legitimate interest, contract, consent) but with retained safeguards including transparency, human review on request, and contest rights; tighter restrictions where special category data is involved.
- Recognised legitimate interests: certain processing activities (national security, defence, public emergencies, safeguarding vulnerable individuals, crime prevention) presumed legitimate, eliminating the balancing-test requirement.
- ICO statutory codes including the Children's Code (Age Appropriate Design Code) and the AI Code of Practice (in development).
- Smart data schemes and digital verification services with sectoral implementation timelines.

Penalties & Enforcement

Aligned with EU GDPR - up to £17.5M or 4% of worldwide turnover; ICO has actively used enforcement powers including against Clearview AI (£7.5M, 2022 - overturned on jurisdictional grounds), TikTok (£12.7M, 2023), and several AI-adjacent matters.

Recent Developments (through 2026)

The ICO's 2024–2025 work on generative AI consultations, Snapchat MyAI investigation, and continuing scrutiny of facial recognition deployments establish the ICO as the most operationally active AI regulator in Europe outside the EU AI Office.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	UK reforms widen the scope for AI training under legitimate interests but retain transparency and accountability obligations at Layer 1.
Layer 2 - Model Governance	Reformed Article 22 maintains explanation and contest rights - Layer 2 explainability remains mandatory for solely automated decisions with legal/similar effects.
Layer 5 - Audit & Evidence	ICO statutory codes (e.g., Children's Code) function as Layer 5 audit benchmarks; non-compliance creates evidentiary presumptions.

UK Sectoral AI Regulation and AI Safety Institute

AI White Paper (March 2023); A pro-innovation approach to AI regulation; AI Opportunities Action Plan (January 2025)

Jurisdiction	United Kingdom
Effective	Ongoing; statutory framework under consideration
Regulator	Sectoral regulators (FCA, ICO, Ofcom, MHRA, CMA, others) coordinated through Cabinet Office and DSIT; UK AI Safety Institute
Scope	AI systems within sectoral remits; voluntary frontier model assessment via the AI Safety Institute

Applicability

The UK has pursued a sectoral, principles-based AI regulatory model rather than horizontal legislation. Five cross-sectoral principles (safety, transparency and explainability, fairness, accountability and governance, contestability and redress) are operationalized by existing sector regulators. The AI Safety Institute (AISI) conducts pre-deployment evaluation of frontier models. The 2025 Action Plan signals an intent to legislate for the most powerful AI models while continuing the sectoral approach for the broader market.

Core Obligations

- Sectoral regulator guidance: FCA on AI in financial services, Ofcom on online safety and AI, MHRA on AI medical devices, CMA on competition aspects, ICO on data protection.
- Voluntary AISI evaluations for frontier models (effectively expected for major frontier model releases targeting the UK).
- Cross-sectoral principles operationalized via regulator-specific guidance and enforcement.

Penalties & Enforcement

Sectoral regulator penalties under existing statutory frameworks; AISI evaluations are voluntary but commercially important.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	AISI evaluations focus on Layer 2 model risk including dangerous capabilities, security uplift, and societal harms.
Layer 5 - Audit & Evidence	Sectoral regulator guidance translates into Layer 5 documentation expectations; compliance evidence must address each applicable regulator's priorities.

UK Online Safety Act

Online Safety Act 2023

Jurisdiction	United Kingdom (extraterritorial)
Effective	Phased from October 2023; child-safety duties effective July 2025
Regulator	Office of Communications (Ofcom)
Scope	User-to-user services and search services with links to the UK

Applicability

The Online Safety Act imposes duties on user-to-user and search services to address illegal content and (for services likely to be accessed by children) content harmful to children. AI implications include risk assessments for AI-generated illegal content; age assurance technology (with attendant privacy considerations); and risk assessment for AI-driven recommender systems.

Core Obligations

- Risk assessments for illegal content, content harmful to children, and (for Category 1 services) certain content harmful to adults.
- Reporting and complaint mechanisms; record-keeping.
- Age assurance for services likely to be accessed by children where appropriate.
- Specific duties on illegal content categories (CSAM, terrorism, fraud, hate, harassment, etc.).
- Transparency reports for Categories 1, 2A, and 2B.

Penalties & Enforcement

Up to £18M or 10% of qualifying worldwide revenue (whichever higher); senior management criminal liability for certain failures.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	AI content moderation systems are Layer 2 deployments subject to risk assessment and accuracy expectations.
Layer 4 - Control & Monitoring	Real-time content monitoring and reporting infrastructure are Layer 4 operational deliverables.

Canada - PIPEDA and the Pending Consumer Privacy Protection Act

Personal Information Protection and Electronic Documents Act, S.C. 2000, c. 5; pending CPPA (Bill C-27)

Jurisdiction	Canada (federal); substantially similar provincial laws in Quebec, BC, Alberta
Effective	PIPEDA April 13, 2000; CPPA pending
Regulator	Office of the Privacy Commissioner of Canada (OPC); provincial privacy commissioners
Scope	Federal works, undertakings, and businesses; commercial activities in provinces without substantially similar laws

Applicability

PIPEDA governs the collection, use, and disclosure of personal information in commercial activities, structured around 10 fair information principles. Canada's long-pending CPPA reform (and the parallel Artificial Intelligence and Data Act, AIDA) would substantially modernize the framework - including AI-specific obligations - but the legislative path is uncertain. Quebec's Law 25 (effective in stages 2022–2024) has become the de facto Canadian privacy floor for AI use, requiring consent before automated decisioning of personal information producing significant effects, transparency, and the right to a human review.

Core Obligations

- Comply with the 10 fair information principles: accountability, identifying purposes, consent, limiting collection, limiting use/disclosure/retention, accuracy, safeguards, openness, individual access, challenging compliance.
- Quebec Law 25: appoint a privacy officer; conduct privacy impact assessments for technology projects; obtain consent before using personal information for automated decisioning producing significant effects; provide explanation and human review on request; data portability; mandatory breach notification.

- Mandatory breach reporting to the OPC for breaches of security safeguards involving "real risk of significant harm."

Penalties & Enforcement

PIPEDA: limited penalties; OPC primarily uses investigation and findings. Quebec Law 25: penalties up to CAD \$25M or 4% of worldwide turnover. CPPA, if enacted, would introduce GDPR-scale penalties.

Recent Developments (through 2026)

OPC and provincial commissioners issued joint guidance on generative AI in December 2023 establishing six principles. The 2024 OpenAI investigation and 2025 Clearview AI follow-up are the most consequential AI enforcement matters. AIDA remains under consideration.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Quebec Law 25 PIA requirement and OPC accountability principle drive Layer 1 documentation.
Layer 2 - Model Governance	Automated decision transparency and human-review obligations parallel GDPR Article 22.
Layer 3 - System Integration	Cross-border transfer obligations under Quebec Law 25 are among the strictest in the Americas.

Brazil - Lei Geral de Proteção de Dados (LGPD) and Marco Legal da Inteligência Artificial

Lei Nº 13.709/2018 (LGPD); PL 2338/2023 (AI Bill, pending)

Jurisdiction	Brazil
Effective	LGPD August 2020; administrative penalties August 2021; pending AI law
Regulator	Autoridade Nacional de Proteção de Dados (ANPD)
Scope	Processing of personal data in Brazil; processing where data subjects are in Brazil at time of collection; processing for offering goods/services in Brazil

Applicability

LGPD is the GDPR-style Brazilian privacy law with broadly parallel principles, lawful bases, data subject rights, and accountability obligations. The pending AI Bill (PL 2338/2023, advancing through Congress) would establish a risk-based AI framework with explicit prohibitions, high-risk categorization, and a national AI authority. ANPD has prioritized AI-related guidance including a 2024 generative AI guide and 2025 high-risk processing definitions.

Core Obligations

- Comply with LGPD lawful bases (10 enumerated, broader than GDPR including legitimate interest, contract, legal/regulatory obligation, public administration, research, credit protection, life protection, and health).
- Honor data subject rights: confirmation of processing, access, correction, anonymization/blocking/deletion, portability, deletion of consent-based processing, information about sharing, information about possibility of refusing consent, revocation of consent.
- Conduct Data Protection Impact Reports (Relatório de Impacto à Proteção de Dados Pessoais) for high-risk processing.
- Appoint a Data Protection Officer (Encarregado).
- Notify ANPD of incidents that may pose risk or relevant damage to data subjects.

Penalties & Enforcement

Administrative penalties up to BRL 50M per infraction or 2% of revenue in Brazil for the prior fiscal year; daily penalties for continuing violations; suspension of processing and publication of infractions.

Recent Developments (through 2026)

ANPD's 2024 generative AI guide and 2025 enforcement plan position Brazil as a leading Latin American AI regulator. The pending AI Bill would create the most comprehensive AI framework in the Americas.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	LGPD's broader lawful basis menu permits more flexible Layer 1 training data postures than GDPR but with comparable documentation.
Layer 5 - Audit & Evidence	DPIR (RIPD) is the principal Layer 5 deliverable; ANPD enforcement increasingly requests these on inquiry.

Mexico - Federal Law on Protection of Personal Data Held by Private Parties

Ley Federal de Protección de Datos Personales en Posesión de los Particulares (LFPDPPP, 2010)

Jurisdiction	Mexico
Effective	2010; substantial 2025 reforms
Regulator	Instituto Nacional de Transparencia, Acceso a la Información y Protección de Datos Personales (INAI), with reforms transferring functions to a successor agency
Scope	Private parties processing personal data in Mexico

Applicability

The LFPDPPP establishes Mexican privacy obligations including notice (privacy notice or aviso de privacidad), consent, data subject ARCO rights (Access, Rectification, Cancellation, Opposition), and

security obligations. AI-specific obligations are emerging through INAI guidance rather than statutory amendment.

Core Obligations

- Provide a privacy notice (aviso de privacidad) at the time of data collection.
- Obtain consent (express or tacit depending on data type and processing).
- Honor ARCO rights and respond to requests within 20 business days.
- Implement security measures consistent with INAI guidance.
- Notify INAI of breaches significantly affecting data subjects' rights.

Penalties & Enforcement

Administrative fines up to approximately MXN 32M (varying by violation type and inflation indexation).

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Privacy notice content requirements demand Layer 1 transparency about AI uses; vague notices fail INAI scrutiny.

Argentina - Personal Data Protection Law and Pending Reform

Ley 25.326 (PDPL, 2000); pending reform Ley de Protección de Datos Personales

Jurisdiction	Argentina
Effective	2000; significant reform under consideration
Regulator	Agencia de Acceso a la Información Pública (AAIP)
Scope	Public and private sector processing of personal data in Argentina

Applicability

Argentina was the first Latin American country with EU adequacy status. The pending PDPL reform would modernize the framework consistent with GDPR principles and add AI-specific obligations. AAIP has issued AI guidance including for automated decision-making.

Core Obligations

- PDPL fair information principles, ARCO rights, and consent obligations.
- Registration with the National Database Registry (in transition under reform).
- Cross-border transfer restrictions; adequacy assessments.

Penalties & Enforcement

Administrative fines under PDPL framework; reform anticipated to add GDPR-scale penalties.

STACK LENS - How this law maps to the AI Governance Stack**Layer 5 - Audit & Evidence**

AAIP guidance requires documentation supporting automated decisioning logic - Layer 5 transparency for explanation rights.

Additional Latin American Regimes

Chile's 2024 reform (effective 2026), Colombia's SIC-led enforcement, and the legacy and emerging frameworks of Peru, Uruguay, Costa Rica, Ecuador, the Dominican Republic, and Panama complete the Latin American privacy landscape.

Chile - Personal Data Protection Law

Law No. 19.628 (1999, as amended by Law No. 21.719 enacted 2024)

Jurisdiction	Republic of Chile
Effective	Original 1999; comprehensive reform Law 21.719 effective December 1, 2026
Regulator	Personal Data Protection Agency (in formation under Law 21.719)
Scope	Public and private bodies processing personal data

Applicability

Chile's 2024 reform (Law 21.719) substantially modernized the country's privacy regime to GDPR alignment, including the creation of an independent data protection authority, expanded data subject rights, lawful basis framework, sensitive data protections, breach notification, and substantial penalties. The reform takes effect December 2026; practitioners with Chilean operations should be in active preparation.

Core Obligations

- Lawful basis for processing.
- Sensitive personal data requires explicit consent or specific exception.
- Data subject ARCO+ rights (access, rectification, cancellation, opposition, portability, blocking, opposition to automated decisions).
- Cross-border transfer restrictions; adequacy and other authorized mechanisms.
- Mandatory breach notification.
- DPO required for certain processing.
- Records of processing activities.

Penalties & Enforcement

Administrative fines up to 4% of annual turnover for serious violations under Law 21.719; current Law 19.628 penalties more limited.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Reform implementation creates immediate Layer 5 documentation work - DPO, RoPA, and breach response infrastructure are 2026 priorities.

Colombia - Personal Data Protection Law

Law 1581 of 2012; Decree 1377 of 2013; Decree 90 of 2018

Jurisdiction	Republic of Colombia
Effective	2012; subsequent decrees
Regulator	Superintendencia de Industria y Comercio (SIC) - Delegatura de Protección de Datos
Scope	Public and private bodies processing personal data of Colombian data subjects

Applicability

Colombia's privacy framework is GDPR-aligned in many respects. SIC has been highly active in enforcement, including against tech and ad-tech companies. Colombia has signaled intent toward AI-specific regulation; SIC has issued guidance on automated decisioning.

Core Obligations

- Constitutional habeas data right; specific principles of legality, purpose, freedom, truthfulness, transparency, restricted access, security, confidentiality.
- Sensitive personal data and children's data heightened protections.
- Registration with SIC RNBD (National Registry of Databases) for certain controllers.
- Cross-border transfer restrictions; adequate countries list maintained by SIC.
- Mandatory breach notification within 15 working days.
- Data subject rights including access, knowledge, update, rectification, revocation of authorization, supersession.

Penalties & Enforcement

Administrative fines up to ~2,000 monthly minimum wages (USD ~1.4M+ depending on indexation); operational suspensions.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	SIC RNBD registration is a Layer 5 administrative obligation; database inventory documentation supports SIC inquiries.
---------------------------------------	--

Peru - Personal Data Protection Law

Law No. 29733 (2011); Regulation Supreme Decree No. 003-2013-JUS

Jurisdiction	Republic of Peru
Effective	May 8, 2013
Regulator	Autoridad Nacional de Protección de Datos Personales (ANPDP) within Ministry of Justice
Scope	Public and private bodies processing personal data

Applicability

Peru's law is structured around personal data protection principles, requires registration of databases, sensitive data protections, and cross-border transfer restrictions. Reform aligning more closely with GDPR has been considered.

Core Obligations

- Registration of databases with ANPDP.
- Consent-based processing; specific exceptions.
- Sensitive personal data requires explicit consent.
- Cross-border transfer restrictions to non-adequate jurisdictions.

Penalties & Enforcement

Administrative fines.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	Database registration is a Layer 5 baseline.
---------------------------------------	--

Uruguay - Personal Data Protection Law

Law No. 18.331 (2008); Regulation Decree No. 414/009

Jurisdiction	Eastern Republic of Uruguay
Effective	August 11, 2008

Regulator	Unidad Reguladora y de Control de Datos Personales (URCDP)
Scope	Public and private bodies processing personal data

Applicability

Uruguay was the first South American country to obtain EU adequacy. The law is comprehensive and GDPR-aligned with provisions on data subject rights, sensitive data, registration, and cross-border transfer.

Core Obligations

- Registration of databases with URCDP.
- Consent or other lawful basis.
- Sensitive data protections.
- Cross-border transfer to non-adequate countries requires consent or contractual safeguards.
- DPO designation requirements (added through 2019 reform).

Penalties & Enforcement

Administrative fines; reprimands and operational consequences.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	URCDP registration and (where applicable) DPO designation are Layer 5 obligations for Uruguayan operations.
---------------------------------------	---

Costa Rica - Protection of the Person against the Treatment of Personal Data

Law No. 8968 (2011)

Jurisdiction	Republic of Costa Rica
Effective	September 5, 2011
Regulator	Agency for Data Protection of the Inhabitants (PRODHAB)
Scope	Public and private databases

Applicability

Costa Rica's law is structured around data subject self-determination, sensitive data protections, registration, and cross-border transfer. PRODHAB enforcement has focused on database registration and consent compliance.

Core Obligations

- Database registration with PRODHAB.
- Sensitive data heightened protections.
- Data subject ARCO rights.
- Cross-border transfer restrictions.

Penalties & Enforcement

Administrative fines under PRODHAB framework.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	PRODHAB registration is a Layer 5 obligation.

Dominican Republic - Personal Data Protection Law

Law No. 172-13 (2013)

Jurisdiction	Dominican Republic
Effective	December 13, 2013
Regulator	No dedicated supervisory authority; sectoral enforcement (Superintendencia de Bancos for financial; Pro-Consumidor for consumer)
Scope	Public and private bodies processing personal data

Applicability

The DR's law is structured around the constitutional right of self-determination over personal data with detailed obligations. Reform to establish a unified regulator has been under consideration.

Core Obligations

- Consent-based processing; sensitive data heightened protections.
- Sectoral registration in financial sector; consumer protection enforcement otherwise.
- Data subject rights.

Penalties & Enforcement

Sectoral penalties; criminal penalties for serious violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Sectoral enforcement model creates regulator-specific Layer 5 documentation expectations.

Panama - Personal Data Protection Law

Law No. 81 of 2019

Jurisdiction	Republic of Panama
Effective	March 29, 2021
Regulator	National Authority for Transparency and Access to Information (ANTAI)
Scope	Public and private bodies processing personal data of Panama residents

Applicability

Panama's 2019 law is comprehensive and includes data subject rights, sensitive data protections, breach notification, and cross-border transfer requirements. ANTAI has issued AI-relevant guidance through 2025.

Core Obligations

- Consent or other lawful basis.
- Sensitive data requires explicit consent.
- Mandatory breach notification within 72 hours.
- Cross-border transfer to non-adequate countries requires safeguards.
- DPO designation for certain processing.

Penalties & Enforcement

Administrative fines up to USD 100,000+; suspension of operations for serious violations.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	DPO and breach notification create Layer 5 baseline for Panama operations.
---------------------------------------	--

Ecuador - Organic Law on Protection of Personal Data

Organic Law on Protection of Personal Data (LOPDP, 2021)

Jurisdiction	Republic of Ecuador
Effective	May 26, 2021; full enforcement May 26, 2023
Regulator	Personal Data Protection Authority (in formation)
Scope	Public and private bodies processing personal data of natural persons

Applicability

Ecuador's LOPDP is GDPR-aligned with provisions on lawful basis, data subject rights including automated decisioning, sensitive data, breach notification, and cross-border transfer.

Core Obligations

- Lawful basis for processing.
- Sensitive data and children's data heightened protections.
- Data subject rights including access, correction, deletion, opposition, automated decisioning.
- Breach notification.
- DPO designation for certain processing.
- Cross-border transfer restrictions.

Penalties & Enforcement

Administrative fines up to 1% of annual revenue.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

LOPDP DPO and DPIA establish Layer 5 baseline.

P A R T V I I

Asia-Pacific, Middle East & Africa

The most consequential national regimes outside Europe and the Americas - China's data trinity and operational AI rules, Japan's APPI, Korea's AI Basic Act, the deeper Asia-Pacific landscape, the GCC and broader Middle East, and Africa.

Leading Asia-Pacific and MENA Regimes

China, Japan, South Korea, Singapore, India, and Australia constitute the heaviest-consequence APAC regimes for AI. Israel, the UAE, and Saudi Arabia anchor the MENA regulatory landscape. South Africa's POPIA leads African comprehensive privacy law. Each regime has distinct enforcement posture and architectural implications for AI deployments.

China - Personal Information Protection Law (PIPL)

中华人民共和国个人信息保护法 (PIPL, 2021)

Jurisdiction	People's Republic of China (extraterritorial under Article 3)
Effective	November 1, 2021
Regulator	Cyberspace Administration of China (CAC); Ministry of Public Security; State Administration for Market Regulation
Scope	Processing of personal information of natural persons in the territory of China; processing outside China for purposes of providing products/services to natural persons in China, analyzing or evaluating activities of natural persons in China, or other circumstances prescribed by law

Applicability

PIPL is the principal Chinese personal information statute, complementing the Cybersecurity Law (CSL, 2017) and the Data Security Law (DSL, 2021). The "trinity" of data laws governs the full data lifecycle; AI-specific regulations layer on top including the Generative AI Measures, the Algorithmic Recommendation Provisions, and the Deep Synthesis Provisions. Cross-border transfer obligations are among the most prescriptive in the world; data localization requirements apply to critical information infrastructure operators (CIIOs) and to large-scale personal information processors.

Core Obligations

- Lawful basis for processing (consent is the principal basis; six other bases under Article 13 are narrower than GDPR).
- Separate consent (single dianxuan tongyi) required for sensitive personal information processing, cross-border transfer, automated decisioning with significant effect, and several other categories.
- Honor data subject rights including access, copy, correction, deletion, transparency, and the right to refuse automated decision-making with significant effects.
- Conduct a Personal Information Protection Impact Assessment (PIPIA) for sensitive personal information, automated decisioning, entrusted processing, cross-border transfer, and other high-risk processing.

- Cross-border transfer requires one of: CAC security assessment (mandatory for CIIOs and large processors); standard contract filing with CAC; certification by a CAC-recognized body; or other CAC-permitted mechanism.
- Data localization for CIIOs and large-scale personal information processors.
- Designate a representative in China for foreign processors in scope.

Penalties & Enforcement

Administrative fines up to RMB 50M or 5% of prior year turnover; suspension of operations; revocation of licenses; personal liability for responsible persons (RMB 100k–1M plus prohibition orders).

Recent Developments (through 2026)

CAC enforcement of cross-border transfer obligations has tightened through 2025; the Standard Contract regime is the practical compliance vehicle for most multinationals. The Generative AI Measures require pre-market security assessment for "public-facing" generative AI and content moderation aligned with PRC content policies.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Separate consent for sensitive personal information and AI-relevant uses is the most demanding consent regime globally; Layer 1 consent infrastructure must be granular.
Layer 2 - Model Governance	PIPL Article 24 obligations on automated decisioning (transparency, fairness, right to refuse) parallel GDPR Article 22.
Layer 3 - System Integration	Cross-border transfer mechanisms shape Layer 3 architecture; many global AI services architect Chinese operations as data-localized variants.
Layer 5 - Audit & Evidence	PIPIA documentation and CAC security assessment artifacts are core Layer 5 deliverables.

China - Generative AI Measures, Deep Synthesis Provisions, and Algorithmic Recommendation Provisions

Interim Measures for the Management of Generative Artificial Intelligence Services (2023); Provisions on the Administration of Deep Synthesis Internet Information Services (2022); Provisions on the Administration of Algorithmic Recommendations of Internet Information Services (2022)

Jurisdiction	People's Republic of China
Effective	Algorithmic Recommendations: March 1, 2022; Deep Synthesis: January 10, 2023; Generative AI: August 15, 2023
Regulator	Cyberspace Administration of China (CAC); coordinated with NDRC, MIIT, MPS, and other authorities

Scope	Providers of generative AI services, deep synthesis services, and algorithmic recommendation services to the public within the PRC
--------------	--

Applicability

China's AI regulatory regime is the most comprehensive functional AI law in operation today. The Generative AI Measures require pre-market filing with CAC for services with "public opinion attributes or social mobilization capabilities," content alignment with PRC values, training data compliance, content labeling, complaint mechanisms, and minor protection. The Deep Synthesis Provisions require labeling of synthetic content and prohibitions on certain categories of content. The Algorithmic Recommendation Provisions require algorithmic transparency, opt-out mechanisms, and protection of vulnerable groups.

Core Obligations

- Generative AI: pre-market filing with CAC (algorithm filing) for public-facing services; security assessment for services with public opinion or social mobilization capabilities; training data lawfulness and quality; output content alignment with PRC content rules; content labeling; minor protection; user identity verification; complaint mechanism.
- Deep Synthesis: prohibitions on creating false news; consent for using personal biometric information; clear and prominent labeling of synthetic content; security assessment for services that may cause significant impact.
- Algorithmic Recommendations: algorithm filing with CAC; transparency about algorithmic principles; user right to disable algorithmic recommendations or set preferences; protection of minors, elderly, workers, and consumers; refusal to provide unfair advantages.

Penalties & Enforcement

Administrative fines, suspension, license revocation, personal liability; aligned with PIPL/CSL/DSL framework.

Recent Developments (through 2026)

The 2024 generative AI labeling regulation (effective September 2025) tightened content marking obligations. CAC algorithm filings (algorithm registry) function as a public AI registry; published filings are a notable practitioner reference.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Training data lawfulness and quality requirements are explicit and externally assessable; CAC may demand evidence on inquiry.
Layer 2 - Model Governance	Algorithm filing requires Layer 2 documentation including model purpose, mechanisms, and risk mitigations.
Layer 4 - Control & Monitoring	Content labeling and complaint mechanisms are Layer 4 operational obligations.

Japan - Act on the Protection of Personal Information (APPI)

個人情報保護に関する法律 (APPI, 2003; substantially amended 2017, 2020, 2022, 2025)

Jurisdiction	Japan (extraterritorial)
Effective	2003 (continuously amended; 2025 amendments effective in stages)
Regulator	Personal Information Protection Commission (PPC)
Scope	Personal information handling business operators handling personal information of individuals in Japan

Applicability

APPI is the principal Japanese privacy law, with EU adequacy and a regulatory regime closer to GDPR than other Asia-Pacific frameworks. The 2025 amendments address AI-specific concerns including transparency for automated decisioning, third-party provision of pseudonymized information, and tightened enforcement. Japan has pursued an AI principles-based approach with the AI Strategy and the Hiroshima AI Process Code of Conduct rather than horizontal regulation.

Core Obligations

- Obtain consent for use beyond originally specified purposes; provide notice of utilization purposes.
- For sensitive personal information: prior consent for acquisition.
- Restrictions on third-party provision; opt-out for general third-party transfers; consent for sensitive data transfers.
- For pseudonymized information (kanou jouhou) and anonymized information: distinct regulatory regimes permitting flexible use including for AI training.
- Cross-border transfer requires consent or transfer to a country with equivalent protection (EU is recognized; many Asia-Pacific jurisdictions are not).
- Honor data subject rights of disclosure, correction, suspension of use, and deletion.
- Mandatory breach notification.

Penalties & Enforcement

Administrative orders and (for non-compliance with orders) criminal penalties up to JPY 100M for legal entities and JPY 1M plus imprisonment for individuals; civil suits available.

Recent Developments (through 2026)

PPC has issued AI-specific guidance through 2025 emphasizing the responsibilities of AI service providers using personal information. The Hiroshima AI Process produced a voluntary Code of Conduct for advanced AI developers, with reporting framework launched in 2025.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Pseudonymized information regime provides a Layer 1 mechanism for permissible AI training without full anonymization; rules on combining datasets must be observed.
Layer 5 - Audit & Evidence	Hiroshima AI Process Code of Conduct reporting framework is a voluntary but commercially significant Layer 5 disclosure regime.

South Korea - Personal Information Protection Act (PIPA) and AI Basic Act

개인정보 보호법 (PIPA, 2011; substantially amended 2020, 2023, 2025); AI Basic Act (2025)

Jurisdiction	Republic of Korea
Effective	PIPA 2011 (continuously amended); AI Basic Act effective in stages 2025–2026
Regulator	Personal Information Protection Commission (PIPC); Ministry of Science and ICT (AI)
Scope	Public institutions, juridical persons, organizations, and individuals processing personal information in Korea

Applicability

PIPA is one of the most stringent privacy regimes in Asia. The 2023 amendments introduced explicit automated decisioning rights (right to explanation, right to refusal). The 2025 AI Basic Act establishes a comprehensive Korean AI regulatory framework including high-risk AI categorization, transparency obligations, and a national AI safety institute.

Core Obligations

- PIPA: lawful basis (consent is principal); separate consent for sensitive information; data subject rights including access, correction, deletion, suspension of processing; automated decisioning rights (explanation, refusal, human review).
- PIPA cross-border transfer requires consent or equivalent-protection determination by PIPC.
- AI Basic Act: high-risk AI obligations including risk management, transparency, human oversight, and post-market monitoring; foundation model obligations including transparency and safety evaluation; AI safety institute oversight.
- Mandatory breach notification within 72 hours.

Penalties & Enforcement

PIPA: administrative fines up to 3% of related sales (since 2023 amendment); criminal penalties for serious violations. AI Basic Act: penalties scaling with violation severity.

Recent Developments (through 2026)

PIPC has been highly active in AI enforcement, including investigations of facial recognition deployments, generative AI services, and ad-tech profiling. The AI Basic Act puts Korea on a similar trajectory to the EU AI Act with adapted obligations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	AI Basic Act high-risk obligations parallel EU AI Act Annex III requirements with Korean-specific definitions.
Layer 3 - System Integration	Cross-border transfer restrictions shape Layer 3 architecture for Korean operations.
Layer 4 - Control & Monitoring	Automated decisioning rights require Layer 4 explanation and override infrastructure.

Singapore - Personal Data Protection Act (PDPA) and Model AI Governance Framework

Personal Data Protection Act 2012; Model AI Governance Framework v2 (2020); Generative AI Model Framework (2024)

Jurisdiction	Singapore
Effective	PDPA 2014 (continuously amended); Model Frameworks ongoing
Regulator	Personal Data Protection Commission (PDPC); Infocomm Media Development Authority (IMDA)
Scope	Organizations collecting, using, or disclosing personal data in Singapore

Applicability

Singapore has pursued a soft-law plus targeted-statute approach to AI governance, with the Model AI Governance Framework and the AI Verify testing toolkit as principal instruments. The 2020 PDPA amendments expanded data subject rights; the 2024 Generative AI Model Framework establishes voluntary guidelines for foundation model providers and deployers.

Core Obligations

- PDPA: notice, consent, purpose limitation, accuracy, protection, retention limitation, transfer limitation, data portability (effective when commenced), accountability.
- Mandatory breach notification (effective 2021).
- Model AI Governance Framework: internal governance, determining the level of human involvement in AI-augmented decision-making, operations management, stakeholder interaction and communication.

- Generative AI Model Framework: nine dimensions including accountability, data, trusted development and deployment, incident reporting, testing and assurance, security, content provenance, safety and alignment R&D, AI for public good.

Penalties & Enforcement

PDPA: financial penalties up to 10% of annual gross turnover in Singapore (organizations with revenue exceeding SGD 10M) or SGD 1M (others).

Recent Developments (through 2026)

AI Verify Foundation has emerged as a significant international testing and certification platform; Singapore positions itself as the AI governance hub for Asia.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	AI Verify provides a Layer 2 testing toolkit that operationalizes fairness, robustness, and explainability assessments.
Layer 5 - Audit & Evidence	Generative AI Model Framework reporting can serve as a Layer 5 cross-jurisdictional transparency artifact.

India - Digital Personal Data Protection Act (DPDP Act)

Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023); Digital Personal Data Protection Rules (anticipated 2026)

Jurisdiction	India (extraterritorial under Section 3)
Effective	Enacted August 2023; full operationalization pending Rules and notification
Regulator	Data Protection Board of India
Scope	Digital personal data processing in India; processing outside India in connection with offering goods/services to data principals in India

Applicability

The DPDP Act establishes India's first comprehensive personal data protection regime. It uses a consent-centric framework with notice obligations, data subject rights, and a Data Protection Board for enforcement. AI-specific obligations are limited; consequential decision-making rights are not explicit. The 2024–2026 Rules and operational notifications will substantially shape practical compliance.

Core Obligations

- Provide notice in English or any of the 22 scheduled Indian languages, in clear and plain terms.
- Obtain consent that is free, specific, informed, unconditional, and unambiguous; consent withdrawal mechanisms.

- Honor data principal rights of access, correction/erasure, grievance redressal, nomination.
- Process children's data only with verifiable parental consent; prohibitions on tracking, behavioral monitoring, and targeted advertising directed at children.
- For Significant Data Fiduciaries (designated by central government based on volume, sensitivity, risk): additional obligations including DPIA, audit, DPO appointment.
- Mandatory breach notification.
- Cross-border transfer permitted except to countries restricted by central government.

Penalties & Enforcement

Penalties up to INR 250 crore (~USD 30M) per instance of non-compliance, with several enumerated penalty categories.

Recent Developments (through 2026)

The DPDP Rules consultation (2024–2025) addressed many operational questions; final Rules and Board operationalization remain the central 2026 milestone for India practitioners.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Consent-centric framework requires substantial Layer 1 consent infrastructure including multilingual notices.
Layer 5 - Audit & Evidence	Significant Data Fiduciary designation creates Layer 5 audit and DPO obligations comparable to GDPR for large processors.

Australia - Privacy Act 1988 and AI Ethics Framework

Privacy Act 1988 (Cth); Privacy and Other Legislation Amendment Act 2024

Jurisdiction	Australia
Effective	1988 (continuously amended); 2024 Amendment Act in stages
Regulator	Office of the Australian Information Commissioner (OAIC)
Scope	Australian Privacy Principles entities - most government agencies, organizations with revenue exceeding AUD 3M, health service providers, traders in personal information, and others

Applicability

Australia's 2024 Privacy Act reform introduces several AI-relevant changes including a statutory tort for serious invasion of privacy, expanded enforcement powers, and (in subsequent reform tranches anticipated 2025–2026) substantive changes to APP 11 (security), automated decision-making transparency, and a children's online privacy code. The voluntary AI Ethics Framework and the Department of Industry's AI guidance establish principles-based expectations.

Core Obligations

- 13 Australian Privacy Principles including notice, consent, use/disclosure limitations, cross-border transfer restrictions, security, access, correction.
- Notifiable Data Breaches scheme: notify OAIC and affected individuals within reasonable time.
- For automated decision-making (per 2024 reform Tranche 2): privacy policy disclosures and (for significant decisions) information about the kinds of decisions made and the kinds of personal information used.
- Statutory tort for serious invasion of privacy (effective June 10, 2025).

Penalties & Enforcement

Civil penalties up to the greater of AUD 50M, three times the benefit obtained, or 30% of adjusted turnover during the contravention period.

Recent Developments (through 2026)

The 2024 reform substantially increased penalties and added the statutory tort. OAIC Clearview AI determination (2021) and 2024–2025 enforcement against Bunnings (facial recognition) and Optus (breach) are leading cases.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Automated decisioning transparency under Tranche 2 reform requires Layer 2 model documentation accessible to data subjects.
Layer 3 - System Integration	Cross-border transfer restrictions under APP 8 shape Layer 3 vendor architecture.

Israel - Privacy Protection Law (PPL) and AI Policy

Protection of Privacy Law, 5741-1981; PPL Amendment 13 (2024); AI Policy of the Ministry of Innovation, Science and Technology (2023)

Jurisdiction	Israel
Effective	PPL 1981 (continuously amended); Amendment 13 effective August 14, 2025
Regulator	Privacy Protection Authority (PPA)
Scope	Database holders processing personal information; biometric and sensitive data subject to enhanced obligations

Applicability

Israel has EU adequacy. Amendment 13 substantially modernizes Israeli privacy law including expanded definitions of personal information, increased PPA enforcement powers, and enhanced obligations for

sensitive data. Israeli AI policy is currently soft-law plus sectoral; horizontal AI legislation is under consideration.

Core Obligations

- Database registration for certain database categories (transitioning to risk-based approach under Amendment 13).
- Notice and lawful basis for collection.
- Data subject rights of access and correction.
- Security obligations under PPA Information Security Regulations (2017).
- Mandatory breach notification.

Penalties & Enforcement

Significantly enhanced under Amendment 13; administrative penalties up to ILS 3.2M plus per-affected-individual penalties for certain violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	PPA database registration and security regulation compliance documentation are Layer 5 deliverables.

United Arab Emirates - Federal Personal Data Protection Law and AI Strategy

Federal Decree-Law No. 45 of 2021 on Personal Data Protection; UAE National AI Strategy 2031

Jurisdiction	United Arab Emirates (federal); separate DIFC and ADGM regimes
Effective	PDPL 2022 (executive regulations pending); DIFC DP Law 2020
Regulator	UAE Data Office; DIFC Commissioner of Data Protection; ADGM Office of Data Protection
Scope	Processing of personal data in UAE (federal); separate jurisdictional regimes for DIFC and ADGM free zones

Applicability

UAE operates parallel privacy regimes: federal PDPL, DIFC DP Law (GDPR-aligned), and ADGM regulations. AI is governed through the National AI Strategy and ministerial AI guidance rather than horizontal legislation, though several emirate-level initiatives address AI deployment in government services.

Core Obligations

- PDPL: lawful basis, consent for sensitive data, data subject rights, cross-border transfer restrictions, data protection officer for certain processing.

- DIFC DP Law: GDPR-aligned obligations including DPIA, controller/processor distinctions, breach notification.
- ADGM: GDPR-aligned regime including extraterritorial application.

Penalties & Enforcement

PDPL: penalties prescribed by executive regulations (pending); DIFC: penalties up to USD 1M per violation; ADGM: penalties up to USD 28M per violation.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	Multi-zone architecture (federal, DIFC, ADGM) creates Layer 3 jurisdictional segmentation for UAE operations.

Saudi Arabia - Personal Data Protection Law (PDPL)

Saudi Arabia PDPL (Royal Decree M/19, 2021)

Jurisdiction	Kingdom of Saudi Arabia
Effective	September 14, 2024
Regulator	Saudi Data and Artificial Intelligence Authority (SDAIA); National Data Management Office
Scope	Processing of personal data in the Kingdom; processing outside the Kingdom of personal data of residents

Applicability

The Saudi PDPL is a comprehensive data protection regime aligned with GDPR principles, with Saudi-specific elements including data localization for certain categories, Sharia-aligned content considerations, and SDAIA registration. SDAIA is also Saudi Arabia's AI authority and has issued AI-specific principles.

Core Obligations

- Lawful basis for processing; consent is principal basis with enumerated exceptions.
- Data subject rights including access, correction, deletion, withdrawal of consent.
- Cross-border transfer requires SDAIA approval, adequacy determination, or other prescribed mechanism.
- Data localization for sensitive personal data and certain processing.
- Mandatory breach notification.
- DPO designation for certain processing.

Penalties & Enforcement

Administrative fines up to SAR 5M per violation; criminal penalties for sensitive personal data violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Data localization requirements substantially shape Layer 1 data architecture for Saudi operations.

South Africa - Protection of Personal Information Act (POPIA)

Protection of Personal Information Act 4 of 2013

Jurisdiction	Republic of South Africa
Effective	July 1, 2020 (most provisions); enforcement July 1, 2021
Regulator	Information Regulator
Scope	Public and private bodies processing personal information of data subjects in South Africa

Applicability

POPIA is the principal South African privacy law, structured around eight conditions for lawful processing similar to OECD privacy principles. Information Regulator has been increasingly active in AI-related enforcement including against credit bureaus, ad-tech, and biometric deployments.

Core Obligations

- Eight conditions: accountability, processing limitation, purpose specification, further processing limitation, information quality, openness, security safeguards, data subject participation.
- Special personal information requires specific legal grounds (consent or enumerated exceptions).
- Cross-border transfer requires consent, contractual safeguards, equivalent protection, or other prescribed basis.
- Mandatory breach notification.
- Information officer designation.

Penalties & Enforcement

Administrative fines up to ZAR 10M per violation; criminal penalties for certain offenses.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Information officer designation and POPIA compliance manual are Layer 5 documentation requirements.

Additional Asia-Pacific Regimes

Vietnam, Indonesia, Thailand, the Philippines, Malaysia, Hong Kong, Taiwan, and New Zealand constitute the deeper APAC regulatory environment. Vietnam's data localization, Indonesia's GDPR-alignment, and the rapid maturation of Hong Kong's AI guidance are particular practitioner priorities.

Vietnam - Personal Data Protection Decree and Cybersecurity Law

Decree No. 13/2023/ND-CP (PDPD); Law on Cyber Information Security 2015; Cybersecurity Law 2018; pending Personal Data Protection Law

Jurisdiction	Socialist Republic of Vietnam
Effective	PDPD July 1, 2023; PDPL pending
Regulator	Ministry of Public Security (MPS); Department of Cybersecurity and High-Tech Crime Prevention
Scope	Domestic and foreign agencies, organizations, and individuals processing personal data of Vietnamese citizens

Applicability

Vietnam's PDPD is the operative comprehensive privacy regulation. It establishes data subject rights, lawful basis for processing, sensitive data protections, cross-border transfer assessment requirements, and mandatory breach notification within 72 hours. The Cybersecurity Law adds data localization requirements for certain categories of data and certain service providers. Comprehensive PDPL legislation is pending. AI implications include the data localization requirements, MPS oversight, and explicit consent requirements that constrain training data acquisition.

Core Obligations

- Consent (express, voluntary, informed) is principal basis; specific exceptions including legal obligation, vital interests, public interest.
- Sensitive personal data requires explicit consent and additional safeguards.
- Cross-border transfer requires Transfer Impact Assessment (TIA) submitted to MPS.
- Data localization for certain data categories and service providers under Cybersecurity Law.
- Breach notification within 72 hours.
- DPO required for processing of basic personal data; data protection role for processing of sensitive personal data.

Penalties & Enforcement

Administrative fines up to VND 100M; potential criminal penalties; loss of operating license; suspension of operations.

STACK LENS - How this law maps to the AI Governance Stack

Layer 1 - Data Governance	Data localization requirements substantially shape Layer 1 architecture for Vietnam operations.
Layer 3 - System Integration	TIA submission for cross-border transfer is a Layer 3 administrative obligation specific to Vietnam.

Indonesia - Personal Data Protection Law (PDP Law)*Law No. 27 of 2022 on Personal Data Protection*

Jurisdiction	Republic of Indonesia
Effective	October 17, 2022; full enforcement October 17, 2024
Regulator	Ministry of Communication and Informatics (Kominfo); Personal Data Protection Authority (in formation)
Scope	Public and private bodies processing personal data of natural persons in Indonesia or processing of data in connection with activities targeting Indonesia

Applicability

Indonesia's PDP Law is GDPR-aligned with adaptations including a strong consent framework, data subject rights, sensitive data protections, mandatory breach notification, and substantial penalties. The independent Personal Data Protection Authority is in formation; enforcement is currently led by Kominfo. AI implications include explicit recognition of automated decisioning rights and broad cross-border transfer restrictions.

Core Obligations

- Consent or other lawful basis (limited compared to GDPR) for processing.
- Sensitive personal data requires explicit consent and DPIA.
- Data subject rights of access, correction, deletion, withdrawal of consent, objection, and rights related to automated decisions.
- Cross-border transfer requires adequacy, binding agreement, or consent.
- Breach notification within 72 hours.
- DPO required for public bodies, large-scale processing, and processing of sensitive data.

Penalties & Enforcement

Administrative fines up to 2% of annual revenue; criminal penalties for sensitive data violations including imprisonment.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

DPO and DPIA obligations create Layer 5 documentation baseline for Indonesia operations.

Thailand - Personal Data Protection Act (PDPA)

Personal Data Protection Act B.E. 2562 (2019)

Jurisdiction	Kingdom of Thailand
Effective	June 1, 2022 (full effective)
Regulator	Personal Data Protection Committee (PDPC)
Scope	Public and private bodies processing personal data in Thailand or processing for offering goods/services or monitoring of behavior in Thailand

Applicability

Thailand's PDPA is GDPR-aligned with Thai-specific adaptations. It includes consent, lawful basis, sensitive data protections, data subject rights, breach notification, DPO requirements, and cross-border transfer restrictions. PDPC has been increasingly active in enforcement, including in AI-related matters.

Core Obligations

- Lawful basis including consent, contract, legal obligation, vital interests, public task, legitimate interests.
- Sensitive personal data requires explicit consent or one of the additional bases.
- Data subject rights including access, rectification, erasure, restriction, portability, objection.
- Cross-border transfer restrictions; PDPC adequacy designations for permitted destinations.
- Breach notification within 72 hours.
- DPO required for large-scale processing, public bodies, and processing of sensitive data.

Penalties & Enforcement

Administrative fines up to THB 5M; criminal penalties for sensitive data violations.

STACK LENS - How this law maps to the AI Governance Stack

Layer 3 - System Integration

Cross-border transfer adequacy designations shape Layer 3 architecture; non-adequacy destinations require additional safeguards.

Philippines - Data Privacy Act

Republic Act No. 10173 (Data Privacy Act of 2012); Implementing Rules and Regulations

Jurisdiction	Republic of the Philippines
Effective	2012; IRR 2016
Regulator	National Privacy Commission (NPC)
Scope	Public and private bodies processing personal information of Philippine residents or processing in the Philippines

Applicability

The Philippines DPA was an early GDPR-style Asian privacy law. It includes general data privacy principles, sensitive personal information, data subject rights, breach notification, DPO requirements, and Privacy Impact Assessments. NPC has issued AI-related guidance through 2025.

Core Obligations

- Lawful processing principles: transparency, legitimate purpose, proportionality.
- Sensitive personal information requires consent or one of enumerated bases including specific protection.
- Data subject rights of information, access, correction, erasure or blocking, damages, data portability.
- PIA for high-risk processing.
- Mandatory breach notification within 72 hours.
- DPO required for personal information controllers and processors.
- NPC registration for systems processing personal information of 1,000+ individuals.

Penalties & Enforcement

Administrative fines under NPC framework; criminal penalties for unauthorized processing, access, or disclosure (1–6 years imprisonment, fines).

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

NPC registration and PIA documentation are Layer 5 baseline obligations.

Malaysia - Personal Data Protection Act (PDPA)

Malaysia Personal Data Protection Act 2010 (Act 709); 2024 amendments

Jurisdiction	Malaysia
Effective	November 15, 2013; 2024 amendments effective in stages 2025

Regulator	Personal Data Protection Department (PDP Department)
Scope	Persons processing personal data in respect of commercial transactions

Applicability

Malaysia's PDPA is sectorally limited to commercial transactions (excluding government processing). The 2024 amendments substantially modernized the act to align more closely with GDPR principles, including breach notification, DPO requirement, and broadened data subject rights. AI implications include the explicit data portability right and tightened cross-border transfer regime.

Core Obligations

- Seven principles: general, notice and choice, disclosure, security, retention, data integrity, access.
- Post-2024 amendments: mandatory breach notification, DPO designation, data portability right, expanded categories of sensitive data, cross-border transfer reform.
- Class registration for certain controllers.
- Sensitive data requires explicit consent.

Penalties & Enforcement

Administrative fines up to MYR 1M (post-2024); criminal penalties.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	Post-2024 DPO and breach notification create Layer 5 obligations comparable to other GDPR-aligned regimes.

Hong Kong - Personal Data (Privacy) Ordinance (PDPO)

Cap. 486, Personal Data (Privacy) Ordinance

Jurisdiction	Hong Kong Special Administrative Region
Effective	1996 (continuously amended)
Regulator	Office of the Privacy Commissioner for Personal Data (PCPD)
Scope	Data users processing personal data of natural persons

Applicability

Hong Kong's PDPO is structured around six Data Protection Principles. The 2021 amendments criminalized doxxing (publication of personal data with intent to cause harm). The PCPD has issued AI-specific guidance (2024 Model Personal Data Protection Framework for AI) addressing governance, risk

assessment, and accountability for AI systems processing personal data. AI ethics-related work continues through 2025–2026.

Core Obligations

- Six DPPs: lawful and fair collection, accuracy, retention, use limitation, security, openness, data subject access.
- Anti-doxxing offenses (post-2021).
- PCPD AI Framework: AI risk assessment, governance, transparency.
- Cross-border transfer restrictions under § 33 (provision not yet in force, but PCPD guidance addresses recommended practice).

Penalties & Enforcement

Criminal penalties for serious violations; PCPD enforcement notices and (post-2021) doxxing penalties.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	PCPD AI Framework provides operational Layer 2 governance template applicable to Hong Kong AI deployments.

Taiwan - Personal Data Protection Act

Personal Data Protection Act (PDPA, 2010, amended 2023)

Jurisdiction	Taiwan
Effective	2012; 2023 amendments effective phased
Regulator	Personal Data Protection Commission (PDPC, established 2024); previously Ministry of Justice
Scope	Public and private bodies processing personal data

Applicability

Taiwan's PDPA is among the more comprehensive Asian regimes. The 2023 amendments and the 2024 establishment of an independent PDPC have substantially raised enforcement profile. Sensitive personal data requires written consent. AI-specific guidance is being developed by the new PDPC.

Core Obligations

- Lawful basis for processing; sensitive personal data requires written consent or specific exception.
- Data subject rights of access, correction, deletion, supplementation, cessation of processing.
- Cross-border transfer restrictions; central authority restrictions for designated countries.
- Mandatory breach notification.
- Security measures consistent with PDPA Implementation Regulation.

Penalties & Enforcement

Administrative fines; criminal penalties for serious violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Written consent for sensitive personal data shapes Layer 1 collection design for Taiwan operations.

New Zealand - Privacy Act 2020

Privacy Act 2020

Jurisdiction	New Zealand
Effective	December 1, 2020
Regulator	Office of the Privacy Commissioner
Scope	Agencies (public and private) processing personal information of natural persons in New Zealand or extraterritorially in connection with NZ activities

Applicability

New Zealand's Privacy Act 2020 modernized the prior 1993 law with a stronger enforcement framework, mandatory breach notification, cross-border transfer restrictions, and 13 Information Privacy Principles. The Office of the Privacy Commissioner has issued AI guidance including for generative AI use by businesses. New Zealand has EU adequacy.

Core Obligations

- 13 Information Privacy Principles.
- Mandatory breach notification (notifiable privacy breach).
- Cross-border transfer restrictions: receiving agency must be subject to comparable safeguards or have other authorized basis.
- Privacy Officer required.
- Information Privacy Statements for sensitive categories.

Penalties & Enforcement

Compliance notices, complaints, and Human Rights Review Tribunal awards including damages up to NZD 350,000 for class proceedings.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

OPC guidance positions New Zealand as soft-law-leading; voluntary alignment with OPC AI guidance is the practical compliance posture.

Deeper MENA and African Regimes

The Gulf Cooperation Council privacy frameworks (Bahrain, Qatar, Oman, Kuwait), Egypt's comprehensive PDPL, the leading African regimes (Nigeria, Kenya, Ghana, Morocco, Tunisia), and Turkey's KVKK with its 2024 transfer reform together complete the Middle East and Africa coverage.

Bahrain - Personal Data Protection Law

Bahrain Law No. 30 of 2018 (PDPL); Decision No. 43 of 2022

Jurisdiction	Kingdom of Bahrain
Effective	August 1, 2019
Regulator	Personal Data Protection Authority
Scope	Persons processing personal data within Bahrain or processing data of Bahrain residents from outside Bahrain by automated means or as part of an organized filing system

Applicability

Bahrain's PDPL is one of the more comprehensive Gulf Cooperation Council (GCC) privacy regimes. It includes data subject rights, sensitive data protections, cross-border transfer restrictions, and a notification regime for certain processing categories. AI implications include the need for prior authorization for processing involving direct marketing, biometric identification, and processing for evaluating personality.

Core Obligations

- Lawful basis for processing; consent is principal basis with limited exceptions.
- Notification to or authorization from PDPA for certain processing including processing of sensitive data, biometric identification, evaluation of personality, and others.
- Cross-border transfer to countries without adequate protection requires PDPA approval, controller-processor contract, or other authorized mechanism.
- Data subject rights including access, rectification, deletion, and restriction.
- Designate Data Protection Guardian (Bahrain's DPO equivalent).

Penalties & Enforcement

Fines up to BHD 20,000 (~USD 53,000) per violation; criminal penalties for sensitive data violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	PDPA notification or authorization for sensitive processing creates a Layer 1 gating step before AI deployment.

Qatar - Personal Data Privacy Protection Law

Qatar Law No. 13 of 2016 (PDPPL); Qatar National Cyber Security Strategy 2024

Jurisdiction	State of Qatar
Effective	December 29, 2016
Regulator	Ministry of Communications and Information Technology (MCIT); Compliance and Data Protection Department (CDP)
Scope	Public and private bodies processing personal data of natural persons in Qatar

Applicability

Qatar's PDPPL applies broadly to processing of personal data and includes notice, consent, sensitive data protections, breach notification, and DPO requirements for certain entities. Qatar has signaled intent to develop AI-specific regulation as part of its National AI Strategy. The Qatar Financial Centre maintains a separate GDPR-aligned data protection regime applicable to QFC entities.

Core Obligations

- Provide notice and obtain consent for processing.
- Comply with sensitive data protections requiring authorization and explicit consent.
- Notify CDP of breaches.
- Honor data subject rights.
- Cross-border transfer restrictions.

Penalties & Enforcement

Administrative penalties up to QAR 5M (~USD 1.4M).

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Sensitive data authorization requirements affect Layer 1 collection design for Qatar operations.

Oman - Personal Data Protection Law

Royal Decree No. 6 of 2022

Jurisdiction	Sultanate of Oman
Effective	February 13, 2023 (one-year transition); enforcement from February 2024
Regulator	Ministry of Transport, Communications, and Information Technology
Scope	Persons processing personal data of natural persons in Oman

Applicability

Oman's PDPL is GDPR-aligned with provisions on lawful basis, consent, sensitive data, data subject rights, and cross-border transfers. Implementing regulations issued in 2024 provide operational specificity. AI-specific guidance is anticipated as part of the National AI Strategy.

Core Obligations

- Lawful basis; consent is principal basis with enumerated exceptions.
- Sensitive data requires explicit consent and additional safeguards.
- Data subject rights of access, correction, deletion, withdrawal of consent, and objection.
- Cross-border transfer requires consent, contractual safeguards, or adequacy.
- Mandatory breach notification.
- DPO required for certain processing.

Penalties & Enforcement

Administrative fines up to OMR 100,000 (~USD 260,000); criminal penalties for serious violations.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	PDPL documentation and DPO designation are Layer 5 obligations for in-scope Oman operations.
---------------------------------------	--

Kuwait - Data Privacy Protection Regulation

Kuwait CITRA Resolution No. 26 of 2024 (DPPR)

Jurisdiction	State of Kuwait
Effective	2024 (telecommunications sector); broader applicability evolving
Regulator	Communication and Information Technology Regulatory Authority (CITRA)
Scope	Initially telecommunications and ICT sector; broader application via subsequent regulations

Applicability

Kuwait's DPPR is a sectoral regulation initially covering telecommunications and ICT. It establishes consent, data subject rights, security obligations, and breach notification. Comprehensive privacy legislation has been considered but not enacted as of 2026.

Core Obligations

- Consent for data collection and processing.
- Data subject rights including access and correction.
- Security measures and breach notification.
- Cross-border transfer restrictions for sensitive data.

Penalties & Enforcement

CITRA enforcement; license-related consequences for telecommunications entities.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Telecommunications-sector AI applications must observe DPPR consent obligations.

Egypt - Personal Data Protection Law

Egypt Law No. 151 of 2020 (PDPL)

Jurisdiction	Arab Republic of Egypt
Effective	October 14, 2020 (general); phased enforcement post-implementing regulations
Regulator	Personal Data Protection Center (PDPC, established 2024)
Scope	Persons processing personal data inside or outside Egypt where the data relates to Egyptian residents

Applicability

Egypt's PDPL is the most comprehensive privacy regime in North Africa. It includes consent, data subject rights, sensitive data protections, cross-border transfer restrictions, breach notification, and DPO requirements. The PDPC was established and began operation in 2024; full enforcement is in early stages.

Core Obligations

- Consent (express, written, or electronic) is the principal basis; limited exceptions for legal obligation, vital interests, public interest, scientific research.
- Sensitive data requires explicit consent and additional security.

- Cross-border transfer requires PDPC license; transfers to countries with equivalent protection may be permitted under streamlined procedures.
- Mandatory breach notification within 72 hours.
- DPO required for processing of sensitive data, large-scale processing, and processing involving regular monitoring.

Penalties & Enforcement

Administrative fines up to EGP 5M (~USD 100,000); criminal penalties up to 5 years imprisonment for sensitive data violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	PDPC licensing for cross-border transfer creates Layer 3 architectural constraint for Egypt operations.

Nigeria - Nigeria Data Protection Act (NDPA)

Nigeria Data Protection Act 2023; Nigeria Data Protection Regulation 2019 (NDPR, predecessor)

Jurisdiction	Federal Republic of Nigeria
Effective	NDPA June 14, 2023
Regulator	Nigeria Data Protection Commission (NDPC)
Scope	Public and private bodies processing personal data of data subjects in Nigeria

Applicability

The NDPA modernized Nigerian privacy law from the prior NDPR to a GDPR-aligned statutory framework. The NDPC has been highly active in enforcement, with substantial fines against tech and financial services companies in 2024–2025. AI implications include explicit recognition of automated decision-making rights and the broad definition of personal data including online identifiers.

Core Obligations

- Lawful basis for processing including consent, contract, legal obligation, vital interests, public interest, legitimate interests.
- Data subject rights including access, correction, deletion, portability, objection, and rights related to automated decisioning.
- Sensitive personal data requires explicit consent and additional safeguards.
- Cross-border transfer requires adequacy decision, standard contract, binding corporate rules, or specific authorization.
- Mandatory breach notification within 72 hours.

- DPO required for public bodies, large-scale processing, and certain other categories.
- Data Controllers and Processors of Major Importance: enhanced obligations including registration with NDPC.

Penalties & Enforcement

Administrative fines up to higher of NGN 10M (data controller) or NGN 2M (processor) and 2% of annual gross revenue (data controller) or 1% (processor) for major violations; criminal penalties for serious violations.

Recent Developments (through 2026)

NDPC enforcement against Meta (2024 settlement), Truecaller, and others established Nigeria as the leading African privacy enforcer. The NDPC has signaled AI-specific guidance development through 2026.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	NDPC registration and DPO documentation are Layer 5 obligations; the Commission's active enforcement posture means audit-readiness matters.

Kenya - Data Protection Act

Kenya Data Protection Act, 2019 (Act No. 24 of 2019)

Jurisdiction	Republic of Kenya
Effective	November 25, 2019
Regulator	Office of the Data Protection Commissioner (ODPC)
Scope	Persons processing personal data of data subjects in Kenya or processing where the data subject is in Kenya

Applicability

Kenya's DPA is GDPR-aligned and is among the more developed African privacy regimes. The ODPC has been increasingly active in enforcement, including in fintech and ad-tech. Kenya is a regional AI hub and has been working toward AI-specific regulation.

Core Obligations

- GDPR-style obligations: lawful basis, consent for sensitive data, data subject rights, DPIA for high-risk processing, cross-border transfer restrictions, mandatory breach notification, DPO requirements.
- Registration with ODPC for certain controllers and processors.

Penalties & Enforcement

Administrative fines up to KES 5M or 1% of annual turnover; criminal penalties.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

ODPC registration and DPIA requirements are core Layer 5 obligations for Kenya operations.

Ghana - Data Protection Act

Ghana Data Protection Act, 2012 (Act 843)

Jurisdiction	Republic of Ghana
Effective	October 16, 2012
Regulator	Data Protection Commission
Scope	Public and private bodies processing personal data in Ghana or relating to data subjects in Ghana

Applicability

Ghana's 2012 DPA was an early African privacy law. Eight data protection principles, registration requirements, and notification of breach. The DPA Commission has been active, and reform to align more closely with GDPR is under consideration.

Core Obligations

- Eight principles: accountability, lawfulness, specification of purpose, compatibility of further processing, quality, openness, security, data subject participation.
- Mandatory registration with the Commission.
- Sensitive personal data heightened protections.
- Cross-border transfer restrictions.

Penalties & Enforcement

Administrative fines and criminal penalties.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence

Ghana DPA registration is a Layer 5 baseline for Ghana operations.

Morocco - Data Protection Law

Loi n° 09-08 (2009)

Jurisdiction	Kingdom of Morocco
---------------------	--------------------

Effective	February 18, 2009
Regulator	Commission Nationale de contrôle de la protection des Données à caractère Personnel (CNDP)
Scope	Public and private bodies processing personal data in Morocco

Applicability

Morocco's 09-08 was the first comprehensive Maghreb privacy law and established the CNDP. Morocco has Convention 108+ adequacy with the Council of Europe and pursues alignment with EU adequacy. Reform legislation has been considered through 2026.

Core Obligations

- CNDP authorization or notification for processing including AI-relevant categories.
- Sensitive personal data requires CNDP authorization.
- Cross-border transfer restrictions; CNDP authorization for transfers to non-adequate countries.
- Data subject rights.

Penalties & Enforcement

Administrative penalties; criminal penalties for serious violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	CNDP authorization for AI-relevant sensitive processing creates Layer 1 gating in Morocco operations.

Tunisia - Personal Data Protection Law

Law No. 2004-63 (2004); pending comprehensive reform

Jurisdiction	Republic of Tunisia
Effective	July 27, 2004; reform pending
Regulator	Instance Nationale de Protection des Données Personnelles (INPDP)
Scope	Public and private bodies processing personal data in Tunisia

Applicability

Tunisia was an early Arab privacy regulator. The 2004 law has been overtaken by GDPR-style standards in much of the region; comprehensive reform has been under consideration since 2018, with progress toward Convention 108+ alignment.

Core Obligations

- INPDP authorization or notification for processing.
- Sensitive data restrictions.
- Cross-border transfer restrictions.

Penalties & Enforcement

Administrative and criminal penalties.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	INPDP authorization is the principal Layer 5 administrative obligation for Tunisia operations.

Turkey - Personal Data Protection Law (KVKK)

Law No. 6698 on the Protection of Personal Data (KVKK, 2016)

Jurisdiction	Republic of Türkiye
Effective	April 7, 2016
Regulator	Personal Data Protection Authority (KVKK)
Scope	Persons processing personal data wholly or partly by automated means and processing forming part of a filing system

Applicability

Turkey's KVKK is GDPR-influenced and includes data subject rights, sensitive data protections, breach notification, and registration with VERBİS (data controllers' registry). 2024 amendments substantially modernized the cross-border transfer framework, introducing standard contractual clauses and binding corporate rules. Turkey has signaled AI-specific regulation as part of broader digital reform.

Core Obligations

- Lawful basis (consent or enumerated exceptions); explicit consent for sensitive data.
- Data subject rights of information, access, correction, deletion/destruction, anonymization, objection.
- Cross-border transfer requires explicit consent, adequacy, standard contractual clauses, or binding corporate rules (post-2024 amendment).
- Mandatory breach notification.
- VERBİS registration for data controllers exceeding thresholds.

Penalties & Enforcement

Administrative fines up to TRY 10M+ (varies with annual indexation); criminal penalties for sensitive data violations.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	Post-2024 SCC framework provides Layer 3 transfer mechanism for Turkey operations.
Layer 5 - Audit & Evidence	VERBİS registration is a Layer 5 baseline for Turkey operations exceeding thresholds.

P A R T V I I I

International Frameworks & Voluntary Standards

Soft-law instruments and voluntary standards that shape compliance expectations - including the Council of Europe AI Convention, OECD AI Principles, UNESCO Recommendation, NIST AI RMF, ISO/IEC 42001, MITRE ATLAS, OWASP, and the foundational security and management frameworks AI governance is built on.

AI-Specific International Frameworks

The most-referenced AI-specific international instruments - the Council of Europe Framework Convention on AI, the OECD AI Principles, UNESCO Recommendation, NIST AI RMF + GenAI Profile, ISO/IEC 42001 and 23894, MITRE ATLAS, and OWASP LLM/ML Top 10 - together provide the operational consensus for AI governance practice. Adoption of these frameworks supports defensibility across multiple regulatory regimes simultaneously.

Council of Europe Framework Convention on Artificial Intelligence

Council of Europe Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225, 2024)

Jurisdiction	International (open for signature)
Effective	Opened for signature September 5, 2024; entry into force after 5 ratifications including 3 Council of Europe Member States
Regulator	State parties; Conference of the Parties
Scope	AI activities by parties (states) within the lifecycle of AI systems that have potential to interfere with human rights, democracy, and the rule of law

Applicability

The Framework Convention is the first international treaty on AI. It establishes binding obligations for state parties to ensure that AI activities are consistent with human rights, democracy, and rule of law. While addressed to states, the Convention drives domestic legislation and provides a normative framework for AI governance globally. Initial signatories include the EU, UK, U.S., Israel, and several others.

Core Obligations

- Adopt or maintain measures to ensure AI activities are consistent with obligations to protect human rights, democracy, and rule of law.
- Risk and impact management framework including identification, assessment, prevention, and mitigation of risks.
- Documentation and transparency obligations.
- Effective remedies for persons affected by AI systems.
- Procedural safeguards including notification when interacting with AI and contestation of decisions.

Penalties & Enforcement

Treaty obligations enforced through state parties' domestic measures and Conference of the Parties review.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	Convention drives domestic AI legislation; cross-jurisdictional Layer 5 documentation should reflect Convention principles where state party obligations apply.
---------------------------------------	---

OECD AI Principles and OECD Recommendation on AI

OECD Recommendation of the Council on Artificial Intelligence (2019, updated 2024)

Jurisdiction	International (38 OECD members and adhering non-members)
Effective	2019 (continuously updated)
Regulator	OECD AI Policy Observatory; AI principles are non-binding
Scope	AI policy across adhering jurisdictions; voluntary principles for AI actors

Applicability

The OECD AI Principles are the most widely adopted international AI principles, endorsed by 47+ jurisdictions. The 2024 update addresses generative AI specifically. Five values-based principles (inclusive growth, human-centered values and fairness, transparency and explainability, robustness/security/safety, accountability) and five recommendations to governments establish the international consensus baseline.

Core Obligations

- For governments: invest in AI R&D; foster a digital ecosystem; shape an enabling policy environment; build human capacity; promote international cooperation.
- For AI actors (voluntary): inclusive growth and well-being; human-centered values and fairness; transparency and explainability; robustness, security, and safety; accountability.

Penalties & Enforcement

Non-binding; influence is normative.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	OECD principles function as a Layer 5 cross-jurisdictional benchmark; alignment with the principles supports defensibility in regulatory inquiries globally.
---------------------------------------	--

UNESCO Recommendation on the Ethics of Artificial Intelligence

UNESCO Recommendation on the Ethics of Artificial Intelligence (November 2021)

Jurisdiction	International (UNESCO Member States)
Effective	November 23, 2021
Regulator	UNESCO; State implementation
Scope	AI policy and practice across UNESCO Member States

Applicability

The UNESCO Recommendation is the first global standard-setting instrument on AI ethics, adopted by 193 Member States. It establishes 10 values and principles, 11 policy areas, and a Readiness Assessment Methodology for Member States to evaluate their AI governance maturity.

Core Obligations

- Member States: adopt policies aligned with the Recommendation; conduct Readiness Assessments; implement Ethical Impact Assessments for AI systems.
- Voluntary alignment by AI developers and deployers with values and principles.

Penalties & Enforcement

Non-binding.

STACK LENS - How this law maps to the AI Governance Stack

Layer 5 - Audit & Evidence	Ethical Impact Assessment methodology is a Layer 5 documentation tool that complements other regulatory assessments.
---------------------------------------	--

NIST AI Risk Management Framework (AI RMF 1.0) and Generative AI Profile

NIST AI 100-1 (AI RMF 1.0, January 2023); NIST AI 600-1 (Generative AI Profile, July 2024)

Jurisdiction	United States - Federal (voluntary; binding for federal contractors via incorporation)
Effective	January 26, 2023; July 26, 2024
Regulator	NIST (publishing authority); referenced by federal agencies and incorporated by state laws (e.g., Colorado AI Act)
Scope	Voluntary framework for AI actors; widely adopted as the U.S. baseline AI governance reference

Applicability

The AI RMF is the U.S. baseline AI risk management framework. It is structured around four functions - Govern, Map, Measure, Manage - applicable to the AI lifecycle. The Generative AI Profile (NIST AI 600-1)

provides cross-sectoral generative AI risk guidance. The Colorado AI Act and several other state and federal frameworks create rebuttable presumptions of "reasonable care" for compliance with the AI RMF.

Core Obligations

- Govern: cultivate a risk culture; establish policies, processes, procedures, and practices for trustworthy AI.
- Map: contextualize the AI system, including risks and impacts.
- Measure: analyze, assess, benchmark, and monitor AI risks and benefits.
- Manage: prioritize and act on AI risks based on assessed potential impact.

Penalties & Enforcement

Voluntary; legal significance arises through statutory references and contractual incorporation.

Recent Developments (through 2026)

The 2024 Generative AI Profile is the most-referenced contemporary AI risk artifact; the 2026 RMF revision (in progress) will reflect updated federal priorities.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Map function explicitly includes Layer 1 considerations: data sources, lineage, privacy, IP.
Layer 2 - Model Governance	Measure function operationalizes Layer 2 evaluation including fairness, robustness, accuracy.
Layer 4 - Control & Monitoring	Manage function aligns with Layer 4 monitoring and incident response.
Layer 5 - Audit & Evidence	Govern function and the Profile playbook align with Layer 5 documentation and accountability.

PRACTITIONER NOTE

Adopt the AI RMF as the cross-jurisdictional governance backbone. Map your internal program controls onto the four functions and 76 sub-categories; the resulting matrix supports compliance demonstrations against multiple regulatory regimes simultaneously.

ISO/IEC 42001 - AI Management Systems

ISO/IEC 42001:2023 - Information technology - Artificial intelligence - Management system

Jurisdiction	International (voluntary)
Effective	December 18, 2023
Regulator	Accredited certification bodies

Scope	Organizations providing or using products or services that utilize AI systems
--------------	---

Applicability

ISO/IEC 42001 is the first formal AI management system standard, structured for certification. It establishes requirements for an AI Management System (AIMS) including context, leadership, planning, support, operation, performance evaluation, and improvement. Annex A provides 38 AI-specific control objectives. ISO 42001 certification has emerged as a significant commercial differentiator and is increasingly referenced in procurement.

Core Obligations

- Establish, implement, maintain, and continually improve an AIMS.
- Define AI policy, objectives, and processes.
- Conduct AI risk assessments and AI impact assessments.
- Implement Annex A controls (38 controls covering AI policies, internal organization, AI lifecycle, third-party relationships, customer issues, system design and development, and operational governance).
- Maintain documented information demonstrating AIMS effectiveness.

Penalties & Enforcement

Voluntary; loss of certification on non-conformity. Commercial impact through customer requirements and procurement weighting.

Recent Developments (through 2026)

Major AI providers and platforms have pursued 2024–2026 certification; ISO 42001 is becoming the practical equivalent of SOC 2 for AI governance assurance.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 5 - Audit & Evidence	AIMS is structurally a Layer 5 governance system; the Annex A controls span all five Stack layers and provide a certifiable evidence framework.

PRACTITIONER NOTE

For B2B AI providers, ISO 42001 certification is rapidly becoming a procurement requirement. Plan certification in concert with SOC 2 and ISO 27001 - the management system structures are compatible and audit windows can be combined.

ISO/IEC 23894 - AI Risk Management Guidance

ISO/IEC 23894:2023 - Information technology - Artificial intelligence - Guidance on risk management

Jurisdiction	International (voluntary)
Effective	February 6, 2023
Regulator	N/A (voluntary)
Scope	Organizations developing, deploying, or using AI systems

Applicability

ISO/IEC 23894 provides AI-specific guidance for implementing ISO 31000 risk management principles. It addresses AI-specific risk sources including data quality, model bias, opacity, and emergent behaviors. ISO 23894 is referenced in the EU AI Act drafting and is widely cited in implementation guidance.

Core Obligations

- Apply ISO 31000 principles in the AI context.
- Identify AI-specific risk sources.
- Implement risk treatment, monitoring, review, and communication aligned with the AI lifecycle.

Penalties & Enforcement

Voluntary.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	AI-specific risk sources catalog is a useful Layer 2 input; many organizations integrate ISO 23894 risk taxonomies into model risk frameworks.

MITRE ATLAS - Adversarial Threat Landscape for AI Systems

MITRE ATLAS (continuously updated)

Jurisdiction	International (voluntary)
Effective	2020 (continuously updated)
Regulator	N/A (community-maintained)
Scope	Adversarial machine learning threat modeling

Applicability

MITRE ATLAS is the leading adversarial ML knowledge base, structured as a STIX-compatible threat matrix. It catalogs tactics, techniques, and procedures (TTPs) used to attack AI systems including reconnaissance, resource development, initial access, ML model access, execution, persistence, defense evasion, discovery, collection, ML attack staging, exfiltration, and impact.

Core Obligations

- Voluntary use as a threat-modeling reference.

Penalties & Enforcement

N/A.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	ATLAS techniques inform Layer 4 detection rules and security monitoring; ATT&CK-equivalent for AI deployments.

OWASP LLM Top 10 and ML Top 10

OWASP Top 10 for LLM Applications (continuously updated); OWASP ML Security Top 10

Jurisdiction	International (voluntary)
Effective	2023 (LLM); 2023 (ML)
Regulator	N/A (community-maintained)
Scope	AI security risk references

Applicability

The OWASP LLM Top 10 enumerates the leading risks for LLM applications: prompt injection, insecure output handling, training data poisoning, model denial of service, supply chain vulnerabilities, sensitive information disclosure, insecure plugin design, excessive agency, overreliance, model theft. The ML Top 10 covers traditional ML system risks. These are the de facto AI security benchmarks for application developers.

Core Obligations

- Voluntary use as a development and security reference.

Penalties & Enforcement

N/A.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 2 - Model Governance	Training data poisoning and model theft are Layer 2 considerations.
Layer 3 - System Integration	Insecure output handling, plugin design, and excessive agency are Layer 3 integration vulnerabilities.
Layer 4 - Control & Monitoring	Prompt injection detection, sensitive information disclosure prevention, and model denial of service mitigation are Layer 4 operational controls.

Foundational Security and Management System Standards

AI governance programs are built on top of foundational security and privacy management frameworks: ISO 27001 (information security), ISO 27701 (privacy), SOC 2 (service organization controls), NIST SP 800-53 (federal baseline), the CSA Cloud Controls Matrix, IEEE 7000 series ethics standards, CIS Controls, and PCI DSS where payment data is involved. ENISA materials provide European baseline expectations. These frameworks predate AI-specific governance but provide the substrate on which AI controls are built.

ISO/IEC 27001 - Information Security Management Systems

ISO/IEC 27001:2022 - Information security, cybersecurity and privacy protection - Information security management systems - Requirements

Jurisdiction	International (voluntary)
Effective	October 25, 2022 (current edition; predecessor 2013)
Regulator	Accredited certification bodies
Scope	Organizations of any size, type, or sector establishing an Information Security Management System (ISMS)

Applicability

ISO 27001 is the foundational international information security standard and the basis for AI-relevant security frameworks (ISO 27701 for privacy, ISO 42001 for AI). The 2022 revision restructured Annex A controls into 93 controls in four themes (organizational, people, physical, technological) and explicitly addressed cloud security, threat intelligence, and secure development. Practitioners building AI security programs typically establish ISO 27001 as the foundation, with ISO 42001 providing the AI-specific overlay.

Core Obligations

- Establish, implement, maintain, and continually improve an ISMS.
- Conduct information security risk assessments and risk treatment.
- Implement Annex A controls (93 controls in 2022 edition) or document justified exclusions.
- Maintain documented information demonstrating ISMS effectiveness.
- Internal audit and management review.
- Continual improvement.

Penalties & Enforcement

Voluntary; loss of certification on non-conformity. Commercial impact through procurement requirements.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	Annex A controls map directly to Layer 4 operational security; AI-specific extensions through ISO 42001 layer cleanly on top.
Layer 5 - Audit & Evidence	ISMS documentation is foundational Layer 5 evidence; audit reports and corrective action records support compliance demonstrations across multiple regulatory regimes.

PRACTITIONER NOTE
 Build ISO 27001 as the foundation, ISO 27701 as the privacy extension, and ISO 42001 as the AI extension. The shared management-system structure permits combined audits and consolidated documentation.

ISO/IEC 27701 - Privacy Information Management Systems

ISO/IEC 27701:2019 - Security techniques - Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management - Requirements and guidelines

Jurisdiction	International (voluntary)
Effective	August 6, 2019; revision in development
Regulator	Accredited certification bodies
Scope	PII controllers and PII processors implementing a Privacy Information Management System (PIMS) as an extension to ISO 27001

Applicability

ISO 27701 extends ISO 27001 with privacy-specific controls and guidance for both controllers and processors. It is mapped to GDPR Articles and other privacy frameworks, supporting demonstration of accountability across multiple regimes. The 2019 edition is being revised to reflect the post-2022 ISO 27001 restructure and developments in privacy regulation. ISO 27701 certification is the most widely recognized privacy management system certification globally and is referenced in the Tennessee TIPA affirmative defense.

Core Obligations

- Implement PIMS as extension to ISO 27001 ISMS.
- For PII controllers: privacy notices, consent management, data subject rights, DPIAs, sub-processor management, breach response.

- For PII processors: processor-controller relationships, sub-processor cascading, data minimization in service delivery.
- Maintain documentation supporting privacy regulatory compliance.

Penalties & Enforcement

Voluntary; loss of certification on non-conformity.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	PII handling controls in Annex A map to Layer 1 governance for AI training data.
Layer 5 - Audit & Evidence	PIMS documentation supports Tennessee TIPA affirmative defense and provides cross-jurisdictional Layer 5 evidence.

SOC 2 - Service Organization Controls

AICPA Trust Services Criteria; SSAE 18 / SOC 2 reporting

Jurisdiction	International (audit framework)
Effective	Continuous; current Trust Services Criteria 2017 (updated 2022)
Regulator	AICPA-licensed CPA firms
Scope	Service organizations providing services to user entities

Applicability

SOC 2 reports are the principal U.S. trust framework for service organizations including SaaS, cloud, and AI providers. Type I reports describe controls at a point in time; Type II reports include effectiveness testing over a period (typically 6–12 months). The five Trust Services Criteria are Security, Availability, Processing Integrity, Confidentiality, and Privacy. AI providers selling to enterprise customers face SOC 2 procurement expectations as standard practice. AI-specific SOC 2+ engagements are emerging.

Core Obligations

- Implement controls satisfying selected Trust Services Criteria (Security is mandatory; others optional).
- Document control descriptions, control tests, and exceptions.
- For Type II: maintain controls operating effectively over the audit period.
- Manage exceptions and remediation.
- Annual recertification typical.

Penalties & Enforcement

Voluntary; commercial impact through procurement and customer contractual requirements.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	SOC 2 controls map closely to Layer 4 operational security; AI-specific testing addresses model lifecycle controls.
Layer 5 - Audit & Evidence	SOC 2 Type II reports are core Layer 5 deliverables shared with enterprise customers; the "complementary user entity controls" section assigns shared responsibility.

NIST SP 800-53 - Security and Privacy Controls

NIST SP 800-53 Rev. 5 (Security and Privacy Controls for Information Systems and Organizations)

Jurisdiction	United States - Federal (mandatory for federal); voluntary internationally
Effective	September 23, 2020 (Rev. 5; continuous updates)
Regulator	NIST (publishing authority); referenced by FedRAMP, FISMA, DOD, and others
Scope	Federal information systems and contractor systems processing federal information; widely adopted as voluntary baseline

Applicability

NIST SP 800-53 is the foundational U.S. federal control catalog with 1,189 controls and control enhancements organized into 20 control families. FedRAMP authorization, FISMA compliance, DOD CMMC, and most federal cybersecurity frameworks reference 800-53. AI-specific overlays are in development; the 2024 NIST AI 100-1 (AI RMF) and 2024 NIST IR 8470 (AI security overlay considerations) provide AI-relevant guidance. AI vendors selling to federal customers face direct 800-53 application; commercial vendors increasingly adopt as baseline.

Core Obligations

- Categorize information systems per FIPS 199.
- Select baseline (Low, Moderate, High) and tailor based on risk assessment.
- Implement selected controls; document implementation.
- Assess control effectiveness.
- Authorize the system to operate (ATO).
- Continuously monitor.

Penalties & Enforcement

Federal: loss of ATO; FAR/DFARS contract performance consequences. Commercial: voluntary.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	Privacy control family (PT) and AI-relevant data controls inform Layer 1 governance for federal AI workloads.
Layer 4 - Control & Monitoring	AC, AU, CA, CM, IR, SC, SI control families align with Layer 4 operational practice.
Layer 5 - Audit & Evidence	System Security Plans, ATO packages, and continuous monitoring evidence are core Layer 5 federal artifacts.

Cloud Security Alliance (CSA) Cloud Controls Matrix and AI Controls

CSA Cloud Controls Matrix v4 (2021); STAR Registry; CSA AI Controls Matrix (development)

Jurisdiction	International (voluntary)
Effective	CCM continuously updated
Regulator	CSA
Scope	Cloud service providers and enterprise cloud users

Applicability

The CSA CCM provides a cloud-specific control framework with 197 controls in 17 domains, mapped to ISO 27001, SOC 2, PCI DSS, FedRAMP, and other frameworks. The CSA STAR Registry provides voluntary public attestations from cloud providers. CSA's AI Controls Matrix is in development through 2026 and will provide the AI-specific overlay.

Core Obligations

- Voluntary control implementation; STAR registry submission optional.
- Self-assessment Level 1 (CAIQ); third-party attestation Level 2; continuous monitoring Level 3.

Penalties & Enforcement

N/A (voluntary).

STACK LENS - How this law maps to the AI Governance Stack	
Layer 3 - System Integration	CCM provides Layer 3 cloud architecture guidance widely referenced in vendor due diligence.

IEEE 7000 Series - Ethics-Driven Standards

IEEE 7000-2021 (Model Process for Addressing Ethical Concerns During System Design); IEEE 7001 (Transparency); IEEE 7002 (Privacy); IEEE 7010 (Wellbeing); IEEE 7014 (Empathic Systems); others in development

Jurisdiction	International (voluntary)
Effective	IEEE 7000 published 2021; continuing series
Regulator	IEEE Standards Association
Scope	Engineering practice for AI and autonomous systems

Applicability

The IEEE 7000 series operationalizes ethical considerations in technical engineering practice. IEEE 7000-2021 is the first international standard for engineering ethics processes; subsequent standards address specific ethical concerns. The series provides engineering-level operational guidance complementing higher-level frameworks like ISO 42001 and the EU AI Act.

Core Obligations

- Voluntary adoption in engineering practice; some procurement specifications reference specific 7000-series standards.

Penalties & Enforcement

N/A (voluntary).

STACK LENS - How this law maps to the AI Governance Stack

Layer 2 - Model Governance	7001 transparency standards inform Layer 2 explainability practice; 7002 privacy standards inform Layer 1 design.
-----------------------------------	---

CIS Controls and CIS Benchmarks

CIS Critical Security Controls v8 (2021); CIS Benchmarks (continuous publication)

Jurisdiction	International (voluntary)
Effective	Continuous
Regulator	Center for Internet Security
Scope	Organizational cybersecurity controls; system-specific configuration benchmarks

Applicability

CIS Controls v8 provides 18 prioritized cybersecurity controls and 153 sub-controls organized into Implementation Groups (IG1, IG2, IG3) for organizations of varying size and risk. CIS Benchmarks provide specific configuration guidance for hundreds of products including operating systems, cloud platforms, containers, and databases. Both are widely adopted as practical operational guidance complementing higher-level frameworks.

Core Obligations

- Voluntary adoption; many organizations use CIS as the operational implementation of ISO/NIST higher-level frameworks.

Penalties & Enforcement

N/A.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	CIS Benchmarks provide Layer 4 configuration baselines for AI infrastructure; container, Kubernetes, and cloud benchmarks are particularly relevant.

ENISA AI Threat Landscape and Multilateral AI Cybersecurity Frameworks

ENISA Threat Landscape for Artificial Intelligence (2024); Multilateral CISA / NCSC-UK / ENISA AI security guidance (2024)

Jurisdiction	International (voluntary; high practical authority)
Effective	ENISA AI threat landscape continuous; multilateral guidance 2024
Regulator	ENISA; national cybersecurity authorities
Scope	AI system developers, deployers, and operators

Applicability

ENISA's AI Threat Landscape report and the joint multilateral guidance establish the European baseline AI cybersecurity expectation. Combined with NIST AI RMF, MITRE ATLAS, and OWASP LLM/ML Top 10, the multilateral guidance provides the operational consensus for AI security practice. Practitioners advising on EU operations should treat ENISA materials as effectively required references.

Core Obligations

- Voluntary; influence is normative.

Penalties & Enforcement

N/A.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 4 - Control & Monitoring	ENISA threat catalog informs Layer 4 detection rules; multilateral guidance provides international baseline expectations.

PCI DSS and AI in Payment Card Processing

PCI Data Security Standard v4.0 (March 2022; v4.0.1 2024)

Jurisdiction	International (contractual; payment card networks)
Effective	v4.0 effective March 31, 2024 (with phased requirements through 2025)
Regulator	PCI Security Standards Council; payment brands; acquiring banks
Scope	Entities that store, process, or transmit cardholder data; certain entities affecting cardholder data security

Applicability

PCI DSS v4.0 introduced significant changes including more flexible compliance approaches (Customized Approach), targeted risk analysis, expanded MFA requirements, and enhanced authentication. AI implications include: (1) AI-driven fraud detection systems often have direct access to cardholder data; (2) AI vendors processing cardholder data are themselves in scope; (3) AI use in cardholder data environment must satisfy applicable requirements; and (4) shadow AI use of cardholder data creates compliance failures.

Core Obligations

- Implement 12 high-level requirements with 320+ sub-requirements at Level 1 (largest merchants).
- For AI in cardholder data environment: apply applicable controls including access management, encryption, change management, and logging.
- Quarterly external vulnerability scans; annual penetration testing; annual on-site assessment for Level 1.

Penalties & Enforcement

Payment card brand fines (significant for non-compliance with breach); loss of payment card processing privileges; contractual penalties from acquiring banks.

STACK LENS - How this law maps to the AI Governance Stack	
Layer 1 - Data Governance	AI training data containing cardholder data is in PCI scope; data inventories must include AI-relevant payment data flows.
Layer 4 - Control & Monitoring	PCI DSS controls map to Layer 4 operational security for payment-relevant AI systems.

P A R T I X

Quick-Reference Tables

At-a-glance comparisons across jurisdictions and frameworks for the obligations practitioners are asked about most often.

Maximum Civil Penalties - Comparative View

Headline penalty exposure for the leading AI, privacy, and cybersecurity regimes. Penalty calculations vary; figures shown reflect the highest available statutory cap.

Regime	Maximum Penalty
EU AI Act - Article 5 prohibitions	€35M or 7% worldwide turnover
EU AI Act - most other violations	€15M or 3% worldwide turnover
GDPR / UK GDPR - Tier 2	€20M / £17.5M or 4% worldwide turnover
GDPR / UK GDPR - Tier 1	€10M / £8.7M or 2% worldwide turnover
EU Digital Services Act	6% worldwide annual turnover
EU Digital Markets Act	10% worldwide annual turnover (20% repeat)
EU Cyber Resilience Act	€15M or 2.5% worldwide turnover
NIS2 - essential entities	€10M or 2% worldwide turnover
DORA - critical ICT third parties	1% average daily worldwide turnover (per day)
China PIPL	RMB 50M or 5% prior-year turnover
Brazil LGPD	BRL 50M per infraction or 2% Brazil revenue
Korea PIPA	3% related sales (post-2023)
Singapore PDPA	10% Singapore turnover (large entities)
Australia Privacy Act (post-2024)	AUD 50M, 3× benefit, or 30% adjusted turnover
India DPDP Act	INR 250 crore (~USD 30M) per instance
Saudi Arabia PDPL	SAR 5M per violation
Nigeria NDPA	Higher of NGN 10M or 2% annual gross revenue
Chile Law 21.719 (eff. 2026)	4% annual turnover
California CCPA/CPRA	\$2,500 / \$7,500 per violation
Colorado AI Act	\$20,000 per violation
Texas TDPSA	\$7,500 per violation
Texas CUBI	\$25,000 per violation
Illinois BIPA (2024 amendment)	\$1,000–\$5,000 per individual per type
Maryland MODPA	\$10,000 / \$25,000 per violation
Florida Digital Bill of Rights	\$50,000 per violation; treble for minors
NY DFS Part 500	Per-violation; settlements \$1.5M–\$40M+
HIPAA	\$137–\$2,067,813 per violation per year
FCRA - willful	\$100–\$1,000 per violation + actual + punitive
VPPA	\$2,500 per violation + actual + punitive
TCPA	\$500–\$1,500 per violation (no aggregate cap)

Regime	Maximum Penalty
FTC Act § 5 (consent order violations)	\$53,088 per violation per day
Title VII (private)	\$50,000–\$300,000 capped by employer size
Copyright Act (willful)	\$150,000 per work + actual + fees
Defend Trade Secrets Act	Damages + 2x exemplary + fees

Breach Notification Timelines

Regulator notification windows for the leading regimes. Affected-individual notification windows are typically distinct and often more flexible.

Regime	Regulator notification
GDPR / UK GDPR	72 hours from awareness
NIS2	24h early warning, 72h notification, 1 month final
EU Cyber Resilience Act	24h actively exploited vulnerability; 72h severe incident
DORA	Initial / intermediate / final classification reports
NY DFS Part 500	72 hours from determination
HIPAA Breach Notification Rule	60 days (individuals); 60 days (HHS, 500+); annually (HHS, <500)
GLBA (FTC Safeguards Rule)	30 days from discovery (500+ consumers)
CISA / CIRCIA (when finalized)	72h cyber incident; 24h ransomware payment
TSA Pipeline / Rail	24 hours to CISA
Brazil LGPD	Reasonable time; ANPD guidance specifies promptness
Egypt PDPL	72 hours
Indonesia PDP Law	72 hours
Vietnam PDPD	72 hours
India DPDP Act	72 hours (per draft Rules)
Korea PIPA	72 hours
Australia NDB scheme	As soon as practicable
China PIPL	Promptly; specifics by sector
Most U.S. state breach laws	30–60 days varying by state
NY SHIELD Act	Without unreasonable delay; AG/police for 5,000+

Automated Decision-Making - Cross-Jurisdictional Comparison

How the leading regimes handle solely or substantially automated decisions affecting individuals.

Regime	Trigger	Core Right
GDPR Art. 22	Solely automated; legal/similar effect	Right not to be subject; explanation; human review
UK GDPR (DUAA reform)	Solely automated; legal/similar effect	Lawful bases broadened; safeguards retained
CCPA ADMT regs	ADMT for significant decisions	Pre-use notice; opt-out; alternative process
Colorado AI Act	High-risk consequential decision	Notice; explanation; appeal with human review
Minnesota MNCDPA	Profiling for legal/similar decisions	Question result; review data; reassessment
Korea PIPA (2023 amendment)	Solely automated; significant effect	Explanation; refusal; human review
Quebec Law 25	Automated processing; significant effect	Consent; explanation; human review
Brazil LGPD Art. 20	Solely automated; affecting interests	Review by natural person; explanation
China PIPL Art. 24	Automated decisioning; significant effect	Transparency; refusal; explanation
Australia Privacy Act (Tranche 2)	Automated decisions used in significant ways	Disclosure of kinds of decisions and information
Indonesia PDP Law	Automated decisions	Rights including objection
Nigeria NDPA	Automated decisioning	Explanation and human review

Stack Layer Cross-Reference - Where to Find Each Layer's Heaviest Hitters

When practitioners are asked "which laws bite hardest at this layer?" - start here, then drill into the entries above.

Stack Layer	Highest-impact regimes
Layer 1 - Data Governance	GDPR; EU AI Act Art. 10; CPRA; Colorado biometrics/neural; Texas CUBI; Illinois BIPA; HIPAA; COPPA 2025; CA AB 2013; China PIPL; Maryland MODPA strict minimization; EU CDSM Art 4 TDM; Japan Art 30-4; trade secret programs
Layer 2 - Model Governance	EU AI Act Arts. 9–15; Colorado AI Act; CFPB Circulars; SR 11-7; NAIC Bulletin; FDA SaMD; ISO 42001; NIST AI RMF; Title VII / ADEA / FHA disparate impact; California SB 53 frontier; Section 1557
Layer 3 - System Integration	EU AI Act provider-deployer chain; DORA; HIPAA BAAs; CCPA service-provider/contractor terms; FERPA "direct control"; SCC and EU-US DPF; eIDAS 2.0 wallet integration; CRA supply chain; CSA CCM
Layer 4 - Control & Monitoring	NY DFS Part 500; NIS2 implementing acts; DORA; CCPA opt-out signals; EU AI Act post-market monitoring; TCPA; Section 1557; CISA AI guidance; NERC CIP; CIS Controls

Stack Layer	Highest-impact regimes
Layer 5 - Audit & Evidence	EU AI Act Annex IV; DPIAs (GDPR, LGPD, PIPA, PIPL); SR 11-7 model documentation; ISO 42001 AIMS; CCPA risk assessments; NIST AI RMF Govern function; Tennessee TIPA NIST PF safe harbor; SOC 2 Type II; FedRAMP ATO packages; AIBOM emerging practice

P A R T X

Alphabetical Index

A reverse lookup - every regime, statute, and framework covered in this Field Guide, with the Part where its full entry appears.

The following index lists every entry in the Field Guide alphabetically by short title or commonly used name. Multiple sub-entries (such as the GDPR appearing in the EU Part and being referenced in the UK Part for UK GDPR) are listed under each name.

A

Age Discrimination in Employment Act (ADEA) and AI.....	Part I (Civil Rights)
Colorado AI Act (CAIA, SB 24-205).....	Part III (State AI)
EU AI Act (Regulation (EU) 2024/1689).....	Part V (EU Core)
EU AI Act Code of Practice for General-Purpose AI.....	Part V (EU Additional)
California AI in Healthcare Disclosure Act (AB 3030).....	Part III (State AI)
EU AI Liability Directive (Pending) and Revised Product Liability Directive.....	Part V (EU Additional)
California AI Provenance Act (SB 942) and California AI Transparency Act.....	Part III (State AI)
Americans with Disabilities Act (ADA) Title III and AI Accessibility.....	Part I (Civil Rights)
Argentina - Personal Data Protection Law and Pending Reform.....	Part VI (UK & Americas)
Australia - Privacy Act 1988 and AI Ethics Framework.....	Part VII (APAC & MEA Core)

B

Bahrain - Personal Data Protection Law.....	Part VII (MENA & Africa)
New York Biometric and Student Data Laws.....	Part II (State Other)
Brazil - Lei Geral de Proteção de Dados (LGPD) and Marco Legal da Inteligência Artificial.....	Part VI (UK & Americas)

C

CAN-SPAM Act and Commercial Electronic Communications.....	Part I (Communications)
Canada - PIPEDA and the Pending Consumer Privacy Protection Act.....	Part VI (UK & Americas)
Texas Capture or Use of Biometric Identifier Act (CUBI).....	Part II (State Privacy)
CFPB Adverse Action and AI - Circulars 2022-03 and 2023-03.....	Part I (Sector-Specific)
The Children's Online Privacy Protection Act (COPPA).....	Part I (Federal)
Chile - Personal Data Protection Law.....	Part VI (Latin America)
China - Generative AI Measures, Deep Synthesis Provisions, and Algorithmic Recommendation Provisions.....	Part VII (APAC & MEA Core)
China - Personal Information Protection Law (PIPL).....	Part VII (APAC & MEA Core)
CIS Controls and CIS Benchmarks.....	Part VIII (Security Standards)
New York City Local Law 144 - Automated Employment Decision Tools (AEDT).....	Part III (State AI)
CLOUD Act and Cross-Border Government Access.....	Part I (Communications)
Cloud Security Alliance (CSA) Cloud Controls Matrix and AI Controls.....	Part VIII (Security Standards)
Colombia - Personal Data Protection Law.....	Part VI (Latin America)
Communications Decency Act Section 230.....	Part I (Communications)
The Computer Fraud and Abuse Act (CFAA).....	Part I (Federal)
Connecticut Data Privacy Act (CTDPA).....	Part II (State Privacy)
Connecticut SB 2 - AI Consequential Decision Bill (Pending) and Other Connecticut AI Initiatives.....	Part III (State AI - Additional)
California Consumer Privacy Act (CCPA) and California Privacy Rights Act (CPRA).....	Part II (State Privacy)
EU Copyright Directive Article 4 (Text and Data Mining Exception).....	Part IV (IP & AI)
U.S. Copyright Office Guidance on AI-Generated Works.....	Part IV (IP & AI)
Costa Rica - Protection of the Person against the Treatment of Personal Data.....	Part VI (Latin America)

Council of Europe Framework Convention on Artificial Intelligence *Part VIII (AI Standards)*
 EU Cyber Resilience Act (CRA) *Part V (EU Additional)*
 EU Cybersecurity Act and ENISA *Part V (EU Additional)*
 Cybersecurity and Infrastructure Security Agency (CISA) Act and Critical Infrastructure AI Guidance . *Part I (Critical Infrastructure)*

D

EU Data Act (Regulation (EU) 2023/2854) *Part V (EU Core)*
 EU Data Governance Act *Part V (EU Additional)*
 Texas Data Privacy and Security Act (TDPSA) *Part II (State Privacy)*
 State Deepfake and AI-Generated Content Laws (Survey) *Part III (State AI)*
 Delaware Personal Data Privacy Act *Part II (State Privacy - Second Wave)*
 Department of Defense Responsible AI Strategy and Procurement *Part I (Critical Infrastructure)*
 EU Digital Markets Act (DMA) *Part V (EU Core)*
 EU Digital Operational Resilience Act (DORA) *Part V (EU Core)*
 EU Digital Services Act (DSA) *Part V (EU Core)*
 Dominican Republic - Personal Data Protection Law *Part VI (Latin America)*
 Driver's Privacy Protection Act (DPPA) and Other Federal Sectoral Privacy Statutes *Part I (Communications)*

E

Ecuador - Organic Law on Protection of Personal Data *Part VI (Latin America)*
 EEOC AI in Employment Guidance and ADA Title I *Part I (Sector-Specific)*
 Egypt - Personal Data Protection Law *Part VII (MENA & Africa)*
 eIDAS 2.0 and the European Digital Identity Wallet *Part V (EU Additional)*
 State Election AI Deepfake Laws (Survey) *Part III (State AI - Additional)*
 Electronic Communications Privacy Act (ECPA) - Comprehensive *Part I (Communications)*
 Federal Energy Regulatory Commission (FERC) Order 901 and Critical Infrastructure *Part I (Critical Infrastructure)*
 ENISA AI Threat Landscape and Multilateral AI Cybersecurity Frameworks *Part VIII (Security Standards)*
 ePrivacy Directive (Cookie Law) *Part V (EU Core)*
 The Equal Credit Opportunity Act (ECOA) and Regulation B *Part I (Federal)*
 EU–U.S. Data Privacy Framework (DPF) and Successor Mechanisms *Part I (Communications)*
 Executive Order 14179 and the Federal AI Action Plan *Part I (Federal)*

F

FAA AI in Aviation and Aviation Cybersecurity *Part I (Critical Infrastructure)*
 The Fair Credit Reporting Act (FCRA) *Part I (Federal)*
 Fair Housing Act (FHA) and HUD AI Guidance *Part I (Civil Rights)*
 The Family Educational Rights and Privacy Act (FERPA) *Part I (Federal)*
 FDA AI/ML-Based Software as a Medical Device (SaMD) Framework *Part I (Sector-Specific)*
 The Federal Information Security Modernization Act (FISMA) and FedRAMP *Part I (Federal)*
 The Federal Trade Commission Act, Section 5 (UDAP) *Part I (Federal)*
 Florida Digital Bill of Rights *Part II (State Privacy - Second Wave)*
 Foreign Intelligence Surveillance Act (FISA) Section 702 and U.S. Surveillance Reform *Part I (Communications)*

G

General Data Protection Regulation (GDPR, Regulation (EU) 2016/679) *Part V (EU Core)*

California Generative AI Accountability Act (AB 2885) and Related California AI Statutes *Part III (State AI - Additional)*
 California Generative AI Training Data Transparency Act (AB 2013).....*Part III (State AI)*
 Genetic Information Nondiscrimination Act (GINA).....*Part I (Civil Rights)*
 Ghana - Data Protection Act..... *Part VII (MENA & Africa)*
 The Gramm-Leach-Bliley Act (GLBA) and the FTC Safeguards Rule..... *Part I (Federal)*

H

EU Health Data Space (EHDS) Regulation..... *Part V (EU Additional)*
 The Health Insurance Portability and Accountability Act (HIPAA) and the HITECH Act *Part I (Federal)*
 Hong Kong - Personal Data (Privacy) Ordinance (PDPO)*Part VII (APAC Additional)*

I

IEEE 7000 Series - Ethics-Driven Standards *Part VIII (Security Standards)*
 Illinois Artificial Intelligence Video Interview Act*Part III (State AI)*
 Illinois Biometric Information Privacy Act (BIPA) *Part II (State Privacy)*
 Illinois HB 3773 - AI in Employment Decisions*Part III (State AI)*
 India - Digital Personal Data Protection Act (DPDP Act)..... *Part VII (APAC & MEA Core)*
 Indiana Consumer Data Protection Act (INCDPA) *Part II (State Privacy - Second Wave)*
 Indonesia - Personal Data Protection Law (PDP Law).....*Part VII (APAC Additional)*
 Iowa Consumer Data Protection Act (ICDPA)*Part II (State Privacy - Second Wave)*
 ISO/IEC 23894 - AI Risk Management Guidance.....*Part VIII (AI Standards)*
 ISO/IEC 27001 - Information Security Management Systems *Part VIII (Security Standards)*
 ISO/IEC 27701 - Privacy Information Management Systems..... *Part VIII (Security Standards)*
 ISO/IEC 42001 - AI Management Systems*Part VIII (AI Standards)*
 Israel - Privacy Protection Law (PPL) and AI Policy..... *Part VII (APAC & MEA Core)*

J

Japan - Act on the Protection of Personal Information (APPI) *Part VII (APAC & MEA Core)*
 Japan Copyright Act Article 30-4 - TDM Exception..... *Part IV (IP & AI)*

K

Kentucky Consumer Data Protection Act *Part II (State Privacy - Second Wave)*
 Kenya - Data Protection Act *Part VII (MENA & Africa)*
 Kuwait - Data Privacy Protection Regulation..... *Part VII (MENA & Africa)*

L

New York Local Law 144 - AEDT Detailed Requirements and NY State AI Bills *Part III (State AI - Additional)*

M

Malaysia - Personal Data Protection Act (PDPA)*Part VII (APAC Additional)*
 Maryland Online Data Privacy Act (MODPA).....*Part II (State Privacy - Second Wave)*
 Massachusetts 201 CMR 17 - Standards for the Protection of Personal Information.....*Part II (State Other)*
 Massachusetts Information Privacy and Security Act (Pending) and Massachusetts AI Initiatives..... *Part III (State AI - Additional)*
 Mexico - Federal Law on Protection of Personal Data Held by Private Parties *Part VI (UK & Americas)*

Minnesota Consumer Data Privacy Act (MNCDPA) *Part II (State Privacy - Second Wave)*
 MITRE ATLAS - Adversarial Threat Landscape for AI Systems *Part VIII (AI Standards)*
 Montana Consumer Data Privacy Act (MCDPA) *Part II (State Privacy - Second Wave)*
 Morocco - Data Protection Law *Part VII (MENA & Africa)*

N

NAIC Model Bulletin on Use of Artificial Intelligence Systems *Part I (Sector-Specific)*
 National Labor Relations Act (NLRA) and AI in the Workplace *Part I (Civil Rights)*
 Nebraska Data Privacy Act *Part II (State Privacy - Second Wave)*
 NERC Critical Infrastructure Protection (CIP) Standards *Part I (Critical Infrastructure)*
 EU Network and Information Security (NIS2) Sectoral Implementing Acts *Part V (EU Additional)*
 New Hampshire Data Privacy Act *Part II (State Privacy - Second Wave)*
 New Jersey Data Privacy Act *Part II (State Privacy - Second Wave)*
 New Zealand - Privacy Act 2020 *Part VII (APAC Additional)*
 Nigeria - Nigeria Data Protection Act (NDPA) *Part VII (MENA & Africa)*
 EU NIS2 Directive *Part V (EU Core)*
 NIST AI Risk Management Framework (AI RMF 1.0) and Generative AI Profile *Part VIII (AI Standards)*
 NIST SP 800-53 - Security and Privacy Controls *Part VIII (Security Standards)*
 State Non-Consensual Intimate Imagery (NCII) and Synthetic Sexual Content Laws *Part III (State AI - Additional)*
 Nuclear Regulatory Commission (NRC) Cybersecurity and AI Frameworks *Part I (Critical Infrastructure)*
 NYT v. OpenAI and the Generative AI Training Litigation Landscape *Part IV (IP & AI)*

O

OECD AI Principles and OECD Recommendation on AI *Part VIII (AI Standards)*
 Office of Federal Contract Compliance Programs (OFCCP) and AI in Federal Contractor Employment *Part I (Civil Rights)*
 Oklahoma AI in Healthcare Act and Related State Sectoral AI Laws *Part III (State AI - Additional)*
 Oman - Personal Data Protection Law *Part VII (MENA & Africa)*
 Open Source AI Licensing - OSI Definition, OpenRAIL, MIT/Apache, and Custom Licenses *Part IV (IP & AI)*
 Oregon Consumer Privacy Act (OCPA) *Part II (State Privacy)*
 OWASP LLM Top 10 and ML Top 10 *Part VIII (AI Standards)*

P

Panama - Personal Data Protection Law *Part VI (Latin America)*
 PCI DSS and AI in Payment Card Processing *Part VIII (Security Standards)*
 Peru - Personal Data Protection Law *Part VI (Latin America)*
 Philippines - Data Privacy Act *Part VII (APAC Additional)*
 Colorado Privacy Act (CPA) *Part II (State Privacy)*
 The Privacy Act of 1974 *Part I (Federal)*

Q

Qatar - Personal Data Privacy Protection Law *Part VII (MENA & Africa)*

R

Federal Reserve SR 11-7 - Model Risk Management *Part I (Sector-Specific)*

Texas Responsible Artificial Intelligence Governance Act (TRAIGA).....Part III (State AI)
 Rhode Island Data Transparency and Privacy Protection Act.....Part II (State Privacy - Second Wave)
 Right of Publicity, NO FAKES Act, and AI Voice/Likeness Protection.....Part IV (IP & AI)

S

Saudi Arabia - Personal Data Protection Law (PDPL).....Part VII (APAC & MEA Core)
 California SB 1047 - Safe and Secure Innovation for Frontier AI Models Act (Vetoed).....Part III (State AI - Additional)
 California SB 53 - Transparency in Frontier AI Act.....Part III (State AI - Additional)
 SEC Predictive Data Analytics Proposal and Existing Disclosure Obligations.....Part I (Sector-Specific)
 Section 1557 of the Affordable Care Act and AI Discrimination.....Part I (Sector-Specific)
 New York SHIELD Act and NY DFS 23 NYCRR Part 500.....Part II (State Privacy)
 Singapore - Personal Data Protection Act (PDPA) and Model AI Governance Framework.....Part VII (APAC & MEA Core)
 SOC 2 - Service Organization Controls.....Part VIII (Security Standards)
 South Africa - Protection of Personal Information Act (POPIA).....Part VII (APAC & MEA Core)
 South Korea - Personal Information Protection Act (PIPA) and AI Basic Act.....Part VII (APAC & MEA Core)
 U.S. State Data Breach Notification Laws (Survey).....Part II (State Other)
 New York Stop Hacks and Improve Electronic Data Security (SHIELD) - Detailed.....Part II (State Other)
 The Stored Communications Act (SCA) and the Wiretap Act.....Part I (Federal)
 Colorado Student Data Transparency and Security Act.....Part II (State Other)
 California Student Online Personal Information Protection Act (SOPIPA).....Part II (State Other)

T

Taiwan - Personal Data Protection Act.....Part VII (APAC Additional)
 The Telephone Consumer Protection Act (TCPA).....Part I (Federal)
 Tennessee Ensuring Likeness, Voice, and Image Security Act (ELVIS Act).....Part III (State AI)
 Tennessee Information Protection Act (TIPA).....Part II (State Privacy - Second Wave)
 Thailand - Personal Data Protection Act (PDPA).....Part VII (APAC Additional)
 Thaler v. Perlmutter and AI Inventorship/Authorship Doctrine.....Part IV (IP & AI)
 Title VII of the Civil Rights Act and AI in Employment.....Part I (Civil Rights)
 Trade Secret Protection of AI Models and Training Data.....Part IV (IP & AI)
 TSA Pipeline and Rail Cybersecurity Directives.....Part I (Critical Infrastructure)
 Tunisia - Personal Data Protection Law.....Part VII (MENA & Africa)
 Turkey - Personal Data Protection Law (KVKK).....Part VII (MENA & Africa)

U

UK GDPR and the Data Protection Act 2018.....Part VI (UK & Americas)
 UK Online Safety Act.....Part VI (UK & Americas)
 UK Sectoral AI Regulation and AI Safety Institute.....Part VI (UK & Americas)
 UK Text and Data Mining and AI Copyright Reform.....Part IV (IP & AI)
 UNESCO Recommendation on the Ethics of Artificial Intelligence.....Part VIII (AI Standards)
 United Arab Emirates - Federal Personal Data Protection Law and AI Strategy.....Part VII (APAC & MEA Core)
 Uruguay - Personal Data Protection Law.....Part VI (Latin America)
 USPTO Inventorship Guidance for AI-Assisted Inventions.....Part IV (IP & AI)
 Utah Artificial Intelligence Policy Act (UAIP).....Part III (State AI - Additional)
 Utah Consumer Privacy Act (UCPA).....Part II (State Privacy)

V

Video Privacy Protection Act (VPPA)	<i>Part I (Communications)</i>
Vietnam - Personal Data Protection Decree and Cybersecurity Law.....	<i>Part VII (APAC Additional)</i>
Virginia Consumer Data Protection Act (VCDPA)	<i>Part II (State Privacy)</i>

W

Washington Biometric Identifier Statute (RCW 19.375).....	<i>Part II (State Other)</i>
Washington My Health My Data Act (MHMDA).....	<i>Part II (State Privacy)</i>
Workforce Surveillance Statutes (State Survey).....	<i>Part I (Civil Rights)</i>

Closing Note

AI governance law is the most rapidly evolving regulatory surface a contemporary compliance practitioner is asked to navigate. The framework chosen - the AI Governance Stack - was designed to outlast specific statutory regimes precisely because the substantive rules will continue to change while the operational architecture must remain coherent.

Three observations are worth committing to long-term practice. First, the substantive convergence across jurisdictions - on consequential-decision impact assessment, training-data documentation, fairness testing, and meaningful human oversight - is now sufficient to design a single global AI governance program with jurisdiction-specific configuration, rather than parallel programs per regime. Second, the layer at which an AI risk is best controlled is rarely the layer at which the regulation appears to bite; the Stack Lens callouts in this Guide are designed to surface that reality. Third, the strongest compliance defenses are built before the question is asked - the program documentation, governance artifacts, and impact assessments that exist when a regulator inquires materially constrain enforcement outcomes.

The companion textbook, "Governing Intelligence," sets out the underlying methodology in full. This Field Guide is the practitioner's working reference. Used together, they support the construction and defense of AI governance programs across the regulatory regimes that practitioners actually face - and the ones that will inevitably follow.

- Noah M. Kenney • Digital 520 • 2026 -