SHADOW AI DISCOVERY · COMPLIANCE FRAMEWORK

AI Regulatory Reference

Umbravi · Kango Labs LLC · umbravi.io

Version 2.0 · May 2026

Covers (8 frameworks): Colorado AI Act (SB26-189) · California's CPPA Automated Decision-Making Technology (ADMT) regulations · EU AI Act · Massachusetts Data Security Regulation (201 CMR 17.00) · HIPAA · MODPA · COPPA · IRS 501(c)(3) (with Federal Grant Compliance).

PURPOSE OF THIS REFERENCE

This document provides the regulatory framework used by Umbravi's AI Auditor to assess Shadow AI risk in client organizations. It defines applicable laws, risk thresholds, compliance obligations, and recommended remediation actions. It is not legal advice. Organizations should consult qualified counsel for formal compliance determinations.

Contents

PART I · REGULATORY FRAMEWORK

PART II · UMBRAVI REFERENCE

PART I

Regulatory framework

Sections 1–8 · Laws and regulations applicable to AI tool usage in US organizations.

SHADOW AI DISCOVERY — FRAMEWORK COVERAGE

Sections 1, 2, 3, and 4 are applied in every Shadow AI Risk Report based on vendor identity and known tool behaviors from your software (i.e., SaaS) spend export. Sections 5, 6, 7, and 8 require additional organizational profile data and are available in extended tiers. FTC Act enforcement history is referenced across all tiers as a supplementary risk signal.

1. Colorado AI Act (SB26-189)

EFFECTIVE DATE **January 1, 2027** · **STATUS** **Signed · Pending effective date**

INCLUDED IN SHADOW AI DISCOVERY

This framework is assessed in every Shadow AI Risk Report.

Overview

The Colorado Artificial Intelligence Act was enacted as SB26-189, signed by Governor Polis on May 14, 2026. SB26-189 repealed and replaced the prior Colorado AI Act (SB24-205). It takes effect January 1, 2027, and applies to developers and deployers of high-risk AI systems that interact with Colorado consumers.

Applicability

- Applies to any developer or deployer of an AI system that uses Automated Decision-Making Technology (ADMT) to make or materially influence a consequential decision affecting a Colorado consumer.
- Consequential decisions are defined statutorily and include employment, education, residential real estate in Colorado, financial or lending services, insurance, healthcare services, and essential government services and public benefits.
- Applies to organizations serving Colorado residents regardless of where the organization is headquartered.
- No employee or revenue threshold carve-out — SMBs are fully in scope.

Key definitions

TERM	DEFINITION
Covered ADMT	Automated Decision-Making Technology that materially influences a consequential decision. The "materially influences" trigger is broader than California's "replaces or substantially replaces" trigger.
Consequential decision	Any decision that has a material legal or similarly significant effect on a consumer's access to or cost of one of the statutorily named domains (employment, education, Colorado residential real estate, finance or lending, insurance, healthcare, or essential government services).
Developer	A person that develops or substantially modifies a covered ADMT for deployment by a deployer.
Deployer	A person that deploys a covered ADMT in the state of Colorado.

Core obligations

- Pre-use consumer notice when covered ADMT is used in a consequential decision.
- Post-adverse-outcome disclosure when an ADMT contributes to a consequential decision that adversely affects the consumer.
- Consumer right to correction of inaccurate data used in the ADMT.
- Meaningful human review to the extent commercially reasonable.
- Annual impact assessment for each covered ADMT.
- Risk management program governing the developer's or deployer's ADMT use, documented and maintained.
- Avoid algorithmic discrimination on the basis of protected characteristics.

General-purpose AI chatbot condition

General-purpose AI chatbots are conditionally excluded from the ADMT definition only when both of the following are met: (a) the chatbot is not configured or marketed for consequential decisions, and (b) the deployer maintains an acceptable use policy (AUP) that prohibits such use. Without a documented AUP, general-purpose chatbots remain in scope.

SHADOW AI RISK INDICATOR · HIGH

Any unapproved AI tool used by HR, finance, or legal teams to screen candidates, evaluate creditworthiness, or make benefits decisions is a potential covered ADMT under SB26-189. This includes AI-assisted resume screeners, productivity scoring tools, and automated scheduling systems used in consequential decisions.

Enforcement and penalties

- Enforced exclusively by the Colorado Attorney General under the Colorado Consumer Protection Act. A violation is a deceptive trade practice.
- 60-day cure period applies; waivable for knowing or repeated violations.
- No private right of action.
- The Colorado Attorney General is conducting rulemaking through 2026 on multiple required topics, including post-adverse-outcome disclosure mechanics, correction procedures, and the "materially influence" presumption framework. Specific implementation requirements may shift as rulemaking concludes.

2. California's CPPA Automated Decision-Making Technology (ADMT) regulations

Cal. Code Regs. tit. 11, Art. 11

EFFECTIVE DATE **Regulations: January 1, 2026 · ADMT obligations: January 1, 2027**

INCLUDED IN SHADOW AI DISCOVERY

This framework is assessed in every Shadow AI Risk Report.

Overview

The California Privacy Protection Agency (CPPA) has issued regulations under the California Consumer Privacy Act (CCPA) governing the use of Automated Decision-Making Technology (ADMT). The regulations themselves took effect January 1, 2026. ADMT-specific obligations (pre-use notice, opt-out rights, access rights) take effect January 1, 2027. Risk-assessment obligations are running now.

Applicability — CCPA business threshold

The CPPA's ADMT regulations apply to businesses meeting the CCPA business threshold:

- $25M+ annual gross revenue, **OR**
- Annual processing of personal information of 100,000+ consumers or households, **OR**
- 50%+ of annual revenue derived from selling or sharing personal information.

For organizations meeting the threshold and processing personal information about California consumers, employees, or applicants, the regulations apply regardless of headquarters location.

ADMT trigger — three-prong human-involvement test

California's ADMT trigger is narrower than Colorado's. California requires the technology to *replace or substantially replace* human decision-making. Human involvement defeats the ADMT trigger only if all three of the following are satisfied:

- The human reviewer knows how to interpret and use the ADMT's output.
- The reviewer considers the output alongside other information relevant to the decision.
- The reviewer has authority to make or change the decision based on that review.

Tools where human review is present but cursory, undocumented, or rubber-stamping still trigger ADMT obligations.

Significant-decision domains

California's regulations apply ADMT-specific obligations to significant decisions in the following domains:

- Financial or lending services.
- Housing.

- Education enrollment or opportunities.
- Employment or independent-contracting opportunities or compensation.
- Healthcare services.

Advertising is **explicitly excluded** from significant-decision domains. Tools used to train ADMT on personal information may still fall under separate risk-assessment obligations even when not used directly in significant decisions.

Risk-assessment obligations

- Businesses processing personal information in "significant-risk" categories must conduct risk assessments. Using or training ADMT for significant decisions is a significant-risk category.
- Risk assessments for ongoing 2026–2027 processing must be completed by **December 31, 2027**.
- A summary attestation must be submitted to the CPPA by **April 1, 2028**.
- Risk-assessment obligations are running now and are independent of the ADMT-specific obligations effective January 1, 2027.

SHADOW AI RISK INDICATOR · HIGH

Unapproved AI tools used in HR and employment workflows where human review is undocumented are the most likely category to trigger California ADMT obligations. This includes AI-assisted resume screening, enterprise AI search indexing HR data, meeting transcription used in recruiting or performance review, and AI writing assistants used in employment decisions.

Enforcement and penalties

- Enforced by the California Privacy Protection Agency and the California Attorney General under the CCPA.
- Civil penalties up to **$2,500 per violation**.
- Civil penalties up to **$7,500 per intentional violation** or per violation involving a minor.
- Private right of action available for certain CCPA breach scenarios.

3. EU AI Act (Regulation 2024/1689)

ENTERED INTO FORCE August 1, 2024 · **STATUS** Phased through 2030

INCLUDED IN SHADOW AI DISCOVERY

This framework is assessed in every Shadow AI Risk Report.

Overview

The EU Artificial Intelligence Act is the world's first comprehensive legal framework on AI. It entered into force August 1, 2024, and is being phased in through 2030. It applies a risk-based approach, classifying AI systems into four categories: Unacceptable Risk, High Risk, Limited Risk, and Minimal Risk.

Applicability — US organizations

- Applies to any organization placing AI systems on the EU market OR whose AI systems affect EU-based individuals.
- US companies with EU customers, EU employees, or EU data subjects are in scope regardless of physical location.
- No minimum size threshold — SMBs that process EU data or serve EU consumers must comply.

Risk classification framework

RISK TIER	EXAMPLES	OBLIGATIONS
Unacceptable	Social scoring by governments, real-time biometric surveillance, subliminal manipulation	Prohibited — banned outright
High risk	CV screening, credit scoring, medical devices, educational assessment, law enforcement	Full compliance: conformity assessment, documentation, human oversight, transparency
Limited risk	Chatbots, deepfake generators, emotion recognition	Transparency obligations — users must be informed they are interacting with AI
Minimal risk	AI-enabled spam filters, AI in video games	No mandatory requirements — voluntary codes of conduct

Key compliance obligations (high risk)

- Maintain an AI system inventory — all high-risk systems must be registered in the EU database.
- Conduct a Fundamental Rights Impact Assessment (FRIA) before deployment.

- Implement a quality management system covering data governance, technical documentation, and monitoring.
- Ensure human oversight mechanisms are built into high-risk AI deployments.
- Log all high-risk AI system decisions for post-market monitoring.
- General-Purpose AI (GPAI) model providers must publish technical documentation and comply with copyright law.

Phase-in timeline

DATE	MILESTONE
Feb 2, 2025	Prohibited AI practices ban took effect — unacceptable-risk systems must be discontinued.
Aug 2, 2025	GPAI model rules took effect — providers of general-purpose AI models (GPT-4, Claude, Gemini, etc.) must comply.
Aug 2, 2026	High-risk AI system obligations fully enforced — most enterprise AI deployments in scope.
Aug 2, 2030	All remaining provisions in force, including legacy embedded AI systems.

Penalties

- Prohibited AI practices: up to **€35 million or 7% of global annual turnover**, whichever is higher.
- High-risk violations: up to **€15 million or 3% of global annual turnover**.
- Incorrect information to authorities: up to **€7.5 million or 1% of global annual turnover**.

SHADOW AI RISK

Employees using unapproved AI tools that perform CV screening, customer scoring, or automated decision-making involving EU data subjects may constitute unauthorized deployment of a high-risk AI system. The organization — not the employee — bears liability.

4. Massachusetts Data Security Regulation (201 CMR 17.00)

EFFECTIVE DATE **March 1, 2010** · STATUS **In force**

INCLUDED IN SHADOW AI DISCOVERY

This framework is assessed in every Shadow AI Risk Report.

Overview

The Massachusetts Data Security Regulation (201 CMR 17.00) establishes minimum standards for the protection of personal information of Massachusetts residents. It is one of the strictest state data security regulations in the US and applies broadly to any organization that holds personal information about Massachusetts residents.

Applicability

- Any person or organization that owns, licenses, stores, or maintains personal information about Massachusetts residents.
- No minimum size threshold — sole proprietors are in scope.
- Personal information defined as first name and last name combined with SSN, driver's license number, financial account number, or credit/debit card number.

Written Information Security Program (WISP) requirements

- Every in-scope organization must maintain a comprehensive Written Information Security Program (WISP).
- The WISP must cover the systems that handle personal information, including AI tools used in workflows that handle protected data. AI tools that access or store personal information must be included in the WISP or prohibited.
- The WISP must include employee training on security awareness; AI tool usage policies must be part of training.
- Third-party service provider oversight is required. AI vendors with access to personal information must be contractually required to maintain appropriate safeguards.

SHADOW AI RISK

Any AI tool that stores, processes, or transmits Massachusetts resident personal information without being included in the WISP is a compliance gap. Browser-extension AI tools that read page content are especially high risk — they may silently capture personal information displayed in web applications.

Penalties

- Civil penalties up to **$5,000 per violation**.
- Attorney General enforcement.
- Private right of action — individuals may sue for data breaches resulting from security failures.

5. HIPAA — Health Insurance Portability and Accountability Act

EFFECTIVE DATE **Privacy Rule: April 14, 2003 · Security Rule: April 20, 2005** · **STATUS** **In force**

> **NOT INCLUDED IN SHADOW AI DISCOVERY TIER**
>
> Available in extended tiers. Requires additional information. Contact hello@umbravi.io to learn more.

Overview

HIPAA establishes national standards for protecting sensitive patient health information (Protected Health Information, or PHI). The Privacy Rule, Security Rule, and Breach Notification Rule together govern how covered entities and business associates handle PHI. AI tools that process, store, or transmit PHI are subject to full HIPAA compliance requirements.

Applicability

- **Covered Entities:** health plans, healthcare clearinghouses, healthcare providers that transmit health information electronically.
- **Business Associates:** any vendor or contractor that creates, receives, maintains, or transmits PHI on behalf of a covered entity — this includes AI tool vendors.
- Shadow AI tools that ingest PHI (for example, meeting transcription tools that record patient discussions, AI writing tools used to draft clinical notes) automatically trigger HIPAA obligations.

AI-specific risk areas

AI TOOL CATEGORY	HIPAA RISK
Meeting transcription (Otter.ai, Fireflies)	If clinical discussions are recorded, PHI is transmitted to a third-party server. Requires BAA with vendor.
AI writing assistants (ChatGPT, Claude, Jasper)	Pasting patient data into any AI tool without a BAA is a HIPAA violation, regardless of intent.
AI coding assistants (GitHub Copilot, Cursor)	If proprietary code references patient data structures, transmission to AI servers may constitute PHI disclosure.
AI search (Glean, Notion AI)	Enterprise AI search tools that index internal data may index PHI if not properly segmented.

Key obligations for AI tool governance

- Execute a Business Associate Agreement (BAA) with every AI vendor that may access PHI.

- Conduct a Security Risk Analysis annually — must include all AI tools in scope.
- Implement minimum-necessary standard — AI tools should only access PHI required for their function.
- Breach Notification: notify HHS and affected individuals within 60 days of discovering a breach.
- Maintain an audit trail of all PHI access including AI system access logs.

Penalties

TIER	PENALTY RANGE
Tier 1 (unknowing)	$100–$50,000 per violation · max $25,000/year
Tier 2 (reasonable cause)	$1,000–$50,000 per violation · max $100,000/year
Tier 3 (willful neglect, corrected)	$10,000–$50,000 per violation · max $250,000/year
Tier 4 (willful neglect, uncorrected)	$50,000 per violation · max $1.9 million/year

6. MODPA — Maryland Online Data Privacy Act

EFFECTIVE DATE October 1, 2025 · **STATUS** In force

NOT INCLUDED IN SHADOW AI DISCOVERY TIER

Available in extended tiers. Requires additional information. Contact hello@umbravi.io to learn more.

Overview

The Maryland Online Data Privacy Act (MODPA), effective October 1, 2025, is one of the most expansive state privacy laws in the US. It covers consumers and applies to organizations that process personal data of Maryland residents. It has notably broad applicability thresholds compared to other state laws.

Applicability thresholds

- Controls or processes personal data of 35,000+ Maryland consumers annually, OR
- Controls or processes personal data of 10,000+ Maryland consumers AND derives more than 20% of gross revenue from selling personal data.
- No revenue floor — smaller organizations serving Maryland consumers may qualify.

AI-specific provisions

- Prohibits processing sensitive data (including biometric, health, and precise geolocation data) without consent.
- Requires Data Protection Assessments for processing activities that present a heightened risk — including profiling that produces legal or similarly significant effects.
- Explicitly covers automated decision-making — consumers have the right to opt out of profiling used for consequential decisions.
- Controllers must disclose if personal data is used to train AI models.

SHADOW AI RISK INDICATOR · HIGH EXPOSURE

AI tools that process behavioral data, browsing history, or interaction data from Maryland consumers without disclosure or consent — including analytics AI, personalization engines, and recommendation systems — are high-exposure items under MODPA.

Penalties

- Civil penalties enforced by the Maryland Attorney General.
- Up to **$10,000 per violation**.
- Up to **$25,000 per intentional violation**.

- 30-day cure period for first violations.

7. COPPA — Children's Online Privacy Protection Act

EFFECTIVE DATE **April 21, 2000 · Updated Rule: July 1, 2013** · **STATUS** **In force**

> **NOT INCLUDED IN SHADOW AI DISCOVERY TIER**
>
> Available in extended tiers. Requires additional information. Contact hello@umbravi.io to learn more.

Overview

COPPA governs the online collection, use, and disclosure of personal information from children under 13. It applies to operators of commercial websites and online services — including AI-powered applications — directed to children or with actual knowledge they are collecting data from children under 13.

AI-specific applicability

- Any AI tool embedded in a product or service directed to children is subject to COPPA.
- AI tutoring tools, educational AI assistants, and AI-powered games used by minors are in scope.
- If an organization serves K-12 schools, all AI tools processing student data may be in scope.
- Actual knowledge standard: if an AI tool's intake form or registration collects age and the user indicates they are under 13, COPPA obligations activate.

Key obligations

- Obtain verifiable parental consent before collecting personal information from children under 13.
- Maintain a clear privacy policy disclosing data collection and use practices.
- Provide parents access to collected information and the ability to delete it.
- No conditioning of services on collection of more personal information than reasonably necessary.

Penalties

- FTC enforcement — civil penalties up to **$51,744 per violation per day**.
- State Attorney General enforcement also available.

8. IRS 501(c)(3) — Nonprofit AI Governance Obligations

EFFECTIVE DATE Ongoing · AI guidance evolving · **STATUS** In force

> **NOT INCLUDED IN SHADOW AI DISCOVERY TIER**
>
> Available in extended tiers. Requires additional information. Contact hello@umbravi.io to learn more.

Overview

Nonprofit organizations exempt under IRS Section 501(c)(3) have fiduciary obligations that extend to technology governance, including AI. While the IRS does not have a specific AI regulation, board fiduciary duty, donor data obligations, and grant compliance requirements create meaningful AI governance exposure for nonprofits.

Board fiduciary duty and AI

- Duty of Care requires board members to make informed decisions — failing to establish AI governance policies when AI is in use may constitute a breach of the duty of care.
- Duty of Loyalty prohibits use of organizational resources — including AI tools — for private benefit.
- Boards should formally adopt an AI usage policy and include AI risk in organizational risk registers.

Donor data and AI

- Donor PII (names, addresses, financial data) processed by unapproved AI tools may constitute a breach of donor trust and potentially a state privacy law violation.
- AI tools used for donor prospecting or wealth screening must be disclosed in privacy policies.
- Donor data must not be used to train third-party AI models without explicit consent.

Federal grant compliance

- Organizations receiving federal grants are subject to 2 CFR Part 200 (Uniform Guidance) — data security requirements apply to all systems processing grant-related data.
- AI tools that process beneficiary data, financial data, or program data funded by federal grants must be disclosed to the administering agency.
- Unauthorized use of AI to generate grant reports or compliance documentation without disclosure may violate grant terms.

PART II

Umbravi reference

Sections 9–12 · Scoring methodology, risk matrix, glossary, and reference notes.

9. Umbravi Risk Scoring Methodology

Regulatory Readiness Score

The Umbravi Regulatory Readiness Score is a 0–100 composite score assessing an organization's AI compliance posture based on its identified AI tool inventory. A higher score indicates greater regulatory readiness. The score is calculated at the time of report generation based on available software (i.e., SaaS) spend data and is not a formal compliance certification.

Score bands

SCORE	BAND	INTERPRETATION
80–100	Low Exposure	Majority of AI tools are approved or pose minimal regulatory risk. Formal AI inventory recommended.
60–79	Moderate Exposure	Several at-risk tools detected. At least one regulatory framework triggered. Action plan required.
40–59	Elevated Exposure	Material compliance gaps identified. Multiple frameworks triggered. Immediate remediation recommended.
0–39	High Exposure	Significant at-risk AI tool footprint. Regulatory exposure is material. Legal counsel recommended.

Tool risk classification

LEVEL	CRITERIA
High	Tool trains on user data by default, stores data on third-party servers in unspecified jurisdictions, has no enterprise privacy mode, or is known to have experienced data breaches. Multiple regulatory frameworks triggered.
Medium	Tool may process sensitive data depending on use case, has privacy controls available but not enabled by default, or is subject to a single regulatory framework.
Low	Tool has strong enterprise privacy controls, does not train on customer data, operates under established BAAs or DPAs, and presents minimal regulatory exposure.
Approved	Tool appears in the organization's known-approved stack (inferred from spend category, vendor type, and known enterprise agreements). No significant regulatory flags identified.

Data sources

- Tool risk classifications are derived from publicly available vendor privacy policies, terms of service, and data processing agreements current as of the report date.
- Regulatory framework applicability is assessed based on tool type and known use case patterns — not client-specific disclosure.
- Umbravi maintains an internal AI Tool Risk Database updated quarterly.
- Regulatory citations reference official published text of applicable laws and regulations.

IMPORTANT LIMITATION

Risk observations based on software (i.e., SaaS) spend data alone reflect probable risk exposure based on known tool behaviors. Actual regulatory exposure depends on how tools are used within the organization, which data they access, and the specific regulatory profile of the organization. Full assessment requires the Umbravi Extended Intake (available in subsequent report tiers).

10. Quick reference — Regulatory framework matrix

Risk level of each AI tool category against each regulatory framework. Based on known tool behaviors and typical use patterns — not client-specific configuration.

TOOL CATEGORY	CO AI ACT	CA CPPA ADMT	EU AI ACT	MA 201 CMR	HIPAA	MODPA	COPPA
HR / Recruiting AI	HIGH	HIGH	HIGH	MED	LOW	MED	LOW
Meeting transcription	MED	MED	MED	MED	HIGH	MED	MED
AI coding assistants	LOW	LOW	MED	MED	MED	LOW	LOW
AI writing / content	LOW	MED	MED	MED	HIGH	MED	MED
Browser-extension AI	MED	MED	HIGH	HIGH	HIGH	HIGH	HIGH
Enterprise AI search	MED	HIGH	HIGH	MED	HIGH	HIGH	MED
AI analytics / BI	MED	MED	MED	MED	MED	HIGH	MED

11. Glossary of key terms

The following terms appear throughout Umbravi's Shadow AI Risk Reports and regulatory reference materials.

ADMT — Automated Decision-Making Technology

Used in both the Colorado AI Act (SB26-189) and California's CPPA regulations to refer to technologies that influence or make consequential decisions. Colorado's trigger is "materially influences"; California's is "replaces or substantially replaces human decision-making." The two definitions are not identical and may surface different findings against the same tool.

AI Manifest

A formal, documented inventory of all AI systems in use within an organization, including their purpose, data inputs, vendor, and applicable regulatory classifications. Required under the EU AI Act for high-risk systems.

Algorithm

A set of rules or instructions followed by a computer system to perform a task or make a decision. In AI systems, algorithms process input data and produce outputs that may influence decisions.

AUP — Acceptable Use Policy

A documented internal policy governing how AI tools may and may not be used within an organization. Under the Colorado AI Act (SB26-189), a documented AUP that prohibits use of general-purpose AI chatbots for consequential decisions is required to qualify for the conditional exclusion from ADMT scope.

Automated Decision-Making

The use of AI or algorithmic systems to make or substantially influence decisions without meaningful human review. Regulated under the EU AI Act, Colorado AI Act, California's CPPA ADMT regulations, and MODPA when decisions have material effects on individuals.

Business Associate Agreement (BAA)

A HIPAA-required contract between a covered entity and a vendor (business associate) that processes Protected Health Information (PHI) on its behalf. Any AI vendor with access to PHI must have a signed BAA in place.

CCPA — California Consumer Privacy Act

The underlying California statute governing consumer privacy rights and business obligations regarding personal information. The CCPA establishes the business threshold for applicability ($25M revenue / 100,000 consumers / 50% data revenue) and delegates rulemaking authority to the California Privacy Protection Agency (CPPA).

Consequential Decision

Under the Colorado AI Act (SB26-189), any decision that has a material legal or similarly significant effect on a consumer's access to or cost of education, employment, residential real estate in Colorado, financial or lending services, insurance, healthcare, or essential government services.

Covered Entity

Under HIPAA, a health plan, healthcare clearinghouse, or healthcare provider that transmits health information electronically. Covered entities are directly subject to all HIPAA requirements.

CPPA — California Privacy Protection Agency

The California state agency that issues regulations under the California Consumer Privacy Act (CCPA). The CPPA's Automated Decision-Making Technology (ADMT) regulations are codified at Cal. Code Regs. tit. 11, Art. 11. Note: CPPA (the agency) is distinct from CCPA (the statute).

Data Processing Agreement (DPA)

A contract between a data controller and a data processor that governs how personal data is handled, stored, and protected. Required under GDPR and recommended for any AI vendor relationship involving personal data.

Deployer

Under the Colorado AI Act and EU AI Act, an organization that deploys an AI system for use. Deployers bear compliance obligations including impact assessments, consumer notices, and risk management programs.

Developer

Under the Colorado AI Act and EU AI Act, an organization that develops or substantially modifies an AI system. Developers must provide documentation, transparency, and risk management support to deployers.

Fiduciary Duty

The legal obligation of board members and officers to act in the best interest of the organization. In the context of AI, this includes a duty of care to understand and govern AI risks, particularly for nonprofit boards.

GDPR

The General Data Protection Regulation (EU Regulation 2016/679). Applies to any organization processing personal data of EU residents. Requires lawful basis for processing, data subject rights, and data protection by design.

General Purpose AI (GPAI) Model

Under the EU AI Act, an AI model trained on broad data that can perform a wide range of tasks. Examples include GPT-4, Claude, and Gemini. GPAI providers must publish technical documentation and comply with copyright obligations.

High-Risk AI System

Under the EU AI Act and Colorado AI Act (SB26-189), an AI system that poses significant risk to health, safety, or fundamental rights. Examples include CV screening tools, credit scoring systems, medical diagnostic tools, and educational assessment AI.

Impact Assessment

A structured evaluation of an AI system's potential risks before deployment. Required under the Colorado AI Act, EU AI Act, California's CPPA ADMT regulations (as a "risk assessment"), and MODPA for systems that present heightened risk. Also called a Fundamental Rights Impact Assessment (FRIA) under the EU AI Act.

Personal Information

Information that identifies or can reasonably be linked to a specific individual. Under the Massachusetts Data Security Regulation (201 CMR 17.00), personal information is first name and last name combined with SSN, driver's license number, financial account number, or credit/debit card number.

Protected Health Information (PHI)

Under HIPAA, any individually identifiable health information created, received, maintained, or transmitted by a covered entity or business associate. Includes diagnoses, treatment records, payment information, and any data that could identify a patient.

Regulatory Readiness Score

Umbravi's 0–100 composite score assessing an organization's AI compliance posture based on identified AI tools. A higher score indicates greater readiness. Not a formal compliance certification.

Shadow AI

AI tools in use within an organization that have not been formally approved, documented, or governed by IT or compliance teams. Also called "unapproved AI" or "unsanctioned AI." Shadow AI is the primary risk category addressed by Umbravi's discovery service.

Software (SaaS) Spend Export

A CSV or spreadsheet export from an organization's accounting, expense management, or IT spend platform listing software vendor names, amounts, and transaction dates. The primary data input for Umbravi's Shadow AI Risk Report.

Significant Decision

Under California's CPPA ADMT regulations, a decision in one of the statutorily named domains: financial or lending services, housing, education enrollment or opportunities, employment or independent-contracting opportunities or compensation, or healthcare services. Advertising is explicitly excluded.

Training Data

The dataset used to develop and improve an AI model. When AI vendors train their models on user-submitted data by default, this creates regulatory exposure — particularly under the EU AI Act, HIPAA, the Colorado AI Act, and California's CPPA ADMT regulations.

Written Information Security Program (WISP)

A documented security program required by the Massachusetts Data Security Regulation (201 CMR 17.00) for any organization that holds personal information of Massachusetts residents. Must include policies, employee training, vendor oversight, and incident response procedures. AI tools that handle personal information must be included in or prohibited by the WISP.

Zero-Retention Terms

Contractual provisions in which a vendor agrees not to store, use, or retain customer data after processing. Umbravi's AI analysis pipeline operates under zero-retention terms with its language model provider, meaning submitted data is not used to train or improve AI models.

12. About this reference

This reference was prepared by Umbravi for use in generating Shadow AI Risk Reports. It reflects the regulatory landscape as of May 2026. Regulatory frameworks evolve. Umbravi updates this reference quarterly to reflect new guidance, enforcement actions, and legislative changes.

For questions about this reference or to discuss your organization's specific regulatory exposure, contact Umbravi at **hello@umbravi.io**.

NOT LEGAL ADVICE

Nothing in this reference constitutes legal advice or creates an attorney-client relationship. Organizations should consult qualified legal counsel for formal regulatory compliance determinations. Umbravi's reports are operational and technical guidance tools only.