



Abstract

US agencies procuring AI systems and evaluating vendor compliance claims are operating without a clear evaluation framework in today's regulatory environment. IT leaders who lack a clear understanding of which external governance requirements apply to that vendor are evaluating an incomplete picture. A vendor subject to EU AI Act high-risk obligations has undergone independent conformity assessments against binding technical standards. A vendor operating under China's CAC generative AI registration regime has completed government-reviewed security assessments and is subject to active administrative enforcement. A vendor operating exclusively under voluntary US frameworks has self-attested.

This article examines the EU AI Act, China's CAC enforcement regime, and the current US federal posture across six governance dimensions. Analysis of research conducted by GIAG draws practical conclusions for federal and state agency CIOs and CAIOs managing AI procurements and deployments. The central finding is that the verification asymmetry across the three models creates an evaluation gap that the proposed Great American Artificial Intelligence Act (GAAIA) will not resolve alone, even if the bill passes in its current form.

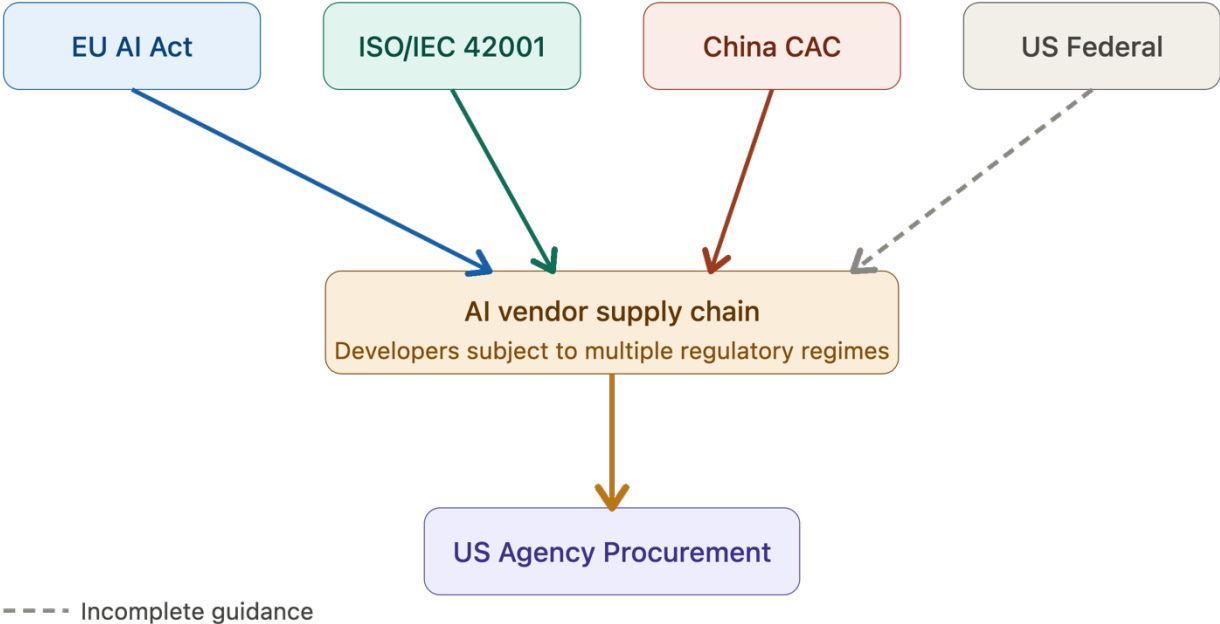
In the last section, I provide five important diagnostic questions that CIOs and CAIOs can apply immediately, without waiting for new policy authority.

Why the Global Comparison Matters Now

Federal and state agency technology leaders are already tracking a substantial domestic requirement set. The AI Action Plan, the GAAIA discussion draft, and CISA's forthcoming binding operational directive create compounding obligations on their own. The case for adding a global comparative layer is operational, not academic, and it turns on a specific procurement reality.

US federal agencies are buyers of AI systems built by developers operating across multiple regulatory jurisdictions. A frontier AI model procured from a major developer may be subject to EU AI Act

conformity assessments, Chinese CAC security review requirements, and US federal procurement standards simultaneously. The developer's governance obligations under those external regimes shape what documentation the agency receives, what audit rights attach to the procurement, and what the developer's liability exposure is for the system's behavior in deployment. Those obligations are neither uniform nor equivalent.



The practical consequence is a procurement evaluation standard that treats structurally different compliance postures as equivalent. This is a gap no current US instrument requires agencies to close. The three sections that follow document what each regime requires.

The EU AI Act: Binding Law with Market Teeth

The EU AI Act entered into force in August 2024 and has implemented in phases. As of August 2, 2025, obligations for general-purpose AI model providers took effect. As of August 2, 2026, full enforcement of high-risk AI system requirements applies. High-risk AI categories include systems used in employment, education, credit, essential public services, law enforcement, and critical infrastructure. This scope encompasses many AI deployment contexts relevant to government procurement.

The Act's architecture is risk-tiered rather than revenue-threshold based. Prohibited practices are banned outright regardless of developer size or revenue. High-risk systems require independent conformity assessment before deployment, ongoing post-market monitoring, incident reporting to national competent authorities, and registration in a public EU database. Transparency requirements apply to limited-risk systems. Minimal-risk systems face no specific obligations. The risk classification is subject to regulatory scrutiny (it is not self-assessed by the developer).

The enforcement mechanism is the element that most distinguishes the EU AI Act from US frameworks. National market surveillance authorities in each member state hold power to order market withdrawal of non-compliant systems, require corrective action, and impose financial penalties. The European AI Office at the Commission level holds enforcement authority over general-purpose AI model providers. Financial penalties reach € 35 million or 7 percent of global annual turnover for violations of prohibited practice provisions, € 15 million or 3 percent of global turnover for high-risk system non-compliance, and € 7.5 million or 1.5 percent for providing misleading information to authorities. These penalties exceed the GDPR's maximum of € 20 million or 4 percent of turnover. This makes the EU AI Act the highest percentage-based penalty regime in EU digital regulation, exceeding the GDPR, which generated € 5.88 billion in cumulative penalties through 2024.

The enforcement toolkit extends beyond financial penalties. Non-compliance findings are public. National authorities maintain records of enforcement actions, and significant cases attract commercial attention. Enterprise commercial customers in regulated sectors are increasingly requiring EU AI Act compliance in procurement contracts, creating market pressure independent of regulatory action. For US agencies procuring systems from EU developers, the practical implication is that those vendors are operating under external governance constraints. The enforcement chain includes binding documentation requirements, independent audit obligations, and market consequences for non-compliance that have no current US federal equivalent.

The prior GIAG research article “Three Frameworks, One Governance Gap” examines the EU AI Act in detail alongside the NIST AI RMF and ISO/IEC 42001, covering the oversight standard, accountability role, and verification dimensions. Practitioners seeking the full framework comparison are referred to that article, available at <https://www.thinkcapital.org/publications.html>. The critical differences are structural: the EU AI Act requires governance to produce evidence. The NIST AI RMF and current US frameworks do not.

China's Model: Administrative Enforcement Without a Comprehensive Law

China's approach to AI governance is architecturally distinct from both the EU and the United States. Between 2021 and 2025, China enacted more sector-specific AI regulations than any other country, taking a path of rapid iterative rulemaking that targeted algorithms, deepfakes, generative AI services, and data security through separate instruments rather than a single comprehensive statute. In late 2025, Beijing removed plans for a unified comprehensive AI law from the legislative schedule, opting instead for continued pilots, targeted measures, and technical standards development. The comprehensive law has been deferred. However, the enforcement infrastructure has not.

The Cyberspace Administration of China (CAC) issued *Interim Administrative Measures for Generative AI Services* in July 2023, jointly with six other central government regulators including the National Development and Reform Commission, Ministry of Education, Ministry of Science and Technology, and Ministry of Public Security. The measures require generative AI service providers to conduct security assessments and file registration with the CAC before offering services to the public. Content labeling requirements for AI-generated material took effect in September 2025 under mandatory

national standard GB 45438-2025. China's amended Cybersecurity Law, effective January 1, 2026, introduces a dedicated AI compliance provision, adding security review requirements and data localization obligations for AI systems processing certain categories of information.

Enforcement is active and immediate. Local CACs have imposed administrative penalties on AI applications providing services without completing required registration procedures. The Shanghai CAC summoned and penalized three AI applications for operating without filing. The Zhejiang CAC ordered mobile app distribution platforms to remove an AI face-swapping application that had not undergone required security assessment. In 2024, the Chongqing CAC shut down a ChatGPT-based service for operating without a security assessment and LLM filing. These are documented enforcement actions using existing authority to suspend or penalize non-compliant AI deployments.

National Technical Committee 260 on Cybersecurity is publishing a pipeline of technical standards that progressively translate framework requirements into binding technical obligations. The *AI Safety Standards System VI.0*, released in January 2025, mapped the complete sequence of forthcoming national standards. Organizations tracking TC260 releases gain early warning of compliance obligations before they become enforceable. Major Chinese AI providers, including DeepSeek and Baidu's Ernie Bot, have completed CAC registration, meaning they operate under an active government visibility regime that includes security assessment documentation accessible to Chinese regulators.

The analytical point for US procurement governance is specific. AI systems built by Chinese developers and deployed in US government contexts exist within a regulatory relationship between those developers and the Chinese government that is not visible through US procurement documentation and is not addressed by any current US federal AI governance instrument. That relationship is a procurement governance consideration independent of any assessment of the political implications of Chinese AI regulation.

For US agencies, the relevant question is not whether to engage with Chinese AI vendors, but whether the procurement evaluation process accounts for the regulatory relationship that engagement creates.

The US Approach: Three Instruments, One Enforcement Gap

The full analysis of the US federal governance posture appears in the companion article in this series, "The Federal AI Governance Stack: What the GAAIA, the AI Action Plan, and CISA's Forthcoming Directive Mean Together." For comparative purposes, the summary follows:

The AI Action Plan operates through executive authority, is effective immediately, and is subject to revision or reversal by subsequent administration. The draft Great American Artificial Intelligence Act proposes the first statutory federal AI governance framework, targets frontier developers above a revenue threshold, and designates a compliance enforcement body that does not yet have operational capacity. The bill has not been formally introduced, its most controversial provision has failed in two prior legislative vehicles at 99-to-1, and its realistic enactment timeline extends at minimum into fall 2026 with passage uncertain within the current Congress. CISA's forthcoming binding operational

directive will create mandatory AI security governance obligations for federal civilian agencies on its own timeline, independent of Congressional action.

The United States has no AI-specific civil penalty regime currently in force at the federal level. No mandatory pre-deployment conformity assessment exists. No independent AI regulator with market surveillance authority and market withdrawal power exists. In the absence of federal action, 45 states had introduced 1,561 AI-related bills as of March 2026, already surpassing the total from all of 2024, and multiple state laws have effective dates in 2026. The federal enforcement posture is primarily litigation: the DOJ AI Litigation Task Force challenged Colorado’s algorithmic discrimination statute in April 2026, the first federal challenge to a state AI law in US history.

Against the EU AI Act and China’s administrative enforcement model, the US federal posture reflects a fundamentally different governance philosophy: innovation primacy, voluntary frameworks, and enforcement through existing legal authorities rather than AI-specific statute. That philosophy has a specific implication for US agencies as buyers in an international procurement environment, which the remaining sections address directly.

A Structured Comparison: Six Governance Dimensions

The following comparison examines the three governance models across six dimensions that CIOs and CAIOs directly control or influence in procurement and deployment decisions. The dimensions are drawn from the framework comparison in “Three Frameworks, One Governance Gap” and extended to the jurisdictional level.

The verification model is the dimension with the most direct implication for procurement governance. Under the EU AI Act, high-risk system compliance requires independent third-party conformity assessment. The developer cannot self-certify. Under China’s CAC regime, compliance requires government-reviewed security assessment and active registration. Under the current US federal framework, compliance is self-attested. The GAAIA would add third-party audit requirements for covered developers, but those audits would be evaluated against CAISI standards that do not yet exist. A US agency accepting vendor compliance documentation today is accepting self-attestation where EU and Chinese models require external verification.

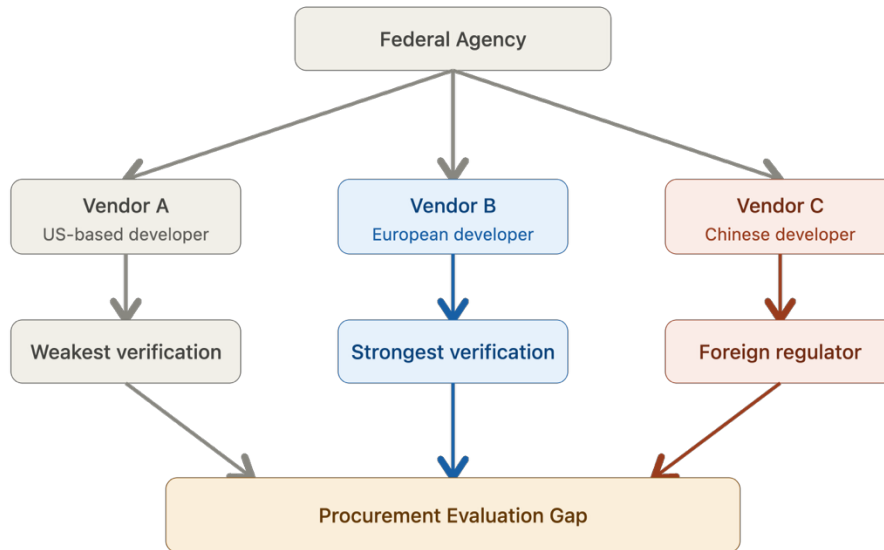
The pre-deployment requirement dimension is the second most consequential. The EU requires conformity assessment before a high-risk system enters service. China requires security assessment and registration before public service. The US requires neither at the federal level. This means that the same AI system, procured by a US agency from a vendor with EU market exposure, may have a more rigorous pre-deployment governance record in its EU compliance file than anything the US procurement process required the vendor to produce. That file is available for evaluation if the agency knows to ask for it.

Dimension	EU AI Act	China (CAC Regime)	US Federal (Current)
Legal authority	Binding regulation, directly applicable in all member states	Binding administrative measures plus technical standards pipeline	Executive directives and voluntary frameworks; no AI-specific statute in force
Risk classification	Statutory risk tiers: prohibited, high-risk, limited, minimal	Sector and service type categories; generative AI treated separately	Revenue threshold (GAAIA, proposed); no tiered classification in force
Pre-deployment requirement	Mandatory conformity assessment for all high-risk systems	Mandatory security assessment and CAC registration before public service	None currently in force at federal level
Enforcement body	National market surveillance authorities plus EU AI Office at Commission level	Local and national CAC with multi-agency coordination	No dedicated AI regulator; CAISI proposed but not yet operational
Maximum financial penalty	€ 35 million or 7% of global annual turnover (prohibited practices)	Administrative penalties, service suspension, and platform removal orders	\$1 million per violation per day (GAAIA, proposed); no federal AI penalty in force today
Verification model	Third-party conformity assessment required; self-certification not sufficient for high-risk	Government security review, registration audit, ongoing CAC oversight	Self-attestation; voluntary framework alignment; GAAIA would add third-party audit (proposed)

Table 1. Three-jurisdiction governance comparison across six dimensions. Sources: EU AI Act (2024); China CAC Interim Measures for Generative AI Services (2023), amended Cybersecurity Law (2026); US AI Action Plan (2025), GAAIA discussion draft (2026), CISA forthcoming BOD.

What the Asymmetry Means for US Agency Planners

The practical consequence of the three-model asymmetry is a procurement evaluation problem that US agencies have not yet been formally asked to solve and for which no current US instrument provides guidance. The following scenario illustrates the gap (the figure below maps the three compliance postures and the evaluation gap they create).



A federal agency issues an RFP for an AI system to support benefits adjudication. Three vendors respond. Vendor A is a US-based developer operating under voluntary NIST AI RMF alignment and self-attestation. Vendor B is a European developer whose system has undergone EU AI Act conformity assessment for high-risk deployment in an employment context, with documented post-market monitoring architecture and incident reporting obligations to a national competent authority. Vendor C is a Chinese developer whose system has completed CAC security assessment and generative AI registration, and whose ongoing compliance is subject to administrative enforcement by Chinese regulators.

Under current US federal procurement governance, the agency has no required methodology for differentiating these three compliance postures. The GAAIA, if enacted, would require Vendors B and C to produce third-party audit documentation if they exceed the revenue threshold. It would not require the agency to evaluate what the external compliance regime means for the system’s behavior in a US deployment context, or what the developer’s obligations to a foreign regulator mean for the agency’s oversight authority over the deployed system.

The governance gap is specific. US agencies need evaluation criteria for AI vendor compliance claims that account for the jurisdiction in which the compliance was established, the independence of the verification mechanism, and the ongoing regulatory relationship between the developer and any foreign government. None of the three US instruments provides those criteria. Building them requires CIOs and CAIOs to exercise procurement judgment that the current framework leaves to their discretion.

The agencies best positioned to close this gap are those that add jurisdictional analysis to their AI vendor evaluation process now. The EU AI Act’s August 2026 full enforcement deadline means that vendors with European market exposure will be producing conformity assessment documentation whether or not US agencies require it. That documentation is available for evaluation if agencies develop criteria for interpreting it. No additional authority is required. It requires that CIOs and CAIOs exercise judgment that the current framework leaves open.

Five Diagnostic Questions for CIOs and CAIOs

These questions are designed for immediate application in ongoing AI procurement and deployment decisions. They do not require waiting for GAAIA enactment or CAISI guidance, and they do not require formal policy authority beyond what CIOs and CAIOs already hold in procurement evaluation.

Question 1. For each AI system currently in procurement or deployment, what jurisdiction does the developer primarily operate in, and what AI governance obligations apply to the developer in that jurisdiction? The answer changes the procurement evaluation standard. A developer subject to EU AI Act high-risk obligations has external documentation requirements that a purely US-based developer operating under voluntary frameworks does not. Knowing this before contract award is more useful than discovering it during an incident review.

Question 2. What verification mechanism supports the vendor's AI compliance claims: self-attestation, voluntary framework alignment, third-party audit, or government-reviewed assessment? The answer determines how much independent weight the documentation carries. Self-attestation against a voluntary framework is the weakest verification signal. Third-party conformity assessment against a binding legal standard is the strongest. US agencies currently have no required methodology for distinguishing these, but nothing prevents building one internally.

Question 3. Does the vendor's system carry ongoing regulatory obligations to a foreign government, and if so, what are those obligations? This question addresses the foreign-developer scenario directly. A system subject to Chinese CAC registration carries an ongoing compliance relationship between the developer and a foreign regulator that is not reflected in US procurement documentation. CIOs and CAIOs need to understand what that relationship means for the agency's oversight authority and the developer's data obligations in a US deployment context.

Question 4. If the EU AI Act would classify your deployed AI use case as high-risk, what does that classification tell you about the governance architecture your oversight program should have? The EU AI Act's risk classification is useful as an independent diagnostic even for US deployments. A system that would be classified as high-risk under the Act is one where EU legal analysis has determined that independent oversight, post-market monitoring, and human intervention capacity are required. US agencies can use that classification as a reference architecture for their own oversight design, independent of any legal obligation to comply.

Question 5. Does your agency's AI vendor evaluation process include jurisdictional analysis, verification mechanism assessment, and foreign regulatory relationship review as standard criteria? If the answer is no, the agency is evaluating vendor AI compliance claims against a standard that the international procurement environment has already rendered incomplete. Building these criteria into standard evaluation does not require new authority. It requires CIOs and CAIOs to exercise judgment the current procurement framework leaves to their discretion.

Research Context and Supporting Advisory

This article is the second in GIAG’s 2026 research article series examining the governance landscape for US federal and state agency AI deployment. The first article, “The Federal AI Governance Stack: What the GAAIA, the AI Action Plan, and CISA’s Forthcoming Directive Mean Together,” covers the domestic governance requirement set in detail. The earlier article “Three Frameworks, One Governance Gap” provides the NIST AI RMF, EU AI Act, and ISO/IEC 42001 comparison that informs the framework analysis here.

GIAG’s Stream One research is documenting how agencies translate AI governance requirements into operational practice. The jurisdictional analysis gap described in this article represents a category of governance requirement that most agencies have not yet encountered as a formal obligation and are not currently building for. Stream Two is examining human oversight quality in agentic AI deployments. The oversight architecture questions the EU AI Act’s high-risk classification raises are directly applicable to Stream Two’s research on what effective oversight looks like at the consequential decision boundaries where autonomous action and human review must interface.

ThinkCapital LLC provides advisory services to federal and state agencies navigating AI procurement governance, jurisdictional compliance analysis, vendor accountability framework development, and oversight architecture design. Structured assessment frameworks and advisory engagements are available for agencies at any stage of this work. Practitioners working through these issues are welcome to engage directly.

Direct research and study engagement inquiries to: research@thinkcapital.org

© 2026 ThinkCapital LLC. All rights reserved. This research article is published for practitioner research and educational purposes. Reproduction with attribution is permitted. Contact: research@thinkcapital.org

Citations

European Union: The EU AI Act

European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union, L 2024/1689. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L_202401689

European Commission. (2024). *EU AI Act: first regulation on artificial intelligence — Timeline and key milestones*. European Commission. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/excellence-and-trust-artificial-intelligence_en

European AI Office. (2025). *Guidance on the application of the AI Act: general-purpose AI models*. European Commission AI Office. <https://digital-strategy.ec.europa.eu/en/policies/ai-office>

Lomas, N. (2025). *EU AI Act enforcement enters new phase: high-risk system obligations apply from August 2026*. TechCrunch. <https://techcrunch.com>

Note: GDPR cumulative penalty figure of €5.88 billion through 2024 sourced from CMS Law GDPR Enforcement Tracker. <https://www.enforcementtracker.com>

China: CAC Regulatory Regime

Cyberspace Administration of China (CAC), National Development and Reform Commission, Ministry of Education, Ministry of Science and Technology, Ministry of Industry and Information Technology, Ministry of Public Security, and National Radio and Television Administration. (2023). *Interim Administrative Measures for Generative Artificial Intelligence Services [生成式人工智能服务管理暂行办法]*. Cyberspace Administration of China. Effective 15 August 2023: https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm

Standardization Administration of China. (2025). *GB 45438-2025: Information Security Technology: Requirements for Identification of AI-Generated Synthetic Content*. National Standards of the People's Republic of China. Effective September 2025. <https://www.sac.gov.cn>

Standing Committee of the National People's Congress. (2025). *Cybersecurity Law of the People's Republic of China (2025 Amendment)*. National People's Congress. Effective January 1, 2026. <http://www.npc.gov.cn>

National Technical Committee 260 on Cybersecurity (TC260). (2025). *AI Safety Standards System Construction Guide V1.0 [人工智能安全标准体系建设指南]*. TC260. Released January 2025. <https://www.tc260.org.cn>

DigiChina (Stanford University Cyber Policy Center). (2024). *Translation and analysis: China's generative AI regulations and enforcement actions*. DigiChina, Stanford Cyber Policy Center. <https://digichina.stanford.edu>

Note: Shanghai, Zhejiang, and Chongqing CAC enforcement actions referenced are documented in DigiChina's enforcement tracker and contemporaneous reporting by Nikkei Asia and Reuters (2024).

United States Federal: Governance Instruments

Executive Office of the President. (2025). *Removing Barriers to American Leadership in Artificial Intelligence: Executive Order 14179*. Federal Register, Vol. 90, January 23, 2025. <https://www.federalregister.gov/executive-order/14179>

Office of Science and Technology Policy (OSTP). (2025). *AI Action Plan*. White House OSTP. <https://www.whitehouse.gov/ostp>

US Senate Commerce Committee. (2026). *Great American Artificial Intelligence Act (GAIA): Discussion Draft*. US Senate Commerce Committee. Released 2026. <https://www.commerce.senate.gov>

Cybersecurity and Infrastructure Security Agency (CISA). (2026). *Forthcoming Binding Operational Directive: AI Security Governance Requirements for Federal Civilian Agencies*. CISA. Timeline: anticipated 2026. <https://www.cisa.gov>

National Institute of Standards and Technology (NIST). (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. <https://doi.org/10.6028/NIST.AI.100-1>

US Department of Justice. (2026). *DOJ AI Litigation Task Force challenges Colorado Artificial Intelligence Act*. US Department of Justice press release, April 2026. <https://www.justice.gov>

Note: The 1,561 state AI bill figure and 45-state count are from the National Conference of State Legislatures AI legislation tracker, as of March 2026. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2026-legislation>

ISO/IEC 42001

International Organization for Standardization / International Electrotechnical Commission. (2023). *ISO/IEC 42001:2023 — Information Technology: Artificial Intelligence — Management System*. ISO/IEC. Published December 2023. <https://www.iso.org/standard/81230.html>

GIAG Research (ThinkCapital LLC)

Bragen, M. (2026). *Three Frameworks, One Governance Gap: NIST AI RMF, EU AI Act, and ISO/IEC 42001 for US Government Practitioners*. Government IT and AI Governance Initiative (GIAG) Research Article. ThinkCapital LLC. <https://www.thinkcapital.org/publications.html>

Bragen, M. (2026). *The Federal AI Governance Stack: What the GAAIA, the AI Action Plan, and CISA's Forthcoming Directive Mean Together*. GIAG Research Article Series, Article 1. ThinkCapital LLC. <https://www.thinkcapital.org/publications.html>

Bragen, M. (2026). *Three Jurisdictions, Three Models: What Global AI Governance Means for US Government Leaders*. GIAG Research Article Series, Article 2. ThinkCapital LLC. <https://www.thinkcapital.org/publications.html>

Bragen, M. (2026). *When Humans Must Intervene: A Decision-Grounded Framework for Human Oversight in Government and Commercial Agentic AI Deployments*. GIAG Working Paper 2. ThinkCapital LLC. <https://www.thinkcapital.org/publications.html>

Additional Reference Sources

OECD. (2024). *OECD AI Policy Observatory: National AI policies and strategies*. Organisation for Economic Co-operation and Development. <https://oecd.ai/en/dashboards/policy-instruments/national-AI-strategies>

International Telecommunication Union. (2024). *ITU AI for Good: Global AI governance tracker*. ITU AI for Good. <https://aiforgood.itu.int>

CMS Law. (2024). *GDPR Enforcement Tracker (cumulative penalty data through 2024)*. CMS Law GDPR Enforcement Tracker. <https://www.enforcementtracker.com>

All URLs verified as of June 2026. Legislative and regulatory references reflect the status of instruments as of the publication date of this article. Readers should verify status of proposed legislation (GAAIA) and forthcoming directives (CISA BOD) at primary government sources.