AI Governance Fragmentation: Multinational Enterprise Compliance Risk

Navigating Divergent US, EU, and China Regulatory Regimes in an Era of Geopolitical Decoupling

Unofficial AI-assisted Research

Cloud Security Alliance AI Safety Initiative

2026-03-16

Executive Summary

In the span of eighteen months, the international coordination infrastructure for AI governance has effectively dissolved. The United States withdrew support from the second International AI Safety Report in February 2026 [1], announced the withdrawal from 66 international organizations in January 2026 [2], and released a national cyber strategy in March 2026 that emphasizes offensive operations and domestic deregulation over multilateral cooperation [3]. Simultaneously, the European Union began enforcing the world's most comprehensive AI regulatory regime, with high-risk AI obligations scheduled to become fully enforceable on August 2, 2026 [4]. China, for its part, embedded AI governance directly into national law through amendments to its Cybersecurity Law effective January 1, 2026, and published an ambitious national AI deployment strategy aimed at 70% AI penetration across key sectors by 2027 [26].

The consequences for multinational enterprises are substantial and immediate. Organizations that operate AI systems across US, EU, and Chinese jurisdictions face compliance requirements that are not merely different in degree but different in kind. The EU framework imposes binding obligations on any AI system affecting EU residents regardless of where the deploying organization is headquartered. The US federal government has moved to suppress state-level AI laws through executive action while simultaneously reducing the institutional capacity of the Cybersecurity and Infrastructure Security Agency (CISA), which lost roughly one-third of its workforce over the past year [6]. China's approach layers cybersecurity obligations, algorithmic recommendation controls, and emerging anthropomorphic AI regulations on top of national security obligations that have no equivalent in Western jurisdictions.

The Cloud Security Alliance's own research confirms that organizations are struggling with this landscape: 50% of surveyed IT and security professionals identified regulatory compliance as a top challenge in AI security, while only 26% reported having comprehensive AI security governance policies [7]. The compliance gap is not merely an administrative inconvenience — it represents a strategic risk that could expose organizations to fines of up to 7% of global annual revenue under the EU AI Act [4], loss of market access in regulated jurisdictions, and an erosion of customer trust across geographies that hold incompatible views of AI accountability.

This paper argues that multinational enterprises require a new governance architecture — one built not around compliance with a single dominant framework but around managing structured divergence. The paper outlines the principal regulatory blocs, analyzes the specific tensions they create, and offers a

practical framework grounded in CSA tools and guidance for building durable, multi-jurisdiction AI governance programs.

---

1. Introduction: The End of Governance Convergence

The premise that international AI governance would converge around a shared set of principles — transparency, fairness, accountability, safety — has not been borne out by events. For much of the period between 2018 and 2024, practitioners could observe genuine movement toward alignment: the OECD AI Principles gathered wide adoption, UNESCO issued its Recommendation on the Ethics of AI, and the G7 launched the Hiroshima AI Process. The EU AI Act, though far more prescriptive than these soft instruments, was developed in explicit dialogue with international norms.

What 2025 and early 2026 have revealed is that the scaffolding of convergence was always thinner than it appeared. When the geopolitical environment shifted — driven by US-China strategic competition, changing US domestic political priorities, and escalating concerns about AI's implications for national sovereignty — the facade of consensus cracked rapidly. In February 2026, the United States declined to endorse the second International AI Safety Report, a document developed by more than 100 independent experts from over 30 countries [1]. One month later, the White House released a cyber strategy that made no mention of multilateral AI governance frameworks and positioned international engagement primarily through the lens of deterrence [3].

This is not an abstract trend. For a financial services firm headquartered in New York, operating data centers in Frankfurt, and deploying AI-assisted lending decisions to customers in Singapore and Beijing, the practical question is not whether international convergence is philosophically desirable. The question is how to operate AI systems legally and responsibly across jurisdictions where the fundamental premises of legitimate AI deployment have diverged. The answer is no longer available from any single framework. It must be constructed by the enterprise itself — and the cost and complexity of doing so is already material.

This whitepaper proceeds in four analytical stages. The first describes the architecture of the three primary governance regimes. The second examines the specific operational tensions those regimes create for multinational enterprises. The third grounds the analysis in specific early-2026 events that illustrate the dynamics at work. The fourth provides actionable guidance, grounded in CSA frameworks, for building governance programs capable of navigating structured divergence.

---

2. The Architecture of Divergence: Three Distinct Governance Models

2.1 The European Model: Binding Regulation with Extraterritorial Reach

The EU AI Act, which entered into force on August 1, 2024, represents the most comprehensive attempt to impose legal obligations on AI systems at the level of national law, applied with global jurisdictional reach. The Act employs a risk-tiered classification system. AI systems presenting unacceptable risk — such as social scoring by public authorities or real-time biometric surveillance in public spaces — are prohibited outright. High-risk AI systems, including those used in employment screening, credit decisions, education access, and law enforcement, face the most extensive obligations: conformity assessments, technical documentation, human oversight requirements, data governance standards, and registration in a public EU database [4].

The compliance deadline that concentrates enterprise attention in 2026 is August 2, when the obligations for Annex III high-risk AI systems become enforceable. The penalty structure is not trivial. Violations of the prohibited-practice provisions can attract fines of up to €35 million or 7% of total global annual turnover, whichever is higher [4]. For an enterprise with €1 billion in annual revenue, that ceiling represents a potential €70 million liability. Even less severe infractions — supplying incorrect information to authorities, for instance — can generate fines of €7.5 million or 1% of global turnover.

The extraterritorial dimension of the Act makes it a global compliance imperative, not a regional one. The Act's obligations apply to providers placing AI systems on the EU market or putting them into service in the EU, and to deployers of AI systems located within the EU, regardless of where the provider is established. An enterprise headquartered in the United States that deploys an AI system affecting EU residents must comply with the Act's requirements, and this includes US-based cloud service providers operating the underlying infrastructure. The Act thus extends the EU's regulatory authority well beyond its geographic boundaries in a manner consistent with the GDPR's jurisdictional approach, which has been tested and upheld in multiple enforcement actions.

2.2 The US Model: Federal Deregulation, State Fragmentation, and Capability Erosion

The United States presents the most internally contradictory AI governance environment of the three major blocs. At the federal level, the current administration has explicitly chosen not to enact comprehensive AI regulation, preferring instead to direct agencies to streamline AI-related requirements. The December 11, 2025 Executive Order "Ensuring a National Policy Framework for Artificial Intelligence" went further: it established an AI Litigation Task Force within the Department of Justice to challenge state-level AI laws, and created mechanisms to condition federal broadband funding on states' compliance with federal AI policy positions [8]. The executive order specifically named Colorado's AI Act and California's AI transparency requirements as targets, signaling that Washington views state-level AI regulation as an obstacle to national competitiveness.

The irony of the federal preemption effort is that it has not replaced the regulatory complexity it is designed to reduce. As of March 2026, there is no comprehensive federal AI regulatory regime [9]. The result is not a deregulated environment — it is a fragmented one in which organizations must simultaneously navigate dozens of state-level laws of varying scope and rigor, respond to voluntary federal standards issued by agencies like NIST, and comply with sector-specific rules from financial, healthcare, and transportation regulators. Organizations operating in California, Colorado, Texas, Illinois, and New York face meaningfully different obligations for AI-assisted employment decisions, consumer disclosures, and algorithmic accountability, with no federal floor to establish a common baseline.

Compounding this structural problem is the erosion of the federal institution most directly responsible for helping enterprises understand and manage cyber risk in AI environments: CISA. Over the past year, CISA has lost approximately one-third of its workforce [6]. The Trump administration's proposed fiscal year 2026 budget called for cutting $495 million from CISA and eliminating nearly 30% of its positions, reducing the agency from 3,292 to 2,324 employees [10]. Congress ultimately approved a compromise cut of $135 million — a 4.6% decrease — bringing total CISA funding to approximately $2.7 billion [11]. The cuts have terminated specific programs with direct relevance to enterprise AI security: the Cyber Defense Education and Training program ($45.4 million eliminated), the Joint Collaborative Environment ($36.5 million eliminated), and the NRMC's critical infrastructure security planning activities ($67.3 million eliminated) [10].

The March 2026 National Cyber Strategy does articulate ambitions for AI-accelerated network defense and offensive cyber operations [3], but the strategy document does not resolve the gap between declared ambitions and institutional capacity. The central tension in US cyber governance is now between a strategy that envisions expanding offensive and defensive AI capabilities and the simultaneous dismantling of the agencies and programs that would execute those capabilities. For

multinational enterprises accustomed to relying on CISA guidance, advisories, and collaborative programs as an input to their own security governance, the agency's reduced capacity constitutes a direct operational impact — not merely a policy concern.

2.3 The Chinese Model: State-Directed Innovation with Layered Controls

China's approach to AI governance differs from both the European and American models in a fundamental respect: it is not primarily designed to balance innovation against harm, or to protect individual rights against algorithmic systems. It is designed to advance national strategic objectives while maintaining party and state control over AI-enabled information environments.

The layered character of China's approach is evident in the regulatory stack as it stood at the beginning of 2026. The State Council's AI Plus Action Plan, issued August 27, 2025, sets targets for AI penetration across six strategic sectors — science and technology development, industrial utilization, consumer services, public welfare, governance and security, and international collaboration [5]. The Cybersecurity Law amendments, effective January 1, 2026, bring AI explicitly within China's national law for the first time, establishing obligations for AI risk assessment and security governance [12]. Separate sectoral regulations govern algorithmic recommendation systems, generative AI services, and deep synthesis (deepfake) technologies, each administered by the Cyberspace Administration of China. The Cyberspace Administration published draft measures for anthropomorphic AI services in late 2025 — focused on human-like emotionally interactive AI — with public comment concluded in January 2026 [13].

For multinational enterprises, China's framework presents compliance challenges that differ categorically from those of the EU. EU compliance obligations are demanding but transparent: they can be assessed against a published text, and they derive legitimacy from democratic processes with rights of appeal. China's obligations intersect with national security laws — including the 2017 National Intelligence Law, which requires Chinese citizens and organizations to "support, assist, and cooperate" with national intelligence work — in ways that create structural conflicts with data governance obligations in other jurisdictions. An AI system that is lawfully operated in China, including its data processing and model behavior requirements, may violate EU data protection rules, US export controls, or confidentiality obligations to European or American customers.

This is not a hypothetical tension. It has already materialized in the cloud storage, semiconductor, and telecommunications sectors, and it is now propagating into AI services. Organizations that deploy AI systems across both Chinese and Western jurisdictions face a structural conflict that cannot be resolved by compliance engineering alone — it requires strategic decisions about architectural separation and data residency that have significant operational and cost implications.

---

3. The Compliance Chasm: Operational Challenges for Multinational Enterprises

3.1 Conflicting Requirements and the Dual-Track Compliance Problem

The most direct manifestation of governance fragmentation is conflicting compliance requirements. Consider an enterprise deploying an AI-assisted human resources system that screens job applicants. Under the EU AI Act, this is a high-risk AI system subject to Annex III obligations: conformity assessments before deployment, technical documentation, automatic logging of decisions, a post-market monitoring system, and mechanisms for human review of significant decisions [4]. Under China's Cyberspace Administration guidelines for algorithmic recommendations, a system that affects the delivery of employment opportunities must comply with separate labeling, explainability, and user right-to-opt-out requirements. Under US state law — California's AB 2930, for instance, which was among the measures targeted by the December 2025 executive order — the enterprise would have faced impact assessment requirements that may be enjoined, or may survive federal challenge.

These requirements do not merely overlap — they conflict. The EU's conformity assessment process and the Chinese security assessment process may require disclosure of technical documentation about the AI system's design that creates mutual vulnerabilities: disclosing architecture details to Chinese authorities could compromise IP protections relevant in the EU, while EU transparency obligations could result in disclosures that Chinese authorities regard as requiring their own approval. The enterprise cannot fully satisfy one regulatory regime without creating exposure in another.

Research indicates that this problem is not a theoretical edge case. Approximately 90 countries have established national AI strategies or formal governance frameworks, and at least 33 countries have enacted binding AI-specific legislation as of 2026 [9]. Analysts project that by 2028, large enterprises will require multiple distinct governance software products to manage fragmented compliance obligations [9]. The EU requirements are likely to function as the de facto ceiling for global operations among enterprises that treat compliance as a minimum floor — but satisfying EU requirements does not immunize an organization from Chinese data localization obligations, US state algorithmic accountability laws, or India's emerging IT rules for AI-generated content.

3.2 Enforcement Asymmetries and Regulatory Arbitrage Risks

The three major governance blocs differ not only in their substantive requirements but in their enforcement intensity and mechanisms. The EU AI Act creates enforcement rights in EU national market surveillance authorities, with the European AI Office overseeing GPAI models at the EU level. EU enforcement has a track record — developed through GDPR enforcement — of imposing large fines on major technology companies. The EU AI Act introduces specific obligations on General Purpose AI model providers, including foundation model providers that are not themselves deploying AI applications, which extends the regulatory perimeter to include the upstream AI supply chain.

China's enforcement is administered through a combination of the Cyberspace Administration, the Ministry of Industry and Information Technology, and national security agencies, with enforcement actions that are largely non-transparent and which may be influenced by factors beyond technical compliance. In practice, foreign enterprises have found Chinese enforcement to be selective: periodic, intense actions against specific categories of non-compliance, with significant discretionary authority vested in administrative bodies. This creates a different kind of risk than EU enforcement — less predictable, less legible, and entangled with diplomatic considerations.

US enforcement at the federal level is currently oriented toward reducing rather than increasing AI-related regulatory burden. The administration's posture is to challenge state enforcement as well as federal regulation, creating a period of unusual regulatory uncertainty where the compliance obligations that existed twelve months ago may be enjoined, pending legal challenge, or superseded by federal action. This is not equivalent to a deregulated environment — civil litigation, FTC enforcement under existing consumer protection authorities, and state attorneys general remain active — but it introduces a volatility that complicates long-term compliance planning.

3.3 The AI Security Governance Gap

Beyond the legal compliance dimension, governance fragmentation creates a security governance gap that is becoming a distinct enterprise risk. When national cybersecurity agencies operate according to divergent threat models and publish advisories that reflect different priorities, enterprises that integrate guidance from multiple jurisdictions receive inconsistent signals.

The CNCERT/CC advisory on OpenClaw, issued March 10, 2026, offers a concrete illustration [14]. China's national internet emergency response team issued a warning on AI agent vulnerabilities — specifically prompt injection attacks enabling data exfiltration, misoperation risks leading to accidental data deletion, and malicious skill uploads through repositories like ClawHub — in advance of comparable advisories from Western cybersecurity authorities [14]. The advisory reflects a legitimate and technically

accurate threat assessment. But it was issued by an authority whose enforcement perimeter, threat actor attributions, and recommended mitigations differ from those that CISA, ENISA, or the UK NCSC would produce for the same underlying vulnerabilities.

For a multinational enterprise operating AI agents that traverse multiple jurisdictions, the practical question is which advisory framework governs their security program. The answer is all of them, applied in parallel, with the enterprise itself responsible for resolving inconsistencies. This is a non-trivial operational requirement that currently falls entirely on enterprise security and compliance teams, without meaningful international coordination to reduce the burden.

---

4. Case Studies in Cross-Regime Friction

4.1 Google's $32 Billion Wiz Acquisition: Multi-Jurisdictional Regulatory Navigation

On March 11, 2026, Google completed its $32 billion acquisition of Wiz, a cloud and AI security platform — Google's largest acquisition in its history [15]. The deal's regulatory journey illustrates the multi-jurisdictional compliance architecture that major technology transactions must now navigate. Google and Wiz reached agreement in March 2025. US regulatory approval followed in November 2025. EU approval came in February 2026 [15]. The twelve-month gap between agreement and EU clearance was not unusual for a transaction of this scale and in a sector with significant platform considerations, but it illustrates the distinct assessment frameworks that US and EU regulators apply to technology consolidation.

The strategic implications for AI security extend beyond the transaction itself. Google has committed to maintaining Wiz's multi-cloud compatibility — Wiz products will continue to operate across AWS, Azure, and Oracle Cloud [15] — a commitment that reflects both commercial reality and regulatory sensitivity about platform lock-in. For enterprises evaluating AI security tooling, the Wiz-Google combination concentrates significant cloud security capabilities within Google's ecosystem while the enterprise market's governance concerns about single-vendor dependence remain acute. The transaction is also occurring precisely when the US federal government has reduced its own institutional capacity to guide enterprise security programs, meaning enterprises are more dependent on vendor-provided security capabilities at the same moment that vendor consolidation is increasing.

From a governance fragmentation perspective, the Wiz acquisition demonstrates that even the consolidation of enterprise security tooling is subject to divergent regulatory evaluation: US antitrust analysis focuses on different competitive considerations than EU antitrust analysis, and a transaction that US regulators approved without conditions required additional months of European review under a different analytical framework.

4.2 CNCERT's OpenClaw Advisory: Divergent Approaches to AI Agent Security

The CNCERT/CC advisory on OpenClaw, published March 10, 2026, deserves attention not only for its technical content but for what it reveals about the divergence in national AI security governance approaches [14]. CNCERT/CC identified a specific set of AI agent vulnerabilities — including prompt

injection enabling system key exfiltration, misoperation risks causing accidental deletion of critical data, and malicious skill uploads through the ClawHub repository — and issued guidance framing these as significant threats to finance, energy, and other critical infrastructure sectors. The advisory was issued by the same authority that administers China's cybersecurity incident response framework and that operates within the governance structure of the Cyberspace Administration.

The advisory's technical content is substantively consistent with vulnerability analyses published by Western security researchers on AI agent architectures. The recommended mitigations — restricting default management port exposure, implementing identity authentication, isolating operating environments through containerization — align with what CISA or ENISA might recommend. But the advisory reflects the priorities, threat models, and disclosure frameworks of a Chinese national authority, operating within a regulatory structure that imposes separate obligations on enterprises for incident reporting, data localization, and security assessment.

For an enterprise operating OpenClaw deployments across Chinese and Western environments, the practical implications differ by jurisdiction. In China, the advisory may trigger specific compliance obligations under the Cybersecurity Law amendments effective January 2026, including potential security review requirements. In the EU, the same vulnerabilities might be analyzed under the AI Act's high-risk system requirements or the NIS2 Directive's cybersecurity obligations. In the United States, in the absence of prescriptive federal AI security guidance, enterprises must derive implications from CISA guidance (which is now less comprehensive given workforce reductions), industry standards, and contractual obligations.

4.3 CISA's Institutional Decline and the Federal Guidance Vacuum

The reduction in CISA's institutional capacity is not primarily a budgetary concern — it is a governance concern with direct enterprise implications. CISA has functioned as the primary channel through which the US federal government translates threat intelligence into actionable guidance for enterprise security programs. The agency's Joint Collaborative Environment, Cyber Defense Education and Training program, and critical infrastructure security planning activities were not administrative overhead. They were the mechanisms through which enterprise security teams and sector-specific ISACs received synthesized guidance on emerging threats, including AI-related vulnerabilities.

The elimination of $45.4 million from CISA's Cyber Defense Education and Training program, $36.5 million from the Joint Collaborative Environment, and $67.3 million from NRMC critical infrastructure security planning [10] represents the removal of specific programs that enterprises in critical infrastructure sectors depended on for AI governance inputs. The $135 million congressional reduction that resulted from compromise — significantly less than the administration's proposed $495 million cut —

still leaves CISA with substantially reduced capacity [11]. The agency's counter-ransomware initiative has been terminated [6]. Expertise in AI security, secure software development, and critical infrastructure protection has been reduced through the attrition of one-third of the workforce.

The consequence is a federal guidance vacuum that enterprises cannot individually fill. Enterprises that previously relied on CISA advisories, CISA-facilitated sector exercises, and CISA coordination with sector-specific agencies as inputs to their AI governance programs must now either replicate that function internally, rely more heavily on commercial providers, or reduce the rigor of their AI security governance. None of these is a satisfactory outcome. The private sector cannot reconstruct the cross-sector intelligence synthesis that a well-funded CISA provides, and commercial providers have commercial incentives that may not align with comprehensive threat intelligence sharing.

---

5. Emerging Frameworks and Partial Harmonization

5.1 The Residual Multilateral Infrastructure

Despite the US withdrawal from 66 international organizations [2] and its refusal to endorse the International AI Safety Report [1], the multilateral AI governance infrastructure has not collapsed entirely. The OECD AI Principles, endorsed by 38 member states plus additional adherents, continue to provide a soft-law baseline that informs national legislation in multiple jurisdictions. The EU AI Act was developed in explicit dialogue with OECD principles, and significant elements of the Act's risk-tiered approach can be traced to that analytical lineage.

The second International AI Safety Report, published February 3, 2026, was completed with contributions from more than 100 independent experts representing more than 30 countries and international organizations, including the EU, OECD, and UN agencies [1]. The report's completion without US endorsement represents both a governance failure — the world's leading AI developer declining to formally participate in the most comprehensive international AI safety analysis — and a governance resilience story, in that the international community continued the work without US leadership.

For enterprises navigating fragmentation, these residual multilateral frameworks provide a useful reference. Compliance programs anchored to the OECD AI Principles or the ISO/IEC 42001:2023 AI management systems standard can provide a defensible baseline across jurisdictions that may not be perfectly aligned with any single national regime but that is broadly coherent with the principles underlying most national frameworks. This approach is analogous to the risk-management architecture that enterprises developed during the fragmented early years of GDPR implementation — where anchoring to a principled framework rather than literal rule-by-rule compliance in each jurisdiction provided both defensibility and operational manageability.

5.2 The ISO/IEC 42001 Bridge

ISO/IEC 42001:2023, the international standard for AI management systems, has emerged as the most credible cross-jurisdictional compliance bridge available to multinational enterprises. The standard establishes requirements for an AI management system — governance structures, risk assessment processes, impact assessments, operational controls — that are jurisdiction-agnostic while mapping closely to the substantive requirements of the EU AI Act, the NIST AI RMF, and national frameworks in several other jurisdictions.

CSA's own guidance explicitly recognizes this. The AI Risk Management document "Thinking Beyond Regulatory Boundaries" references ISO/IEC 42001:2023 as the primary standard for AI management system auditing, alongside the NIST AI RMF and the EU AI Act [16]. The STAR for AI program, which enables organizations to demonstrate AI governance maturity through the AI-CAIQ self-assessment and independent certification, maps directly to AICM control domains that themselves reflect ISO/IEC 42001 architecture [17].

Enterprises that build their AI governance programs on ISO/IEC 42001 are not immunized from jurisdiction-specific compliance obligations — EU conformity assessments, Chinese security assessments, and US sector-specific requirements still apply — but they are positioned to demonstrate principled governance to regulators in multiple jurisdictions, which reduces enforcement risk and provides a defensible record of due diligence.

---

6. Enterprise Risk Management in a Fragmented Landscape

6.1 A Framework for Structured Divergence

The governance challenge for multinational enterprises is not to achieve full compliance with every applicable framework — an increasingly unachievable goal as the number of binding AI regulatory frameworks grows. It is to manage divergence in a structured way that reduces legal exposure, maintains operational continuity, and demonstrates good-faith governance to regulators across jurisdictions. This requires a governance architecture with at least three distinct layers.

The first layer is inventory and classification. An enterprise cannot manage compliance obligations it cannot see. AI systems must be inventoried, classified by risk tier against the most demanding applicable framework (which will typically be the EU AI Act for any system affecting EU residents), and mapped to the specific compliance obligations they trigger in each relevant jurisdiction. This is an ongoing operational function, not a one-time assessment, because AI systems are modified, repurposed, and extended in ways that can change their risk classification without formal change management oversight.

The second layer is jurisdictional architecture. Some conflicts between regulatory regimes cannot be resolved through compliance engineering — they require architectural separation. An enterprise that must comply with both EU data protection obligations and Chinese data localization requirements for the same data set may have no option except to operate separate systems, with separate data stores, under separate governance processes. This is expensive and operationally complex, but the alternative — attempting to satisfy conflicting obligations with a single system architecture — generates legal exposure in both jurisdictions. The jurisdictional architecture layer defines where data resides, which systems process which data, and which regulatory regimes apply to which system components.

The third layer is a horizon-scanning and regulatory intelligence function. The AI regulatory landscape is changing faster than annual compliance review cycles can accommodate. The EU AI Act's full enforcement timeline extends through August 2026 and beyond, with GPAI model obligations, designation procedures for high-risk AI systems, and Code of Practice development all active in 2026. US state AI laws are proliferating even as federal preemption efforts proceed through courts. China's draft anthropomorphic AI regulations will reach final form in 2026. India's IT Rules amendments have

already taken effect. An enterprise without a dedicated function monitoring these developments will find itself in reactive compliance mode — responding to enforcement actions rather than building ahead of them.

6.2 The Role of Governance Maturity

CSA's 2025 State of AI Security Survey found that organizations with comprehensive AI governance policies are meaningfully different in their security postures: they are twice as likely to be confident in protecting their AI systems, twice as likely to have adopted agentic AI (suggesting governance enables rather than constrains innovation), and three times as likely to be providing AI security training to staff [7]. These findings reinforce a case that CSA has made consistently: governance maturity is not a compliance cost — it is a capability multiplier.

The survey also found that 72% of respondents were neutral or not confident in their organization's ability to execute an AI security strategy [7]. This statistic, read alongside the governance fragmentation described in this paper, suggests that the problem is not primarily technical. Enterprises that have invested in technical AI capabilities have not always built the governance structures needed to deploy those capabilities responsibly across multiple jurisdictions. The compliance chasm is, in significant part, a governance maturity gap.

Addressing that gap requires investment in roles and processes that do not map cleanly to existing organizational structures: AI governance officers who understand both technical AI systems and regulatory frameworks across multiple jurisdictions; legal and compliance functions with genuine AI technical literacy; security teams with specific training in AI threat modeling, including frameworks like MAESTRO for agentic AI threats; and board-level engagement with AI governance risk that goes beyond quarterly dashboards.

6.3 Supply Chain and Vendor Governance in a Fragmented Environment

The governance fragmentation problem does not stop at the enterprise boundary — it propagates through the AI supply chain. Enterprises that deploy AI systems built on foundation models provided by US cloud hyperscalers are indirectly subject to those providers' compliance postures in EU, Chinese, and other jurisdictions. The Google/Wiz acquisition [15] concentrates significant cloud security capability within Google's ecosystem at exactly the moment when enterprises need maximum visibility into the security posture of their AI infrastructure. An enterprise whose AI security monitoring depends on Wiz now has a material interest in Google's compliance posture in every jurisdiction where the enterprise operates.

This is not an argument against using Wiz or against the acquisition — Wiz's commitment to multi-cloud compatibility is a meaningful governance safeguard — but it illustrates that vendor governance requires the same jurisdictional analysis as direct operations. Third-party AI providers that offer services in multiple jurisdictions must comply with the requirements of each, and an enterprise that relies on those providers inherits exposure to failures in that compliance. Due diligence on AI vendors should now include assessment of their compliance posture against EU AI Act obligations, Chinese Cybersecurity Law amendments, and applicable US sector-specific requirements, not just traditional information security controls.

---

7. Conclusions and Recommendations

The convergence of several developments in early 2026 — the US rejection of international AI safety governance, the approaching enforcement deadlines of the EU AI Act, the embedding of AI regulation in Chinese national law, and the reduction of CISA's institutional capacity — marks a structural break in the AI governance environment. Multinational enterprises can no longer manage AI governance risk by following a single leading framework and trusting that other jurisdictions will converge toward it. The frameworks are diverging, the enforcement bodies are activating, and the costs of non-compliance are material.

The following recommendations are grounded in this analysis and aligned with CSA's AI governance frameworks.

Conduct an immediate AI system inventory and jurisdictional mapping. Before the EU AI Act's August 2, 2026 enforcement deadline, organizations operating AI systems affecting EU residents must complete a comprehensive inventory of those systems, classify them against the EU AI Act's risk tiers, and identify the compliance obligations each system triggers. This inventory should be extended to Chinese and other major jurisdictions to identify conflicts requiring architectural resolution.

Build a jurisdictional architecture for AI systems handling sensitive or regulated data. For data and AI systems that must satisfy conflicting requirements across EU and Chinese jurisdictions — particularly data localization, security assessment, and surveillance-related requirements — enterprises should evaluate and, where necessary, implement architectural separation. Attempting to satisfy conflicting requirements with a single system architecture generates legal exposure in all relevant jurisdictions.

Anchor AI governance programs to ISO/IEC 42001 and CSA AICM. These frameworks provide cross-jurisdictional coherence, map to the EU AI Act and NIST AI RMF, and are designed to support independent certification through CSA STAR for AI. Governance programs anchored to these standards are better positioned to demonstrate principled compliance to regulators across jurisdictions.

Establish dedicated regulatory intelligence capacity. The AI regulatory environment is changing faster than annual review cycles can accommodate. Enterprises should establish a function — whether internal or through external advisory relationships — specifically dedicated to monitoring AI regulatory developments across the EU, US (federal and state), China, India, and other relevant jurisdictions, and translating those developments into operational compliance implications.

Reassess federal guidance dependency in US cyber governance. With CISA operating at significantly reduced capacity and specific enterprise-support programs eliminated, US enterprises that relied on CISA guidance, sector exercises, and threat intelligence sharing as inputs to their AI security governance programs must identify alternative sources. This includes deeper engagement with sector-specific ISACs, CSA working groups, and industry associations that can partially compensate for reduced federal capacity.

Apply AI vendor due diligence to jurisdictional compliance posture. Third-party AI providers — including infrastructure providers, foundation model providers, and security tooling vendors — should be assessed for their compliance posture across the major jurisdictions where the enterprise operates. Vendor concentration in AI security tooling (illustrated by the Wiz-Google acquisition) warrants particular attention to terms of service, data access rights, and multi-cloud portability commitments.

Prepare board-level AI governance reporting. CSA research indicates that organizations whose boards fully understand AI security implications are significantly more likely to have comprehensive governance policies [7]. Board-level reporting on AI governance risk should include jurisdictional analysis, compliance readiness against the EU AI Act's August deadline, and an assessment of enterprise exposure to CISA capacity reduction and federal guidance changes.

---

8. CSA Resource Alignment

The Cloud Security Alliance has developed a portfolio of frameworks, tools, and guidance that directly address the AI governance challenges described in this paper.

AI Controls Matrix (AICM) and AI-CAIQ. The AICM provides 18 control domains governing AI security, governance, and compliance across the AI lifecycle. The AI Consensus Assessments Initiative Questionnaire (AI-CAIQ) enables organizations to self-assess against these controls and publish their posture in the CSA STAR Registry. For multinational enterprises building cross-jurisdictional governance programs, the AICM provides a control vocabulary that is coherent across regulatory environments and maps to both the EU AI Act and NIST AI RMF [17].

STAR for AI Program. The STAR for AI program's Level 1 self-assessment and Level 2 independent certification provide a defensible, publicly verifiable record of AI governance posture. The Valid-AI-ted advanced processing service, available to CSA members through the STAR Registry, enables enhanced assessment validation. For enterprises seeking to demonstrate AI governance maturity to EU regulators, enterprise customers, and other stakeholders across jurisdictions, STAR for AI certification provides a recognized credential [17].

MAESTRO (Multi-Agent Environment Safety and Trust Risk Ontology). As enterprises deploy agentic AI systems — with 40% of organizations already running AI agents in production and another 31% conducting pilots or tests, per CSA's "Securing Autonomous AI Agents" survey [27] — the security challenges identified by CNCERT/CC in the OpenClaw advisory (prompt injection, misoperation risks, malicious skill uploads) become directly relevant. MAESTRO provides a threat modeling framework specifically designed for agentic AI architectures, enabling enterprises to apply structured security analysis to AI agent deployments across jurisdictions.

AI Organizational Responsibilities Series. CSA's three-part series on AI Organizational Responsibilities — covering core security responsibilities, governance and risk management, and AI tools and applications — provides practical implementation guidance for building the organizational structures required for multi-jurisdiction AI governance [18]. The RACI models and cross-cutting concern frameworks in this series are directly applicable to the governance architecture challenges described in this paper.

AI Risk Management: Thinking Beyond Regulatory Boundaries. This CSA guidance document, developed by the AI Governance and Compliance Working Group, explicitly addresses the challenge of auditing AI systems across multiple regulatory frameworks. Its methodology for holistic AI risk

assessment — covering the full AI lifecycle and addressing cross-border data regulations as a specific concern — provides a principled approach to governance that is not dependent on any single national framework [16].

Cloud Controls Matrix (CCM) v4 and STAR. The foundational CCM v4 and STAR program provide cloud security governance infrastructure on which AI-specific controls can be layered. Enterprises building AI governance programs should integrate AI-specific controls from the AICM with existing CCM v4 compliance programs to ensure that AI security is not siloed from broader cloud security governance.

"Don't Panic! Getting Real About AI Governance." This CSA working group paper provides a practical, accessible treatment of AI governance connected to established risk management practices, offering a bridge for organizations transitioning from traditional GRC frameworks to AI-specific governance structures. It is particularly relevant for enterprises confronting the AI governance maturity gap identified in CSA survey data.

---

9. References

[1] Time, "U.S. Withholds Support From Global AI Safety Report," February 2026. https://time.com/7364551/ai-impact-summit-safety-report/

[2] Focus 2030, "United States withdrawal from 66 international organizations: a new step in a disengagement initiated in 2025," January 2026. https://focus2030.org/en/united-states-withdrawal-from-66-international-organizations-a-new-step-in-a-disengagement-initiated-in-2025/

[3] White House, "President Trump's Cyber Strategy for America," March 6, 2026. https://www.whitehouse.gov/articles/2026/03/white-house-unveils-president-trumps-cyber-strategy-for-america/

[4] SecurePrivacy, "EU AI Act 2026 Compliance Guide: Key Requirements Explained," 2026. https://secureprivacy.ai/blog/eu-ai-act-2026-compliance

[5] IAPP, "Notes from the Asia-Pacific region: Strong start to 2026 for China's data, AI governance landscape," 2026. https://iapp.org/news/a/notes-from-the-asia-pacific-region-strong-start-to-2026-for-china-s-data-ai-governance-landscape

[6] TechCrunch, "US cybersecurity agency CISA reportedly in dire shape amid Trump cuts and layoffs," February 25, 2026. https://techcrunch.com/2026/02/25/us-cybersecurity-agency-cisa-reportedly-in-dire-shape-amid-trump-cuts-and-layoffs/

[7] Cloud Security Alliance, "The State of AI Security and Governance," co-published with Google Cloud, December 2025. https://cloudsecurityalliance.org/artifacts/the-state-of-ai-security-and-governance

[8] White House, "Ensuring a National Policy Framework for Artificial Intelligence," Executive Order, December 11, 2025. https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/

[9] Airia, "AI Compliance Takes Center Stage: Global Regulatory Trends for 2026," 2026. https://airia.com/ai-compliance-takes-center-stage-global-regulatory-trends-for-2026/

[10] Cybersecurity Dive, "Trump's CISA budget lays out deep job cuts, program reductions," 2026. https://www.cybersecuritydive.com/news/cisa-trump-2026-budget-proposal/749539/

[11] CyberScoop, "House committee sets CISA budget cut at $135M, not Trump's $495M," 2026. https://cyberscoop.com/cisa-budget-135m-dhs-2026-house-appropriations/

[12] Mayer Brown, "China Finalises Amendments to the Cybersecurity Law: What Businesses Need to Know Before 1 January 2026," December 2025. https://www.mayerbrown.com/en/insights/publications/2025/12/china-finalises-amendments-to-the-cybersecurity-law-what-businesses-need-to-know-before-1-january-2026

[13] China Law Translate, "Provisional Measures on the Administration of Human-like Interactive Artificial Intelligence Services (Draft)," December 2025. https://www.chinalawtranslate.com/en/chatbot-measures-draft/

[14] The Hacker News, "OpenClaw AI Agent Flaws Could Enable Prompt Injection and Data Exfiltration," March 14, 2026. https://thehackernews.com/2026/03/openclaw-ai-agent-flaws-could-enable.html

[15] TechCrunch, "Google wraps up $32B acquisition of cloud cybersecurity startup Wiz," March 11, 2026. https://techcrunch.com/2026/03/11/google-completes-32b-acquisition-of-wiz/

[16] Cloud Security Alliance, "AI Risk Management: Thinking Beyond Regulatory Boundaries," 2024. https://cloudsecurityalliance.org/artifacts/ai-risk-management-thinking-beyond-regulatory-boundaries

[17] Cloud Security Alliance, "STAR for AI Level 1 Submission Guide," 2025. https://cloudsecurityalliance.org/star/ai

[18] Cloud Security Alliance, "AI Organizational Responsibilities: Governance, Risk Management, Compliance and Cultural Aspects," 2024. https://cloudsecurityalliance.org/artifacts/ai-organizational-responsibilities

[19] Nextgov/FCW, "CISA projected to lose a third of its workforce under Trump's 2026 budget," June 2025. https://www.nextgov.com/cybersecurity/2025/06/cisa-projected-lose-third-its-workforce-under-trumps-2026-budget/405726/

[20] International AI Safety Report, "International AI Safety Report 2026," February 3, 2026. https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026

[21] MeitY (Ministry of Electronics and Information Technology), "India AI Governance Guidelines," November 5, 2025. https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc2025115685601.pdf

[22] Axios, "Trump calls for more offensive cyber operations, streamlined regulations in new cyber strategy," March 6, 2026. https://www.axios.com/2026/03/06/trump-administration-cyber-strategy-2026

[23] Wiz Blog, "It's Official: Wiz Joins Google!" March 11, 2026. https://www.wiz.io/blog/google-closes-deal-to-acquire-wiz

[24] Legalnodes, "EU AI Act 2026 Updates: Compliance Requirements and Business Risks," 2026. https://www.legalnodes.com/article/eu-ai-act-2026-updates-compliance-requirements-and-business-risks

[25] South China Morning Post, "China issues second warning on OpenClaw risks amid adoption frenzy," 2026. https://www.scmp.com/tech/tech-trends/article/3346138/china-issues-second-warning-openclaw-risks-amid-adoption-frenzy

[26] State Council of the People's Republic of China, "AI Plus Action Plan (人工智能+行动计划)," August 27, 2025. http://english.www.gov.cn/

[27] Cloud Security Alliance, "Securing Autonomous AI Agents," February 2026. https://cloudsecurityalliance.org/artifacts/securing-autonomous-ai-agents

---

This document was produced by the Cloud Security Alliance AI Safety Initiative. It represents a point-in-time analysis as of March 16, 2026. Readers should verify regulatory developments independently, as the AI governance landscape is evolving rapidly.