CSAI Foundation | Cloud Security Alliance

Global AI Governance Divergence: Compliance Bifurcation

Navigating the Regulatory Chasm Between US Deregulation, EU Mandates, and Chinese Standards

2026-05-19

AI-assisted Rapid Research

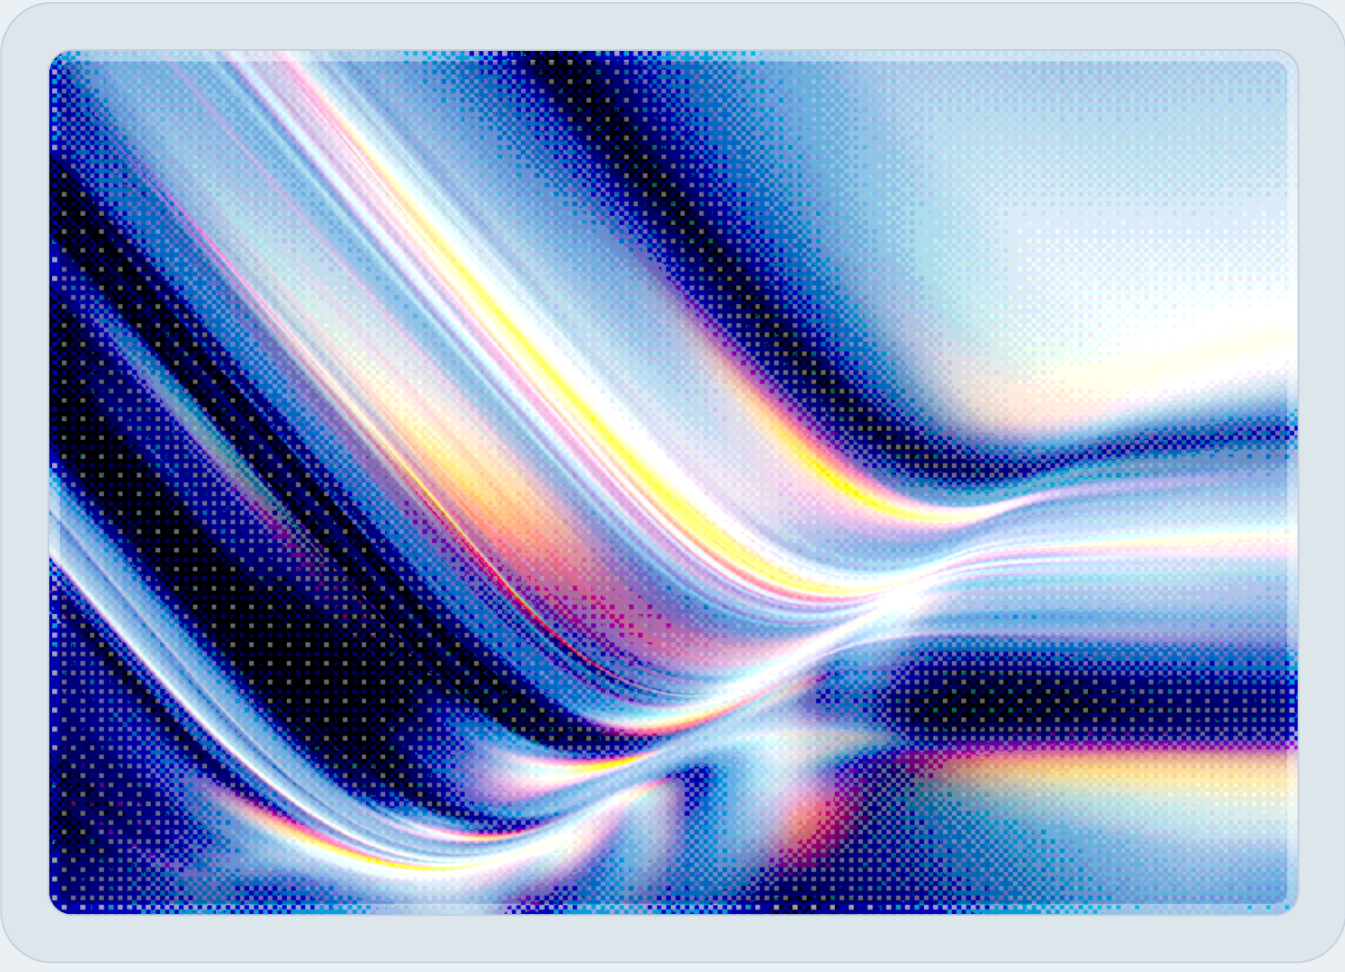

CSAI | CSA cloud security alliance®

Table of Contents

Executive Summary

The global AI regulatory landscape has fractured into three structurally incompatible regimes. In the United States, a deregulatory pivot led by successive executive orders dismantled the prior administration's AI safety requirements, consolidated rulemaking authority at the federal level, and framed AI governance primarily as a competition instrument against China. In the European Union, the AI Act is entering its most consequential phase: general-purpose AI model obligations became legally enforceable in August 2025, and the Commission's full enforcement powers over GPAI providers activate on August 2, 2026, backed by penalties reaching €35 million or 7% of global annual turnover [1]. In China, a dual-drive governance model combining binding law with detailed technical standards has been accelerating rapidly, producing more sector-specific AI regulations between 2021 and 2025 than any other jurisdiction in the world [2], while simultaneously advancing a global standards agenda through ISO, IEC, and the International Telecommunication Union.

These three regimes are not merely different in their details—they reflect fundamentally different theories of AI risk, different assumptions about who should bear compliance burdens, and different relationships between AI governance and national strategic interest. For cloud service providers, enterprise AI operators, and multinational organizations building AI systems that must function across all three jurisdictions, the practical consequence is that unified compliance programs are no longer feasible. An AI system classified as limited-risk under US guidance may face high-risk obligations under the EU AI Act, while also triggering China's algorithmic filing requirements and training data security standards. Compliance teams face the prospect of designing three parallel governance frameworks, or making deliberate jurisdictional choices about where specific capabilities are deployed.

The security implications extend beyond compliance cost. Divergent frameworks create incentives for regulatory arbitrage, complicate supply chain accountability, and introduce ambiguity into AI export controls at precisely the moment when AI systems are becoming instruments of strategic competition. Security professionals responsible for AI governance programs need to understand not only the content of each regime, but how the gaps between them create exposures that none of the regimes individually anticipates.

Introduction: Three Regimes, One Market

For most of the 2010s, AI governance was an aspirational exercise. Governments published principles, frameworks, and voluntary guidelines at a pace that far exceeded their ability—or willingness—to translate those commitments into enforceable obligations. The pattern changed materially in 2021 when China began

enacting binding sector-specific AI regulations, and again in 2024 when the EU AI Act entered into force as the world's first comprehensive AI law. The US departure from Biden-era AI safety requirements in January 2025 did not simplify this landscape; it added a third, structurally distinct posture to an already complex international environment.

The organizations most exposed to this complexity are those operating at scale across multiple jurisdictions: hyperscale cloud providers, AI platform companies, enterprise software vendors deploying AI-embedded products in global markets, and regulated industries—financial services, healthcare, critical infrastructure—where sector-specific AI requirements layer on top of horizontal AI law. For these organizations, the question is no longer whether AI regulation affects them, but how to build governance architectures capable of simultaneously satisfying requirements that, in some cases, directly conflict with one another.

This paper examines each of the three primary regulatory regimes in depth, analyzes the compliance gaps and security implications that emerge at their intersections, surveys the middle-ground positions adopted by the UK, Singapore, Canada, and other jurisdictions, and concludes with practical guidance for security and governance professionals building programs that must function across all three regimes. Throughout, the analysis is grounded in the CSA AI Controls Matrix (AICM) and the MAESTRO threat modeling framework, which together provide a jurisdiction-neutral foundation for AI security governance that can absorb regime-specific requirements without requiring fundamental program redesign each time a regulatory development occurs.

Dimension	United States	European Union	China
Regulatory model	Deregulatory; voluntary frameworks	Risk-based; binding horizontal law	Layered: binding sectoral laws + mandatory technical standards
Primary instrument	Executive orders; NIST AI RMF (voluntary)	EU AI Act (Regulation 2024/1689)	CAC Generative AI Measures + TC260 standards (GB/T)
Risk classification	No mandatory federal system	Four-tier (prohibited / high / limited / minimal)	Content attributes + scope of public reach
Enforcement authority	None at federal AI level; state AG enforcement varies	EU AI Office (full GPAI authority from Aug 2026)	Cyberspace Administration of China (CAC) + co-regulators

Dimension	United States	European Union	China
Maximum penalty	No federal AI-specific fines	€35M or 7% global turnover	Regulatory orders; market access withdrawal
Algorithm registration	Not required	Not required (documentation to AI Office)	Required within 10 business days of launch (for public-opinion AI)
Training data obligations	No federal mandate	Copyright compliance + provenance documentation (GPAI)	GB/T 45652-2025 security specification; content screening
National security exclusion	Broad (security AI exempted from oversight)	Explicit exclusion for military/defense-only AI	Security review required for all AI, including national security applications
State/subnational overlay	Yes — CA, CO, NY, IL active legislation	Member state enforcement of national high-risk rules	Uniform national implementation
International standards push	AI export packages; chip controls	GPAI code of practice; AI standards via ISO/IEC	ISO/IEC/ITU engagement; 13-point Global AI Governance Action Plan

---

The United States: Deregulation as Competition Strategy

Executive Action and the Dismantling of Predecessor Requirements

The current US posture on AI governance is the product of a deliberate reversal executed through executive action in the first days of the Trump administration's second term. On January 20, 2025, President Trump revoked Executive Order 14110, the Biden administration's foundational AI safety directive that had

established mandatory reporting requirements for frontier AI developers, directed NIST to develop AI safety standards, and created a series of interagency coordination mechanisms centered on AI risk [3]. Three days later, on January 23, 2025, the administration signed Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence," which directed federal agencies to identify and revise regulatory frameworks perceived as constraining AI innovation [4].

The framing of EO 14179 was explicitly competitive: the administration characterized the Biden-era requirements as regulatory overreach that risked ceding AI leadership to China, and directed NIST to revise the AI Risk Management Framework to eliminate references to misinformation, DEI, and climate change—content categories deemed politically contentious rather than technically relevant to AI safety [5]. This revision represents a meaningful departure from the previous RMF's scope, which had positioned AI safety as a multidimensional concept encompassing societal as well as technical risks. The revised framework narrows the governance lens toward performance, reliability, and security in the conventional cybersecurity sense.

On July 23, 2025, the White House released "Winning the Race: America's AI Action Plan," a 90-point strategic document organized around three pillars: accelerating AI innovation domestically, building AI infrastructure at scale, and leading international AI diplomacy and security [6]. The third pillar is particularly significant for security professionals: it directs the Department of Commerce to develop more granular export controls on components and sub-systems in semiconductor manufacturing, establish an "American AI Exports Program" to displace rivals' products in international markets, and enhance chip location verification technology. The Action Plan also directs NIST's Center for AI Standards and Innovation to evaluate frontier AI systems for national security risks, positioning AI governance as an instrument of geopolitical competition rather than a technical regulatory matter.

The most expansive executive action came on December 11, 2025, when the administration signed EO 14365, "Ensuring a National Policy Framework for Artificial Intelligence." This order directed the Attorney General to establish an AI Litigation Task Force tasked specifically with challenging state AI laws deemed inconsistent with the federal policy of a "minimally burdensome national AI framework" [7]. The order also conditioned certain federal funding streams—including Broadband Equity Access and Deployment (BEAD) program funds—on state compliance with the federal framework, creating financial leverage to suppress sub-federal AI regulation [7].

The State-Federal Fault Line

EO 14365 did not produce the compliance simplification the administration sought. Governors in California, Colorado, and New York publicly declared the order would not prevent enforcement of their state AI statutes. Colorado's comprehensive AI Act remained on track for applicability beginning June 30, 2026. California and Illinois have enacted significant AI legislation that took effect at the start of 2026. The result

is a layered and contested domestic landscape in which federal deregulatory intent sits alongside a growing body of state-level AI requirements, creating compliance uncertainty that falls heaviest on organizations operating nationally.

For security practitioners, this intra-US fragmentation matters. A financial services organization operating across major US states cannot simply assume that federal deregulation eliminates its AI compliance obligations domestically. State laws addressing biometric data, automated employment decisions, and algorithmic credit scoring remain in force. The AI Litigation Task Force created under EO 14365 may ultimately narrow this gap through legal challenges, but organizations cannot defer governance program development on the assumption that litigation outcomes will resolve in the federal direction.

NIST AI RMF: Continuity Amid Revision

Despite the political revisions to its scope, the NIST AI Risk Management Framework remains the de facto technical baseline for AI governance among US federal contractors and in sectors where federal oversight cascades through the supply chain. The RMF's core structure—governing, mapping, measuring, and managing AI risk—continues to underpin agency-specific guidance and procurement expectations. An April 2026 concept note from NIST introduced a forthcoming AI RMF Profile on Trustworthy AI in Critical Infrastructure, which will establish specific risk management practices for AI-enabled capabilities in energy, finance, transportation, and related sectors [8]. SP 800-53 Control Overlays for AI and the Cyber AI Profile finalization remain in progress through 2026.

For organizations subject to Federal Risk and Authorization Management Program (FedRAMP) requirements or operating within the defense industrial base, NIST alignment remains a practical necessity regardless of the administration's broader deregulatory orientation. The security floor for federally procured AI has not been eliminated; it has been narrowed and reprioritized, with the emphasis shifting from societal risk mitigation to technical security and supply chain integrity.

The European Union: Mandatory Risk Classification and Enforced Accountability

The AI Act Architecture and Enforcement Timeline

The EU AI Act entered into force on August 1, 2024, establishing the world's first comprehensive binding AI law. Its risk-tiered structure classifies AI systems across four categories—unacceptable risk (prohibited), high-risk, limited-risk, and minimal-risk—and imposes compliance obligations that scale with the potential for

harm [9]. The Act's implementation has proceeded in phases. Prohibitions on unacceptable-risk AI practices and AI literacy obligations became applicable on February 2, 2025. Governance rules and the full suite of obligations for General Purpose AI (GPAI) model providers became legally applicable on August 2, 2025 [10].

The most consequential threshold arrives on August 2, 2026, when the European Commission's enforcement powers over GPAI providers fully activate. As of that date, the Commission—operating through the EU AI Office established within DG CONNECT—will hold authority to request documentation, conduct evaluations, require compliance measures, restrict market access, and impose fines against GPAI model providers [11]. The penalty structure is among the most aggressive of any technology regulation globally: violations involving prohibited AI practices can result in fines of up to €35 million or 7% of worldwide annual turnover, whichever is greater, while breaches of high-risk AI system requirements carry penalties of up to €15 million or 3% of worldwide annual turnover [1].

A significant modification arrived through the May 7, 2026 political agreement on the "AI Omnibus" legislative package. High-risk AI systems embedded in regulated products (Annex I of the AI Act) received an extended transition period: compliance obligations for those systems now apply from August 2, 2028, rather than the original August 2, 2026 date [12]. Annex III high-risk applications—including AI systems used in biometric identification, critical infrastructure management, employment decisions, education, essential services, law enforcement, migration control, and judicial processes—are now subject to compliance obligations beginning December 2, 2027. This extension provides additional runway for organizations deploying AI in those specific contexts, but GPAI provider obligations are not affected by the Omnibus extension and remain on the original 2026 timeline.

National Security Exclusions and Their Boundaries

The AI Act's scope explicitly excludes AI systems placed on the market or put into service exclusively for military, defense, or national security purposes, regardless of the type of entity conducting those activities [13]. This exclusion is broad in its face: an AI system used solely for national security functions falls entirely outside the Act's requirements. However, the boundary between national security applications and general-purpose deployments is not always clear. AI systems used for both law enforcement and public security purposes remain subject to the Act's requirements, and the exclusion does not extend to AI developed for national security but deployed in civilian applications.

The national security exclusion has attracted criticism from civil liberties organizations, who argue it creates a governance blind spot for state AI deployments that may have significant impacts on fundamental rights. From a security governance perspective, the exclusion also creates compliance complexity for defense technology companies that develop AI with dual-use potential: products sold to both government defense customers and commercial markets may straddle the boundary between the exclusion and the Act's requirements, requiring careful product segmentation and compliance mapping.

GPAI Obligations: Transparency, Copyright, and Systemic Risk

For GPAI model providers—organizations that develop and make available general-purpose foundation models—the EU AI Act's requirements are detailed and operationally demanding. All GPAI providers must maintain technical documentation, comply with EU copyright law in their training data practices, publish model cards summarizing training methodology and capabilities, and implement policies to enforce copyright compliance throughout their training pipelines [14].

GPAI models with "systemic risk"—defined by training compute thresholds exceeding 10^25 floating-point operations or by capability evaluations triggering the threshold—face additional obligations. These include adversarial testing programs (red-teaming), serious incident reporting to the AI Office within 72 hours, enhanced cybersecurity measures meeting "state-of-the-art" standards, and annual assessments of systemic risk [11]. The cybersecurity provisions specifically call for resilience against adversarial attacks including prompt injection, data poisoning, and model extraction attempts. For organizations operating large foundation models that may meet the systemic risk threshold, these requirements are not merely compliance exercises—they constitute an enforceable AI security standard with corresponding liability.

The EU AI Office has published voluntary codes of practice to help GPAI providers navigate these obligations, offering practical guidance on transparency, copyright compliance, and safety and security measures while enforcement powers remain in a preparatory phase [15]. Organizations that participate in and adhere to approved codes of practice may gain some procedural protection in enforcement proceedings, creating an incentive structure for proactive engagement with the EU's voluntary governance mechanisms ahead of August 2026.

China: Layered Standards, Content Control, and Global Standards Ambition

The Regulatory Architecture

China's AI governance system is architecturally distinct from both the US and EU approaches. Rather than a single horizontal law, China has built a layered regulatory structure in which binding sectoral regulations interact with a growing body of national technical standards to produce detailed, verifiable compliance requirements. Between 2021 and 2025, Chinese regulators enacted more sector-specific AI regulations than any other jurisdiction globally [2], covering algorithmic recommendations, deep synthesis (deepfake) technology, generative AI, autonomous vehicles, and AI in financial services, among other domains.

The foundational instrument for generative AI governance is the Interim Measures for the Management of Generative AI Services, issued jointly by the Cyberspace Administration of China (CAC) and six co-regulatory agencies, effective August 15, 2023 [16]. The Interim Measures apply to organizations providing generative AI services to users within China, requiring that generated content align with core socialist values, respect intellectual property rights, and not contain prohibited content categories including material endangering national unity or social stability. Providers whose services carry "public opinion attributes or the capacity for social mobilization"—a broadly construed category that encompasses any AI-generated content with wide public reach—must conduct security assessments and register their algorithms with the CAC [16]. By the end of 2024, 302 generative AI services had completed the national-level filing process [17].

The CAC's algorithm registration requirement creates a significant operational obligation for foreign providers seeking to offer services in the Chinese market. Filing requires disclosure of algorithm architecture, training data characteristics, and content moderation mechanisms—information that organizations may consider competitively sensitive. The requirement also functions as a market access control: services that do not complete filing cannot legally provide generative AI to Chinese users.

Technical Standards as Governance Instruments

The "Law + Standard" dual-drive model that characterizes Chinese AI governance transforms principled regulatory requirements into measurable, verifiable technical indicators through mandatory national standards. The National Technical Committee 260 on Cybersecurity (TC260) published TC260-003, the Basic Security Requirements for Generative AI Services, in March 2024 [18]. This standard provides detailed guidance on security assessments across three domains: training data safety, model security, and content safety evaluation. It is not voluntary guidance—compliance with TC260-003 is required to demonstrate conformance with the Interim Measures during security assessment processes.

In 2025, TC260 published GB/T 45652-2025, the Cybersecurity Technology—Generative Artificial Intelligence Pre-training and Optimization Training Data Security Specification [19]. This standard substantially extends the training data security requirements established under TC260-003, requiring data provenance documentation, content filtering, sensitive information screening, and bias detection across training corpora. For organizations using foundation models trained outside China to power services offered within China, GB/T 45652-2025 creates a documentation obligation that may require contractual cooperation from upstream model providers.

The overarching policy framework is the AI Safety Governance Framework, which TC260 originally published in 2024 and updated to version 2.0 in September 2025 [20]. Where version 1.0 articulated principles and directions, version 2.0 functions as an operational manual specifying how to govern AI across its lifecycle, at what stages particular interventions apply, and how to respond when problems arise. The

Framework 2.0 is not itself binding, but it feeds directly into the binding standard-setting process, effectively serving as the pre-standard for requirements that will become enforceable through future GB/T publications.

The State Council's August 27, 2025 issuance of the AI Plus Action Plan further elevated AI governance as a national strategic priority, establishing a blueprint for China's AI development that integrates governance requirements into the broader industrial policy framework [21]. Implementing regulations flowing from the AI Plus Action Plan are expected throughout 2026, indicating that the pace of new binding requirements will not slow.

China's Global Standards Agenda

Alongside its domestic regulatory program, China has pursued a deliberate strategy of shaping international AI standards through engagement with ISO, IEC, and the ITU. Premier Li Qiang's announcement of China's Global AI Governance Action Plan at the 2025 World AI Conference proposed a 13-point roadmap for international AI coordination, explicitly calling for accelerated work through standard-setting bodies and positioning China as a constructive participant in multilateral AI governance [22]. China issued more national AI standards in the first half of 2025 than in the prior three years combined [2], establishing a domestic standards base from which to project positions into international standardization processes.

The strategic purpose of this standards engagement is to establish technical defaults that favor Chinese AI governance approaches—particularly around content safety, data localization, and algorithm transparency requirements—in international frameworks that may otherwise be shaped by US or European preferences. For organizations navigating the global standards landscape, the emergence of competing standards blocs represents a long-term structural risk: AI products designed to comply with ISO standards developed primarily under Western influence may face additional certification requirements in jurisdictions where Chinese-influenced standards have been adopted.

Compliance Bifurcation: Where the Regimes Collide

Risk Classification Mismatches

The most operationally significant point of divergence among the three regimes is the treatment of risk. The EU AI Act's risk classification system is systematic and legally defined: specific application categories (biometric identification, employment decisions, credit scoring) are categorized as high-risk regardless of the specific AI system's technical characteristics, and compliance obligations attach to that classification. The US framework, post-2025, applies no comparable mandatory risk classification system to commercial

AI. NIST's AI RMF provides voluntary risk management guidance, but conformance is not legally required outside of specific federal procurement contexts. China's regime does not classify AI by risk tier in the EU sense; instead, it applies content requirements uniformly to all services with public opinion attributes, with additional security assessment requirements triggering based on the service's reach and sensitivity.

This divergence means that a multinational AI system assessed as minimal-risk under NIST guidance—because it is deployed in a narrow enterprise context with limited human impact—may still qualify as high-risk under the EU AI Act's categorical definitions if its function touches employment screening, credit assessment, or another Annex III category. The same system, if offered as a service to Chinese users, may trigger CAC registration and TC260-003 compliance regardless of its risk profile under either Western framework. There is no single compliance mapping that resolves these overlaps; organizations must maintain parallel risk assessments calibrated to each jurisdiction's classification system.

Data Architecture and Sovereignty Conflicts

Data governance requirements across the three regimes create structural conflicts that cannot be resolved through policy alone—they require architectural choices made at the time AI systems are designed. The EU AI Act's GPAI transparency requirements mandate disclosure of training data sources and copyright compliance mechanisms, which in practice requires that training data provenance be documented and auditable. China's GB/T 45652-2025 imposes training data security requirements that include provenance documentation, sensitive data screening, and bias detection across corpora. US frameworks, by contrast, impose no comparable federal training data documentation requirements, and the administration's deregulatory orientation makes near-term federal mandates unlikely.

The conflict emerges for organizations that train models in the US (under light-touch governance), then offer those models in both EU and Chinese markets. The EU AI Act requires that training data comply with EU copyright law and be transparently documented. China's standards require training data security assessments calibrated to Chinese content standards. These requirements may be simultaneously satisfiable in practice—documentation of training data provenance serves both—but they can also conflict when training data includes content protected under one jurisdiction's copyright regime but not another's, or when data that is compliant in the US contains material that would be flagged as sensitive under China's content standards.

Algorithm Transparency and Algorithmic Filing

China's algorithmic filing requirement—under which generative AI services with public opinion attributes must register algorithm characteristics with the CAC within 10 business days of launch—has no analog in US or EU frameworks. The EU AI Act requires technical documentation for GPAI models, but that documentation is provided to the AI Office on a confidential basis and is not publicly registered. The US

does not require algorithm registration of any kind. For organizations designing AI products for global deployment, China's registration requirement creates a distinct operational process that must be planned in advance: launching a generative AI service to Chinese users without completing the CAC filing is a regulatory violation, not a grace-period issue.

The disclosure required by CAC filing—including algorithm architecture, training data characteristics, and content moderation mechanisms—raises competitive sensitivity concerns that US-developed AI companies have found difficult to reconcile. Some organizations have responded by creating separate model variants for the Chinese market, with architectures that minimize proprietary disclosure while meeting filing requirements. This approach embeds a product bifurcation decision early in the development cycle, with implications for maintenance costs, security testing, and incident response across the resulting product variants.

Multinational Compliance Program Architecture

The aggregate effect of these divergences is that multinational organizations cannot build a single AI compliance program that satisfies all three regimes simultaneously. The practical choices are to build a highest-common-denominator program—designing to EU AI Act requirements as the strictest baseline, then supplementing with China-specific filing and documentation requirements, and accepting that the resulting program will exceed what US law mandates—or to explicitly segment products and services by jurisdiction, maintaining separate governance stacks for EU-facing, US-facing, and China-facing offerings.

The highest-common-denominator approach has the advantage of governance simplicity and reduces the risk of regulatory gap-filling errors when a product's geographic reach expands. Its disadvantage is cost: EU AI Act compliance for GPAI providers involves sustained investment in technical documentation, copyright compliance infrastructure, adversarial testing programs, and incident reporting capabilities that are not legally required for US-only operations. Organizations that choose the highest-common-denominator approach are effectively subsidizing their US operations with compliance infrastructure designed for EU requirements.

Product segmentation distributes compliance costs but creates its own risks. Separate product variants maintained in parallel over time may develop security property divergences: a security fix applied to the EU variant may not be automatically propagated to the US or China variant, creating a window of differential vulnerability. Incident response across segmented products requires understanding which users are affected by which variant, adding complexity to breach notification obligations that themselves vary significantly across the three regimes.

Security Implications of Governance Divergence

Regulatory Arbitrage and Jurisdiction Shopping

Governance divergence creates structural incentives for regulatory arbitrage. When compliance with the EU AI Act's GPAI requirements imposes substantial engineering and legal costs that are not required for market access in the US or other less-regulated jurisdictions, organizations face a decision about where to locate AI development operations. The UAE, Singapore, and Gulf states have explicitly positioned their regulatory frameworks to attract AI companies seeking lighter compliance environments than the EU provides. Singapore's approach—voluntary but practically detailed frameworks without mandatory legislation—is specifically designed to offer governance structure without creating the liability exposure that binding EU requirements carry [23].

For security professionals, this arbitrage creates a different kind of risk than conventional jurisdiction shopping. When AI development migrates to environments with lighter safety and security governance requirements, the resulting systems may carry security properties that do not meet the standards required for deployment in regulated markets. If those systems are subsequently offered to EU or regulated-market users through indirect channels—via API access, embedded in third-party products, or through subsidiary arrangements—the compliance gap may not be discovered until after deployment.

Supply Chain Security and Governance Accountability

The three regimes approach supply chain security accountability differently, creating gaps in the chain of responsibility for AI safety. The EU AI Act's operator and deployer obligations create accountability downstream of the original model developer: organizations that deploy a third-party AI model in a high-risk application take on compliance obligations for that deployment even if they did not train the underlying model. This cascading accountability model creates pressure for due diligence at each link in the AI supply chain—operators must assess whether the models they deploy satisfy the Act's requirements, which requires transparency from upstream providers.

US frameworks do not create comparable downstream accountability. An enterprise that deploys a commercially available AI model in an employment decision context faces no federal obligation to assess that model's compliance with any AI-specific standard. This absence of accountability creates a dynamic in which the EU AI Act functions as a de facto global standard for supply chain due diligence: US-headquartered enterprises operating in EU markets must perform the assessments the EU requires, even if their US operations are not subject to comparable scrutiny.

China's security assessment requirements for generative AI services create a third accountability layer. Foreign AI developers whose models are used by Chinese service providers to build generative AI applications may find themselves in a chain of accountability they did not anticipate: the Chinese service provider's compliance with TC260-003 may require technical documentation from the upstream model developer about training data characteristics and model security properties. Organizations that have not designed their documentation processes to support downstream compliance obligations may face requests they cannot readily satisfy without significant retrospective engineering effort.

AI Export Controls and Technology Bifurcation

The US AI Action Plan's direction to develop new export controls on AI components and sub-systems, combined with China's own AI governance and standards agenda, is accelerating a structural bifurcation of the AI technology ecosystem along geopolitical lines [6]. This bifurcation has direct security implications: when AI systems in Western and Chinese supply chains are developed under incompatible standards, trained on different data corpora, and subject to different security requirements, the resulting systems may have security properties that are not directly comparable. Incident response, vulnerability disclosure, and coordinated remediation across this divide become substantially more difficult.

The chip export controls enacted since 2022, which restrict Chinese access to advanced GPU and semiconductor manufacturing equipment, have already begun creating divergent AI hardware development trajectories. As the US AI Action Plan directs further granularity in these controls—including at the component and sub-system level—the technical gap between AI systems developed with access to the full range of advanced chips and those developed under restricted access is likely to widen [6]. Security researchers studying AI vulnerabilities may find that findings in one ecosystem do not translate cleanly to the other, requiring parallel research programs to maintain security coverage across both.

The Middle-Ground Coalition

Principled Pragmatism from Secondary Regulators

Several major jurisdictions have positioned themselves between the US deregulatory and EU mandatory poles, offering governance frameworks that seek innovation latitude alongside meaningful accountability structures. These middle-ground approaches are significant because they represent regulatory options that technology companies may find more accommodating than either extreme, and because several of these jurisdictions exercise substantial AI supply chain influence.

The United Kingdom announced a pro-innovation, outcome-based approach to AI governance in 2024, focusing on the characteristics of adaptivity and autonomy as guides for domain-specific regulatory interpretation rather than adopting the AI Act's categorical risk taxonomy. The UK's approach assigns existing sectoral regulators responsibility for AI oversight in their respective domains, avoiding a new dedicated AI authority while maintaining accountability through existing institutional structures. This distributed model has the advantage of regulatory familiarity but creates potential coordination gaps when AI systems operate across multiple sectoral contexts.

Singapore has emerged as a particularly sophisticated example of the middle-ground approach. The Infocomm Media Development Authority (IMDA) released in January 2026 the Model AI Governance Framework for Agentic AI—described as the world's first governance framework specifically addressing the risks of multi-agent AI systems—providing detailed practical guidance on accountability, testing, and incident management without creating binding legal obligations [23]. Singapore's approach explicitly positions the city-state as an AI governance laboratory: frameworks are developed, refined through industry engagement, and offered as models for other jurisdictions to adapt. For organizations with Asia-Pacific operations, Singapore's frameworks offer compliance alignment that does not carry the legal exposure of EU requirements.

Canada's proposed Artificial Intelligence and Data Act (AIDA), which would establish a risk-based regulatory framework for high-impact AI applications, remains under parliamentary consideration and has not yet been enacted [24]. Canada faces structural pressure from both directions: close economic integration with the US creates incentives to avoid requirements that create friction with American AI providers, while domestic values alignment with European approaches to digital rights creates political pressure for mandatory accountability structures. The resolution of this tension will significantly affect compliance requirements for AI systems deployed across the North American market.

Conclusions and Recommendations

The Near-Term Compliance Priority

The most urgent compliance action for organizations with EU exposure is preparation for the August 2, 2026 GPAI enforcement threshold. As of that date, the EU AI Office will hold full authority to investigate, demand documentation, and impose penalties against GPAI model providers. Organizations that have not established technical documentation processes, copyright compliance mechanisms, and adversarial testing programs should treat the next three months as an implementation sprint, not a planning phase. The AI

Office has signaled that it will prioritize systemic-risk GPAI providers—those training above the 10^25 FLOP threshold—for initial enforcement attention, but organizations below that threshold are not exempt from the Act's documentation and transparency requirements.

Governance Architecture for Divergent Regimes

Security and governance professionals should design AI governance programs around a jurisdiction-neutral core that can absorb regime-specific requirements as modular overlays. The CSA AI Controls Matrix (AICM) provides such a foundation: its 18 security domains and 243 controls cover the threat categories—model manipulation, data poisoning, supply chain integrity, governance and compliance—that each of the three regimes addresses, even when the specific requirements differ [25]. An AICM-aligned program can accommodate the EU AI Act's GPAI documentation requirements, China's training data security standards under GB/T 45652-2025, and NIST AI RMF alignment as supplements to a common baseline, rather than requiring three separate program architectures.

Organizations should implement explicit jurisdiction tagging for AI systems and their components—tracking which systems face obligations under which regimes and maintaining this mapping as a living document updated when regulatory developments occur. This tagging should integrate with the change management process so that modifications to a system trigger an automated review of whether the change affects the system's compliance status in any jurisdiction. Without this integration, compliance mapping documents become stale within months of creation.

Data Architecture Decisions Made Early

Training data governance decisions made during model development have compliance consequences that cannot be easily remediated after the fact. Organizations should establish training data provenance documentation as a standard engineering practice—not a compliance afterthought—capturing data sources, licensing terms, content screening methodologies, and bias detection results at the time of data collection and curation. This documentation satisfies EU AI Act transparency requirements, supports GB/T 45652-2025 compliance in the Chinese market, and provides the evidentiary base for responding to regulatory inquiries without retrospective reconstruction.

Data residency and transfer architecture should be explicitly evaluated against the requirements of each target jurisdiction before deployment. The EU's data protection framework, China's data localization requirements under the Personal Information Protection Law (PIPL) and Data Security Law, and US export control implications for AI training data collectively constrain which data can be used to train models intended for global deployment. Legal and engineering teams need to coordinate on these constraints before training begins, not after a regulatory inquiry surfaces the issue.

Regulatory Intelligence as an Operational Discipline

The pace of regulatory change across all three regimes—the US executive orders arriving in sequence, the EU Omnibus modifications to the AI Act timeline, China's implementation regulations expected throughout 2026—means that a compliance mapping completed in January 2026 is already partially outdated. Organizations should establish regulatory intelligence as a continuous operational discipline: monitoring official sources in each jurisdiction, tracking legislative and regulatory developments, and maintaining a cadence of compliance mapping updates that keeps pace with the rate of regulatory change.

For most organizations, this monitoring is best maintained through a combination of internal regulatory affairs staff, external legal counsel with jurisdiction-specific expertise, and participation in industry associations that engage directly with regulators in each jurisdiction. The CSA's STAR for AI certification program and participation in the AICM's Consensus Assessment Initiative provide an additional channel for regulatory alignment, connecting organizational governance practices to a community of peers navigating the same divergent landscape.

Security Testing Across Regulatory Dimensions

Adversarial testing requirements under the EU AI Act's systemic-risk provisions—and the security assessment obligations under China's TC260-003—create an opportunity to build a security testing program that satisfies both requirements from a common test infrastructure. Red-teaming programs designed to meet EU AI Office expectations (covering prompt injection, data poisoning, model extraction, and jailbreaking attacks) substantially overlap with the security evaluation categories in TC260-003. Organizations that build these programs to EU standards and document results against TC260-003 criteria can use a single set of security evaluations to satisfy both regulatory requirements, reducing duplication.

NIST's forthcoming AI RMF Profile on Trustworthy AI in Critical Infrastructure, expected to finalize in late 2026, will provide additional guidance for organizations in regulated sectors [8]. Critical infrastructure operators should track this profile development and plan to align their AI RMF implementations with the profile's requirements as it reaches final publication, particularly given that federal procurement conditions are likely to reference the profile once available.

CSA Resource Alignment

The Cloud Security Alliance has developed a portfolio of frameworks and resources directly applicable to organizations navigating the global AI governance divergence described in this paper.

The **CSA AI Controls Matrix (AICM)** is the most directly applicable resource for building a jurisdiction-neutral AI governance foundation. The AICM's 18 security domains and 243 controls address the threat categories and governance requirements that all three major regulatory regimes target, providing a common baseline from which regime-specific requirements can be implemented as overlays. Organizations seeking to demonstrate AI governance maturity across multiple jurisdictions should begin AICM implementation before any single regulatory deadline forces reactive compliance, as the AICM's design anticipates cross-jurisdictional governance needs [25]. The AICM's Consensus Assessment Initiative Questionnaire (CAIQ for AI) supports both internal self-assessment and third-party vendor due diligence, directly supporting the supply chain accountability obligations that the EU AI Act imposes on operators [26].

The **MAESTRO framework** for agentic AI threat modeling addresses the multi-agent, multi-step AI architectures that are increasingly prevalent in enterprise deployment and that pose the most challenging governance questions under all three regimes. MAESTRO's layered architecture covers the model, agent, orchestration, data, and infrastructure layers that must each be assessed for risk under a comprehensive AI governance program, and provides a threat taxonomy applicable to the agentic AI architectures that Singapore's MAIGF and the EU AI Office's systemic risk provisions are beginning to address [27].

The **CSA STAR for AI** program, building on the Cloud Controls Matrix's established STAR assurance framework, provides a structured path to third-party validated AI governance assurance. For organizations seeking to demonstrate compliance to EU deployers under the AI Act's due diligence expectations, or to position their AI governance practices for prospective assessment under future mandatory assurance schemes, STAR for AI certification provides a recognized and auditable credential [28].

The **CCM (Cloud Controls Matrix)** remains foundational for the infrastructure layer of AI governance. Cloud security controls for data protection, identity and access management, infrastructure security, and incident response are prerequisites for AI-specific governance: a GPAI provider cannot satisfy EU AI Act security requirements for its model if the cloud infrastructure on which the model runs does not meet baseline cloud security standards [29]. CCM alignment ensures that AI-specific governance programs are built on a sound infrastructure security foundation rather than addressing AI risks in isolation.

References

[1] Aqua Security. "Penalties of the EU AI Act." Aqua Security, 2026.

[2] Global AI Governance and Compliance Center. "China AI Governance Framework: What Global Businesses Need to Know in 2026." GAICC, 2026.

[3] White House. "Initial Rescissions of Harmful Executive Orders and Actions." Presidential Actions, January 20, 2025.

[4] White House. "Removing Barriers to American Leadership in Artificial Intelligence." Executive Order 14179, January 23, 2025.

[5] NIST. "AI Risk Management Framework." National Institute of Standards and Technology, updated 2025.

[6] White House. "Winning the Race: America's AI Action Plan." July 2025.

[7] White House. "Ensuring a National Policy Framework for Artificial Intelligence." Executive Order 14365, December 11, 2025.

[8] NIST. "Artificial Intelligence – Standards, Frameworks, and Guidance." NIST, April 2026.

[9] European Commission. "AI Act: Shaping Europe's Digital Future." European Commission, 2024.

[10] DataGuard. "EU AI Act Timeline: Key Compliance Dates & Deadlines Explained." DataGuard, 2025.

[11] European Commission. "Guidelines for Providers of General-Purpose AI Models." EU AI Office, 2025.

[12] Council of the European Union. "Artificial Intelligence: Council and Parliament Agree to Simplify and Streamline Rules." May 7, 2026.

[13] European Commission. "Regulation (EU) 2024/1689 – Artificial Intelligence Act (Full Text)." Official Journal of the European Union, 2024.

[14] Orrick. "The EU AI Act: 6 Steps to Take Before 2 August 2026." Orrick, November 2025.

[15] EU Artificial Intelligence Act. "Responsibilities of the European Commission (AI Office)." artificialintelligenceact.eu, 2025.

[16] Library of Congress. "China: Generative AI Measures Finalized." Global Legal Monitor, July 2023.

[17] White & Case. "AI Watch: Global Regulatory Tracker – China." White & Case, 2025.

[18] Center for Security and Emerging Technology (CSET). "China: Safety Requirements for Generative AI (Final)." CSET Translation, 2024.

[19] GeopolitEchs. "What's in China's New National Standard on GAI Service Safety?." GeopolitEchs, 2025.

[20] GeopolitEchs. "China Releases Upgraded AI Safety Governance Framework to Tackle Emerging AI Risks." GeopolitEchs, 2025.

[21] State Council of the People's Republic of China. "Full Text: AI Plus Action Plan." gov.cn, August 27, 2025.

[22] Ministry of Foreign Affairs of the People's Republic of China. "Global AI Governance Action Plan." July 2025.

[23] AskAjay.ai. "How 8 Countries Regulate AI in 2026: The Executive Comparison Guide." 2026.

[24] GDPR Local. "AI Regulations Around the World: Everything You Need to Know in 2026." 2026.

[25] Cloud Security Alliance. "AI Controls Matrix." CSA, 2025.

[26] Cloud Security Alliance. "Introductory Guidance to AICM." CSA, 2025.

[27] Cloud Security Alliance. "Agentic AI Threat Modeling Framework: MAESTRO." CSA, February 2025.

[28] Cloud Security Alliance. "Introducing the CSA AI Controls Matrix: A Comprehensive Framework for Trustworthy AI." CSA, July 2025.

[29] Cloud Security Alliance. "Cloud Controls Matrix." CSA, 2024.

