CSAI Foundation | Cloud Security Alliance

Divergent AI Regulation: EU Enforcement Meets US Deregulation

Enterprise Compliance Implications of the 2026 EU-US Policy Split

2026-05-15

AI-assisted Rapid Research

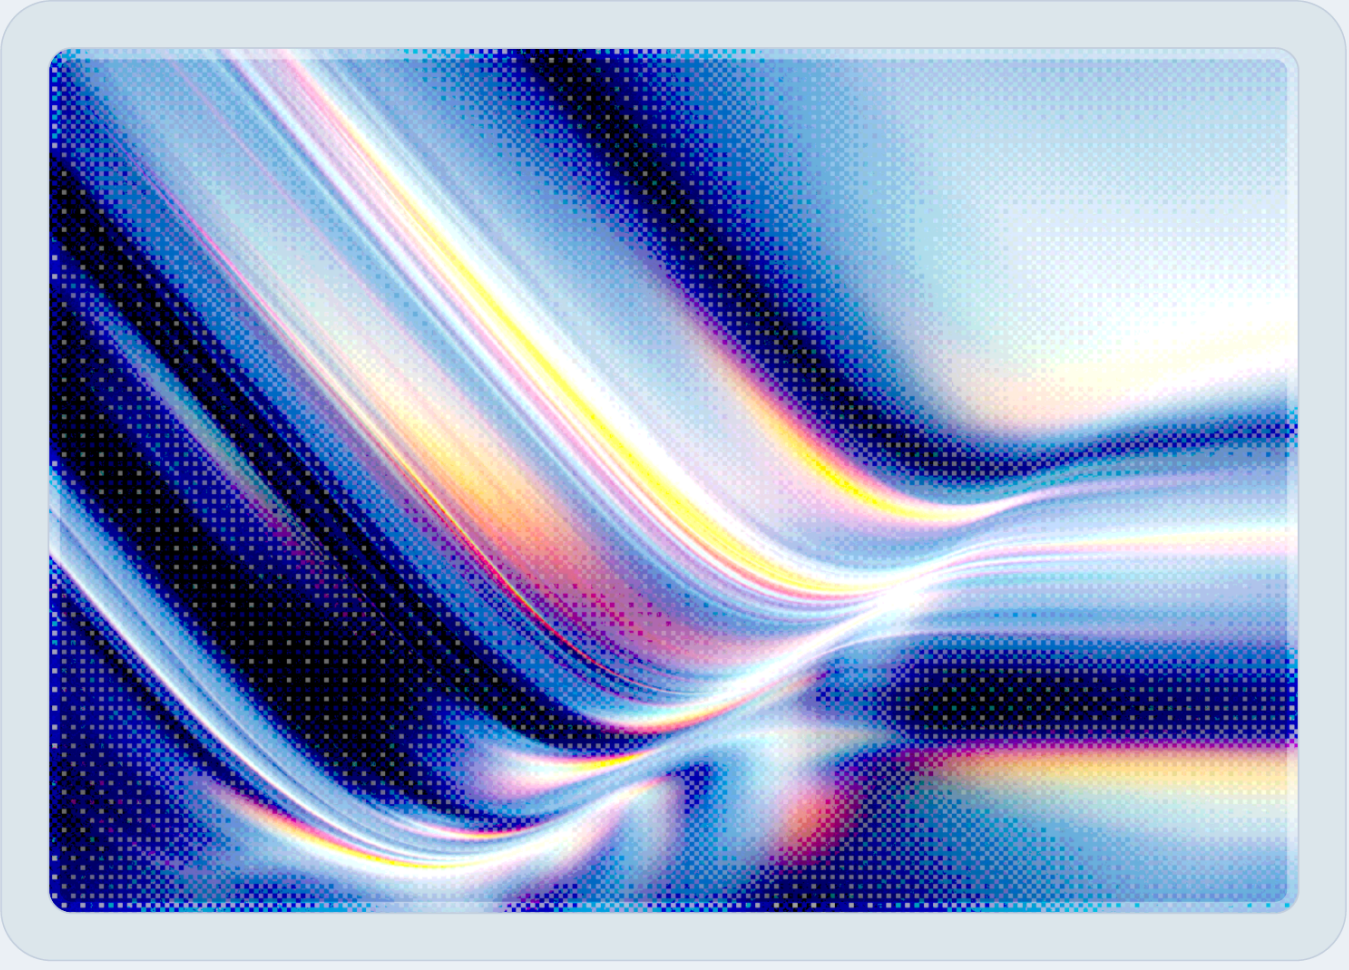

CSAI | cloud security alliance®

Key Takeaways

- The EU AI Act's August 2, 2026 enforcement deadline for high-risk AI systems carries penalties up to €35 million or 7% of global annual turnover [1]; US federal policy has simultaneously moved in the opposite direction, revoking safety mandates and actively working to preempt state-level AI laws [2][3].
- The EU AI Act's territorial scope, defined under Article 2, extends to providers placing systems on the EU market and deployers operating in EU contexts, creating compliance obligations for US-headquartered organizations regardless of where those systems are built or operated [17]; organizations uncertain about whether their specific deployment scenarios trigger these obligations should seek EU-qualified legal counsel.
- Prohibited AI practices under the EU Act—including real-time biometric identification in public spaces, social scoring, and certain predictive policing systems—have been enforceable since February 2, 2025 [5]; organizations running non-compliant systems today face active legal exposure.
- The US deregulatory pivot does not eliminate compliance obligations. It reshapes them into an increasingly fragmented landscape of state laws—California, New York, Colorado, and others—with incompatible definitions, enforcement timelines, and penalty structures [6][7].
- According to CSA's March 2026 readiness analysis, over half of organizations lack systematic inventories of AI systems currently in production, making risk classification under either regime practically infeasible without foundational preparatory work [8].
- Compliance programs designed for a single regulatory environment will not scale. Enterprises with EU market exposure and US operations need a dual-track governance architecture that satisfies EU pre-market obligations while accommodating the evolving US state patchwork.

Background

The regulatory environments governing artificial intelligence on either side of the Atlantic diverged sharply in 2025 and are now pulling in structurally incompatible directions. The European Union completed a multi-year legislative effort with the EU AI Act entering into force on August 1, 2024, establishing the world's most comprehensive binding AI regulatory framework. The Act phases in requirements over a four-year window, with each wave targeting a different tier of risk. Prohibited

practices became enforceable first, on February 2, 2025, followed by obligations for general-purpose AI (GPAI) model providers on August 2, 2025, and culminating in the binding enforcement of high-risk AI system requirements on August 2, 2026 [9]. The EU AI Office—seated within the European Commission—holds sole enforcement authority over GPAI rules and will begin exercising its full enforcement powers, including the authority to impose fines, on that same August 2026 date [10].

The United States took the opposite trajectory. On January 20, 2025, his first day in office, President Trump revoked Executive Order 14110—the Biden administration's foundational AI safety directive—through an inaugural-day omnibus rescission order [16]. That order eliminated requirements for frontier AI developers to share red-teaming results with the federal government, Chief AI Officer mandates across federal agencies, and more than one hundred agency-level safety actions. Three days later, EO 14179 ("Removing Barriers to American Leadership in Artificial Intelligence") [2] established a new federal posture explicitly oriented around AI dominance, economic competitiveness, and deregulation, directing agencies to identify and rescind rules deemed to hinder AI development [11]. The July 2025 AI Action Plan formalized this into a ninety-point policy framework organized around accelerating innovation, building AI infrastructure, and leading international diplomacy—with the Office of Management and Budget tasked with a government-wide effort to identify and repeal regulations obstructing AI deployment [12].

The December 2025 executive order on eliminating state law obstruction of national AI policy [3] added a preemptive dimension, establishing a DOJ AI Litigation Task Force charged with challenging state AI laws in federal court on commerce and preemption grounds beginning January 10, 2026, and directing the Secretary of Commerce to publish a comprehensive review of state AI laws by March 11, 2026, identifying those deemed incompatible with federal policy. Despite this federal pressure, the US Senate declined to pass a House-approved provision that would have barred states from enforcing AI-specific regulations for ten years, leaving state enforcement authority intact [6]. Governors in California, New York, and Colorado have publicly stated their intent to continue enforcing state AI statutes regardless of federal directives [6].

Security Analysis

The EU AI Act's Compliance Architecture

The EU AI Act imposes a tiered compliance architecture based on risk classification. At the apex, a small set of AI capabilities are prohibited outright under Article 5: real-time remote biometric identification in publicly accessible spaces for law enforcement purposes (with narrow exceptions), emotion recognition systems in workplaces and educational institutions, social scoring systems operated by public authorities,

and AI systems that predict crime risk based solely on profiling [5]. These prohibitions have been in force since February 2025, meaning any organization operating such systems today is already out of compliance—an exposure that does not await August 2026.

The enforcement inflection point arriving in August 2026 pertains to what the Act calls high-risk AI systems, defined primarily by reference to the applications listed in Annex III. This category encompasses AI used in employment and HR decisions, creditworthiness assessments, educational access and outcomes, administration of justice, critical infrastructure management, law enforcement, biometric categorization, and essential private and public services [4]. Providers of these systems must complete documented risk management processes, data governance reviews, conformity assessments, CE marking, and registration in the EU's AI database before placing systems on the market [13]. Deployers—organizations using high-risk AI built by third parties—must implement human oversight mechanisms, monitor performance, and maintain records adequate to demonstrate compliance with Article 26 [14]. Both roles carry independent legal obligations, and under the Act's territorial scope provisions a US enterprise using a third-party HR screening tool or credit decision engine in EU-facing contexts may qualify as a deployer within the Act's scope. The precise reach for US-operated systems that only incidentally affect EU residents involves ongoing legal interpretation, and organizations should confirm their exposure with EU-qualified counsel.

For GPAI model providers—companies offering foundation models as a service—the compliance timeline is already running. Obligations attached on August 2, 2025, giving providers one year to reach compliance before the Commission's enforcement powers activate in August 2026. The maximum fine structure for GPAI violations is €15 million or 3% of global annual turnover, whichever is greater [10][15]. For higher-tier violations involving prohibited practices or systemic non-compliance by high-risk system providers, penalties reach €35 million or 7% of global annual turnover [1]. These figures are not theoretical: the AI Office has already begun information-gathering exercises ahead of its full activation. The enforcement architecture has structural similarities to the GDPR's—in both cases, the responsible authorities built enforcement capacity and gathered information in advance of the compliance window closing. Whether EU AI Act enforcement will follow the GDPR's pattern of an extended initial ramp followed by consequential action is a reasonable expectation, but not a certainty.

The practical demands on enterprise compliance programs are significant. The Act requires technical documentation capturing a system's general description, development process, data requirements, testing procedures, lifecycle change history, and conformity assessment evidence [13]. Harmonized standards that would simplify this work have been delayed; the key standard prEN 18286 entered its enquiry phase in October 2025—eight months behind schedule—leaving most enterprises without a finished standards-based compliance pathway ahead of the August 2026 deadline [8].

The US Federal Deregulatory Posture

The US federal deregulatory posture removes rather than imposes obligations. Following the revocation of EO 14110, federal AI safety requirements that had applied to frontier AI developers—including mandatory reporting of red-team results and safety testing documentation—are no longer in effect at the federal level [2]. The AI Action Plan directs agencies to review and eliminate procurement rules that constrain AI adoption, signals that the FTC's AI-related consent decrees may be revisited, and conditions $42 billion in broadband infrastructure funding under the BEAD program on state repeal of AI regulations deemed onerous [3]. For enterprises that primarily serve US federal customers and have no EU market exposure, this creates a materially lighter compliance environment in the near term.

The security risk embedded in this posture is less regulatory and more operational. The US approach to AI governance has historically relied on a combination of sector-specific regulation (financial services, healthcare, defense), litigation risk, and voluntary industry frameworks. Without a federal floor equivalent to the EU Act's baseline requirements, organizations operating in the US are not compelled to conduct the kind of pre-market conformity assessment, data governance review, or human oversight documentation that the EU Act mandates. Systems that would require formal risk classification and mitigation in Europe can be deployed in the US with minimal governance overhead. This means that multinationals maintaining different compliance postures across geographies are effectively operating higher-risk AI pipelines in their US operations than their EU operations—a condition that can introduce material inconsistencies in safety practices, audit trails, and incident response readiness.

The State-Level Patchwork

The absence of comprehensive federal AI legislation leaves US operations subject to a growing and legally contested patchwork of state laws. California's package of AI laws taking effect January 1, 2026—including the Transparency in Frontier Artificial Intelligence Act (SB 53), training data transparency requirements under AB 2013, and AI-generated content disclosure requirements under SB 942—applies to companies that develop or distribute AI systems to California residents regardless of where those companies are headquartered [7]. New York's RAISE Act, signed December 19, 2025 and amended March 27, 2026, establishes a transparency and reporting framework for frontier AI with an effective date of January 1, 2027 [18]. Colorado's SB 205, which created high-risk AI obligations for consequential decision systems, had its enforcement stayed by a federal court pending litigation, though its core risk management requirements remain legally operative [6].

The practical compliance problem is not that any single state law is unmanageable. It is that the aggregate demands are structurally incompatible. Colorado's framework uses different risk classification criteria than California's transparency obligations; neither maps cleanly to the EU Act's Annex III

categories. Definitions of "consequential decision," "high-risk," and "deployer" vary across jurisdictions. Enterprises operating across all three states plus EU markets face at minimum four overlapping and partially conflicting compliance regimes, each with distinct documentation requirements, disclosure timing, and enforcement authority.

Jurisdiction	Framework Type	Key Obligations	Enforcement Status (May 2026)
European Union	Pre-market conformity, risk-tiered	Conformity assessment, technical docs, CE marking, EU database registration	Binding August 2, 2026
California	Transparency, disclosure	Frontier model risk frameworks, training data summaries, watermarking	Effective January 1, 2026
New York	Transparency, reporting	Safety reporting, whistleblower protections	Signed December 2025; effective January 1, 2027
Colorado	Consequential decision oversight	Risk management, impact assessments, AG reporting	Enforcement stayed by federal court
Federal (US)	Deregulatory, voluntary	No binding AI-specific requirements; sectoral rules apply	No comprehensive federal AI legislation

Recommendations

Immediate Actions

The most urgent task for any enterprise with EU market exposure is completing an AI system inventory. Without a systematic catalog of what AI systems are in production or development, including their intended purpose, the populations they affect, the decisions they inform or automate, and the data they process, risk classification under the EU Act is practically infeasible. This inventory should be treated as a

compliance prerequisite, not a compliance deliverable. Organizations that have not yet begun this work face a narrow window before the August 2026 enforcement deadline makes incomplete inventories a direct enforcement risk.

Organizations should immediately audit all AI systems deployed in EU-facing contexts against Article 5's prohibited practices list. Real-time biometric identification, emotion recognition in workplaces, and crime risk prediction based on profiling have been prohibited since February 2025. Any system in production that touches these categories should receive immediate legal review. The enforcement clock on prohibited practices is already running, and continued operation of non-compliant systems extends the period of potential exposure; organizations should seek legal guidance on how duration of non-compliance may affect penalty calculations under Article 99.

Enterprises using third-party AI in HR screening, credit underwriting, educational access, or other Annex III categories must review their vendor agreements and assess their deployer obligations under Article 26. A compliance gap worth examining is the assumption that deployer obligations are the vendor's responsibility alone—under the Act, they are not. The deployer's obligation to implement human oversight, monitor system performance, and retain documentation is independent of the provider's own compliance status.

Short-Term Mitigations

Over the next three to six months, organizations should build a parallel compliance tracking function capable of monitoring both EU AI Act implementation milestones and US state legislative developments simultaneously. The EU timeline is fixed; the US landscape is not. Colorado's enforcement stay may lift, California's laws are already in effect, and additional states are moving legislation. A compliance function that tracks only one jurisdiction introduces structural blind spots.

For AI systems that fall under the EU Act's high-risk categories, organizations should begin conformity assessment work now rather than waiting for harmonized standards to finalize. The absence of completed standards does not suspend the obligation—providers can conduct conformity assessments under Annex VII using the Act's own requirements as the reference. Technical documentation, risk management process records, and data governance documentation should be drafted against Article 11 and Annex IV requirements directly. The standards, when finalized, are likely to offer alternative pathways rather than replacing these obligations.

Enterprises should conduct a gap analysis of AI governance programs against both the EU Act's requirements and the substantive obligations in applicable state laws, using a common underlying framework to avoid building parallel compliance bureaucracies. The CSA AI Controls Matrix provides a

role-differentiated control structure that can be mapped to external regulatory obligations and serves as a practical baseline for this exercise.

Strategic Considerations

Over a twelve-month horizon, multinational enterprises need a governance architecture that can sustain different compliance postures across regions without creating incoherence in underlying AI safety practices. The EU posture demands prospective, pre-market governance: risk assessment before deployment, conformity assessment before market placement, human oversight by design. The US posture, particularly at the federal level, imposes no equivalent pre-market gate. The strategic risk—one observed in analogous regulatory contexts such as financial services and environmental compliance, where compliance-floor divergence across jurisdictions has contributed to material safety inconsistencies—is that organizations operating at the minimum permissible level in each jurisdiction may develop systematically weaker governance practices for their US AI pipelines. Those weaker practices increase the probability of incidents, enforcement actions under state laws, or EU enforcement action when systems cross jurisdictional boundaries.

The more defensible long-term approach is to treat the EU Act's substantive governance requirements—AI inventory, risk classification, data governance, human oversight documentation, and post-deployment monitoring—as a global baseline applied regardless of the deployment geography. This creates operational consistency, reduces the burden of maintaining regime-specific documentation variants, and positions the enterprise for whatever regulatory trajectory the US eventually takes. Among binding regulatory frameworks currently in force in a major economic bloc, the EU AI Act is the most comprehensive and operationally demanding; compliance with its core governance requirements addresses many—though not all—requirements under applicable US state frameworks, and organizations should conduct a gap analysis against specific state obligations rather than assuming full coverage.

Organizations should also assess their supply chains. Third-party AI vendors serving EU markets are subject to provider obligations under the Act, and deployers inherit compliance dependencies from their vendors. Vendors that fail to meet provider obligations—completing conformity assessments, maintaining technical documentation, registering in the EU AI database—expose their deployer customers to enforcement risk. Procurement and vendor management programs should incorporate AI Act compliance as a supplier evaluation criterion before the August 2026 enforcement date.

CSA Resource Alignment

The CSA AI Controls Matrix (AICM) v1.0.3 provides the most directly applicable CSA framework for organizations navigating the EU AI Act compliance obligations described in this note. The AICM is structured around eighteen control domains spanning AI supply chain security, data governance, model security, and runtime monitoring, and addresses the shared responsibility model across model providers, application providers, orchestrated service providers, and cloud service providers—roles that map closely to the EU Act's provider and deployer distinction. Organizations beginning conformity assessment preparation under the EU Act should use the AICM implementation guidelines as a starting point for their technical documentation and risk management processes. CSA has published role-specific implementation guidelines for each provider type, as well as auditing guidelines that align with the EU Act's conformity assessment structure.

The CSA Cloud Controls Matrix (CCM) addresses underlying infrastructure governance obligations that remain in scope for high-risk AI systems, particularly in areas of data governance, access control, logging, and incident response. Enterprises using cloud-hosted AI should map their CCM controls to the EU Act's Article 15 accuracy, robustness, and cybersecurity requirements, which apply to high-risk AI systems and require documented technical safeguards against adversarial inputs and distributional shift.

The CSA STAR program provides the trust assurance and third-party assessment infrastructure that enterprises will need when evaluating GPAI model providers and other AI vendors for EU Act compliance. STAR-certified cloud providers have undergone independent assessment against the CCM; extending this assessment to AI-specific controls using the AICM's AI-CAIQ questionnaire supports the deployer's due diligence obligation and provides documented evidence of supply chain compliance review.

CSA's previous research note on the EU AI Act high-risk compliance deadline [8] provides deployment-focused technical guidance on conformity assessment procedures and the Annex III classification exercise. Organizations should treat this note as a companion document addressing the technical compliance pathway, while this note addresses the broader governance architecture required when operating across both the EU and US regulatory environments simultaneously.

References

[1] European Parliament and Council. "Regulation (EU) 2024/1689 — Artificial Intelligence Act, Article 99: Penalties." EU Artificial Intelligence Act, August 2024.

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

[3] The White House. "Eliminating State Law Obstruction of National Artificial Intelligence Policy." December 11, 2025.

[4] European Commission. "AI Act — Shaping Europe's Digital Future." European Commission Digital Strategy, 2024.

[5] European Parliament and Council. "Regulation (EU) 2024/1689 — Article 5: Prohibited AI Practices." EU Artificial Intelligence Act, August 2024.

[6] King & Spalding. "New State AI Laws are Effective on January 1, 2026, But a New Executive Order Signals Disruption." King & Spalding Client Alert, January 2026.

[7] Baker Botts. "U.S. Artificial Intelligence Law Update: Navigating the Evolving State and Federal Regulatory Landscape." Baker Botts Client Alert, January 2026.

[8] Cloud Security Alliance. "EU AI Act High-Risk Deadline: Enterprise Readiness Gap." CSA Labs Research Note, March 2026.

[9] Artificial Intelligence Act. "Implementation Timeline." artificialintelligenceact.eu, 2024.

[10] European Parliament and Council. "Regulation (EU) 2024/1689 — Article 101: Fines for Providers of General-Purpose AI Models." EU Artificial Intelligence Act, August 2024.

[11] Center for Security and Emerging Technology (CSET). "The Executive Order on Removing Barriers To American Leadership In Artificial Intelligence." Georgetown CSET, February 2025.

[12] Sidley Austin LLP. "The Trump Administration's 2025 AI Action Plan." Data Matters Privacy Blog, July 2025.

[13] European Parliament and Council. "Regulation (EU) 2024/1689 — Annex IV: Technical Documentation." EU Artificial Intelligence Act, August 2024.

[14] European Parliament and Council. "Regulation (EU) 2024/1689 – Article 26: Obligations of Deployers of High-Risk AI Systems." EU Artificial Intelligence Act, August 2024.

[15] Holland & Knight. "U.S. Companies Face EU AI Act's Possible August 2026 Compliance Deadline." Holland & Knight Insights, April 2026.

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

[17] European Parliament and Council. "Regulation (EU) 2024/1689 – Article 2: Scope." EU Artificial Intelligence Act, August 2024.

[18] New York State Senate. "S7583-A: Responsible AI Safety and Education Act (RAISE Act)." New York State Legislature, 2025.