CSAI Foundation

Cloud Security Alliance AI Safety Initiative

US–EU AI Governance Divergence: Enterprise Compliance Guide

Navigating a Fragmented Regulatory Landscape in 2026

Unofficial AI-assisted Research

2026-04-06

Key Takeaways

- The EU AI Act is actively enforcing requirements in 2026: prohibited AI practices have been banned since February 2025, General-Purpose AI (GPAI) obligations entered force in August 2025, and full high-risk AI system enforcement begins August 2, 2026, with penalties up to €35 million or 7% of global annual turnover [1][2].
- The US has moved in the opposite direction. President Trump revoked the Biden administration's AI executive order on January 20, 2025, and replaced it with a deregulatory framework emphasizing innovation over oversight [3][4]. No binding federal AI statute exists, and the NIST AI Risk Management Framework remains entirely voluntary [5].
- On current trajectory, the regulatory gap appears to be widening rather than narrowing. The US and UK declined to sign the Paris AI Action Summit's inclusive AI declaration in February 2025 [6], and the US Trade Representative has signaled potential Section 301 investigations targeting the EU AI Act as an unfair trade barrier [7].
- US enterprises operating in EU markets face binding EU AI Act obligations regardless of their home jurisdiction. The Act's extraterritorial scope applies to any provider whose AI system is placed on the EU market or whose AI outputs are used within the Union [8].
- NIST AI RMF serves as a widely adopted governance bridge: its governance controls map to EU AI Act conformity obligations while satisfying US voluntary compliance postures, enabling enterprises to build dual-regime compliance on a common governance foundation [5][9].

Background

For enterprises operating across US and EU markets, the regulatory landscape for artificial intelligence has divided along two fundamentally different models over the past eighteen months. The European Union has pursued a risk-tiered statutory framework that imposes binding obligations across defined AI use categories. The United States has pursued a deregulatory posture that prioritizes national competitiveness and innovation, relying on voluntary frameworks and sector-specific guidance in place of broad-based regulation. For enterprises operating across both jurisdictions, this divergence has become a material compliance and operational risk — one that is accelerating rather than resolving.

The EU AI Act entered into force on August 1, 2024, with a phased enforcement schedule staggered across 2025 through 2027 [1]. Its first major enforcement milestone arrived on February 2, 2025, when prohibitions on the highest-risk AI practices took effect across all 27 member states. These prohibited practices include AI systems that deploy subliminal or deceptive manipulation, social scoring mechanisms used by public authorities, predictive policing systems based on profiling, and real-time remote biometric identification in publicly accessible spaces [25]. AI literacy obligations under Article 4 also entered application at the same time, requiring organizations to ensure that personnel interacting with AI systems possess adequate AI literacy — encompassing awareness of AI capabilities, limitations, and appropriate use — and can demonstrate that competency [10].

August 2, 2025, marked the second major milestone: GPAI obligations entered application, and the EU AI Office became fully operational as the primary enforcement body for foundation model providers [11] [12]. All providers of general-purpose AI models — including large language models, multimodal systems, and foundation models offered via API — are now subject to transparency obligations, technical documentation requirements, and, for models deemed to pose "systemic risk" (those trained on compute exceeding 10^{25} FLOPs), enhanced requirements including adversarial testing, incident reporting, and cybersecurity safeguards [13].

In the United States, the trajectory has been sharply different. On January 20, 2025, President Trump revoked Executive Order 14110, the Biden administration's "Safe, Secure, and Trustworthy AI" order that had directed agencies to develop AI safety standards and mandated red-teaming for high-risk AI systems [3]. Three days later, the administration issued Executive Order 14179, "Removing Barriers to American Leadership in AI," which directed agencies to review all Biden-era AI policies and suspend or rescind those deemed inconsistent with an innovation-first orientation [4]. A December 11, 2025, executive order further directed the Department of Justice to establish an AI Litigation Task Force authorized to challenge conflicting state AI laws on commerce clause and preemption grounds, signaling an effort to establish a unified — and minimally burdensome — federal floor [14].

Security Analysis

The Compliance Gap Is Now Operational

For enterprises with EU market exposure, the EU AI Act's requirements are no longer theoretical planning considerations — they are active obligations with enforcement consequences. The Act's extraterritorial scope is broad: any US-based provider whose AI system is placed on the EU market, or whose AI outputs

are used within the Union, falls within scope regardless of where the provider is incorporated or where its systems are hosted [8][15]. Penalties are calculated on global annual turnover, meaning that a US company's EU-facing AI products can expose its entire global revenue base to EU regulatory risk.

High-risk AI system obligations — the most demanding tier — become fully enforceable on August 2, 2026 [24]. Systems classified as high-risk under Annex III include AI used in biometric identification, critical infrastructure management, employment and human resource decisions, access to essential services such as credit scoring, law enforcement applications, migration and border management, and administration of justice [16]. For enterprises in financial services, healthcare, human resources technology, and critical infrastructure, an August 2026 deadline assessment is not premature — conformity assessments, technical documentation, EU database registration, and human oversight mechanisms all require extended implementation timelines.

The compliance burden is compounded by overlapping obligations for AI systems that also process personal data. EU AI Act Article 27 requires Fundamental Rights Impact Assessments (FRIAs) for high-risk AI deployed by public authorities or by private entities in specific contexts, while GDPR Article 35 independently requires Data Protection Impact Assessments (DPIAs) for high-risk data processing. The two assessments are not identical in scope or methodology, requiring enterprises to conduct and maintain parallel documentation regimes under different legal instruments [17].

Divergent Risk Philosophies Create Dual-Stack Risk

The philosophical divergence between US and EU frameworks extends beyond documentation requirements to substantive questions about which AI applications are permissible at all. The EU has enacted hard prohibitions on specific practices — prohibitions enforced through severe administrative penalties of up to €35 million or 7% of global annual turnover. There is no US federal equivalent. A US company that deploys a workplace AI system for behavioral monitoring may face no federal restrictions in the United States while simultaneously triggering high-risk classification, mandatory conformity assessment, and human oversight obligations under the EU AI Act if the system is used by EU-based employees.

This divergence is pushing some multinationals toward architecturally separate AI deployments for EU and non-EU markets — a "dual-stack" model that carries its own technical and security risks [18]. Maintaining parallel systems can increase the attack surface, complicate incident response, and create governance blind spots at the integration layer between stacks. From a security architecture perspective, the dual-stack approach trades regulatory risk for operational complexity. Organizations that instead pursue a unified governance architecture calibrated to the higher EU standard may incur greater upfront compliance cost but reduce long-term operational fragmentation.

The GPAI Regime Affects the Entire AI Supply Chain

The GPAI obligations that entered force in August 2025 affect not only frontier model providers but also enterprises that build on top of foundation models via API or fine-tuning. The EU AI Act establishes a downstream propagation mechanism: GPAI providers must supply technical documentation and usage policies to downstream system integrators, and downstream integrators acquire residual obligations to ensure that their deployed systems do not violate prohibited practice provisions [11]. This means that an enterprise deploying a third-party large language model for an EU-facing customer-service application inherits compliance obligations even if it did not train the underlying model.

Enterprises must therefore assess their AI supply chains, not just their internally developed systems. Standard contractual representations from GPAI providers regarding EU AI Act compliance status should be treated as necessary but not sufficient: downstream deployers remain independently liable for ensuring their particular use case does not constitute a prohibited practice and, if the deployment qualifies as high-risk, for conducting their own conformity assessments.

Geopolitical Volatility Amplifies Compliance Uncertainty

The regulatory divergence has acquired a geopolitical dimension that amplifies compliance planning uncertainty. The US government's refusal to sign the Paris AI Action Summit declaration in February 2025 indicated that the divergence is not incidental but reflects deliberate policy [6]. The USTR's signaling of potential Section 301 investigations targeting the EU AI Act and Digital Markets Act as unfair trade barriers introduces the possibility of retaliatory trade measures that could affect the operating environment for enterprises on both sides of the Atlantic [7]. If a Section 301 investigation is formally initiated, the resulting determination could significantly reshape the compliance calculus for US companies with EU operations.

The Turnberry Agreement reached in July 2025, under which the EU committed to purchasing approximately $40 billion in US-made AI chips, demonstrates that AI governance has become intertwined with broader trade and economic negotiations [19]. Compliance planning that assumes regulatory stability in either jurisdiction carries increasing tail risk; enterprises should structure their governance programs to accommodate regulatory change rather than optimizing for current requirements.

Recommendations

Immediate Actions

- **Conduct an EU AI Act scope assessment now.** Determine whether any AI systems currently deployed — or planned for deployment — process data involving EU-based individuals or are offered on the EU market. The extraterritorial scope is broad; enterprises that have not conducted a formal scope assessment may be in-scope without having recognized it.
- **Audit third-party AI providers for GPAI compliance status.** Request documentation from foundation model providers regarding their EU AI Act GPAI obligations status, the Code of Practice adherence, and the technical documentation they are providing to downstream deployers.
- **Map current AI deployments against Annex III high-risk categories.** If any systems fall within high-risk categories, begin the conformity assessment process immediately — the August 2, 2026, enforcement date requires significant lead time for documentation, technical controls, and registration [23][24].
- **Ensure Article 4 AI literacy obligations are being met for EU-facing personnel.** Training programs should document that staff interacting with AI systems have received structured competency instruction; this obligation has been enforceable since February 2025.

Short-Term Mitigations

Organizations should adopt NIST AI RMF as a dual-regime governance baseline. The NIST AI RMF's GOVERN, MAP, MEASURE, and MANAGE functions provide a widely adopted governance bridge that maps to EU AI Act conformity obligations while fulfilling voluntary US compliance expectations [5][9]. Implementing NIST AI RMF as a common framework reduces the risk of maintaining entirely parallel governance programs. ISO/IEC 42001 provides a complementary international standard that has been positioned by standards bodies as a tool for demonstrating EU AI Act compliance alignment; acceptance by national competent authorities remains developing and varies by jurisdiction, but the standard is increasingly referenced in conformity planning.

For systems that may qualify as high-risk under Annex III, enterprises should complete FRIA and DPIA assessments in parallel rather than sequentially, using a harmonized assessment template that satisfies both EU AI Act and GDPR requirements. Legal and compliance teams should develop a shared methodology with clear ownership of each requirement to avoid duplicative effort and documentation

gaps. Technical teams should begin implementing human oversight mechanisms and logging infrastructure — requirements that are prerequisites for conformity assessment and that cannot be retrofitted quickly into production systems.

Strategic Considerations

The current trajectory of US-EU AI regulatory divergence does not suggest near-term harmonization. Enterprises should plan for a persistent dual-regime compliance environment and architect their AI governance programs accordingly. This means establishing governance structures with the capacity to track regulatory developments in both jurisdictions — including the ongoing Digital Omnibus simplification proposal from the European Commission, the potential outcomes of any USTR Section 301 investigation, and federal versus state AI regulation dynamics in the United States [20][21].

A unified AI governance program calibrated to the EU standard — which imposes more demanding horizontal obligations across most AI use cases — will generally satisfy US federal requirements as well, since US requirements are currently less stringent, though sector-specific US requirements (such as FDA AI/ML guidance or OCC model risk management) should be assessed separately. The incremental cost of building to the EU standard is, in many cases, likely lower than the organizational cost of managing two separate, divergent governance programs over a multi-year compliance horizon, though the actual cost differential will vary significantly by organization size and existing governance maturity. Where true architectural separation is unavoidable (for example, to accommodate EU-prohibited practices that are lawful in the US), those separation points should be documented explicitly, monitored continuously, and reviewed as regulations evolve.

Finally, enterprises should engage with industry coalitions and standards bodies that are working to develop transatlantic AI governance bridges [22]. The EU AI Office's GPAI Code of Practice process, NIST's ongoing AI RMF development, and ISO/IEC AI standards work all present opportunities to influence harmonization and to signal institutional commitment to responsible AI governance — a posture that reduces regulatory risk and supports procurement relationships with EU public-sector clients.

CSA Resource Alignment

The US–EU AI governance divergence maps directly to several Cloud Security Alliance frameworks and publications that provide actionable guidance for enterprises navigating dual-regime compliance.

MAESTRO (Multi-Agent Environment Safety and Threat Reasoning for Organizations): The MAESTRO framework provides threat modeling guidance specifically for agentic and AI systems. Organizations deploying agentic AI within the EU market should apply MAESTRO's layer-by-layer threat analysis to identify which threat categories — particularly unauthorized actions, data exfiltration, and human oversight bypass — are relevant to EU AI Act high-risk classification assessments and FRIA requirements.

AI Incident Management and Control Matrix (AICM): The AICM extends the Cloud Controls Matrix to address AI-specific risk domains including data governance, model management, and operational AI controls. The AICM control domains map to several EU AI Act conformity requirements: technical documentation, data governance, accuracy and robustness testing, and human oversight implementation. Enterprises using AICM as an internal control framework can cross-reference control status against EU AI Act Annex IV technical documentation requirements to identify compliance gaps.

STAR (Security Trust Assurance and Risk): The CSA STAR program provides a public registry of security assurance postures for cloud and AI providers. For enterprises assessing GPAI provider compliance, STAR registry entries combined with EU AI Act GPAI documentation requests provide a structured method for third-party supply chain due diligence.

Zero Trust Guidance: CSA's Zero Trust architecture principles align with EU AI Act human oversight and access control requirements. High-risk AI systems must implement access controls, audit logging, and human override capabilities that are consistent with Zero Trust design — enterprises implementing Zero Trust architecture in parallel with AI Act compliance work should treat these as synergistic, not sequential, initiatives.

AI Organizational Responsibilities: CSA's guidance on organizational accountability for AI systems addresses the governance structures — board oversight, designated AI officers, risk committee ownership — that are prerequisites for EU AI Act compliance governance. The appointment of an EU AI Act compliance officer or equivalent internal role is recommended for enterprises with substantive in-scope deployments.

References

[1] EU AI Act. "EU Artificial Intelligence Act – Implementation Timeline." artificialintelligenceact.eu, 2024–2026.

[2] EU AI Act. "Article 99 – Penalties." artificialintelligenceact.eu, 2024.

[3] Wiley Rein LLP. "President Trump Revokes Biden Administration's AI EO – What To Know." Wiley, January 2025.

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

[5] NIST. "AI Risk Management Framework." NIST, January 2023 (updated 2025).

[6] TechCrunch. "As US and UK refuse to sign the Paris AI Action Summit statement, countries fail to agree on the basics." TechCrunch, February 11, 2025.

[7] ROIC News. "USTR Official Warns of Section 301 Action Over EU Tech Regulations in Escalating Trade Dispute." ROIC, December 18, 2025.

[8] Morgan Lewis. "The EU Artificial Intelligence Act Is Here – With Extraterritorial Reach." Morgan Lewis, July 2024.

[9] RSI Security. "NIST AI Risk Management Framework and ISO/IEC 42001 Crosswalk." RSI Security, 2025.

[10] Latham & Watkins. "Upcoming EU AI Act Obligations: Mandatory Training and Prohibited Practices." Latham & Watkins, 2025.

[11] Baker McKenzie. "General-Purpose AI Obligations Under the EU AI Act Kick in From 2 August 2025." Baker McKenzie, August 2025.

[12] Skadden. "EU's General-Purpose AI Obligations Are Now in Force." Skadden, August 2025.

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

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

[15] William Fry. "A Practical Guide to the Extraterritorial Reach of the AI Act." William Fry, 2024–2025.

[16] EU AI Act. "Annex III – High-Risk AI Systems." artificialintelligenceact.eu, 2024.

[17] Parloa. "The Global AI Privacy Maze: GDPR, DMA, and U.S. Rules." Parloa, 2026.

[18] SecurePrivacy. "AI Risk and Compliance in 2026." SecurePrivacy, 2026.

[19] CEPA. "What the US-EU $40 Billion Chip Deal Means." Center for European Policy Analysis, 2025.

[20] Cooley. "EU AI Act: Proposed 'Digital Omnibus on AI' Will Impact Business AI Compliance Roadmaps." Cooley, November 24, 2025.

[21] Skadden. "White House Launches National Framework Seeking To Preempt State AI Regulation." Skadden, December 2025.

[22] King & Spalding. "Transatlantic AI Governance: Strategic Implications for US-EU Compliance." King & Spalding, 2025–2026.

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

[24] LegalNodes. "EU AI Act 2026 Updates: Compliance Requirements and Business Risks." LegalNodes, 2026.

[25] EU AI Act. "Article 5 – Prohibited Artificial Intelligence Practices." artificialintelligenceact.eu, 2024.