CSAI Foundation | Cloud Security Alliance

EU AI Act Risk Tiers: The ISO 42001 Compliance Path

Enterprise AI Governance Under the 2026 Deadline

2026-05-05

AI-assisted Rapid Research

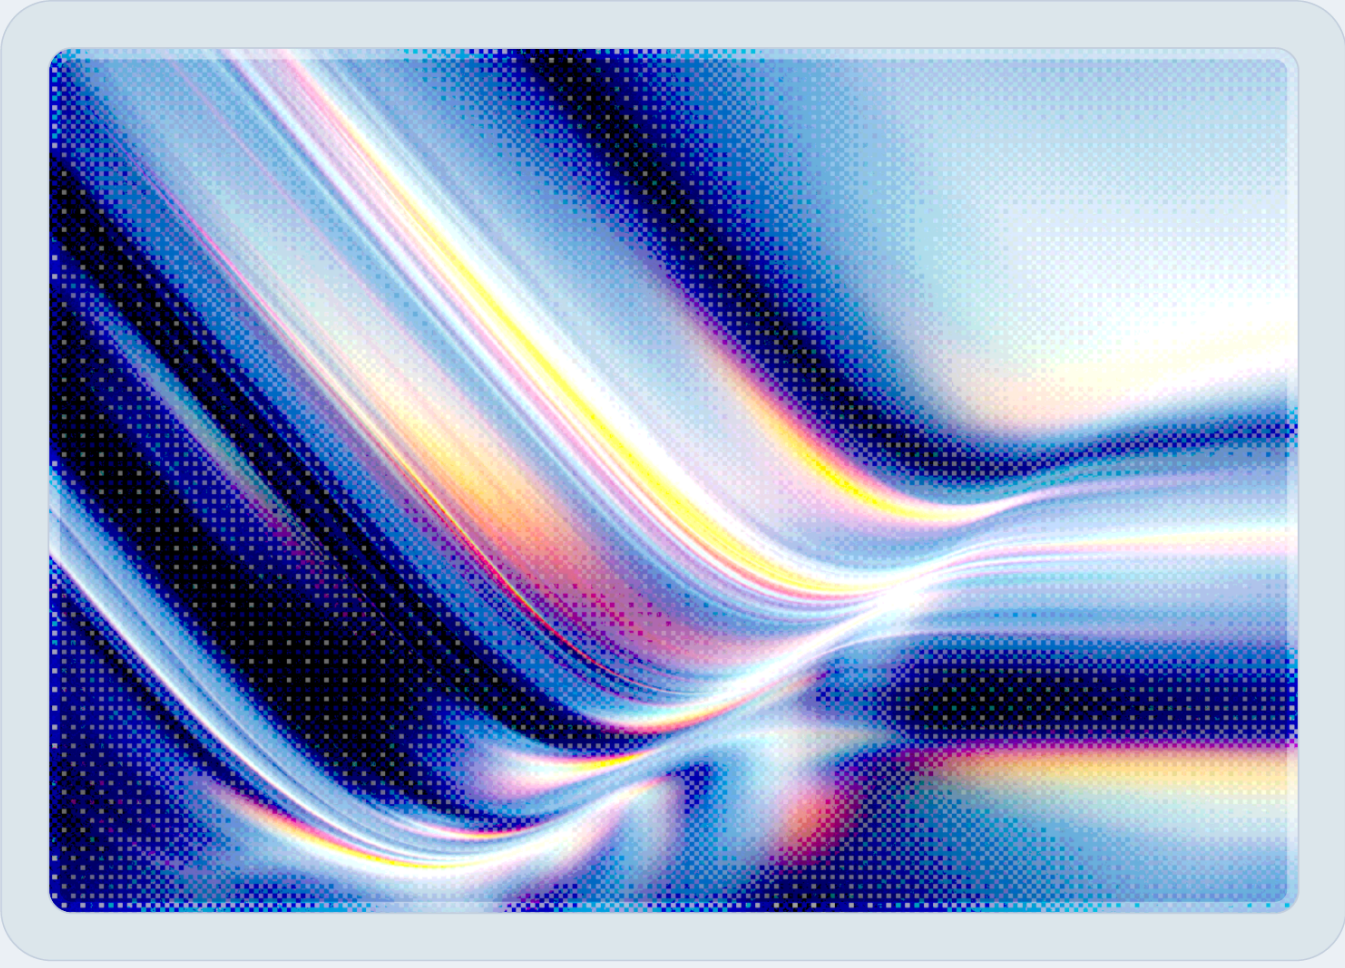

CSAI | CSA cloud security alliance®

Table of Contents

Executive Summary

The European Union's Artificial Intelligence Act entered into force on August 1, 2024, inaugurating the world's first comprehensive legal framework for AI systems. Its risk-tiered structure imposes graduated obligations that range from outright prohibition of certain AI practices to lightweight transparency requirements for low-risk deployments. For the majority of enterprises deploying AI in consequential business contexts, the pivotal deadline is August 2, 2026, when the full suite of high-risk AI system requirements becomes enforceable, including mandatory risk management systems, data governance controls, technical documentation, human oversight mechanisms, and conformity assessment procedures.

The compliance question most enterprises face is not whether to act, but how to build a governance infrastructure that satisfies the Act's legal obligations without duplicating effort across multiple regulatory regimes. ISO/IEC 42001:2023, the international AI management system standard, provides approximately 70–80% of the organizational foundation needed to meet EU AI Act high-risk obligations [1]. However, ISO 42001 was not developed specifically for EU regulatory compliance, and legal analysis has confirmed that the standard is not part of the EU harmonization process and is insufficient to grant presumption of conformity under the Act [2]. The gap is being closed by prEN 18286, a European pre-standard for AI quality management systems that entered public enquiry in October 2025 and is designed specifically to operationalize Article 17 of the Act [3].

This paper examines each risk tier's obligations in detail, analyzes the coverage and limitations of ISO 42001, explains the role and status of prEN 18286, and presents a practical layered compliance architecture. The architecture draws on the CSA AI Controls Matrix (AICM) and MAESTRO threat-modeling framework as technical overlays that bridge organizational governance to operational AI security. Enterprise security leaders, AI governance teams, and compliance officers will find in these pages both a conceptual map of the regulatory landscape and a roadmap for building durable, audit-ready AI governance programs.

Introduction and Background

The Regulatory Imperative

For nearly three decades, AI development in the European Union proceeded under a patchwork of sector-specific rules—medical device regulations, financial services directives, data protection law—without any governance framework specific to the capabilities and risks that AI systems introduce. The EU AI Act marks a

significant departure from that sector-specific model: a horizontal, product-safety-style regulation that applies across sectors and establishes AI-specific obligations based on the risk that each application poses to health, safety, and fundamental rights [4].

The Act was formally published in the Official Journal of the EU in July 2024 and entered into force on August 1, 2024. Its obligations apply in phases calibrated to the severity of the risk and the complexity of compliance preparation required. Prohibitions on unacceptable-risk AI practices took effect on February 2, 2025 [4]. Obligations for providers of general-purpose AI (GPAI) models became applicable on August 2, 2025 [5], [20]. The full high-risk AI system provisions—the most operationally intensive requirements—become enforceable on August 2, 2026, with a further extended deadline of August 2, 2027, for systems already on the market before that date [4], [17]. Separately, the European Commission's Digital Omnibus proposal of late 2025 introduced the possibility of linking Annex III high-risk rule application to the availability of finalized harmonized standards, establishing backstop dates of December 2, 2027 for Annex III systems and August 2, 2028 for Annex I-integrated systems—though this proposal had not yet been finalized as of this writing [3].

For enterprise AI teams, the practical implication is that 2026 is not a future planning horizon; it is an approaching enforcement date. Organizations that deploy AI systems in recruiting, lending, insurance underwriting, medical diagnosis support, infrastructure operations, or law enforcement contexts—to name only the most prominent Annex III categories—must have compliance programs operational before that date or face penalties, mandatory market withdrawal, and reputational damage [18].

Why Standards Matter

The EU AI Act establishes legal obligations but leaves considerable technical discretion to organizations in how those obligations are met. Article 40 grants organizations a "presumption of conformity" with the Act's requirements when they demonstrate compliance with harmonized European standards that the Commission has published in the Official Journal [6]. In practical terms, an organization demonstrating compliance with a harmonized standard benefits from a presumption of conformity, meaning market surveillance authorities must provide specific evidence of non-compliance before imposing sanctions—a meaningful procedural advantage compared to organizations relying on documented internal assessments alone.

ISO/IEC 42001 is among the most prominently adopted international AI management system standards, with certifications achieved across multiple sectors and jurisdictions [9], but it was developed for international applicability rather than EU regulatory compliance specifically, and legal analysis has confirmed that the standard is insufficient to grant presumption of conformity under the Act [2]. prEN 18286 fills this gap by providing a European standards-track document designed explicitly to satisfy Article 17's quality management system requirements. Until prEN 18286 is finalized and published in the Official Journal,

however, no harmonized standard currently grants presumption of conformity for AI Act obligations. This creates a period of compliance uncertainty that demands a deliberate, documented governance approach rather than simple certification reliance.

EU AI Act Risk Tier Obligations

The Tier Architecture

The Act organizes AI systems into four risk tiers: prohibited (unacceptable risk), high risk, limited risk, and minimal risk. The tier an AI system occupies determines the complete set of legal obligations that its provider and deployer must satisfy. A single enterprise may have AI systems in multiple tiers simultaneously, requiring parallel compliance workstreams.

Tier 1: Prohibited AI Practices

Article 5 enumerates eight categories of AI practice that the Act treats as incompatible with EU fundamental rights, banning them outright. These prohibitions became enforceable on February 2, 2025. Prohibited practices include AI-based social scoring of persons by public authorities, subliminal or manipulative techniques that exploit psychological vulnerabilities, real-time remote biometric identification in publicly accessible spaces for law enforcement purposes (with narrow exceptions for specified serious crimes), biometric categorization based on sensitive protected attributes, AI systems used to infer emotions in workplace or educational settings in ways that are not medically or safety-justified, and AI systems that create or expand facial recognition databases through untargeted scraping [4].

Violations of Article 5 carry the Act's most severe sanctions: administrative fines of up to €35 million or 7% of total worldwide annual turnover for the preceding financial year, whichever is higher [7]. Enterprises must assess their existing AI portfolio against these prohibitions and retire or redesign any system that falls within them; the prohibitions are not subject to any compliance timeline extension.

Tier 2: High-Risk AI Systems

High-risk AI systems represent the Act's most extensively regulated category. Article 6 provides two routes to high-risk classification. The first applies to AI systems that serve as safety components in products already subject to EU product safety legislation listed in Annex II—including medical devices, machinery, and civil aviation components. The second applies to any AI system performing one of the functions listed in Annex III, regardless of the product context [8].

Annex III identifies eight domains in which AI applications are classified as high-risk by default: biometric identification and categorization; management and operation of critical infrastructure such as power, water, and digital networks; education and vocational training (including admissions and evaluation systems); employment and workforce management (including recruitment screening, performance monitoring, and termination decisions); access to essential public and private services such as credit scoring and health insurance underwriting; law enforcement, including risk assessment and evidence evaluation; migration, asylum, and border control; and the administration of justice and democratic processes [8].

The obligations attached to high-risk AI systems are substantive and operational, spanning the entire system lifecycle. Article 9 requires a continuous risk management system that identifies, analyzes, and mitigates risks to health, safety, or fundamental rights, updated throughout the system's operational life [6]. Article 10 mandates data governance practices ensuring that training, validation, and testing data are relevant, sufficiently representative, free of errors, and complete, with explicit requirements to examine data for biases [6]. Articles 11 and 12 require technical documentation comprehensive enough to demonstrate regulatory compliance and automated logging sufficient to enable post-deployment audits [6]. Article 13 requires transparency obligations toward deployers, including clear instructions for use [6]. Article 14 mandates human oversight measures enabling humans to monitor, understand, and intervene in system outputs [6]. Article 15 requires adequate levels of accuracy, robustness, and cybersecurity, with specific provisions for systems operating in adversarial conditions [6].

Beyond these technical requirements, providers must also establish a quality management system meeting Article 17's thirteen-element specification, conduct a conformity assessment procedure under Article 43 before placing the system on the market, affix a CE marking, and register the system in the EU's AI database [6]. Deployers of high-risk systems carry their own obligations under Article 26, including conducting fundamental rights impact assessments for certain applications, assigning human oversight responsibilities, and maintaining post-deployment monitoring [6].

Penalties for noncompliance with high-risk AI obligations reach €15 million or 3% of global annual turnover, whichever is higher [7].

Tier 3: General-Purpose AI Models

The Act's GPAI provisions represent a distinct regulatory regime targeting foundation models and large-scale AI systems that can serve multiple purposes across many downstream applications. These obligations became applicable on August 2, 2025, making GPAI the earliest-enforced tier after prohibited practices [5].

Providers of GPAI models must maintain and update technical documentation, publish summaries of training data used to develop the model, comply with EU copyright law in training data sourcing, and provide downstream providers with sufficient information to meet their own compliance obligations [5]. For providers of GPAI models that the AI Office determines pose systemic risk—currently defined by reference to

training compute thresholds—additional obligations apply: adversarial testing and model evaluation before and after deployment, cybersecurity protections for model weights, and reporting of serious incidents to the AI Office within defined timeframes [5].

The European Commission published guidelines for GPAI model providers in July 2025 to clarify the scope of these obligations [5]. Full Commission enforcement powers against GPAI providers, including the authority to demand documentation, order model withdrawals, and impose fines, are scheduled to take effect on August 2, 2026. GPAI model noncompliance can result in fines of up to €15 million or 3% of global revenue, whichever is higher [7].

For most enterprise AI teams, the GPAI tier creates an important upstream dependency: when an enterprise deploys a high-risk AI system built on a third-party foundation model, compliance requires verifying that the GPAI provider has met its own Act obligations, since the provider's technical documentation forms part of the enterprise's own compliance evidence chain.

Tier 4: Limited and Minimal Risk

Limited-risk AI systems are subject to transparency obligations designed to ensure that users understand they are interacting with AI. Chatbots and conversational AI must identify themselves as such; AI-generated or AI-manipulated audio, image, video, or text content must be labeled as such, with specific requirements for deepfakes [4]. These obligations are operationally lightweight but create documentation requirements: providers must implement technical measures enabling users to recognize AI-generated content, and deployers must ensure disclosure practices are in place.

Minimal-risk AI systems—including most AI-enabled productivity tools, recommendation engines, and spam filters—carry no legally mandated obligations under the Act. The Commission has encouraged the development of voluntary codes of conduct for this tier [4]. For enterprise AI security purposes, the absence of mandatory controls does not imply the absence of risk, but it does mean that compliance resources should be concentrated on the higher-risk tiers.

The Article 17 Quality Management System Mandate

Scope and Structure

Article 17 requires providers of high-risk AI systems to establish a quality management system before placing a system on the market or putting it into service. Unlike the technical requirements of Articles 9–15, which apply to individual AI systems, the Article 17 QMS is an organizational-level requirement: a documented,

implemented, and maintained management system that encompasses all of the provider's high-risk AI activities.

The Act specifies thirteen elements that every compliant QMS must address. These include a documented strategy for regulatory compliance; procedures governing design and development; testing and validation procedures; technical specifications and standards applied; data management provisions; risk management procedures and their integration with design; post-market monitoring plans; procedures for serious incident reporting; communication and documentation protocols; resource management, including personnel competency; accountability assignments; and provisions governing AI systems that continue to learn after deployment [6]. The functional significance of this enumeration is substantial: Article 17 does not permit a provider to point to a generic quality program or an existing ISO certification as sufficient. The QMS must demonstrably address each element in the context of AI system development and deployment.

Why Article 17 Is the Central Compliance Challenge

Article 17 is strategically important beyond its literal requirements because it serves as the anchor point for the harmonized standards pathway. prEN 18286 was developed specifically to operationalize Article 17, and once published in the Official Journal, compliance with prEN 18286 will grant presumption of conformity with Article 17 specifically. An enterprise that builds its QMS against prEN 18286 therefore positions itself for the strongest available legal defense against regulatory challenge.

The thirteen-element structure of Article 17 also reveals an important design choice in the Act: the legislature intended compliance to be a continuous management discipline, not a point-in-time assessment. Organizations that treat Article 17 as a documentation exercise—assembling the required paperwork without embedding the underlying processes—risk satisfying initial audits while leaving material operational gaps that typically surface through system updates, data drift, or incident response situations where documented procedures are not followed in practice.

ISO 42001: The Organizational Foundation

Standard Architecture and Purpose

ISO/IEC 42001:2023, published in December 2023, is the first international management system standard specifically designed for AI [9]. Its structure follows the ISO high-level structure (HLS) common to ISO 27001, ISO 9001, and other management system standards, organizing requirements across ten clauses: scope; normative references; terms and definitions; organizational context; leadership; planning; support; operation; performance evaluation; and improvement [9].

Clause 6 (Planning) requires organizations to identify AI risks and opportunities, establish AI policies, and define objectives with measurable targets. Clause 8 (Operation) covers the operational planning and control of AI activities, including the development and deployment of AI systems and the management of the AI supply chain. Clause 9 (Performance Evaluation) requires internal audits, management reviews, and monitoring of the AI management system's effectiveness. Annex A provides 38 controls across eight domains—AI system context, AI system data, AI system knowledge, human resources, security and privacy, AI system processes, AI system transparency, and AI use and impact—with guidance in Annex B on how to apply them [9].

ISO 42001 certification is conducted by accredited third-party certification bodies following the same two-stage audit process used for ISO 27001 and ISO 9001. By late 2025, major professional services firms across multiple sectors had achieved ISO 42001 certification, and organizations with active AI governance programs are increasingly pursuing the standard as a structured foundation for their AI compliance obligations [10]. The standard's global applicability makes it the natural anchor for multinational enterprises that must satisfy both EU AI Act requirements and comparable regulations in other jurisdictions.

Coverage of EU AI Act High-Risk Obligations

ISO 42001 maps to EU AI Act high-risk requirements with meaningful but incomplete coverage. Risk management activities required by Article 9 align closely with ISO 42001's Clause 6 planning and risk assessment provisions, though the standard's risk language is less prescriptive than the Act's continuous lifecycle requirement. Data governance requirements in Article 10 find corresponding controls in ISO 42001's Annex A data domain, though the Act specifies additional obligations around bias examination and statistical representativeness that the standard addresses at a higher level of abstraction. Technical documentation and logging requirements under Articles 11 and 12 are supported by ISO 42001's documentation requirements, though not at the granularity the Act requires. Human oversight (Article 14) and supply chain management provisions find reasonable analogues in ISO 42001's Annex A controls [1].

The CSA AI Controls Matrix mapping to ISO 42001, published in August 2025, provides a control-level crosswalk demonstrating this coverage in operational terms. The AICM's 243 controls across 18 domains include explicit mappings to both ISO 42001 clauses and EU AI Act articles, enabling organizations to use a single control set as evidence across both frameworks [11], [21].

---

prEN 18286: The Harmonized Compliance Standard

Development and Status

prEN 18286, formally titled "Artificial Intelligence – Quality Management System for EU AI Act Regulatory Purposes," entered CEN public enquiry on October 30, 2025, making it the first AI-specific harmonized standard to progress through the European standardization process under the EU AI Act [3], [19]. The enquiry period closed on December 27, 2025; the standard is undergoing resolution of national body comments with finalization and Official Journal publication anticipated in late 2026 [3].

The standard was developed by CEN/CLC/JTC 21, the joint technical committee for AI standardization established by CEN (the European Committee for Standardization) and CENELEC (the European Committee for Electrotechnical Standardization). Its development was mandated by the European Commission as part of the broader standardization program supporting the AI Act, which encompasses dozens of standards across safety, robustness, transparency, cybersecurity, and testing [12].

Until prEN 18286 is finalized and cited in the Official Journal, no harmonized standard currently provides presumption of conformity for AI Act Article 17 obligations. Organizations that wish to demonstrate compliance before the standard's publication must document their QMS against Article 17's thirteen elements directly, using the draft standard as a technical guide to the expected structure and evidence requirements while acknowledging that the final standard may differ from the enquiry text.

Structural Design

prEN 18286 mirrors ISO 42001's high-level structure while adding EU AI Act-specific requirements at each stage. Clause 4 addresses organizational context with an explicit requirement to map the organization's AI activities to the Act's risk classification framework, identifying which systems are high-risk and why. Clause 5 (Leadership) requires executive accountability structures for AI compliance specifically, not merely for AI governance in the abstract. Clause 6 (Planning) requires a regulatory compliance strategy as a first-class QMS component, distinguishing it from the general risk management strategy ISO 42001 requires [3].

The standard's operational clauses are where its specificity most clearly exceeds ISO 42001's scope. Clause 8 requires lifecycle controls that explicitly address each phase of AI system development—data collection, model design, training, validation, deployment, and monitoring—with documentation requirements tied to the Act's technical documentation articles. Annex D of prEN 18286 provides a normative mapping to ISO 42001, identifying the clauses and controls in each standard that address the same underlying governance requirements. This mapping serves a practical purpose: organizations that have already implemented ISO 42001 can use Annex D as a gap analysis guide, identifying precisely where they must extend or supplement their existing management system to satisfy prEN 18286 [3].

Legal Effect of Harmonized Standard Status

The legal significance of prEN 18286's eventual Official Journal citation is substantial. Under Article 40 of the AI Act, a provider that demonstrates conformity with a harmonized standard cited in the Official Journal is presumed to conform with the corresponding Act requirements, shifting the burden of proof in any regulatory proceeding [6]. A conformity assessment audit against prEN 18286 conducted by a qualified third party provides auditable evidence of this presumption. Organizations without a harmonized standard to rely on must instead affirmatively demonstrate compliance by other means—a more resource-intensive and legally uncertain posture.

The Compliance Gap: Where ISO 42001 Falls Short

Five Material Gaps

Legal and technical analysis of the relationship between ISO 42001 and EU AI Act high-risk obligations has identified five areas where ISO 42001 does not adequately address the Act's requirements, even when fully implemented [2], [14], [15], [16]. Understanding these gaps is essential for enterprises that have invested in ISO 42001 certification and wish to extend that investment to cover EU AI Act compliance.

The first gap concerns the absence of a per-system regulatory compliance strategy. ISO 42001 establishes organizational-level governance for AI activities but does not require providers to develop and document a compliance strategy at the level of individual AI systems. Article 17(1)(a) of the AI Act explicitly requires such a strategy as a component of the QMS, meaning a certified ISO 42001 program provides no direct evidence of this Article 17 element [2].

The second gap involves change management for continuously learning systems. Many enterprise AI systems—particularly those in fraud detection, personalization, and predictive analytics—update their parameters or retrain on new data after deployment. ISO 42001 does not include provisions specifically governing the controls required when a post-deployment learning event may change a system's risk profile. Article 17(1)(k) of the Act specifically requires QMS procedures addressing this scenario [2].

The third gap is the absence of incident reporting timelines aligned with Article 73 of the Act. The Act requires providers of high-risk AI systems to report serious incidents to national market surveillance authorities within specific timeframes. ISO 42001's incident management provisions address general AI incident governance without specifying the regulatory notification obligations or timelines the Act requires [2].

The fourth gap concerns supply chain provisions. ISO 42001 includes general supply chain requirements for AI governance but does not address the AI Act's specific obligations regarding information flows between providers and deployers, technical documentation sharing, and the distribution of compliance responsibilities across the value chain. For enterprise deployers that rely on third-party model providers, this gap has direct practical consequences [2].

The fifth gap involves fundamental rights impact assessments. Article 26 of the Act requires deployers of certain high-risk AI systems to conduct fundamental rights impact assessments before deployment. ISO 42001 does not include a corresponding requirement, and its risk management provisions use a safety-and-reliability framing that does not fully capture the fundamental rights analysis the Act expects [2].

Accelerating Without Duplicating

Despite these gaps, ISO 42001 remains the most practical organizational starting point for EU AI Act high-risk compliance. Compliance practitioners report that organizations with ISO 42001 certification can significantly accelerate EU AI Act program development—the governance infrastructure, documentation discipline, and management system mindset that ISO 42001 establishes are directly reusable—though the degree of acceleration varies with organizational maturity and system complexity [1]. The strategic posture for 2026 is to implement ISO 42001 as the organizational AI management system, address the five identified gaps through supplementary controls and documentation, and then re-baseline the resulting system against prEN 18286 when the final standard is published.

A Layered Compliance Architecture

The Three-Tier Model

CSA recommends a layered architecture that assigns distinct responsibilities to three complementary frameworks, each operating at a different level of abstraction in the governance stack. This model enables evidence reuse, reduces duplication, and provides clear organizational roles for each compliance instrument [11].

At the organizational level, ISO/IEC 42001 establishes the management system infrastructure: governance policies, risk management processes, resource allocation, internal audit, and continual improvement. Every high-risk AI activity the enterprise undertakes operates within this system. ISO 42001 certification provides an independently verified signal of organizational AI governance maturity that is recognized globally and maps to multiple regulatory regimes simultaneously [9].

At the per-system level, prEN 18286 provides the quality management framework that maps each individual high-risk AI system's lifecycle controls to Article 17's thirteen-element requirements. Where ISO 42001 establishes how the organization governs AI in general, prEN 18286 establishes how each regulated system specifically is designed, deployed, monitored, and updated within that organizational context. Until prEN 18286 is finalized, the Article 17 elements themselves serve as the per-system compliance specification [3].

At the technical control level, the CSA AI Controls Matrix provides 243 granular controls across 18 domains with explicit mappings to ISO 42001, prEN 18286, and EU AI Act articles [11], [21]. The MAESTRO framework (Multi-layer Architecture for Evaluating Security of AI Systems with Threat Reasoning and Oversight) provides a threat-modeling methodology for identifying and mitigating AI-specific security risks at the system level, including adversarial robustness, data poisoning, and model inversion threats that Article 15 of the Act requires providers to address. Organizations may also evaluate complementary frameworks such as the NIST AI Risk Management Framework and MITRE ATLAS depending on their existing tooling and jurisdictional requirements.

Practical Implementation Sequence

For enterprises that have not yet started a formal AI governance program, the following sequence minimizes effort while maximizing compliance coverage:

- Conduct an AICM-based gap assessment against current practices, and use the findings to scope the ISO 42001 implementation
- Implement ISO 42001 with the five EU AI Act gaps addressed as supplementary requirements
- Classify all AI systems by Act risk tier and document the classification rationale
- Build per-system compliance artifacts using the prEN 18286 draft as a structural guide
- Complete conformity assessments for high-risk systems before the August 2026 deadline
- Re-baseline the entire program against prEN 18286 once the final standard is published

For enterprises already certified to ISO 42001, the sequence compresses to the last three steps, making accelerated compliance substantially more achievable.

Conformity Assessment and Market Entry

Assessment Routes

Article 43 of the Act establishes the conformity assessment procedures that high-risk AI providers must complete before placing a system on the market. The route available depends on which Annex the system falls under [6].

For high-risk AI systems listed in points 2 through 8 of Annex III—which covers the majority of enterprise AI use cases in employment, lending, essential services, and similar domains—the required procedure is the internal control procedure defined in Annex VI. Under this route, the provider itself assesses the system against the Act's requirements, prepares technical documentation, and issues an EU declaration of conformity. No involvement of a notified body is required [6]. This is a significant practical relief: it means most enterprise deployers do not depend on the availability or capacity of designated third-party assessment bodies to complete their compliance procedures.

For remote biometric identification systems listed in Annex III point 1, and for high-risk AI systems that are safety components in Annex II-regulated products, the provider must choose between internal control with notified body quality management assessment (Annex VII) or the product-specific conformity assessment procedure applicable to the regulated product [6]. Notified body designation under the Act commenced in August 2025, but as of early 2026, the number of bodies fully designated specifically for AI Act conformity assessment remains limited, and harmonized standards are not yet available in finalized form [13]. Enterprises requiring notified body assessment should begin the engagement process early, as capacity constraints are likely to create scheduling delays approaching the August 2026 deadline.

Following a successful conformity assessment, providers must affix the CE marking, register the system in the EU AI database, and provide the declaration of conformity to deployers and market surveillance authorities on request [6].

Post-Market Monitoring and Incident Reporting

Conformity assessment marks the beginning of an ongoing compliance obligation, not its conclusion. Article 72 requires providers to establish and maintain post-market monitoring systems that actively gather and analyze data on high-risk AI system performance in real-world conditions, identifying emerging risks and unintended impacts not visible in pre-deployment testing [6]. The monitoring plan required by Article 17 must specify what data will be collected, at what frequency, with what analysis methods, and what thresholds will trigger action.

Article 73 requires providers to report serious incidents—defined as incidents that result or may result in death, serious harm to health, serious property damage, or significant disruption to critical services—to national market surveillance authorities. Providers must also report cases where a high-risk AI system behaves in an unexpected manner that represents a significant risk [6]. These notification obligations come with specific timelines that the QMS must operationalize, creating a direct dependency on the incident management gap identified in ISO 42001.

Implementation Roadmap

2025–2026 Priority Actions

Enterprise AI teams have a narrowing window to establish compliant governance programs ahead of the August 2026 deadline. The following table maps the key compliance activities to their associated Act requirements and recommended timing.

Activity	Act Requirement	Target Completion
AI system inventory and risk tier classification	Article 6, Annex I/III	Q2 2026
Prohibited practices audit and remediation	Article 5	Complete (Feb 2025 deadline passed)
ISO 42001 gap assessment or certification initiation	Article 17 foundation	Q2 2026
Per-system QMS documentation for each high-risk system	Article 17 (all 13 elements)	Q3 2026
Risk management system implementation	Article 9	Q3 2026
Data governance controls for training/validation data	Article 10	Q3 2026
Technical documentation preparation	Articles 11–12	Q3 2026

Activity	Act Requirement	Target Completion
Human oversight mechanism design and testing	Article 14	Q3 2026
Conformity assessment (internal or with notified body)	Article 43	Q3 2026
CE marking and EU database registration	Articles 49, 71	Before market entry
Post-market monitoring system deployment	Article 72	August 2026
Incident reporting procedures	Article 73	August 2026
Pilot audit against prEN 18286 draft	Article 17 future-proofing	Q4 2026
Re-baseline against finalized prEN 18286	Article 17 (presumption of conformity)	2027 (after publication)

Managing GPAI Dependencies

Enterprises deploying high-risk AI systems built on third-party GPAI models face a compound compliance obligation: they must meet their own Article 17 QMS requirements while also verifying that their GPAI model provider has met the GPAI obligations that became applicable in August 2025. The technical documentation that high-risk AI providers must prepare under Article 11 must reference or incorporate information from the GPAI provider's own documentation [5]. Enterprises should establish contractual requirements with GPAI providers for documentation access and timely notification of model changes, since a post-deployment model update may alter a high-risk system's risk profile in ways that trigger a revised conformity assessment.

Avoiding Common Implementation Pitfalls

Several implementation patterns tend to produce compliance gaps that audits subsequently expose. Treating Article 17 as a documentation project rather than a management discipline produces a QMS that is formally complete but operationally inert—policies exist but are not followed, and monitoring data is collected but not analyzed. Auditors experienced with management system reviews are trained to identify this pattern through worker interviews and control testing that reveals disconnects between documented procedures and actual practices, following the evidence-based auditing methodology codified in ISO 17021.

Relying on ISO 42001 certification alone, without addressing the five identified gaps, represents a second common pitfall. The certification provides valuable organizational governance infrastructure and may satisfy auditors in some jurisdictions, but it will not provide presumption of conformity under Article 40 for any EU AI Act article, and it will leave material compliance exposures in incident reporting, continuous learning governance, and fundamental rights assessment.

Finally, deferring high-risk system classification decisions on the grounds that the system's use case is ambiguous creates a compliance risk that compounds over time. The Act's Article 6 classification rules require affirmative analysis, and undocumented classification decisions—even correct ones—leave organizations without the documentary evidence needed to support those decisions in a regulatory inquiry.

CSA Framework Alignment

AI Controls Matrix

The CSA AI Controls Matrix (AICM) v1.0 provides a structured bridge between organizational AI governance and the technical control requirements that the EU AI Act and ISO 42001 share. The AICM's 243 controls span 18 domains including AI system security, data governance, model risk management, supply chain security, human oversight, incident management, and compliance monitoring [11], [21]. The August 2025 mapping release explicitly cross-references AICM controls to ISO 42001 clauses and EU AI Act articles, enabling compliance teams to generate consolidated evidence packages that satisfy multiple frameworks from a single control implementation.

For the five compliance gaps identified in ISO 42001, specific AICM controls provide direct coverage. The AICM's regulatory compliance management domain addresses the per-system strategy requirement of Article 17(1)(a). Its continuous learning governance controls address Article 17(1)(k). Incident management controls in the AICM's security and operations domain address the Article 73 notification timeline gap. Supply chain security controls address the cross-party documentation and information-sharing requirements. And the AICM's impact assessment domain addresses the fundamental rights assessment requirements of Article 26 [11].

MAESTRO for AI Threat Modeling

The MAESTRO framework provides a seven-layer model of AI system architecture that enables structured threat identification and risk assessment at a granularity the EU AI Act's Article 9 risk management requirements demand but do not specify [11]. MAESTRO's layers—from foundation models and data operations through agent orchestration, tool integration, and deployment infrastructure—align with the

lifecycle stages that both prEN 18286 and Article 17 require QMS controls to address. Applying MAESTRO to high-risk AI systems enables organizations to produce technically credible, auditable risk assessments that satisfy the "identify and analyze risks" obligation of Article 9 while generating the threat-model documentation that Article 11 technical documentation requirements encompass.

AI Security Governance and Zero Trust

CSA's Zero Trust guidance applies to AI deployments through the principle that no AI component—model, tool, data source, or third-party service—should receive unconditional trust. In EU AI Act terms, zero trust principles operationalize the human oversight obligations of Article 14 and the robustness and cybersecurity requirements of Article 15. Enterprises deploying high-risk AI systems should incorporate zero trust network and access controls into their technical architecture documentation, demonstrating that the system's design actively constrains the blast radius of a model compromise or adversarial attack. CSA's Cloud Controls Matrix (CCM) provides the cloud security control mapping that supports this integration, and the CCM-to-AICM alignment enables unified evidence collection across cloud security and AI governance control domains.

Conclusions and Recommendations

For enterprise AI teams in 2026, compliance preparation is no longer optional—prohibited practices have been enforceable since February 2025, GPAI obligations since August 2025, and the high-risk AI system provisions take effect on August 2, 2026. The financial penalties and reputational consequences attached to noncompliance make the approaching deadline a governance priority that demands action now [18]. The regulatory architecture is clear: risk tier classification drives obligation scope, Article 17's QMS mandate creates the organizational governance anchor, and the conformity assessment procedure closes the compliance loop before market entry.

ISO 42001 provides a strong organizational foundation for this compliance architecture, with approximately 70–80% coverage of high-risk AI system requirements and a globally recognized certification process that supports multinational compliance programs [1]. However, ISO 42001 alone is insufficient, and enterprises that treat certification as an endpoint will carry material compliance gaps into enforcement. The five structural gaps—per-system regulatory strategy, continuous learning governance, incident reporting timelines, AI-specific supply chain provisions, and fundamental rights assessment—must be addressed through supplementary controls and documentation.

prEN 18286, once finalized, will provide the harmonized compliance pathway that grants presumption of conformity with Article 17. Organizations that build against the draft now, document their rationale, and commit to re-baselining against the final standard will be well-positioned for the regulatory landscape that

emerges as prEN 18286 is finalized and enforcement patterns develop [3]. The Digital Omnibus backstop dates provide some relief for the most complex compliance scenarios, but waiting for those backstops to materialize is a higher-risk strategy than building compliant governance programs within the original deadlines.

The layered architecture—ISO 42001 at the organizational level, prEN 18286 at the per-system QMS level, and AICM/MAESTRO as the technical control and threat-modeling overlay—provides a practical, auditable, and maintainable governance structure. Enterprises that implement this architecture with the operational discipline that management systems require, rather than as a documentation project, will satisfy regulatory requirements while building AI governance capabilities that compound in value as AI systems become more central to business operations.

References

[1] GLACIS. "ISO 42001 vs EU AI Act: Framework Crosswalk Guide." GLACIS, 2025.

[2] CMS Law. "The first draft AI Act standard for public consultation: what prEN 18286 signals for providers, users and regulators." CMS Law-Now, December 2025.

[3] Cloud Security Alliance. "EU AI Act Compliance: prEN 18286 and ISO 42001." CSA Labs, April 2026.

[4] European Union. "Regulation (EU) 2024/1689 of the European Parliament and of the Council (EU AI Act)." Official Journal of the European Union, July 2024.

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

[6] EU Artificial Intelligence Act Repository. "High-level summary of the AI Act." artificialintelligenceact.eu, 2024.

[7] EU Artificial Intelligence Act Repository. "Article 99: Penalties." artificialintelligenceact.eu, 2024.

[8] EU Artificial Intelligence Act Repository. "Annex III: High-Risk AI Systems Referred to in Article 6(2)." artificialintelligenceact.eu, 2024.

[9] ISO. "ISO/IEC 42001:2023 — Artificial intelligence — Management system." International Organization for Standardization, December 2023.

[10] Cloud Security Alliance. "6 Key Steps to ISO 42001 Certification Explained." CSA Blog, July 2025.

[11] Cloud Security Alliance. "Announcing the AI Controls Matrix & ISO 42001 Mapping." CSA Blog, August 2025.

[12] European Commission. "Standardisation of the AI Act." European Commission, 2025.

[13] Future of Privacy Forum. "Conformity Assessments under the EU AI Act: A Step-by-Step Guide." FPF, April 2025.

[14] ISACA. "ISO/IEC 42001 and EU AI Act: A Practical Pairing for AI Governance." ISACA, 2025.

[15] Schellman. "Building EU AI Act Compliance with prEN 18286 and ISO 42001." Schellman, 2025.

[16] Lumenova AI. "A Summary of prEN 18286: Quality Management System for EU AI Act Compliance." Lumenova AI, 2025.

[17] Trilateral Research. "EU AI Act Compliance Timeline: Key Dates for 2025–2027 by Risk Tier." Trilateral Research, 2025.

[18] Legal Nodes. "EU AI Act 2026 Updates: Compliance Requirements and Business Risks." Legal Nodes, 2026.

[19] Cloud Security Alliance. "Building EU AI Act Compliance with prEN 18286 and ISO 42001." CSA Blog, April 2026.

[20] DLA Piper. "Latest wave of obligations under the EU AI Act take effect: Key considerations." DLA Piper, August 2025.

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