CSAI Foundation | Cloud Security Alliance

Fragmented Governance, Misread Mandate: ISO 42001 and the EU AI Act

Eighteen Uncoordinated Coalitions and a Growing AI Compliance Misconception

2026-09-08

AI-assisted Rapid Research

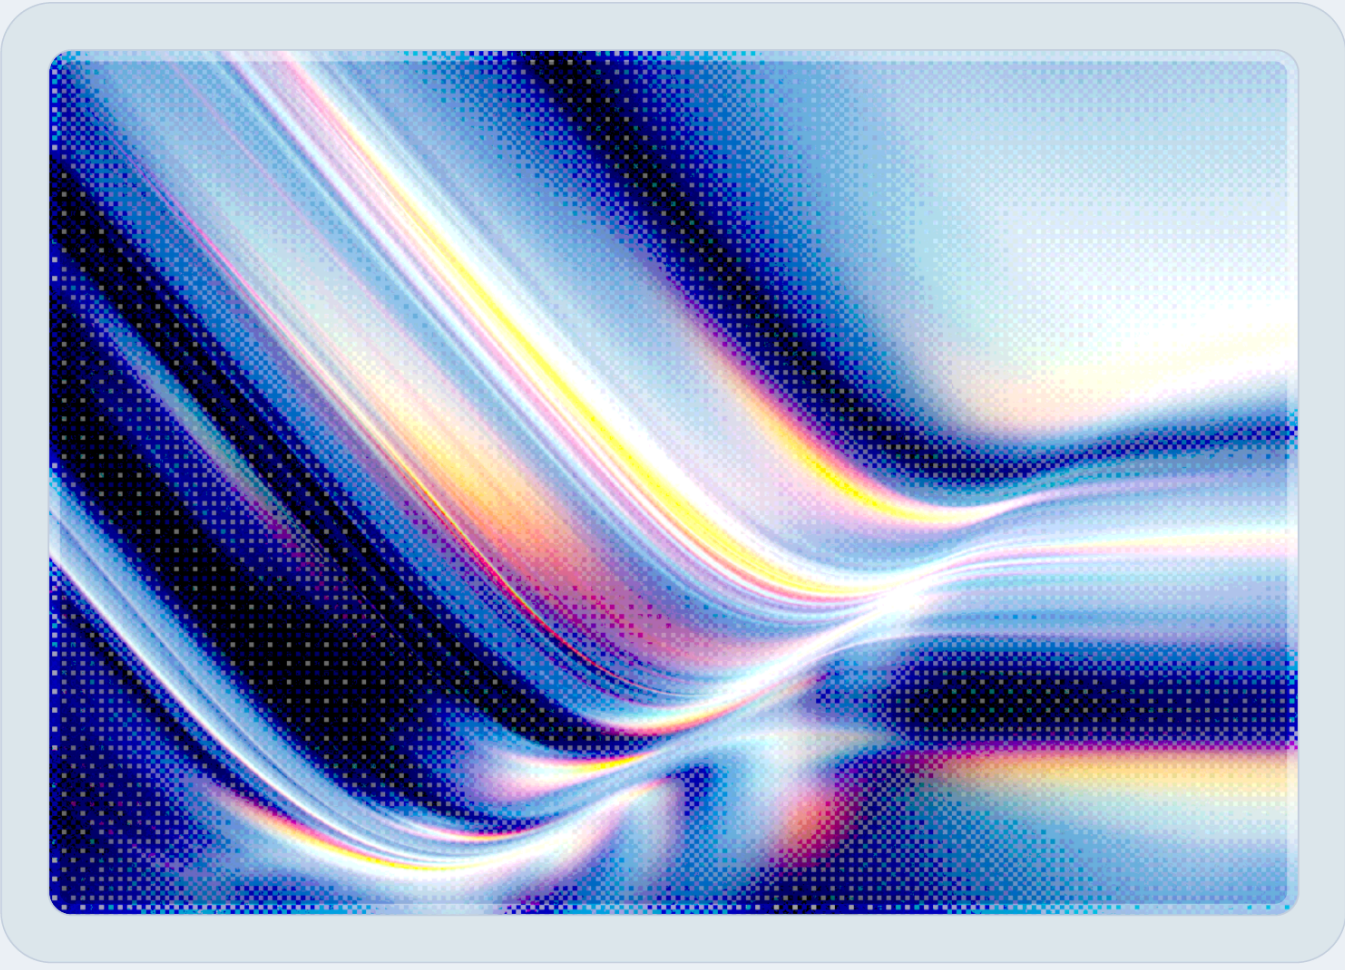

CSAI

CSA cloud security alliance®

Key Takeaways

- An August 2026 analysis by AI agent security firm Zenity counted at least 18 distinct AI security governance initiatives launched between April and August 2026 alone, with seven of them addressing agentic AI security independently and "no visible coordination between most of them" [1].
- A small number of organizations — chiefly Google, Microsoft, and Anthropic — sit inside a large share of these 18 efforts simultaneously, meaning the appearance of broad industry consensus often reflects the overlapping membership of a handful of companies rather than independent convergence [1].
- Against that backdrop, construction ERP vendor CMiC announced ISO/IEC 42001 certification of its NEXUS AI chatbot on September 1, 2026, framing the achievement in terms of "robust frameworks for security, transparency, and accountability" — language that, absent additional context, invites the inference that the certification also signals EU AI Act readiness, even though CMiC's announcement makes no such claim [3].
- That inference is not accurate: ISO/IEC 42001 has not been cited in the Official Journal of the European Union and does not itself create a legal presumption of conformity under the EU AI Act. The harmonized standard actually built for the Act's Article 17 quality-management requirement, prEN 18286, cleared CEN-CENELEC's formal vote and was approved as EN 18286:2026 on July 12, 2026 — the first AI Act-supporting standard to reach that milestone — but as of this note's September 2026 date it still awaits the Official Journal citation that would trigger the Article 40 presumption [4][6].
- A second Zenity piece published September 2, 2026 argues that the security industry has separately hollowed out the word "governance" itself, reducing it to point-in-time, permission-based access controls rather than continuous evaluation of whether an agent's unfolding actions still serve its intended purpose — a conceptual gap that compounds the practical one created by overlapping coalitions and misread certifications [2].

Background

Two developments converged in late summer 2026 to illustrate how crowded and confusing the AI governance landscape has become for the enterprises trying to navigate it. The first is structural: Taylor Roberts' analysis for Zenity, published August 11, 2026, catalogued at least 18 separate AI security governance initiatives that had launched or materially advanced in just the four months prior, spanning infrastructure defense efforts such as Anthropic's Project Glasswing, standards bodies such as the Coalition for Secure AI (CoSAI) and the OWASP Gen AI Security Project, multilateral government efforts such as the International Network of AI Safety Institutes, and open-source initiatives such as the NVIDIA-founded Open Secure AI Alliance [1]. Roberts found meaningful overlap in some areas — model and system security, and agentic AI security, are each addressed by seven of the 18 efforts — but almost no visible coordination between them, and warned that this pattern is "likely to produce incompatible standards rather than a converged one unless something changes" [1]. Other risk categories fared worse: trust and safety content moderation is addressed by only one initiative (ROOST), and CBRN or catastrophic-risk coordination and certification/compliance work each show up in only two or three of the 18 [1].

The second development is narrower but more concrete, and it shows what fragmentation looks like from inside a single vendor's compliance decision. On September 1, 2026, CMiC — a provider of AI-powered enterprise resource planning software for the construction industry — announced that its "AL" analytics and reporting chatbot, part of the company's NEXUS platform, had earned ISO/IEC 42001:2023 certification of its AI management system from the accredited certification body Schellman [3]. CMiC's chief information security officer, Joseph Kwok, described the milestone as reflecting "AI capabilities that are not only innovative, but also governed by robust frameworks for security, transparency, and accountability," and Schellman's Danny Manimbo called it evidence of "the kind of rigorous AI governance that construction technology providers need as AI becomes embedded in mission-critical workflows" [3]. Nothing in the announcement itself claims EU AI Act conformity. But certification language of this kind, similar to phrasing other AI vendors have used in 2026 ISO 42001 announcements, sits close to a misunderstanding this note has observed repeatedly in enterprise governance conversations this year: that certifying an AI management system to ISO/IEC 42001 satisfies, or substantially advances, compliance with the EU AI Act's binding legal requirements for high-risk AI systems.

It does not, and understanding why requires separating two things that AI governance coalitions and marketing copy alike tend to conflate: an organizational management-system standard, and a legally recognized harmonized standard tied to a specific regulation's conformity-assessment mechanism.

Security Analysis

Why 18 coalitions is a coordination problem, not a coverage problem

The proliferation Zenity documented is not, on its own, evidence that AI governance is under-resourced. If anything, the opposite risk is more pressing: incompatible standards can lock in before anyone has the chance to reconcile them, particularly in agentic AI security, where seven initiatives are working in parallel with little visible cross-reference to one another [1]. Because a handful of large model developers hold seats across most of these efforts, an enterprise security team surveying the landscape can reasonably mistake broad institutional participation for genuine technical consensus, when in practice the same three or four companies are simply repeating similar positions inside different rooms. Zenity's recommended remedy is deliberately modest: rather than calling for consolidation into a single body, Roberts argues that existing groups should "clearly articulate intended objectives, outcomes, and coordination channels between efforts." Roberts' analysis counts seven initiatives touching agentic AI security in aggregate, but the specific coordination recommendation targets a narrower set: "the six initiatives working on agentic AI security" should converge on a joint industry position "before their standards harden" into mutually incompatible technical requirements [1]. For enterprise governance teams, the practical implication is that no single coalition or framework announcement should be treated as authoritative simply because of who is in the room; cross-referencing which specific technical requirements a given initiative actually produces, and how those requirements map to controls a security program can implement, matters more than counting member logos.

The ISO 42001 mandate misconception

The EU AI Act does not designate ISO/IEC 42001 as a route to legal compliance, and the distinction between what the standard actually provides and what buyers and even some vendors assume it provides has become one of the more consequential points of confusion in enterprise AI governance conversations during 2026. ISO/IEC 42001, published in December 2023, specifies requirements for an organizational AI management system (AIMS) — the governance structures, roles, risk processes, and continual-improvement mechanisms an organization uses to manage AI responsibly across its portfolio [3]. The growth in certifications against it throughout 2026 reflects real organizational investment in AI governance maturity, even where that investment does not translate into EU AI Act conformity. The EU AI Act's presumption-of-conformity mechanism, set out in Article 40, applies only to standards that the European Commission has requested, that CEN and CENELEC have developed and formally voted through, and that the Commission has subsequently cited in the Official Journal of the European Union; ISO/IEC 42001 has not gone through that process and carries no such legal status under the Act [4].

The standard actually built for that purpose is prEN 18286, "Artificial Intelligence – Quality Management System for EU AI Act Regulatory Purposes," which entered public enquiry in late October 2025 as the first AI-specific standard in the EU's harmonization process [4][5]. Its consultation window closed in early 2026, and it cleared CEN-CENELEC's formal vote on July 12, 2026, publishing as EN 18286:2026 [6]. What has not yet happened, as of this note's September 2026 date, is citation in the Official Journal of the European Union, the specific event that triggers the Article 40 presumption of conformity; until that citation occurs, EN 18286:2026 carries the same non-binding status ISO/IEC 42001 does, even though it was purpose-built for the Act [6]. Where ISO/IEC 42001 addresses organizational governance in general terms, EN 18286:2026 is built specifically around Article 17 of the AI Act, which enumerates thirteen distinct mandatory elements a high-risk AI provider's quality management system must cover – including a regulatory compliance strategy tied to change management, integration with the Article 9 risk-management system, and adherence to the Article 73 serious-incident reporting windows of two, ten, or fifteen days depending on severity [9][10]. ISO/IEC 42001 does not natively address several of these Article-17-specific obligations, meaning that an organization holding only ISO 42001 certification retains the full evidentiary burden of demonstrating AI Act conformity if challenged, rather than benefiting from any legal presumption.

That regulatory timeline has also been in motion this year. The EU's Digital Omnibus on AI, which the European Parliament endorsed on June 16, 2026 by a vote of 423 to 57 with 174 abstentions and the Council of the EU approved on June 29, 2026, deferred the compliance deadline for standalone high-risk AI systems under Annex III from August 2, 2026 to December 2, 2027, and for AI embedded in already-regulated products under Annex I to August 2, 2028 [7]. The resulting Regulation (EU) 2026/1744 entered into force on July 27, 2026, and the deferral was granted explicitly because the harmonized standards and conformity-assessment tools that high-risk providers need – prEN 18286/EN 18286:2026 chief among them – were not yet ready, not because the underlying obligations were abandoned [8]. This deferral applies only to the high-risk-system deadlines it names; it does not touch obligations that took effect on an entirely earlier schedule. The Article 5 prohibited-practices regime became applicable on February 2, 2025, and general-purpose AI provider requirements under Articles 51-56 took effect August 2, 2025 – both well before the Digital Omnibus was even proposed [11][12]. Article 50 transparency duties, by contrast, did remain on their original August 2, 2026 date and are unaffected by the delay. Vendors and buyers who read the deadline extension as a general reprieve covering every AI Act obligation, or who treat an ISO 42001 certificate earned in the interim as a substitute for the forthcoming harmonized standard, are misreading both documents.

When "governance" becomes a permission check

A related and more conceptual critique came from Zenity's Ben Hanson on September 2, 2026, arguing that the security industry has quietly narrowed the meaning of "governance" itself, particularly in the context of autonomous AI agents [2]. Hanson's argument is that most tools marketed as agentic governance — identity platforms, network segmentation, API gateways — answer only whether an agent technically *can* take an action, evaluated as an isolated, point-in-time transaction, rather than whether it *should*, evaluated across the full chain of an unfolding execution and in light of its stated objective [2]. He illustrates the gap with an incident in which a sequence of individually permitted actions combined to produce an outcome no single access check would have flagged, describing genuine agentic governance as "continuously determining whether an unfolding execution remains consistent with what the system is trying to accomplish" [2]. This is a useful frame for evaluating both problems described above: a coalition's published standard and a vendor's management-system certificate can each satisfy a checklist — a "can" question — while leaving the harder "should" question, about whether an agent's behavior over time still serves its intended purpose, effectively unaddressed. Fragmentation among coalitions and conflation of governance standards are not separate problems from Hanson's critique; they are two ways the same underlying tendency — mistaking a point-in-time compliance artifact for continuous, substantive oversight — expresses itself at the industry and product level, respectively.

Recommendations

Immediate Actions

Security and compliance teams should audit any internal or vendor communications that describe ISO/IEC 42001 certification as satisfying EU AI Act obligations, and correct that framing before it reaches customers, auditors, or regulators; the certificate demonstrates organizational AI governance maturity, not legal conformity under Article 40. Teams operating or procuring high-risk AI systems in scope for the AI Act should confirm their actual compliance deadline against the Digital Omnibus's revised schedule — December 2, 2027 for standalone Annex III systems and August 2, 2028 for Annex I embedded systems — rather than assuming a blanket delay covers every AI Act obligation. Article 50 transparency duties remained on their original August 2, 2026 date, and the Article 5 prohibited-practices regime and general-purpose AI provider requirements were already in force well before the Digital Omnibus was even proposed, having taken effect in February and August 2025 respectively [7][8][11][12].

Short-Term Mitigations

Organizations should track EN 18286:2026's remaining path to Official Journal citation — now that CEN-CENELEC has approved the standard, that citation is the specific event that will trigger the Article 40 presumption of conformity that ISO/IEC 42001 alone cannot provide [4][6]. Where an organization already holds or is pursuing ISO/IEC 42001 certification, it should map that certification's coverage against Article 17's thirteen mandatory elements to identify gaps — particularly around Article 73 incident-reporting timelines and Article 9 risk-management integration — rather than treating the ISO certificate as a finished compliance artifact [9][10]. When evaluating claims from any of the 18-plus coalitions Zenity identified, security teams should ask what specific, implementable technical requirement a given initiative has actually produced, rather than weighting the initiative by its list of member organizations, since a small set of companies participates across most of them.

Strategic Considerations

At a program level, enterprises should build AI governance around a control framework that can absorb multiple external standards as they mature, rather than betting a compliance program on any single coalition's output or any single certification's assumed legal weight. CSA's AI Controls Matrix is designed for exactly that layered role, mapping to ISO/IEC 42001, prEN 18286's emerging structure, and the EU AI Act simultaneously so that a single control set can satisfy multiple external frameworks as they each reach maturity. Security leaders should also treat Hanson's "can versus should" distinction as a procurement question: before adopting an agentic governance tool, or crediting a vendor's compliance certification as sufficient assurance, ask whether it evaluates ongoing agent behavior against intended purpose, or merely confirms that a permitted action was technically permitted at a single point in time.

CSA Resource Alignment

CSA has already produced two research notes that speak directly to the regulatory confusion this document describes. **EU AI Act Compliance: prEN 18286 and ISO 42001** [13], published April 28, 2026, lays out the exact gap analyzed above in detail — the thirteen mandatory elements of Article 17, the five specific areas where ISO/IEC 42001 does not satisfy them, and a recommended three-tier architecture combining ISO/IEC 42001 for organizational governance, prEN 18286 for per-system conformity, and CSA's AI Controls Matrix (AICM) as the technical control overlay tying both together. **EU AI Act's High-Risk Deadline: Deferred, Not Cancelled** [14], published August 1, 2026, provides the authoritative account of the Digital Omnibus deadline changes summarized in this note's Background section, including which obligations remained on their original schedule. Organizations evaluating a

certification pathway — the specific decision CMiC's announcement represents — can also consult CSA's **Decision Tree Workflow for ISO 27001 and ISO 42001 Paths** [15], published March 5, 2026, which maps how an existing ISO 27001 or ISO 42001 certification feeds into CSA's own STAR for AI certification track.

More broadly, this note's core recommendation — that a layered, mappable control set is more durable than reliance on any single coalition or certification — is the organizing premise of the **AI Controls Matrix (AICM v1.1)** [16], which maps its controls across ISO/IEC 42001, NIST AI frameworks, and the EU AI Act so that governance teams have a stable technical baseline even as the 18-plus coalitions and evolving harmonized standards described in this note continue to shift around it.

References

[1] Zenity. "Coalition Chaos." Zenity Blog, August 11, 2026.

[2] Zenity. "Governance Strikes Back: The Most Used, Most Abused Word in the Galaxy." Zenity Blog, September 2, 2026.

[3] GlobeNewswire. "CMiC Earns ISO 42001 Certification, Reinforcing the Value and Rigor Behind Its AI-Powered NEXUS Platform." GlobeNewswire, September 1, 2026.

[4] European Commission. "Standardisation of the AI Act." Shaping Europe's Digital Future, 2026.

[5] CEN-CENELEC. "Update on CEN and CENELEC's Decision to Accelerate the Development of Standards for Artificial Intelligence." CEN-CENELEC News, October 23, 2025.

[6] Modulos. "EN 18286: Quality Management System for EU AI Act." Modulos Docs, July 2026.

[7] European Parliament. "AI Act: EP Approves Simplification Measures and 'Nudifier' App Ban." European Parliament Press Room, June 16, 2026.

[8] Hunton Andrews Kurth. "EU Digital Omnibus on AI Enters Into Force." Privacy and Cybersecurity Law Blog, July 28, 2026.

[9] EU Artificial Intelligence Act. "Article 17: Quality Management System." artificialintelligenceact.eu, 2026.

[10] EU Artificial Intelligence Act. "Article 73: Reporting of Serious Incidents." artificialintelligenceact.eu, 2026.

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

[12] Mayer Brown. "EU AI Act: Ban on Certain AI Practices and Requirements for AI Literacy Come Into Effect." Mayer Brown Insights, January 2025.

[13] Cloud Security Alliance AI Safety Initiative. "EU AI Act Compliance: prEN 18286 and ISO 42001." CSA Labs, April 28, 2026.

[14] Cloud Security Alliance AI Safety Initiative. "EU AI Act's High-Risk Deadline: Deferred, Not Cancelled." CSA Labs, August 1, 2026.

[15] Cloud Security Alliance. "Decision Tree Workflow for ISO 27001 and ISO 42001 Paths." Cloud Security Alliance, March 5, 2026.

[16] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.