CSAI Foundation | Cloud Security Alliance

# Who Assures the Assurers? Third-Party AI Assessment Standards

Security Implications and Guidance

2026-09-27

AI-assisted Rapid Research

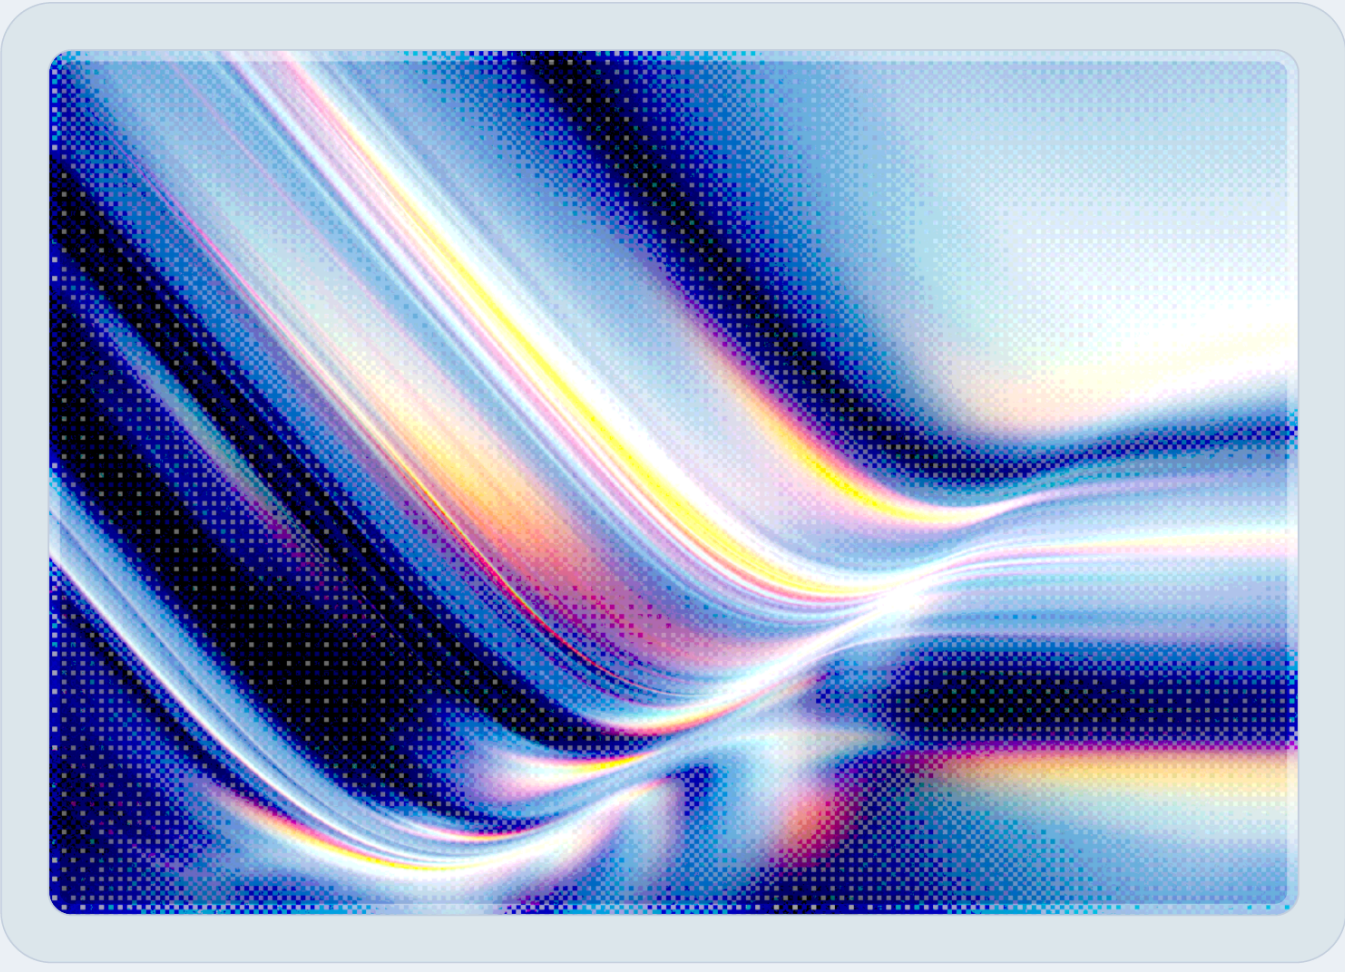

CSAI

CSA cloud security alliance®

## Key Takeaways

Over the past eighteen months, the AI industry has built the visible half of a third-party assurance system — management-system standards such as ISO/IEC 42001, certification schemes such as CSA's STAR for AI, and agent-specific programs such as AIUC-1 — faster than it has built the infrastructure that assures the assurers themselves [1][2]. That infrastructure gap is now closing in specific, documented ways: ISO published ISO/IEC 42006 in 2025 to set competence and independence requirements for the bodies that audit AI management systems, CSA's own STAR Certification requirements govern which certifying bodies may issue a STAR credential, and AIUC-1 named Schellman as its first accredited auditor in February 2026, giving that agent-security standard its first outside check [3][4][5]. The European Union's AI Act, meanwhile, illustrates how far behind an accreditation ecosystem can lag its underlying law: as of early 2026, the compliance-automation firm eyreACT reported that only three EU member states had fully designated both the notifying and market-surveillance authorities the Act requires, and that very few, if any, bodies had been fully designated specifically to perform conformity assessments for high-risk AI systems ahead of the Act's August 2026 application date, leaving providers with few qualified third parties to turn to even where the law already requires one [6][15].

This matters because every one of these schemes ultimately asks an enterprise to trust a certificate, and a certificate is only as trustworthy as the body that issued it. The UK government's 2024 market study found the AI assurance supplier base had grown from 17 specialized firms in 2023 to 84 in a single year, embedded within a broader ecosystem of 524 firms offering some form of AI assurance service [7]. In this note's assessment, that growth rate outpaces the capacity of any single accreditation body to vet new entrants. Security and compliance leaders evaluating a vendor's ISO/IEC 42001 certificate, STAR for AI listing, or AIUC-1 badge should treat the identity and accreditation status of the underlying auditor as a first-order due diligence question, not an afterthought, because the assurance market is still assembling the mechanisms that would let a buyer take that trust for granted.

## Background

The current generation of AI assurance standards emerged to answer a narrow but urgent question: how does an enterprise verify a vendor's claims about how it develops, deploys, and governs an AI system, when the enterprise itself cannot inspect the vendor's models, data pipelines, or internal controls directly? ISO/IEC 42001, published in December 2023, answered this by defining an AI management

system (AIMS) — a documented structure of policies, risk assessments, and control processes an organization operates and that an accredited third-party auditor can inspect and certify, following the same conformity-assessment logic long used for information security management under ISO/IEC 27001 [1]. Certification is not a one-time event: it requires a two-stage initial audit followed by annual surveillance audits over a three-year certificate cycle [1]. Market data on ISO/IEC 42001 engagements indicates that organizations starting from scratch commonly face six-to-twelve-month timelines and certification-body audit fees in the tens of thousands of dollars, reflecting the depth of documentation and evidence an AIMS audit expects, though organizations with an existing ISO/IEC 27001 certification can often shorten this to four to six months by reusing existing documentation and processes [16].

CSA extended this model with STAR for AI, which layers a public registry and a graduated assurance structure on top of existing standards rather than replacing them. Organizations can earn STAR for AI Level 1 by submitting a self-assessment against the AI Consensus Assessments Initiative Questionnaire (AI-CAIQ), with an enhanced Level 1 available once that self-assessment passes CSA's automated Valid-AI-ted scoring system; Level 2, introduced in November 2025, requires both a Valid-AI-ted AI-CAIQ and an accredited third-party ISO/IEC 42001 certification, combining a structured self-attestation with independent audit evidence [8]. The AI-CAIQ and the STAR for AI structure sit on top of CSA's AI Controls Matrix (AICM), which maps to ISO/IEC 42001, ISO/IEC 27001, NIST's AI Risk Management Framework, and other national and sectoral frameworks, giving the registry a common reference point across the fragmented standards landscape [8]. In June 2026, CSA extended STAR's coverage further into agentic systems by adding AIUC-1 — a certification developed with Orrick, MITRE, and a consortium of enterprise security leaders that has grown from roughly 60 members at its 2025 launch to more than 250 Fortune 1000 CISOs and security leaders by mid-2026, and covering security, safety, reliability, data privacy, and accountability for AI agents — to the STAR Registry, and CSA is separately developing a STAR for AI Agentic Certification Scheme targeting pilot enrollment and registry integration during 2026 [9][10][17].

What distinguishes this current wave from earlier compliance frameworks is that several of the new schemes were explicitly designed with the auditor problem in mind from the outset, rather than treating it as an afterthought once certificates began circulating. AIUC-1 pairs its certification with underwriting: a certified agent's conformance can be backed by Lloyd's of London-affiliated insurance covering losses up to $50 million, as in the policy backing ElevenLabs' AIUC-1 certification, which gives the auditor's work a financial consequence beyond reputational risk if a certified system fails in the field [17]. NIST's AI RMF took a different path, deliberately declining to define a formal certification or conformity-assessment process at all, leaving enterprises seeking certifiable assurance to combine RMF-aligned practices with a certifiable standard such as ISO/IEC 42001 [11]. These design choices reflect a genuine unresolved question in the field: whether AI assurance should follow the accredited-auditor model

inherited from information security and financial audit, or whether newer mechanisms such as insurance-backed certification and automated scoring systems can substitute for or supplement traditional third-party audit.

## Security Analysis

The clearest evidence that the industry recognizes an auditor-quality gap is the direct standards response to it. ISO published ISO/IEC 42006 in 2025 specifically to govern the bodies that audit and certify AI management systems against ISO/IEC 42001, building on the existing ISO/IEC 17021-1 framework for management-system certification bodies and adding AI-specific competence requirements for auditors, clearer rules on liability and audit time calculation, and stricter access requirements to an organization's AI documentation [2][12]. The standard is moving into national accreditation programs during 2026, meaning that for most of the period since ISO/IEC 42001's 2023 publication, organizations claiming certification were being audited by bodies without a uniform, AI-specific competence standard governing their auditors — a roughly two-year sequencing gap between the management-system standard and its audit-quality standard [2]. CSA's own certification ecosystem shows a parallel structure: its March 2025 "Requirements for Bodies Providing STAR Certification" document sets out the competency standards a certifying body must meet to issue STAR credentials aligned with the Cloud Controls Matrix as part of an ISO/IEC 27001 assessment, explicitly positioning STAR certification as supplementary to, rather than a replacement for, the ISO/IEC 27006 accreditation framework that governs ISO/IEC 27001 certifying bodies [3].

The regulatory picture is less mature than the voluntary standards picture, and the gap carries direct compliance risk for organizations subject to the EU AI Act. The Act requires conformity assessment — in many cases through an independent notified body — for high-risk AI systems, including those embedded in products that already require third-party certification under EU product safety law or that perform remote biometric identification [6][13]. Notified bodies must be independent of the AI system providers whose products they assess, and the governance infrastructure surrounding this requirement, including the notification and designation process itself, was required to be operational from August 2025 ahead of the broader Act provisions applying from August 2026 [6]. As of early 2026, however, the compliance-automation firm eyreACT reported that very few, if any, bodies had been fully designated specifically for AI Act conformity assessment work, with designation processes still active across most member states and assessment capacity limited relative to the hundreds of high-risk systems expected to require review [6][15]. An organization that determines its system requires notified-body assessment under the Act may therefore find that qualified capacity does not yet exist at the scale demand requires, a supply-side constraint that has no equivalent yet in the more mature information-security certification market.

The market-structure evidence points to a related but distinct risk: rapid, largely unregulated entry into AI assurance services. The UK Department for Science, Innovation and Technology's November 2024 study identified 524 firms in the UK supplying some form of AI assurance good or service, including 84 specialized AI assurance providers — up from 17 specialized firms identified in a 2023 government study — with the specialized segment alone generating an estimated portion of a combined £1.01 billion in market revenue [7]. That growth rate is an indicator of market demand, but it also means the accreditation and competence infrastructure described above is chasing a rapidly expanding supplier base rather than gating entry to it in advance, and academic research on AI audit tooling has separately identified structural gaps in the methods and tools available to auditors even at established firms, independent of any individual auditor's credentials [14]. Enterprises relying on a vendor's assurance credential should therefore distinguish between certificates issued under a mature, accredited audit regime — such as ISO/IEC 42001 certificates issued by a body operating under ISO/IEC 42006 once that accreditation is in force — and certificates or attestations issued by firms whose own qualifications have not been independently verified against any comparable standard.

# Recommendations

## Immediate Actions

Security and procurement teams evaluating a vendor's AI assurance claims should verify the accreditation status of the certifying body itself, not only the existence of the certificate, by confirming through the relevant national accreditation body or the certification scheme's own registry (such as CSA's STAR Registry) that the auditor was authorized to issue that specific credential. Where a vendor presents an ISO/IEC 42001 certificate, teams should ask which accredited certification body performed the audit and whether that body's accreditation already reflects ISO/IEC 42006 requirements, since accreditation against the older, more general ISO/IEC 17021-1 baseline alone does not guarantee AI-specific auditor competence. Organizations should also confirm whether a presented credential reflects a genuine third-party audit — such as STAR for AI Level 2 or an accredited AIUC-1 certification — versus a self-assessment such as STAR for AI Level 1, since the two carry materially different assurance weight despite appearing in the same registry.

## Short-Term Mitigations

Enterprises with AI systems in scope for EU AI Act conformity assessment should begin engaging with candidate notified bodies now rather than waiting until closer to enforcement deadlines, given the documented capacity constraints in the current designation landscape, and should build contingency

time into compliance timelines to account for the possibility that a preferred notified body has a backlog or has not yet been designated for the specific system category involved [6][15]. Compliance and vendor-risk teams should update AI vendor questionnaires to explicitly request auditor accreditation details as a required field, rather than accepting a certificate name or logo as sufficient evidence, and should periodically re-verify that a previously accredited auditor's status remains current given how recently these accreditation programs have stood up. Where a vendor's assurance rests on an insurance-backed model such as AIUC-1, risk teams should request the specific policy terms and coverage limits rather than treating "insurance-backed" as a self-explanatory assurance signal.

### Strategic Considerations

Security leaders should expect the accreditation layer of AI assurance to continue maturing on a lag behind the underlying management-system and certification standards, mirroring the roughly two-year gap already observed between ISO/IEC 42001 and ISO/IEC 42006, and should build periodic reassessment of vendor certifications into ongoing third-party risk programs rather than treating a certificate obtained at onboarding as durable proof. Organizations that themselves seek AI certification — whether STAR for AI, ISO/IEC 42001, or AIUC-1 — should factor the auditor-selection decision into program planning, since choosing a certifying body accredited under the most current relevant standard is likely to keep the resulting certificate more durable as accreditation requirements continue to tighten. Finally, as CSA's STAR for AI Agentic Certification Scheme and comparable programs move from pilot to production during 2026, enterprises should track which certifying and accreditation bodies are approved to issue agentic-specific credentials, since agent-specific assurance is likely to inherit the same auditor-quality questions this note describes for management-system certification, on a compressed timeline.

## CSA Resource Alignment

CSA's Requirements for Bodies Providing STAR Certification is the most directly relevant CSA artifact to this note's central question, since it is CSA's own answer to "who assures the assurers" within the STAR ecosystem: it sets the competency standards a certifying body must meet to issue STAR credentials aligned with the Cloud Controls Matrix and positions that requirement as supplementary to the ISO/IEC 27006 accreditation framework governing ISO/IEC 27001 certifying bodies generally. Organizations evaluating any STAR-related credential, including STAR for AI, should treat this document as the baseline for what "accredited" should mean within the registry, and CSA's continued maintenance of this requirements document as STAR for AI and the agentic certification scheme mature will determine how well the registry avoids the auditor-quality gap this note identifies in the broader market.

CSA's STAR Certification Guidance Document: Auditing the Cloud Controls Matrix (CCM) and the related Evolution of STAR: Introducing Continuous Auditing framework provide the operational detail behind how STAR auditors are expected to conduct assessments, including the management-capability scoring model that STAR for AI's AI-CAIQ and Valid-AI-ted mechanisms build on, and the Level 3 continuous-auditing model these documents introduce is a plausible template for how STAR for AI could eventually reduce reliance on point-in-time audits as the AI assurance market matures. Enterprises should also connect this note's findings to CSA's AI Controls Matrix (AICM) v1.1, which underlies AI-CAIQ and STAR for AI and maps to ISO/IEC 42001, NIST's AI RMF, and other frameworks referenced throughout this note; an organization's AICM implementation is a reasonable internal reference point for evaluating whether a vendor's third-party-audited claims are actually consistent with the controls the vendor is being assessed against.

# References


[1] International Organization for Standardization. "ISO/IEC 42001:2023 – Information technology – Artificial intelligence – Management system." ISO, December 2023.

[2] International Organization for Standardization. "ISO/IEC 42006:2025 – Information technology – Artificial intelligence – Requirements for bodies providing audit and certification of artificial intelligence management systems." ISO, 2025.

[3] Cloud Security Alliance. "Requirements for Bodies Providing STAR Certification." Cloud Security Alliance, March 31, 2025.

[4] AIUC. "AIUC-1: AI Agent Security, Safety, and Reliability Standard." AIUC, 2025.

[5] Schellman. "Schellman Becomes the First Accredited Auditor for AIUC-1, the Security Standard for AI Agents." Schellman, February 3, 2026.

[6] European Artificial Intelligence Act. "Article 31: Requirements Relating to Notified Bodies." artificialintelligenceact.eu, accessed September 2026.

[7] UK Department for Science, Innovation and Technology. "Assuring a Responsible Future for AI: The UK's AI Assurance Market." DSIT, November 6, 2024.

[8] Cloud Security Alliance. "CSA STAR for AI." Cloud Security Alliance, 2026.

[9] Cloud Security Alliance. "Cloud Security Alliance Extends AI Assurance Leadership Into Agentic AI with Addition of AIUC-1 Certification to STAR Registry." Cloud Security Alliance, June 30, 2026.

[10] Cloud Security Alliance. "STAR for AI Agentic Certification Scheme." Cloud Security Alliance Labs, 2026.

[11] National Institute of Standards and Technology. "AI Risk Management Framework." NIST, accessed September 2026.

[12] Standards Council of Canada. "Transition to ISO/IEC 42006:2025 for Bodies Providing Audit and Certification of Artificial Intelligence Management Systems." SCC, 2025.

[13] European Artificial Intelligence Act. "Article 43: Conformity Assessment." artificialintelligenceact.eu, accessed September 2026.


 

[14] Ojewale, Victor, et al. "Towards AI Accountability Infrastructure: Gaps and Opportunities in AI Audit Tooling." arXiv, February 2024.

[15] eyreACT. "Notified Bodies Under the EU AI Act: The Gatekeepers You Haven't Met Yet." eyreACT, February 2, 2026.

[16] Elevate Consulting. "ISO 42001 Certification Cost Breakdown: What Enterprise AI Teams Pay in 2026." Elevate Consulting, March 17, 2026.

[17] Forkast News. "AIUC Raises $40M to Build the Certification and Insurance Layer That Makes Agent Governance Auditable." Forkast News, September 16, 2026.