CSAI Foundation | Cloud Security Alliance

ISO 42001 Stalls When Organizations Skip Role Mapping

Producer, Provider, and User Distinctions Drive Scope, Controls, and Audit Readiness

2026-07-25

AI-assisted Rapid Research

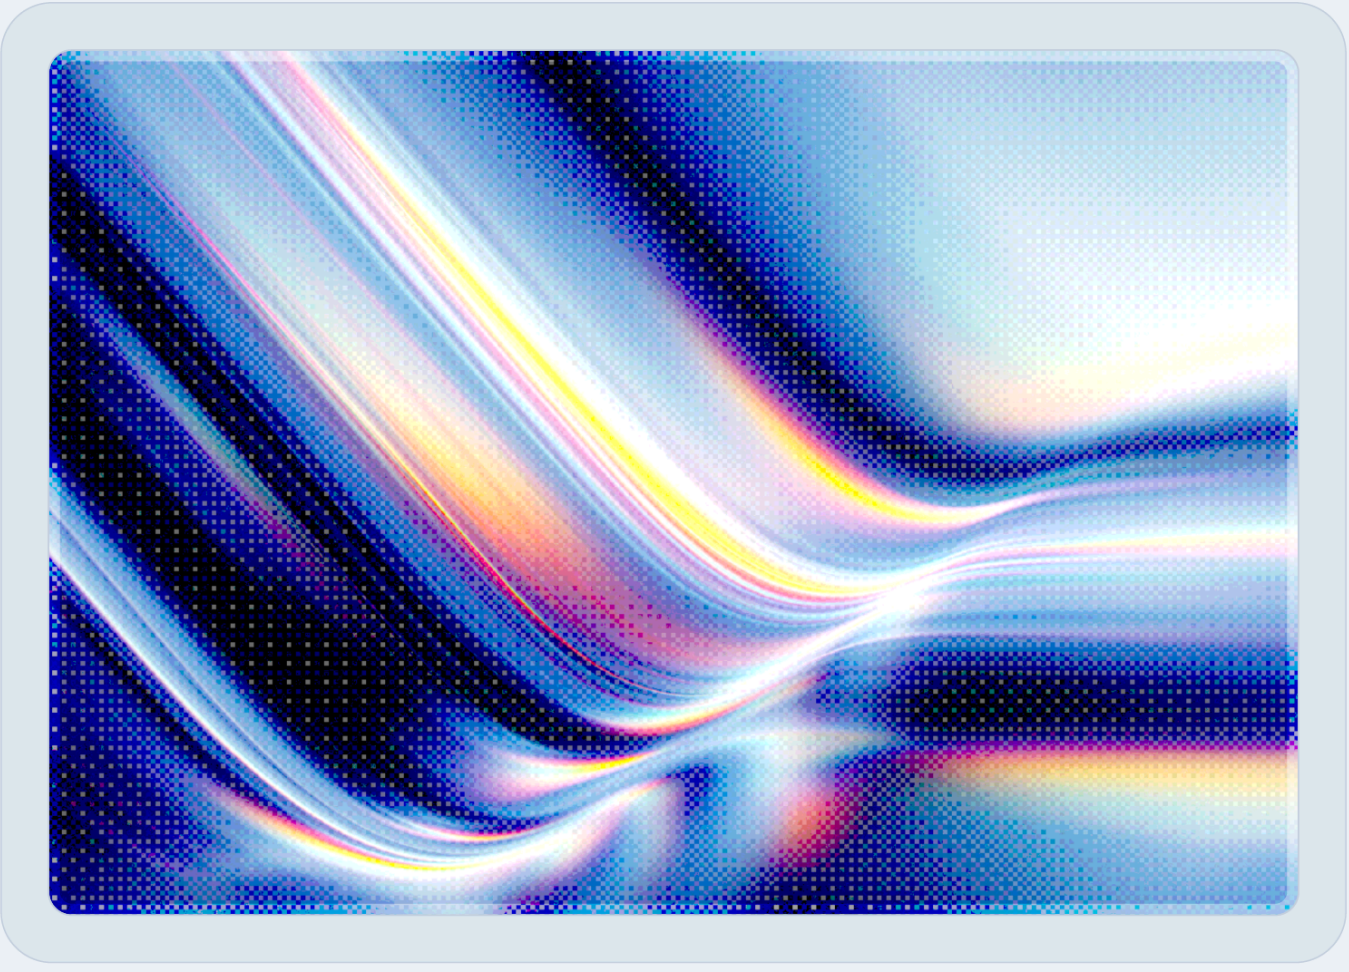

CSAI | CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

- ISO/IEC 42001 certification work can stall when organizations never formally establish which of the standard's roles — Producer, Provider, or User — they actually occupy; Clause 4.3 requires the certification scope to reflect every role the organization holds [1].
- Many organizations hold more than one role simultaneously — a firm that trains its own models while also reselling third-party AI tools is both a Producer and a Provider — and treating the organization as a single role produces a scope statement that omits activities auditors expect to see governed [1].
- Role confusion cascades into concrete audit findings, including incomplete AI system inventories and undocumented risk assessments [2][3]; a related risk not directly documented in these sources is that governance responsibilities end up assigned to senior stakeholders on paper rather than the people who actually operate the AI systems in scope.
- The EU AI Act's provider and deployer designations do not map one-to-one onto ISO 42001's roles. This note's own reading of the two frameworks suggests a Provider under the Act generally aligns with a Provider under ISO 42001, while a Deployer generally aligns with a User, though no published crosswalk formally establishes this — so organizations pursuing both frameworks should treat them as complementary systems requiring one integrated role determination rather than two parallel exercises [4].
- CSA's STAR for AI certification path is built on top of ISO/IEC 42001 and AI-CAIQ; because the STAR for AI submission inherits whatever scope and role determination an organization made for its underlying ISO 42001 certificate, an unresolved role question at the ISO 42001 stage propagates directly into the STAR for AI assessment [5][6].

Background

ISO/IEC 42001:2023 is the first internationally recognized management-system standard for governing artificial intelligence, and it follows the same ten-clause structure used by ISO 27001 and other ISO management-system standards: clauses 1 through 3 establish scope and definitions, while clauses 4 through 10 contain the auditable requirements covering context, leadership, planning, support, operation, performance evaluation, and improvement [7]. Clause 6.1.3 requires organizations to compare their selected controls against Annex A and justify any omissions, and clause 4.3 requires the

certification scope statement to name every organizational role the AI management system (AIMS) covers. That second requirement is where organizations that skip role determination tend to run into trouble.

ISO/IEC 22989, the terminology standard that ISO 42001 draws on, defines three roles an organization can hold with respect to an AI system: a Producer that designs and trains AI models, a Provider that offers AI systems as products or services to others, and a User that acquires and operates AI tools built by someone else [1]. These roles are not mutually exclusive, and CSA's own guidance on the subject notes that an organization building its own machine learning models while also licensing third-party tools functions as both a Provider and a User at once [1]. The practical difficulty is that most organizations begin their ISO 42001 project by inventorying AI systems and assigning owners, without first stepping back to ask which of these three roles applies to each system — and, as a result, without asking whether the scope statement they eventually submit to a certification body actually reflects the full set of roles their AI activities require them to govern.

CSA's earlier guidance on ISO 42001 requirements explains the standard's clause structure and its Clause 6.1.3 requirement to justify control selections against Annex A [7]. Applying that structure to the role taxonomy above, which controls are relevant, and how much evidence an auditor expects for each, depends heavily on whether the organization sits on the build side or the operate side of an AI system. An organization that never resolves its role tends to apply a generic subset of controls to everything, which produces the kind of incomplete inventory and undocumented risk assessment that can surface as a nonconformity at Stage 2 audit [2][3].

Security Analysis

Why Role Determines Scope, Not the Other Way Around

CSA's July 2026 guidance frames role identification as a precondition for scope definition rather than a byproduct of it: "Clause 4.3 requires the scope to reflect every role the organization holds" [1]. In practice, this means an organization cannot productively start by drawing a boundary around "the AI systems we use" and then work out what controls apply — it has to first determine, system by system, whether it is producing the model, providing it to others, or using someone else's, because each of those roles pulls in a different set of Annex A controls and a different evidentiary burden. An organization that skips this step risks discovering the gap only when an auditor asks for evidence — a data lineage trail, a human-oversight log, a supplier contract clause — that was never collected because nobody had determined which role made that evidence necessary in the first place.

For organizations acting as Producers or Providers, the center of gravity sits with the AI system itself. Risk assessments need to address algorithmic bias, training-data provenance, model drift, adversarial inputs, and explainability gaps, and the lifecycle documentation has to span design, data collection, training, validation, monitoring, retraining, and decommissioning [1]. CSA's guidance identifies data governance as a persistent audit gap specific to this side of the role split: Providers often cannot produce a documented bias analysis or a data lineage trail sufficient to satisfy an auditor, even when the underlying engineering practice is sound [1]. One plausible explanation is that the documentation obligation was never tied explicitly to the Provider role in the first place, though the sources reviewed here do not state that causal link directly.

For organizations acting as Users, the center of gravity shifts to operational use rather than system construction. Human oversight — Annex A control A.8.4 — becomes the primary control category, requiring documented procedures for reviewing AI-generated decisions, defined escalation paths, and retained override records [1]. CSA's guidance flags missing supplier contract provisions as the User-side counterpart to the Provider-side data-governance gap: an organization that has not identified itself as a User in scope may never think to negotiate contractual assurances from its AI vendor covering incident notification, model change disclosure, or audit cooperation, and that absence becomes visible only when an auditor asks how the organization manages a risk it does not control directly [1][2].

The EU AI Act Does Not Solve the Role Question — It Adds a Second One

Organizations operating in or selling into the European Union often pursue ISO 42001 certification alongside EU AI Act compliance, and it is tempting to assume the Act's provider and deployer designations settle the ISO 42001 role question by analogy. They do not, entirely. The EU AI Act assigns obligations based on an organization's position in the AI supply chain — provider or deployer — while ISO 42001's roles describe stakeholder responsibility within a management system that is, by design, role-agnostic in its clause structure [4]. Reading the two frameworks together, in most circumstances a Provider under the EU AI Act will also be a Provider under ISO 42001, and a Deployer under the Act generally corresponds to a User under ISO 42001, but "generally corresponds" is doing real work in that sentence: CSA has not published an official crosswalk between the two role systems, the two frameworks were not built by the same body, and an organization that assumes automatic equivalence risks importing gaps from one framework's role logic into the other's compliance obligations.

CSA's guidance on prEN 18286 — the draft European pre-standard developed by CEN-CENELEC Joint Technical Committee 21 to operationalize the EU AI Act's Article 17 requirements for high-risk systems — reinforces that both standards expect enterprise-wide governance with clear executive accountability rather than responsibility siloed inside a technical team, and that organizations should treat ISO 42001 and prEN 18286 as complementary systems to integrate rather than parallel initiatives to run separately

[4]. The practical implication for role determination is that an organization pursuing both frameworks should perform a single, sufficiently rigorous role-mapping exercise and apply its output to both compliance efforts, rather than letting each framework's audit deadline drive an independent, rushed determination.

How Role Ambiguity Surfaces at Audit

A related risk, though not directly documented in the sources reviewed for this note, is that vaguely defined roles can lead to governance accountability being assigned to senior stakeholders on paper while the people who actually operate and monitor the AI systems in scope have no documented role at all. That gap between documented accountability and operational control is exactly what an ISO 42001 Stage 2 auditor is trained to probe, because the standard is built around the premise that role clarity should be traceable from the scope statement down to individual system owners.

Scope problems compound the issue. Guidance on ISO 42001 scope definition identifies internal AI tools, AI features embedded in purchased software, and shadow AI usage as categories organizations should explicitly account for during inventory [3] — categories that are easy to overlook when nobody has first asked which roles the organization's AI footprint requires it to hold. As an illustration of that risk, an organization that has identified itself only as a User has less reason to look for the internally built models that would make it a Producer as well, and that blind spot can surface only when an external auditor asks a question the organization was not prepared for.

A structured role-determination process addresses both problems at once. Guidance on determining an organization's ISO 42001 role recommends mapping every AI-related activity across development, deployment, operation, and use; defining interested parties per Clause 4.1; cross-referencing each activity against the standard's role definitions; consulting stakeholders across departments rather than relying on a single team's inventory; and documenting the resulting roles directly in the scope statement and the AIMS itself [8]. Because roles can change as an organization's AI portfolio evolves — a User that begins fine-tuning a vendor model, for example, edges toward also being a Producer — this determination is not a one-time exercise but a periodic review tied to the AIMS's continuous-improvement cycle under Clause 10.

Recommendations

Immediate Actions

Organizations that have started or are about to start ISO 42001 implementation work should pause before finalizing a scope statement and conduct an explicit role-mapping exercise: list every AI system, tool, or feature the organization builds, resells, or consumes, and classify each against the Producer, Provider, and User definitions rather than assigning the organization a single label. Where an activity plausibly fits more than one role — a common outcome — document that overlap explicitly in the scope statement rather than resolving it by omission. Organizations already mid-audit and encountering unexpected control gaps should treat that experience as a signal to revisit the underlying role determination rather than patching the immediate finding in isolation, since a single unresolved role question tends to produce more than one downstream gap.

Short-Term Mitigations

Within the next audit cycle, organizations should close the two gaps CSA's guidance identifies as most persistent: Providers should establish a documented bias-analysis process and a data lineage trail for any model they design or train, and Users should review AI vendor contracts for provisions covering incident notification, model change disclosure, and audit cooperation. Organizations pursuing both ISO 42001 and EU AI Act alignment should confirm that a single role-mapping output feeds both compliance tracks, and should not assume that a Provider or Deployer designation made for the EU AI Act automatically satisfies the separate role determination ISO 42001's Clause 4.3 requires.

Strategic Considerations

As ISO 42001 certification becomes a more common procurement expectation and a prerequisite for CSA's STAR for AI designation, organizations should treat role determination as a governance capability that needs periodic review rather than a one-time scoping exercise completed at project kickoff. An organization's AI portfolio changes — new internally built models, new vendor relationships, new embedded AI features in purchased software — and each change has the potential to shift which roles apply, which controls are relevant, and what evidence an auditor will expect. Building role review into the AIMS's existing continuous-improvement cycle, rather than treating it as a pre-certification checkbox, is the more durable way to keep scope statements accurate as the organization's AI footprint grows.

CSA Resource Alignment

CSA's Decision Tree Workflow for ISO 27001 and ISO 42001 Paths is the most directly applicable CSA artifact to this topic: it is a visual decision-tree tool built specifically to guide organizations through certification pathway choices based on their existing ISO 27001 and ISO 42001 status, which is a natural companion to the role-determination process this note describes — an organization cannot select the right certification path until it knows which roles its scope statement needs to cover. CSA's Getting Ready for STAR for AI Level 2 makes the stakes of role ambiguity concrete for CSA's own certification program: STAR for AI Level 2 is earned by submitting an ISO/IEC 42001 certificate alongside a completed AI-CAIQ self-assessment, which means an ISO 42001 scope statement that omits a role the organization actually holds will produce an incomplete or inaccurate STAR for AI submission, not just an ISO 42001 nonconformity.

CSA's Requirements for Bodies Providing STAR for AI Certification specifies the mandatory procedures certification bodies must follow when assessing organizations against the AI Controls Matrix (AICM) on top of an ISO/IEC 27001 and/or ISO/IEC 42001 foundation, underscoring that the AICM audit inherits whatever role and scope determination the underlying ISO 42001 certificate encodes — a further reason to get that determination right at the source rather than treating it as an ISO-specific detail. Where an organization's questions extend beyond certification pathway mechanics into broader AI governance control selection, CSA's AI Controls Matrix (AICM) v1.1 remains the applicable control framework for mapping Producer-side and User-side obligations — such as data governance and human-oversight controls, respectively — onto a concrete control set once the role determination is complete.

References

[1] Cloud Security Alliance. "ISO 42001: The Importance of Knowing Your Role Before Building Your AI System." CSA Blog, July 21, 2026.

[2] Schellman. "The Roles and Responsibilities in ISO 42001 Explained." Schellman, 2026.

[3] Glocert International. "ISO 42001 Scope Definition: Inclusions & Exclusions." Glocert, 2026.

[4] Cloud Security Alliance. "Building EU AI Act Compliance with prEN 18286 and ISO 42001." CSA Blog, April 27, 2026.

[5] Cloud Security Alliance. "Getting Ready for STAR for AI Level 2." Cloud Security Alliance, 2026.

[6] Cloud Security Alliance. "Requirements for Bodies Providing STAR for AI Certification." Cloud Security Alliance, 2026.

[7] Cloud Security Alliance. "ISO 42001 Requirements Explained: What You Need for Compliance." CSA Blog, May 14, 2025.

[8] risk3sixty. "Determining Your ISO 42001 Role: A Guide for Organizations." risk3sixty, 2026.

© 2026 Cloud Security Alliance. AI-assisted Rapid Research.