Reliable Context

RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

ISO/IEC 42001 Explained: What an AI Management System Is

> Reliable Context is a kernel of truth: a system of record for organizations that fuse private and public information and must defend every value they act on. One canonical record per real thing, a cited source and an observation date behind every value, unknowns declared rather than left blank, and any past state of the record reproducible.

Published by Reliable Context · 2026-10-29 · iso 42001 · iso/iec 42001 · ai management system · iso 42001 certification · iso 42001 requirements

Original article: https://reliablecontext.com/articles/iso-iec-42001-ai-management-system-explained



In this article

1. What ISO/IEC 42001 is
2. What a management system standard is
3. The structure: clauses 4 to 10
4. The AI-specific content
5. Documented information and records
6. Relationship to other ISO/IEC standards
7. Certification versus conformity
8. Worked example: what an auditor might ask a data platform for
9. Common misreadings

What ISO/IEC 42001 is

ISO/IEC 42001:2023, *Information technology — Artificial intelligence — Management system*, was published in December 2023 by the joint ISO and IEC subcommittee on AI, ISO/IEC JTC 1/SC 42 (ISO, standard page). It "specifies the requirements and provides guidance for establishing, implementing, maintaining and continually improving an AI (artificial intelligence) management system within the context of an organization" (ISO OBP preview, clause 1). It applies to any organization "that provides or uses products or services that utilize AI systems", whatever its size or type (same source).

The full text is paywalled. This article uses only public material: ISO's product and explainer pages and the free preview on the ISO Online Browsing Platform (OBP), which shows the introduction, scope, definitions, table of contents and bibliography. For clauses 4 to 10 it reports titles, not requirement text.

What a management system standard is

A **management system** is a "set of interrelated or interacting elements of an organization to establish policies and objectives, as well as processes to achieve those objectives" (OBP, 3.4). A **policy** is the intentions and direction of an organization as formally expressed by its top management. An **objective** is a result to be achieved (3.5, 3.6, same source).

A management system standard (MSS) therefore governs how an organization decides, assigns, records and corrects. It does not test a product. ISO says of 42001: "Rather than looking at the details of specific AI applications, it provides a practical way of managing AI-related risks and opportunities across an organization", and it is implemented "using the Plan-Do-Check-Act methodology" (ISO, standard page).

For a procurement or risk team, this is the key point: a 42001 certificate describes a management system. It does not say a particular model is accurate, fair or safe.

The harmonized structure

ISO management system standards share a **harmonized structure** (HS). The ISO/IEC Directives describe it as "identical clause numbers in the same sequence, clause titles, text, common terms, core definitions" (ISO/IEC Directives, consolidated, Annex Q). This lets one organization run a single "integrated" system

against several standards (ISO, management system standards). The 42001 introduction states that it "applies the harmonized structure" (OBP, Introduction).

The structure: clauses 4 to 10

Clauses 1 to 3 cover scope, normative references and terms; the requirements sit in clauses 4 to 10. The titles below come from the public table of contents (OBP). The text under them is in the paid document.

Clause	Title	Subclauses (titles only)
4	Context of the organization	4.1 Understanding the organization and its context; 4.2 Understanding the needs and expectations of interested parties; 4.3 Determining the scope of the AI management system; 4.4 AI management system
5	Leadership	5.1 Leadership and commitment; 5.2 AI policy; 5.3 Roles, responsibilities and authorities
6	Planning	6.1 Actions to address risks and opportunities; 6.2 AI objectives and planning to achieve them; 6.3 Planning of changes
7	Support	7.1 Resources; 7.2 Competence; 7.3 Awareness; 7.4 Communication; 7.5 Documented information
8	Operation	8.1 Operational planning and control; 8.2 AI risk assessment; 8.3 AI risk treatment; 8.4 AI system impact assessment
9	Performance evaluation	9.1 Monitoring, measurement, analysis and evaluation; 9.2 Internal audit; 9.3 Management review
10	Improvement	10.1 Continual improvement; 10.2 Nonconformity and corrective action

Because the standard applies the harmonized structure, much of this skeleton is shared with other management system standards. AI appears explicitly in the titles of 4.3, 5.2, 6.2 and 8.2 to 8.4. ISO's explainer adds "data governance and system lifecycle controls" to its summary of the requirements (ISO, "ISO/IEC 42001 explained").

The AI-specific content

Why AI gets its own management system

The introduction gives three reasons AI "potentially raises specific considerations": AI can make automatic decisions "in a non-transparent and non-explainable way"; machine learning changes how systems are "developed, justified and deployed"; and systems that learn continuously "change their behaviour during use" (OBP, Introduction). Organizations "are expected to focus their application of requirements on features that are unique to AI", and they "can elect to apply these requirements using a risk-based approach" (same source).

The processes it names include managing trustworthiness concerns "such as security, safety, fairness, transparency, data quality and quality of AI systems throughout their life cycle", and managing "suppliers, partners and third parties that provide or develop AI systems for the organization" (same source).

Impact assessment

42001 defines an **AI system impact assessment** as a "formal, documented process by which the impacts on individuals, groups of individuals, or both, and societies are identified, evaluated and addressed" (OBP, 3.24). It has its own subclause, 8.4, separate from AI risk assessment in 8.2, and it points outward at people and society rather than at the organization. ISO/IEC 42005:2025 gives guidance on carrying out these assessments "throughout the AI system lifecycle" (ISO, ISO/IEC 42005).

Controls, Annex A and the statement of applicability

42001 has four annexes: Annex A, "Reference control objectives and controls"; Annex B, "Implementation guidance for AI controls"; Annex C, "Potential AI-related organizational objectives and risk sources"; and Annex D, "Use of the AI management system across domains or sectors" (OBP, table of contents). The Annex A control table is not in the public preview, so its controls are not listed here. The Annex B subclause titles show its guidance topics:

- B.2 Policies related to AI
- B.3 Internal organization
- B.4 Resources for AI systems
- B.5 Assessing impacts of AI systems
- B.6 AI system life cycle
- B.7 Data for AI systems
- B.8 Information for interested parties
- B.9 Use of AI systems
- B.10 Third-party and customer relationships

A **control** is a "measure that maintains and/or modifies risk" (OBP, 3.21). A **statement of applicability** (SoA) is "documentation of all necessary controls ... and justification for inclusion or exclusion of controls" (3.26). Its notes add that organizations "may not require all controls listed in Annex A or may even exceed the list", and that "all identified risks and the risk management measures (controls) established to address them shall be reflected in the statement of applicability" (3.26, Notes 1 and 2).

Data

42001 defines **data quality** as the "characteristic of data that the data meet the organization's data requirements for a specific context", taken from ISO/IEC 5259-1 (OBP, 3.25). The definition is relative: data are good when they meet requirements the organization has stated for a specific use. The introduction says 42001 "avoids specific guidance on management processes" and leaves organizations to draw on other frameworks and standards for risk, life cycle and data quality management (OBP, Introduction).

Documented information and records

Documented information is "information required to be controlled and maintained by an organization and the medium on which it is contained". It "can be in any format and media and from any source". It can refer to the management system itself, to documentation created to operate, and to "evidence of results achieved

(records)" (OBP, 3.10). Clause 7.5 carries that title.

The distinction decides what an audit can check. A policy says what should happen; a record shows what did happen, when, and on whose authority. The introduction states: "An organization conforming with the requirements in this document can generate evidence of its responsibility and accountability regarding its role with respect to AI systems" (same source). Such evidence depends on records that can be retrieved, dated and attributed: in data terms, on an audit trail and on provenance for the values a decision relied on.

Relationship to other ISO/IEC standards

Standard	What ISO says it does	Relation to 42001
ISO/IEC 22989:2022	Establishes AI terminology and describes AI concepts (ISO)	The only normative reference in 42001. Its terms apply alongside those in clause 3 (OBP)
ISO/IEC 23894:2023	Guidance on managing risk specifically related to AI (ISO)	Guidance, not certifiable requirements. Listed in the 42001 bibliography
ISO/IEC 42005:2025	Guidance for AI system impact assessments (ISO)	Supports the impact assessment in 8.4
ISO/IEC 27001	Requirements for an information security management system (ISO)	A sibling MSS on the same harmonized structure. ISO sells the two together
ISO/IEC 5259-1:2024	Overview, terminology and examples for data quality in analytics and ML (ISO)	Source of the 42001 data-quality definition
ISO/IEC 42006:2025	Additional requirements for bodies that audit and certify AI management systems (ISO)	Governs the certifiers, not the certified organization

ISO/IEC 27001 covers information security, the "preservation of confidentiality, integrity and availability of information" (OBP, 3.23), while 42001 covers AI. Sharing a structure, they can run as one integrated system; Annex D.2 of 42001 is titled "Integration of AI management system with other management system standards". 23894 and 42005 are guidance: certification is against 42001, not them.

The 42001 bibliography also lists the NIST AI Risk Management Framework 1.0 (OBP, Bibliography). NIST released it on 26 January 2023 for voluntary use (NIST); mapping the two is outside this article's scope.

Certification versus conformity

Conformity is "fulfilment of a requirement". A **nonconformity** is its non-fulfilment. An **audit** is a "systematic and independent process for obtaining evidence and evaluating it objectively to determine the extent to which the audit criteria are fulfilled", and it can be internal (first party) or external (second or third party) (OBP, 3.15, 3.16, 3.18). An organization can conform without being certified, and 42001 has its own internal audit subclause (9.2) either way.

Certification is a separate, voluntary step. ISO states that "ISO does not certify organizations" and that certification "is carried out by independent certification bodies, which may be accredited by national accreditation bodies" (ISO, "ISO/IEC 42001 explained"). **Accreditation** is formal recognition that a certification

body "operates according to international standards"; ISO points to the IAF CertSearch database to verify accredited certificates (ISO, Certification). Two standards govern the certifiers:

- ISO/IEC 17021-1:2015 sets requirements for the "competence, consistency and impartiality" of management system certification bodies (ISO, ISO/IEC 17021-1).
- ISO/IEC 42006:2025 "builds on ISO/IEC 17021-1" with AI-specific requirements for bodies certifying against 42001 (ISO, ISO/IEC 42006).

Reviewing a supplier's claim, ask: which body issued the certificate, is it accredited for this scheme, and what is the certified scope? Scope matters because "organization" in 42001 refers only to the part of a larger entity "that is within the scope of the AI management system" (OBP, 3.1, Note 2). A certificate may cover one business unit.

The EU AI Act

In the European Union, Regulation (EU) 2024/1689 (the AI Act) requires providers of high-risk AI systems to "put a quality management system in place". Article 17(1) lists aspects including "systems and procedures for data management", the risk management system, post-market monitoring, and "systems and procedures for record-keeping of all relevant documentation and information" (Regulation (EU) 2024/1689, Art. 17). Under Article 40(1), a presumption of conformity attaches to harmonised standards "the references of which have been published in the Official Journal of the European Union" (same source). As of 28 October 2026, the first harmonised standard referenced in the Official Journal under the AI Act is EN 18286:2026, *Artificial intelligence - Quality management system for EU AI Act regulatory purposes*, published through Commission Implementing Decision (EU) 2026/2274 and announced on 19 October 2026; it gives a presumption of conformity with the Article 17 requirements it covers, and the Commission describes it as designed to be layered on top of standards such as ISO/IEC 42001 rather than replace them (European Commission). ISO/IEC 42001 is not a harmonised standard under the AI Act, so 42001 certification should not be read as a presumption of conformity.

On timing, Regulation (EU) 2026/1744 (the Digital Omnibus on AI) entered into force on 27 July 2026. It sets the application of the high-risk requirements in Chapter III, Sections 1 to 3, to 2 December 2027 for systems classified under Article 6(2) and Annex III, and to 2 August 2028 for those under Article 6(1) and Annex I (Regulation (EU) 2026/1744, amending Art. 113). This describes the legal text and is not legal advice.

Worked example: what an auditor might ask a data platform for

Fictional example. Halden Therapeutics AG uses a machine-learning model to prioritize incoming adverse-event reports for human review. It sets the scope of its AIMS (4.3) to cover that system. Its AI risk assessment (8.2) identifies a risk that reports with missing fields are ranked too low. The treatment (8.3) is a control that sends any report below a completeness threshold to a reviewer, and the control appears in the SoA with its justification. The impact assessment (8.4) records the possible effect on patients if a serious report is delayed.

Six months later, an internal auditor (9.2) samples one delayed report. The table shows records a data layer could supply; it is an illustration mapped to clause titles, not the standard's requirements.

Auditor's question	Relevant clause title	Record a data layer can supply
Was this system inside the AIMS scope on that date?	4.3 Determining the scope	Dated scope document with version history
Which control applied, and was it in force?	8.3 AI risk treatment; SoA (3.26)	Control identifier, effective-from date, the SoA version that listed it
What input did the model see?	7.5 Documented information	The report as received, with source, receipt time and every later correction
Who changed the completeness threshold, and when?	6.3 Planning of changes	Append-only change log with author, timestamp and approval reference
Who could see or edit the record?	5.3 Roles, responsibilities and authorities	Access grants by role, consistent with need-to-know access
Was the nonconformity found and fixed?	10.2 Nonconformity and corrective action	Finding, root cause, corrective action and closure date, all linked

Each row asks: can the record be reproduced as it stood on that date, with source and author? If the platform overwrote the input or threshold, the audit can confirm only the current state.

Common misreadings

Misreading	What the public sources say
"ISO certified us."	ISO does not certify. Independent certification bodies do (ISO)
"42001 certifies our model."	It is a management system standard and does not look at "the details of specific AI applications" (ISO)
"Every Annex A control is mandatory."	Organizations may not need all Annex A controls and may add their own. Inclusion and exclusion are justified in the SoA (3.26)
"It only applies to AI developers."	It applies to organizations that provide or use AI-based products or services (clause 1)
"42001 certification equals AI Act compliance."	The AI Act's presumption attaches to harmonised standards published in the Official Journal (Art. 40(1)). The first such standard is EN 18286:2026 (QMS, Article 17); 42001 is not one

A management system is only as checkable as its records; the practical test for a data platform under an AIMS is whether a reader can trace a stated fact to its dated source and see who recorded it.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.iso.org/standard/81230.html
2. https://www.iso.org/obp/ui/en
3. https://www.iso.org/sites/directives/current/consolidated/index.html

4. https://www.iso.org/management-system-standards.html
5. https://www.iso.org/home/insights-news/resources/iso-42001-explained-what-it-is.html
6. https://www.iso.org/standard/42005
7. https://www.iso.org/standard/74296.html
8. https://www.iso.org/standard/77304.html
9. https://www.iso.org/standard/27001
10. https://www.iso.org/standard/81088.html
11. https://www.iso.org/standard/42006
12. https://www.nist.gov/itl/ai-risk-management-framework
13. https://www.iso.org/certification.html
14. https://www.iso.org/standard/61651.html
15. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
16. https://digital-strategy.ec.europa.eu/en/library/first-harmonised-standard-supporting-ai-act-published-official-journal
17. https://eur-lex.europa.eu/eli/reg/2026/1744/oj

THE PUBLISHER

About Reliable Context

Reliable Context is a system of record for organizations that fuse private and public information and must defend every value they act on.

The guarantees

One canonical record per real thing: duplicates resolve by rule, ambiguity is shown rather than hidden, and related organizations relate without ever being merged. A verbatim, cited source and an observation date sit behind every value. Unknowns are declared as not found, blocked or not applicable with a reason, never left as an empty cell. Any past state of the record is reproducible, together with what changed since.

Where it runs

It runs inside the customer's perimeter, one database per organization. Permitted sources, model routing, reader clearance and retention are set as the customer's own policies, and every value is exportable with its chain of custody for an auditor or regulator.

What it is not

It is not a vector database or a RAG framework, not a data catalog, MDM or data-quality suite, not a CRM or data-enrichment vendor, and not a chatbot or AI assistant. Where a reader needs one of those, the honest answer is to say so and point them there.

Visit

See Reliable Context.

Reliable Context is in production use in two sectors; more than six thousand of its identity decisions were made by rule and none has been reversed. Articles published here explain general practice, cite their primary and academic sources, and name the jurisdiction of any regulatory rule they describe. They are not legal, regulatory or compliance advice for a specific organization, and requirements change.

Website: https://reliablecontext.com

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.