Reliable Context

RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

ISO 42001 vs NIST AI RMF vs EU AI Act: One Evidence Base

> 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 vs nist ai rmf · ai governance frameworks comparison · ai act vs iso 42001 · ai rmf crosswalk · ai governance evidence

Original article: https://reliablecontext.com/articles/ai-governance-frameworks-compared-iso-42001-nist-ai-rmf-eu-ai-act



In this article

1. Three instruments that are often treated as one
2. What each instrument is
3. Scope and timing
4. Common ground: five things all three ask for
5. Where they diverge
6. Designing one evidence register
7. What none of them tests
8. Sequencing a programme
9. Summary matrix
10. Sources

Three instruments that are often treated as one

AI governance programmes in multinational organisations usually cite three references together: ISO/IEC 42001, the NIST AI Risk Management Framework, and the EU AI Act. They are routinely presented as a set of interchangeable "frameworks". They are not. One is an international management system standard, one is a voluntary framework from a US federal agency, and one is a regulation that creates legal obligations with dates of application. They differ in who they bind, what they ask for, and what evidence would satisfy them.

This article compares them on those three axes and then proposes a single evidence register: one set of records, kept once, that can answer the questions all three ask. It is educational, describes the texts as published, and is not legal advice.

What each instrument is

	ISO/IEC 42001:2023	NIST AI RMF 1.0 (AI 100-1)	EU AI Act, Regulation (EU) 2024/1689
Kind of instrument	International management system standard (MSS)	Voluntary framework	EU regulation, directly applicable in Member States
Publisher	ISO and IEC	US National Institute of Standards and Technology	European Parliament and Council
Date	Edition 1, December 2023 (ISO)	January 2023 (NIST AI 100-1)	In force 1 August 2024; amended by the Digital Omnibus on AI, Regulation (EU) 2026/1744 (EUR-Lex)
Unit of concern	The organisation's AI management system	AI risks across the AI lifecycle	AI systems and general-purpose AI models placed on the EU market or used in the EU
Binding force	None by itself; adopted voluntarily or required by contract	None; NIST calls it "voluntary, rights-preserving, non-sector-specific, and use-case agnostic"	Legal obligations on providers, deployers and others, enforced by authorities

	ISO/IEC 42001:2023	NIST AI RMF 1.0 (AI 100-1)	EU AI Act, Regulation (EU) 2024/1689
Structure	Plan-Do-Check-Act management system (ISO)	Four functions: Govern, Map, Measure, Manage, each with categories and subcategories	Risk tiers: prohibited practices, high-risk systems, transparency obligations, general-purpose AI models

Three definitions keep the rest of the comparison precise:

- **Management system.** ISO describes an AI management system as "a set of interrelated or interacting elements of an organization intended to establish policies and objectives, as well as processes to achieve those objectives, in relation to the responsible development, provision or use of AI systems" (ISO). It governs the organisation, not a single model.
- **Function, category, subcategory.** In the NIST AI RMF, the Core "is composed of four functions: GOVERN, MAP, MEASURE, and MANAGE", broken into categories and subcategories; NIST adds that the actions "do not constitute a checklist, nor are they necessarily an ordered set of steps" (AI 100-1, section 5).
- **Provider and deployer.** Under the AI Act, the provider develops or places a system on the market under its name; the deployer uses it under its authority. Article 2(1) applies the Act to providers placing systems on the EU market "irrespective of whether those providers are established or located within the Union or in a third country", to deployers in the Union, and to third-country providers and deployers "where the output produced by the AI system is used in the Union" (AI Act, Article 2).

Scope and timing

ISO/IEC 42001 is written for "organizations of any size involved in developing, providing, or using AI-based products or services" (ISO). The NIST AI RMF is addressed to organisations "of all sizes and in all sectors" (AI 100-1). NIST added a Generative AI Profile in July 2024, NIST AI 600-1, which lists risks "unique to or exacerbated by" generative AI and suggested actions mapped to the RMF subcategories.

The AI Act is the only one of the three with legal deadlines. After the Digital Omnibus on AI, which was published in the Official Journal on 24 July 2026 and entered into force on 27 July 2026, Article 113 sets these dates (AI Act, Article 113 as amended; Regulation (EU) 2026/1744):

Provisions	Apply from
Chapters I and II (general provisions, AI literacy, prohibited practices)	2 February 2025 (certain new Article 5 points from 2 December 2026)
General-purpose AI models (Chapter V), governance, penalties	2 August 2025
General date of application	2 August 2026
High-risk requirements, Chapter III Sections 1 to 3, for Annex III systems	2 December 2027
High-risk requirements for Annex I (product) systems	2 August 2028

The practical consequence is that, for most Annex III high-risk use cases, the requirements on risk management, data governance, logging and human oversight are not yet applicable on the date of this article. Programmes building evidence now are building it ahead of the deadline, which is also the cheapest time to build it.

Common ground: five things all three ask for

Read side by side, the three instruments converge on five topics. The table states, for each, the closest provision that was read for this article.

Topic	ISO/IEC 42001	NIST AI RMF	EU AI Act (high-risk)
Inventory of AI systems	Implied by an organisation-wide management system	GOVERN 1.6: "Mechanisms are in place to inventory AI systems"	Registration in the EU database (Article 49, Article 71) for Annex III systems
Risk management	Risk and opportunity management within the MSS	MAP and MEASURE functions; GOVERN 1.5 periodic review	Article 9: a risk management system as "a continuous iterative process" over the lifecycle
Data	Not detailed on the public page	MEASURE 2.1: "Test sets, metrics, and details about the tools used during TEVV are documented"	Article 10: data governance for training, validation and test data, including "data collection processes and the origin of data"
Logging and records	Documented information as part of an MSS	MEASURE 2.4: behaviour "monitored when in production"; MANAGE 4.3: incidents "followed and documented"	Article 12: systems "shall technically allow for the automatic recording of events (logs)"; Articles 19 and 26(6): logs kept at least six months
Human oversight	Roles and responsibilities in the MSS	MAP 3.5: "Processes for human oversight are defined, assessed, and documented"	Article 14: designed to be "effectively overseen by natural persons"

Sources: NIST AI 100-1, section 5; AI Act Article 9, Article 10, Article 12, Article 14, Article 19, Article 26.

The ISO column is deliberately thin. The text of ISO/IEC 42001 is sold by ISO, and this article cites only what ISO publishes openly. Organisations that hold the standard can extend the column with clause references.

Where they diverge

Legal consequence. Only the AI Act carries enforcement. ISO/IEC 42001 and the NIST AI RMF matter legally only where a contract, a regulator or a procurement rule refers to them.

Unit of evidence. ISO/IEC 42001 asks whether the organisation runs a management system. The NIST AI RMF asks whether outcomes in each subcategory are achieved, to the degree an organisation chooses. The AI Act asks, system by system, whether specific requirements are met, and assigns them to named roles: Article 12 to the system's design, Article 19 to providers, Article 26 to deployers.

Specificity about data. The AI Act is the most specific. Article 10(2) lists the data governance practices that must be addressed, from "relevant design choices" through "data-preparation processing operations, such as annotation, labelling, cleaning, updating, enrichment and aggregation" to "the identification of relevant data gaps or shortcomings" (Article 10). NIST asks for documentation of test sets and metrics. NIST's generative AI profile adds a more concrete suggestion for inventories: action GV-1.6-003 lists "Data provenance information (e.g., source, signatures, versioning, watermarks)" among items to consider in inventory entries (NIST AI 600-1).

Standards and presumption of conformity. Under Article 40, systems that conform to harmonised standards whose references are published in the Official Journal "shall be presumed to be in conformity" with the high-risk requirements, "to the extent that those standards cover those requirements" (Article 40). 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 itself is not cited in that way, so conformity with it gives no presumption; a management system certificate is not the same thing as a conformity assessment of a high-risk AI system.

Literacy. The AI Act imposes an AI literacy obligation in Article 4, which the Digital Omnibus amended to state that it "does not require providers or deployers to guarantee any specific level of AI literacy of any individual" (Article 4). NIST covers training under GOVERN 2.2. Neither requirement is phrased the same way.

Designing one evidence register

An evidence register is a structured record of the facts an organisation would need to produce if asked "show me". The design principle is to record each fact once, at the level of the AI system, with its source and date, and to map it outwards to each instrument. The alternative, three parallel compliance workbooks, guarantees that the same fact will be stated three ways and drift.

Register field	Example value (fictional)	ISO/IEC 42001 use	NIST AI RMF	EU AI Act
System identifier and version	`claims-triage` v3.2	Scope of the AIMS	GOVERN 1.6 inventory	Technical documentation, Article 11
Intended purpose	Prioritise incoming health claims for manual review	Objectives	MAP 1.1 intended purposes	Basis for Annex III classification
Role of the organisation	Deployer	Context of the organisation	AI actor roles	Provider vs deployer obligations
Risk classification and rationale	Not Annex III point 5(c): no pricing or risk assessment of persons	Risk assessment	MAP 5.1 likelihood and magnitude	Article 6 classification

Register field	Example value (fictional)	ISO/IEC 42001 use	NIST AI RMF	EU AI Act
Data sources with origin and collection purpose	Claims system extract 2024-2025; origin and legal basis recorded	Documented information	MEASURE 2.1	Article 10(2)(b)
Data preparation steps	De-identification, labelling protocol v2	Documented information	MEASURE 2.1	Article 10(2)(c)
Test set and metrics, dated	Hold-out set, recall 0.94 on 2026-06-30	Performance evaluation	MEASURE 2.1, 2.4	Article 9 testing; Article 15 accuracy
Log specification and retention	Event log fields; retained 13 months	Operational control	MEASURE 2.4	Articles 12, 19, 26(6)
Human oversight assignment	Named role with override authority	Roles and responsibilities	MAP 3.5	Article 14, Article 26(2)
Incidents and changes	Change ticket, incident ID, outcome	Nonconformity and corrective action	MANAGE 4.1, 4.3	Article 72, 73

Each value in the register is itself a fact that may be challenged. A recall figure without the test set version, the date, and who ran the evaluation is a number, not evidence. The same applies to "data source" entries: a field that says "claims system" does not identify which extract, as of which date, with which filters. Recording provenance in the sense of the site's provenance glossary entry, meaning where a value came from and through which steps, is what turns a register into something an auditor can test. The difference between provenance and a pipeline-level data lineage arrow matters here: lineage says which table fed which; provenance says which specific record supports which specific claim.

A worked example

The following organisation is fictional. Halden Therapeutics AG (fictional) is a Swiss company with EU affiliates that uses a third-party model to screen adverse-event reports and an internal model to forecast clinical-supply demand. It wants one register to answer an ISO/IEC 42001 audit, a US partner's request to show alignment with the NIST AI RMF, and an EU market surveillance question.

1. **Inventory.** Two entries are created. The supply model's outputs are used only internally and do not fall under Annex III; it is recorded as not high-risk, with the rationale. The adverse-event model is also outside Annex III, but its intended purpose and the human review step are recorded because both ISO and NIST ask for them.
2. **Data.** For each model, the register holds the training and test data set identifiers, the date range, origin, and the preparation steps. For the vendor model, it records what the vendor disclosed and what it did not. That gap becomes a risk item.
3. **Logging.** Halden specifies which events are logged per prediction: input reference, model version, output, reviewer decision. The same specification answers NIST MEASURE 2.4 and would answer AI Act Article 12 if a future use case were high-risk.

4. **Oversight.** Each model has a named oversight role with the authority to stop it, which answers MAP 3.5 and NIST MANAGE 2.4 ("Mechanisms are in place ... to supersede, disengage, or deactivate AI systems").

The point of the example is that no record is created for a single framework. Each record exists because the organisation needs it to explain a system; the mappings are views.

What none of them tests

All three instruments ask that records exist. None of them, as published, specifies a test that a record is correct. A register can state that training data came from a given source while the source extract was replaced last month. A log can capture a model's output and the documents it retrieved, without capturing which statement in which document supported the output. An inventory can list a vendor model under one name while procurement lists it under another. These are data quality and identity problems, and the instruments leave the method to the organisation.

Three checks fill some of that gap:

- **Re-derivation.** Can a reviewer regenerate a reported metric from the recorded test set version and code version?
- **Source resolution.** Does every data source in the register resolve to a retained, dated artefact rather than a system name?
- **Identity consistency.** Is each AI system, vendor and data set identified by one identifier across the register, contracts and logs?

Sequencing a programme

A defensible order follows the dependencies rather than the frameworks:

1. Build the inventory with intended purpose and role (provider or deployer) for each system. Everything else hangs from it.
2. Classify each system under the AI Act and record the rationale; this determines which systems need the high-risk evidence before 2 December 2027 or 2 August 2028.
3. Define the register fields once, and populate them for high-risk and high-impact systems first.
4. Specify logs and retention per system, in line with Article 12 and the six-month minimum in Articles 19 and 26(6) where applicable.
5. Wrap the register in a management system (policy, roles, internal audit, review), which is what ISO/IEC 42001 is designed to assess.
6. Map the register to NIST subcategories for counterparties who ask for it.

Summary matrix

Question	ISO/IEC 42001	NIST AI RMF	EU AI Act
Is it law?	No	No	Yes
Who is assessed?	The organisation's management system	The organisation's practices, self-selected	Each AI system and each operator's obligations
What is the key evidence?	Documented management system and its operation	Documented outcomes per subcategory	Technical documentation, logs, data governance, conformity assessment
When does it bite?	When adopted or required by contract	When adopted or required by contract	Dates in Article 113; high-risk from 2 December 2027 (Annex III)
What does it say about data provenance?	Not stated on the public page	AI 600-1 suggests provenance information in inventory entries	Article 10(2)(b): origin of data and original purpose of collection

A single evidence register, keyed by system and version and populated with sourced, dated facts, can serve all three. Reliable Context is designed to hold that kind of sourced, dated fact for organisations that need to defend it.

Sources

1. ISO, *ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system* (standard page). https://www.iso.org/standard/81230.html
2. NIST, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, January 2023. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
3. NIST, *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1, July 2024. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
4. Regulation (EU) 2024/1689 (Artificial Intelligence Act), EUR-Lex. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
5. Regulation (EU) 2026/1744 (Digital Omnibus on AI), OJ L, 24 July 2026, EUR-Lex. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
6. European Commission, AI Act Service Desk, consolidated text as at 27 July 2026: Articles 2, 4, 9, 10, 12, 14, 19, 26, 40 and 113. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-12
7. European Commission, "First harmonised standard supporting the AI Act published in the Official Journal" (EN 18286:2026, Implementing Decision (EU) 2026/2274), 19 October 2026. https://digital-strategy.ec.europa.eu/en/library/first-harmonised-standard-supporting-ai-act-published-official-journal

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://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
3. https://eur-lex.europa.eu/eli/reg/2026/1744/oj

4. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-2
5. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
6. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-113
7. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-9
8. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-10
9. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-12
10. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-14
11. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-19
12. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-26
13. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-40
14. https://digital-strategy.ec.europa.eu/en/library/first-harmonised-standard-supporting-ai-act-published-official-journal
15. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-4
16. https://eur-lex.europa.eu/eli/reg/2024/1689/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.