FREE PRIMER · 24 PAGES

■ **ISO/IEC 42001:2023 · Standards Explainer**

The ISO/IEC 42001:2023 Primer

The world's first AI Management System standard — explained in plain English. Structure, clauses, Annex A, and family map.

Structure	Clauses	Annex A	Family map
10 clauses + Annex A	Plain-language 4–10	All 37 controls	27001 · 27701 · EU AI Act

Inside the primer

- ✓ Plain-language explanation of every clause (4–10)
- ✓ Annex A control summary — all 37 controls
- ✓ How ISO 42001 maps to ISO 27001 / 27701
- ✓ Comparison with EU AI Act and NIST AI RMF

PUBLISHED	PUBLISHER	CLAUSES	CONTROLS
December 2023	**ISO & IEC**	**10 main + Annex A**	**37 (Annex A)**

Contents

A 24-page walkthrough of the world's first AI Management System standard.

01 · The plain-English answer

What ISO/IEC 42001:2023 is — and isn't.

Most explanations of ISO 42001 get tangled in standards jargon. Here's the clear answer in two columns.

■ What it IS

ISO/IEC 42001:2023 is a **certifiable management system standard** that specifies requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS) within an organization.

It is designed for entities providing or using products or services that utilise AI systems — to ensure responsible development and use.

If you know ISO 27001 (security) or ISO 9001 (quality), the format is identical: governance structure, policies, controls, audits, continuous improvement.

■■ What it ISN'T

It is **not** a technical AI standard. It does not specify algorithms, architectures, or specific AI techniques to use or avoid.

It is **not** a regulation. ISO 42001 is voluntary, but increasingly required by enterprise buyers, regulators (EU AI Act), and partners as evidence of responsible AI governance.

It is **not** the same as the EU AI Act. The Act is binding law; ISO 42001 is a voluntary management framework that helps organizations meet many EU AI Act obligations.

STANDARD	TYPE	PUBLISHED	ADOPTION
ISO/IEC 42001:2023	**Management standard**	**December 2023**	**Voluntary**

02 · Why this standard now

Why ISO created ISO 42001 in late 2023.

Three converging pressures made an AI Management System standard inevitable. Each pressure on its own would have prompted a framework; together, they forced ISO and IEC to act.

01

The AI governance vacuum

Every organization deploying AI faced the same question: 'How do we govern this responsibly?' Without a standard, every company invented their own framework — incompatible, unauditable, and impossible to validate externally. Boards, regulators, and customers had no common reference point. ISO 42001 fills that gap with a vendor-neutral, internationally agreed framework that any organization can adopt.

02

Regulatory acceleration

The EU AI Act (effective 2024–2026), the US Executive Order on AI, Chinese AI regulations, and sector-specific rules in finance and healthcare created an urgent need for a recognized governance framework — one that auditors and regulators could converge on. ISO 42001 acts as a meeting point: a structure regulators can recognise and assessors can audit against.

03

Enterprise buyer demand

Fortune 500 procurement teams started demanding 'AI governance evidence' from vendors. Without a standard like ISO 27001 to point to, AI vendors had no clean way to demonstrate responsible practices. Selling AI into regulated industries became a slow, custom-questionnaire exercise. ISO 42001 fills that gap and turns AI governance into a procurement-ready credential.

03 · Standard Structure

The 10 clauses of ISO/IEC 42001:2023.

The standard follows the High-Level Structure (HLS) shared by ISO 27001, 9001, and 14001 — making integrated management systems easier to design and audit.

1–3	**Scope, References, Definitions** What the standard covers, normative references, and core vocabulary.
4	**Context of the organization** Internal/external issues, interested parties, AIMS scope.
5	**Leadership** Top management commitment, AI policy, roles and responsibilities.
6	**Planning** AI risk and opportunity assessment, AI impact assessment, objectives.
7	**Support** Resources, competence, awareness, communication, documented information.
8	**Operation** Operational planning, AI system lifecycle, third-party AI relationships.
9	**Performance evaluation** Monitoring, measurement, internal audit, management review.
10	**Improvement** Nonconformity, corrective action, continual improvement.

+ Annex A — 37 Reference Controls
Implementation guidance organized into 9 control categories — the operational heart of the standard. See pages 13–16.

04 · Clause 4

Context of the organization.

Clause 4 asks the organization to define *why* and *where* its AI Management System exists before anything else is built. Get this wrong and every subsequent clause inherits the error.

What you must determine

- **External and internal issues** relevant to the organization's purpose and that affect its ability to achieve the intended outcome(s) of the AIMS.
- **Interested parties** — regulators, customers, employees, affected end-users, data subjects, suppliers — and their requirements relevant to AI.
- **The scope of the AIMS** — which AI systems, which business units, which jurisdictions are inside the boundary, and which are outside.
- **The AI Management System itself** — establish, implement, maintain, and continually improve it, including the processes needed and their interactions.

Practical implementation tip

Treat the scoping document as a living contract. Start narrow — one product line or one business unit — and prove the model works. Expanding scope later is easier than retro-fitting governance onto an over-promised certificate boundary.

Common pitfall. Organizations often scope "all AI everywhere," then discover undocumented shadow-AI tools during the certification audit. A defensible scope statement names systems, owners, and exclusions explicitly.

05 · Clause 5

Leadership.

Clause 5 forces top management to own AI governance. Without a signed AI policy, explicit accountability, and visible commitment, no AIMS is auditable.

5.1 Leadership and commitment

Top management must demonstrate leadership by ensuring the AI policy and AI objectives are compatible with strategic direction, integrating AIMS requirements into business processes, providing the resources needed, and communicating the importance of effective AI management.

5.2 AI policy

Establish a policy that is appropriate to the purpose of the organization, provides a framework for setting AI objectives, includes a commitment to satisfy applicable requirements, and includes a commitment to continual improvement of the AIMS. The policy must be documented, communicated, and available to interested parties as appropriate.

5.3 Roles, responsibilities & authorities

Top management must assign and communicate responsibility and authority for ensuring the AIMS conforms to ISO/IEC 42001 and for reporting on AIMS performance. In practice, mature programmes formalize an AI ethics or AI governance committee, a designated AIMS owner (often the CISO or Chief AI Officer), and clear escalation paths for AI incidents.

What auditors look for

- Signed and dated AI policy, visibly endorsed by the executive team.
- Documented AI governance structure with named individuals — not just job titles.
- Evidence of leadership review (board minutes, steering-committee notes).
- Allocation of budget, headcount, and tooling to the AIMS.

06 · Clause 6

Planning.

Clause 6 turns intent into measurable action. It is the single clause most implementers under-invest in — and the single clause auditors probe hardest.

6.1.1 Actions to address risks and opportunities

Determine the risks and opportunities that need to be addressed to give assurance the AIMS can achieve its intended outcomes, prevent or reduce undesired effects, and achieve continual improvement.

6.1.2 AI risk assessment

Define and apply an AI risk assessment process that establishes and maintains AI risk criteria, ensures repeatable and comparable results, identifies risks associated with AI, analyses them, and evaluates them against the criteria.

6.1.3 AI risk treatment

Define an AI risk treatment process to select appropriate options, determine necessary controls, compare them with Annex A, produce a Statement of Applicability (SoA), and formulate a treatment plan.

6.1.4 AI system impact assessment

Assess potential impacts of AI systems on individuals, groups, and society — covering fairness, safety, privacy, transparency, and accountability.

6.2 AI objectives and planning to achieve them

Establish measurable AI objectives at relevant functions and levels. Plan what will be done, what resources are required, who is responsible, when it will be completed, and how results will be evaluated.

6.3 Planning of changes

When changes to the AIMS are needed, carry them out in a planned manner — considering purpose, consequences, integrity, resources, and responsibilities.

07 · Clause 7

Support.

Clause 7 is about the resources, people, and documentation that make the AIMS actually function in practice — beyond paper policy.

7.1 Resources
Determine and provide the resources needed for the establishment, implementation, maintenance, and continual improvement of the AIMS — covering people, infrastructure, computing power, data, and tooling.

7.2 Competence
Determine the necessary competence of persons doing AI-related work, ensure they are competent on the basis of education, training, or experience, and retain documented information as evidence of competence.

7.3 Awareness
Persons doing work under the organization's control must be aware of the AI policy, their contribution to the AIMS, the implications of not conforming, and the broader context of responsible AI use.

7.4 Communication
Determine the internal and external communications relevant to the AIMS — what, when, with whom, how, and who communicates. Critical for incident response and interested-party engagement.

7.5 Documented information
Required documented information must be controlled — distributed, accessed, retrieved, used, stored, preserved (including legibility), controlled through changes, and properly retained and disposed of.

08 · Clause 8

Operation.

Clause 8 is where the AIMS meets day-to-day reality. It applies operational control to the design, development, deployment, and retirement of every AI system in scope.

8.1 Operational planning and control
Plan, implement, and control the processes needed to meet AIMS requirements. Establish criteria for the processes, implement control in accordance with those criteria, and keep documented information sufficient to have confidence the processes have been carried out as planned.

8.2 AI risk assessment — operational
Perform AI risk assessments at planned intervals or when significant changes are proposed. Retain documented information of the results.

8.3 AI risk treatment — operational
Implement the AI risk treatment plan. Retain documented information of the results of the AI risk treatment.

8.4 AI system impact assessment — operational
Perform AI system impact assessments at planned intervals and when significant changes occur. Retain documented information of the assessment results, including the methodology used.

Operation is the clause that consumes the most effort post-certification. Build sustainable workflows here, not heroics.

09 · Clause 9

Performance evaluation.

Clause 9 is how you prove the AIMS works. Three mechanisms — monitoring, internal audit, and management review — feed evidence to the certification body and to your board.

9.1 Monitoring, measurement, analysis and evaluation

Determine what needs to be monitored and measured, the methods used to ensure valid results, when monitoring shall be performed, and when results shall be analysed and evaluated. Retain documented information as evidence.

9.2 Internal audit

Conduct internal audits at planned intervals to provide information on whether the AIMS conforms to the organization's own requirements and to the requirements of ISO/IEC 42001, and is effectively implemented and maintained. Plan, establish, implement, and maintain an audit programme.

9.3 Management review

Top management must review the AIMS at planned intervals to ensure its continuing suitability, adequacy, and effectiveness — covering the status of actions from previous reviews, changes in external/internal issues, feedback on AIMS performance, results of risk assessment and treatment, and opportunities for continual improvement.

10 · Clause 10

Improvement.

Clause 10 closes the Plan-Do-Check-Act loop. It is what separates a ceremonial AIMS from one that actually reduces AI risk year over year.

10.1 Continual improvement
The organization must continually improve the suitability, adequacy, and effectiveness of the AI Management System. This is not optional or aspirational — auditors expect evidence.

10.2 Nonconformity and corrective action
When a nonconformity occurs, react to it, evaluate the need for action to eliminate its cause(s), implement any action needed, review the effectiveness of corrective action taken, and make changes to the AIMS if necessary. Retain documented information as evidence.

The PDCA loop in practice

Plan the AIMS (Clauses 4–6) · **Do** operate it (Clauses 7–8) · **Check** performance (Clause 9) · **Act** to improve (Clause 10). Repeat annually, with management review as the formal pivot point.

11 · Annex A · Operational Controls

Annex A — the operational heart of the standard.

Annex A organizes 37 reference controls into 9 categories. Implementers spend most of their time mapping these to existing controls (often inherited from ISO 27001) and building new ones where gaps appear.

A.2
Policies related to AI
Approval, communication, review of AI policies.

A.3
Internal organization
Roles, responsibilities, AI ethics committee structure.

A.4
Resources for AI systems
Computing, data, tooling, system documentation.

A.5
AI system impact assessment
Identify, document, mitigate stakeholder impacts.

A.6
AI system lifecycle
Design, development, validation, deployment, retirement.

A.7
Data for AI systems
Sources, quality, lineage, privacy, retention.

A.8
Information for interested parties
Communication to users, regulators, affected parties.

A.9
Use of AI systems
Intended use, monitoring, deviation handling.

A.10
Third-party relationships
Vendor AI controls, embedded LLMs, supply chain.

12 · Annex A — first half

A.2 to A.5 in detail.

The first four control categories cover governance fundamentals — policy, people, resources, and impact assessment.

A.2 Policies related to AI

Establish and maintain AI policies addressing acceptable use, ethics, transparency, human oversight, and review cadence. Sub-controls cover approval, communication, and periodic review of those policies.

A.3 Internal organization

Define AI-related roles and responsibilities. Establish accountability for AI risk owners, AI system owners, and an AI ethics or governance committee where appropriate. Ensure segregation of duties for AI development, testing, and deployment.

A.4 Resources for AI systems

Identify and document the resources needed across the AI lifecycle — including compute infrastructure, training data, tooling, and human expertise. Ensure each AI system has documented inventory, owners, and metadata sufficient for audit.

A.5 AI system impact assessment

Establish a process to identify, analyse, document, and mitigate the impacts of AI systems on individuals, groups, and society — including fairness, safety, privacy, and human autonomy. Update assessments when AI systems materially change.

13 · Annex A — second half

A.6 to A.10 in detail.

The second set of categories cover the operational mechanics — lifecycle, data, communications, end-use, and third-party relationships.

A.6 AI system lifecycle

Govern each phase: design, development, validation, deployment, operation, and retirement. Include controls for objective-setting, requirements specification, testing, deployment approval, monitoring, and decommissioning.

A.7 Data for AI systems

Manage data quality, lineage, provenance, retention, and privacy across the lifecycle. Document training, validation, and test datasets; address bias; ensure data minimization and lawful basis of processing.

A.8 Information for interested parties

Provide appropriate information to users, regulators, affected end-users, and the public — including system purpose, capabilities, limitations, contact points for queries, and channels for raising concerns.

A.9 Use of AI systems

Define intended use boundaries, monitor for deviation, and govern human-AI interaction. Establish processes for handling out-of-distribution inputs, automation bias, and override authority.

A.10 Third-party & customer relationships

Govern AI obtained from suppliers (LLM vendors, model marketplaces, embedded AI features) and AI provided to customers. Include contractual controls, shared-responsibility models, and supplier review cadence.

14 · Quick reference

The 37 controls at a glance.

A flat index of every Annex A control — useful for gap-analysis spreadsheets and Statement-of-Applicability templates.

A.2.2	Documented AI policy	A.6.2.5	Operation & monitoring
A.2.3	Alignment with other policies	A.6.2.6	AI system technical documentation
A.2.4	Review of AI policy	A.6.2.7	Event logging
A.3.2	AI roles & responsibilities	A.6.2.8	AI system disposal
A.3.3	Reporting of concerns	A.7.2	Data management process
A.4.2	Resource documentation	A.7.3	Acquisition of data
A.4.3	Data resources	A.7.4	Quality of data
A.4.4	Tooling resources	A.7.5	Provenance of data
A.4.5	System & compute resources	A.7.6	Data preparation
A.4.6	Human resources	A.8.2	System info for users
A.5.2	AI impact assessment process	A.8.3	External reporting
A.5.3	Documentation of impacts	A.8.4	Communication of incidents
A.5.4	Assessing impact to individuals	A.8.5	Information to interested parties
A.5.5	Assessing societal impact	A.9.2	Processes for responsible use
A.6.1.2	Objectives for the AI system	A.9.3	Objectives for responsible use
A.6.1.3	Documentation of AI system	A.9.4	Intended use of AI system
A.6.2.2	AI system requirements	A.10.2	Allocation of responsibilities
A.6.2.3	Verification & validation	A.10.3	Suppliers
A.6.2.4	Deployment of AI system	A.10.4	Customers

Control identifiers based on ISO/IEC 42001:2023 Annex A. Use as an SoA starter — refer to the official standard for normative wording.

15 · ISO 42001 in the family

How ISO 42001 fits the bigger picture.

ISO 42001 doesn't replace existing management systems — it complements them. Most enterprises layer it on top of the ISO 27001 programme they already operate.

AI MANAGEMENT **ISO/IEC 42001**	Governance, risk, and lifecycle management for AI systems. The newest member of the family.	ADOPTION **2023**
INFORMATION SECURITY **ISO/IEC 27001**	Information Security Management System. The most widely adopted ISO standard globally.	ADOPTION **71K+ certs**
PRIVACY **ISO/IEC 27701**	Privacy Information Management. Extends ISO 27001 for privacy/PIMS controls.	ADOPTION **Growing**
QUALITY **ISO 9001**	Quality Management System. The grandfather of all management-system standards.	ADOPTION **1M+ certs**

16 · Family mapping

Mapping ISO 42001 to 27001 / 27701 / 9001.

Because all four standards share the High-Level Structure, controls and clauses align cleanly. Where you already have controls under 27001 or 27701, you can usually inherit a large portion of the AIMS evidence base.

ISO 42001	ISO 27001	ISO 27701	ISO 9001
Clause 4 — Context	Clause 4 — Context	Clause 4 — Context	Clause 4 — Context
Clause 5 — Leadership	Clause 5 — Leadership	Clause 5 — Leadership	Clause 5 — Leadership
Clause 6 — Planning + AI risk	Clause 6 — Planning + InfoSec risk	Clause 6 — Privacy risk	Clause 6 — Quality planning
A.6 — AI system lifecycle	A.8 / A.14 — Asset, SDLC	Privacy by design controls	Clause 8 — Operation
A.7 — Data for AI systems	A.8 / A.5 — Asset & info classification	PII processing controls	Clause 7.5 — Documented info
A.10 — Third-party relationships	A.15 / A.5.19 — Supplier security	Processor / sub-processor controls	Clause 8.4 — External providers
Clause 9 — Performance eval	Clause 9 — Performance eval	Clause 9 — Performance eval	Clause 9 — Performance eval
Clause 10 — Improvement	Clause 10 — Improvement	Clause 10 — Improvement	Clause 10 — Improvement

Organizations with mature ISO 27001 programmes typically achieve 40–60% time savings when implementing ISO 42001, by extending existing GRC infrastructure rather than rebuilding it.

17 · Regulatory comparison

ISO 42001 vs the EU AI Act.

They overlap heavily — but they are not interchangeable. Most regulated organizations need both: the Act for legal compliance, ISO 42001 for management-system rigour and external validation.

Dimension	ISO/IEC 42001:2023	EU AI Act
Type	International voluntary standard	Binding regulation in EU law
Scope	Any organization providing or using AI	AI systems placed on the EU market or affecting EU persons
Approach	Management-system framework — process, governance, lifecycle	Risk-tier rules: unacceptable / high / limited / minimal risk
Enforcement	Third-party certification (voluntary)	Fines up to 7% of global annual turnover
Timeline	Published Dec 2023; certifications now available	Phased entry into force 2024–2026
Geography	Global — recognized in 100+ countries	European Union (with extraterritorial reach)
Best used for	Demonstrating responsible-AI governance to customers, regulators, boards	Achieving legal compliance for products sold in the EU

18 · Framework comparison

ISO 42001 vs NIST AI RMF.

Both promote responsible AI. They differ in nature: ISO 42001 is a certifiable management system; NIST AI RMF is a voluntary guidance framework. Most US-headquartered enterprises adopt both — RMF as the internal North Star, 42001 as the external certificate.

Dimension	ISO/IEC 42001:2023	NIST AI RMF 1.0
Issuer	ISO & IEC (international)	US NIST (national)
Type	Certifiable management-system standard	Voluntary guidance framework (not certifiable as-is)
Core structure	10 clauses + Annex A controls	Four functions: Govern, Map, Measure, Manage
Audit	Independent third-party audit possible	Self-assessment based; no formal certificate
Mapping	Aligns with EU AI Act, ISO 27001 family	Aligns with NIST CSF and Privacy Framework
Best used for	Procurement evidence, regulator readiness, global recognition	Internal risk-management discipline, especially in US public sector

*Practical rule of thumb: use NIST AI RMF to **build** your governance programme, then implement ISO/IEC 42001 to **certify** it.*

19 · Adoption signals

Who is implementing ISO 42001 first.

Adoption is concentrated in four categories — all dealing with high-stakes AI deployments where governance evidence is now part of doing business.

FIN

Regulated financial services
Banks, insurers, and fintechs deploying AI for credit, fraud, underwriting, and AML. ISO 42001 provides governance evidence for regulators (EU AI Act, OCC, FCA, MAS, RBI).

MED

Healthcare & pharmaceuticals
AI-assisted diagnosis, drug discovery, claims processing. Aligns with FDA AI/ML guidelines and EU MDR for AI medical devices, and supports ISO 13485 quality programmes.

GOV

Government & public sector
AI in benefits decisions, predictive policing, citizen services, and procurement. Public-sector buyers increasingly require independent governance evidence from suppliers.

ENT

Enterprise SaaS & AI vendors
B2B AI products selling into Fortune 500. ISO 42001 alignment is becoming a procurement question — much like ISO 27001 became a decade ago.

20 · Implementation roadmap

A realistic 9–14 month roadmap.

For a mid-sized organization (500–5,000 employees, moderate AI footprint), this is the typical path from kick-off to certificate. Organizations with mature ISO 27001 programmes can move 40–60% faster.

Phase 1 · Months 1–2 **Mobilize**	Executive sponsorship secured. AIMS scope defined. AI policy drafted. Steering committee stood up. Existing ISO 27001 / 27701 controls inventoried for reuse.
Phase 2 · Months 3–5 **Design**	Risk-assessment methodology adopted. AI system inventory completed. Impact-assessment process documented. Statement of Applicability drafted against Annex A's 37 controls.
Phase 3 · Months 6–9 **Implement**	Controls deployed across in-scope AI systems. Lifecycle gates introduced. Data-governance uplifts complete. Training rolled out to AI builders, reviewers, and operators.
Phase 4 · Months 10–12 **Operate & internal audit**	AIMS runs in production for at least one operating cycle. Internal audit completes. Management review held. Corrective actions tracked to closure.
Phase 5 · Months 13–14 **Certify**	Stage 1 audit (documentation review) with accredited certification body. Stage 2 audit (operational evidence). Certificate issued. Surveillance audit cycle begins.

21 · FAQ

ISO 42001 explained — common questions.

Q. Is ISO 42001 mandatory?
No, it is voluntary. But it is increasingly required by enterprise procurement teams, fits well with EU AI Act compliance, and is becoming the de facto governance evidence for organizations selling AI-enabled products.

Q. How is ISO 42001 different from the EU AI Act?
The EU AI Act is binding regulation. ISO 42001 is a voluntary standard. They overlap heavily — implementing ISO 42001 helps satisfy many EU AI Act obligations — but they are not the same. Some organizations need both: the Act for legal compliance, ISO 42001 for management-system rigour and external validation.

Q. How long does implementation take?
For a mid-sized organization (500–5,000 employees, moderate AI footprint): typically 9–14 months from kick-off to certificate. Organizations with mature ISO 27001 programmes can move 40–60% faster by extending existing GRC infrastructure.

Q. Where do I get the official standard document?
Published by ISO and sold through national standards bodies (BSI, ANSI, DIN, BIS, etc.). This primer covers structure, key clauses, and Annex A — useful before you decide whether to buy the full normative document.

Q. Who should learn the standard?
Auditors (internal and external), GRC professionals, AI/ML practitioners moving into governance, CISOs and security leaders extending into AI risk, risk managers, compliance officers, and consultants serving regulated industries.

22 · Where to from here

From understanding to credential.

If reading this primer is the first step, auditing the standard is where careers compound. GSDC's globally-recognized ISO 42001 Lead Auditor programme turns this knowledge into a portable credential.

100+	16+	30	∞
Countries recognize the credential	Hours of expert-led learning	Learn-by-Doing audit projects	Lifetime access, SME mentorship

What you'll get inside the certification

- ✓ 30 hands-on Learn-by-Doing audits
- ✓ SME-reviewed enterprise audit capstone
- ✓ Daily live sessions + 1-on-1 mentorship
- ✓ Lifetime access to materials & updates

In one sentence

ISO/IEC 42001:2023 is a certifiable international standard that gives organizations a structured framework to govern AI systems responsibly — covering policies, risk, lifecycle controls, and continual improvement. Think "ISO 27001 for AI."

This primer is informational. For the normative wording of ISO/IEC 42001:2023, refer to the official ISO publication. All CTAs in this PDF link to gsdcouncil.org/what-is-iso-42001.