RSAA

RS ASSURANCE & ADVISORY

ISO/IEC 42001 Readiness

What certification to the AI management system standard actually requires — and how to be ready for it

WHITE PAPER · OCTOBER 2026

CONTENTS

What this paper covers

1	Why AI governance became a procurement question	3
2	What ISO/IEC 42001 is, and what it is not	3
3	What certification actually requires	4
4	The five decisions that determine readiness	6
5	Governing AI agents under ISO 42001	8
6	If you are already certified to ISO 27001	9
7	A readiness sequence in dependency order	9
8	How RS Assurance & Advisory supports readiness	10
A	Mandatory requirements: clauses 4–10	11
B	Annex A: the 38 controls	13
C	Agent controls mapped to ISO 42001	15

EXECUTIVE SUMMARY

The short version

ISO/IEC 42001:2023 is the first certifiable management system standard for artificial intelligence. Certification turns on two things: meeting every requirement in clauses 4 through 10, and operating each Annex A control that your Statement of Applicability says applies. Clauses are never optional. Annex A controls can be excluded only with a written justification.

The hard part of 42001 is rarely technical. Four of its requirements carry most of the risk, and none of them concern models: a complete inventory of where AI is actually used, an impact assessment that is genuinely an impact assessment rather than a risk assessment, exclusions in the Statement of Applicability that are justified rather than asserted, and hosted-model providers run through vendor due diligence.

The fifth obstacle is time, and it is structural rather than procedural. A certification auditor samples records to confirm the management system operates. A system stood up last month has no records to sample. The pacing item is

therefore not writing the documentation — it is accumulating enough operating evidence, then completing an internal audit and a management review, before Stage 2 begins.

This paper sets out what the standard requires, which of its requirements are easiest to under-build, and a sequence that gets an organization to Stage 1 without rework. The complete clause and control checklists are included as appendices.

SECTION 1

Why AI governance became a procurement question

For most of the last decade, the question a buyer asked a vendor about AI was whether they had any. That has changed. AI capability now arrives inside software that organizations already own — summarization in a ticketing system, drafting inside a CRM, transcription in a meeting platform, code assistance in a developer toolchain. It is rarely procured as "an AI project," which is precisely why it is so often ungoverned.

Three pressures have converged. Buyers have added AI questions to security questionnaires and vendor reviews, and they increasingly want an answer backed by something other than a policy PDF. Regulators in multiple jurisdictions have moved from guidance toward enforceable obligations for AI systems that affect people. And boards, having watched the security function mature under ISO 27001 and SOC 2, are asking the obvious follow-up question about AI.

ISO/IEC 42001:2023 is the response to that. It gives an organization a way to demonstrate, through independent certification, that it manages AI deliberately rather than incidentally.

> **A POINT WORTH BEING PRECISE ABOUT**
>
> ISO 42001 certification is not a legal conformity assessment. In particular, it does not by itself create a presumption of conformity with the EU AI Act; the harmonized standards for that regulation are a separate workstream. What certification does demonstrate is governance maturity — which is what most customer questionnaires, contractual commitments, and board oversight requirements are actually asking about. Confirm the current regulatory position with counsel before making claims in contracts or marketing.

SECTION 2

What ISO/IEC 42001 is, and what it is not

ISO/IEC 42001 specifies requirements for an **AI management system**: the set of policies, roles, processes, and records through which an organization directs and controls its use and provision of AI. It follows the same Harmonized Structure as ISO 27001 and ISO 9001, so the clause architecture will look familiar to anyone who has been through either.

© 2026 RS Assurance & Advisory | rsassure.com

The standard applies to any organization that develops, provides, or uses AI systems, and clause 4.1 requires you to determine which of those roles you occupy — often several at once. A firm that builds a model, resells a model-backed product, and uses a third party's assistant internally has three roles and three different sets of obligations.

What it governs

42001 governs the *management* of AI. It asks whether you know where AI is used, whether someone owns it, whether you assessed what it could do to people, whether you decided which controls apply and why, and whether you can show all of that operating over time.

What it does not do

It does not certify that a model is accurate, safe, unbiased, or fit for a given purpose. It does not substitute for sector regulation. A certificate says the management system meets the standard — not that every output of every system is sound. Organizations that market it as the latter create a credibility problem for themselves at the first bad output.

> **SCOPE IS A CHOICE, AND IT IS CONSEQUENTIAL**
>
> The scope statement names the AI systems, products, and locations covered. A narrow scope is faster to certify and easier to defend, but a certificate that excludes the AI system your customers actually care about will not survive their first careful read of it. Set scope against what the market is asking you to prove, not against what is convenient to prove.

SECTION 3

What certification actually requires

Two layers of obligation

The first layer is clauses 4 through 10. These are the management system requirements — context, leadership, planning, support, operation, performance evaluation, and improvement. Every one is mandatory. A missing clause requirement is a major nonconformity, and a major nonconformity blocks certification until it is corrected.

The second layer is Annex A, which contains 38 controls organized into nine areas covering policy, internal organization, resources, impact assessment, the AI system life cycle, data, information for interested parties, responsible use, and third-party relationships. You decide which apply, and you record that decision — with justification for both inclusions and exclusions — in a **Statement of Applicability**. Annex B of the standard provides implementation guidance for each control.

© 2026 RS Assurance & Advisory | rsassure.com

The audit itself

Stage 1 — documentation review

The auditor examines your scope, AI policy, risk and impact assessment methods, and Statement of Applicability, then decides whether you are ready to proceed. Findings here are a gift: they are cheap to fix compared with findings at Stage 2.

Stage 2 — implementation audit

The auditor samples records, interviews control owners, and confirms the system actually operates as documented. This is where a management system that exists only on paper comes apart.

A certificate runs for three years, with surveillance audits in years one and two and a recertification audit in year three. Minor nonconformities require a corrective action plan; major ones must be corrected before a certificate is issued. Use a certification body accredited for ISO 42001 — ISO/IEC 42006 sets the requirements for bodies that audit AI management systems, and an unaccredited certificate will not carry weight with a sophisticated buyer.

© 2026 RS Assurance & Advisory | rsassure.com

SECTION 4

The five decisions that determine readiness

A small set of decisions, taken early, determines how much of this work has to be redone later. None of them are technical — each follows directly from how the standard is written.

1. Whether your AI inventory is actually complete

This is the requirement with the widest surface area, and the difficulty is structural: AI arrives through channels that do not look like procurement. A team enables an AI feature in a tool the company already licenses. Someone installs a browser extension. A meeting note-taker joins calls because a single employee invited it. None of these pass through a vendor review, and none of them appear on an inventory built by asking engineering what models they run.

Build the inventory from multiple directions — expense and SaaS spend, identity provider sign-in logs, network egress to model-provider domains, and a direct survey of teams — and treat embedded AI features in existing applications as in-scope systems, not as features of something already covered.

2. Whether you understand that an impact assessment is not a risk assessment

Clause 6.1.4 requires an AI system impact assessment, and it is the requirement most easily satisfied with the wrong document. A risk assessment asks what could happen to the organization. An impact assessment asks what could happen to *individuals, groups, and society* — fairness, safety, rights, and broader effects. An auditor who reads an impact assessment and finds only business risk will record a finding, because the clause has not been met. ISO/IEC 42005 provides guidance specific to this assessment and is worth obtaining alongside the standard.

3. Whether your exclusions are justified or merely asserted

Exclusions are permitted. Unexplained exclusions are not. If the organization only uses AI and does not build it, some A.6 life cycle controls will be narrower in scope — but the auditor still expects the reasoning written down, tied to the documented role from clause 4.1. "Not applicable" is not a justification. "We do not develop or fine-tune models; our role under clause 4.1 is AI user; development controls are therefore limited to configuration of third-party systems" is.

4. Whether hosted-model providers went through supplier due diligence

Control A.10.3 covers suppliers, and hosted-model providers are suppliers. They are also the suppliers most likely to have been onboarded by a product team without a vendor review. Due diligence should establish, in writing, whether your data is used for training, what the retention and deletion terms are, who the subprocessors are, how model versions are deprecated and with what notice, and what the provider commits to on availability and incident notification.

5. Whether the system has operated long enough to be auditable

Clause 9.1 requires monitoring and measurement, 9.2 requires an internal audit by a competent and independent auditor, and 9.3 requires a management review covering every required input. Each of these produces records, and

© 2026 RS Assurance & Advisory | rsassure.com

Stage 2 samples records. An organization that writes excellent documentation and books Stage 2 for the following month has nothing to sample. Plan for a period of genuine operation before the audit, not a documentation sprint.

> **TWO FURTHER REQUIREMENTS THAT ARE EASY TO UNDER-BUILD**
>
> **Objectives with no measure or owner.** Clause 6.2 requires measurable AI objectives. "Use AI responsibly" is not measurable. Attach a metric, an owner, and a date to each.
>
> **No internal audit or management review before Stage 2.** These are clause requirements in their own right, and they are also the mechanism by which you find your own nonconformities before the auditor does. Skipping them converts cheap findings into expensive ones.

SECTION 5

Governing AI agents under ISO 42001

An AI agent is any system that takes actions on its own rather than only producing text for a person to review — sending email, moving money, changing records, opening tickets, modifying infrastructure. Agents concentrate AI risk because they remove the human checkpoint that made a wrong output harmless.

Five controls carry most of the weight. They are not exotic; they are the controls an organization already applies to any privileged automated process, applied to a new class of actor.

1. **An agent inventory with a named owner.** Every agent, its purpose, what it can access, and the person accountable for it. An agent nobody owns is an agent nobody will retire.
2. **Its own least-privilege identity.** Not a shared service account, and never a person's credentials. A distinct identity is what makes the logs meaningful and the access review possible.
3. **Human approval for high-impact actions.** Payments, deletions, external messages, production changes. Define the threshold deliberately and keep the approval records.
4. **Complete, tamper-resistant logs.** What the agent did, when, under whose authority, with retention long enough to investigate and to satisfy an audit period.
5. **A kill switch you have actually tested.** A documented stop procedure, with a last-tested date and result. An untested kill switch is a design intention, not a control.

Under ISO 42001 these land mostly in A.6 and A.9, with the approval requirement also flowing through risk treatment at clause 6.1.3. The standard has no access-control section of its own, so identity controls are drawn from ISO 27001.

> ON THE ISO 27001 CROSS-REFERENCE
>
> If you are certified to ISO 27001, those access controls are already in your scope and the mapping is direct. If you are pursuing ISO 42001 alone, they are not automatically in scope — bring them in explicitly through clause 6.1.3 risk treatment, and say so in the Statement of Applicability, so the auditor can see where the identity requirement is being satisfied.

Appendix C maps each of the five controls to the specific ISO 42001 clauses and Annex A controls, with the evidence an auditor will ask for.

© 2026 RS Assurance & Advisory | rsassure.com

SECTION 6

If you are already certified to ISO 27001

A substantial part of the work is already done, and it is worth being specific about which part.

Carries over from ISO 27001	Genuinely new for ISO 42001
Clause structure 4–10 and the management system mechanics	AI system impact assessment (clause 6.1.4) — effects on people and society
Document control, versioning, approvals, retention (7.5)	AI life cycle controls (A.6): responsible development, verification and validation, deployment, monitoring for drift and misuse
Internal audit, management review, corrective action (9.2, 9.3, 10.2)	Data controls specific to AI (A.7): provenance, quality criteria, preparation and labeling, rights to use
Access control, identity management, privileged access	Responsible-use objectives, intended use, and misuse monitoring (A.9)
Supplier management framework (extended to model providers)	AI-specific roles and the clause 4.1 role determination
Risk assessment methodology (extended, not replaced)	Information for interested parties and external adverse-effect reporting (A.8)

The management system carries; the AI-specific substance does not.

Because the clause structure is shared, the two systems can be operated as one and audited together. Where a certification body offers a combined audit, the efficiency is real — but confirm that the body is accredited for ISO 42001 specifically, not only for 27001.

SECTION 7

A readiness sequence in dependency order

The order matters more than the pace. Each step produces an input the next step depends on, and skipping forward almost always produces rework.

1. **Inventory your AI systems and set the scope.** Everything downstream — risk, impact, applicability, evidence — is bounded by this. Get it complete before anything else.
2. **Publish the AI policy and assign roles.** Named owner for the management system, role assignments for every AI system, and a channel for employees to report concerns.
3. **Run the risk assessment and the impact assessment.** Two separate exercises with two separate questions. Do not merge them.

4. **Write the Statement of Applicability.** Every inclusion and every exclusion justified against the documented role and the assessment results.
5. **Implement controls and build operating records.** This is where elapsed time accumulates, and it cannot be compressed by writing faster.
6. **Complete an internal audit and a management review.** Independent, competent, documented. Fix what they find.
7. **Book Stage 1, correct the findings, then Stage 2.** Treat Stage 1 findings as the last cheap feedback you will get.

> THE PLANNING MISTAKE TO AVOID
>
> It is tempting to plan backward from a customer deadline and allocate the time to documentation, because documentation is the visible work. The binding constraint is steps 5 and 6: records that span a period of real operation, and an internal audit and management review conducted on those records. Build the schedule around that constraint and the documentation will fit inside it.

SECTION 8

How RS Assurance & Advisory supports readiness

RSAA supports the readiness path, from first scoping conversation through the point at which an accredited certification body takes over.

Scoping and gap assessment

A complete AI inventory, a defensible scope statement, and a gap assessment against clauses 4–10 and all 38 Annex A controls.

Statement of Applicability

Inclusions and exclusions justified against your documented role, so the exclusions survive scrutiny at Stage 1.

Risk and impact assessment design

Two distinct methodologies built to meet 6.1.2 and 6.1.4 on their own terms, with the documentation an auditor expects.

Internal audit and management review

Clause 9.2 and 9.3 execution by an independent, competent auditor — finding your nonconformities before the certification body does.

Our unified methodology maps controls across SOC 1, SOC 2, HIPAA, ISO 27001, ISO 42001, and CMMC, so evidence gathered once can serve multiple examinations. For organizations pursuing ISO 42001 alongside an existing SOC 2 or ISO 27001 program, that overlap is where most of the cost comes out.

© 2026 RS Assurance & Advisory | rsassure.com

APPENDIX A

Mandatory requirements: clauses 4–10

Every row below is required. Missing any one is a major nonconformity. The right-hand column is what an auditor will ask to see.

Clause	Requirement	What to show the auditor
4.1	Understand your context, including your role (AI provider, developer, user, or several)	Context analysis; documented AI role(s)
4.2	Needs of interested parties	Register of customers, regulators, and others, with their requirements
4.3	Scope of the AI management system	Scope statement naming the AI systems, products, and locations covered
4.4	Establish and maintain the management system	The system itself, operating
5.1	Leadership and commitment	Management review minutes; resourcing decisions
5.2	AI policy	Approved, published, communicated AI policy
5.3	Roles, responsibilities, authorities	Named owner of the management system; RACI for AI roles
6.1.2	AI risk assessment	Documented method and results
6.1.3	AI risk treatment	Risk treatment plan; Statement of Applicability justifying every Annex A inclusion and exclusion
6.1.4	AI system impact assessment	Documented process for effects on individuals, groups, and society
6.2	AI objectives	Measurable objectives with owners and dates
6.3	Planning of changes	Records of planned changes to the system
7.1–7.3	Resources, competence, awareness	Training records; policy acknowledgments
7.4	Communication	Who communicates what, to whom, and when
7.5	Documented information	Document control: versions, approvals, retention
8.1	Operational planning and control	Procedures, including control of outsourced AI work
8.2–8.4	Risk assessment, risk treatment, and impact assessment in operation	Completed, current assessments at planned intervals
9.1	Monitoring and measurement	Metrics tracked against objectives
9.2	Internal audit	Audit program and report from an independent, competent auditor
9.3	Management review	Minutes covering every required input

© 2026 RS Assurance & Advisory | rsassure.com

Clause	Requirement	What to show the auditor
10.1–10.2	Continual improvement; nonconformity and corrective action	Corrective action log with root causes

© 2026 RS Assurance & Advisory | rsassure.com

APPENDIX B

Annex A: the 38 controls

Plan to implement all 38. Exclusions are allowed only in the Statement of Applicability with a reason. If you only use AI and don't build it, some A.6 life cycle controls may be narrower in scope, but the auditor will still expect the reasoning written down.

ID	Area	Control	What to show
A.2.2	Policy	AI policy	Approved AI policy
A.2.3	Policy	Alignment with other policies	Cross-references to security, privacy, and HR policies
A.2.4	Policy	Review of the AI policy	Dated review record
A.3.2	Organization	AI roles and responsibilities	Role assignments for every AI system
A.3.3	Organization	Reporting of concerns	A channel employees can use, and its records
A.4.2	Resources	Resource documentation	AI system inventory with components
A.4.3	Resources	Data resources	Data sources documented per system
A.4.4	Resources	Tooling resources	Tools, models, and libraries listed
A.4.5	Resources	System and computing resources	Infrastructure and hosting documented
A.4.6	Resources	Human resources	Competence needed per role, and who has it
A.5.2	Impact	Impact assessment process	Written procedure
A.5.3	Impact	Documentation of impact assessments	Completed assessments, retained
A.5.4	Impact	Impact on individuals or groups	Assessment covering fairness, safety, rights
A.5.5	Impact	Societal impacts	Assessment of broader effects
A.6.1.2	Lifecycle	Objectives for responsible development	Documented objectives
A.6.1.3	Lifecycle	Responsible design and development processes	Development procedure
A.6.2.2	Lifecycle	Requirements and specification	Requirements per system
A.6.2.3	Lifecycle	Design and development documentation	Design records
A.6.2.4	Lifecycle	Verification and validation	Test and evaluation results
A.6.2.5	Lifecycle	Deployment	Deployment plan and approval
A.6.2.6	Lifecycle	Operation and monitoring	Monitoring for errors, drift, and misuse
A.6.2.7	Lifecycle	Technical documentation	Documentation for users and operators

© 2026 RS Assurance & Advisory | rsassure.com

ID	Area	Control	What to show
A.6.2.8	Lifecycle	Recording of event logs	Log settings and retained logs
A.7.2	Data	Data for development and enhancement	Data management process
A.7.3	Data	Acquisition of data	Sources and rights to use them
A.7.4	Data	Quality of data	Quality criteria and checks
A.7.5	Data	Data provenance	Where each dataset came from
A.7.6	Data	Data preparation	Preparation and labeling methods
A.8.2	Interested parties	Information for users	User-facing documentation
A.8.3	Interested parties	External reporting	A way for outsiders to report adverse effects
A.8.4	Interested parties	Communication of incidents	Incident communication plan
A.8.5	Interested parties	Information for interested parties	What you disclose, to whom
A.9.2	Use	Processes for responsible use	Acceptable use policy and procedures
A.9.3	Use	Objectives for responsible use	Documented use objectives
A.9.4	Use	Intended use	Intended use stated per system; misuse monitored
A.10.2	Third parties	Allocating responsibilities	Contracts splitting duties with partners
A.10.3	Third parties	Suppliers	Due diligence on AI vendors and hosted-model providers
A.10.4	Third parties	Customers	Customer requirements and commitments documented

Control titles are paraphrased from ISO/IEC 42001:2023. Check them against your licensed copy; Annex B gives implementation guidance for each.

© 2026 RS Assurance & Advisory | rsassure.com

APPENDIX C

Agent controls mapped to ISO 42001

Agent control	ISO 42001	What to show
1. Agent inventory with an owner	A.3.2, A.4.2, A.9.4	Inventory listing each agent's purpose, access, and named owner
2. Its own least-privilege identity	A.6.2.5; ISO 27001 A.5.15–A.5.18, A.8.2	Distinct identity per agent; access reviews
3. Human approval for high-impact actions	A.9.2, A.9.3; risk treatment (6.1.3)	Approval rules and approval records
4. Complete, tamper-resistant logs	A.6.2.6, A.6.2.8	Log configuration, retention, and review evidence
5. A kill switch you've tested	A.6.2.6, A.8.4	Documented stop procedure; last test date and result

© 2026 RS Assurance & Advisory | rsassure.com

ABOUT

RS Assurance & Advisory

RS Assurance & Advisory is an assurance and advisory firm headquartered in Southlake, Texas. We work exclusively with seasoned practitioners — no junior staff on engagements — and our unified methodology maps controls across frameworks so your team answers each evidence request once.

Service lines: SOC 1, SOC 2, and SOC 3 · SOX · ISO 27001 · ISO 42001 · CMMC 2.0 · HITRUST · HIPAA · GDPR · Nacha ACH Rules Compliance · PCI DSS readiness

Richard Stevenson
Managing Partner, RS Assurance & Advisory
richards@rsassure.com · (903) 229-0341 · rsassure.com

This paper is an educational starting point. It is not a substitute for ISO/IEC 42001:2023 itself, for the guidance of your certification body, or for legal advice. Control titles are paraphrased and should be checked against your licensed copy of the standard. Regulatory positions change; confirm current requirements with counsel before relying on any statement here in a contract or public claim. *© 2026 RS Assurance & Advisory. All rights reserved.*