

Practitioner guide

The Bridge

From ISO 9001 and ISO/IEC 27001 to ISO/IEC 42001.
What carries over, what must be rethought, and what
never existed before.

Michel Fotsing, CISSP
ISO/IEC 42001 Lead Auditor, PECB certified

Volunteer, ISC2 Exam Review
Author of *L'Architecte Numérique*, distributed by Hachette

September 2026
Edition 1.0 · Free to share

CONTENTS

What you will find here.

01	Why this guide	3
02	What ISO/IEC 42001 actually asks for	4
03	Your hidden advantage: the harmonized structure	5
04	Three things that did not exist before	6
	Impact on people	7
	The delegation boundary	8
	Drift surveillance	9
05	Annex A through ISO/IEC 27001 eyes	10
06	Where the standard meets regulation	11
07	The 90-day plan	12
08	Five mistakes I keep seeing	13
09	Before you talk about certification	14

Your current standards ask you to control your processes and protect your information. ISO/IEC 42001 also asks you to answer for what your automated systems do to people.

CHAPTER 01

Published in December 2023, ISO/IEC 42001 is the world's first certifiable standard for an artificial intelligence management system.

It was deliberately written on the same structure as ISO 9001 and ISO/IEC 27001.

Why this guide

This document is written for one specific person: the quality director, information security lead or compliance manager who has already built a management system, has kept it alive for years, and has just learned that the organization is deploying artificial intelligence.

That person has two reactions at almost the same moment. The first is that everything has to be rebuilt. The second is a quieter suspicion that the existing system may already cover a good part of the ground.

The second instinct is the right one. This guide shows it, clause by clause.

An organization that already runs ISO 9001 or ISO/IEC 27001 has the frame, the internal audit culture, the management review and the document control. It does not need to learn how to run a management system. It needs to know precisely what AI adds.

This guide does three things. It maps what carries over as is. It names what must be adapted. And it isolates the three requirements that existed in no management system standard before 2023, the ones where most organizations discover they have nothing.

What this guide is not

It is not an accredited course, not a substitute for the text of the standard, and not a guarantee of certification. It is a practitioner's field notes, written after doing the exercise myself.

CHAPTER 02

4

structural additions compared with ISO 9001 and ISO/IEC 27001

Clauses 4 to 10 carry the same numbering you already know.

What ISO/IEC 42001 actually asks for

One confusion keeps coming back, even among serious people: ISO/IEC 42001 does not certify an AI system. It certifies how an organization manages its AI systems.

The distinction changes everything. A mediocre model managed with rigor can conform. An excellent model deployed without a framework never will. As with ISO 9001 and ISO/IEC 27001, the standard requires the organization to understand its context, commit its leadership, plan, resource, operate, evaluate and improve. What sets it apart comes down to four additions.

Impact assessment, on top of risk assessment

Clause 6.1.4 and Annex A objective A.5 introduce a notion absent from ISO 9001 and ISO/IEC 27001: the impact of AI systems on individuals, groups and society. Risk, in your current standards, is risk to the organization. Impact, here, is an effect on third parties who never signed anything. It is the deepest shift in perspective in the standard.

The lifecycle as an object of governance

Objective A.6 covers the full lifecycle of an AI system, from requirements to retirement. A well-tuned industrial process stays tuned. A model does not. The standard draws the consequences.

Data as a governed resource

Objective A.7 treats training, validation and test data as a resource in its own right: acquisition, quality, provenance, preparation. For an ISO/IEC 27001 organization, this is a natural extension of information classification. For an ISO 9001-only organization, it is new work.

Responsible use, including by third parties

Objectives A.8 to A.10 cover information for interested parties and relationships with suppliers and customers. The standard applies as much to those who buy and use AI as to those who build it. Most small and mid-sized organizations are in the first group, and they are just as concerned.

CHAPTER 03

1

clause out of seven introduces a requirement with no equivalent

Your hidden advantage: the harmonized structure

Since 2012, every ISO management system standard has shared the same skeleton: ten clauses, in the same order, with the same intent. If your organization runs a living ISO 9001 or ISO/IEC 27001 system, the elements below already exist and do not need to be reinvented. They need to be extended to cover AI, which is weeks of work, not months.

Exhibit 1

What your current standards already gave you, clause by clause

Clause	ISO 9001 or 27001 gave you	ISO/IEC 42001 adds
4. Context ADAPT	Internal and external issues, interested parties, scope.	The organization's role regarding AI (developer, provider, user), and the parties affected by AI systems, not only by the organization.
5. Leadership ADAPT	Top management commitment, policy, roles and responsibilities.	An AI policy, standalone or integrated (A.2), and named roles for AI governance (A.3).
6. Planning NEW	Risk-based approach, objectives, planning of changes.	AI system impact assessment (6.1.4). The only clause in the common structure that carries a genuinely new requirement.
7. Support ADAPT	Resources, competence, awareness, communication, documented information.	AI-specific competence; documented data, tooling and computing resources (A.4).
8. Operation ADAPT	Operational planning and control.	Repeat risk and impact assessment in operation (8.2 to 8.4); govern the lifecycle (A.6).
9. Evaluation ADAPT	Monitoring, measurement, internal audit, management review.	Monitor the performance and drift of deployed systems, not only management-system indicators.
10. Improvement REUSABLE	Nonconformity, corrective action, continual improvement.	Nothing structurally new. AI incidents enter the process you already run.

Source: ISO/IEC 42001:2023, clauses 4 to 10, read against the ISO harmonized structure. Author's analysis.

The real work sits in Annex A, and in three requirements the rest of this guide examines one by one.

Three things that did not exist before.

Strip away everything that belongs to the harmonized structure and everything that naturally extends ISO/IEC 27001, and three requirements remain that I had never seen formalized in any management system standard.

They take the most thought, and often the most nerve.

FIRST ADDITION

Clauses 6.1.4 and 8.4

Objective A.5

The key word is *separate*. This is not one more line in your risk register. It is its own exercise, with its own method and documentation.

The privacy bridge. If you have run privacy impact assessments, you have done the equivalent, limited to personal data. Reuse the method; widen the scope.

Impact on people

Your current standards assess risk to the organization: financial loss, disruption, reputational damage, noncompliance. ISO/IEC 42001 requires a separate assessment: the impact of the AI system on the individuals and groups exposed to it, and on society.

Take a plant that deploys predictive maintenance on sensor data. The organizational risk is clear: a false negative costs a production stop. But the tool also flags equipment behavior anomalies, and some of them correlate with specific operators. The day the flag becomes an individual performance indicator, the question changes in nature.

That is impact on people, and the standard asks for it to be assessed before deployment, not after the first complaint.

What the assessment must contain

At a minimum: the people and groups affected, the nature of the effect, its severity, its reversibility, the mitigations, and a reasoned decision to proceed or not. It must be repeated when the context or the system changes significantly.

An ISO/IEC 42001 impact assessment is broader than a privacy assessment: it includes effects that have nothing to do with privacy, such as access to a service, discrimination or physical safety.

SECOND ADDITION**Objectives A.6 and A.9****Intended use, A.9.4**

Sorting decisions into the three zones is not a philosophical preference. It is an engineering and governance decision that must be written, dated and approved.

The delegation boundary

No management system standard has ever asked you to write down which decisions a machine may take on its own. ISO/IEC 42001 does, without using those words, through its requirements on intended use, human oversight and responsible use processes.

In my own work, I use three zones to draw that boundary.

Exhibit 2**The three delegation zones**

Zone	Who decides	The human's role
Automated	The system decides.	Review by sampling.
Augmented	The system proposes.	The human decides.
Sanctuary	The system does not enter.	Regardless of how well the system performs. An organizational choice, made in advance and written down.

Source: the Three Zones framework, *L'Architecte Numérique*, 2026.

What the standard expects in practice: every AI system has a documented intended use, out-of-scope uses are identified, and human oversight mechanisms are defined and proportionate. An organization that cannot answer "which decisions does this tool make on its own" does not meet the standard, even if everything else is in order.

The most useful deliverable I have seen is a one- to two-page delegation charter, signed by leadership, listing decisions by zone. It becomes the centerpiece of internal audit.

THIRD ADDITION

Clause 9.1

Control A.6.2.6

75%

The threshold I held for four years, written down nowhere.

Acceptance criteria, A.6.2.4.

Before any deployment, the standard expects criteria that are defined, documented and approved, and evidence they were met. If you have an informal threshold, formalize it first.

Drift surveillance

A well-tuned process stays tuned. A model does not. The world's data changes, behavior changes, and a system's accuracy slides without anyone noticing.

The standard requires that operating an AI system include continuous monitoring of its performance, with thresholds that trigger a response.

For an ISO 9001 organization, the closest analogy is statistical process control.

You already know how to watch a dimension drift on a production line. Apply the same discipline to a model: a performance indicator defined before deployment, a measurement frequency, an alert threshold, and a procedure when the threshold is crossed.

What is usually missing is not the measurement. It is the threshold written in advance. For four years I held a simple rule on a computer vision system: below 75% accuracy on the test set, we do not deploy. The rule was written nowhere. It lived in my head, and it left with me.

That is exactly what the standard comes to fix.

CHAPTER 05

38

controls across nine objectives, applied through a statement of applicability

Annex A through ISO/IEC 27001 eyes

As with ISO/IEC 27001, you justify the controls you include and those you exclude. For each objective, what an ISO/IEC 27001 organization already has and what it must build.

Exhibit 3

The nine Annex A objectives: in place and to build

Objective	No.	If you have ISO/IEC 27001	To build
A.2 Policies related to AI	3	An information security policy and review cycle.	An AI policy, aligned, reviewed periodically.
A.3 Internal organization	2	Formal roles; a channel for reporting concerns.	Extend roles to AI.
A.4 Resources for AI systems	5	Asset inventory; human and technical resources.	Document data, tooling, compute and competence per system.
A.5 Assessing impacts	4	Nothing directly equivalent. A privacy assessment is the closest starting point.	The full process: method, documentation, individual and societal impacts. The main build.
A.6 AI system lifecycle	9	Secure development, change management, logging.	Responsible development, verification and validation, monitoring in operation, event logs.
A.7 Data for AI systems	5	Classification, asset ownership, protection.	Acquisition, quality, provenance, preparation. An extension of classification.
A.8 Information for interested parties	4	Communication, incidents, notification duties.	User documentation, external reporting, AI-related incidents.
A.9 Use of AI systems	3	Acceptable use policy.	Responsible use, documented intended use per system. The delegation boundary lives here.
A.10 Third-party and customer relationships	3	Suppliers, contract clauses, third-party assessment.	Allocation of responsibilities across the AI value chain.

Source: ISO/IEC 42001:2023, Annex A, read against ISO/IEC 27001:2022. Author's analysis.

Six objectives out of nine extend controls you already run. Two are partly covered. Only one starts from zero. For an ISO 9001 organization without ISO/IEC 27001, the shortest path is often a joint implementation of both standards.

CHAPTER 06**Voluntary, but not alone.**

ISO/IEC 42001 is a voluntary standard. It increasingly sits next to frameworks and laws that are not.

A delegation boundary drawn once answers several of these obligations at the same time.

Where the standard meets regulation

Few organizations adopt ISO/IEC 42001 in a vacuum. The value of a management system is that one set of evidence can answer several expectations at once.

NIST AI Risk Management Framework

In the United States, the NIST AI RMF is the reference most enterprises and federal contractors already use. It describes good practice across its govern, map, measure and manage functions, but it certifies nothing. ISO/IEC 42001 gives that same practice a certifiable management system structure, which is what a customer, an auditor or an insurer can actually verify. The two map to each other closely; running one makes the other cheaper.

The EU AI Act

The AI Act has applied its transparency obligations since August 2026, with high-risk requirements phasing in from December 2027. It does not require ISO/IEC 42001. But a conforming AI management system produces much of what the Act expects: risk management, technical documentation, human oversight and post-market monitoring.

Automated decisions and privacy law

Privacy laws increasingly regulate decisions made by machines. Quebec's Law 25, for example, requires organizations to inform a person when a decision about them is based exclusively on automated processing, and to explain the main factors on request. A decision placed in the automated zone of your delegation charter is exactly the kind of decision these rules target. An organization that has drawn its zones knows immediately which decisions are in scope.

Run your privacy impact assessment and your ISO/IEC 42001 impact assessment as one exercise with two parts, rather than two documents that contradict each other six months later.

CHAPTER 07

90

days to move from ungoverned to governed AI use, with evidence

This plan assumes a living ISO 9001 or ISO/IEC 27001 system and one to five AI systems in use or in development.

After 90 days. Certification becomes an option to weigh against business need, not an end in itself. Many organizations stop here, rightly, until a customer or regulator asks for the certificate.

The 90-day plan

01 Days 1 to 30 · See clearly

Inventory the AI systems in use, including those employees use without formal approval. The inventory almost always surprises. For each system: who uses it, on what data, for which decisions. Extend the scope of the existing management system. Name an owner for AI governance, even part-time.

Deliverable: a one-page inventory per system, and a scope decision approved by leadership.

02 Days 31 to 60 · Draw the boundaries

For each system, write the intended use and sort its decisions into the three zones. Draft the delegation charter. Run a first impact assessment on the most exposed system, reusing an extended privacy assessment method. Formalize acceptance criteria and monitoring thresholds.

Deliverable: the signed delegation charter, one complete impact assessment, and written thresholds.

03 Days 61 to 90 · Make it hold

Bring AI into the existing internal audit program. Add AI to the management review agenda. Put drift monitoring in place on at least one system. Write the AI policy, or amend the existing security policy. Train the process owners concerned.

Deliverable: a first internal audit covering AI, and an approved policy.

CHAPTER 08

Five mistakes observed in serious organizations, with good systems, approaching ISO/IEC 42001 for the first time.

Five mistakes I keep seeing

01 Treating impact assessment as a risk register line

The most common mistake, and an understandable one: the register exists, it is comfortable, you add a line. But impact on people is not risk to the organization, and an auditor will see it immediately.

02 Confusing governing AI with securing AI

Security teams read ISO/IEC 42001 as a second ISO/IEC 27001. Much of it is familiar, but objectives A.5 and A.9 are not security questions. They are questions of accountability and organizational choice. They need leadership at the table, not only IT.

03 Documenting without deciding

An organization can produce a policy, an inventory and a charter without ever settling a single hard boundary. A charter that lists only obvious cases is worthless. The contested cases, where leadership hesitates, are what give it value.

04 Forgetting purchased systems

The standard covers use, not only development. The assistant in your office suite, the recommendation module in your CRM, the resume screening tool from your HR vendor: all are in scope. Most small organizations build no AI in-house and are still fully concerned.

05 Chasing certification before the need

Certification costs money and attention. It is worth it when a customer, regulator or partner requires it, or when it opens a market. A living, uncertified AI management system beats a certified one that sleeps in a binder.

CHAPTER 09

Going further. The Three Zones framework is developed in full in *L'Architecte Numérique : Orchestrer les intelligences à l'ère de l'IA*, distributed by Hachette (in French). It takes up the question this guide only touches: not how to document the boundary, but where to draw it.

The author. Cybersecurity architect, CISSP, PECB-certified ISO/IEC 42001 Lead Auditor, ISC2 exam review volunteer. Four years as product owner of a computer vision system in production.

michel@larchitectenumerique.com

Before you talk about certification

There is something reassuring in discovering that a standard published in 2023 formalizes practices that reasonable people had already found, each on their own, out of necessity.

Thresholds before deployment, decisions we refuse to hand to a machine, monitoring after go-live: none of it fell from the sky with the standard. But all of it rested on people. On their judgment, their rigor, their presence. The day they leave, the practice leaves with them.

That is what a management system really fixes. Not making people better. Making sure the good things they do survive their departure.

If your organization already runs ISO 9001 or ISO/IEC 27001, you have done half the journey without knowing it. The other half is three requirements and ninety days. This guide gives you the map. The rest is a matter of decision.

Pass it on.

This guide may be shared, quoted and distributed freely, in full and unmodified, with credit to the author. If it is useful, send it to the person in your organization who will need it.

For use in training or advisory work, write to me.