InfoSecAI

A BRIEF FOR CISOS · BISOS · GRC LEADERSHIP

FIELD NOTE · 2026

The ISO 42001 Documentation That Decides Your AIMS Certification

A practitioner field note for CISOs, Heads of AI Governance, AIMS managers and lead implementers preparing for certification.

AUTHORED BY

Paul Jolliffe

Paul Jolliffe, Founder, InfoSecAI · MBA · CISSP · ISO 27001:2022 LA / LI / IA · PRINCE2 Practitioner

01 · THE INVENTORY AND THE IMPACT ASSESSMENTS, NOT THE POLICY

The inventory and the impact assessments, not the policy

ISO/IEC 42001:2023 is new enough that most organisations approach it as a writing exercise. Produce an AI policy, draft a handful of procedures, book the audit. That misframes it from the first day. The document that fails a 42001 audit is rarely the policy. It is the AI system inventory and the impact assessments, because those two artefacts require you to actually know what AI you are running, who owns it, and what it could do to the people on the other side of its outputs. Most organisations do not, and the gap shows the moment an auditor asks to trace one system end to end.

42001 certifies an AI Management System, the AIMS. It shares the harmonised structure of ISO 27001: the same Clauses 4 to 10, the same Plan-Do-Check-Act backbone, the same auditable mechanics. If you have been through 27001, the management-system scaffolding will feel familiar. What is genuinely different is the spine.

In 27001 the integrating artefact is the risk assessment. In 42001 it is the AI system impact assessment, the requirement at Clause 6.1.4 and control A.5.2. Every AI system in scope receives one. The impact assessment drives control selection, the lifecycle gates, and the post-deployment monitoring regime. If the impact assessments are thin, everything downstream is built on sand. Treat that document as the one the rest of the AIMS hangs from.

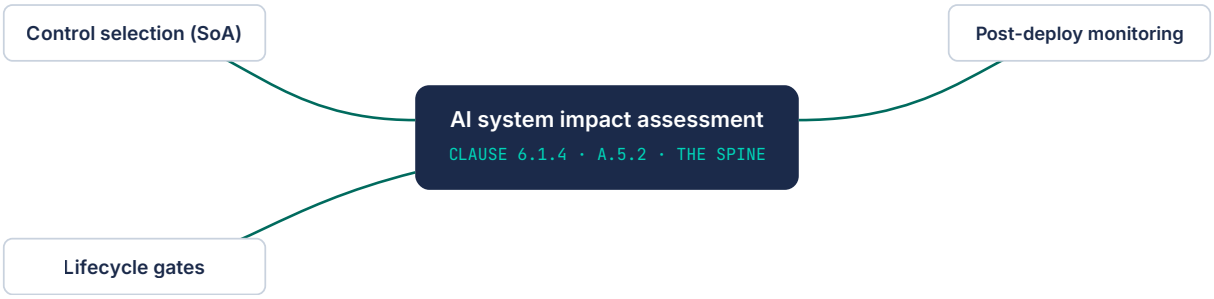

Figure 1. The impact assessment is the artefact the rest of the AIMS hangs from.

As with any management system, the documentation splits into mandatory and supporting. Confusing the two produces the same two failure modes: teams that under-document because they did not realise a record was required, and teams that over-document because they treated all 38 Annex A controls as demanding a separate policy each.

02 · MANDATORY DOCUMENTED INFORMATION

Mandatory documented information

These are the documents and records ISO 42001 explicitly requires by clause. A certification body will ask for every one.

DOCUMENT	CLAUSE	WHAT THE AUDITOR IS CHECKING
AIMS scope statement	4.3	Boundaries of the AIMS. Which AI systems are in, which roles the organisation plays (provider, deployer, or both), which markets.
AI policy	5.2	Approved by top management, communicated, and it references the AI systems in scope and the impact assessment process rather than reading as generic.
AI risk assessment process	6.1.2	A repeatable method with defined risk-acceptance criteria, covering AI-specific impact dimensions, not just confidentiality, integrity, and availability.
AI risk treatment process	6.1.3	Treatment options and the logic for selecting Annex A controls.
AI system impact assessment process	6.1.4	The method for assessing consequences to individuals, groups, and society. This is the spine of the AIMS.
Statement of Applicability	6.1.3	All 38 Annex A controls accounted for: applicable or not, with justification and implementation status.
AI objectives	6.2	Measurable, with owners, timelines, and resources.
Evidence of competence	7.2	The people governing, building, and deploying AI can be shown to be competent for it.
AI risk assessment results	8.2	The actual output: a completed, scored AI risk register.
AI risk treatment results	8.3	Actions, owners, dates, residual risk, and risk-owner sign-off.
AI system impact assessment results	8.4	A completed impact assessment for every AI system in scope. An absent one is a major nonconformity on its own.
Monitoring and measurement results	9.1	Evidence the AIMS and the AI systems are measured, including model performance and drift.
Internal audit programme and results	9.2	A schedule, reports, and evidence the auditor was independent of what they audited.
Management review results	9.3	Minutes covering every required input and output, with decisions and actions.
Nonconformities and corrective actions	10.2	A log with root cause analysis, corrective action, and an effectiveness check.

Where 42001 audits break most often: missing or incomplete impact assessments, an AI system inventory that does not match what is actually in production, and a Statement of Applicability with controls marked applicable but never implemented.

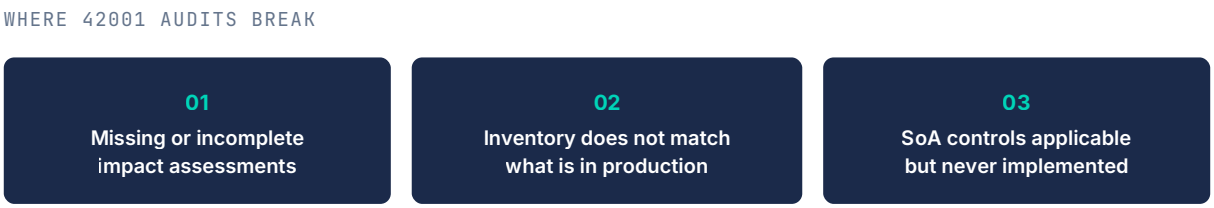

Figure 2. The three places certification most often fails.

03 · SUPPORTING DOCUMENTATION

Supporting documentation

These are not named by the clauses. They are how you demonstrate that the Annex A controls you marked applicable genuinely operate. The Annex A category is shown rather than a sub-control number, because the precise control numbering belongs in your Statement of Applicability, not in a reading list.

DOCUMENT	ANNEX A CATEGORY
AI roles, responsibilities and authorities	A.3 Internal organisation
AI resources documentation: data, tooling, compute, people	A.4 Resources for AI systems
AI system impact assessment methodology and template	A.5 Assessing impacts of AI systems
AI system inventory	A.6 AI system life cycle
Data management for AI policy	A.7 Data for AI systems
Data quality and provenance procedure	A.7 Data for AI systems
Model development and training procedure	A.6 AI system life cycle
Model evaluation procedure: bias, fairness, robustness	A.6 AI system life cycle
Deployment approval gate procedure	A.6 AI system life cycle
AI system monitoring procedure	A.6 AI system life cycle
AI incident response procedure	A.6 AI system life cycle
AI system decommissioning procedure	A.6 AI system life cycle
Transparency notices for AI systems	A.8 Information for interested parties

DOCUMENT	ANNEX A CATEGORY
AI system documentation pack, the system card	A.8 Information for interested parties
Human review and right-to-object procedure	A.8 Information for interested parties
AI acceptable use procedure	A.9 Use of AI systems
AI supplier policy and due diligence questionnaire	A.10 Third-party and customer relationships

04 · ONE DOCUMENT, SEVERAL CONTROLS

One document, several controls

There is no requirement for 38 separate policies. A single AI lifecycle procedure can carry development, evaluation, deployment, monitoring, and decommissioning under the A.6 category. Auditors accept consolidated documents, provided every control referenced inside them is genuinely addressed. As with 27001, the real risk is not too few documents but too many, each thinner than the control it claims to cover.

CONSOLIDATION UNDER A.6 AI SYSTEM LIFE CYCLE

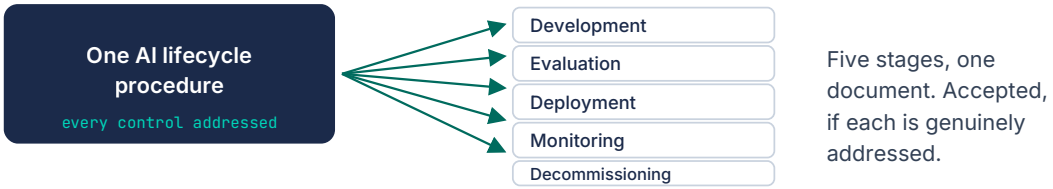

Figure 3. Consolidation under A.6, accepted when every stage is genuinely addressed.

05 · THE POINT THAT CATCHES PEOPLE: 42001 IS NOT THE EU AI ACT

The point that catches people: 42001 is not the EU AI Act

This is the distinction that does the most damage when it is missed. ISO 42001 is a voluntary management-system standard. It tells you how to govern AI responsibly. It does not, on its own, discharge a binding regulatory obligation. An organisation that is an EU AI Act provider or deployer of a high-risk system still owes that regulation its specific requirements: the fundamental rights impact assessment, conformity assessment, post-market monitoring, the transparency duties. The AIMS is the vehicle that operationalises those obligations, but the certificate is not a substitute for them.

Figure 4. A certificate is not compliance. State the regulatory overlay in your AIMS scope.

State the regulatory overlay explicitly in your AIMS scope. An auditor reviewing a 42001 system in a regulated context will expect to see the AI Act, NIST AI RMF, or the relevant sector regulator named and mapped, not assumed.

06 · WHAT GOOD LOOKS LIKE BEFORE THE AUDIT

What good looks like before the audit

Run the mandatory list as a hard gate: every item exists, is approved, is current, and an owner can speak to it. Then take the AI system inventory and pick one system at random. Walk it from the AI policy, to its impact assessment, to the controls selected in the Statement of Applicability, to the operational evidence that those controls run, to any corrective action raised against it. If you can do that walk in under five minutes, you are close to ready. If you cannot, you have found the gap before the auditor did.

The standard is not asking you to be impressive about AI. It is asking you to be able to account for it.

InfoSecAI

About InfoSecAI

InfoSecAI is an independent UK consultancy helping organisations turn security, regulatory, resilience and AI governance requirements into practical operating models, stronger controls and robust delivery.

We work across strategy, governance, risk, compliance, AI security, assurance, operations and engineering. Our services help leadership teams assess their current position, align to standards and regulation, define the target operating model, and deliver the governance, controls, artefacts and ways of working needed to move from intent to implementation.

Our toolkit capability accelerates structured work across ISO 27001, ISO 22301, ISO 42001, NIST CSF, NIST AI RMF, CIS Controls, Cyber Essentials, DORA, NIS 2, the EU AI Act, GDPR, UK GDPR, SOC 1 and SOC 2. The approach combines AI-enabled workflow support with senior practitioner judgement, so outputs remain proportionate, usable and connected to the way the organisation actually operates.

InfoSecAI was founded in **2025** by **Paul Jolliffe**. The company is built for organisations that need clarity, senior leadership and hands-on delivery across information security and AI governance, without adding unnecessary complexity or treating compliance as a paperwork exercise.

infosecai.net · paul.jolliffe@infosecai.net