

ISO 42001 Lead Auditor practice questions and answers

Original practice questions written from ISO/IEC 42001:2023 and ISO/IEC 23894:2023. Not real exam questions.

Ten questions, then the answer key with the correct option, a plain-English rationale and the clause each answer comes from. Answer all ten before you turn to the answers.

Prefer instant scoring? The same questions are online at mindsetprep.com.au/free-practice-questions/iso-42001-lead-auditor/.

QUESTION 1

Per ISO/IEC 42001:2023, an organisation determines that a recruitment-screening AI may affect candidates who do not interact with the organisation directly. Which assessment is the appropriate vehicle to document this concern?

- A. The organisation's privacy impact assessment, separate from the AIMS
- B. The Risk Treatment Plan (Clause 6.1.3) once treatment options are selected
- C. AI system impact assessment (Clause 6.1.4) – because the concern is about effects on individuals and groups
- D. AI risk assessment (Clause 6.1.2) – because all 42001 concerns flow through the risk register

QUESTION 2

An organisation is defining its ISO/IEC 42001 AIMS scope. They operate three AI systems: a chatbot, a fraud detector, and an internal HR-screening tool. Which statement correctly distinguishes AIMS scope from AI system scope?

- A. The AIMS scope (Clause 4.3) defines the organisational boundary of the management system; each AI system has its own scope describing purpose, inputs, outputs and operational context
- B. The AIMS scope IS the list of AI systems in scope – they are the same artefact
- C. Individual AI systems do not need separate scoping once the AIMS scope is defined
- D. Each AI system's scope must match the AIMS scope exactly to maintain conformance

QUESTION 3

ISO/IEC 42001:2023 introduces specific role concepts to AIMS governance. An auditor reviewing role assignments should expect to see explicit assignment of:

- A. Risk owners for AI risks and parties responsible for AI impact assessment, alongside the generic top-management role under Clause 5.1
- B. Only the generic top-management role from Clause 5.1, as 42001 inherits roles from 27001 without addition
- C. The Information Security Officer role, repurposed to cover AI governance
- D. Three independent AI ethics officers, one per AI system in scope

QUESTION 4

An ISO/IEC 42001 implementer is preparing to onboard a new AI system into the AIMS. Per the standard, which assessment(s) must be completed before the AI system is deployed into operation?

- A. Only the AI system impact assessment – the risk assessment is performed at the AIMS level, not the system level
- B. Both the AI risk assessment AND the AI system impact assessment, because they address distinct concerns and neither substitutes for the other
- C. Only the AI risk assessment – the impact assessment is required only for AI systems with consumer-facing outputs
- D. Neither at deployment – both are performed during the first management review after operation begins

QUESTION 5

Mid-cycle, the organisation decommissions one of three AI systems within its AIMS scope. Per ISO/IEC 42001:2023, the appropriate response is:

- A. Update the AI system inventory and the relevant AI system impact assessments; the AIMS scope itself only changes if decommissioning materially alters the organisational boundary
- B. Take no documented action – the AIMS scope is set at certification and does not change
- C. Update the AIMS scope statement to remove the decommissioned system
- D. Trigger a full AIMS recertification because scope has changed

QUESTION 6

Per ISO/IEC 42001:2023, who has formal responsibility for approving residual AI risks after treatment?

- A. The certification body's lead auditor, before issuing the certificate
- B. The AI system's product manager, because they own the system's outputs
- C. Designated management, from whom Clause 6.1.3 requires the organisation to obtain approval
- D. Top management, because the AI policy under Clause 5.2 is a top-management artefact

QUESTION 7

Per ISO/IEC 42001:2023 Clause 6.1.4, an AI system impact assessment should address:

- A. Compliance impact under privacy regulations only
- B. Potential consequences of the AI system for individuals, groups of individuals, and societies, in identified contexts of use
- C. The likelihood and severity of organisational risks the AI system creates
- D. Financial impact to the organisation if the AI system underperforms or fails

QUESTION 8

An organisation operates an AI system trained on third-party data and now runs inference on customer inputs. Per ISO/IEC 42001:2023 (with ISO/IEC 23894:2023 guidance), the AI system scope should cover:

- A.** Only the inference phase, because training was completed by a third party and is out of scope
- B.** Only the training phase, because that is where the model's behaviour is determined
- C.** Both the training phase (including the third-party data source) and the inference phase across the AI lifecycle the organisation operates
- D.** Only the deployed model artefacts, not the data pipelines around them

QUESTION 9

A new AI governance role appears in an organisation's AIMS documentation: 'AI Ethics Officer'. Per ISO/IEC 42001:2023, how should this role be evaluated by the auditor?

- A.** The role is equivalent to the ISO/IEC 27001 Information Security Officer and inherits its responsibilities
- B.** The role is mandatory under Clause 5.3 and the auditor should raise a nonconformity if absent
- C.** The standard does not prescribe an AI Ethics Officer specifically; verify the role's responsibilities map to standard roles (risk owners, impact owners, top management) and that gaps or duplications are documented
- D.** The role is informative only and out of audit scope

QUESTION 10

An ISO/IEC 42001 auditor reviews the organisation's risk register and notes that 'societal harm severity' is being used as a factor in the risk-scoring formula. Per the standard, this is:

- A.** Correct practice – Clause 6.1.4 requires the impact assessment results to be considered in the risk assessment, and 6.1.2 permits using the impact assessment when assessing consequences
- B.** Required – risk-scoring under 42001 must include societal harm to be conformant
- C.** A mis-mapping – societal harm severity belongs in the AI system impact assessment (Clause 6.1.4), not the risk assessment (Clause 6.1.2); flag the conflation
- D.** Optional – 42001 does not specify the inputs to risk-scoring

Answers

Q1 Correct answer: **C**. AI system impact assessment (Clause 6.1.4) – because the concern is about effects on individuals and groups

ISO/IEC 42001:2023 separates two distinct artefacts: AI risk assessment (Clause 6.1.2) addresses risks to the organisation's own AI objectives, and AI system impact assessment (Clause 6.1.4) addresses impacts on individuals, groups, and societies. Effects on candidates external to the organisation belong in the impact assessment. Option D is the textbook T-01 trap – sending the concern to the risk register when impact assessment is the correct artefact. Option A bypasses the AIMS, which is incorrect when AI is in scope. Option B mis-orders the artefact chain.

Cited: ISO/IEC 42001:2023, Clauses 6.1.2 and 6.1.4

Trap pattern: T-01 AI impact vs AI risk assessment

Q2 Correct answer: **A**. The AIMS scope (Clause 4.3) defines the organisational boundary of the management system; each AI system has its own scope describing purpose, inputs, outputs and operational context

ISO/IEC 42001:2023 Clause 4.3 defines the AIMS scope – the organisational and functional boundary of the management system. Each AI system within scope has its own context-of-use scope describing what it does, the data it processes, and where it operates. Option B is the T-02 trap: collapsing two distinct scoping decisions into one. Options C and D either remove the system-level scope or force it to match the AIMS scope – both corrupt the standard's two-level model.

Cited: ISO/IEC 42001:2023, Clause 4.3

Trap pattern: T-02 AIMS scope vs AI system scope

Q3 Correct answer: **A**. Risk owners for AI risks and parties responsible for AI impact assessment, alongside the generic top-management role under Clause 5.1

ISO/IEC 42001:2023 introduces AI-specific responsibilities – risk owners for AI risks under Clause 6.1.3, and parties responsible for AI impact assessment under Clause 6.1.4 – in addition to generic management-system roles, because AI governance touches both technical and societal concerns that the standard treats distinctly. Option B is the T-03 trap: claiming 42001 inherits 27001 roles verbatim understates the AI-specific structure. Option C wrongly repurposes a 27001 role. Option D invents a structural rule (one ethics officer per system) the standard does not prescribe.

Cited: ISO/IEC 42001:2023, Clauses 5.3 and 6.1

Trap pattern: T-03 AI roles vs generic management roles

Q4 Correct answer: **B**. Both the AI risk assessment AND the AI system impact assessment, because they address distinct concerns and neither substitutes for the other

ISO/IEC 42001:2023 requires both assessments. Clause 6.1.2 risk assessment covers risks to the organisation's AI objectives; Clause 6.1.4 impact assessment covers impacts on individuals, groups, and societies. The two are distinct artefacts and neither substitutes for the other (T-01 trap).

Option C introduces a consumer-facing carve-out the standard does not contain. Option A wrongly disaggregates the two assessments by level. Option D defers both, contradicting the planning-before-operation logic of Clauses 6 and 8.

Cited: ISO/IEC 42001:2023, Clauses 6.1.2 and 6.1.4

Trap pattern: T-01 AI impact vs AI risk assessment

Q5 Correct answer: **A**. Update the AI system inventory and the relevant AI system impact assessments; the AIMS scope itself only changes if decommissioning materially alters the organisational boundary

ISO/IEC 42001:2023 Clause 4.3 sets the AIMS scope at the organisational and functional level.

Removing one AI system from operation is a change to the AI system inventory (a Clause 7.5 documented-information matter) and to the impact and risk assessments for that system, but it does not automatically alter the AIMS scope unless the organisational boundary is materially affected. Option B is the T-02 trap: confusing the AIMS scope with the inventory of AI systems within it. Option C overreacts; Option D under-reacts.

Cited: ISO/IEC 42001:2023, Clauses 4.3 and 7.5

Trap pattern: T-02 AIMS scope vs AI system scope

Q6 Correct answer: **C**. Designated management, from whom Clause 6.1.3 requires the organisation to obtain approval

ISO/IEC 42001:2023 Clause 6.1.3 requires the organisation to obtain approval from designated management for the AI risk treatment plan and for acceptance of the residual AI risks. Option D is the T-03 trap: Clauses 5.1 and 5.2 say 'top management' explicitly, but 6.1.3 says 'designated management' – answering 'top management' substitutes a Clause 5 term for the one 6.1.3 uses, and the AI policy being a top-management artefact does not settle who accepts residual risk. Clause 6.1.2's AI risk assessment process, steps (a) to (e), does not require a risk owner to be identified, and 6.1.3 assigns residual-risk acceptance to designated management. Option B invents a product-manager role for residual-risk approval. Option A mis-frames the auditor's role – certification bodies review for conformance, they do not approve risk decisions.

Cited: ISO/IEC 42001:2023, Clauses 6.1.3, 5.1 and 5.2

Trap pattern: T-03 AI roles vs generic management roles

Q7 Correct answer: **B**. Potential consequences of the AI system for individuals, groups of individuals, and societies, in identified contexts of use

Clause 6.1.4 directs the impact assessment at consequences for individuals, groups, and societies, in identified contexts of use. This is the differentiator from the AI risk assessment (Clause 6.1.2), which addresses organisational risk. Option D re-frames the impact assessment as financial-to-organisation, which is risk territory. Option C is verbatim risk-assessment language – the T-01 swap. Option A narrows impact to one compliance domain when the standard scopes broader societal impact.

Cited: ISO/IEC 42001:2023, Clause 6.1.4

Trap pattern: T-01 AI impact vs AI risk assessment

Q8 Correct answer: **C**. Both the training phase (including the third-party data source) and the inference phase across the AI lifecycle the organisation operates

ISO/IEC 42001:2023 (with ISO/IEC 23894:2023 guidance) treats the AI system scope as covering the AI lifecycle the organisation operates: design, development, deployment, operation. A model trained externally then deployed internally is still in scope for the phases the organisation is responsible for, training-data lineage included. Option A is the T-02 trap: confining system scope to the operational moment and excluding the training source the organisation relies on. Option B is the inverse mistake. Option D is too narrow.

Cited: ISO/IEC 42001:2023 with ISO/IEC 23894:2023 (AI lifecycle guidance)

Trap pattern: T-02 AIMS scope vs AI system scope

Q9 Correct answer: **C**. The standard does not prescribe an AI Ethics Officer specifically; verify the role's responsibilities map to standard roles (risk owners, impact owners, top management) and that gaps or duplications are documented

ISO/IEC 42001:2023 names role concepts (risk owners, impact owners, top management) but does not prescribe specific job titles. Organisations frequently invent AI Ethics Officer roles; the auditor's job is to map declared roles to the standard's role concepts, not to require or reject specific titles. Option C captures this. Option B over-applies the standard (T-03 trap: inventing a mandatory role the standard does not name). Option A wrongly imports a 27001 role. Option D removes the role from audit scope, which is incorrect once the organisation has assigned responsibilities to it.

Cited: ISO/IEC 42001:2023, Clause 5.3

Trap pattern: T-03 AI roles vs generic management roles

Q10 Correct answer: **A**. Correct practice – Clause 6.1.4 requires the impact assessment results to be considered in the risk assessment, and 6.1.2 permits using the impact assessment when assessing consequences

The artefacts do have different scopes – Clause 6.1.2 assesses risks to the organisation's AI objectives, Clause 6.1.4 assesses consequences for individuals, groups and societies – but 42001 bridges them rather than walling them apart. Clause 6.1.4 closes: 'The organization shall consider the results of the AI system impact assessment in the risk assessment (6.1.2)', and 6.1.2 closes: 'When assessing consequences, the organization can utilize an AI system impact assessment as indicated in 6.1.4.' Feeding societal harm severity into the consequence limb of 6.1.2 (d) is that bridge working, so Option C raises a finding the standard's own text forecloses – the T-01 trap here is over-separating, not conflating. Option B overstates: the shall is to consider impact results, while the scoring mechanism is a 'can', so no particular formula is mandated. Option D fails against 6.1.1, which requires AI risk criteria to be established and maintained.

Cited: ISO/IEC 42001:2023, Clauses 6.1.2 and 6.1.4

Trap pattern: T-01 AI impact vs AI risk assessment

Want the full bank?

Mindset Prep has 500+ clause-cited ISO 42001 Lead Auditor questions, each with the cited clause and a rationale, and an adaptive engine that brings back the clauses you miss. Start with a 3-day free trial.

[Start the 3-day free trial](https://app.mindsetprep.com.au/start?source=free-questions-pdf&exam=iso-42001-la)

<https://app.mindsetprep.com.au/start?source=free-questions-pdf&exam=iso-42001-la>