

ISO 42001 Lead Implementer 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-implementer/.

QUESTION 1

ISO/IEC 42001:2023 Annex A is labelled normative. An auditor reviewing the SoA finds the implementer has selected 23 of the listed Annex A controls. Per the standard, this is:

- A.** A nonconformity – 'normative' means all listed controls must be implemented in conformant AIMSES
- B.** A nonconformity – the standard requires at least 30 of the Annex A controls to be in scope
- C.** Acceptable only if the certification body has approved the selection
- D.** Acceptable – Annex A is a reference list; selection is driven by the AI risk and impact assessments under Clause 6.1.3, with exclusions justified in the SoA

QUESTION 2

An implementer is mapping the organisation's MLOps practices onto the ISO/IEC 42001 AIMS. The MLOps team operates in two-week sprints with continuous deployment. Per the standard, the implementer should:

- A.** Map MLOps activities onto the AIMS AI lifecycle phases (design, development, deployment, operation); the sprint cadence is an internal practice, not an AIMS lifecycle phase
- B.** Adopt the two-week sprint as the formal AIMS lifecycle cycle, since 42001 inherits whatever delivery cadence the implementing team uses
- C.** Document MLOps separately from the AIMS – they are not compatible
- D.** Replace MLOps with a waterfall AI lifecycle to align with the standard

QUESTION 3

An organisation has two policy documents: an 'AI Ethics Policy' (board-approved, sets the organisation's ethical principles for AI use) and an 'AI Management Policy' (top-management-approved, operationalises Clause 5.2). Per ISO/IEC 42001:2023, the auditor should treat these as:

- A. Distinct documents with different scope, signatories, and review cadence; verify both are governed and that the Clause 5.2 AI policy commitments are explicitly addressed in at least one
- B. Only the AI Ethics Policy matters for audit, since 42001 inherits ethics framing from ISO/IEC 23894
- C. The same document with two titles; raise a finding if both exist
- D. Only the AI Management Policy matters; ethics policies are out of audit scope

QUESTION 4

An organisation has excluded three Annex A controls from its 42001 implementation because they are not relevant to the AI systems in scope. Per the standard, the exclusions:

- A. Are not permitted – Annex A is normative, so all listed controls must be included
- B. Must be approved by the certification body in advance
- C. Are only permitted for AI systems with no consumer-facing outputs
- D. Must be justified in the SoA with reasons referencing the risk and impact assessment outcomes

QUESTION 5

An implementer is placing a specific activity into a lifecycle phase. The activity is 'monitoring of model performance against agreed thresholds, with retraining triggered when thresholds are breached'. Per ISO/IEC 42001:2023 (with ISO/IEC 23894:2023 guidance), this activity belongs to the:

- A. Deployment phase – monitoring is a deployment-time concern
- B. Operation phase – the AI system is deployed and monitored against operational criteria
- C. Design phase – thresholds are set in the design specification
- D. Development phase – retraining is part of model development

QUESTION 6

Per ISO/IEC 42001:2023 Clause 5.2, the AI policy shall:

- A. Be appropriate to the organisation's purpose, provide a framework for AI objectives, include a commitment to satisfy applicable requirements, and include a commitment to continual improvement of the AIMS
- B. Whatever the organisation chooses – Clause 5.2 leaves AI policy content to the organisation's discretion
- C. Define the organisation's ethical principles for AI use, including fairness, accountability, transparency, and explainability
- D. Both – the AI policy must combine management-system commitments and ethical principles in a single document

QUESTION 7

An implementer is reviewing the structure of ISO/IEC 42001:2023. Which statement correctly describes Annex A and Annex B?

- A.** Annex A is normative; Annex B is normative, and B items must appear in the SoA with the same inclusion/exclusion justification as Annex A
- B.** Annex A is normative (reference controls); Annex B is normative implementation guidance, and Clause B.1 specifies that B items do not require SoA justification
- C.** Both Annex A and Annex B are informative; the SoA covers only Clause 6 selections
- D.** Annex A is normative; Annex B is informative implementation guidance

QUESTION 8

An implementer is documenting how the organisation handles training–data collection. Per ISO/IEC 23894:2023 lifecycle guidance referenced by ISO/IEC 42001:2023, training–data collection belongs primarily to the:

- A.** Deployment phase – data is loaded when the system is deployed
- B.** Operation phase – data is collected continuously during system use
- C.** Design phase – data sources are specified in the design
- D.** Development phase – data collection is part of building the AI system

QUESTION 9

An organisation has both an AI Policy (top–management–approved, Clause 5.2) and an AI Ethics Policy (board–approved, sets the organisation's ethical commitments). The AI Policy commits the AIMS to operate in accordance with the AI Ethics Policy. Per ISO/IEC 42001:2023, this arrangement is:

- A.** Out of scope – the AI Ethics Policy is voluntary and the auditor should ignore it
- B.** Non-conformant – the AI Ethics Policy must be top–management–approved, not board–approved
- C.** Non-conformant – 42001 requires a single AI Policy, so the Ethics Policy should be merged into it
- D.** Acceptable – the two documents are distinct and the cross-reference makes the relationship explicit; the auditor should verify both are governed

QUESTION 10

An implementer has selected an Annex A control into the SoA but has not yet implemented it. The auditor sees this during a Stage 2 audit. Per ISO/IEC 42001:2023, the correct framing is:

- A.** Each Annex A control included in the SoA must actually be implemented under Clauses 8.1 and 8.3; 'selected but not implemented' is permitted only where it is carried in the management–approved AI risk treatment plan – otherwise it is a nonconformity
- B.** Acceptable in all cases – Annex A controls are guidance, not requirements, so non-implementation is not a finding
- C.** Acceptable – the SoA is a planning document, so implementation status is reviewed in surveillance audits only
- D.** Always non-conformant – any Annex A control in the SoA must be fully implemented by the time of certification audit

Answers

Q1 Correct answer: **D**. Acceptable – Annex A is a reference list; selection is driven by the AI risk and impact assessments under Clause 6.1.3, with exclusions justified in the SoA

ISO/IEC 42001:2023 Annex A is normative as a reference list, meaning the controls listed are the canonical set the organisation considers – not the set it must implement universally. Clause 6.1.3 drives selection from Annex A based on risk and impact assessment outcomes, with exclusions justified in the SoA. Option A is the textbook T-01 trap: 'normative' misread as 'mandatory in entirety'. Option B invents a quorum. Option C invents an approval gate the standard does not assign to the certification body.

Cited: ISO/IEC 42001:2023, Clause 6.1.3 (with Annex A)

Trap pattern: T-01 Annex A treated as mandatory

Q2 Correct answer: **A**. Map MLOps activities onto the AIMS AI lifecycle phases (design, development, deployment, operation); the sprint cadence is an internal practice, not an AIMS lifecycle phase

ISO/IEC 42001:2023 with ISO/IEC 23894:2023 guidance frames the AIMS AI lifecycle as design, development, deployment, and operation – a governance framework, not a delivery cadence. MLOps and sprint cycles are internal delivery practices that operate within those phases. Option A maps the team's practices onto the standard's phases. Option B is the T-02 trap: confusing delivery cadence with the AIMS lifecycle. Option D wrongly forces a waterfall model. Option C severs MLOps from governance, which is precisely what 42001 prevents.

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

Trap pattern: T-02 Software lifecycle vs AIMS lifecycle

Q3 Correct answer: **A**. Distinct documents with different scope, signatories, and review cadence; verify both are governed and that the Clause 5.2 AI policy commitments are explicitly addressed in at least one

ISO/IEC 42001:2023 Clause 5.2 requires the organisation to establish an AI policy at the management–system level (top–management–approved, addressing the AIMS commitments). Many organisations also publish AI ethics principles – these are distinct artefacts, with different scope and signatory. Option A keeps them distinct, which is correct. Option C collapses them – the T-03 trap. Option B bypasses the Clause 5.2 obligation. Option D removes the ethics document from scope when it sits within the AIMS.

Cited: ISO/IEC 42001:2023, Clause 5.2

Trap pattern: T-03 AI ethics policy vs AI management policy

Q4 Correct answer: **D**. Must be justified in the SoA with reasons referencing the risk and impact assessment outcomes

Annex A is a normative reference list, not a blanket–mandatory set. Clause 6.1.3 d) requires the SoA to justify both inclusions AND exclusions of Annex A controls, with reasons tied to the risk and impact outcomes. Option D captures this. Option A is the T-01 trap: 'normative' misread as 'all controls mandatory'. Option B invents a certification–body approval gate. Option C introduces a consumer–output carve–out the standard does not contain.

Cited: ISO/IEC 42001:2023, Clause 6.1.3 d)

Trap pattern: T-01 Annex A treated as mandatory

Q5 Correct answer: **B**. Operation phase – the AI system is deployed and monitored against operational criteria

ISO/IEC 42001:2023 with ISO/IEC 23894:2023 guidance frames operation as the phase covering deployed–system monitoring, performance evaluation, and the trigger logic for further action. Retraining triggered by performance breach is an operational control loop. Option D is the T–02 trap: placing an operational monitoring activity into development because retraining 'feels like' a development task. Option C misplaces threshold–evaluation into design (where thresholds are set, but not where they are evaluated in production). Option A treats deployment as a phase that includes ongoing monitoring, conflating deployment–the–event with operation–the–phase.

Cited: ISO/IEC 42001:2023 with ISO/IEC 23894:2023

Trap pattern: T–02 Software lifecycle vs AIMS lifecycle

Q6 Correct answer: **A**. Be appropriate to the organisation's purpose, provide a framework for AI objectives, include a commitment to satisfy applicable requirements, and include a commitment to continual improvement of the AIMS

Clause 5.2 lists specific content requirements for the AI policy (alignment to organisational purpose, framework for AI objectives, commitment to applicable requirements, continual improvement). Option A captures these. Option C describes ethical principles, which may sit in a separate AI Ethics document with different scope – the T–03 trap. Option D forces a combined document, which is not what the standard requires. Option B under–specifies the standard's actual requirements.

Cited: ISO/IEC 42001:2023, Clause 5.2

Trap pattern: T–03 AI ethics policy vs AI management policy

Q7 Correct answer: **B**. Annex A is normative (reference controls); Annex B is normative implementation guidance, and Clause B.1 specifies that B items do not require SoA justification

ISO/IEC 42001:2023 labels both Annex A and Annex B as normative. Annex A is the reference control catalogue (selection driven by risk and impact). Annex B is implementation guidance; Clause B.1 specifies that B items do not need SoA justification because they are guidance, not selectable controls. Option D is the common error – mislabelling B as informative because B.1 waives SoA justification (the waiver is a property of B.1's content, not a status change on Annex B). Option C wrongly demotes both. Option A wrongly extends SoA justification to B items.

Cited: ISO/IEC 42001:2023, Annexes A and B (with Clause B.1)

Trap pattern: T–01 Annex A treated as mandatory

Q8 Correct answer: **D**. Development phase – data collection is part of building the AI system

ISO/IEC 23894:2023 AI lifecycle guidance, referenced by ISO/IEC 42001:2023, places training–data collection in the development phase – the phase where the model and its training corpus are built. Option B confuses inference–time data flow with training–data collection. Option C confuses specifying data sources (a design activity) with collecting them (a development activity). Option A conflates loading–into–production with collection. The T–02 trap is candidates mapping AI lifecycle phases onto whatever delivery–pipeline phase happens to touch the data at a given moment, rather than to the governance lifecycle the standard defines.

Cited: ISO/IEC 42001:2023 with ISO/IEC 23894:2023

Trap pattern: T–02 Software lifecycle vs AIMS lifecycle

Q9 Correct answer: **D**. Acceptable – the two documents are distinct and the cross-reference makes the relationship explicit; the auditor should verify both are governed

Clause 5.2 requires the organisation to establish an AI policy at the management–system level. The standard does not prevent the organisation from having additional policies (such as a board-approved Ethics Policy) – it requires the AIMS policy to exist and to meet the Clause 5.2 content requirements. A cross-reference where the AI Policy commits the AIMS to operate in accordance with the Ethics Policy is good governance. Option C is the T-03 trap: collapsing two distinct documents because both contain 'AI'. Option B invents an approval-level rule. Option A removes a referenced document from audit scope.

Cited: ISO/IEC 42001:2023, Clause 5.2

Trap pattern: T-03 AI ethics policy vs AI management policy

Q10 Correct answer: **A**. Each Annex A control included in the SoA must actually be implemented under Clauses 8.1 and 8.3; 'selected but not implemented' is permitted only where it is carried in the management-approved AI risk treatment plan – otherwise it is a nonconformity

Clause 6.1.3 f) requires the SoA to contain the necessary controls with justification for inclusion and exclusion; it carries no implementation–status element. The implementation duty sits in Clause 8.1, which requires the controls determined according to 6.1.3 to be implemented, and Clause 8.3, which requires the AI risk treatment plan to be implemented and its effectiveness verified. A 'selected but not implemented' entry is therefore defensible only where it is carried in the management-approved AI risk treatment plan formulated under 6.1.3 g) and tracked to closure; otherwise it is a nonconformity, because the organisation has identified the control as necessary but is not operating it. Option B mis-frames Annex A as guidance (the inverse of the T-01 trap). Option C defers implementation to surveillance audits, but 8.1 and 8.3 carry no such deferral – effectiveness verification is a routine Stage 2 focus. Option D is too absolute – 6.1.3 g) and 8.3 contemplate treatment actions that are planned and tracked to closure.

Cited: ISO/IEC 42001:2023, Clauses 6.1.3 f)-g), 8.1 and 8.3

Trap pattern: T-01 Annex A treated as mandatory

Want the full bank?

Mindset Prep has 450+ clause-cited ISO 42001 Lead Implementer 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-li)

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