FREE 2026 TOOLKIT · 38 PAGES

ISO 42001 Compliance Toolkit · 2026

The ISO 42001 Compliance Toolkit

The honest 38-page checklist and cost guide. Clause-by-clause readiness, implementation cost ranges, 12-month timeline, and every Annex A control with audit-evidence guidance.

Checklist	Cost guide	Timeline	Annex A
57 readiness points	$25K–$150K range	12-month roadmap	All 37 controls

Inside the toolkit

- ✓ 57-point compliance readiness checklist
- ✓ Full cost breakdown ($25K–$150K range)
- ✓ 12-month implementation roadmap
- ✓ All 37 Annex A controls — with audit evidence
- ✓ Stage 1 + Stage 2 audit prep
- ✓ Integration tips with ISO 27001

EDITION	STANDARD	PAGES	CONTROLS
2026	**ISO/IEC 42001:2023**	**38**	**37 (Annex A)**

Contents

A 38-page walk through ISO/IEC 42001:2023 compliance — checklist, cost, timeline, controls, audit prep.

01 · The plain-English answer

What ISO 42001 compliance actually means.

Most explanations of ISO 42001 compliance get tangled in standards jargon. Here's the clean version — what compliance means, what it doesn't, and why it matters in 2026.

✓ What compliance IS

A documented, auditable AI Management System (AIMS) that meets the requirements of ISO/IEC 42001:2023 — policies, risk assessments, lifecycle controls, monitoring, and continual improvement.

Certifiable. An accredited audit body issues the certificate after Stage 1 and Stage 2 audits, with annual surveillance.

Holistic. Compliance touches policy, people, data, models, vendors, monitoring — every place AI lives in your organization.

✗ What compliance ISN'T

Not a one-time project. It is a management system — it operates continuously, with internal audits, management reviews, and corrective action.

Not the EU AI Act. The Act is binding regulation; ISO 42001 is voluntary and helps satisfy many — not all — Act obligations.

Not just paperwork. Auditors test that the documented controls are actually operating in production, not just written down.

The four compliance pillars

- **Governance** — board-approved AI policy, named accountabilities, ethics committee where needed.
- **Risk & impact** — repeatable AI risk methodology covering bias, drift, robustness, fairness, and societal impact.
- **Lifecycle controls** — design, development, validation, deployment, monitoring, and retirement of every AI system.
- **Data & vendors** — lineage and quality of training data, contractual controls on AI vendors and embedded LLMs.

02 · How to use this toolkit

Why this toolkit — and how to use it.

Most ISO 42001 cost answers ignore internal time, surveillance audits, and the real implementation drag. This toolkit gives the honest numbers and the steps to compress them.

Three ways to use this document

As a readiness assessment
Print the 57-point checklist. Walk it with your AI, security, and risk leads. Score each item Red / Amber / Green. The result is your initial gap analysis.

As a budget builder
Use the cost-line table on pages 5–6 to model Year-1 spend across audit, tooling, internal time, and training. Adjust by org size using page 7.

As a project plan
The 12-month roadmap on pages 17–21 is structured so you can lift each phase directly into your project tooling — milestones, deliverables, and owners.

A note on numbers. Cost ranges in this toolkit are 2026 market estimates aggregated from public consultancy rate cards, audit-body fee schedules, and industry surveys. Your actual cost depends on org size, AI footprint, ISO 27001 maturity, and audit body. Treat these as planning ranges, not quotes.

03 · Cost breakdown · 2026

The honest cost picture.

What ISO 42001 alignment actually costs in 2026 — across audit, certification, ongoing surveillance, and the biggest hidden line: internal team time.

$30K–$130K	9–14 mo	40–60%	$3K–$8K
Year-1 total — typical mid-size org	Kick-off to first certificate	Saved if ISO 27001 already in place	Annual surveillance cost

The five cost buckets

Gap analysis (external)
A pre-audit readiness assessment by a consultancy or accredited body. 3–4 weeks. Optional but recommended — surfaces issues before paid audit time begins.

Internal implementation time
By far the biggest line item, and the most under-budgeted. Documentation, control build-out, training, and tooling integration across 3–9 months.

Stage 1 + Stage 2 audit
External audits by an accredited certification body. Stage 1 reviews documentation; Stage 2 verifies operational evidence. Both required for certificate issuance.

Annual surveillance
After certification, the audit body runs lighter-touch surveillance audits annually for the 3-year cycle, with a re-certification audit at year 3.

Internal Lead Auditor training
Every certified AIMS needs internal auditors. Training costs vary widely by provider — GSDC at the affordable end, large global bodies at the premium end.

04 · Cost by line item

Cost — by line item.

Industry-aggregated 2026 ranges. Use these as planning ranges, not quotes. Internal time is the line most teams under-estimate.

Cost item	Range (USD)	When	Notes
Gap analysis External consultancy	**$8K – $20K**	3–4 weeks	Pre-audit readiness assessment
Internal implementation Team time across 3–9 months	**$15K – $80K**	3–9 months	Documentation, controls, training — biggest hidden cost
Tooling & GRC If extending existing stack	**$2K – $25K**	Months 3–6	Often free if ISO 27001 GRC already in place
Stage 1 + Stage 2 audit Accredited certification body	**$8K – $25K**	Months 10–12	Documentation + on-site / virtual audits
Surveillance audits Annual lighter-touch	**$3K – $8K/yr**	Years 2–3	Plus re-certification audit at year 3
Internal Lead Auditor Per person, varies by provider	**$400 – $3K**	Any time	Globally-recognized credentials at the lower end of this range exist
Year-1 total — typical mid-size org	**$30K – $130K**	Year 1	Wide range — depends on org maturity + AI footprint

05 · Cost by organization size

How cost scales with org size.

Three illustrative scenarios. The biggest variable is AI footprint, not headcount — an SMB with 20 production AI systems will cost more than an enterprise with three.

Small / SMB

Up to 500 employees · 1–3 AI systems

YEAR-1 RANGE
$25K – $50K

Often handled with one part-time AI compliance lead and an external consultant for gap analysis. Existing ISO 27001 controls cover most of Annex A. Internal audit cycles are lighter. Certificate scope is typically single business unit.

Mid-size

500–5,000 employees · moderate AI footprint

YEAR-1 RANGE
$50K – $130K

The baseline scenario this toolkit is calibrated for. Dedicated AI governance lead, cross-functional steering committee, formal AI ethics committee. Often layered onto an existing ISO 27001 programme. Scope is org-wide or major business unit.

Enterprise

5,000+ employees · heavy AI footprint

YEAR-1 RANGE
$100K – $400K+

Multiple AI products in production, regulated industry, multiple jurisdictions. Dedicated AI governance function, formal AI Risk Officer role, complex vendor estate. Audits longer, surveillance higher, internal team larger. Scope is typically org-wide.

06 · Cost-reduction levers

Cost-reduction levers that actually work.

Five proven levers to compress Year-1 ISO 42001 cost without compromising the certificate.

Lean on existing ISO 27001
Annex A overlaps significantly with ISO 27001 Annex A. Most mature ISMS programmes can extend existing GRC tooling, audit calendar, and policy structure — saving 40–60% on Year-1 implementation cost.

Train internal auditors early
An internal Lead Auditor pays for themselves in the first audit cycle. The cost delta between an internal hour and an external consultancy hour is large; compound it across a 12-month build.

Narrow scope first
Start with one business unit or product line. Demonstrate the model works. Expand scope at year-2 recertification. Easier path than retro-fitting a broad scope.

Use templates, not bespoke documents
Policies, risk methodology, model cards, vendor questionnaires — all available as templates. Resist the urge to bespoke every artifact; auditors do not care about branding.

Pre-empt findings with a gap analysis
A $10K gap analysis before Stage 1 is cheaper than a failed Stage 2 audit. Surface and remediate findings before paid audit clock starts.

07 · The 57-point checklist

The compliance readiness checklist.

A 57-point self-assessment organised by ISO 42001 clause and Annex A category. Use Red / Amber / Green scoring on each item — the result is your initial gap analysis.

12	**Governance** Policy, accountabilities, leadership
14	**Risk & impact** Methodology, assessments, treatment
15	**Lifecycle** Design, build, validate, deploy, retire
8	**Data & vendors** Lineage, quality, supplier controls
8	**Audit & improve** Internal audit, monitoring, corrective action

How to score it

- **Green** — control is documented, deployed, and evidenced. Auditors would find no gap.
- **Amber** — control is partly in place but missing documentation, evidence, or coverage.
- **Red** — control is absent or fundamentally non-conforming. Will block certification.

Recommended approach. Walk the checklist with three people: an AI/ML lead, a security/risk lead, and a compliance lead. The disagreements between them are where the real gaps usually hide.

08 · Checklist · items 1–6

Checklist items 1–6 · Governance.

The highest-impact items most organizations miss in their first gap analysis. Walk each line — print this page if useful.

- [ ] 01 **AI scope & system inventory documented**
  Every AI use case across the org logged with risk classification.
- [ ] 02 **AI policy approved at executive level**
  Board-signed AI policy — not just an IT-level document.
- [ ] 03 **AI risk assessment methodology defined**
  Repeatable framework covering bias, drift, robustness, fairness.
- [ ] 04 **AI roles & responsibilities (RACI)**
  Who owns model risk, who approves deployment, who can halt production.
- [ ] 05 **Third-party AI vendor controls**
  Procurement gates, vendor questionnaires, contracts.
- [ ] 06 **Data governance for AI training**
  Lineage, consent, retention, quality controls on training data.

09 · Checklist · items 7–12

Checklist items 7–12 · Lifecycle & audit.

Items 7–12 cover model documentation, validation, monitoring, and the internal audit programme.

- [ ] 07 **Model documentation (Model Cards)**
  Standardized doc of purpose, training, performance, limits.
- [ ] 08 **Pre-deployment testing & validation**
  Accuracy, fairness, robustness tests with documented results.
- [ ] 09 **Production monitoring — drift & performance**
  Continuous monitoring, alerting, retraining triggers.
- [ ] 10 **Incident response for AI failures**
  What happens when a model behaves unexpectedly in production.
- [ ] 11 **Explainability evidence per AI use case**
  How decisions are explained to users, regulators, auditors.
- [ ] 12 **Internal audit programme for AI**
  Scheduled audits, qualified auditors, management review cadence.

10 · Remaining checklist

The remaining 45 items — by category.

Items 13–57 are organized by Annex A category. Each item maps to a specific control and includes audit-evidence guidance.

Items 13–20 · AI policies & internal organization
Policy review cadence, exception handling, ethics committee charter, reporting lines, segregation of duties for development vs deployment.

Items 21–28 · Resources & impact assessment
Compute & data resource governance, AI impact assessment methodology, stakeholder identification, documented impact mitigation.

Items 29–38 · AI system lifecycle
Design objectives, requirements specification, verification, validation, deployment approval gates, event logging, retirement procedures.

Items 39–46 · Data for AI systems
Data acquisition, quality controls, provenance, preparation, bias testing, retention, privacy by design integration.

Items 47–52 · Information & use of AI
User-facing information, regulator reporting, incident communication, intended-use boundaries, deviation handling.

Items 53–57 · Third-party & monitoring
Supplier AI controls, embedded LLM governance, customer-facing AI controls, ongoing monitoring evidence, surveillance audit prep.

11 · Annex A · Control categories

The 37 Annex A controls — grouped.

ISO 42001's Annex A defines 37 reference controls grouped into 9 categories. Here's the structure most consultants don't explain clearly.

A.2
Policies related to AI
Approval, communication, review, exceptions of AI policies.

A.3
Internal organization
Roles, responsibilities, reporting lines, AI ethics committee.

A.4
Resources for AI systems
Data, tooling, computing, system documentation.

A.5
AI system impact assessment
Identify, document, and mitigate stakeholder impacts.

A.6
AI system lifecycle
Objectives, design, development, validation, deployment, retirement.

A.7
Data for AI systems
Data sources, quality, lineage, integrity, privacy, retention.

A.8
Information for interested parties
Communication to users, regulators, affected parties.

A.9
Use of AI systems
Intended use, monitoring during operation, deviation handling.

A.10
Third-party relationships
Vendor AI controls, embedded LLMs, supply chain.

12 · Annex A · first half

Annex A.2–A.5 in detail.

Governance fundamentals — policy, people, resources, and impact assessment.

A.2 Policies related to AI

Establish and maintain AI policies addressing acceptable use, ethics, transparency, human oversight, and review cadence. Sub-controls cover approval (executive sign-off), communication across the org, and periodic review.

A.3 Internal organization

Define AI-related roles and responsibilities. Establish accountability for AI risk owners, AI system owners, and an AI ethics or governance committee where appropriate. Ensure segregation of duties for AI development, testing, and deployment.

A.4 Resources for AI systems

Identify and document resources across the AI lifecycle — compute infrastructure, training data, tooling, and human expertise. Each AI system must have documented inventory, owners, and metadata sufficient for audit.

A.5 AI system impact assessment

Establish a process to identify, analyse, document, and mitigate impacts of AI systems on individuals, groups, and society — including fairness, safety, privacy, and human autonomy. Update assessments when AI systems materially change.

13 · Annex A · second half

Annex A.6–A.10 in detail.

Operational mechanics — lifecycle, data, communications, end-use, third parties.

A.6 AI system lifecycle

Govern each phase: design, development, validation, deployment, operation, and retirement. Include controls for objective-setting, requirements specification, testing, deployment approval, monitoring, and decommissioning.

A.7 Data for AI systems

Manage data quality, lineage, provenance, retention, and privacy across the lifecycle. Document training, validation, and test datasets; address bias; ensure data minimization and lawful basis of processing.

A.8 Information for interested parties

Provide appropriate information to users, regulators, affected end-users, and the public — system purpose, capabilities, limitations, contact points for queries, and channels for raising concerns.

A.9 Use of AI systems

Define intended use boundaries, monitor for deviation, and govern human-AI interaction. Establish processes for handling out-of-distribution inputs, automation bias, and override authority.

A.10 Third-party & customer relationships

Govern AI obtained from suppliers (LLM vendors, model marketplaces, embedded AI features) and AI provided to customers. Include contractual controls, shared-responsibility models, and supplier review cadence.

14 · Audit-evidence guide

Audit-evidence quick reference.

What auditors look for, per control category. Have these artifacts ready before Stage 1 begins.

Control area	Audit evidence the auditor expects
A.2 Policies	Approved AI policy with signatures and dates; policy review minutes; communication evidence (email, intranet, training records)
A.3 Internal org	RACI matrix; AI ethics committee charter and meeting minutes; named AI system owners in inventory
A.4 Resources	AI system inventory; compute / data resource registers; documented tooling stack; training records for AI practitioners
A.5 Impact assessment	Completed impact assessments per AI system; mitigation actions; stakeholder consultation records
A.6 Lifecycle	Design docs; validation test results; deployment approval records; event logs; retirement procedures
A.7 Data	Data lineage diagrams; dataset cards; quality test reports; retention schedules; lawful-basis records
A.8 Information	User-facing AI disclosures; regulator notifications; incident communications; feedback channels
A.9 Use of AI	Intended-use documentation; production monitoring dashboards; deviation handling logs; override evidence
A.10 Third party	Vendor assessments; contracts with AI-specific clauses; shared-responsibility documentation; supplier reviews

15 · Implementation roadmap

The realistic 12-month ISO 42001 roadmap.

Mid-sized organizations (500–5,000 employees, moderate AI footprint) reach Stage 2 audit readiness at month 10–12. Organizations with mature ISO 27001 programmes compress this by 30–50%.

Months 1–2
Foundation

Gap analysis & scope. External or internal gap assessment. Decide certification scope (org-wide vs business unit). Form steering committee. Train internal Lead Auditor.

Months 3–6
Build

Documentation & controls. Build AI policy, risk methodology, control library across all 37 Annex A controls. Data governance, model documentation, vendor controls. Biggest investment phase.

Months 7–9
Operate

Internal audits & refinement. Run internal audits. Identify nonconformities. Remediate. Train staff on AI policies. Operate the management system for at least 90 days before external audit.

Months 10–12
Certify

Stage 1 + Stage 2 audit. External audit body conducts Stage 1 (documentation review) then Stage 2 (operational evidence). Address findings. Receive certificate. Surveillance audits annually.

16 · Phase 1 · Months 1–2

Foundation — gap analysis & scope.

The two-month foundation phase sets every downstream decision. Get scope wrong here and you'll spend months recovering.

Key deliverables

AI system inventory
Every AI use case across the org logged with risk classification (high / medium / low) and a named owner. This becomes the master register the rest of the programme operates against.

Certification scope statement
Which business units, AI systems, and jurisdictions are inside the scope. Defensible exclusions documented. Get this approved by the executive sponsor before any documentation begins.

Gap analysis report
External or internal assessment against the 57-point checklist. Each gap labelled Red / Amber / Green with remediation owner and target date.

Steering committee charter
Cross-functional committee (AI, security, risk, legal, product) with named members, meeting cadence, decision rights, escalation path.

Internal Lead Auditor named
At least one named internal auditor with appropriate training started. Internal auditors are required by the standard — don't leave this to month 7.

Common pitfall. Teams skip the gap analysis to "save money," then discover fundamental issues during Stage 1 — when remediation costs more than the gap analysis would have. The $10K gap analysis is the highest-ROI line item in the entire programme.

17 · Phase 2 · Months 3–6

Build — documentation & controls.

The biggest investment phase. Documentation, control build-out, training. This is where most of the internal-time budget goes.

AI policy & supporting policies
Executive-signed AI policy. Supporting policies for acceptable use, third-party AI, model documentation standards, incident response. Communicated and acknowledged by relevant staff.

AI risk methodology
Repeatable framework covering bias, drift, robustness, fairness, safety, and societal impact. Tie back to enterprise risk appetite. Document risk criteria, scoring, and treatment options.

Annex A control library
Document and deploy controls covering all 37 Annex A reference controls. Statement of Applicability (SoA) finalised. Each control has an owner, evidence source, and review cadence.

Data governance for AI
Lineage, consent, retention, quality controls on training data. Bias testing protocols agreed. Dataset cards completed for each material training set.

Model documentation standard
Model Card template adopted org-wide. Each in-scope AI system has a current Model Card with purpose, training data, performance metrics, and known limitations.

Training rollout
Targeted training for AI builders, reviewers, operators, vendors. Records retained as evidence of competence.

18 · Phase 3 · Months 7–9

Operate — internal audits & refinement.

The management system must operate in production for at least 90 days before external audit. This is where ceremonial AIMS programmes get found out.

Internal audit programme runs
Scheduled internal audits against the AIMS. Audit findings logged, remediated, and tracked to closure. Auditor independence maintained — not the same people who built the controls.

Management review held
Formal management review with executive sponsor. Reviews AIMS performance, audit findings, risk landscape, resource adequacy, opportunities for improvement. Documented minutes retained.

Nonconformities & corrective action
Each nonconformity has a root-cause analysis, corrective action, effectiveness review. The corrective action register itself becomes audit evidence.

90 days of live operation
AIMS demonstrably operating across in-scope AI systems for the 90-day window before Stage 2. Monitoring data, incident logs, decision records all accumulate.

Pre-Stage-1 readiness review
Internal walk-through of all documentation against ISO 42001 clauses and SoA. Last chance to surface issues before external auditors arrive.

19 · Phase 4 · Months 10–12

Certify — Stage 1 and Stage 2.

The certification audits happen in two stages, typically 4–8 weeks apart. Findings from Stage 1 are remediated before Stage 2 begins.

Stage 1 — Documentation review
Audit body reviews your documentation against ISO/IEC 42001 requirements. Confirms scope appropriateness, policy alignment, SoA completeness, internal audit programme adequacy. Findings issued.

Stage 1 remediation
Address findings between Stage 1 and Stage 2. Major findings may delay Stage 2 until remediated. Minor findings can sometimes be resolved during Stage 2.

Stage 2 — Operational evidence
Audit body verifies that documented controls are operating in production. Interviews staff, reviews records, samples AI systems, examines monitoring evidence. Typically 3–10 days on-site or virtual.

Certificate decision
Audit body reviews findings, issues recommendation. Certification body's independent reviewer confirms. Certificate issued with 3-year validity, subject to annual surveillance.

Year-1 surveillance audit
Lighter-touch audit ~12 months after certification. Verifies the AIMS continues to operate. Sample of controls reviewed. Findings handled in same way as main audit.

20 · Stage 1 audit prep

Stage 1 — what to prepare.

Stage 1 focuses on whether your documentation says what ISO 42001 requires. Bring complete, signed, current documents.

The Stage 1 evidence pack

- ✓ AI policy — current version, with executive signatures and dates
- ✓ Scope statement — defensible boundaries, named systems and exclusions
- ✓ Statement of Applicability (SoA) — all 37 Annex A controls addressed
- ✓ AI system inventory — every in-scope AI system with owner and risk class
- ✓ AI risk assessment methodology — with worked examples per risk class
- ✓ Internal audit programme — schedule, auditor qualifications, sample reports
- ✓ Management review minutes — at least one full review cycle
- ✓ AI policies cascade — acceptable use, third-party AI, model documentation

Common Stage 1 findings

Scope too broad or too narrow
Scope claims more AI than the org actually controls, or excludes high-impact systems.

SoA inconsistencies
Controls marked applicable in SoA aren't reflected in policy; controls excluded lack justification.

Risk methodology too generic
Generic infosec risk methodology copied across — doesn't address AI-specific risks like drift, bias, automation bias.

Missing management review
Management review not yet held, or held without proper inputs (audit results, risk landscape, resource review).

21 · Stage 2 audit prep

Stage 2 — what to prepare.

Stage 2 verifies the documented controls are actually operating. Auditors interview staff and sample AI systems. Prepare for live, not paper.

Stage 2 operational evidence

- ✓ AI system records — Model Cards, validation results, deployment approvals
- ✓ Monitoring evidence — production dashboards, alerts, drift detection logs
- ✓ Incident response records — at least one incident handled end-to-end
- ✓ Internal audit reports — recent cycle with findings and corrective actions
- ✓ Training records — competence evidence for AI practitioners
- ✓ Vendor reviews — supplier AI assessment evidence for material vendors
- ✓ Data governance evidence — lineage, quality reports, retention enforcement
- ✓ Impact assessments — completed and updated where systems materially changed

How auditors sample

Risk-weighted
Higher-risk AI systems sampled more heavily. Customer-facing or decision-impacting AI gets priority.

Across the lifecycle
Auditors will sample at design, deployment, and monitoring stages — not just one phase. Be prepared with end-to-end evidence.

Interview-driven
Document review is paired with staff interviews. Ensure named owners can explain the controls they own — auditor's red flag is 'I don't know, ask compliance.'

Trace-back exercises
Auditor picks one AI system at random and traces: scope → impact assessment → design → validation → deployment → monitoring. Have the chain ready.

22 · Internal audit programme

Internal audit — what 'good' looks like.

Internal audits are required by ISO 42001. Done well, they surface issues before external auditors do; done badly, they create their own findings.

Auditor independence
Internal auditors must not audit work they themselves performed. Maintain independence between control owners and auditors — even informally.

Auditor competence
Internal auditors must be qualified. ISO 19011 sets out competence requirements. Most orgs train auditors with a recognized Lead Auditor credential.

Audit programme planning
Risk-based plan covering all clauses and Annex A controls across a defined cycle (typically annual). Higher-risk areas audited more frequently.

Audit-cycle reporting
Findings logged, classified (major / minor / observation), tracked to closure. Trends reported up to management review.

Continual improvement loop
Findings feed risk register and improvement backlog. Repeat findings flag systemic issues — investigate root cause, not just surface symptom.

23 · ISO 27001 integration

Integrating with ISO 27001.

Most enterprises with ISO 27001 in place can extend rather than rebuild. The two standards share the High-Level Structure (HLS) and significant control overlap.

What you can reuse from ISO 27001

- ✓ **Clauses 4, 5, 9, 10** — Context, Leadership, Performance, Improvement are nearly identical between the two standards.
- ✓ **Information classification** (A.5 of 27001) — extends naturally to AI training data.
- ✓ **Supplier security** (A.5.19–A.5.23 of 27001) — augment with AI-specific clauses for A.10 of 42001.
- ✓ **Asset management** (A.5 of 27001) — extend asset inventory with AI systems as a class.
- ✓ **Incident management** (A.5.24–A.5.26 of 27001) — extend playbooks with AI-specific scenarios.
- ✓ **SDLC controls** (A.8.25–A.8.29 of 27001) — overlap heavily with A.6 AI lifecycle.

What you need to add

- + **AI policy** — distinct from infosec policy; covers ethics, fairness, transparency.
- + **AI risk methodology** — covers bias, drift, robustness — distinct from infosec risk.
- + **AI impact assessment** — A.5 has no direct equivalent in ISO 27001.
- + **Model lifecycle gates** — A.6 introduces controls infosec SDLC doesn't fully cover.
- + **Model documentation** — Model Cards have no 27001 analogue; net-new artifact.
- + **AI-specific vendor questionnaires** — extend, don't replace, infosec vendor reviews.

24 · Common pitfalls

Common ISO 42001 compliance pitfalls.

Patterns that show up repeatedly in failed Stage 1 audits and weak Stage 2 outcomes. Pre-empt them in build, not in audit.

Shadow-AI surprises
Scope says 'all AI in the org' — but the audit discovers undocumented AI tools (LLM browser extensions, embedded vendor AI, marketing automation). Fix: explicit scope statement, shadow-AI discovery sweep before Stage 1.

Generic risk methodology
Infosec risk methodology copied unchanged. Auditors expect AI-specific risk types: drift, fairness degradation, automation bias, training-data poisoning. Fix: extend methodology with AI risk categories and worked examples.

Missing impact assessments
A.5 impact assessments present only for the marquee AI use case. Auditors expect an impact assessment for every in-scope system at proportionate depth. Fix: tiered impact assessment process — short form for low-risk, full form for high-risk.

Model Cards without owners
Model Cards filled in by data scientists, then abandoned. Auditor asks 'who reviews this Card when the model changes?' and the answer is silence. Fix: named Model Card owner per system, with review trigger tied to retraining.

Vendor questionnaires without follow-through
Vendor sent AI questionnaire; response received; no review documented; no contract clauses updated. Fix: vendor review evidence chain — questionnaire → review → decision → contract.

Internal audit done by control owners
Internal audit performed by the same team that built the controls — independence compromised. Fix: rotate auditors or use a separate team; even a small org can achieve separation.

25 · AI risk methodology

AI risk — methodology essentials.

What auditors expect to see in your AI risk methodology — beyond a copy of the infosec methodology with 'AI' added in the margin.

AI risk categories you must address

Fairness & bias
Differential outcomes across protected groups. Includes both training-data bias and deployment-context bias.

Drift
Statistical drift in input data or performance degradation over time. Triggers for retraining or pause.

Robustness
Behaviour on out-of-distribution inputs, adversarial inputs, edge cases. Includes prompt injection for LLMs.

Explainability
Ability to explain individual decisions to affected users, regulators, auditors. Required varies by use case.

Privacy & data leakage
Models leaking training data; inference attacks; PII exposure in outputs. Includes LLM memorization risks.

Misuse & abuse
Use of AI outside intended purpose; jailbreaking; automation bias by operators.

Societal impact
Effect on individuals beyond the immediate user — affected end-users, communities, labour markets.

26 · Documentation pack

The documentation pack — what you need.

The minimum documented artifacts required for ISO 42001 certification. Templates for most of these are in the GSDC toolkit appendix.

Tier 1 — Top-level documents
AI policy · AIMS scope statement · Statement of Applicability · Risk methodology · Internal audit programme

Tier 2 — Process documents
AI lifecycle procedure · Impact assessment procedure · Vendor AI procedure · Data governance procedure · Incident response playbook

Tier 3 — Per-system artifacts
Model Card · Impact assessment record · Validation report · Monitoring dashboard · Deployment approval record

Tier 4 — Records
Internal audit reports · Management review minutes · Training records · Vendor reviews · Incident reports · Corrective action log

Templates appendix. The GSDC Lead Auditor toolkit (module 14) ships with ready-to-use templates for every Tier 1 and Tier 2 artifact above, plus a Model Card template, Top-100 audit-finding playbook, and AI audit checklist.

27 · ISO 27001 ↔ 42001 mapping

ISO 27001 ↔ 42001 control mapping.

Where existing ISO 27001 controls give you a head start on ISO 42001 — and where you'll need to build something net-new.

ISO 42001 control	Maps to ISO 27001	Reuse level
Clauses 4, 5, 9, 10	Clauses 4, 5, 9, 10	**High** — near-identical
A.2 AI policies	A.5.1 Policies for infosec	**Medium** — same structure, new content
A.3 Internal organization	A.5.2–A.5.4 Roles, segregation	**Medium** — extend RACI
A.4 Resources for AI	A.5.9–A.5.10 Asset mgmt	**Medium** — extend asset registers
A.5 Impact assessment	None — net-new	**Low** — build from scratch
A.6 AI lifecycle	A.8.25–A.8.29 SDLC	**Medium** — extend SDLC gates
A.7 Data for AI	A.5.12–A.5.14 Classification	**Medium** — extend data controls
A.8 Info for parties	A.5.34 Privacy	**Low** — mostly new
A.9 Use of AI	None — net-new	**Low** — build from scratch
A.10 Third party	A.5.19–A.5.23 Supplier	**Medium** — extend supplier controls

28 · A.10 · Third-party AI

Vendor & third-party AI controls.

A.10 is where most orgs have the biggest hidden gap. Every embedded LLM, every AI-powered SaaS, every model marketplace introduces obligations.

Vendor AI inventory
Maintain a register of every supplier providing AI capability or AI components — including embedded LLMs, AI features in SaaS tools, model APIs, and AI-augmented consulting services.

AI-specific vendor questionnaire
Augment standard infosec vendor questionnaires with AI-specific items: training data sources, model evaluation, bias testing, drift monitoring, incident response, AI sub-processors, model update notification.

Contractual clauses
AI-specific contract clauses: notification of material model changes, audit rights, training-data warranties, performance metrics, indemnities for AI-caused harm, data residency and processing constraints.

Shared-responsibility model
Document who is responsible for which controls in vendor-provided AI. The customer is rarely responsible for the model itself — but is responsible for use, monitoring, human oversight, and impact assessment.

Ongoing vendor review
Annual review of material AI vendors. Re-run questionnaire. Update risk rating. Track open issues to closure. Maintain evidence chain for auditors.

29 · Model Card template

Model Card — minimum content.

A standardised one-page artifact per AI system, kept current and reviewed at each material change. The audit-evidence backbone for A.6 and A.9.

Identification
System name · version · owner · risk class · in-scope status · last review date

Purpose & intended use
What the system is for · what decisions it informs · who the users are · what's explicitly out of scope

Training data
Sources · time range · size · preprocessing steps · known biases · lawful basis · retention

Performance
Headline metrics (accuracy, precision, recall, fairness measures) · benchmark dataset · evaluation date · acceptance thresholds

Limitations
Known failure modes · out-of-distribution behaviour · adversarial robustness · groups where performance degrades

Monitoring
Drift detection method · alerting thresholds · review cadence · retraining triggers

Human oversight
Where humans review · override authority · escalation path · audit-trail location

Change log
Material changes since previous version · re-validation evidence · approver and date for each change

30 · AI incident response

AI incident response playbook.

An infosec incident playbook is necessary but not sufficient for AI. AI failures look different — and bad ones can run silently for weeks.

Detection
Sources: drift alerts, user complaints, internal staff escalation, regulator inquiry, press coverage. Define minimum detection mechanisms per risk class.

Containment
Options: disable feature, fall back to non-AI baseline, revert to previous model version, throttle, route to human review. Document the decision tree.

Investigation
Root-cause analysis with AI-specific tooling: explainability, dataset analysis, input drift analysis, prompt logs. Preserve evidence — don't redeploy before RCA.

Communication
Internal: executives, affected business units. External: affected users, regulators, customers. Pre-agreed templates and decision authority for each.

Recovery & remediation
Fix root cause, re-validate, redeploy with extra monitoring. Update Model Card. Update risk register. Feed lessons into training data and design controls.

Post-incident review
Formal review within 30 days. Findings feed corrective action log. Material incidents reported to management review.

31 · FAQ · cost & integration

Cost, time, and integration — answered.

Q. How much does ISO 42001 certification really cost?
For organizations: $30K–$130K Year-1, depending on org size and AI footprint. Organizations with mature ISO 27001 programmes typically save 40–60% on Year-1 implementation cost. Internal Lead Auditor training varies widely by provider — from a few hundred dollars at the accessible end to several thousand at the premium end.

Q. How long does ISO 42001 certification take?
For organizations: 9–14 months from kick-off to certificate, with mid-sized organizations averaging 12 months. ISO 27001-mature orgs can compress this by 30–50%. Lead Auditor training for an individual takes 30–60 days at roughly 6 hrs/week.

Q. Can ISO 42001 integrate with our existing ISO 27001?
Yes — and it should. Annex A controls overlap significantly with ISO 27001 Annex A. Most mature ISMS programmes can extend their existing GRC tooling, audit calendar, and policy structure rather than creating parallel infrastructure.

Q. Do we need an external auditor or internal?
Both. To receive the ISO/IEC 42001 certificate, you need an external accredited audit body. But you also need internal auditors for your management system to function — internal audits are required by the standard itself. Most organizations train 2–3 internal auditors to manage ongoing compliance.

32 · FAQ · auditors & credentials

Auditors, credentials, scope — answered.

Q. Is the GSDC Lead Auditor credential accepted by external audit bodies?
Yes — the GSDC certification is ISO/IEC 17024-aligned and recognized for internal audit roles in 100+ countries. For becoming an external lead auditor employed by an accredited certification body, additional CB-specific certification may be required (varies by body).

Q. Should we scope org-wide or by business unit?
Start narrow. A single business unit or product line lets you prove the model works before expanding scope at year-2 recertification. Org-wide first-time scope is a common failure pattern.

Q. Can we use the same audit body for ISO 27001 and ISO 42001?
Yes, and most large bodies (BSI, BV, DNV, LRQA, SGS, TÜV) offer integrated audits. This can compress audit time and cost, especially if scope boundaries align.

Q. What about the EU AI Act?
ISO 42001 is voluntary; the EU AI Act is binding regulation. They overlap heavily — ISO 42001 helps satisfy many EU AI Act obligations, but they're not the same. Most regulated organizations need both: the Act for legal compliance, ISO 42001 for management-system rigour and procurement evidence.

Q. How often must we recertify?
ISO 42001 certificates have a 3-year validity. Year 1 and Year 2 have surveillance audits (lighter touch). Year 3 has a re-certification audit (deeper, similar in scope to Stage 2).

33 · Templates & artifacts

Templates & artifacts index.

What you need to build vs what you can borrow. Don't reinvent — auditors don't care about brand styling, they care about content.

Artifact	Purpose	Source
AI Policy	Top-level governance document	Template + customise
Scope Statement	Defines AIMS boundary	Custom — bespoke per org
Statement of Applicability	Annex A control inclusion / exclusion	Template + customise
Risk Methodology	AI risk identification & treatment	Template + customise
Impact Assessment	Per-AI-system societal impact	Template + customise
Model Card	Per-system technical documentation	Template + standardise
Vendor AI Questionnaire	Third-party assessment	Template + customise
Incident Playbook	AI failure response	Template + customise
Internal Audit Programme	Schedule, scope, auditor list	Custom — per org cadence
Management Review Template	Inputs, outputs, minutes	Template + customise

34 · Quick-reference cheat sheet

ISO 42001 cheat sheet.

One page to print. Pin to the wall behind the implementation lead's desk.

STANDARD ISO/IEC 42001:2023 — published Dec 2023	**STRUCTURE** 10 clauses (HLS) + Annex A (37 controls in 9 categories)
ADOPTION Voluntary, certifiable, internationally recognised	**TIMELINE** 9–14 months kick-off → certificate
VALIDITY 3-year certificate cycle	**SURVEILLANCE** Annual (Year 1 + Year 2)
RECERTIFICATION Year 3 — deeper audit	**AUDITS** Stage 1 (docs) + Stage 2 (operations)
INTERNAL AUDIT Required by the standard	**ISO 27001 REUSE** 40–60% control overlap
BIGGEST HIDDEN COST Internal time (3–9 months)	**BIGGEST PRE-AUDIT WIN** $10K gap analysis

Key dates rhythm

- **Month 1** — Steering committee stood up, scope agreed, gap analysis kicked off
- **Month 6** — All 37 Annex A controls have a documented control and owner
- **Month 7** — First internal audit cycle starts
- **Month 9** — Management review held; 90-day operate window starts
- **Month 11** — Stage 1 audit
- **Month 12** — Stage 2 audit; certificate issued

35 · Glossary

Key terms — quick definitions.

The terms you'll hear most often in ISO 42001 implementation.

AIMS	AI Management System. The full set of policies, processes, controls, and records that satisfy ISO 42001.
Annex A	The 37 reference controls grouped into 9 categories. Implementers cross-reference these against their Statement of Applicability.
SoA	Statement of Applicability. The document that lists which Annex A controls the organization includes (with justification) and excludes (with rationale).
Stage 1 audit	Documentation review by the external certification body. Confirms readiness for Stage 2.
Stage 2 audit	Operational evidence audit. Verifies that documented controls are operating in production. Required for certificate issuance.
Surveillance audit	Lighter annual audit between certification audits, confirming the AIMS continues to operate.
Nonconformity	A failure to meet an ISO 42001 requirement. Classed as major (blocks certification) or minor (remediated in agreed timeframe).
Internal audit	Audit of your own AIMS by qualified internal auditors. Required by the standard.
Impact assessment	A.5 control — assessment of how an AI system affects individuals, groups, and society. Distinct from privacy impact assessment.
Model Card	Standardised one-page artifact per AI system documenting purpose, training data, performance, limitations.
Drift	Statistical change in input data distribution or model performance over time. A core AI-specific risk category.
Lead Auditor	Qualified individual who can lead audits against ISO 42001 — either internal audits or, with additional CB credentialing, external audits.

36 · Next steps

Where to from here.

Three concrete moves you can make in the next 30 days, regardless of where you are on the ISO 42001 journey.

TODAY
Print the 57-point checklist. Walk it with one AI, one security, and one risk person in a 90-minute session. Score Red / Amber / Green per item. The disagreement patterns are your highest-priority gaps.

THIS MONTH
Build the AI system inventory. Every AI use case across the org with risk class and owner. This single document drives every downstream decision in the programme.

THIS QUARTER
Train at least one internal Lead Auditor. Whether the broader programme starts in Q1 or Q3, the certification trains the AI compliance mindset across the team and pays for itself in the first audit cycle.

In one paragraph

ISO/IEC 42001:2023 is a certifiable international standard for managing AI responsibly. For a mid-sized organization, Year-1 cost runs $30K–$130K and the path takes 9–14 months. ISO 27001-mature organizations compress both. The single biggest cost reduction is starting with a gap analysis; the single biggest schedule reduction is training internal auditors early. This toolkit gives you the checklist, the cost ranges, the timeline, the controls, and the audit-evidence guidance to do all three.

This toolkit is informational. For the normative wording of ISO/IEC 42001:2023, refer to the official ISO publication. All CTAs in this PDF link to gsdcouncil.org/iso-42001-compliance.