# ISO/IEC 42001:2023 Statement of Applicability

All 38 Annex A controls — worked for a 100-staff Australian B2B SaaS

*Every control is applicable here, and that is the finding — this organisation both develops and procures AI, and ISO 42001 is far harder to scope down than ISO 27001. The final column shows which controls an organisation that only procures AI could argue down.*

| CONTROL | THEME | CONTROL TITLE | APPLICABLE | JUSTIFICATION | IMPLEMENTATION STATUS | AI SYSTEMS IN SCOPE | IF YOU ONLY PROCURE AI |
|---|---|---|---|---|---|---|---|
| A.2.2 | A.2 Policies related to AI | AI policy | Yes | The organisation develops and uses AI systems, so a documented AI policy approved by top management is the foundation of the AIMS. | Implemented — approved by the exec 2026-02-14, published internally and on the trust page | All | Applies to any organisation using AI, however it was obtained |
| A.2.3 | A.2 Policies related to AI | Alignment with other organisational policies | Yes | An ISO 27001 ISMS, a privacy policy and an acceptable-use policy already exist; the AI policy must not contradict them. | Implemented — cross-referenced against the ISMS policy set at the 2026-02 review | All | Applies to any organisation using AI, however it was obtained |
| A.2.4 | A.2 Policies related to AI | Review of the AI policy | Yes | Clause 9.3 requires management review; the policy is an input to it and must be reviewed at planned intervals. | Implemented — annual review scheduled, plus on significant change | All | Applies to any organisation using AI, however it was obtained |
| A.3.2 | A.3 Internal organisation | AI roles and responsibilities | Yes | Accountability must be assigned to named individuals; guardrail 1 of the Australian Voluntary AI Safety Standard asks for the same thing. | Implemented — CTO accountable, system owners named per row of the AI inventory | All | Applies to any organisation using AI, however it was obtained |
| A.3.3 | A.3 Internal organisation | Reporting of concerns | Yes | Staff and affected people need a route to raise AI concerns without going through the system owner. | Implemented — routed through the existing whistleblower channel with an AI category added | All | Applies to any organisation using AI, however it was obtained |
| A.4.2 | A.4 Resources for AI systems | Resource documentation | Yes | The resources supporting each AI system must be identified and documented to make the AIMS auditable. | Implemented — captured in the AI system inventory | All | Applies to any organisation using AI, however it was obtained |
| A.4.3 | A.4 Resources for AI systems | Data resources | Yes | Both developed systems train on internal data; both procured systems process customer data at inference. | Implemented — data sources recorded per system with classification | All | Applies to any organisation using AI, however it was obtained |
| A.4.4 | A.4 Resources for AI systems | Tooling resources | Yes | Model training, evaluation and monitoring tooling is in scope because the organisation builds two systems. | Implemented — tooling inventory maintained alongside the ISMS asset register | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.4.5 | A.4 Resources for AI systems | System and computing resources | Yes | Compute used for training and inference is a resource the AIMS must account for. | Implemented — cloud accounts tagged by AI system for cost and access traceability | All | Applies to any organisation using AI, however it was obtained |
| A.4.6 | A.4 Resources for AI systems | Human resources | Yes | Competence for AI-related roles must be determined and evidenced under clause 7.2. | Implemented — competence matrix covers system owners and the review roles in A.9 | All | Applies to any organisation using AI, however it was obtained |
| A.5.2 | A.5 Assessing impacts of AI systems | AI system impact assessment process | Yes | Clause 8.4 requires it. This is the requirement with no ISO 27001 equivalent and it cannot be excluded by any organisation using AI. | Implemented — documented process, applied to all eight systems in the inventory | All | Applies to any organisation using AI, however it was obtained |
| A.5.3 | A.5 Assessing impacts of AI systems | Documentation of AI system impact assessments | Yes | The assessment must be retained as documented information, or an assessor cannot verify it happened. | Implemented — one assessment per system, version-controlled, review dates set | All | Applies to any organisation using AI, however it was obtained |
| A.5.4 | A.5 Assessing impacts of AI systems | Assessing AI system impact on individuals or groups | Yes | AI-003 ranks job applicants and directly affects individuals; this control is the reason its risk tier is High. | Implemented — quarterly disparity testing on ranked-vs-hired outcomes for AI-003 | AI-003, AI-006, AI-007 | Applies to any organisation using AI, however it was obtained |
| A.5.5 | A.5 Assessing impacts of AI systems | Assessing societal impacts of AI systems | Yes | Impact assessment must look beyond the organisation. Guardrail 10 asks for the same broader view. | Implemented — societal section completed in each assessment; most return "limited", which is itself a recorded finding | All | Applies to any organisation using AI, however it was obtained |

| CONTROL | THEME | CONTROL TITLE | APPLICABLE | JUSTIFICATION | IMPLEMENTATION STATUS | AI SYSTEMS IN SCOPE | IF YOU ONLY PROCURE AI |
|---|---|---|---|---|---|---|---|
| A.6.1.2 | A.6 AI system life cycle | Objectives for responsible development of AI systems | Yes | The organisation develops two AI systems in-house, so the development controls are in scope. | Implemented — responsible-development objectives set and measured at management review | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.6.1.3 | A.6 AI system life cycle | Processes for responsible AI system design and development | Yes | Applies to the two in-house systems; the design process must be documented rather than tacit. | Implemented — design review gate added to the existing SDLC | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.6.2.2 | A.6 AI system life cycle | AI system requirements and specification | Yes | Requirements must be specified before build, including performance and fairness criteria. | Implemented — acceptance criteria defined per system before development begins | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.6.2.3 | A.6 AI system life cycle | Documentation of AI system design and development | Yes | Design decisions must be traceable; an assessor will ask why a model was chosen over an alternative. | Implemented — design records retained per model version | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.6.2.4 | A.6 AI system life cycle | AI system verification and validation | Yes | Guardrail 4 requires testing model performance before deployment. Without it there is no evidence the system works as claimed. | Implemented — pre-deployment test results retained against the stated acceptance criteria | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.6.2.5 | A.6 AI system life cycle | AI system deployment | Yes | Deployment must be controlled and reversible; a model rollback path is required. | Implemented — staged rollout with a documented rollback procedure | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.6.2.6 | A.6 AI system life cycle | AI system operation and monitoring | Yes | Guardrail 4 also requires monitoring after deployment. Drift is the failure mode a pre-deployment test cannot catch. | Implemented — drift monitoring on both in-house models; vendor-reported metrics reviewed for procured systems | All | Applies to any organisation using AI, however it was obtained |
| A.6.2.7 | A.6 AI system life cycle | AI system technical documentation | Yes | Technical documentation supports both internal assurance and, for EU-facing products, AI Act Article 11. | Implemented — maintained per in-house system; requested from vendors for procured systems | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.6.2.8 | A.6 AI system life cycle | AI system recording of event logs | Yes | Logging supports guardrail 9 — records that let a third party assess compliance. | Implemented — inference and decision logs retained per the ISMS retention schedule | All | Applies to any organisation using AI, however it was obtained |
| A.7.2 | A.7 Data for AI systems | Data for development and enhancement of AI systems | Yes | Training and enhancement data must be identified and controlled for the in-house models. | Implemented — training datasets registered and access-controlled | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.7.3 | A.7 Data for AI systems | Acquisition of data | Yes | Data acquisition must be lawful and documented, including data acquired through the product itself. | Implemented — acquisition basis recorded per dataset alongside the privacy assessment | All | Applies to any organisation using AI, however it was obtained |
| A.7.4 | A.7 Data for AI systems | Quality of data for AI systems | Yes | Guardrail 3 requires data quality management. Poor data quality is the most common cause of biased output. | Implemented — quality criteria defined and checked before each training run | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.7.5 | A.7 Data for AI systems | Data provenance | Yes | Provenance is required even where the answer is uncomfortable. For procured systems the vendor does not disclose training data, and that gap is recorded rather than hidden. | Partially implemented — full lineage for in-house models; vendor training data undisclosed for AI-001, AI-003, AI-004, AI-005, AI-007 and logged as a residual risk | All | Applies to any organisation using AI, however it was obtained |
| A.7.6 | A.7 Data for AI systems | Data preparation | Yes | Preparation steps materially change model behaviour and must be reproducible. | Implemented — preparation pipelines version-controlled with the model code | AI-002, AI-006 | Development control — a pure-deployer may scope this down with a documented justification |
| A.8.2 | A.8 Information for interested parties | System documentation and information for users | Yes | Users need to know when they are interacting with AI. Guardrail 6 requires it. | Implemented — in-product disclosure for AI-004; careers page disclosure for AI-003 | All | Applies to any organisation using AI, however it was obtained |

| CONTROL | THEME | CONTROL TITLE | APPLICABLE | JUSTIFICATION | IMPLEMENTATION STATUS | AI SYSTEMS IN SCOPE | IF YOU ONLY PROCURE AI |
|---|---|---|---|---|---|---|---|
| A.8.3 | A.8 Information for interested parties | External reporting | Yes | Guardrail 7 requires a route for affected people to challenge an outcome. | Implemented — published review-request route for AI-003, with requests and outcomes logged | AI-003 | Applies to any organisation using AI, however it was obtained |
| A.8.4 | A.8 Information for interested parties | Communication of incidents | Yes | AI incidents must be communicable to affected parties; the existing security incident process did not cover model failures. | Implemented — AI incident category added to the ISMS incident procedure | All | Applies to any organisation using AI, however it was obtained |
| A.8.5 | A.8 Information for interested parties | Information for interested parties | Yes | Customers ask about AI use in security questionnaires; the answer must be prepared rather than improvised. | Implemented — AI section added to the trust page and the standard questionnaire response set | All | Applies to any organisation using AI, however it was obtained |
| A.9.2 | A.9 Use of AI systems | Processes for responsible use of AI systems | Yes | Applies to every organisation using AI regardless of who built it, which is why a pure-deployer cannot exclude A.9. | Implemented — acceptable-use process covering all eight systems including procured ones | All | Applies to any organisation using AI, however it was obtained |
| A.9.3 | A.9 Use of AI systems | Objectives for responsible use of AI systems | Yes | Objectives must be measurable under clause 6.2, not aspirational. | Implemented — objectives set with metrics and reviewed at management review | All | Applies to any organisation using AI, however it was obtained |
| A.9.4 | A.9 Use of AI systems | Intended use of the AI system | Yes | Intended use must be defined so that use outside it can be recognised. Shadow AI (AI-008) is what happens when it is not. | Implemented — intended use recorded per system; AI-008 brought into scope at discovery | All | Applies to any organisation using AI, however it was obtained |
| A.10.2 | A.10 Third-party and customer relationships | Allocating responsibilities | Yes | Five of eight systems are procured, so responsibility between provider and deployer must be explicit. | Implemented — responsibility split documented per vendor in the supplier register | AI-001, AI-003, AI-004, AI-005, AI-007 | Applies to any organisation using AI, however it was obtained |
| A.10.3 | A.10 Third-party and customer relationships | Suppliers | Yes | Guardrail 8 requires transparency across the AI supply chain. Supplier assurance is where the provenance gap in A.7.5 has to be chased. | Implemented — AI clauses added to vendor agreements; annual supplier assurance review | AI-001, AI-003, AI-004, AI-005, AI-007 | Applies to any organisation using AI, however it was obtained |
| A.10.4 | A.10 Third-party and customer relationships | Customers | Yes | AI-004 surfaces third-party AI to the organisation's own customers, so downstream obligations apply. | Implemented — AI use disclosed in customer documentation and the DPA | AI-004 | Applies to any organisation using AI, however it was obtained |