ISO/IEC 27090 Readiness Checklist

A self-assessment aligned to the AI-security themes of ISO/IEC 27090:2026 — the first international standard for security threats and failures in AI systems. PlayCISO · playciso.com/tools/iso-27090-readiness

This is a practical readiness aid, not a reproduction of the ISO text and not a certification. ISO/IEC 27090 is guidance; there is no "ISO 27090 certificate". Buy the official standard from ISO or your national body.

1. Governance & scope

- ☐ An owner is accountable for AI security across the AI lifecycle (not just IT security generally).
- ☐ In-scope AI systems are enumerated: models, agents, RAG pipelines, and embedded third-party AI.
- ☐ AI security expectations are written into policy and supplier contracts.

2. AI asset inventory

- ☐ Every model, dataset and material dependency is inventoried (an AIBOM / ML-BOM).
- ☐ Provenance and licence are recorded for each model and training/eval dataset.
- ☐ Inventory is refreshed when a model, dataset or agent changes.

3. Threat identification (AI-specific)

- ☐ Adversarial inputs / evasion are considered for each deployed model.
- ☐ Data & model poisoning risks in training and fine-tuning are assessed.
- ☐ Model theft / extraction and membership-inference exposure are assessed.
- ☐ Prompt injection and tool-abuse are assessed for LLM and agentic systems.
- ☐ AI-specific failure modes (hallucination, unsafe autonomy) are treated as security-relevant.

4. Controls

- ☐ Input validation / output filtering appropriate to each model are in place.
- ☐ Access to models, weights, prompts and vector stores is least-privilege.
- ☐ Human-in-the-loop or guardrails gate high-impact autonomous actions.
- ☐ Sensitive data in prompts, context and logs is minimised and protected.

5. Monitoring & response

- ☐ AI systems emit security-relevant telemetry (abuse, drift, anomalous tool use).
- ☐ The incident-response plan names AI-specific scenarios and who is paged.
- ☐ Model/agent changes are logged for audit and rollback.

6. Supply chain

- ☐ Third-party models and AI vendors are risk-assessed continuously, not once.
- ☐ Known-vulnerability and advisory feeds cover AI components you depend on.
- ☐ You can produce an AIBOM for a customer or regulator on request.

7. Documentation for audit

- ☐ Risk assessments, control decisions and residual-risk sign-off are written down.
- ☐ Evidence maps to the themes above so a reviewer can follow it.
- ☐ The record is current — reviewed on a set cadence, not left to go stale.

Run the interactive version and track your score over time at playciso.com/tools/iso-27090-readiness. Read the explainer at playciso.com/blog/iso-iec-27090-ai-security-standard-what-ciso-should-do.