CWS
CYBERWERKSUITE

ISO 42001 Checklist

45 checks — AI governance, data, AI security and autonomous agents

A working list for organisations that use AI — not only for those seeking certification. Every item says what it rests on: the standard, Annex A, separate law, or recognised practice.

ISO/IEC 42001:2023 · Annex A · Regulation (EU) 2024/1689 (AI Act) only where it actually applies

45	7	4
Checks	Sections	Levels

CyberWerkSuite
cyberwerksuite.com

Dr. Sait Yalazay
CISO · DPO · Author | CISM · CIPP · AAISM · LA 27001, 27701, 22301, 42001
Architect of Automated Compliance Systems for NIS2, GDPR, ISMS, BCM, DORA, Cloud Security (C5), Tisax & AI Act

A standard with no deadline — and questions that get asked anyway.

ISO/IEC 42001 is voluntary. There is no deadline, no supervisory authority and no penalty. The questions get asked regardless: by customers in procurement, by the auditor of your ISMS, and by your own board on the day an agent does something that cannot be undone.

The four levels

WHAT SITS BEHIND EACH ITEM

N — a clause of ISO/IEC 42001. Binding for those who want certification.
A — a control from Annex A. Applicability is justified in the Statement of Applicability, not assumed.
R — a separate legal duty (GDPR, AI Act). It applies independently of this standard; 42001 neither creates it nor removes it.
H — recognised practice and state of the art. **Not in the standard.** Included because an auditor asks for it and an attacker uses it.

What this list is not

BOUNDARY

Not a certification audit and not a Statement of Applicability. It tests whether the questions are **answerable** — not whether a certification body accepts the answer. Where an item carries **H**, it is not a requirement of the standard; deleting it breaches nothing.

One pass per **AI system**, not per organisation. A firm with four models in production has four answers to the same question — and that is exactly where the assessment falls apart.

The AI management system

The frame first: scope, policy, roles, inventory.

☐	1	The scope of the AI management system is set down in writing and approved by top management.	**N** · 42001 · 4.3
☐	2	A signed **AI policy** exists; it names the purpose, the reach and the limits of AI use, and staff can actually find it.	**N** · 42001 · 5.2
☐	3	Roles and authorities are assigned and communicated — including named owners for AI operations, AI risk and **ethical questions**.	**N** · 42001 · 5.3
☐	4	A complete **AI inventory** is maintained and kept current; for every system the purpose, owner and classification are recorded.	**A** · Annex A.4.4
☐	5	Introducing a new AI system passes through a defined **approval process** in lifecycle management.	**A** · Annex A.6.2.2
☐	6	A process exists that finds **uninventoried AI (“shadow AI”)** through procurement, SaaS, network and discovery controls and brings it into scope.	**H** · Practice
☐	7	An internal **audit programme** for the AI management system is established; findings are documented and tracked to closure.	**N** · 42001 · 9.2

Risk and impact assessment

What gets assessed, what it is compared against, and when it is done again.

☐	8	AI risk criteria — acceptance thresholds and risk tolerance — are defined, approved by management and maintained.	**N** · 42001 · 6.1.2
☐	9	The risk assessment procedure is repeatable and expressly covers **bias, explainability, robustness and security**.	**N** · 42001 · 6.1.4
☐	10	The risk treatment process reconciles the determined measures against **Annex A**, so that no necessary control is left out.	**N** · 42001 · 6.1.3
☐	11	For every Annex A control it is recorded whether it is applied, why the decision went that way, how far implementation has got and who owns it — in **one** document, not scattered.	**N** · 42001 · 6.1.3 d
☐	12	Before deployment it is examined and written down **who the system can affect** — the individual, whole groups of people, and the world outside the organisation.	**N** · 42001 · 8.3
☐	13	Risk assessments are repeated at planned intervals and on **material change** — model, data, purpose, context, regulation — and the triggers are documented.	**N** · 42001 · 8.2

CWS
CYBERWERKSUITE

Software

AIGovernanceSuite

a tool from **CyberWerkSuite**

45 IS THE PUBLIC CUT OF 111

The list you are reading is the public extract.

The tool holds 111 controls — each with its assessment question, applicability decision, justification, implementation status and owner. What comes out at the end is not a table but a Statement of Applicability you can hand over.

111 controls

Governance, risk, data, AI security, agents, supply chain — one catalogue instead of six files.

SoA as output

Applicability, justification, status and owner per control — as a document, not a screenshot.

Annex A reconciliation

Risk treatment is checked against Annex A so that no required control goes quietly missing.

Several AI systems

One pass per system, one picture for the board.

111	38	1	DE/EN
Controls	Annex A links	SoA output	Languages

See AIGovernanceSuite cyberwerksuite.com

NOTE

The tool does not replace a certification body. It produces the evidence a certification body wants to see.

Where Compliance Meets Automation

CyberWerkSuite cyberwerksuite.com

Data and data provenance

Where the data came from, what was done to it, and what the model kept.

☐	14	For development and improvement data it is settled who releases it, how its quality is checked, and how anyone would notice that it **represents the target population unevenly**.	**A** · Annex A.7.4
☐	15	Collection and selection are documented per dataset: categories, sources, demographics, rights in personal data and copyright, and the **lawfulness of web scraping**.	**A** · Annex A.7.3
☐	16	**Data provenance** is recorded traceably across the whole lifecycle — creation, update, transfer.	**A** · Annex A.7.3
☐	17	Criteria for data preparation — cleaning, imputation, normalisation, encoding, labelling — are defined and documented.	**A** · Annex A.7.6

What the model kept.

The part that only hurts once somebody files a subject access request.

☐	18	Production feedback and **ground-truth labelling pipelines** are quality-assured and protected against poisoning before they feed retraining.	H · Practice
☐	19	The model is tested for **memorisation and regurgitation** of personal data — membership inference, extraction, prompt recall.	H · Practice
☐	20	Data subject rights are honoured — as far as technically feasible — against training data and the trained model too: **unlearning**, output suppression, retraining triggers.	R · GDPR Art. 15–21
☐	21	Retention periods for training data, prompts, outputs and AI logs are defined and **technically enforced** with automatic deletion or review triggers.	H · Practice

AI security: model, chain, guardrails

The part management-system checklists skip.

☐	22	Before an AI system is built or bought, it is written down **what it is for**, what data it needs, and what will be measured to tell whether it does its job.	**A** · Annex A.6.2.2
☐	23	Changes to model, data, prompts or configuration trigger, before release, a fresh risk and impact assessment, **re-validation**, and a release decision with a **rollback plan**.	**A** · Annex A.6.2.3
☐	24	Every production model version is held in a register linked to training-data version, code commit, hyperparameters, metrics and release, so that it can be **rebuilt reproducibly**.	**H** · Practice
☐	25	A machine-readable **AI/ML bill of materials (ML-BOM)** is produced per system — base models, datasets, weights, adapters, plugins, libraries, with origin, licence and vulnerability tracking.	**H** · Practice

Input, output, and what happens in between.

Four items that appear in no management-system checklist — and in every incident report.

☐	26	Model weights and artefacts are **cryptographically signed** with verifiable build provenance; unsigned artefacts are rejected at the deployment gate.	H · Practice
☐	27	Tested defences against **prompt injection and jailbreaks** are in place — input filtering, system-prompt hardening, bypass testing.	H · Practice
☐	28	Untrusted content from RAG, tools, the web, files and e-mail is isolated and marked so that it **cannot be read as instructions**.	H · Practice
☐	29	A model-independent **guardrail layer** moderates input and output, **fails closed** on error, and logs what it blocked.	H · Practice

CWS
CYBERWERKSUITE

Consulting

Two rulebooks, one engagement.

AI Act (EU) 2024/1689 and ISO/IEC 42001 — assessed separately, implemented together.

CWS AI COMPLIANCE

Classification first, then the legal duties, then the management system.

✓ **Classification**	Which AI systems you run, which role you hold — provider, deployer, importer — and which the AI Act reaches.
✓ **Legal duties**	What applies regardless of any standard: transparency, human oversight, AI literacy, the GDPR interface.
✓ **Management system**	ISO/IEC 42001 as the carrying structure — policy, risk, impact assessment, Annex A reconciliation, SoA.
✓ **AI security**	Prompt injection, guardrails, model provenance, autonomous agents.

For organisations already running AI that must now produce the evidence.

Book an intake call cyberwerksuite.com

ALSO IN THE PORTFOLIO

AIGovernanceSuite

AI MANAGEMENT SYSTEM

111 controls, SoA output, multi-system operation.

cyberwerksuite.com

Where Compliance Meets Automation

Autonomous agents

When the system stops answering and starts acting.

☐	30	Autonomous agents run on **least privilege** with narrow, explicitly approved tools (allow-lists, deny by default); every capability is authorised individually.	H · Practice
☐	31	**Irreversible or consequential agent actions** — payments, deletions, external communication, code merge or deploy — require human approval with a complete audit trail.	H · Practice
☐	32	Token, compute, concurrency, depth and cost quotas are enforced per tenant with **circuit breakers**, agent loop limits and budget alarms.	H · Practice
☐	33	Model-generated or agent-executed code runs in **isolated, network-restricted, short-lived sandboxes** with resource limits and no access to production secrets.	H · Practice
☐	34	All LLM output is treated as untrusted and is sanitised or encoded appropriately before every downstream sink — browser, shell, database, code execution, third-party systems.	H · Practice
☐	35	**Human oversight** is in place for AI-assisted decisions, and the person overseeing can actually intervene.	A · Annex A.6.2 · AI Act Art. 14

Supply chain and third parties

Someone else's model is still your responsibility.

☐	36	A structured **AI procurement process** with fixed assessment questions is implemented and applied consistently.	**A** · Annex A.10.2
☐	37	Third-party and base models are risk-assessed before adoption — lawfulness of training data, capabilities, limitations — and checked for **licence and terms-of-use conformity**.	**A** · Annex A.9.2
☐	38	The **subcontractors** of AI suppliers along the AI supply chain are known and assessed against the relevant requirements.	**A** · Annex A.10.4
☐	39	The terms and conditions of every AI service in use are analysed and assessed **before use**, not after.	**H** · Practice
☐	40	Responsibilities across the AI lifecycle are assigned and documented between the organisation, partners, suppliers, customers and third parties — including the roles of **controller and processor**.	**R** · GDPR Art. 26, 28

Transparency, people and operation

Who gets told what — and who notices when the model slips.

☐	**41** Whoever operates the system **knows it is AI** — and can find out what it can and cannot do without having to ask anyone.	**A** · Annex A.8.2 · AI Act Art. 50
☐	**42** The use of AI is communicated to external stakeholders **transparently and intelligibly**.	**A** · Annex A.10.4
☐	**43** Affected persons can raise **complaints** about AI-assisted decisions, and an accessible route for doing so exists.	**A** · Annex A.8.2
☐	**44** A confidential, **retaliation-free reporting channel** for concerns about AI systems exists, works in practice, and its results feed the management review.	**H** · Practice
☐	**45** AI systems are **monitored continuously in operation** — in particular for drift in model behaviour and for degrading performance.	**A** · Annex A.10.2

How you can tell this list is honest.

The same three rules we build every CWS catalogue on.

1

Provenance sits on every item

Standard, Annex A, separate law or practice — visible on the item itself, not in a footnote. Delete an **H** item and you breach nothing. You should be able to see that.

2

The standard's text is not reprinted

No wording of ISO/IEC 42001 is reproduced verbatim. Every item is a CWS formulation of the assessment question. The standard itself you buy from ISO or your national body.

3

The status is dated

Edition ISO/IEC 42001:2023. AI Act references are as at **20 September 2026** and appear only where they actually apply — the standard creates no legal duty.

These 45 are an extract from 111. The extract is the selection we check first in projects — not the easier half.

CWS
CYBERWERKSUITE

AI Governance

From the list to a system.

You have the gaps. What is missing is the order.

CWS CONSULTANCY · AI GOVERNANCE

Not a paper project — the work happens on real systems.

✓ **Where you stand**	Your AI systems held against the full control catalogue — applicability, justification and status per control.
✓ **Order of work**	What comes first, what can wait, what is legal duty anyway — with dates and effort, not traffic lights.
✓ **Building it**	Policy, roles, inventory, impact assessment, Annex A reconciliation, guardrails — built, not described.
✓ **Proving it**	Statement of Applicability, internal audit, management review — with or without an intention to certify.

For organisations already using AI that now have to catch up on the evidence.

Book an intro call cyberwerksuite.com

THE TOOL BEHIND IT

AIGovernanceSuite

AI MANAGEMENT SYSTEM

The full control catalogue, Statement of Applicability as a document.

cyberwerksuite.com

Where Compliance Meets Automation