Assurance without theatre

A certificate is a statement that a management system exists, not evidence that a model works. The board's task is to tell the difference, and then to ask for the second thing.

THE MACHINERY

Dec 2023

42001

The AI management system standard. It certifies a system of management, not a model.

Jul 2025

42006

Requirements for the bodies that audit and certify against it, eighteen months later.

READINESS

9 of 27

AUTHORITIES

Member states with both market surveillance and notifying authorities designated, June 2026.

MOVING TARGET

12

RISKS

Categories named in the generative profile of the American framework, itself now under revision.

The certificate covers the management system

ISO/IEC 42001 was published in December 2023 as a management system standard for artificial intelligence. It sets requirements for establishing, implementing, maintaining and improving an AI management system, and addresses management system aspects rather than the details of any AI application.[1] Certification is voluntary, and the standards body does not certify organisations at all: independent certification bodies do.[1]

The infrastructure behind those certificates arrived later. ISO/IEC 42006, published in July 2025, sets requirements for the bodies that audit and certify AI management systems, including the competence the audit team must hold collectively and how audit duration is calculated.[2] National accreditation followed in 2026, with the first United Kingdom accreditation granted in January and the same certification body accredited in the United States that March.[2]

One limitation deserves the board's attention. A 42001 certificate carries no presumption of conformity with the European AI Act, because that presumption requires harmonised standards developed at the Commission's request and cited in the Official Journal, and as at mid-2026 none had been cited.[3] A certificate is useful, but it is not a permit.

The standards are late, and the law says so

The regulation that deferred the high risk obligations of the AI Act gives its reasons in its own recitals: the delayed preparation of standards, and the delayed establishment of governance and conformity assessment frameworks at national level.[4] The legislator has recorded, in the text of the law, that the assurance machinery was not ready.

The machinery is now being built at speed. The European standards bodies adopted acceleration measures in October 2025, including direct publication after a positive enquiry vote, and described them as exceptional and temporary.[5] The first harmonised European standard for AI Act purposes, covering the quality management system required of high risk providers, was approved in July 2026 and is a deliberate architectural break from both ISO 9001 and ISO/IEC 42001.[5] National readiness lags further behind: in June 2026, nine member states had designated both market surveillance and notifying authorities, twelve had designated some, and six had designated none.[5]

The frameworks themselves are moving

The other reference points a supplier may cite are also in motion. The American risk management framework for AI, published in January 2023 with a generative profile added in July 2024, remains voluntary and organises work around four functions: govern, map, measure and manage. The generative profile names twelve risk categories, from confabulation and information integrity to value chain and component integration.[6] It is now being revised under a federal action plan that directs the removal of references to misinformation, to diversity, equity and inclusion, and to climate change, and that revision was still in progress through 2026.[6] Alignment with a named framework therefore describes a moment rather

THE ROUTE

Annex VI

INTERNAL

The assessment route for most high risk categories. No notified body is involved.

WHAT TO ASK FOR

1 **The tester**
Who ran the evaluation, and are they independent of the builder?

2 **The hours**
How long the testing ran. A number, not an adjective.

3 **The residue**
What failed, what was not tested, and what remains open.

THE MEASURE

18 of 391

AUDITED

Employers posting the bias audit required by New York City's hiring algorithm law.

than a permanent state, and a supplier claiming it should be asked which framework, which version, and as at what date.

Most high risk AI is assessed by the people who built it

Article 43 of the AI Act sets the conformity assessment routes. For biometrics, a provider may choose between internal control and a notified body. For the other Annex III categories, which include education, employment, access to essential services and creditworthiness, migration and the administration of justice, the route is conformity assessment based on internal control, which does not involve a notified body at all.[7] A provider that decides its own system is not high risk must register that determination in the European database.[7]

This is the design rather than a scandal, and the consequence is direct: the attestation arriving with a supplier's product is usually the supplier's own, and independent assurance has to be arranged rather than assumed.

A usable document names who tested it

The regulation is specific about what a high risk system's documentation must contain: a general description, the development process, monitoring and control arrangements, performance metrics and why they are appropriate, the risk management system, lifecycle changes, the standards applied, the declaration of conformity, and the post-market monitoring plan.[8] That monitoring plan is not decoration; it must collect and analyse performance data across the system's life.[8]

Below the regulation sits a practice worth borrowing. The model card format, proposed in 2019, asks for intended use, factors, metrics, evaluation and training data, quantitative analyses disaggregated by those factors, and stated caveats.[9] Current system cards go further, naming who tested a model and for how long: one published in July 2026 records a national AI security institute running multi-stage agentic cyber ranges, one external group spending roughly a hundred hours red teaming safeguards and another about sixteen, with the provider stating plainly that it does not expect its classifiers to be perfectly robust.[10]

That is the shape of a document a board can use: the tester, the duration, the failures, the residual risk. A document naming none of those is marketing in a serif typeface. Registers show the same distance between intent and practice. The United Kingdom's algorithmic transparency standard was made mandatory across central government in December 2024; by May 2025 the total number of published records stood at fifty-nine.[11]

Assurance fails in public

The gap between a mandated audit and an actual audit can be measured. Under New York City's hiring algorithm law, a study of 391 employers found 18 that had posted the required bias audit, and 11 that had posted both the audit and the notice in compliant form.[12] Where the audit is the control, the control was mostly absent.

The assurance firms have supplied their own examples. One agreed a partial refund to a national government in October 2025 over a report containing apparent AI generated errors; another withdrew a published report in June 2026 after only five of its forty-five citations were found to point where they claimed.[13] In the public sector, a statistics regulator required a health body to caveat published claims about a data platform's benefits, including a statement that no conclusion about cause could be drawn, and a hospital trust paused an ambient documentation trial that had not passed its own governance.[13]

Internal audit knows the position. In the 2026 risk surveys, technology and AI risk rose sharply in chief audit executives' rankings, while far fewer named

it among the areas receiving audit time.[14] A risk that is recognised but unaudited is a finding waiting for an author.

The assurance market is young

There is a market in third party AI assurance, and it is small. The United Kingdom government sized it at just over one billion pounds of gross value added, across 524 supplying firms of which 84 are specialised.[15] Its roadmap for trusted third party assurance, published in September 2025, sets out to build a profession: a skills and competencies framework, a consortium and a fund.[15] A profession is precisely what the field does not yet have.

Professional bodies are moving at a similar pace. The audit regulator issued guidance on the use of AI in audit in June 2025 and further guidance on generative and agentic tools in March 2026, holding to one rule throughout: the human auditor remains accountable, and the tools sit inside existing quality management obligations.[16] A certification for AI audit was launched in 2025 for holders of existing audit credentials.[16] Assurance capacity is therefore scarce and unevenly spread, and the organisation that specifies what it wants early will get better work than the one going to market during an incident.

CAPACITY

84

SPECIALISTS

Specialised AI assurance firms counted in the UK market assessment, within 524 suppliers overall.

Segregation is the whole point

The standard for certification bodies forbids them from consulting for their certification clients on AI, information security, data protection or risk management, from performing those clients' internal audits, and from recommending specific solutions.[2] The three lines model makes the same distinction inside the organisation: second line functions that support and challenge are part of management, and only the third line is independent of it.[17]

Boards can apply that rule to their own arrangements without waiting to be told. The team that builds the model does not attest to it. The function that bought the tool does not evaluate it. The committee that approved the programme does not mark its own homework. Where an organisation is too small to hold all three apart, the board says so in the minutes and names the compensating arrangement. That is a defensible position. Silence is not.

THE RULE

3

LINES

Build, challenge, assure. Only the third is independent of management.

Ninety days of useful work

1. Ask for one AI system's full evidence pack, chosen by the board rather than the executive, and read it. Its quality is the quality of the assurance.
2. Establish whether each supplier's conformity claim rests on a notified body or on internal control, and record which.
3. Require named third party evaluation for systems touching customers, employment decisions or safety, with the tester, the hours and the findings stated.
4. Separate building from attesting on paper: who builds, who challenges, who assures, and where independence sits.
5. Ask internal audit what proportion of its plan covers AI, and compare that with where AI sits in the risk register.
6. Fix what the standing report contains: evaluations run, failures found, systems withdrawn, and issues still open.

THE TEST

1

PACK

One evidence pack, chosen by the board. Its quality is the quality of the assurance.

Good assurance is uncomfortable to read

Assurance that is working is uncomfortable to read. It names what failed, gives numbers rather than adjectives, and leaves questions open. Assurance that is theatre is smooth, certificated and silent about failure. A board that has never had an uncomfortable AI report is being managed, not assured.

Twelve questions, and the answers that should worry you

Ask the executive	Worry if the answer is
What exactly does our supplier's certificate cover?	"They are ISO certified." The certificate covers a management system, not a model.
Was this system assessed by a notified body, or by its provider?	"It is fully compliant." For most high risk categories the assessment is the provider's own.
Who has independently tested this system, for how long, and what failed?	"The documentation is thorough." Ask for the tester, the hours and the findings.
What does our post-market monitoring collect?	"We monitor performance." Which metrics, at what frequency, reviewed by whom?
Who inside the organisation attests to the systems we build?	"The AI team signs those off." That is the first line marking its own homework.
What proportion of internal audit's plan covers AI?	"It is on the radar." Compare that with where AI sits in the risk register.
When did a report last tell us that something had failed?	"We have had no issues." Either nothing is deployed, or nothing is being examined.
Which AI claims of ours were checked by someone who did not write them?	"The team reviewed them." Two firms refunded or withdrew published work in 2025 and 2026.
Does our inventory reconcile with procurement and the ledger?	"There is a list." Ask who signed it, and when it was last reconciled.
What would we show a regulator tomorrow morning?	"We would pull it together." Evidence assembled after the question is not evidence.
Which decisions can we explain to the person they affect?	"The model is inherently complex." Then the complaint will be decided without us.
What is still unresolved, and who owns it?	"Nothing outstanding." Every honest assurance report has an open list.

NOTES AND SOURCES

1 ISO/IEC 42001:2023, *Artificial intelligence — Management system*, December 2023; ISO, *ISO/IEC 42001 explained*, accessed 17 September 2026.

2 ISO/IEC 42006:2025, 7 July 2025, clauses 5.2.2 and 7.1.3; UKAS accreditation 15 January 2026 and ANAB accreditation 10 March 2026.

3 lawandtechnology.eu, *ISO/IEC 42001 and the AI Act*, 4 July 2026; AI Act Article 40(1).

4 Regulation (EU) 2026/1744, recitals 2 and 40; in force 27 July 2026.

5 CEN and CENELEC, acceleration measures, 23 October 2025, and EN 18286:2026, approved 12 July 2026; AI Act national implementation tracker, 17 June 2026.

6 NIST AI Risk Management Framework 1.0, 26 January 2023, and AI 600-1 generative profile, July 2024; *America's AI action plan*, July 2025; revision in progress, June 2026.

7 Regulation (EU) 2024/1689 (AI Act), Article 43, Annex VI and Article 49(2).

8 AI Act, Annex IV and Article 72.

9 Mitchell et al., *Model cards for model reporting*, FAT* 2019.

10 Anthropic, *Claude Opus 5 system card*, 24 July 2026.

11 DSIT, algorithmic transparency recording standard mandatory scope policy, 17 December 2024; Data in Government, 8 May 2025.

12 Wright et al., *Null compliance: NYC Local Law 144*, FAccT 2024.

13 CFO Dive, 7 October 2025, on the Deloitte refund; *The Register*, 12 June 2026, on the withdrawn KPMG report; Office for Statistics Regulation on NHS England, 22 July 2026; Health Service Journal, 1 April 2026.

14 ECIIA, *Risk in Focus 2026*, September 2025; Institute of Internal Auditors global risk survey, 15 September 2026.

15 DSIT, *Assuring a responsible future for AI*, 2024, and *Trusted third-party AI assurance roadmap*, 3 September 2025.

16 Financial Reporting Council guidance on AI in audit, 26 June 2025, and on generative and agentic AI, 30 March 2026; ISACA Advanced in AI Audit certification, 19 May 2025.

17 Institute of Internal Auditors, *Three Lines Model*, updated September 2024, and *AI Auditing Framework*, 13 September 2024.

AltLibre Papers, board series, paper 04 of five. Adrian Hollister, AltLibre, 12 September 2026. altlibre.com