This document was translated from English by machine and may contain mistakes. The English original is at https://www.novaracg.com/2026/09/05/the-object-of-assurance/.

The Object of Assurance

Heather Grizzle · September 5, 2026

The AI assurance market certifies management systems, not systems. What it measures, and what it still leaves to the buyer of sign language AI.

Cov Ntsiab Lus

1. What the new AI assurance market measures, and what it still leaves to the buyer
2. The management system is not the system
3. The deferral moved the deadline, not the difficulty
4. What the horizontal market cannot see
5. The position Novara Consulting Group takes
6. What a buyer should require between now and December 2027
7. Cov chaw siv

What the new AI assurance market measures, and what it still leaves to the buyer

Heather M. Grizzle Novara Consulting Group September 2026

Artificial intelligence assurance is becoming an industry. That sentence would have been speculative two years ago and is now simply descriptive. The United Kingdom's Department for Science, Innovation and Technology counted over five hundred firms and roughly one billion pounds of gross value added in the domestic assurance market as of 2024, projected a market approaching nineteen billion pounds by 2035, and in September 2025 published a Trusted Third-Party AI Assurance Roadmap setting out how the state intends to help that market mature. In June 2026 the roadmap produced an institution: a consortium chaired by BCS, with techUK convening its industry advisory group, carrying a first-year mandate to draft a

voluntary code of ethics for assurance practitioners, build a skills and competencies framework, and map what information developers must release to the people auditing them. Alongside this, the certification route has been consolidating. ISO/IEC 42001 provides a management-system standard for artificial intelligence, ISO/IEC 42006:2025 sets the requirements for the bodies that audit against it, and national accreditation bodies including UKAS in the United Kingdom and ANAB in the United States have been granting accreditation to certification bodies through 2025 and 2026.

THE ASSURANCE MARKET AT A GLANCE

500+
firms in the UK AI assurance market (2024)

~£1bn
gross value added by UK AI assurance (2024)

~£19bn
projected UK market by 2035

~350
organisations certified to ISO/IEC 42001 worldwide (mid-2026)

1M+

organisations certified to ISO 9001, for comparison

Source: Department for Science, Innovation and Technology (2025); industry certification compilations (2026)

This is what the beginning of a profession looks like. It has standards, accreditation, a nascent competency framework, an emerging code of ethics, and a defensible claim on public confidence. Institutions that have been waiting for a recognised answer to the question of how to tell whether an artificial intelligence system can be trusted now have something to point at.

They should look closely at what it points to.

The management system is not the system

ISO/IEC 42001 is structurally parallel to ISO/IEC 27001. It is a management-system standard, which means it certifies that an organisation has established, documented, and maintained a set of processes for governing artificial intelligence, and that those processes are subject to internal audit and continual improvement. It does not certify that any particular system built inside that management system produces adequate output. This is not a defect in the standard. It is what a management-system standard is. The confusion arises downstream, in procurement, where a certificate is read as a statement about the product rather than a statement about the producer.

The distinction matters most where the output is the whole point of the purchase. A vendor selling automated sign language translation can hold a properly accredited certificate attesting to a well-run artificial intelligence management system while shipping output that a fluent Deaf user cannot follow. Nothing in the certificate is false. Nothing in it addresses the question the buying institution actually needs

answered, which is whether the person standing at the counter, or in the examination room, or in front of the disciplinary panel, receives information they can understand.

Scale confirms the point rather than contradicting it. Industry compilations put worldwide certification against ISO/IEC 42001 at roughly three hundred and fifty organisations by mid-2026. For comparison, ISO 9001 is held by more than a million. The standard is young and adoption is early, which is unremarkable. What is remarkable is how much interpretive weight a small number of certificates are already being asked to carry in procurement conversations, and how rarely a buyer asks what the certificate is a certificate of.

What a management-system certificate attests, and what the buyer needs to know

Question	Answered by an ISO/IEC 42001 certificate	Needed by the procuring institution
Has the vendor documented processes for governing AI?	Yes	Useful but not sufficient
Are those processes internally audited and improved?	Yes	Useful but not sufficient
Does this system at this version produce adequate output?	No	Yes
Can the person relying on the output understand it?	No	Yes

Source: Author's analysis of ISO/IEC 42001:2023

The deferral moved the deadline, not the difficulty

The European Union's timetable has now shifted in a way that will shape institutional behaviour for the next eighteen months. The Digital Omnibus on artificial intelligence

was adopted as Regulation (EU) 2026/1744, published on 24 July 2026 and in force from 27 July. It defers the high-risk obligations of the AI Act for standalone Annex III systems to 2 December 2027, and for artificial intelligence embedded in regulated products under Annex I to 2 August 2028. The transparency obligations under Article 50 were not deferred and applied from 2 August 2026, with a short grace period running to 2 December 2026 for generative systems already on the market before the earlier date.

The reason for the deferral is the substantive part. The original schedule assumed a working ecosystem of harmonised technical standards from CEN-CENELEC, designated and resourced national market surveillance authorities, and notified bodies able to perform conformity assessment. By late 2025 that ecosystem visibly did not exist. Legislators chose fixed new dates rather than a conditional mechanism tied to standards availability, on the reasoning that predictability serves compliance better than a floating trigger.

Read plainly, this is an admission that the machinery for producing evidence about high-risk systems is not yet built. The regulatory clock moved because the evidentiary infrastructure was late. That is a very different situation from one in which the evidence exists and the deadline was merely inconvenient.

EU AI ACT OBLIGATIONS AFTER THE DIGITAL OMNIBUS

24 Jul 2026
Regulation (EU) 2026/1744 published

27 Jul 2026
Digital Omnibus in force

2 Aug 2026
Article 50 transparency obligations apply

2 Dec 2026
Grace period ends for generative systems already on the market

2 Dec 2027
High-risk obligations apply to standalone Annex III systems

2 Aug 2028
High-risk obligations apply to AI embedded in Annex I products

Procurement, however, does not defer. Institutions buy on budget cycles, on grant timelines, on the schedule set by a service gap somebody needs closed this fiscal year. Between now and December 2027, a public body acquiring an accessibility artificial intelligence system will be offered three things: a management-system certificate if the vendor has one, a self-declaration of conformity, and an accessibility conformance report. The first speaks to the producer. The second speaks for itself, which is the problem. The third, in the form the market has normalised it, is a supplier-authored document mapping a product against a checklist of technical criteria that were not written with generated signed language in view. None of the three is evidence about whether the output works.

The three documents a buyer will be offered before December 2027

Document	Authored by	Speaks to	Evidence that the output works
Management-system certificate (ISO/IEC 42001)	Accredited certification body	The producer's governance processes	No
Self-declaration of conformity	The vendor	The vendor's own claim	No
Accessibility conformance report	The vendor	Technical criteria not written for generated signed language	No

Source: Author's analysis

What the horizontal market cannot see

The roadmap's own diagnosis is instructive here. It identifies four constraints on a trustworthy third-party assurance market: the quality of assurance activity itself, including the observation that a substantial share of available artificial intelligence governance tools rely on metrics that may be actively misleading; a shortage of

practitioners combining technical, legal, ethical and sector-specific competence; poor access for assurance providers to the training data, model details and internal documentation they need; and weak incentives for collaborative development of new assurance methods.

Three of those four bite hardest at the domain level. A horizontal assurance framework can tell an institution whether a vendor keeps a risk register. It cannot tell the institution whether the system reproduces the non-manual grammatical features that carry negation, question formation, topicalisation and conditional marking in a signed language, because horizontal assurance has no vocabulary for that question and no competent assessor assigned to it. It cannot evaluate whether a training corpus was assembled with the consent of the signers in it, under a licence permitting commercial deployment, from signers whose fluency was verified. It cannot distinguish a curated demonstration from a deployment. It cannot say whether the automated metric a vendor cites as proof of signing quality measures anything relevant, and for generated sign language there is currently no validated automated quality metric at all, which makes translation-adjacent scores offered as evidence of comprehensibility a category error rather than weak evidence.

The competency framework now under construction may eventually reach some of this. It has not yet, and the roadmap's own timeline leaves the certification of practitioners undated. In the interim, the gap sits precisely where consequence is highest. The systems being deployed into healthcare intake, emergency notification, transit announcement and employment processes are the ones whose failures are least reversible, and they are the ones the general assurance market is least equipped to examine.

What horizontal assurance can see, and what it leaves unexamined

Question the buyer needs answered	Horizontal assurance	Domain-level evaluation
Does the vendor maintain a documented AI management system?	Examined	Examined
Does the output carry non-manual grammar (negation; questions; topicalisation; conditionals)?	No vocabulary or competent assessor	Examined by fluent Deaf reviewers
Was the training corpus licensed for commercial use and collected with signer consent from fluency-verified signers?	Not examined	Examined
Is the evidence from a deployment or a curated demonstration?	Cannot distinguish	Examined
Does the cited quality metric measure comprehensibility?	Cannot say	Examined; no validated automated metric exists for generated signing

Source: Author's analysis; DSIT Trusted Third-Party AI Assurance Roadmap (2025)

The position Novara Consulting Group takes

Assurance requires an object. The question is not whether a vendor is well governed in the abstract but whether a specified system, at a specified version, produces output adequate to a specified use, judged by the people who must rely on it. Everything else is a proxy, and proxies accumulate quietly until an institution believes it has bought evidence when it has bought paperwork.

The Sign Language Access Trust Index is built on that premise. It is a criterion-based evaluation instrument organised around five

Proxies accumulate

trust domains and forty scored indicators, applying separate scales to evidence adequacy and to system trust so that a well-documented weak system and a poorly documented unknown system do not collapse into the same rating. Its deployment stakes classification attaches to deployments rather than to systems, on the reasoning that the same translation engine carries entirely different risk in a museum caption and in a discharge instruction. The Evaluation Standard is published and citable, deposited with a digital object identifier under a Creative Commons attribution licence, so that a buyer, a vendor or a critic can read the method and argue with it.

> quietly until an institution believes it has bought evidence when it has bought paperwork.
>
> Heather M. Grizzle

Three ratings the SLAT Index keeps separate

Rating	Attaches to	Question it answers
Evidence adequacy	The system	Is there enough evidence to judge the system at all?
System trust	The system	Does that evidence show the output is fit for use?
Deployment stakes	The deployment	How much harm does an output failure cause in this setting (a museum caption is not a discharge instruction)?

Source: Grizzle, SLAT Index Evaluation Standard v2.0 (2026)

Three commitments distinguish the instrument from a conformance exercise. Findings are not issued without the concurrence of a Deaf reviewer, on the position that the community bearing the consequence of an inadequate system holds the competence to judge adequacy and that consultation offered after a design decision is not governance. The firm holds no financial interest in any sign language artificial intelligence vendor, accepts no contingent, referral or equity payment from an evaluated party, and does not vary its fee on the basis of an adverse finding, which addresses directly the conflict problem that has dogged assurance markets in every

prior industry that built one. And the method states its own limits: a public-evidence review reports what was located and what was not, records the search locations checked whether or not material was found there, and identifies the indicators unlikely to be publicly evidenced at all, so that a reader can distinguish a failure of the vendor from a limit of the method.

Two things about the current position should be stated rather than implied. Novara Consulting Group is not an accredited certification body, and the Index is not a certification scheme; no accreditation scheme presently exists for domain-specific evaluation of signed language artificial intelligence, and it would be a misrepresentation to suggest otherwise. What the Index produces is a documented, reproducible evidence position that a procuring institution can rely on, defend in a solicitation record, and hand to counsel. Second, applying the instrument across the vendors evaluated to date has produced a consistent finding: none has reached the highest trust tier, and the reason is uniform rather than incidental. Independent validation of output, conducted by parties without a commercial interest in the result, is absent across the market. That is a finding about the field, not about any single company in it.

What a buyer should require between now and December 2027

The deferral has produced a window, and windows are used well or badly. Used badly, it becomes eighteen months of acquisition on unexamined vendor claims, with a compliance reckoning arriving after the systems are already load-bearing. Used well, it is the period in which institutions build the evidentiary habits that the later obligations will assume they already have.

The practical requirement is narrow enough to write into a solicitation. Ask what the certificate certifies, and require the vendor to state in writing whether it covers the management system or the performance of the system being sold. Require evidence generated in the proposed setting rather than in a demonstration environment, with the operating conditions named. Require the provenance of the training corpus, the licence under which it is held, and the basis on which the people recorded in it consented to that use. Require comprehension evidence from fluent users of the language in question, reported with its method, and treat an automated quality score offered in its place as an unanswered question rather than an answer. Require the



vendor to name the accountable role for an output failure. Where the deployment carries legal weight or irreversible consequence, require that a qualified human remain in the path, because on the current evidence no system in this category is ready to be the last line.

☐ The vendor states in writing whether its certificate covers the management system or the performance of the system being sold

☐ Evidence is generated in the proposed setting with operating conditions named, not in a demonstration environment

☐ Training corpus provenance, licence, and the basis of signer consent are disclosed

☐ Comprehension evidence comes from fluent users of the language, with method reported; an automated score in its place is an unanswered question

☐ The vendor names the accountable role for an output failure

☐ Where output carries legal weight or irreversible consequence, a qualified human remains in the path

None of that waits on a harmonised standard, a notified body, or a professional certification scheme that does not yet exist. It waits only on the institution deciding that assurance should be about the thing it is buying.

Cov chaw siv

Department for Science, Innovation and Technology, *Trusted Third-Party AI Assurance Roadmap* (September 2025).

BCS and techUK, AI Assurance Consortium launch materials (June 2026).

Regulation (EU) 2026/1744 (Digital Omnibus on Artificial Intelligence), published 24 July 2026, in force 27 July 2026, amending Regulation (EU) 2024/1689.

ISO/IEC 42001:2023, *Information technology, Artificial intelligence, Management system*.

ISO/IEC 42006:2025, *Requirements for bodies providing audit and certification of artificial intelligence management systems*.

Grizzle, H. M., *SLAT Index Evaluation Standard*, version 2.0, Novara Press, 24 July 2026. DOI 10.5281/zenodo.21537271.

Yuav Hais Qhov Chaw Li Cas

Grizzle, H. M. (2026, September 5). The Object of Assurance. Novara Consulting Group. https://www.novaracg.com/hmn/2026/09/05/the-object-of-assurance/