



Standard
Intelligence

White paper

The Register Nobody Has

Discovery Before Compliance, and Why the Estate Is Larger Than You Think

AUTHOR

Michael Clark
Standard Intelligence

PUBLISHED

August 2026

EDITION

Cross-Sector Paper

WEB

standardintelligence.com

AUDIENCE

For AI governance programme owners starting AI Act compliance

CONTENTS

Summary	3
1. Nearly every AI Act duty is answered per system, so the register comes first	4
2. Procurement and IT records are accurate, but they answer a different question	5
3. Three categories of AI bypass the channels compliance watches	5
3.1 Embedded AI	5
3.2 Departmental adoption	6
3.3 Pilots that became production	6
4. Wherever those routes have been used, the estate is larger than the record	6
5. Four discovery channels, started together, find what the records miss	7
5.1 Spend analysis at service level	7
5.2 Repository and pipeline scanning	7
5.3 Suppliers and the procurement pipeline	8
5.4 The structured survey	8
6. The survey should ask about function, not technology	8
7. Count one register entry per distinct purpose and obligation set	9
8. The fields that turn a list into a compliance record	10
9. Every entry needs a named owner who has accepted ownership in writing	11
10. Discovery ends on a defined test, and intake then continues	11
Where to start, and what it costs	12
Basis and Standing	13
Sources	14

Summary

Most of the work the AI Act requires is answered one system at a time. Classification, the role the organisation holds, the conformity route, the oversight assignment, the registration duty and the technical file all attach to a particular system, and a system nobody has listed cannot be classified, assessed or evidenced. That makes the register, a per-system record of the facts those duties turn on, the first substantive step of an AI Act programme rather than an administrative preliminary to it.

A few duties are organisational, and one set attaches to models rather than systems. The Article 4 duty to take measures supporting AI literacy and a provider's quality management system under Article 17 do not depend on a per-system analysis, and Chapter V places obligations on providers of general-purpose AI models. A register should provide for models as well as systems, but what sets a programme's cost, timing and exposure still resolves per system.

Organisations know what they bought, which is less than what they run. Procurement records and IT asset lists are usually accurate, but they answer questions about purchase and installation, and the register asks about operation. Three categories of AI fall outside both because none passes through the channels compliance already watches: AI embedded in software bought for another purpose, tooling adopted by departments below the threshold that triggers review, and pilots that became production without anyone marking the change.

Wherever any of those routes has been used, the estate is larger than the record. The estate is every AI system and AI-enabled component the organisation uses, including those inside purchased platforms. That conclusion follows from how the three categories arise, not from published surveys, which are vendor and analyst work rather than peer-reviewed research and should not be relied on alone to size an organisation's gap.

Four discovery channels, started together, find what the records miss. Spend analysis at service level, repository and pipeline scanning, a written question to each major supplier alongside a watch on the procurement pipeline, and a structured survey each find something the others do not. The survey depends on asking about function and consequence rather than technology: departments that built something rarely describe it as AI and, asked whether they use AI systems, answer honestly and inaccurately.

The result is a compliance record, not an inventory. Each entry covers one distinct purpose and obligation set, and carries the purpose, where the system operates, where the people affected by it are, the role the organisation holds and a named owner who has accepted ownership in writing. An IT asset list carries few of these fields, so extending one saves less work than it appears to.

Discovery ends on a defined test, and intake then continues. The initial discovery is complete when two consecutive cycles find no new systems above a defined materiality threshold and every entry has an owner who has accepted ownership. Procurement, change control and periodic re-runs keep the register current after that.

Deferring discovery is mainly a planning cost, and the first step costs a letter.

No provision requires a register in this form, and the high-risk obligations apply from 2 December 2027 (Annex III) and 2 August 2028 (Annex I), subject to Article 111(2) for systems already on the market. As at 11 August 2026, however, the literacy duty, the Article 5 prohibitions and the Article 50 transparency duties already apply, including to systems the organisation has not listed. A programme budgeted before discovery is budgeted against a guess that is likely to be low, and the written question to suppliers can go this week.

Basis and Standing, at the end, sets out which propositions are binding law and which are our own recommended controls.

1. Nearly every AI Act duty is answered per system, so the register comes first

An AI governance programme works through a sequence: classify each system, determine which obligations bind it, design controls that satisfy them, produce evidence that the controls operated, and monitor for change. Every step takes a system as its input, and so do the conformity route, the oversight assignment, the registration duty and the technical file.

The Act also allocates most duties by the role a party holds for a particular system rather than to the organisation at large. The two roles that matter most here are the provider, which develops a system (or has it developed) and places it on the market or puts it into service under its own name or trademark, and the deployer, which uses a system under its authority (Article 3(3) and (4)).

Two kinds of duty do not attach to a single system. First, some duties are organisational. Article 4, as amended by Regulation (EU) 2026/1744, requires providers and deployers to take measures to support the AI literacy of their staff and others operating systems on their behalf, and Article 17 requires a provider of high-risk systems to put a quality management system in place.

Second, Chapter V places obligations on providers of general-purpose AI models, which are models rather than systems. A register built only around deployed systems will not show where the organisation holds those duties unless its schema provides for models.

Neither moves the starting point. The organisational duties cannot be discharged without knowing what the estate contains, since Article 4 takes into account “the context the AI systems are to be used in” and “the persons or groups of persons on whom the AI systems are to be used”. What sets the programme’s cost, timing and exposure resolves per system.

The register is the input to classification, not a workstream alongside it. A programme that starts classifying before the register is complete is classifying a sample of unknown size and unknown representativeness.

Completeness is also an assurance question. An organisation can show sound governance over every system in its register and still be exposed, because the exposure

sits in the systems that are not in it. Coverage of the known estate says nothing about coverage of the whole estate, and internal audit that tests the controls without testing whether the population is complete has left the larger question untested.

An authority opening an inquiry is likely to ask early for the organisation's list of systems, because that list establishes what the inquiry covers.

2. Procurement and IT records are accurate, but they answer a different question

Procurement records are usually accurate, and organisations that rely on them are not being careless. An organisation with mature procurement knows what it contracted for, from whom, at what cost and on what terms, and can produce that list quickly. The difficulty is that the list answers a question about purchasing, and the register asks a question about operation.

Three things separate the two. Software is often bought for one capability and delivers several, of which only one was specified; departments acquire tooling below the value that triggers central involvement; and systems change what they do over time without changing what they cost.

An IT asset register has the same limitation from a different direction. It knows what is installed and licensed, which is closer to operation, but almost nothing about purpose or about who is affected. Neither list is wrong; each answers the question it was built to answer.

3. Three categories of AI bypass the channels compliance watches

The gaps in section 2 leave three categories of AI that the existing records do not show.

3.1 Embedded AI

The hardest category to find is AI inside software bought for something else. A case management system with a triage module, a service desk that routes and categorises tickets, a finance platform with anomaly detection, a human resources system with a screening or ranking feature, a customer relationship platform with propensity scoring and a document system that classifies and redacts are all examples.

None of these was bought as AI, and none appears as AI in the procurement record. The capability may have arrived in a release note, switched on by default, months or years after the original purchase. Where the module analyses and filters job applications, the

platform now contains a use that Annex III lists as high-risk (point 4(a)), although the contract still describes the same platform.

The capability may therefore be invisible from the contract and from usage alike, because the module presents as a feature of a system the user already knows.

The route in is the supplier rather than the software. Many enterprise platform suppliers publish documentation of their AI features, and a written question to each major supplier about AI functionality in the licensed modules (the supplier letter, section 5.3) is likely to surface more than an internal search.

3.2 Departmental adoption

The second category is tooling a team adopts below the price that triggers procurement review, often on a corporate card or an existing platform subscription. A marketing team subscribes to a content tool; a recruitment team adds a screening add-on to the applicant tracking system; an analyst builds something on a general-purpose model with an API key issued for a different project; a service manager finds that a tool the organisation already licenses can be pointed at a new problem.

The last of these is the most consequential and the least visible, because nothing was acquired. Something was repurposed, and changing a system's intended purpose (the use for which its provider intended it) so that it becomes high-risk is one of the routes by which an organisation becomes its provider under Article 25(1)(c), once Article 25 applies (section 8).

3.3 Pilots that became production

The third category involves no acquisition and no repurposing: a system approved as an experiment that quietly became something the organisation depends on. Nobody marks the transition, because there is no single moment to mark. Reliance grows gradually, the pilot governance that suited an experiment stays in place, and the review that would have applied to a production deployment never happens because, on paper, nothing has been deployed.

The signal to look for is dependency rather than status: how long the system has been running, how many people use it, and what would happen if it stopped tomorrow. A pilot nobody would miss is a pilot; a pilot whose failure would interrupt a service is production with an inaccurate label.

4. Wherever those routes have been used, the estate is larger than the record

Expect a first discovery to find more than the existing records hold. The reason lies in how the three categories arise, not in published survey figures.

Embedded capability arrives through a procurement decision made about a platform, not about AI. Departmental adoption sits below the review threshold or, where a tool is repurposed, leaves no procurement trace at all, and a pilot becomes production without

any decision that triggers review. None of the three passes through a channel compliance watches, so wherever any of them has been used the estate is larger than the record.

Published figures on the size of the gap come from vendor and analyst surveys rather than peer-reviewed research. Those that measure employee adoption of generative tools count at most part of the departmental category, and none of the embedded capability or repurposing that sits on top of it.

The reliable version of the claim is directional: expect the estate to exceed the record, expect the excess to sit in categories nobody has counted, and do not size a programme from any published multiple. The supplier letter tests the claim in the organisation's own estate; it can be sent this week, and the replies answer it for the embedded category in that organisation rather than in a survey population.

5. Four discovery channels, started together, find what the records miss

Each category is found by a different means, so no channel substitutes for another and none needs to wait for another to finish. Start all four together, within the first month. The quicker channels report within weeks, and what they find can sharpen the survey's follow-up questions while it is still running.

5.1 Spend analysis at service level

Read financial data at the level of the service rather than the supplier. A single supplier line covers a platform; the service detail beneath it shows which capabilities are being consumed.

Three signals repay attention:

- inference and model API charges, which indicate systems built internally on external models, often by teams outside engineering;
- consumption-based charges that grew without a contract change, which often indicate a capability switched on;
- subscriptions at values low enough to have avoided procurement, which map to departmental adoption.

This channel is cheap, available immediately and needs no one's cooperation, which makes it, with the supplier letter, a natural first task.

5.2 Repository and pipeline scanning

Version control and deployment pipelines show what was built internally. Search for model and inference library imports, API client configuration, model artefacts, and training or evaluation code.

The value lies in the date as well as the finding, because repositories carry history and this channel can show when a capability was introduced or changed. That bears on Article 25, since the role follows the conduct. It also bears on Article 111(2), under which a high-

risk system placed on the market or put into service before the high-risk rules apply (2 December 2027 for Annex III) is caught only if its design changes significantly from that date, although systems intended for use by public authorities must comply by 2 August 2030 in any case.

5.3 Suppliers and the procurement pipeline

The supplier letter looks back. It is one written question to each major platform supplier asking which licensed modules score, rank, classify, predict or generate, and which were enabled by default, and it is the quickest route to the embedded category (section 3.1).

The procurement pipeline looks forward: what is in evaluation, in negotiation, or approved and not yet delivered. A register that only looks backwards is out of date on the day it is completed. Building the intake process during the initial discovery costs little.

The pipeline is also the natural control point. A procurement gate that asks the register's questions at the point of purchase keeps the register current for everything acquired after it is built.

5.4 The structured survey

The survey finds what the other three channels miss, and it is the channel most often designed badly. Its main targets are repurposing, which leaves no financial or technical trace, and departmental adoption that the spend data does not show, such as a capability added to an existing subscription. How much it finds depends on the wording, which section 6 sets out.

6. The survey should ask about function, not technology

Asking whether a department uses AI systems produces a small number of accurate answers and misses much of what the department actually uses. The reason is not evasion. People who adopted these tools classify them by what they do (a scheduling tool, a scoring model, a triage assistant, a way of getting the first draft written), and someone who built a scoring spreadsheet that calls a model API answers no in good faith.

The question that works asks about function and consequence:

- whether any process in the team now has software deciding, drafting, ranking, scoring, prioritising, matching, flagging or recommending something a person used to do;
- what would be affected if that software gave a wrong answer;
- whether anyone outside the organisation is affected by its output.

Three further design points matter.

Ask about tools rather than systems. “System” implies something formal and installed; “tool” captures the subscription, the spreadsheet with an API call and the vendor module.

Ask what changed recently. A question about new capabilities in existing software surfaces the switched-on module that nobody bought.

Do not run it as a compliance exercise. A survey that reads as an audit produces defensive answers and slow returns. One sent as an attempt to build a complete picture, with an assurance that disclosure is not a finding, is far more likely to produce the estate.

The first pass is the cheapest point at which to get honest answers. An organisation that treats early disclosures as failures is unlikely to get candid answers next time, and will find the remaining systems later and at greater cost. The assurance covers the disclosure, not the system: a disclosed system that engages one of the Article 5 prohibitions, which have applied since 2 February 2025, must still be withdrawn from use.

7. Count one register entry per distinct purpose and obligation set

Discovery produces a list of things. Turning it into a register requires deciding what counts as one entry, and the decision has consequences.

The working rule is one entry per distinct purpose and obligation set, meaning the set of duties that attach to that use. A platform with four capabilities serving different purposes is four entries, because each may classify differently and carry different obligations.

The temptation runs the other way, towards one entry per product or supplier, because it produces a shorter list. A single entry for a customer service platform, classified on the average of its parts, can hide a high-risk component inside a low-risk classification where later reviews are unlikely to find it.

Splitting is the cheaper error to correct. Two entries that turn out to share a classification and an obligation set can be combined in an afternoon; one entry that turns out to contain two systems has to be unpicked from every downstream artefact that referenced it.

Three boundary cases arise repeatedly.

One model, several uses. A single model serving three business processes is three entries, because purpose drives classification and the purposes differ.

A pipeline of components. Several models in sequence producing one output are usually one entry, defined by the output and its purpose, unless a component is separately deployed elsewhere.

A platform with modules. One entry per module with a distinct purpose; the platform itself is not the unit.

8. The fields that turn a list into a compliance record

The field list is where the difference between an inventory and a register becomes concrete.

Identity and ownership. System name, internal reference, business unit and a named accountable owner (section 9).

Purpose. What the system is for and what it outputs, described in terms a non-technical reader can evaluate. Classification depends on this field, and it is the one most often written too vaguely to use.

Deployment geography. Where the system operates.

Geography of affected persons. Where the people affected by its outputs are located. In our experience this field is often omitted, and it can change the answer: the Act reaches providers and deployers established outside the Union “where the output produced by the AI system is used in the Union” (Article 2(1)(c)), which usually depends on where the affected persons are rather than where the system runs.

Role. Which role the organisation holds for this system: provider, deployer, importer or distributor.

The role field is the hardest to populate and the most consequential. The role follows conduct rather than the procurement decision. Under Article 25(1), a deployer, distributor, importer or other third party becomes the provider of a high-risk system if it puts its name or trademark on a high-risk system already on the market or in service, substantially modifies such a system so that it remains high-risk, or changes the intended purpose of a system so that it becomes high-risk.

None of these requires a decision to change role, and once Article 25 applies (from 2 December 2027 for Annex III systems) each takes effect when the conduct occurs. A role taken from the supplier contract describes the intended arrangement rather than the operating position. Where the register cannot answer the question, it has identified work rather than avoided it.

Data categories. What the system processes, with particular attention to special category data and data concerning children.

Autonomy characteristics. Whether the system acts without a human decision at each step, and what it is authorised to do.

Component manifest reference. A pointer to the record of what the system is built from (models, libraries and data sources) and which versions.

Lifecycle state. In evaluation, in pilot, in production or in wind-down, with the date of each change.

An IT asset list carries name, supplier, cost, licence count and installation state, and few of the fields above. Extending it therefore saves less work than it appears to, which is

worth establishing before a programme is scoped on the assumption that the existing inventory is a starting point.

9. Every entry needs a named owner who has accepted ownership in writing

Every entry needs a named accountable owner, and the name has to be a person rather than a function. When an authority asks why a system was classified as it was, or why a control was not operating, the answer comes from an individual who either can explain it or cannot. An entry attributed to a department has no such individual, so at the moment it matters nobody is responsible.

In our experience, acceptance is the step organisations skip. Assigning ownership is an act by the compliance function; accepting it is an act by the owner, and accountability depends far more on the second.

The practical form is minimal: the owner confirms in writing that they own the system, that the register entry is accurate, and that they will notify defined changes. The confirmation's main effect is behavioural rather than legal. An owner who has confirmed a purpose description is more likely to notice when the purpose drifts, because they have read it and know what it says.

An owner who declines is a finding worth having early. It usually means the system is orphaned, which is a governance problem in its own right and common in the pilot category (section 3.3).

10. Discovery ends on a defined test, and intake then continues

The initial discovery is complete when two consecutive cycles surface no new systems above a defined materiality threshold, and every entry has an owner who has accepted ownership in writing. Two cycles are needed because a single clean pass cannot distinguish a complete register from an ineffective pass.

Define materiality in advance and in terms of consequence rather than cost, since the cheapest tools in the estate can affect the most people.

After that, the register is a live record rather than a project output, and it needs three maintenance routes:

- a procurement gate that captures anything acquired (section 5.3);
- a change process that captures modification, repurposing and rebranding, the events that can change the role (section 8);

- a re-run of the discovery channels at least annually, because departmental adoption does not stop.
-

Where to start, and what it costs

The register comes first because the rest of the programme is scoped from it, and the first step costs a letter.

This week, one person: the supplier letter. One written question to each major platform supplier, with the replies retained. It needs no budget, approval or tooling, and it is likely to be the most productive single step in the programme.

This month, one person part-time: spend analysis and repository scanning. Both use data the organisation already holds. Launch the survey in the same month, since it takes longest to return.

A quarter, a small team: the survey, the boundary decisions and the first population of the register fields. The constraint is the goodwill of the functions being surveyed, which is why it should not be run as an audit (section 6).

A funded programme, only if the estate justifies it: tooling. Most organisations should complete a first discovery in a spreadsheet and decide about tooling once they know the size of what they are tooling for.

Ongoing: intake and re-runs. The procurement gate, the change process and the annual re-run in section 10 keep the register current.

On timing. As at 11 August 2026 no fine under the AI Act had been confirmed by any regulator or official source. The high-risk obligations apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems (Article 113, as amended by Regulation (EU) 2026/1744), subject to Article 111(2) for systems already on the market (section 5.2). No provision requires a register in the form described here, so an organisation without one is not, for that reason alone, in breach.

Some duties already apply, however: the Article 4 literacy duty, the Article 5 prohibitions and the Article 50 transparency duties (applying from 2 August 2026, with a transitional period for some existing systems), together with any sectoral obligations attaching to the same systems. Whether the organisation meets them for systems it has not listed is a question the register exists to answer.

The larger cost of deferral falls on planning. Every downstream activity (classification, role determination, conformity work, oversight assignment, contractual remediation) is scoped by the population, so a programme budgeted before discovery is budgeted against a guess, and for the reasons in section 4 the guess is likely to be low. A programme that discovers its estate late is likely to rebuild its plan and lose the time in which the work could have been done without pressure.

An organisation that has run discovery can say how many systems it runs, what each is for, which role it holds for each and who owns it. Most other AI Act questions it will be asked start from those answers.

This paper describes a discovery and register method, states the position as at 11 August 2026 and is not legal advice. Role determination is summarised only as far as the register depends on it; a full determination for any system requires the facts of that deployment and, where consequential, advice.

Basis and Standing

Binding law. The AI Act (Regulation (EU) 2024/1689) as amended by Regulation (EU) 2026/1744, in force 27 July 2026:

- Article 2(1) on territorial scope, and Article 3(3), (4) and (12) on provider, deployer and intended purpose;
- Article 4 on AI literacy, applying since 2 February 2025 and in amended form from 27 July 2026, and the Article 5 prohibitions, applying since 2 February 2025;
- Article 17 on providers' quality management systems, Article 25(1) on when another party becomes the provider of a high-risk system, Article 26 on deployers' duties and Annex III point 4(a);
- Article 50 transparency duties, applying from 2 August 2026, with a transitional period for some existing systems under Article 11(4);
- Chapter V obligations of providers of general-purpose AI models, applying since 2 August 2025;
- Article 11(2) on high-risk systems placed on the market or put into service before the high-risk rules apply;
- Article 113, under which the high-risk obligations, including Article 25, apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems.

What the Act does and does not require. Article 49 requires providers of Annex III high-risk systems, providers that conclude under Article 6(3) that an Annex III system is not high-risk, and deployers that are public authorities or act on their behalf, to register in the EU database referred to in Article 71; Annex III point 2 systems are registered at national level. No provision requires an enterprise-wide inventory of every AI system or prescribes its fields. A register is the practical means of knowing which duties apply to which system, and this paper argues for it on that basis rather than as a statutory formality.

Regulator guidance. None is relied on.

Recommended controls, not regulatory requirements. The supplier letter, the four discovery channels, the survey design, the boundary rule and boundary cases, the register fields, the named-owner acceptance, the exit test and materiality threshold, and the maintenance routes are our design, not prescribed formalities. Several fields reflect legal provisions: the purpose field reflects classification under Article 6 and Annex III, the two geography fields reflect Article 2(1), the role field reflects Articles 3(3), 3(4) and 25(1), and the lifecycle dates reflect Article 11(2).

Research, and what it does and does not support. The paper cites no survey figures for the size of the gap. Its claim that the estate is larger than the record rests on the structural argument in section 4, which shows where the excess will arise but not how large it is; vendor and analyst surveys are not peer-reviewed and should not be relied on alone to size it.

Do not size the programme from a published multiple. Size it from the organisation's own discovery.

Enforcement position. As at 11 August 2026, no fine under the AI Act had been confirmed by any regulator or official source.

Sources

1. Regulation (EU) 2024/1689 (AI Act), consolidated version of 27 July 2026, Articles 2, 3, 4, 5, 6, 17, 25, 26, 49, 50, 71, 111 and 113, Chapter V and Annex III. <https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng>
2. Regulation (EU) 2026/1744 (Digital Omnibus on AI), amending among other provisions Articles 4, 111 and 113. <https://eur-lex.europa.eu/eli/reg/2026/1744/oj>

Provisions are cited by article, rule or paragraph in the text above.



Standard
Intelligence

GET IN TOUCH

Find the AI your organisation is already running.

EMAIL

hello@standardintelligence.com

WEB

standardintelligence.com

LOCATION

Manchester, United Kingdom

AI Use: AI tools were used in research, editorial review and drafting support. All sources were retrieved and checked against the published record by the author, who holds editorial responsibility for the contents and for the conclusions drawn. This disclosure is made voluntarily. To the extent the Article 50(4) obligation in Regulation (EU) 2024/1689 applies to this paper as text published to inform the public on matters of public interest, the exception for content that has undergone human review and for which a person holds editorial responsibility applies.