Reliable Context

RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

EU AI Act Article 12: Record-Keeping and Logging Explained

> Reliable Context is a kernel of truth: a system of record for organizations that fuse private and public information and must defend every value they act on. One canonical record per real thing, a cited source and an observation date behind every value, unknowns declared rather than left blank, and any past state of the record reproducible.

Published by Reliable Context · 2026-10-29 · eu ai act article 12 · ai act record keeping · ai act logging requirements · ai act automatic logs high-risk · ai act article 19 logs retention

Original article: https://reliablecontext.com/articles/eu-ai-act-article-12-record-keeping-logs



In this article

1. The short answer
2. Scope: which systems Article 12 covers
3. Article 12(1) and (2): automatic recording and its three purposes
4. Article 12(3): the one place the Act lists fields
5. Article 19: provider retention
6. Article 26: deployer retention and monitoring
7. When it applies: the Digital Omnibus dates
8. A worked example: one log event
9. Logs are not source records
10. Design checklist
11. Summary

The short answer

Article 12 of the EU Artificial Intelligence Act (Regulation (EU) 2024/1689) is a design requirement, not a retention rule. It says that a high-risk AI system must be *technically capable* of recording events automatically over its whole lifetime, and it names what those records are for. Who must keep the resulting logs, and for how long, is set by two other provisions: Article 19 for providers and Article 26(6) for deployers. Both set a floor of at least six months, subject to other Union or national law. Under the AI Act as amended by the Digital Omnibus on AI, these obligations apply to Annex III high-risk systems from 2 December 2027 and to Annex I product-embedded systems from 2 August 2028 (Article 113, consolidated text).

This article reads the three provisions together, shows what a log record built to satisfy them could look like, and explains why a log is evidence that the system ran, not evidence that the facts it used were true. It is an educational reading of EU law, not legal advice.

Scope: which systems Article 12 covers

Article 12 sits in Chapter III, Section 2 of the Act, the section titled "Requirements for high-risk AI systems". It therefore applies only to systems classified as high-risk under Article 6, which has two routes (Regulation (EU) 2024/1689, Article 6):

- **Article 6(1), Annex I:** the AI system is a safety component of a product, or is itself a product, covered by the Union harmonisation legislation listed in Annex I (machinery, medical devices and so on), and that product must undergo third-party conformity assessment.
- **Article 6(2), Annex III:** the AI system falls into one of the use areas listed in Annex III, such as biometrics, critical infrastructure, education, employment, access to essential services, law enforcement, migration and the administration of justice.

A general-purpose chatbot used internally, a document summarizer or a coding assistant is not high-risk by default. It becomes relevant to Article 12 only if it is placed on the market or used for an Annex III purpose or as an Annex I safety component.

Two roles matter. A **provider** is the person or body that develops an AI system, or has one developed, and places it on the market or puts it into service under its own name or trademark. A **deployer** is a person or body using an AI system under its authority, other than for personal non-professional activity (Article 3(3) and 3(4)). A bank buying a credit-scoring system is a deployer; the vendor is the provider. A bank that builds its own scoring system and puts it into service is both.

Article 12(1) and (2): automatic recording and its three purposes

Paragraph 1 is a single sentence: "High-risk AI systems shall technically allow for the automatic recording of events (logs) over the lifetime of the system" (Article 12). Three words carry the weight.

- **Technically allow.** The obligation is on the system's capability. It is satisfied, or not, at design time. A deployer who switches logging off does not change whether the system "allows" it, but Article 26 then creates the deployer's own problem (see below).
- **Automatic.** The recording must not depend on an operator remembering to write something down.
- **Lifetime.** Logging is not a launch-phase or validation-phase activity. It must be possible for as long as the system is in use.

Paragraph 2 explains why the capability exists. Logging must ensure "a level of traceability of the functioning of a high-risk AI system that is appropriate to the intended purpose of the system", and the capabilities must enable the recording of events relevant for three things (Article 12(2)):

Point	Purpose	Cross-reference	What it implies for event selection
12(2)(a)	Identifying situations that may result in the system presenting a risk, or in a substantial modification	Article 79(1); Article 3(23)	Events that show drift, out-of-range inputs, failures, version changes
12(2)(b)	Facilitating post-market monitoring	Article 72	Events that let the provider evaluate continuous compliance over the lifetime
12(2)(c)	Monitoring the operation of the system by deployers	Article 26(5)	Events a deployer needs to detect risk or serious incidents in use

The Act does not publish a list of mandatory fields for the general case. The phrase "appropriate to the intended purpose" leaves the event schema to the provider, who must document it. Article 13(3)(f) requires the instructions for use to include, where relevant, "a description of the mechanisms included within the high-risk AI system that allows deployers to properly collect, store and interpret the logs in accordance with Article 12" (Article 13). In other words, a provider's logging design is something a deployer should be able to read about, not reverse-engineer.

Recital 71 places logging inside a broader traceability objective: keeping records and technical documentation so that compliance can be verified and operation and post-market monitoring can be followed, and it repeats that high-risk systems "should technically allow for the automatic recording of events, by means of logs, over the duration of the lifetime of the system" (Recital 71).

Article 12(3): the one place the Act lists fields

For remote biometric identification systems, point 1(a) of Annex III, Article 12(3) sets a minimum. Logging must provide at least (Article 12(3)):

1. the period of each use, with start date and time and end date and time;
2. the reference database against which input data has been checked;
3. the input data for which the search has led to a match;
4. the identification of the natural persons involved in verifying the results, as referred to in Article 14(5).

Article 14(5) in turn requires that, for these systems, no action or decision is taken on an identification unless it has been separately verified and confirmed by at least two natural persons with the necessary competence, training and authority, with a carve-out for law enforcement, migration, border control or asylum where Union or national law considers the requirement disproportionate (Article 14(5)).

The list is narrow, but it is instructive for every other high-risk system because it shows the shape regulators had in mind: *when* the system was used, *against what reference data*, *on which input*, and *which humans* confirmed the output. Those are the four questions most audits of an automated decision eventually ask.

Article 19: provider retention

Article 19(1) turns capability into custody. Providers "shall keep the logs referred to in Article 12(1), automatically generated by their high-risk AI systems, to the extent such logs are under their control". The period is "appropriate to the intended purpose of the high-risk AI system, of at least six months, unless provided otherwise in the applicable Union or national law, in particular in Union law on the protection of personal data" (Article 19).

Four features of that wording are worth separating:

- **"To the extent such logs are under their control."** A provider that ships software running entirely on the deployer's infrastructure may not hold the logs at all. The obligation follows control, not authorship.
- **"At least six months."** This is a floor. The ceiling is "appropriate to the intended purpose". A system whose decisions can be contested years later may need far longer.
- **"Unless provided otherwise ... in particular in Union law on the protection of personal data."** Logs that contain personal data are also subject to the GDPR's storage-limitation and data-minimisation principles. The AI Act defers to that law rather than overriding it.
- **Financial institutions.** Article 19(2) says providers that are financial institutions subject to internal-governance requirements under Union financial services law keep the logs as part of the documentation kept under that law.

Article 16(e) makes keeping logs, when under the provider's control, one of the listed provider obligations, and Article 21(2) requires providers, on a reasoned request by a competent authority, to give access to the automatically generated logs to the extent they control them (Articles 16 and 21). Logs are therefore something a provider may be asked to produce, not only to retain.

Compare Article 18, which requires providers to keep technical documentation, quality-management documentation and the EU declaration of conformity for ten years after the system is placed on the market or put into service (Article 18). The Act thus distinguishes documentation about the system (ten years) from logs of its operation (at least six months).

Article 26: deployer retention and monitoring

Deployers have a parallel duty. Article 26(6) says deployers "shall keep the logs automatically generated by that high-risk AI system to the extent such logs are under their control, for a period appropriate to the intended purpose of the high-risk AI system, of at least six months, unless provided otherwise in applicable Union or national law, in particular in Union law on the protection of personal data", with the same financial-institution rule as Article 19(2) (Article 26).

Article 26(5) explains what the logs are for on the deployer side. Deployers must monitor operation on the basis of the instructions for use, inform providers where relevant under Article 72, and, if they have reason to consider that use may result in a risk within the meaning of Article 79(1), inform the provider or distributor and the market surveillance authority without undue delay and suspend use. Serious incidents must be reported immediately (Article 26(5)). A deployer cannot detect any of this without the events Article 12(2)(c) exists to capture.

Article 99(4) provides for administrative fines of up to EUR 15 million or 3% of total worldwide annual turnover, whichever is higher, for non-compliance with provider obligations under Article 16 and deployer obligations under Article 26, with the lower of the two caps applying to SMEs under Article 99(6) (Article 99).

Question	Provider	Deployer
Must the system be able to log?	Yes, Article 12 (design requirement, Article 16(a))	Not a deployer duty; relies on the provider's design
Must it keep logs?	Yes, Article 19(1) and 16(e), if under its control	Yes, Article 26(6), if under its control
Minimum period	At least six months, appropriate to intended purpose	At least six months, appropriate to intended purpose
Overriding law	Union or national law, especially data protection	Union or national law, especially data protection
Produce on request	Article 21(2), to a competent authority	Article 26(12), cooperate with competent authorities
What the logs are used for	Post-market monitoring (Article 72), risk identification	Operational monitoring and incident reporting (Article 26(5))

When it applies: the Digital Omnibus dates

The AI Act entered into force on 1 August 2024 and has staggered application dates. The Digital Omnibus on AI, Regulation (EU) 2026/1744 of 8 July 2026, was published in the Official Journal on 24 July 2026 and entered into force on the third day after publication, 27 July 2026 (Regulation (EU) 2026/1744, Article 4). It replaced point (c) of the third paragraph of Article 113. The Commission's AI Act Service Desk displays the consolidated text "as at 27 July 2026" (Article 113).

Provision	Applies from
General date of application of the Act	2 August 2026
Chapter III Sections 1-3 (including Articles 12, 19, 26) for Annex III high-risk systems (Article 6(2))	2 December 2027
Chapter III Sections 1-3 for Annex I high-risk systems (Article 6(1))	2 August 2028
Articles 102 to 110 (amendments to other Union acts)	27 July 2026

The Omnibus did not change the wording of Articles 12, 19 or 26: the Service Desk marks no passage of Article 12 as new or amended. It also amended Article 111(2): high-risk systems placed on the market or put into service before the Chapter III application date are covered only if they are subject to significant changes in their design from that date, while providers and deployers of high-risk systems intended for use by public authorities must comply by 2 August 2030 (Regulation (EU) 2026/1744).

On technical standards: the international standard ISO/IEC 24970, *Artificial intelligence — AI system logging*, describes "common capabilities, requirements and a supporting information model for logging of events in AI systems". On 28 October 2026 the ISO catalogue listed it at stage 60.00, "International Standard under publication", reached on 24 October 2026 (ISO/IEC 24970). A standard gives a presumption of conformity with AI Act requirements only once it is a harmonised standard referenced in the Official Journal. ISO/IEC 24970 is not; as of 28 October 2026 the only such reference is EN 18286:2026, which covers the quality management system of Article 17, not logging (European Commission, 19 October 2026).

A worked example: one log event

The Act does not prescribe a format. The record below shows one way to satisfy the three purposes in Article 12(2) for a hypothetical Annex III employment-screening system deployed by **Halden Therapeutics AG, a fictional company**. Every value is invented.

```
{
  "event_id": "evt_01JD7Q9C4ZB1",
  "event_type": "inference.completed",
  "system": {"name": "candidate-screening", "provider_version": "3.4.1",
             "model_id": "rank-2026-09-30", "config_hash": "sha256:9f2c...e1"},
  "deployer": "Halden Therapeutics AG (fictional)",
  "use_period": {"start": "2027-12-14T09:02:11Z", "end": "2027-12-14T09:02:13Z"},
  "input_ref": {"type": "application", "id": "app_88412",
                "content_hash": "sha256:41ab...7d", "stored_by": "deployer"},
```

```
  "reference_data": [{"name": "role-requirements", "version": "2027-11-30"}],
  "output": {"score": 0.62, "band": "review", "threshold_version": "t-12"},
  "human_oversight": {"reviewer_id": "hr-0193", "action": "override",
                      "final_band": "advance", "reason_code": "missing-credit-for-equivalent-degree",
                      "at": "2027-12-14T11:40:05Z"},
  "anomalies": ["input_language_not_in_validated_set"],
  "retention_class": "ai-act-art26-min-6m"
}
```

How each block maps to the text:

- `system.provider_version`, `model_id` and `config_hash` let the provider tell whether a change amounted to a substantial modification, which is purpose 12(2)(a).
- `anomalies` records a situation that may indicate risk, also 12(2)(a), and feeds post-market monitoring, 12(2)(b).
- `human_oversight` records who intervened and why, which is what a deployer needs for 12(2)(c) and Article 26(5). It mirrors the Article 12(3)(d) idea of identifying the people who verified results, although that point is mandatory only for remote biometric identification.
- `input_ref` stores a reference and a hash rather than the application itself. That keeps personal data in the deployer's governed store and lets the log be retained on its own schedule, which matters because Articles 19 and 26(6) defer to data protection law.

Logs are not source records

A log answers the question "what did the system do, when, with which version, and who intervened". It does not answer "were the facts the system relied on true, and where did they come from". In the example above, `reference_data` names a version of the role requirements. If that table said a degree was required when the published job description said "or equivalent experience", the log faithfully records a decision made on a wrong premise.

Property	Event log (Article 12)	Source record for a fact
Unit	An event in the system's operation	A value asserted about a real thing
Answers	When, which version, which input, which human	What the value is, where it was read, when, in what words
Typical retention driver	AI Act floor of six months plus purpose	The life of the decision it supports, plus sector rules
Can show the system ran as designed	Yes	No
Can show an input fact was correct	No	Yes, if the source is retained and cited

For audit purposes the two need to be joinable. If a log entry carries a stable identifier and version for every reference fact or document it used, an investigator can move from "the system ranked this candidate at 09:02" to "on the basis of these requirements, as they read on 30 November". Without that join, an Article 12 log can be complete and still leave the most important question open. The site's audit-trail glossary entry and the handoff fault-injection test treat the operational side of the same problem.

Design checklist

The Act sets purposes rather than fields. The following list is a practical reading of those purposes, not a list of legal requirements.

1. **Event catalogue.** Document which events are logged and map each to 12(2)(a), (b) or (c). Put the description in the instructions for use (Article 13(3)(f)).
2. **Time.** Record start and end timestamps in UTC with a declared clock source; Article 12(3)(a) shows the expected granularity.
3. **Version identity.** Log model, configuration and threshold versions on every event so that changes and possible substantial modifications are visible.
4. **Input reference.** Store a stable identifier and a content hash for each input; keep personal data out of the log where a reference will do.
5. **Reference data.** Identify every reference dataset or fact used, by version, so that the log can be joined to the source record.
6. **Human oversight.** Record who reviewed, confirmed or overrode an output, when, and on what stated ground.
7. **Anomalies.** Log out-of-distribution inputs, errors, fallbacks and refusals; these are the events post-market monitoring needs.
8. **Control and custody.** Decide, in the contract, which party controls which logs, since both Article 19 and Article 26(6) attach the duty to control.
9. **Retention.** Set a period of at least six months, justified by intended purpose, reconciled with GDPR storage limitation and any sector rule (for financial institutions, the documentation regime named in Articles 19(2) and 26(6)).
10. **Integrity and production.** Make logs append-only or tamper-evident and exportable, so they can be produced to a competent authority under Article 21(2) without reconstruction.

Summary

Article 12 requires high-risk AI systems to be built so that their operation leaves a lifetime trail fit for three named purposes. Articles 19 and 26(6) make providers and deployers keep that trail, for at least six months and as long as the purpose requires, within the limits of data protection law. Since the Digital Omnibus, those duties apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems. A compliant log still proves only that the system ran; the facts it ran on need their own provenance. The site's guarantees page describes that second layer as a cited, verbatim, dated source behind every value.

This article is an educational summary of EU law as of 28 October 2026 and is not legal advice.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-113
2. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
3. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-12
4. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-19
5. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-26
6. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
7. https://www.iso.org/standard/88723.html
8. https://digital-strategy.ec.europa.eu/en/library/first-harmonised-standard-supporting-ai-act-published-official-journal

THE PUBLISHER

About Reliable Context

Reliable Context is a system of record for organizations that fuse private and public information and must defend every value they act on.

The guarantees

One canonical record per real thing: duplicates resolve by rule, ambiguity is shown rather than hidden, and related organizations relate without ever being merged. A verbatim, cited source and an observation date sit behind every value. Unknowns are declared as not found, blocked or not applicable with a reason, never left as an empty cell. Any past state of the record is reproducible, together with what changed since.

Where it runs

It runs inside the customer's perimeter, one database per organization. Permitted sources, model routing, reader clearance and retention are set as the customer's own policies, and every value is exportable with its chain of custody for an auditor or regulator.

What it is not

It is not a vector database or a RAG framework, not a data catalog, MDM or data-quality suite, not a CRM or data-enrichment vendor, and not a chatbot or AI assistant. Where a reader needs one of those, the honest answer is to say so and point them there.

Visit

See Reliable Context.

Reliable Context is in production use in two sectors; more than six thousand of its identity decisions were made by rule and none has been reversed. Articles published here explain general practice, cite their primary and academic sources, and name the jurisdiction of any regulatory rule they describe. They are not legal, regulatory or compliance advice for a specific organization, and requirements change.

Website: https://reliablecontext.com

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.