Reliable Context

RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

EIOPA AI Opinion: Data Governance and Records for Insurers

> 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 · eiopa ai opinion · insurance ai governance · solvency ii ai data quality · insurers ai record keeping · eiopa artificial intelligence

Original article: https://reliablecontext.com/articles/eiopa-ai-opinion-data-governance-record-keeping-insurers



In this article

1. What the Opinion is and why it matters
2. Legal nature: an Opinion, not a new rulebook
3. Scope: the AI systems the AI Act leaves to sectoral law
4. The two-step method: impact assessment, then proportionate measures
5. Data governance
6. Documentation and record keeping
7. Transparency, explainability and human oversight
8. Third-party AI
9. A worked example: one claims-triage model
10. Interplay with the AI Act and other law
11. Checklist
12. Summary
13. Sources

What the Opinion is and why it matters

On 6 August 2025 the European Insurance and Occupational Pensions Authority (EIOPA) published its *Opinion on Artificial Intelligence Governance and Risk Management*, reference EIOPA-BoS-25-360 (EIOPA publication page). It is addressed to national supervisors and sets out how existing insurance legislation, principally the Solvency II Directive, the Insurance Distribution Directive (IDD) and the Digital Operational Resilience Act (DORA), should be read when insurers and intermediaries use AI systems.

For data and compliance leaders in EU insurance, the Opinion is useful for a narrow reason: it translates general sectoral duties, such as keeping "adequate and orderly records" and using data that is "complete, accurate and appropriate", into expectations about AI training data, test data, model documentation and reproducibility. This article reads it section by section with that focus. It is educational and is not legal advice.

Legal nature: an Opinion, not a new rulebook

An EIOPA Opinion under Article 29(1)(a) of Regulation (EU) No 1094/2010 is a supervisory convergence tool. That article, as the Opinion recites, "mandates EIOPA to play an active role in building a common Union supervisory culture and consistent supervisory practices" (Opinion, para 1.1). The Opinion is "addressed to the competent authorities" and "covers the activities of both insurance undertakings and intermediaries" when they use AI systems (para 1.3).

It is explicit that it adds no new law: "This Opinion does not set out new requirements and does not seek to alter the scope of the AI Act by extending the requirements of the AI Act for high-risk AI systems to all AI systems used in insurance" (para 2.8). What it does is interpret existing provisions "in the context of AI systems which were either non-existent or not widely used when that legislation was approved" (para 2.9). The

provisions it relies on are listed in para 1.2: Articles 17, 20 and 25 of the IDD; Articles 41, 46 and 82 of Solvency II; Articles 4, 5, 6, 11 and 12 of DORA; Articles 258 and 260 of Delegated Regulation (EU) 2015/35; and Articles 6 to 9 of Delegated Regulation (EU) 2017/2358 on product oversight and governance.

EIOPA also states how it will follow up: "Two years following the publication of this Opinion, EIOPA will look into the supervisory practices of competent authorities with a view to evaluate supervisory convergence" (para 4.1).

Scope: the AI systems the AI Act leaves to sectoral law

The Opinion's scope is defined by reference to the EU AI Act, Regulation (EU) 2024/1689. The AI Act lists as high-risk "AI systems intended to be used for risk assessment and pricing in relation to natural persons in the case of life and health insurance" (AI Act, Annex III, point 5(c)). Those systems, and prohibited practices, are governed by the AI Act's own requirements. Everything else, which EIOPA notes continues to operate "subject to existing sectoral legislation without new requirements" apart from certain transparency duties, AI literacy and voluntary codes of conduct (para 2.5), is where the Opinion applies.

AI use in insurance (examples)	AI Act category	Covered by the Opinion?
Risk assessment and pricing of natural persons in life and health insurance	High-risk, Annex III point 5(c)	No (para 2.6)
Prohibited practices under Article 5	Prohibited	No
Motor or property pricing, claims triage, fraud detection, chatbots	Generally neither prohibited nor high-risk	Yes
Classical models outside the AI system definition (e.g. linear or logistic regression used for optimisation, per the Commission's guidelines quoted at footnote 11)	Not AI systems	No, but sectoral legislation "requires adequate and proportionate governance and risk management measures when using mathematical models, regardless of whether they are considered AI systems or not" (para 2.10)

Timing matters for the high-risk line. Following the Digital Omnibus on AI, Regulation (EU) 2026/1744, in force since 27 July 2026, the high-risk requirements apply to Annex III systems from 2 December 2027 (AI Act, Article 113 as amended). Until then, a life and health pricing model sits under sectoral law alone in practice, although the Opinion's text excludes it by category rather than by date.

The two-step method: impact assessment, then proportionate measures

The Opinion asks undertakings to proceed in two steps. First, "assess the risk of the different AI systems used" (para 3.2), with the depth of the assessment proportionate to potential impact. The criteria in para 3.3 include the scale of data processing, sensitivity of data, number of customers affected including vulnerable customers, autonomy, consumer-facing versus internal use, and potential adverse effects on individuals, plus insurance-

specific factors such as reliance on categories like age, or use in lines of business important for financial inclusion or compulsory by law. Prudential factors follow in para 3.4: business continuity, financial position, legal obligations and reputational risk.

Second, develop "a set of proportionate measures" (para 3.5). For low-impact systems, expectations "would be very limited"; for higher-risk systems, more stringent (para 3.6). The measures are organised under six headings (para 3.7):

1. Fairness and ethics
2. Data governance
3. Documentation and record keeping
4. Transparency and explainability
5. Human oversight
6. Accuracy, robustness and cybersecurity

EIOPA stresses that these are complementary: where a system processing images, video or text cannot be comprehensively explained, "complementary risk management measures such as data governance or human oversight may be developed to compensate for a lack of explainability" (para 3.6).

Data governance

The legal anchors, as quoted in para 3.17, are Article 260(1)(a)(ii) of Delegated Regulation 2015/35 (policies on "the sufficiency and quality of relevant data" for underwriting and reserving) and Article 82 of Solvency II (data for technical provisions "shall be complete, accurate and appropriate"). The Opinion carries those three adjectives over to AI:

> "The data used to train and test the AI system should be complete (e.g. sufficient historical information), accurate (e.g. no material errors) and appropriate (e.g. consistent with the purposes for which it is to be used). Any limitations of data in this regard should be duly documented and addressed." (para 3.19)

Three further expectations follow:

- Data governance applies "throughout the AI system life cycle for data collection, data processing and post processing" (para 3.20).
- External data bought from a third party is held to "the same data quality standards" (para 3.21).
- Undertakings should make "reasonable efforts to remove biases in the data, including potential unlawful proxy discriminatory variables" (para 3.14). Footnote 13 defines proxy discrimination as sensitive characteristics being "indirectly inferred from other customer characteristics (e.g. location)".

In data terms, "complete, accurate and appropriate" are three different tests. Completeness is about coverage of the population and period. Accuracy is about each value matching the fact it records. Appropriateness is about fitness for the purpose, which is why the same claims history can be appropriate for reserving and inappropriate for pricing if it embeds a proxy.

Documentation and record keeping

The anchors here are Article 258(1)(i) of Delegated Regulation 2015/35, which requires "adequate and orderly records of the insurance undertaking's business and internal organisation", and Article 9 of Delegated Regulation 2017/2358, which requires product approval actions to be "duly documented, kept for audit purposes and made available to the competent authorities upon request" (para 3.22). The Opinion's expectation is one sentence:

> "Following a risk-based and proportionate approach, undertakings should keep appropriate records of the training and testing data and the modelling methodologies to enable their reproducibility and traceability." (para 3.23)

Annex I gives an example list of records for higher-risk uses, drawn from the 2021 work of EIOPA's stakeholder group on digital ethics and labelled "for illustrative purposes only and ... not intended to be considered a prescriptive guidance". Its rows are summarised below.

Annex I record	What it describes
Reasons for using AI	Business objective and its consistency with strategy
Integration into IT infrastructure	How the model sits in the IT estate, and significant changes
Staff involved	Roles, responsibilities and training needs
Data collection	How the ground truth was built, how input data were selected, collected and labelled, and how bias was considered
Data preparation	Variables and their domain ranges; construction of training, test and prediction data sets; how engineered features were built
Data post processing	Processes for continuous data improvement, with timing and frequency
Technical choices	Why the algorithm was chosen, "associated libraries with exact references", limitations and trade-offs
Code and data	Code that goes to production; "exclusively for high impact applications", the training data and "all the associated hyper parameters, including pseudo-random seeds"
Model performance	KPIs, satisfactory thresholds, review and retraining frequency
Model security	Protection against manipulation
Ethics and trustworthy assessment	Impact assessment and how governance measures address it

Two details stand out. The pseudo-random seed requirement is a reproducibility requirement in the strict sense: without the seed, retraining on the same data may not produce the same model. And the Annex allows an escape: if recording training data and hyperparameters "proved to be too burdensome", undertakings "may put in place alternative measures that ensure the auditability of the AI model".

Transparency, explainability and human oversight

Explainability is tied to IDD Article 20(1) on comprehensible product information and to Article 258(h) of Delegated Regulation 2015/35 on information systems producing "complete, reliable, clear, consistent, timely and relevant information" (para 3.25). Undertakings "should ensure that the outcomes of AI systems can be meaningfully explained" (para 3.26). Explanations are to be adapted by audience: supervisors and auditors should get "a global and comprehensive explanation about the functioning of the AI system", and customers, on request, an explanation of "the influence of the AI system on the decision that has a material impact on them" in "simple, clear and non-technical language" (para 3.28). Model-agnostic tools such as LIME and SHAP are mentioned, with the caveat that their "assumptions and limitations ... should be duly documented".

Human oversight rests on Solvency II Article 46 (internal control). Para 3.30 assigns roles: the administrative, management or supervisory body is "responsible for the overall use of AI systems"; compliance and audit functions verify legal compliance; the Data Protection Officer verifies personal data processing; and the actuarial function controls AI systems within its remit. Undertakings may add an AI officer or an AI or data committee (para 3.31).

Third-party AI

Many insurers buy rather than build. The Opinion is direct: "Undertakings are ultimately responsible for the AI systems that they use, regardless of whether the AI systems are developed in-house or in collaboration with third party service providers" (para 3.11). They should obtain "adequate information and assurances" from vendors about characteristics, capabilities, training and test data, and limitations. Where intellectual property limits disclosure, they should compensate with contract clauses, service level agreements, external audits, or due diligence testing and monitoring.

A worked example: one claims-triage model

The following insurer is fictional. A mid-sized motor and property insurer deploys a vendor model that scores incoming property claims for likely fraud and routes high scores to a special investigations unit. It is not in Annex III; the Opinion applies.

Opinion paragraph	Question an examiner might ask	Record that answers it
3.2 to 3.4 impact assessment	How many customers does this affect, and can it deny a claim on its own?	Assessment noting 40,000 claims a year, consumer-facing, no autonomous denial
3.19 data quality	Is the training data complete, accurate and appropriate for 2026 claims?	Data set version, period covered, known gaps (no flood claims before 2022), documented limitation
3.14 proxies	Could postcode act as a proxy for a protected characteristic?	Proxy analysis and decision, with date and owner
3.21 external data	Was third-party data used, and to what standard?	Vendor data sheet; internal checks on a sample

Opinion paragraph	Question an examiner might ask	Record that answers it
3.23 reproducibility	Can you reproduce the score for claim C-58812 from 3 March 2026?	Model version, input snapshot, feature values, score, threshold in force
3.28 explanation	Why was this customer's claim referred?	Plain-language reason derived from the top contributing features
3.11 third party	What did the vendor disclose about training data?	Contract clause, disclosed summary, gap noted with compensating test

The row on reproducibility is where most programmes are weakest. A score is reproducible only if the inputs as they were on the day can be retrieved, not the inputs as they are now after the claim file was updated. That requires keeping the inputs, or references to versioned records, alongside each decision; the site's glossary entry on the audit trail describes the record structure that makes this possible.

Interplay with the AI Act and other law

Three points of contact are worth recording:

- **High-risk systems.** For Annex III point 5(c) systems, AI Act obligations apply instead of the Opinion, including deployer duties such as keeping automatically generated logs "for a period appropriate to the intended purpose of the high-risk AI system, of at least six months" (AI Act, Article 26(6)). Deployers that are financial institutions under Union financial services law "shall maintain the logs as part of the documentation kept pursuant to the relevant Union financial service law". Deployers of 5(c) systems must also carry out a fundamental rights impact assessment under Article 27(1).
- **All AI systems.** AI literacy under Article 4 and certain transparency duties apply regardless of risk category, as para 2.5 of the Opinion notes.
- **GDPR.** Footnote 9 notes that the General Data Protection Regulation "may also include provisions relevant to the use of AI systems". The accuracy principle in GDPR Article 5(1)(d) requires personal data to be "accurate and, where necessary, kept up to date".

Checklist

1. Is every AI system classified against the AI Act (prohibited, Annex III point 5(c), other), with the rationale recorded?
2. For in-scope systems, is there a dated impact assessment using the para 3.3 and 3.4 criteria?
3. Are training and test data sets versioned, with completeness, accuracy and appropriateness limitations documented?
4. Has a proxy-discrimination review been done and recorded?
5. Are third-party data and models held to the same data standards, with disclosure gaps and compensating controls recorded?
6. For higher-impact systems, can a past individual output be reproduced from recorded model version, inputs and parameters, including seeds where retraining matters?

7. Are explanations available at two levels: comprehensive for supervisors, plain-language for customers on request?
8. Are oversight roles (management body, compliance, audit, DPO, actuarial) assigned in policy?
9. Are performance and fairness metrics monitored over time for drift (paras 3.35 and 3.36)?

Summary

EIOPA's Opinion does not create new obligations. It reads Solvency II, the IDD and DORA as already requiring, for AI systems outside the AI Act's high-risk and prohibited categories, proportionate data governance and records sufficient for "reproducibility and traceability". For data teams, that means versioned training and test data, documented limitations, and the ability to reconstruct an individual output. The insurance sector note on this site discusses the same requirement for the facts an insurer acts on.

Sources

1. EIOPA, *Opinion on Artificial Intelligence governance and risk management*, publication page, 6 August 2025. https://www.eiopa.europa.eu/publications/opinion-artificial-intelligence-governance-and-risk-management_en
2. EIOPA, *Opinion on AI Governance and Risk Management*, EIOPA-BoS-25-360 (PDF). https://www.eiopa.europa.eu/document/download/88342342-a17f-4f88-842f-bf62c93012d6_en?filename=Opinion on Artificial Intelligence governance and risk management.pdf
3. Regulation (EU) 2024/1689 (Artificial Intelligence Act), EUR-Lex. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
4. Regulation (EU) 2026/1744 (Digital Omnibus on AI), EUR-Lex. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
5. European Commission, AI Act Service Desk, consolidated text as at 27 July 2026: Annex III, Articles 4, 26, 27 and 113. https://ai-act-service-desk.ec.europa.eu/en/ai-act/annex-3
6. Regulation (EU) 2016/679 (General Data Protection Regulation), EUR-Lex. https://eur-lex.europa.eu/eli/reg/2016/679/oj

Source links

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

1. https://www.eiopa.europa.eu/publications/opinion-artificial-intelligence-governance-and-risk-management_en
2. https://www.eiopa.europa.eu/document/download/88342342-a17f-4f88-842f-bf62c93012d6_en?filename=Opinion%20on%20Artificial%20Intelligence%20governance%20and%20risk%20management.pdf
3. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
4. https://ai-act-service-desk.ec.europa.eu/en/ai-act/annex-3
5. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
6. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-113
7. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-26

8. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-27
9. https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-4
10. https://eur-lex.europa.eu/eli/reg/2016/679/oj

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.