# Lightweight AI Governance (LAIG) Framework for SMEs

Aleksander Młodawski[1], Aleksandra Wolniak[1,*]

[1]Kozminski University, Jagiellońska 59, 03-301 Warsaw, Poland

**Abstract**

Small and medium sized enterprises (SMEs) increasingly deploy artificial intelligence solutions, yet they must comply with the European Union's Artificial Intelligence Act. High-risk systems, such as credit scoring tools or automated hiring screeners, must meet stringent requirements for risk management, technical documentation, human oversight, and post-market monitoring. Medium and low-risk systems also require ongoing observation to verify that their initial classification remains accurate and does not escalate. Although legislators introduced proportionality measures for smaller businesses, recent analyses suggest that implementation costs remain prohibitive for SMEs lacking in-house compliance staff. Commission estimates indicate that a 50-person start-up could incur roughly EUR 216 000–319 000 in first-year compliance costs for a single AI system. To close this gap, we propose the Lightweight AI Governance (LAIG) framework, a pragmatic, risk-based governance model expressly designed for SMEs. LAIG distils best practices from ISO/IEC 42001 and the NIST AI Risk Management Framework into modular procedures that can be embedded in familiar DevOps workflows and maintained with limited resources. Core elements include clear role assignment, inventory-centred risk classification, checklist-driven impact assessments, targeted mitigation controls, and concise Markdown and Git documentation templates optionally drafted with large language model assistance and always subject to human verification. An illustrative fintech scenario demonstrates how a company with forty employees can address bias, transparency, and oversight obligations without hiring a dedicated compliance department. By lowering the organisational and financial thresholds for trustworthy AI compliance, LAIG empowers European SMEs to continue innovating while satisfying both the letter and the spirit of the AI Act.

**Keywords**

AI governance, SMEs, EU Artificial Intelligence Act, risk management, TRUST-AI

## 1. Introduction

Artificial intelligence now reaches deeply into everyday operations of European small and medium sized enterprises, yet most of these firms still operate without formal governance structures. The absence of clear oversight exposes them to legal, ethical, and reputational risk. During the past four years the share of companies with fewer than 250 employees experimenting with AI doubled, according to Eurostat [1], while the European Union adopted Regulation (EU) 2024/1689, the "Artificial Intelligence Act" (AIA), which places high-risk systems such as credit-scoring engines, résumé screeners, and medical-triage tools under demanding rules covering risk management, data quality, technical documentation, human oversight, transparency, robustness, cybersecurity, and post-market monitoring [2]. The regulation promises simplified forms, fee reductions, and sandbox access for smaller businesses, but the Commission's impact assessment estimates that a 50-person enterprise would incur approximately EUR 216 000–319 000 in first-year compliance costs for a single AI system [3]. For resource-constrained teams these figures represent an existential barrier.

Voluntary standards have emerged in parallel. ISO/IEC 42001 sets out a Plan–Do–Check–Act management system that presumes enterprise-level resources [4], while the NIST AI Risk Management Framework (AI RMF) encourages an iterative Govern–Map–Measure–Manage cycle and likewise assumes dedicated compliance staff [5]. Transparency artefacts such as Model Cards [6] and Datasheets [7] enhance documentation yet address only part of the legal obligations and require multidisciplinary expertise. Consequently, SMEs confront a patchwork of ambitious frameworks without a practical roadmap, and the gap between guidance and daily engineering practice continues to widen. While

*TRUST-AI: The European Workshop on Trustworthy AI. Organized as part of the European Conference of Artificial Intelligence – ECAI 2025. October 2025, Bologna, Italy.*

*Corresponding author.

EMAIL: 46413@kozminski.edu.pl (A. Młodawski); 54156@kozminski.edu.pl (A. Wolniak)

some web-based compliance tools exist, they yield static reports external to development workflows and therefore offer limited help to fast-paced SME teams.

In response, we present the Lightweight AI Governance framework, an approach tailored to Git-centred development teams that converts every clause of Annex IV into modular checklists, Markdown templates, and automated gap analysis scripts. The framework builds on lessons from agile software engineering and aligns compliance checkpoints with familiar commit and review rituals. By grounding governance in existing developer workflows, we aim to minimise friction while maximising traceability. Our study therefore investigates whether the Act's obligations can be divided into tasks suited to limited resources, which parts of ISO/IEC 42001 and the NIST framework remain essential, how governance artefacts can coexist with source code in the same repository and continuous integration pipeline, and whether the resulting workflow can reduce compliance overhead while preserving auditability and the core principles of trustworthy AI. Unlike standalone checklist tools, LAIG integrates compliance steps directly into DevOps practices and automates verification, which we hypothesise can reduce overhead without compromising rigour.

## 2. Related work

Scholars and standard setters agree that trustworthy AI demands both technical safeguards and organisational controls, yet their guidance differs in level of detail, which complicates implementation for smaller firms. The European AI Act is the most comprehensive legal instrument to date. Chapter III mandates risk management, data governance, documentation, transparency, accuracy, robustness, cybersecurity, and human oversight for every high-risk system, while Annex IV lists the technical documentation required to demonstrate conformity [2]. Voluntary frameworks follow a similar direction. The NIST AI RMF maps trustworthy AI across four lifecycle functions—Govern, Map, Measure, and Manage [5]. ISO/IEC 42001 supplies a certifiable management system that embeds ethics, risk assessment, and continuous improvement through the Plan–Do–Check–Act cycle [4]. These efforts are supported by transparency artefacts including Model Cards that summarise intended use, datasets, metrics, and limitations [6], Datasheets that document data provenance and quality [7], and AI FactSheets that adopt a supplier's declaration approach [8]. A growing set of national and international policy initiatives is tracked by the OECD.AI Policy Observatory [9]. With respect to cybersecurity assurance, the EU adopted the EUCC certification scheme in January 2024 [10]. Regulatory approaches draw on these artefacts, although their production requires time and multidisciplinary expertise that typical SMEs do not possess [11].

Focused research on SME governance remains scarce. Analyses warn that the Act's proportionality measures may prove illusory without practical templates, shared tools, and subsidised audits [3]. Legal analyses clarify the structure and implications of the Act for SMEs, yet practical, publicly documented workflows remain limited [12]. Moreover, surveys confirm that small firms see standards as valuable yet overwhelming and therefore postpone governance until late in development when changes cost more [11]. Recent studies on DevOps pipelines demonstrate that embedding risk management directly into CI/CD loops improves defect detection, reduces incident response times, and ensures continuous evidence generation [13]. Similarly, AI-driven test case optimisation integrated into CI/CD reduces redundant testing and accelerates compliance-relevant validation cycles [14]. No open workflow translates Annex IV obligations and leading standards into proportional, low-overhead practices suitable for lean engineering teams. LAIG stores governance artefacts in the same repository as the code and enforces coverage through continuous integration, avoiding separate platforms. The framework therefore aims to close this gap by coupling modular documentation templates with risk-tiered governance steps that integrate naturally into DevOps pipelines and deliver trustworthy AI without enterprise-level bureaucracy.


**Figure 1:** Repository architecture embedding governance artefacts alongside source code.

# 3. Proposed methodology

The LAIG framework rests on two intertwined investigations. A clause-by-clause reading of the Artificial Intelligence Act supplied an exhaustive inventory of legal obligations, while a matching exercise linked each obligation with the management functions defined in ISO/IEC 42001 and the NIST AI RMF [2, 4, 5]. Complementing this top-down study, we conducted semi-structured interviews with five Polish startups, each employing between ten and sixty staff. The interviews followed a short protocol with informed consent and anonymisation, and the transcripts were analysed using template-based coding. Together these strands revealed both what SMEs must do and what they can realistically sustain, and produced a practical mapping table that guided the artefacts and checks implemented in LAIG within developer workflows.

## 3.1. Design objectives

Interview data showed that small product teams weigh governance steps against tight sprints, quarterly burn rates, and investor deadlines. In response the framework commits to six principles. First, every activity must create compliance evidence or mitigate risk in a way that exceeds its effort cost. Second, the workload is sliced into small autonomous tasks so that a developer can complete a compliance item between ordinary feature tickets. Third, traceability is guaranteed because every file carries explicit tags that point to one or more Annex IV clauses and the entire history lives in Git, which offers diff-based evidence for auditors. Fourth, the architecture embeds governance artefacts in the same repository, continuous integration system, and code review flow already familiar to engineers, so no one needs to log into a second platform. When a separate repository is unavoidable, a release build in the product repository must reference a signed governance commit, and the build fails if that commit does not pass coverage checks. Fifth, language models may propose text, yet a human must read, approve, and merge the content, which reduces the risk of hallucination and keeps accountability clear. Sixth, the framework relies on familiar tools, including Markdown editors, Git command line utilities, and spreadsheets, so adoption does not mandate new software procurement. These principles aim to preserve velocity, protect product quality, and satisfy regulators without turning a ten-person team into a paperwork factory.

## 3.2. Repository architecture

LAIG treats documentation as source material. A dedicated repository contains seven Markdown files that mirror the order of Annex IV, beginning with system description and ending with post-market monitoring. Each file opens with a YAML header that records system identifier, version, and author, then continues with section text marked by short comment tags that reference clause numbers. A lightweight linter runs on every pull request to highlight any tag that still lacks narrative. Because each commit message must quote the clause identifier, reviewers can follow the compliance trail from first model prototype to production deployment with no separate log.

## 3.3. Workflow

Daily operation unfolds in five human-controlled stages. A developer completes a short Markdown intake form that captures purpose, stakeholders, data sources, and maturity stage together with intended use, decision context, affected users, input data origin, model family, oversight thresholds, and the expected deployment pathway. A helper script transforms the answers into prompts that reference the

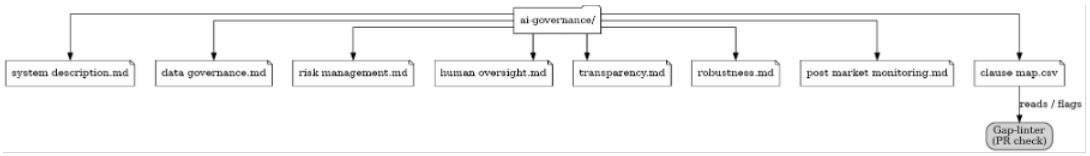


**Figure 2:** Clause mapping matrix aligning Annex IV requirements with governance artefacts.

legal clauses, so authors do not start from a blank page. A language model drafts text and inserts VERIFY tokens wherever confidence is low. A subject matter expert resolves tokens, edits numbers, checks links, and merges the pull request, which locks the draft into Git history. A continuous integration job then renders the repository to PDF and, optionally, to DOCX, so the latest technical file is always downloadable in a single click. The same job applies two gates. First, all applicable Annex IV items are complete or explicitly marked not applicable with a justification in one sentence, and no VERIFY tokens remain. Second, if any condition fails, the release is blocked and a remediation ticket is opened. When governance and product live in different repositories, the product build reads the referenced governance commit and blocks the release whenever the gates are not satisfied. Folding the entire loop into the familiar pull request rhythm converts compliance from an end-of-project silo into an everyday habit.

### 3.4. Clause mapping mechanism

To operationalise the principle that every action must generate evidence or mitigation, the governance repository is linked to the implementation repository through a shared CI/CD pipeline [13]. Each change to the source code follows a merge request procedure in which approval depends on completing a checklist that references the relevant clauses of Annex IV. During merging, an automatic compliance linter verifies the presence and consistency of compliance metadata. If the required documentation or mitigating measures are missing, the system blocks integration with the main branch. This solution eliminates the risk of purely declarative treatment of obligations and ensures that the compliance evidence trail is created in parallel with software development.

The clause mapping table pairs every Annex IV requirement with a heading path and a checkbox. If a requirement is irrelevant, such as photographs of a physical device for a pure software product, the author marks the row not applicable and writes a one-sentence justification. Because the table itself is version-controlled, auditors can review how coverage improves over time and developers can run automated diff checks to catch accidental deletions. Closely related sub-requirements are grouped to reduce boilerplate, and each group links to concrete artefacts and tests. This repository-first mapping contrasts with questionnaire tools because it is versioned alongside the code and enforced through automated checks during pull requests and releases.

## 4. Illustrative scenario

FinPay is a hypothetical fintech based in Warsaw that employs forty people and provides AI-driven credit scoring services. As creditworthiness assessment falls within Annex III of the AI Act as a high-risk use, FinPay must demonstrate comprehensive risk management, maintain technical documentation, ensure human oversight, transparency, robustness, and cybersecurity, and implement post-market monitoring. Before adopting LAIG, compliance evidence was scattered across unversioned documents and ad hoc spreadsheets, hindering audit readiness, obscuring accountability, and undermining traceability. Adoption proceeded in three compact sprints that embedded governance in everyday engineering. First, the team created a dedicated governance repository, assigned named responsibility for each relevant legal clause, and completed an intake questionnaire capturing purpose, stakeholders, data lineage, model family, oversight thresholds, and the intended deployment path. A machine-readable clause map linked Annex IV requirements to specific document headings and checkboxes, ensuring that each item was either completed or explicitly marked as not applicable with justification. Next, developers and data scientists drafted seven Markdown sections mirroring Annex IV using language

model assistance for first drafts with explicit verification markers, while legal and risk specialists resolved markers, corrected figures, validated links, and tightened claims. In parallel, the modelling team addressed dataset imbalance with balanced resampling and recorded the rationale for calibration and fairness thresholds alongside the data governance narrative. Finally, engineers implemented continuous integration that automatically renders the repository into a technical file with each change and enforces two release gates: full clause coverage with no unresolved verification markers, and a signed governance commit referenced by the product build. If either gate fails, the release is blocked and a remediation ticket is generated. The first end-to-end execution of this workflow produced a seventeen-page technical file, ready for engagement with a regulatory sandbox and for supporting early dialogue with assessors. Developers reported that handling governance artefacts through pull requests felt natural, while managers valued the living audit trail in Git and the clear line from requirement to mitigation. Although formal time and cost measurements are forthcoming, the team reports higher confidence and improved readiness for its first conformity assessment under the AIA.

## 5. Discussion

The LAIG framework gives small enterprises a practical route to trustworthy AI by folding risk checks, human oversight, and continuous monitoring into the Git-based routines developers already follow. This arrangement addresses Annex IV documentation duties and supports the operationalisation of key Chapter III obligations within development workflows, while creating a minimum viable Plan–Do–Check–Act loop that aligns with the Govern and Manage functions of the NIST AI RMF. By slicing Annex IV into bite-sized tasks, the method lets teams focus on the highest risks first, a need often voiced in founder interviews. Governance files sit beside source code so engineers can move from feature branch to compliance update without changing context, and each commit records a searchable audit trail that managers and examiners value. Automated gap tests flag missing clauses the moment a pull request appears, and language models generate draft text that experts refine, a pairing that speeds writing yet keeps the final judgement human.

These gains come with caveats. The simplified nature of the framework, which is its main advantage for SMEs, simultaneously creates risks that must be managed. LAIG in its basic form is best suited for low- and moderate-risk systems. For AI systems classified as high-risk under the AI Act, such as tools for credit scoring or recruitment screening, a simplified governance approach alone may prove insufficient. SMEs implementing such systems will need to invest in more rigorous control, risk management, and post-market monitoring mechanisms in line with the strict requirements of the regulation. In this context, LAIG should be interpreted as a transitional instrument of proportional governance. Its role is not to replace comprehensive conformity assessment procedures prescribed for high-risk AI systems, but rather to enable SMEs to initiate systematic preparation of Annex IV documentation within their existing development workflows. The artefacts created in this process—Git-based audit trails, clause-referenced checklists, and modular records—constitute a preparatory layer of evidence that can be expanded into the full compliance structures required under ISO/IEC 42001 or NIST-aligned frameworks and presented to notified bodies during formal assessments.

A lightweight approach omits some formalities, which can lead to overlooking hidden issues such as undiscovered bias in training data or security vulnerabilities [15]. Although LAIG promotes bias mitigation through balanced resampling, its effectiveness depends on team diligence. There is also a risk that automated templates and checklists will be treated as a formality rather than an opportunity for critical risk assessment. The framework therefore depends on an organisational culture that supports critical evaluation of model outputs and continuous interrogation of compliance artefacts [11, 16]. By embedding governance routines directly into software engineering processes rather than relegating them to a separate compliance silo, LAIG fosters incremental institutional capacity-building. This incrementalism is consistent with the proportionality principle embedded in the AI Act and ensures that organisations can scale governance maturity progressively without disruptive restructuring when high-risk classification applies.

Using language models to assist with documentation speeds drafting but introduces concerns about data confidentiality and the accuracy of generated content. Although the framework requires human verification, the risk associated with sending sensitive information to external service providers remains significant. The FinPay case remains illustrative rather than empirical, hardware-rich products will require extra safety documentation, and sector-specific laws may layer additional duties on top of the AI Act. Future pilot deployments are expected to yield empirical data on time and cost efficiency, while the development of open-source tools for automated clause verification and dossier generation, combined with cooperation with notified bodies, may support the emergence of a "light audit track" in which the repository itself serves as primary evidence [15]. Accordingly, LAIG should be regarded not as a substitute for the comprehensive compliance architecture required under European Union law, but as an instrument of progressive compliance that operationalises the transition from minimum viable governance to the full spectrum of obligations triggered by deployment of high-risk systems.

## 6. Conclusions

The European AI Act asks SMEs to deliver governance that rivals large corporations even though their budgets are far smaller. The LAIG framework responds by turning every Annex IV duty, along with the intent of ISO/IEC 42001 and the NIST AI RMF, into modular routines that fit naturally within existing engineering practice. Documentation resides alongside source code in a version-controlled repository, and every change generates verifiable evidence that can be independently reviewed. Clause coverage is enforced through structured checklists and machine-readable mappings, ensuring that each requirement is either fulfilled or explicitly marked as not applicable with justification. Automated CI/CD gates prevent releases in which required artefacts or reviews are incomplete. Optional language model assistance accelerates drafting but human approval remains decisive, which preserves accountability and limits the risk of error. Teams define acceptance thresholds for accuracy, calibration, and fairness in line with the system's risk profile and track them across releases. In this way LAIG establishes a disciplined cycle of planning, doing, checking, and improving that protects developer velocity while creating a durable audit trail.

Next steps are empirical and collaborative. Planned pilots in finance and health analytics will measure documentation effort, auditor revision cycles, developer usability, and outcome quality including calibration and fairness. Findings will feed improvements to templates, coverage checks, and acceptance criteria and will inform engagement with regulators and notified bodies toward a credible light audit path in which a well-maintained repository can serve as primary evidence of compliance. If the expected gains are confirmed, LAIG will lower the organisational and financial threshold for trustworthy AI and help European SMEs sustain innovation while meeting both the letter and the spirit of the Act.

## Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT to generate Figures 1 and 2. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

# References


[1] Eurostat. Use of artificial intelligence in enterprises. Statistics Explained. Data extracted January 2025, planned update January 2026. Available at: ec.europa.eu/eurostat/statistics-explained.

[2] European Parliament and Council. Regulation (EU) 2024/1689 — Artificial Intelligence Act. Official Journal of the European Union, 2024. Available at: eur-lex.europa.eu/eli/reg/2024/1689/oj.

[3] European Commission (DG CONNECT et al.). Study to support an impact assessment of regulatory requirements for Artificial Intelligence in Europe – Final report. Publications Office of the European Union, 2021. Available at: artificialintelligenceact.eu/AIA-COM-Impact-Assessment.

[4] ISO/IEC. ISO/IEC 42001:2023 — Artificial intelligence — Management system — Requirements. International Organization for Standardization, 2023.

[5] National Institute of Standards and Technology. AI Risk Management Framework (AI RMF 1.0). NIST, 2023. doi:10.6028/NIST.AI.100-1.

[6] Margaret Mitchell, Simone Wu, Andrew Zaldivar, and colleagues. Model Cards for Model Reporting. In *Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency (FAccT '19)*, pages 220–229. ACM, 2019. doi:10.1145/3287560.3287596.

[7] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, and colleagues. Datasheets for Datasets. *Communications of the ACM*, 64(12):86–92, 2021. doi:10.1145/3458723.

[8] Rachel K. E. Bellamy, Kush R. Dey, Michael Hind, and colleagues. AI FactSheets: Increasing trust in AI services through supplier's declarations of conformity. *IBM Journal of Research and Development*, 63(4/5):6:1–6:13, 2019. doi:10.1147/JRD.2019.2942288.

[9] OECD.AI Policy Observatory. Policies & initiatives — global navigator (overview). Living repository of national and international AI policies, 2024–2025. Available at: oecd.ai/en/dashboards/overview.

[10] European Union Agency for Cybersecurity (ENISA). An EU Prime! The EU adopts the first cybersecurity certification scheme (EUCC). News release, 31 January 2024. Available at: enisa.europa.eu/news/EUCC.

[11] Mehdi S. Soudi and Michel Bauters. AI guidelines and ethical readiness inside SMEs. *Digital Society*, 3:49, 2024. doi:10.1007/s44206-024-00149-9.

[12] Michael Veale and Frederik Zuiderveen Borgesius. Demystifying the Draft EU Artificial Intelligence Act. *Computer Law Review International*, 22(4):97–112, 2021. doi:10.9785/cri-2021-220402.

[13] Pavan M. Katikireddi. Smart risk management in DevOps using AI. *International Journal of Scientific Research in Science and Technology*, 10(3):1248–1253, 2023. doi:10.32628/IJSRST523103169.

[14] Alice John, Isaac John, and Tiberius Dion. Integrating AI-driven test case optimization into CI/CD pipelines. SSRN preprint, May 2025. doi:10.2139/ssrn.5252630.

[15] Patrick Guldimann, Anton Spiridonov, Ralf Staab, and colleagues. COMPL-AI Framework: A Technical Interpretation and LLM Benchmarking Suite for the EU Artificial Intelligence Act. arXiv preprint arXiv:2410.07959, 2024. doi:10.48550/arXiv.2410.07959.

[16] Manikandan V. Krishnamoorthy. Meta-Sealing: A revolutionizing integrity assurance protocol for transparent, tamper-proof, and trustworthy AI systems. arXiv preprint arXiv:2411.00069, 2024. doi:10.48550/arXiv.2411.00069.