ANNEX IV | ARTICLE 11 | EU AI ACT 2024/1689

The Annex IV Technical Documentation Guide for Machine Learning Teams

Operationalizing EU AI Act Article 11 Across the ML Lifecycle

A Premium AI Governance Desk Briefing for ML Team Leads, MLOps Engineers, AI Compliance Officers, and Enterprise Risk Teams

Published by AIGovernanceDesk.com

AI Governance Desk Editorial Team

June 2026 Edition

aigovernancedesk.com | contact@aigovernancedesk.com

Table of Contents

(c) 2026 AI Governance Desk. All rights reserved.

Executive Summary & Key Governance Points

EU AI Act Article 11 converts technical documentation from a legal artefact into an engineering deliverable. For providers of high-risk AI systems, Annex IV requires eight categories of evidence — architecture diagrams, dataset cards, risk registers, validation reports, and change logs — that must be produced and version-controlled across the full machine learning lifecycle, not assembled retrospectively before a compliance deadline. This briefing translates the statutory text of Article 11 and Annex IV into an operational checklist for ML team leads, MLOps engineers, AI compliance officers, and enterprise risk teams, and maps the requirement against ISO/IEC 42001, ISO/IEC 23053, and the NIST AI Risk Management Framework.

The following points summarize the governance, regulatory, and compliance takeaways developed throughout this briefing.

- **The high-risk compliance timeline has moved.** Co-legislators reached a provisional political agreement on 7 May 2026, as part of the Digital Omnibus on AI, to defer the application of Annex III high-risk obligations from 2 August 2026 to 2 December 2027, and obligations for AI systems embedded in regulated products under Annex I from 2 August 2027 to 2 August 2028. This agreement requires formal adoption and publication in the Official Journal before it takes legal effect, and 2 August 2026 remains the operative date on the statute books until that happens.

- **Annex IV's eight sections are interdependent, not parallel.** Sections 2, 4, and 5 depend on evidence produced during data engineering, model development, and validation; Sections 1 and 3 depend on deployment and monitoring artefacts; Section 6 depends on disciplined change management. Technical documentation cannot be produced as a single post-development exercise.

- **Contemporaneity is the evidentiary standard auditors apply.** Dataset cards, risk assessments, and bias examinations written after model deployment — rather than when the underlying decisions were made — create evidentiary gaps that retrospective narrative cannot close.

- **Article 18 imposes a ten-year retention obligation** on technical documentation, quality management system records, and conformity assessment files, applying to every superseded version, not only the current one. This has direct implications for documentation version control and tooling continuity.

- **Model cards and data sheets typically cover only 30 to 40 percent of Annex IV.** They generally substantiate portions of Sections 2 and 4 but do not address risk management documentation (Section 5), change management (Section 6), harmonised standards (Section 7), or the declaration of conformity (Section 8).

- **The SME simplified form changes presentation, not substance.** Article 11(1) permits small and micro enterprises to use a streamlined template, which notified bodies must accept for conformity assessment purposes, but the underlying evidentiary requirements under Articles 9, 10, 12, 13, and 14 are unchanged.

- **Existing ISO/IEC 42001 or NIST AI RMF investments reduce, but do not eliminate, the documentation burden.** Organizational-level management system controls map well to Annex IV Sections 1, 5, 7, and 8, but ML-specific technical detail for architecture, data, and validation (Sections 2, 3, 4, 6)

(c) 2026 AI Governance Desk. All rights reserved.

generally requires dedicated supplementation.

REGULATORY CAUTION

This briefing reflects publicly available information as of June 2026. The Digital Omnibus amendments described above were politically agreed but not yet formally adopted at the time of writing. Readers should confirm the current legal status of the high-risk AI timeline through the official EU AI Act service desk or counsel before making compliance decisions based on any specific date in this document.

(c) 2026 AI Governance Desk. All rights reserved.

1

Why Technical Documentation Is an Engineering Deliverable, Not Legal Paperwork

For machine learning teams building high-risk AI systems, the EU AI Act's technical documentation requirements represent a structural shift in how compliance evidence is produced. Article 11 and Annex IV of Regulation (EU) 2024/1689 do not request legal memoranda or policy statements. They require engineering artefacts — dataset cards, architecture diagrams, risk registers, validation reports, and change logs — generated, version-controlled, and maintained across the full ML lifecycle. Without this evidence, a provider cannot obtain CE marking, cannot lawfully place a system on the EU market, and cannot satisfy the conformity assessment obligations linked to Article 43.

Most ML teams already maintain documentation practices of some kind: experiment tracking, model cards, data sheets, and pipeline metadata. Published guidance from the European Commission's AI Act service desk indicates that the gap most providers face is not an absence of documentation but a lack of regulatory framing and completeness. Annex IV specifies eight sections, each carrying evidentiary requirements that extend beyond what typical ML tooling produces by default.

Governance implication. Treating Annex IV as a compliance checklist to be completed at the end of a development cycle misreads its structure. The sections are built from evidence generated throughout data engineering, model development, validation, deployment, and change management. Organizations that defer documentation planning until shortly before a conformity assessment will find gaps that cannot be closed under time pressure, because the underlying contemporaneous evidence simply does not exist.

This briefing is organized around the actual ML development lifecycle rather than the legal numbering of Annex IV, and provides explicit cross-references to ISO/IEC 42001:2023, ISO/IEC 23053:2022, and the NIST AI Risk Management Framework so that teams can extend existing management-system investments rather than building parallel regulatory documentation. The intended audience is ML team leads, MLOps engineers, AI compliance officers, technical program managers, and enterprise risk teams responsible for operationalizing EU AI Act compliance inside existing development workflows. Readers seeking statutory interpretation should consult the official consolidated text of Regulation (EU) 2024/1689 directly, or a qualified legal advisor.

2

The Legal Foundation: What Article 11 and Annex IV Actually Require

2.1 Article 11: The Core Obligation

Article 11 establishes four obligations for providers of high-risk AI systems. First, technical documentation must be drawn up before the system is placed on the market or put into service. Second, it must be kept up to date throughout the system's lifecycle. Third, it must demonstrate compliance with the requirements set out in Chapter III, Section 2 of the Act. Fourth, it must be made available to national competent authorities and, where applicable, to notified bodies upon request.

(c) 2026 AI Governance Desk. All rights reserved.

Article 11(1) also contains a specific accommodation for small and medium-sized enterprises and start-ups: they may provide the Annex IV information in a simplified manner, using a form the European Commission is required to establish, and notified bodies must accept that form for conformity assessment purposes. As Section 8 of this briefing addresses, simplification changes the presentation of evidence, not the underlying substantive requirement that the evidence exist.

DOCUMENTATION REQUIREMENT

Article 11(2) requires that where a high-risk AI system relates to a product already covered by Union harmonisation legislation (for example, machinery or medical device regulation), a single combined set of technical documentation must be produced covering both the AI Act requirements and the sectoral product requirements, rather than two parallel files.

2.2 Annex IV: The Eight Mandatory Sections

Annex IV specifies the sections that must appear in the technical documentation. The table below maps each section to its legal basis, its practical content, and the ML lifecycle phase in which the underlying evidence is typically produced.

Section	Legal Basis	What It Covers	ML Lifecycle Phase
1. General description	Annex IV(1)	Intended purpose, provider identity, system version, hardware/software environment, integration interfaces	Deployment
2. Elements & development process	Annex IV(2)	Architecture, algorithms, data requirements, training methodology, human oversight measures, validation, cybersecurity	Data engineering, model development, validation
3. Monitoring, functioning & control	Annex IV(3)	Capabilities and limitations, foreseeable unintended outcomes, oversight capabilities, input specifications	Deployment, monitoring
4. Performance metrics	Annex IV(4)	Rationale for the performance metrics selected for the specific system	Validation
5. Risk management system	Annex IV(5)	Documentation of the risk management system established under Article 9	Risk management (cross-cutting)
6. Lifecycle changes	Annex IV(6)	Change log, modification impact assessments, substantial-modification determinations	Change management (cross-cutting)
7. Harmonised standards	Annex IV(7)	Standards applied, or alternative solutions adopted where no harmonised standard exists	Governance

(c) 2026 AI Governance Desk. All rights reserved.

Section	Legal Basis	What It Covers	ML Lifecycle Phase
8. EU declaration of conformity	Annex IV(8)	The declaration of conformity drawn up under Article 47	Governance

Table 1. Annex IV section mapping. Source: Regulation (EU) 2024/1689, Annex IV; European Commission AI Act service desk, Article 11.

This mapping reveals a structural point that high-level legal summaries frequently miss: the eight sections are not independent deliverables. Sections 2, 4, and 5 depend on evidence produced during data engineering, model development, and validation. Sections 1 and 3 depend on deployment and monitoring artefacts. Section 6 depends on disciplined change management across every phase. Treating Annex IV as a post-development checklist misunderstands its architecture and tends to produce documentation that fails conformity assessment scrutiny.

2.3 The Ten-Year Retention Requirement

A separate but critical obligation concerns document retention. Article 11, read together with Article 18, requires providers of high-risk AI systems to keep technical documentation, quality management system documentation, conformity assessment records, and declarations of conformity available for at least ten years after the system is placed on the market or put into service. For systems with long deployment horizons, this implies documentation governance infrastructure that must outlast individual team members, organizational changes, and tooling migrations.

COMPLIANCE RISK

The retention obligation applies to every documentation version, not only the current one. Organizations that retrain models on a monthly or quarterly cadence can accumulate a substantial volume of superseded documentation versions, each of which must remain accessible for audit purposes for up to a decade. Documentation systems that overwrite prior versions, rather than archiving them, create a retention gap that is difficult to remediate after the fact.

3

Regulatory Timeline Update: The Digital Omnibus Deferral

The original application date for most Annex III high-risk AI obligations, including the Article 11 technical documentation duty, was 2 August 2026. On 19 November 2025, the European Commission proposed the Digital Omnibus on AI, a package of targeted amendments addressing implementation challenges that had emerged around the designation of national competent authorities and the finalisation of harmonised standards needed to support high-risk AI system compliance.

On 7 May 2026, negotiators from the Council of the European Union, the European Parliament, and the European Commission reached a provisional political agreement on the Digital Omnibus on AI. Under the agreed text, the application of high-risk obligations for stand-alone Annex III systems is deferred from 2 August 2026 to 2 December 2027. Obligations for high-risk AI systems embedded in regulated products under Annex I — including machinery, medical devices, and certain vehicles — are deferred from 2

(c) 2026 AI Governance Desk. All rights reserved.

August 2027 to 2 August 2028. Available reporting also indicates the agreement reinstates a fixed-date structure rather than the conditional, standards-availability trigger the Commission had originally proposed.

REGULATORY CAUTION

As of this briefing's publication date, the provisional agreement requires formal endorsement by the Council and the European Parliament and subsequent publication in the Official Journal before it takes legal effect. Published legal commentary notes that 2 August 2026 remains an active compliance date on the statute books until formal adoption occurs, and that formal adoption was expected before that date. Article 50 transparency and labelling obligations are reported to largely remain on their original schedule and are not affected by the high-risk deferral. Teams should verify the current status directly with the European Commission's AI Act service desk before adjusting internal compliance timelines.

Obligation	Original Date	Status as of June 2026
Annex III stand-alone high-risk systems (Article 11 documentation, Articles 9, 10, 12-15)	2 August 2026	Provisionally agreed deferral to 2 December 2027, pending formal adoption
Annex I high-risk systems embedded in regulated products	2 August 2027	Provisionally agreed deferral to 2 August 2028, pending formal adoption
Article 50 transparency and AI-content labelling obligations	2 August 2026	Reported to remain largely on original schedule
National AI regulatory sandbox establishment	2 August 2026	Reported deferral to 2 August 2027

Table 2. High-risk AI Act timeline status, compiled from public reporting on the Digital Omnibus on AI provisional agreement of 7 May 2026. Verify against the Official Journal before relying on any date for compliance planning.

Why this matters for documentation planning. A deferred legal deadline does not reduce the practical argument for early documentation investment. The dependency chain described in Section 4 of this briefing means that data governance, risk management, and validation evidence must be produced as systems are built, not assembled in the weeks before any compliance date, whichever date ultimately applies. Teams that pause documentation work while awaiting final legal certainty risk having to reconstruct evidence retrospectively, which, as later sections discuss, undermines the contemporaneity that auditors and notified bodies expect.

4

The Documentation Dependency Chain: Why You Cannot Start at the End

Technical documentation under Article 11 is not a standalone deliverable. It is a dependent output of upstream obligations that must be satisfied before the Annex IV package can be considered complete. Understanding this dependency chain is essential because it shows why technical documentation cannot be produced retrospectively at the end of development; it must be built incrementally as each upstream

(c) 2026 AI Governance Desk. All rights reserved.

requirement is satisfied.

4.1 Article 10 (Data Governance) Feeds Annex IV Section 2

Article 10 requires providers to establish data governance practices for training, validation, and testing datasets. These practices generate records that become prerequisites for Annex IV Section 2. Specifically, Annex IV(2)(d) requires descriptions of the datasets used, including their origin, size, and main characteristics — descriptions that are substantiated by Article 10 compliance records such as dataset cards, provenance documentation, bias examination reports, and data processing records.

These records must exist before technical documentation can be considered complete. Reconstructing dataset provenance or bias examination methodology after the model has already been trained is a pattern that raises concern during audit, because it cannot establish that decisions were documented when they were made. The dependency operates at two levels: the existence of data governance records is a prerequisite for completing Annex IV Section 2, and the quality of those records determines whether the resulting technical documentation can withstand conformity assessment scrutiny.

4.2 Article 9 (Risk Management) Feeds Annex IV Section 5

Article 9 requires providers to establish a risk management system that identifies, estimates, and evaluates known and foreseeable risks throughout the AI system's lifecycle. Annex IV Section 5 requires the technical documentation to include documentation of this risk management system. The risk management system must be established before technical documentation is finalized, with risk identification, estimation, mitigation measures, and residual risk assessments documented and referenced.

For ML systems specifically, data-related risks — bias, representativeness gaps, distribution shift, label noise — are among the most significant sources of operational risk. A risk management file that does not address these risks, or that fails to link data risk findings to specific system design decisions in a traceable way, is incomplete for the purposes of Annex IV Section 5.

4.3 Article 14 (Human Oversight) Feeds Annex IV Sections 2 and 3

Article 14 requires providers to design high-risk AI systems with appropriate human oversight measures. Annex IV(2)(e) requires an assessment of these measures as part of development-process documentation, and Annex IV(3) requires a description of the oversight capabilities and limitations built into the system. Technical measures that facilitate interpretation of system outputs, required under Article 13(3)(d), must also be documented.

For ML teams, this means human oversight is not solely a post-deployment operational concern. The design of oversight mechanisms — monitoring dashboards, interpretation aids, override controls — must be documented during development, with rationale for why specific measures were selected and how they address the system's foreseeable failure modes.

4.4 Article 12 (Record-Keeping) Feeds Annex IV Section 3

Article 12 requires high-risk AI systems to have logging capabilities that enable tracing of the system's functioning throughout its lifecycle. Annex IV Section 3 requires the technical documentation to describe

(c) 2026 AI Governance Desk. All rights reserved.

these capabilities, including what is logged, how logs are retained, and who has access. For ML systems, this extends beyond standard application logging to model inference logs, input/output pairs, confidence scores, and anomaly flags that enable post-hoc analysis of model behavior.

4.5 Article 13 (Transparency) Feeds Annex IV Section 1

Article 13 requires providers to supply instructions for use with high-risk AI systems. Annex IV Section 1 requires the general description to reference these instructions and ensure consistency across documents. The information on capabilities, limitations, and risks described in the instructions for use must align with the technical documentation; inconsistencies between the two are a recurring source of conformity assessment findings.

GOVERNANCE NOTE

Because Annex IV Sections 1 through 6 each depend on upstream Article 9-14 evidence, a useful governance control is to gate progression between ML lifecycle phases on documentation completeness, in the same way engineering gates often require passing tests before a merge. A phase that produces no contemporaneous documentation output should not be considered complete for compliance purposes, regardless of whether the underlying engineering work is finished.

5

The ML Team Checklist: Mapping Annex IV to the Machine Learning Lifecycle

Most existing guidance on Annex IV organizes content around the legal structure of the eight sections. For ML teams, this structure is backwards: engineers need to know what to document at each stage of development, not which legal paragraph corresponds to which artefact. The checklist below inverts the Annex IV structure and maps requirements to six phases of the actual ML lifecycle: data engineering, model development, validation, risk management integration, deployment with human oversight, and change management.

Phase 1 — Data Engineering and Data Governance (Annex IV Section 2)

Data governance under Article 10 is the foundation of defensible technical documentation. The records produced during this phase substantiate the dataset descriptions required under Annex IV(2)(d). The following items should be complete before the model development phase proceeds with confidence that downstream documentation will be sound:

- Dataset cards produced for all training, validation, and testing datasets, including origin, collection methods, selection criteria, and intended use context
- Data provenance documented with full lineage from raw sources to processed training inputs, including any intermediaries or transformations
- Data quality assessment completed, covering relevance, representativeness, error detection, completeness, and statistical adequacy for the intended deployment context

(c) 2026 AI Governance Desk. All rights reserved.

- Statistical properties documented and matched explicitly to the intended deployment context, with gaps identified and justified
- Bias examination conducted with documented methodology, metrics, findings, and mitigations applied; negative results retained, not discarded
- Data processing pipeline documented, including cleaning, normalization, augmentation, feature engineering, and any synthetic data generation
- Personal data handling documented, including legal basis under GDPR and a Data Protection Impact Assessment where applicable
- Data governance policy in place and referenced in technical documentation, with clear ownership and review schedules
- Version control for datasets, or equivalent provenance records, enabling reconstruction of any training run from documented inputs

Key deliverables: Data catalogue · bias assessment report · data processing record · data lineage diagram · DPIA (if applicable)

Phase 2 — Model Development and Architecture (Annex IV Section 2)

Annex IV(2)(a)-(c) and (f) require detailed documentation of system architecture, algorithms, model types, and development methodology. For ML teams, this is not a post-hoc description task. Architecture decisions should be recorded as they are made, with rationale for why specific approaches were selected over alternatives.

- System architecture diagrams produced for both the training pipeline and the inference pipeline, showing data flow, component boundaries, and external interfaces
- Algorithms and model types documented with selection rationale, including why the chosen approach suits the intended purpose and what alternatives were evaluated
- Model architecture specified in sufficient detail for reproducibility, including network topology, layer configurations, ensemble methods, or other structural choices
- Development methodology documented — supervised learning, reinforcement learning, rule-based, hybrid, or other — with justification
- Computational resources documented, including training and inference infrastructure and, where relevant, energy consumption estimates
- Third-party components inventoried with version numbers, licenses, provenance, and any modifications made, including libraries, pre-trained models, APIs, and cloud services
- Design trade-offs documented — accuracy versus fairness, interpretability versus efficiency, latency versus complexity — with explicit decision rationale
- Key technical decisions recorded with rationale and alternatives considered, maintained in a decision log with timestamps and responsible individuals

(c) 2026 AI Governance Desk. All rights reserved.

- Version control enforced for code, model weights, training configurations, and test datasets, with tags linking artefacts to specific documentation versions

Key deliverables: Architecture design document · AI bill of materials (BOM) · design decision log

Phase 3 — Validation and Testing (Annex IV Sections 2, 3 and 4)

Validation and testing documentation serves two purposes under Annex IV: it demonstrates that the system performs as intended, and it provides evidence that the provider understands the system's limitations. Annex IV(4) specifically requires rationale for the selection of performance metrics, a requirement that standard ML practice frequently overlooks when teams default to accuracy or F1 score without justifying why those metrics suit the specific system and use case.

- Performance metrics defined with explicit appropriateness rationale for the specific system, intended purpose, and affected population; default metrics without justification are insufficient
- Validation and testing procedures documented, including methodology, datasets used, test scenarios, environmental conditions, and acceptance criteria
- Test results reported with confidence intervals and subgroup disaggregation, showing performance across demographic or operational subgroups where relevant
- Robustness testing completed, covering input perturbations, distribution shifts, adversarial inputs, edge cases, and out-of-distribution scenarios
- Cybersecurity assessment conducted, including threat modeling, vulnerability assessment, and penetration testing where appropriate for the system architecture
- Validation against foreseeable conditions of use documented, including expected operational environments, user skill levels, and integration contexts
- Test logs and reports dated and signed by responsible persons, with clear traceability to the system version and dataset versions tested
- Human oversight measures assessed and documented, including how oversight capabilities were validated during testing

Key deliverables: Performance evaluation report · robustness test report · security assessment report · test logs

Phase 4 — Risk Management Integration (Annex IV Section 5)

The risk management system required under Article 9 must be documented and referenced in Annex IV Section 5. For ML teams, this phase is cross-cutting: risk management activity should begin during data engineering and continue through deployment and monitoring.

- Risk management process documented, including methodology for risk identification, estimation, evaluation, and treatment
- Known and foreseeable risks catalogued, covering health and safety risks, fundamental rights risks, and discrimination risks specific to the intended use context

(c) 2026 AI Governance Desk. All rights reserved.

- Risk mitigation measures documented per identified risk, with evidence that the measures were tested and are effective
- Testing to identify the most appropriate risk management measures documented, including why specific mitigations were selected over alternatives
- Residual risk assessment completed with acceptability justification, including rationale for why remaining risks are acceptable given the system's intended purpose and benefits
- Risk monitoring plan in place, specifying how risks will be tracked post-deployment and what triggers a reassessment
- Linkage between risk findings and system design changes documented, showing that risks identified during development led to specific design or implementation modifications

Key deliverables: Risk management file · risk register · risk assessment reports

Phase 5 — Deployment and Human Oversight (Annex IV Sections 1 and 3)

Deployment documentation defines the system's identity, purpose, and operational envelope. The intended purpose statement must be precise, specific, bounded, and testable; vague or expansive purpose statements undermine the rest of the documentation package because they make it impossible to define the boundaries of foreseeable use and misuse.

- Intended purpose statement precisely defined, specific, bounded, and testable — narrow enough to support risk assessment, clear enough to guide deployment
- Provider identification and system version documented, including version numbering scheme and relationship to documentation versions
- Hardware and software requirements specified, including minimum specifications, dependencies, and integration interfaces
- Human oversight design documented, including monitoring capabilities, interpretation aids, override mechanisms, and escalation procedures
- Instructions for use prepared in accordance with Article 13, covering capabilities, limitations, risks, and correct operation
- Logging capabilities specified in accordance with Article 12, including what is logged, retention periods, access controls, and audit trail integrity measures
- Post-market monitoring plan included, referencing Article 72 requirements and specifying how performance and risk will be tracked after deployment

Key deliverables: Intended purpose statement · system overview document · instructions for use · oversight design assessment

Phase 6 — Change Management and Lifecycle Updates (Annex IV Section 6)

(c) 2026 AI Governance Desk. All rights reserved.

Annex IV(6) requires technical documentation to include a description of changes made to the system throughout its lifecycle. For ML systems, which may be retrained, updated, or modified continuously, this requires disciplined change management that many ML teams do not yet practice as a matter of course.

- Change log established with comprehensive modification records, including date, nature of change, responsible party, and impact assessment
- Change classification framework defined, with clear criteria distinguishing non-substantial, potentially substantial, and substantial modifications
- Modification impact assessments conducted for all changes classified as potentially substantial or substantial
- Re-assessment records maintained where substantial modifications trigger a new conformity assessment under Article 43(4)
- Version history linking system versions to documentation versions maintained, so that any deployed system can be matched to its complete documentation package

Key deliverables: Change management log · modification impact assessments · version traceability matrix

6

The Substantial Modification Decision Framework for ML Systems

General coverage of Annex IV Section 6 tends to treat change documentation as a generic record-keeping exercise. For ML teams, this misses the operational significance of a question they face continuously: when does a model update, retraining run, or architecture change trigger a ‘substantial modification’ under Article 43(4), requiring updated technical documentation and potentially a new conformity assessment?

6.1 What Triggers “Substantial Modification” Under Article 43(4)

Article 43(4) provides that a substantial modification occurs when changes go beyond what was foreseen in the initial technical documentation and may affect compliance with the Chapter III, Section 2 requirements. For ML systems, this legal standard translates into several operational triggers that teams must evaluate:

- Model retraining with new data, particularly where the data distribution shifts significantly or new data types are introduced
- Algorithm changes, including architecture modifications, hyperparameter changes that alter model behavior, or switching between model families
- Intended purpose scope changes, including expansion to new use cases, user populations, or deployment contexts
- New data types or domains, such as adding multimodal inputs or processing data from previously unrepresented sources

(c) 2026 AI Governance Desk. All rights reserved.

- Fundamental architecture changes, including changes to the inference pipeline, integration interfaces, or oversight mechanisms

The critical distinction is not whether a change improves performance, but whether it alters the risk profile established during the initial conformity assessment. A model retrained on the same architecture with updated data from the same distribution may be non-substantial. The same retraining using data from a new demographic or geographic distribution may be substantial, because it introduces fairness and representativeness risks that were not evaluated in the original assessment.

6.2 The Three-Tier Classification System

ML teams benefit from a practical framework for classifying changes. The three-tier system below provides decision criteria and the corresponding documentation action.

Tier	Examples	Documentation Action
Non-substantial	Bug fixes, minor performance improvements within established thresholds, UI changes, logging enhancements, infrastructure scaling within documented parameters	Update change log only; no documentation revision required
Potentially substantial	Model retraining on updated data, algorithm modifications, oversight mechanism changes, new integration interfaces, threshold adjustments for decision boundaries	Conduct a formal impact assessment; update technical documentation if the assessment confirms substantial impact
Substantial	New intended purpose, significant performance changes that alter the risk profile, new data domains or types, fundamental architecture changes, removal or weakening of oversight measures	Trigger a new conformity assessment; update all affected Annex IV sections and the declaration of conformity

Table 3. Three-tier change classification framework for ML systems, derived from Article 43(4).

6.3 Common ML Team Misconceptions

COMPLIANCE RISK

Misconception: Retraining with the same architecture on updated data is always non-substantial. **Reality:** If the data distribution shifts significantly — through temporal drift, geographic expansion, or demographic change — the retraining may introduce risks not evaluated in the original conformity assessment. Classification depends on the nature of the data change, not the stability of the architecture.

COMPLIANCE RISK

Misconception: A/B testing a new model version in production does not require documentation updates because it is experimental. **Reality:** Production deployment of any model version, even for a limited A/B test, requires the technical documentation to reflect that version. If the test version is substantial, documentation must be updated before deployment, not after the test concludes.

(c) 2026 AI Governance Desk. All rights reserved.

COMPLIANCE RISK

Misconception: Automated retraining pipelines operating within predefined parameters do not require individual change assessments. **Reality:** While the pipeline itself may be documented as a predetermined change under Annex IV(2)(f), each retraining output must still be evaluated against the established parameters. A retraining run that produces a model falling outside documented performance thresholds must be classified and documented accordingly.

7

Cross-Framework Mapping: ISO/IEC 42001, ISO/IEC 23053 & NIST AI RMF

Organizations that have invested in AI management systems or risk management frameworks should not treat EU AI Act technical documentation as a parallel compliance exercise. Mapping Annex IV requirements against existing standards allows teams to extend and formalize artefacts they already maintain, rather than build separate regulatory documentation from scratch. The crosswalks below indicate where alignment is strong and where ML-specific gaps remain.

STANDARDS ALIGNMENT

ISO/IEC 42001:2023 is a licensed standard published by ISO; its full clause and control text is not freely publicly accessible. The control references below reflect commonly cited correspondences in published third-party compliance analysis and should be verified against the licensed standard text or a qualified ISO 42001 lead auditor before being relied upon for certification or conformity assessment purposes.

7.1 ISO/IEC 42001 to Annex IV Alignment

ISO/IEC 42001:2023 specifies requirements for establishing, implementing, maintaining, and continually improving an AI management system. Organizations pursuing or holding ISO 42001 certification typically have management-system infrastructure that supports several Annex IV sections, though gaps remain that require ML-specific supplementation.

Annex IV Section	ISO 42001 Coverage (Indicative)	Gap for ML Teams
1. General description	Purpose definition and AI system information controls	Strong alignment; ISO purpose-definition and system-information requirements map well to Annex IV(1).
2. Development process	Verification, validation, and AI system documentation controls	Strong alignment for process documentation; ISO 42001 generally does not require the algorithmic and architectural granularity Annex IV(2) demands for ML systems.

(c) 2026 AI Governance Desk. All rights reserved.

Annex IV Section	ISO 42001 Coverage (Indicative)	Gap for ML Teams
3. Monitoring and control	Monitoring and human oversight controls	Moderate alignment at the management-system level; ML-specific detail on input specifications, drift detection, and continuous-learning monitoring is typically not explicit.
4. Performance metrics	Verification and validation controls	Moderate alignment; ISO 42001 does not explicitly require the metric-selection rationale Annex IV(4) mandates, so teams generally must supplement with documented appropriateness justification.
5. Risk management system	AI risk assessment and treatment controls	Strong alignment with Article 9 and Annex IV(5); the AI Act's specific risk categories (fundamental rights, health and safety) may require extending the ISO risk taxonomy.
6. Lifecycle changes	Operational planning and control	Moderate alignment; ISO 42001 does not include the Article 43(4) substantial-modification framework, so teams generally must add ML-specific change classification and impact-assessment procedures.
7. Harmonised standards	Policy controls	Strong alignment; standards adoption and alternative-solution documentation are generally covered by ISO 42001 policy requirements.
8. EU declaration of conformity	Not covered	Gap; the EU declaration of conformity is a regulatory-specific requirement with no direct ISO 42001 equivalent and must be produced separately.

Table 4. Indicative ISO/IEC 42001 to Annex IV crosswalk. Verify control references against the licensed ISO/IEC 42001:2023 text.

The most significant gap for ML teams is the absence of ML-specific technical detail in ISO 42001, which addresses AI management at the organizational level rather than the algorithmic level. An organization with ISO 42001 certification will generally have governance infrastructure supporting Annex IV Sections 1, 5, 7, and 8, but will still need to produce the detailed architecture, data, and validation documentation required by Sections 2, 3, 4, and 6.

7.2 ISO/IEC 23053 as a Technical Foundation

ISO/IEC 23053:2022 provides a framework for AI systems using machine learning, specifying the functional blocks and interfaces that constitute an ML system lifecycle. For ML teams, it provides structural vocabulary for documenting system architecture and development process in a way that aligns with both engineering practice and regulatory expectations relevant to Annex IV Section 2.

(c) 2026 AI Governance Desk. All rights reserved.

The standard's functional blocks for data ingestion, preprocessing, training, evaluation, and deployment map naturally to the six ML lifecycle phases described in Section 5 of this briefing. Organizations using ISO/IEC 23053 as a technical architecture reference will generally find their existing pipeline documentation — data flow diagrams, component specifications, interface definitions — can be extended to satisfy Annex IV(2) with comparatively modest additions, primarily around bias examination methodology, risk linkage, and human oversight design.

7.3 NIST AI RMF Documentation Crosswalk

The NIST AI Risk Management Framework (AI RMF 1.0) organizes AI governance around four functions: Govern, Map, Measure, and Manage. These functions align with Annex IV sections in ways that enable organizations already using the NIST framework to extend their existing documentation for EU AI Act purposes.

NIST AI RMF Function	Annex IV Alignment	Documentation Extension Generally Needed
Govern	Annex IV Section 5 (risk management integration)	NIST Govern covers organizational risk culture and policy; extend with explicit linkage to Article 9 risk management requirements and EU-specific risk categories such as fundamental rights.
Map	Annex IV Sections 1 and 2 (system context, intended purpose, development process)	NIST Map covers system context and intended use; extend with detailed algorithmic and architectural documentation, data governance records, and third-party component inventory.
Measure	Annex IV Sections 2 and 4 (validation, metrics, testing)	NIST Measure covers evaluation and metrics; extend with metric-appropriateness rationale, subgroup disaggregation, robustness testing, and cybersecurity assessment documentation.
Manage	Annex IV Section 6 and post-market monitoring (change management, lifecycle updates)	NIST Manage covers risk response and monitoring; extend with a substantial-modification classification framework, version traceability, and an explicit post-market monitoring plan.

Table 5. NIST AI RMF function to Annex IV crosswalk. Source: NIST AI Risk Management Framework (AI RMF 1.0).

The crosswalk indicates that the NIST AI RMF provides a solid organizational and methodological foundation for EU AI Act technical documentation, but the regulatory specificity of the Act — particularly the Annex IV section structure, the ten-year retention requirement, and the conformity assessment connection — generally requires documentation extensions beyond the NIST framework's scope.

8

The SME and Start-up Simplified Documentation Pathway

(c) 2026 AI Governance Desk. All rights reserved.

8.1 What the AI Act Actually Says About Simplification

Article 11(1), second subparagraph, states that SMEs and start-ups may provide the information required under Annex IV in a simplified manner, using a form to be established by the European Commission. Notified bodies must accept this simplified form for conformity assessment purposes. As of this briefing's publication, the Commission's definitive simplified form had not yet been published, though implementing measures were anticipated.

The provision contains two important limitations. First, simplification applies to the *form* of documentation, not its *substance*. Second, the simplified form is optional: SMEs may choose to provide full Annex IV documentation instead, depending on the organization's documentation maturity and the complexity of the system in question.

8.2 What Simplification Does and Does Not Mean

For ML teams at SMEs and start-ups, the distinction between form and substance is the critical operational question. Simplification generally means:

- Condensed format using streamlined templates rather than extensive narrative documents
- Reduced administrative burden through structured forms with predefined fields
- Integrated presentation, where a single template captures multiple related requirements
- Automated generation potential, where tooling can populate fields from existing pipeline metadata

Simplification does **not** mean:

- Reduced substantive requirements for data governance, risk management, or testing evidence
- Exemption from bias examination, risk assessment, or validation documentation
- Permission to omit Annex IV sections the team finds inconvenient
- Acceptance of incomplete or retrospectively assembled documentation

A practical approach for SME ML teams is to use existing tooling — model cards, data cards, experiment-tracking reports, and automated pipeline documentation — to populate the simplified form. Experiment-tracking platforms can generate structured outputs that, with modest extension, satisfy the substantive requirements of Annex IV while fitting a simplified format.

REGULATORY CAUTION

The most consequential misconception among SME teams is that the simplified form represents a reduced compliance burden. The underlying evidence requirements under Articles 9, 10, 12, 13, and 14 are not reduced for SMEs. The simplified form is a packaging mechanism, not a substantive exemption, and a national competent authority or notified body evaluating a simplified form will still verify that the evidence behind each field exists and is defensible.

Organizations best positioned to use the simplified form effectively are those that already maintain disciplined documentation practices — model cards, data lineage, experiment tracking — and can map those practices to the simplified template fields with minimal additional effort. Teams should treat simplification as an opportunity to integrate compliance into existing workflows, not as a reason to defer

(c) 2026 AI Governance Desk. All rights reserved.

documentation investment.

9

Common Compliance Mistakes and Governance Gaps

Even teams that understand Annex IV requirements operationally fall into predictable patterns that produce incomplete or indefensible documentation. The following five gaps recur in conformity assessment preparation and audit readiness reviews.

9.1 Documentation as Afterthought

The most pervasive mistake is treating technical documentation as a post-development compliance exercise: teams complete model development, validation, and deployment, then assign a junior engineer or compliance officer to “write up” the documentation. The result is generally incomplete, inconsistent, and vulnerable to scrutiny.

Operational impact. Retroactive documentation cannot establish contemporaneity, the standard that national competent authorities apply when evaluating whether records were created when decisions were made. A dataset card written months after model training provides no assurance that the documented selection criteria or bias examination methodology reflect the actual decisions made during development. **Correction.** Documentation should be treated as a lifecycle deliverable with the same scheduling, resourcing, and quality gates as code, models, and data, with each ML lifecycle phase gated on a documentation completion criterion.

9.2 The “We Have Model Cards” Fallacy

Model cards and data sheets are excellent starting points: they capture intended use, performance metrics, training data characteristics, and ethical considerations in a structured format. The fallacy is believing they cover the full scope of Annex IV.

Operational impact. Model cards typically address only portions of Annex IV Sections 2 and 4 — model architecture, performance metrics, dataset summary. They do not generally cover risk management (Section 5), change management (Section 6), harmonised standards (Section 7), the declaration of conformity (Section 8), or the detailed logging, oversight, and monitoring documentation required by Sections 1 and 3. A team with robust model cards and data sheets has typically covered roughly 30 to 40 percent of Annex IV requirements. **Correction.** Conduct a gap analysis mapping every field of existing model cards and data sheets to Annex IV sections, identify uncovered requirements, and build supplementary documentation processes for the gaps before assuming compliance.

9.3 Version Control Gaps

ML teams typically version-control code and sometimes model weights, but fewer version-control documentation with the same rigor. This creates a traceability problem that is consequential given Article 11's ten-year retention requirement and Annex IV(6)'s change documentation requirement.

Operational impact. Teams may be unable to demonstrate which documentation version corresponds to which deployed system version. If documentation has been overwritten, reorganized, or lost during tooling

(c) 2026 AI Governance Desk. All rights reserved.

migrations, the provider cannot satisfy a request for documentation covering a system placed on the market several years prior. **Correction.** Version-control documentation artefacts — dataset cards, architecture documents, risk registers, test reports — using the same systems and practices applied to code, tagged with the same release identifiers as model and code versions.

9.4 Weak Bias Examination Documentation

Many ML teams conduct bias testing but fail to document methodology, metrics, negative results, and mitigations with sufficient rigor. Article 10 requires bias examination with regard to the intended purpose, and Annex IV(2)(d) requires the technical documentation to describe the main characteristics of the datasets, including their appropriateness for the intended purpose.

Operational impact. Weak bias examination documentation under Article 10 produces an incomplete Annex IV Section 2, undermining the entire technical documentation package. Additional testing conducted during a conformity assessment does not remedy a deficiency in the contemporaneous documentation of the original examination. **Correction.** Treat bias examination as a documented research process: record the hypothesis tested, the metrics selected and why they suit the specific system and population, the methodology applied, the results including negative or inconclusive findings, the mitigations implemented, and the residual risk assessment — created when the examination is conducted, not reconstructed afterward.

9.5 Overlooking the “Foreseeable Misuse” Requirement

Annex IV(1) and (3) require documentation of intended purpose, capabilities, limitations, and foreseeable unintended outcomes. Teams frequently document intended purpose narrowly and precisely but omit consideration of reasonably foreseeable misuse scenarios.

Operational impact. A risk management system that does not address foreseeable misuse is incomplete under Article 9, and technical documentation that omits misuse scenarios fails to satisfy Annex IV(3)'s requirement to describe foreseeable unintended outcomes. **Correction.** Conduct a structured misuse analysis during the design phase, documented as part of the risk management file and referenced in the technical documentation, considering deployment contexts outside the intended environment, user populations with different skill levels or incentives, integration with systems that alter the risk profile, and adversarial use cases that exploit system limitations.

10

Operationalizing Documentation: Tools and Workflow Integration

Knowing what to document is insufficient. ML teams need workflows that generate documentation as a byproduct of engineering activity, not as a separate compliance task. The following approaches embed documentation requirements into existing development practices.

10.1 Documentation as Code

(c) 2026 AI Governance Desk. All rights reserved.

The principle of “documentation as code” treats technical documentation with the same version control, review, and automation standards as software. For ML teams, this means integrating documentation generation into CI/CD pipelines and using automated tooling to produce structured outputs from pipeline metadata.

Practical implementation includes configuring experiment tracking platforms to export model cards and experiment reports in formats that map to Annex IV fields, automating dataset card generation from data pipeline metadata, triggering documentation review gates when model versions are promoted to staging or production, and maintaining documentation templates in the same repository as code, subject to the same pull-request and review processes. The objective is to make documentation generation the default path rather than an opt-in activity: when a model is promoted to production, the deployment pipeline should validate that the documentation version matches the model version before allowing the promotion to proceed.

10.2 The Compliance Traceability Matrix

A compliance traceability matrix is a master document mapping each Annex IV requirement to the specific evidence artefacts that satisfy it, the location of those artefacts, and the individual responsible for maintaining them. This matrix serves as the single source of truth for audit readiness and conformity assessment preparation.

The matrix should include each Annex IV section and subsection, the corresponding article or regulatory requirement, the evidence artefact, the storage location, the named individual responsible, the date of last review, and the trigger for the next review — scheduled, change-driven, or incident-driven. Minimum review schedules should be established: annual reviews for stable systems, with ad-hoc reviews triggered by changes, incidents, or regulatory updates. The matrix itself should be reviewed quarterly to confirm that ownership remains current and that artefact locations have not changed due to tooling migrations or organizational restructuring.

10.3 Documentation Tool Categories

No single tool satisfies every Annex IV documentation requirement. The categories below address specific documentation needs and can be integrated into a unified workflow.

Category	Purpose	Representative Tool Types
Experiment tracking	Capture training runs, hyperparameters, metrics, and artefacts with reproducibility	ML experiment tracking platforms
Data versioning	Version datasets and track lineage from raw sources to training inputs	Data version control and lineage tools
Model cards	Generate structured model documentation from pipeline metadata	Model card toolkits and custom templates
Documentation generation	Produce and publish technical documentation with version control	Static-site documentation generators, automated pipeline documentation

(c) 2026 AI Governance Desk. All rights reserved.

Category	Purpose	Representative Tool Types
Risk management	Maintain risk registers, assessment records, and treatment documentation	Governance, risk, and compliance (GRC) platforms

Table 6. Documentation tooling categories relevant to Annex IV compliance workflows. Specific product selection should be driven by existing infrastructure and integration requirements rather than by this table alone.

11

Governance Implications and Compliance Reality

Technical documentation is the primary evidence package that national competent authorities and notified bodies evaluate when determining whether a high-risk AI system may lawfully remain on the EU market. Understanding what these bodies look for, and how documentation quality affects assessment outcomes, shapes how ML teams should prioritize their documentation investments.

11.1 What National Competent Authorities Will Look For

When a national competent authority requests technical documentation under Article 11, evaluation typically focuses on five areas that test the defensibility of the ML team's documentation practices.

- **Documentation existence and completeness.** Verification that all eight Annex IV sections are present and contain substantive, system-specific content rather than placeholder text or generic descriptions.
- **Documentation contemporaneity.** Examination of whether records were created when decisions were made or reconstructed afterward, using timestamped experiment logs, dated dataset cards, version-controlled architecture documents, and signed test reports as evidence of contemporaneous practice.
- **Bias examination depth and methodology.** Assessment of whether bias examination was conducted with rigor appropriate to the system's intended purpose and affected population, including whether negative results and mitigation rationale were documented.
- **Linkage between data and deployed model versions.** Tracing whether the datasets documented in Annex IV Section 2 are the actual datasets used to train the deployed model, with an unbroken chain from data sources to deployed model.
- **Special categories handling.** Where a system processes sensitive data, including special categories under GDPR used for bias examination or other purposes, examination of whether the legal basis, DPIA, and data minimization measures are documented with appropriate rigor.

11.2 The Conformity Assessment Connection

Technical documentation is the primary evidence package for both internal assessment under Annex VI and third-party assessment under Annex VII. The quality and completeness of the documentation directly affects the duration, cost, and outcome of the assessment process.

(c) 2026 AI Governance Desk. All rights reserved.

IMPLEMENTATION CHALLENGE

Incomplete documentation blocks CE marking and market access. A notified body conducting a third-party assessment under Annex VII cannot issue a certificate of conformity where the technical documentation package contains gaps. The provider must then suspend market placement, remediate the deficiencies, and undergo reassessment, incurring direct costs (assessment fees, legal review, engineering remediation) and indirect costs (market delay, reputational exposure, competitive disadvantage) that are frequently larger than the cost of building sound documentation practice in the first place.

Documentation quality also affects assessment duration. A complete, well-organized, and clearly traceable documentation package supports an efficient assessment; a fragmented or incomplete package tends to trigger iterative requests for clarification, extending the assessment timeline. For teams tracking the high-risk obligations described in Section 3 of this briefing, whichever date ultimately applies, assessment delays represent a direct threat to market readiness.

12

Conclusion: Why Regulatory Readiness Now Depends on Documentation Discipline

The EU AI Act's technical documentation requirements are not a legal overlay on engineering practice. They are an engineering discipline in their own right, requiring the same rigor, tooling, and workflow integration that ML teams apply to model development and deployment. Teams that treat documentation as a first-class deliverable — embedding it into data engineering, model development, validation, and monitoring workflows — will be positioned for compliance. Teams that treat it as a post-hoc compliance task will face gaps that cannot be closed under assessment pressure.

Whether the high-risk compliance timeline ultimately settles at 2 August 2026, 2 December 2027, or another date following formal adoption of the Digital Omnibus on AI, the documentation dependency chain described in Section 4 means compliance cannot be achieved in a single sprint. Data governance records must be established before model training. Risk management files must be completed before technical documentation is finalized. Validation reports must be produced before deployment. Change management frameworks must be operational before the first post-deployment update. Each of these activities requires time, tooling, and organizational discipline that cannot be compressed into a final compliance push.

The checklist provided in this briefing is a starting point, not a universal template. Each organization's systems, risk profiles, and governance maturity differ: a medical imaging system requires different validation documentation than a credit scoring model, and a system using continuous learning requires different change management documentation than a static model. The framework should be adapted to the specific context, but the underlying principle — that documentation is built incrementally across the ML lifecycle, not assembled at the end — applies broadly.

Teams should begin by conducting a gap analysis: map existing documentation practices against the six-phase checklist in Section 5, identify missing artefacts, and build the processes and tooling needed to generate them. They should establish a compliance traceability matrix before the first conformity assessment request arrives, integrate documentation generation into CI/CD pipelines, version-control

(c) 2026 AI Governance Desk. All rights reserved.

documentation with the same rigor as code and models, and review documentation practices on a regular schedule, treating documentation governance as a continuous operational discipline rather than a one-time compliance project.

GOVERNANCE NOTE

Organizations that invest in documentation discipline now are likely to find the benefit extends beyond satisfying the AI Act's specific requirements. The operational transparency and traceability built through this discipline supports faster incident response, more effective risk monitoring, and stronger accountability to the individuals and communities affected by deployed AI systems. The regulatory requirement is the immediate driver; the governance benefit is the more durable outcome.

13

Governance Reference Tables

The two reference tools below consolidate material from earlier sections into formats teams can use directly in audit preparation: a compliance readiness checklist organized by Annex IV section, and a documentation-to-evidence matrix template that operationalizes the compliance traceability matrix described in Section 10.2.

13.1 Compliance Readiness Checklist

Annex IV Section	Readiness Question	Primary Evidence Artefact
1. General description	Is the intended purpose statement specific, bounded, and testable, with system version and provider identity documented?	System overview document
2. Development process	Do architecture diagrams, data governance records, and design decision logs exist and tie to the deployed model version?	Architecture document, data catalogue, decision log
3. Monitoring and control	Are logging capabilities, human oversight design, and foreseeable misuse scenarios documented?	Oversight design assessment, logging specification
4. Performance metrics	Is there documented rationale for why the chosen metrics suit this system and population?	Performance evaluation report
5. Risk management	Does the risk management file address data risks, link findings to design decisions, and assess residual risk?	Risk management file, risk register
6. Lifecycle changes	Is there a change classification framework and a change log covering all modifications?	Change management log, impact assessments
7. Harmonised standards	Are applied standards (or alternative solutions) identified and current?	Standards reference list

(c) 2026 AI Governance Desk. All rights reserved.

Annex IV Section	Readiness Question	Primary Evidence Artefact
8. Declaration of conformity	Is the EU declaration of conformity drafted and consistent with the rest of the documentation package?	Declaration of conformity

Table 7. Compliance readiness checklist, organized by Annex IV section. Intended as a working audit-preparation tool, not a substitute for legal or conformity assessment review.

13.2 Documentation Evidence Matrix Template

The structure below is the field set teams can use to build the compliance traceability matrix described in Section 10.2, populated per evidence artefact.

Field	Purpose
Annex IV section / subsection	Identifies which regulatory requirement the row addresses
Corresponding article	Links the requirement to Articles 9, 10, 12, 13, or 14 where applicable
Evidence artefact	Names the specific document, report, log, or record that satisfies the requirement
Storage location	Records where the artefact is stored, including system and repository path
Responsible owner	Names the individual accountable for maintaining the artefact
Date of last review	Tracks recency of the evidence
Next review trigger	Specifies whether review is scheduled, change-driven, or incident-driven

Table 8. Documentation evidence matrix field template.

14

Frequently Asked Questions

Q. Does the Digital Omnibus deferral mean ML teams can pause Annex IV documentation work?

A. No. Even under the provisionally agreed deferral described in Section 3, the underlying engineering discipline this briefing describes — contemporaneous dataset cards, risk registers, validation reports — takes time to build correctly and cannot be reconstructed credibly after the fact. Teams that pause documentation work risk facing the same contemporaneity gap described in Section 9.1 when a compliance date eventually arrives, regardless of which date applies.

Q. What is the difference between a model card and Annex IV technical documentation?

A. A model card is a structured summary, typically covering intended use, performance metrics, and high-level training data characteristics. Annex IV technical documentation is a regulatory evidence package spanning eight sections, including risk management documentation, change management records, harmonised standards references, and the declaration of conformity. As discussed in Section 9.2, a model

(c) 2026 AI Governance Desk. All rights reserved.

card generally substantiates only part of Annex IV Sections 2 and 4.

Q. Who is legally required to produce Annex IV technical documentation?

A. Article 11 places the obligation on providers of high-risk AI systems — the entity that develops the system or has it developed and places it on the market under its own name. Deployers and other actors in the AI value chain have separate obligations under other articles of the Act, but Article 11 itself does not require them to draw up the Annex IV package.

Q. Can a single technical documentation package cover multiple versions of a continuously retrained model?

A. Annex IV(6) and the version traceability requirements described in Section 5, Phase 6 of this briefing indicate that each system version should be matched to its own documentation version, with a change log linking them. A single static package is unlikely to remain accurate for a continuously retrained system unless it is paired with disciplined, version-controlled change documentation for every retraining run.

Q. What happens if a notified body finds the documentation incomplete during conformity assessment?

A. As discussed in Section 11.2, an incomplete technical documentation package generally prevents a notified body from issuing a certificate of conformity under Annex VII. The provider must then remediate the deficiencies and undergo reassessment, which delays market placement and adds direct and indirect costs.

Q. Does ISO/IEC 42001 certification by itself satisfy EU AI Act technical documentation requirements?

A. Not on its own. As the crosswalk in Section 7.1 indicates, ISO/IEC 42001 certification generally provides strong organizational coverage for Annex IV Sections 1, 5, and 7, but ML-specific technical detail for architecture, data, validation, and change management (Sections 2, 3, 4, and 6) typically requires dedicated supplementation beyond what the management-system standard alone produces.

Q. How does the ten-year retention requirement interact with GDPR data minimization principles?

A. The Article 11/18 retention obligation applies to technical documentation, quality management records, and conformity assessment files — not necessarily to the raw personal data used in training, which remains separately subject to GDPR's data minimization and storage limitation principles. Teams should distinguish between retaining documentation artefacts about a dataset (generally required for ten years) and retaining the underlying personal data itself (governed by the applicable GDPR legal basis and retention schedule), and should consult a data protection specialist where the two obligations appear to conflict.

(c) 2026 AI Governance Desk. All rights reserved.

15

Sources and Disclaimer

Official Regulatory and Government Sources

- European Commission, AI Act Service Desk — “Article 11: Technical documentation,” ai-act-service-desk.ec.europa.eu/en/ai-act/article-11. Accessed June 2026.
- EUR-Lex — “Regulation (EU) 2024/1689 of the European Parliament and of the Council (Artificial Intelligence Act),” eur-lex.europa.eu, CELEX:32024R1689. Accessed June 2026.
- European Parliament, Legislative Train Schedule — “Digital Omnibus on AI,” europarl.europa.eu/legislative-train. Accessed June 2026.

Standards and Framework Sources

- International Organization for Standardization — “ISO/IEC 42001:2023, Information technology — Artificial intelligence — Management system,” iso.org/standard/42001. Accessed June 2026.
- International Organization for Standardization — “ISO/IEC 23053:2022, Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML),” iso.org/standard/74438.html. Accessed June 2026.
- National Institute of Standards and Technology — “AI Risk Management Framework (AI RMF 1.0),” nist.gov/itl/ai-risk-management-framework. Accessed June 2026.

Professional and Legal Commentary (Digital Omnibus Status)

The Digital Omnibus on AI was politically agreed but not yet formally adopted at the time of writing; the following secondary commentary was used to describe its status and should be treated as commentary rather than primary law until the Official Journal text is published.

- Hogan Lovells — “EU legislators agree to delay for high-risk AI rules,” hoganlovells.com. Accessed June 2026.
- Gibson Dunn — “EU AI Act Omnibus Agreement — Postponed High-Risk Deadlines and Other Key Changes,” gibsondunn.com. Accessed June 2026.
- DLA Piper — “The Digital AI Omnibus: Proposed deferral of high risk AI obligations under the AI Act,” knowledge.dlapiper.com. Accessed June 2026.

Other Authoritative References

- artificialintelligenceact.eu — consolidated unofficial text and article-by-article reference for Regulation (EU) 2024/1689, including Articles 9-14, 18, 43 and Annex IV. Accessed June 2026. Readers should confirm article text against the official EUR-Lex publication.

(c) 2026 AI Governance Desk. All rights reserved.

REGULATORY CAUTION

This guide is for informational purposes only and reflects publicly available information as of June 2026. AI laws, regulatory interpretations, technical standards, enforcement priorities, and institutional governance practices may change. Readers should verify current requirements through official regulatory sources, standards bodies, legal advisors, compliance professionals, or qualified governance specialists before making decisions. AIGovernanceDesk.com and its editorial team are not liable for decisions made based on this material. This document does not constitute legal, regulatory, technical, or compliance advice.

(c) 2026 AI Governance Desk. All rights reserved.

Thank You for Reading

This briefing was prepared by the AI Governance Desk editorial team to help ML and compliance practitioners operationalize EU AI Act Article 11 inside real engineering workflows.

AIGovernanceDesk.com

contact@aigovernancedesk.com

© 2026 AI Governance Desk. All rights reserved.