

CS-LEG-0005 · COMPLIANCESUITE DOCUMENTATION

EU AI Act Compliance Statement.

FIT-only redaction. Effective 2026-04-28.

DOCUMENT ID	VERSION	EFFECTIVE	OWNER
CS-LEG-0005	v1.0	2026-04-28	Legal & Compliance

Public — Documentation · Review cycle: On change

Control block and metadata anchor.

The control block identifies the document, its current revision, the regulated process it supports, and the people accountable for its lifecycle. Every value below is the source of truth for any downstream record, audit trail entry, or signature block.

DOCUMENT ID	CS-LEG-0005
TITLE	EU AI Act Compliance Statement
VERSION	v1.0
STATUS	FIT-CLEAN
EFFECTIVE DATE	2026-04-28
REVIEW CYCLE	On change
DOCUMENT OWNER	Legal & Compliance
CLASSIFICATION	Public — Documentation
RELATED RECORDS	/output/CS-LEG-0005_EU_AI_Act_Compliance_Statement.pdf
SUPERSEDES	— (initial release)

Sign-off table, ready for ink or e-signature.

The signatures below confirm review and authorisation of this document. Approvals must be recorded in chronological order. If the document is signed electronically, the e-signature record on the ComplianceSuite platform supersedes any handwritten entry on this page and carries the same legal weight under 21 CFR Part 11 and EU GMP Annex 11.

Role	Name	Function	Date	Signature
Author		Validation Lead		
Reviewer		Quality Assurance		
Reviewer		Process / System Owner		
Approver		Head of Quality		
Approver		Regulatory Affairs		

What's in this document.

01 — Document Control	—
02 — Approvals	—
03 — Contents	—
01 — What this edition covers	—
02 — What this edition does NOT cover	—
03 — Disclaimer	—
04 — AI Components in Scope	—
05 — Risk Classification under EU AI Act	—
06 — Transparency Obligations (Article 50)	—
07 — Risk & Quality Management	—
08 — Customer Responsibilities	—
09 — Code Reference	—
Revision History	—
Glossary & Abbreviations	—

What this edition covers.

- **AI Composer:** Multi-provider integration (OpenAI, Anthropic, Google AI) for drafting content fragments under explicit human review
- **Citation Retrieval:** RAG-based approach to candidate-finding from tenant SOP library
- **Transparency Requirements:** AI line markers in audit trail, inspection view, PDF watermarks
- **Data Protection:** No model training on customer data, inference in customer residency region
- **Composer Configuration:** Per-tenant and per-deliverable enable/disable possible

What this edition does **NOT** cover.

- **Every EU AI Act compliance qualification is customer-specific.** The customer is responsible for: (1) written confirmation of AI Act risk classification for its own use case; (2) manual human review of every AI output before sign-off; (3) documented composer configuration per tenant; (4) user training.

Disclaimer.

This compliance statement describes how ComplianceSuite intends to meet Regulation 2024/1689 (EU AI Act). It is **not legal advice**. Customers operating under this regulation remain responsible for confirming against their own legal assessment and use cases. The regulation phases in over multiple tranches; this edition follows the regulation text in adopted form, with provisions per the transitional plan.

AI Components in Scope.

AI Composer

Aspect	Implementation
Purpose	Drafts content fragments in regulated deliverables, marked as AI-drafted, pending explicit accept / edit / reject by author
Human in the Loop	Mandatory. Platform rejects submission of a deliverable while AI-drafted lines remain undecided.
Audit Trail	Prompt, model identifier & version, response, human decision are written to change audit trail (in the same transaction as deliverable edit).
Functional Boundaries	Composer can be disabled per tenant and per deliverable. No automatic blocking logic on professional-judgment sections.

Citation Retrieval

Aspect	Implementation
Purpose	Surfaces candidate citations from tenant SOP library or regulatory corpus
Human Confirmation	Every retrieval-based citation becomes part of the regulated record only after author confirmation (verified status) or — where tenant policy permits — through implicit acceptance upon sign-off of the parent deliverable.
Audit Trail	Query, retrieved candidates, selection are captured.

Risk Classification under EU AI Act.

EU AI Act Category	Applicability	Reasoning
Prohibited (Art. 5)	Not applicable	ComplianceSuite AI systems do not engage in prohibited practices (subliminal manipulation, vulnerable-group exploitation, social scoring, real-time biometric identification in public, etc.).
High-Risk (Annex III)	Usually not applicable for platform default use	Platform AI is decision support for regulated authoring, not autonomous decision-making. Platform's high-risk contexts (employment, law enforcement) are not supported use cases. Customers in high-risk contexts must conduct their own classification and potentially disable composer.
Limited Risk (Transparency, Art. 50)	Applicable	AI-drafted content carries explicit indication; users are informed that they interact with AI assistance; synthetic content is identifiable in audit trail and inspection view.
Minimal Risk	Default for non-composer features	Platform's non-AI features are subject to the regulation's general product expectations only.

Transparency Obligations (Article 50).

- **Disclosure:** Users are informed that they interact with an AI system through (i) composer's explicit invocation; (ii) AI-drafted line markers in the deliverable; (iii) citation provenance disclosure.
- **Identification of Synthetic Content:** AI-drafted lines carry explicit machine-readable markers in audit trail (`composer.draft`, `composer.accept`) and visual markers in inspection view.
- **Watermarking:** PDF exports of AI-assisted records carry inspection-ready watermark with AI-assisted-line ratio in the footer.
- **Information to Deployer:** This statement, AI Composer Guide (CS-DOC-0005), and platform documentation provide the transparency information referenced in Art. 13.

Risk & Quality Management.

- 01 **Risk Management System.** Composer and citation-retrieval pipeline are subject to ComplianceSuite QMS. Risks are documented and reviewed at each model upgrade.
- 02 **Data & Data Governance.** Models are not trained on customer data; inference occurs in the customer's data-residency region; inference logs are deleted after audit trail write.
- 03 **Technical Documentation.** Internal technical documentation of AI components is maintained per EU AI Act expectations and shared with customers under NDA where required.
- 04 **Record-Keeping.** Audit trail captures every prompt, model version, human decision — the record-keeping that the regulation expects.
- 05 **Accuracy, Robustness, Cybersecurity.** Model performance is monitored across releases; validation evidence is reproduced at each composer model upgrade; platform's general security posture (ISO 27001, SOC 2) covers AI components.
- 06 **Human Oversight.** Mandatory human review of every AI-drafted line; platform rejects deliverable advance with undecided AI lines.

Customer Responsibilities.

- **Classify the use case.** Written confirmation of how platform AI assistance fits into your own AI Act risk classification.
- **Configure the Composer.** Decide whether AI assistance is enabled per tenant; optionally pin model version; disable on specific deliverables.
- **Train users.** Users must understand what accept / edit / reject means, and that they remain the authors of the regulated record.
- **Maintain own technical documentation** if use case is high-risk; platform provides supporting evidence under NDA.
- **Notify ComplianceSuite** of AI Act-driven configuration changes via account-level audit trail entries; platform records the configuration decision with rationale.

Disable is always available. Customers at any stage of AI Act qualification can disable composer per tenant or per deliverable. The disable switch is a controlled action, recorded in account audit trail; re-enable requires the same controlled action.

Code Reference.

- **Server Actions for Composer:** `app/actions/change/` (composer.accept, composer.reject, composer.draft audit trail events)
- **Audit Trail Model:** `prisma/schema/prisma` — `AuditLog` with sequence number, action type, JSON payload
- **Multi-Provider Setup:** `.env` with `OpenAI_KEY`, `Anthropic_KEY`, `Google_AI_KEY` per account/tenant
- **Inspection View:** `components/change/InspectionView.tsx` — shows AI-drafted markers

REVISION HISTORY

Every change, tracked and signed.

Add one row for every controlled revision. Minor changes (typos, formatting) increment the patch version; substantive edits trigger a fresh review cycle and a new approver round.

Version	Date	Author	Summary of Change	Approver
1.0	2026-04-28	Documentation Team	FIT-only redaction limited to codebase-verified functionality.	Head of Documentation
—	—	—	Reserved for next revision. Do not delete this row.	—

GLOSSARY

Shared language, **no ambiguity.**

Definitions used throughout this document. Where a term has a specific meaning inside ComplianceSuite, the platform-specific definition takes precedence over the generic regulatory term.

CSV	Computerized Systems Validation
GAMP 5	Good Automated Manufacturing Practice, Edition 5 (2nd edition, 2022)
GxP	Good 'x' Practice — covers GMP, GLP, GCP, GDP, GVP
IQ / OQ / PQ	Installation / Operational / Performance Qualification
Part 11	21 CFR Part 11 — US FDA rule on electronic records and electronic signatures
Annex 11	EU GMP Annex 11 — EU rule on computerised systems
URS	User Requirements Specification
FRS	Functional Requirements Specification
RTM	Requirements Traceability Matrix
SOP	Standard Operating Procedure
ALCOA+	Attributable, Legible, Contemporaneous, Original, Accurate (+ Complete, Consistent, Enduring, Available)
ICH Q9	International Council for Harmonisation Quality Risk Management guideline

— End of document —