

CS-LEG-0005 · GXP-DESK DOCUMENTATION

EU AI Act Compliance Statement.

FIT-only redaction. Effective 2026-06-04.

DOCUMENT ID	VERSION	EFFECTIVE	OWNER
CS-LEG-0005	v1.0	2026-06-04	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	2026-06-04
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 GxP-Desk 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	2
02 — Approvals	3
03 — Contents	4
01 — What this version covers	5
02 — What this version does NOT cover	6
03 — Disclaimer	7
04 — AI Components in Scope	8
05 — Risk Classification under the EU AI Act	9
06 — Transparency Obligations (Article 50)	10
07 — Risk & Quality Management	11
08 — Customer Responsibilities	12
09 — Code Reference	13
Revision History	14
Glossary & Abbreviations	15

What this version 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 discovery from the tenant SOP library
- **Transparency requirements:** AI line markers in the audit trail, inspection view, PDF watermark
- **Data protection:** No model training on customer data, inference in the customer's residency region
- **Composer configuration:** Per-tenant and per-deliverable enable/disable possible

What this version does **NOT** cover.

- **Any EU AI Act compliance qualification is customer-specific.** The customer is responsible for: (1) written confirmation of the 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) training of users.

Disclaimer.

This compliance statement describes how GxP-Desk **intends** to meet Regulation 2024/1689 (EU AI Act). It is **not legal advice**. Customers operating under this regulation remain responsible for confirming compliance against their own legal assessment and use cases. The regulation phases in across multiple tranches; this version is based on the text of the regulation as adopted, with provisions following a 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 the author
Human in the Loop	Mandatory. The platform refuses to submit a deliverable while AI-drafted lines remain undecided.
Audit Trail	Prompt, model identifier & version, response, and human decision are written to the change audit trail (in the same transaction as the deliverable edit).
Functional Limits	The Composer can be disabled per tenant and per deliverable. No automatic lockout logic on professional-judgment sections.

Citation Retrieval

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

Risk Classification under the EU AI Act.

EU AI Act Category	Applicability	Reasoning
Prohibited (Art. 5)	Not applicable	GxP-Desk 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)	Generally not applicable for default platform use	The platform's AI is decision support for regulated authoring, not autonomous decision-making. The platform's high-risk contexts (employment, law enforcement) are not supported use cases. Customers in high-risk contexts must perform their own classification and, where necessary, disable the Composer.
Limited Risk (Transparency, Art. 50)	Applicable	AI-drafted content carries an explicit indication; users are informed that they are interacting with AI assistance; synthetic content is identifiable in the audit trail and inspection view.
Minimal Risk	Default for non-Composer features	The platform's non-AI features are subject only to the regulation's general product expectations.

Transparency Obligations (Article 50).

- **Disclosure:** Users are informed that they are interacting with an AI system through (i) the 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 the audit trail (`composer.draft`, `composer.accept`) and visual markers in the inspection view.
- **Watermarking:** PDF exports of AI-assisted records carry an inspection-ready watermark with the AI-assisted line ratio in the footer.
- **Information to Deployers:** This statement, the AI Composer Guide (CS-DOC-0005), and the platform documentation provide the transparency information referenced in Art. 13.

Risk & Quality Management.

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

Customer Responsibilities.

- **Classify the use case.** Written confirmation of how the platform's AI assistance fits into the customer's own AI Act risk classification.
- **Configure the Composer.** Decide whether AI assistance is enabled per tenant; pin the model version where appropriate; 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** where the use case is high-risk; the platform provides supporting evidence under NDA.
- **Notify GxP-Desk** of AI-Act-driven config changes via account-level audit-trail entries; the platform records the config decision with a rationale.

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

Code Reference.

- **Server actions for the Composer:** `app/actions/change/` (composer.accept, composer.reject, composer.draft audit-trail events)
- **Audit trail model:** `prisma/schema.prisma` — `AuditLog` with `sequenceNumber`, `actionType`, 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	2026-06-04	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 GxP-Desk, 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 —