VaultLM

VaultLM security overview

Data flow and security principles for confidential document AI

WHAT VAULTLM IS

A controlled workspace between confidential documents and AI. VaultLM applies access rights, data masking and pseudonymisation before external AI processing, then returns answers with cited sources and an audit trail.

1. Documents

Original files enter a governed vault with permissions and retention.

2. Protect

Sensitive identifiers are masked and pseudonymised.

3. Minimise

Only relevant and permitted excerpts are selected.

4. AI service

The chosen external AI service receives the minimum protected context.

5. Evidence

The answer returns with cited passages and recorded usage.

STAYS INSIDE VAULTLM

Original files, identity mappings, permissions and retention rules.

SENT FOR AI PROCESSING

Only the minimum relevant, permitted and pseudonymised excerpts.

SECURITY PRINCIPLES

Permission-bound

AI access never exceeds permission to open the source documents.

Fail-closed

If required masking is unavailable, the external AI request stops.

Document-grounded

External AI requests are tied to permitted vault sources, not an unrestricted prompt proxy.

Reviewable

Sources, route details and audit events support human review.

HONEST LIMITS

Reversible masking is pseudonymisation, not anonymisation. Automated detection is not infallible, so human review remains necessary for sensitive or high-impact decisions. VaultLM supports governance and evidence; it does not automatically guarantee legal compliance.

AI ACT TIMING

More AI Act provisions apply from 2 August 2026. High-risk rules have later deadlines: 2 December 2027 for stand-alone systems and 2 August 2028 for systems embedded in regulated products.