Confidencial.io Securely Enable AI

Turn AI from a compliance headache into a competitive weapon.

USE CASE

As enterprises adopt AI, unstructured data fuels everything from agentic workflows to LLM training - but most lack control and governance over its use, access, and safety. Once sensitive documents enter pipelines like RAG or fine-tuning, protections vanish, exposing them to insider misuse, external threats, and regulatory risk.

Confidencial's AI Guard protects knowledge bases and data stores before it enters AI with semantic-preserving encryption, multi-layer embedding obfuscation, and identity-aware access policies that travel with it everywhere.

THE CHALLENGE

AI is outpacing governance, with sensitive data entering pipelines and workflows unchecked.

- No guardrails at ingestion – Sensitive data flows into vector databases and embeddings without fine-grained controls.
- Embeddings can talk – A breached VDB can be inverted back to near-original text.
- Governance gaps – Policies that work at rest disappear the second data enters an AI workflow.
- Regulators are watching – EU AI Act, ISO 42001, HIPAA, GDPR... compliance now requires provable, enforceable controls inside AI.

Once data is in the model, it's too late to protect it.

Why it Matters

- **Invisible AI Risk:** Once sensitive data is embedded, misuse or leakage can happen without detection
- **Breach Fallout:** A single exposed vector database can reveal regulated or crown-jewel data, triggering massive fines and brand damage.
- **AI Adoption Roadblocks:** Without provable controls inside AI workflows, compliance teams can stall or block innovation.
- **Rising Regulatory Heat:** New mandates like the EU AI Act and ISO 42001 require enforceable, auditable governance within AI pipelines.

Confidencial

How Confidencial Mitigates AI Risk

Confidencial protects sensitive data at retrieval and generation time, the most common source of live AI data leaks.

Stop the Silent Leaks

Apply scalpel-level encryption and embedding obfuscation to keep regulated and crown-jewel data out of unauthorized AI outputs.

Unlock More Data for AI

Safely use sensitive content without coarse redaction or risky anonymization, preserving accuracy and context.

Prove You're in Control

Enforce identity-aware policies across files, embeddings, and prompts, with audit trails built in—not bolted on.

Protect Trust Everywhere

Security travels with your data across clouds, vendors, and pipelines, so speed never comes at the expense of safety.

BUSINESS IMPACT

Accelerate AI adoption by removing compliance barriers and enabling safe use of sensitive data in AI workflows.

Reduce risk and cost by preventing breaches, avoiding fines, and eliminating expensive rework.

Maximize data value by safely leveraging more unstructured content for AI-driven insights and automation.

Strengthen market trust by delivering AI that's secure by design and proven to protect customer and partner data.

Case in Point:

Microsoft AI Data Leak (2023)

Security researchers discovered a misconfigured Azure Blob Storage endpoint exposing 38TB of Microsoft AI training data, including private keys, passwords, unstructured content such as internal team messages and documentation, and other sensitive files. The exposed data could have been ingested into AI workflows without controls, risking unauthorized access and leakage of sensitive information.

***The lesson?** Without persistent, identity-aware protections, a single misconfiguration can expose both structured and unstructured crown-jewel data to the world.*

sales@confidencial.io