This document is for DPOs, CISOs, and legal teams evaluating TokenVeil before a deployment. It describes, point by point, the technical measures actually in place and how they connect to your GDPR obligations, the CNIL's guidance on generative AI, and the EU AI Act. We also say plainly what stays your responsibility or falls outside the product's scope: a compliance dossier that only says what's true is more useful than one that says everything you'd like to hear.

1. GDPR: DATA MINIMIZATION AND SECURITY OF PROCESSING

Article 5(1)(c): data minimization

TokenVeil intercepts every message before it reaches an AI provider and replaces identifying data (names, emails, IPs, IBANs, customer references, technical secrets, and so on) with neutral tokens. The AI provider never receives the real data, under any circumstance: it isn't an option you could forget to turn on, it's the system's default behavior.

Article 25: data protection by design and by default

The product is built to run entirely on your own infrastructure (on-premises server or a private cloud of your choice). No data, anonymized or not, passes through any infrastructure belonging to the vendor. Protection doesn't depend on a configuration someone has to maintain: it's structural, from the moment of deployment.

Article 32: security of processing

Measure	Implementation
Encryption at rest	The link between a token and the real value it replaces is encrypted (Fernet/AES) before being stored. Without the encryption key, unique to each deployment, the database is unreadable.
Pseudonymization	Within the meaning of GDPR Article 4(5): stored messages contain only the anonymized version, never the original text.
Access control	Local authentication or LDAP/Active Directory, admin and user roles, sessions stored in the database (survive a service restart).
Logging	Every anonymization is tracked (who, when, how many items per category) without ever recording the real value involved.

Article 30: records of processing activities

TokenVeil doesn't replace your records of processing activities, but it gives you usable facts to fill them in: the categories of data processed, the purpose (exchanging with a third-party AI model), and the security measures applied. Your DPO remains responsible for keeping the records up to date.

2. DO YOU NEED A DATA PROCESSING AGREEMENT (DPA) WITH US?

Most AI compliance tools are SaaS: your documents get uploaded to their servers, which makes the vendor a data processor under GDPR Article 28, and you need a signed DPA before you can use the tool.

TokenVeil works differently. It runs entirely on your own infrastructure. Joopin's Lab never receives, stores, or processes your end users' messages or documents: there is nothing on our side to process. That means TokenVeil itself doesn't put Joopin's Lab in the role of a GDPR processor for your data, and a DPA isn't needed for the core product.

The only thing that touches our infrastructure is license verification (an instance ID and a license key, no business data) and, if you use the contact form on this site, the email address you choose to leave. We can still sign a standard mutual NDA or fill out a vendor security questionnaire if your procurement process asks for one, just nothing that GDPR Article 28 makes mandatory.

3. CNIL GUIDANCE ON GENERATIVE AI

The CNIL recommends that any professional use of an external generative AI tool keep the data sent to a strict minimum, and favor architectures where sensitive data never leaves the organization's perimeter. That's exactly what TokenVeil does: it sits between your teams and the AI model, and only sends the model what has already been anonymized.

In practice, the categories of data that get anonymized are configurable by your administrator (names, IP addresses, customer references, internal identifiers, and so on), so they match the actual sensitivity of your business data instead of a generic list imposed by the vendor.

4. THE EU AI ACT

The AI Act distinguishes the provider of an AI system (here, Anthropic, OpenAI, Google, Mistral, and so on, depending on the model you choose) from the deployer, meaning your organization, which uses that system in the course of its activity. The obligations around transparency, human oversight, and traceability (notably Article 26 on deployer obligations) fall on you, not on TokenVeil or the model's vendor.

TokenVeil gives you concrete means to meet those obligations:

- **Human oversight**: an administrator defines and can change the anonymization rules at any time, without any involvement from TokenVeil's vendor.
- **Traceability**: an audit log of every exchange (user, timestamp, data categories processed), usable to demonstrate controlled use.
- **Reversibility of control**: you keep the ability to cut access to an AI provider, switch models, or disable the tool, without depending on TokenVeil's vendor.

5. DOCUMENTS AND ATTACHED FILES

An attached document (MD, TXT, DOCX, XLSX, PDF, including a scanned PDF via OCR) follows the same pipeline as a typed message: text extraction in memory, anonymization, then sending to the model. The original binary content is never stored in the database or on the data volume, and the file name (which may contain a real name) is never transmitted or displayed.

A redacted copy of a Word or Excel file can be generated to hand off to a third party: same format, sensitive content replaced by tokens, and Office metadata (author, last editor, embedded thumbnail) stripped systematically, not just the visible text.

Position on ISO 27001

Joopin's Lab, the publisher of TokenVeil, does not hold ISO 27001 certification and does not claim it: that would be inaccurate. But the question plays out differently here than for a typical SaaS service. TokenVeil does not store or host any data on its own behalf: everything runs on your infrastructure, within your existing certification scope. If you are already ISO 27001 certified, that certification covers TokenVeil the same way it covers your other internal applications, with no extra audit on the vendor's side. That's a direct consequence of the self-hosted architecture, not a certification we're claiming for ourselves.

What remains your responsibility

- Keeping your records of processing activities and, where relevant, carrying out a data protection impact assessment (DPIA).
- Informing your staff and, where applicable, your customers about the use of an AI tool.
- Choosing and governing the downstream AI model (Claude, GPT, Gemini, Mistral, or another): its own terms of use and GDPR/AI Act status remain yours to evaluate.
- The physical and logical security of the infrastructure TokenVeil is deployed on, since it remains yours.