




Your teams already use AI. The question is **what they send it.**

Production logs, customer data, infrastructure secrets, internal documents: the moment a message or a file goes out to Claude, ChatGPT, or Gemini, you lose control of that information. TokenVeil intercepts it before it's sent, replaces it with neutral tokens, then restores the response with the real values, without ever letting sensitive data reach the AI provider.

HOW IT WORKS



WHY TOKENVEIL

- **Self-hosted on your side**

Deployed on your own infrastructure. No data, even anonymized, leaves your perimeter without your consent.

- **Compatible with your existing AI**

Claude, Gemini, GPT, Mistral, or enterprise cloud billing via Vertex AI, Bedrock, and Azure OpenAI: your teams keep their usual tool.

- **Enterprise authentication**

LDAP/Active Directory, multi-team management with separate seat limits, admin and user roles.

- **Fine-grained settings per deployment**

Choose which categories of data to anonymize (internal IPs, names, customer references...) based on your business needs.

- **Your documents are covered too**

Attach a Word, Excel, or PDF file (even scanned, OCR built in): the content is anonymized before it reaches the AI. Need to hand a document to a third party? Export a redacted copy, same format, hidden metadata removed.

COMPATIBLE AI PROVIDERS

Claude

Gemini

OpenAI / GPT

Mistral AI

GitHub Models

Vertex AI

Amazon Bedrock

Azure OpenAI

Let's see if it fits your needs.

A 20 minute demo is enough to see TokenVeil running on a real case.

Marc Bourrel, Joopin's Lab
contact@tokenveil.eu
tokenveil.eu