Generative AI for your customers, without sending their sensitive data to the model.

Teams already use ChatGPT, Copilot or Claude, with or without permission. A ban does not hold, and letting everything through does not either. NodeGW sits between users and models: what is confidential is replaced by tokens before it leaves, then restored in the answer; what the question needs goes through.

WHAT THE GATEWAY DOES

Reversible pseudonymisation

Emails, phone numbers, IBANs, names, addresses, identifiers and company-specific patterns become tokens, restored in the answer.

Rules per profile

Per data type: pseudonymise, block, log or let through. Images and files handled per extension.

Vocabulary learning

The Learner proposes terms to protect or release based on real usage. An administrator decides.

Multi-model routing

OpenAI, Anthropic, Mistral or in-house models (Ollama, vLLM). Allowed targets per profile, automatic fallback.

Access and quotas

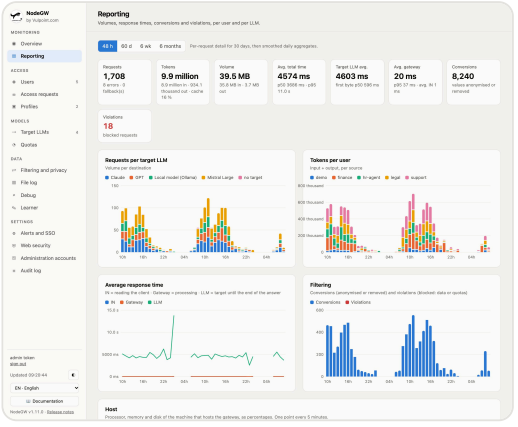
Sign-in with Microsoft Entra ID, Google or OpenID Connect, profiles by group, quotas per profile, user and model.

Traceability

Every request is measured: user, model, detections per type. No sensitive value is kept. Audit log of administration.

HOW IT WORKS

1. The user or application sends the request to the gateway.
2. The gateway identifies the user and applies their profile's rules.
3. Sensitive data becomes tokens; the rest goes through unchanged.
4. The request goes to the model allowed for that profile.
5. The answer comes back with the original values; the exchange is counted, without its content.



Reporting: who uses which model, how much, and what was filtered. Demonstration instance, fictitious data.

DEPLOYMENT

Linux server

A single binary, on an internal server or a European VPS. HTTPS with automatic certificate. Multi-core: capacity follows the load and the number of cores.

Mac

A menu bar app with the gateway built in. For a freelancer or a small team.

Windows workstations

Nothing to install on the workstation: the chat opens in the browser, and tools point to the Linux gateway.

WHAT IS IN IT FOR YOU

A demand that already exists

Your customers use generative AI and have to answer for what leaves their premises. NodeGW gives a technical, visible and measurable answer.

Recurring services

Setup, hosting and operations, rules per profile, follow-up of traceability and quotas.

Tuning specific to each customer

An organisation's confidential vocabulary is like no other. Keeping it up to date with the Learner is a value-added service.

A way into governance

A concrete starting point to discuss AI governance with the CISO, the DPO and management, and records to add to NIS2, GDPR, DORA and AI Act files.

No dependency on a model

The customer keeps the choice of provider, cloud or in-house, and can change it without touching the workstations.

Belgian detectors

National register numbers, INAMI / RIZIV numbers, towns, hospitals and care networks, FR / NL / EN first names. Interface in English, French and Dutch.

WHAT YOU SHOULD KNOW

- Names of people and companies are detected from lists or titles (Mr, Mrs, Dr). The lists are kept up to date, with the Learner's help.
- Image content is not analysed: per profile, an image goes through, is replaced by a blank square, or is refused.
- The gateway sees the data before processing it: it is hosted on the customer's premises or in the European Union, under a data processing agreement.
- Debug mode, off by default, keeps the original values for 24 hours to tune the rules.

THEY TRUST US

H.U.B — HÔPITAL UNIVERSITAIRE DE BRUXELLES — ACADEMISCH ZIEKENHUIS BRUSSEL

INSTITUT JULES BORDET INSTITUUT

Hôpital Erasme — ULB

Hôpital Universitaire des Enfants Reine Fabiola — Universitair Kinderziekenhuis Koningin Fabiola

NEXT STEPS

- A demonstration on a real case from your customer.
- A short pilot, with the reporting of what was filtered.

nodegw.com

WhatsApp