PRIVARO

Privacy Infrastructure for Enterprise AI

Enterprise AI Governance Layer

Intercept · Protect · Certify

Technical Architecture Overview — How Privaro intercepts, protects, and certifies every AI interaction your organization makes.

GDPR Art.5(2) | Blockchain Evidence | Zero-Trust PII | EU AI Act

Version 1.0 · March 2026

iCommunity Labs

privaro.io

Table of Contents

1 Executive Summary
2 The PII Gap in Enterprise AI
2.1 The AI Agent Risk
2.2 Threat Model Overview
3 Architecture Overview
3.1 Architecture Diagram
4 Data Flow — Request Lifecycle
5 Database Schema
6 Blockchain Evidence Layer (iBS)
7 Regulatory Compliance
8 Enterprise Use Cases
9 Role & Access Model
10 Competitive Landscape
11 SLAs & Technical Guarantees
12 Roadmap
13 Architecture Principles

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable

1 Executive Summary

Privaro is an enterprise AI governance infrastructure that operates as a transparent privacy proxy between enterprise applications and Large Language Models (LLMs). Every prompt passes through Privaro's semantic detection engine before reaching any AI provider — sensitive personal data is detected, tokenized, and replaced with pseudonymous placeholders. The LLM never sees real PII.

Every interception event is recorded as an immutable audit entry, cryptographically certified on the Polygon blockchain via iBS (iCommunity Blockchain Solutions), providing organizations with verifiable proof of GDPR Art.5(2) accountability by design — not by assertion.

~50ms	100%	SHA-512	GDPR · HIPAA PSD2 · FCA
Proxy latency added per request	PII coverage before reaching the LLM	Audit hash certified on Polygon	Regulations covered out of the box

2 The PII Gap in Enterprise AI

Enterprise adoption of AI is accelerating rapidly. However, organizations face a fundamental tension: the most valuable AI use cases — clinical notes, legal documents, HR data, financial analysis — are precisely those most likely to contain sensitive personal data.

The Three Failure Modes

Data Exposure

Prompts sent directly to LLM providers (OpenAI, Anthropic, Google) may contain PII that is logged, used for training, or stored in third-party infrastructure outside EU jurisdiction.

Compliance Blindness

No systematic record of what personal data was sent to which AI model, when, and by whom. Impossible to respond to GDPR Art.15 data subject requests or demonstrate Art.5(2) accountability.

Ungoverned Access

Individual employees connect personal API keys to LLM providers, bypassing organizational controls. No visibility into sensitive data flows across the organization.

Privaro's Position

Privaro addresses all three failure modes simultaneously: it intercepts every prompt at the network layer (exposure), creates an immutable blockchain-certified audit trail (compliance), and enforces organizational policies centrally for all AI pipelines (governance). No changes required to existing applications — Privaro is transparent by design.

2.1 The AI Agent Risk

The next phase of enterprise AI adoption is the deployment of autonomous AI agents capable of executing tasks across internal systems, APIs, and enterprise databases. Unlike simple prompt-based interactions, AI agents can initiate actions, retrieve records, generate reports, or trigger workflows autonomously.

This introduces a new governance challenge: organizations must ensure that AI agents cannot expose or process sensitive data outside approved policies.

Privaro extends its interception model beyond user prompts to agent pipelines. Any request generated by an AI agent passes through the same privacy and compliance controls:

- PII detection before execution
- Policy enforcement per data category
- Tokenization or anonymization when required

- Audit logging and blockchain evidence

This ensures that both human users and autonomous AI systems operate within the same governance framework.

Strategic Positioning

AI agent governance is the most critical frontier for enterprise AI compliance in 2025–2026. Privaro's architecture — designed as an interception layer rather than an application-level SDK — is uniquely positioned to govern both user-driven and autonomous AI interactions without requiring changes to the underlying agent frameworks.

2.2 Threat Model Overview

Threat Actor: Internal User

Risk: Sending sensitive data (PII, financial records, health data) directly to external LLM providers via enterprise applications.

Threat Actor: Shadow AI Usage

Risk: Employees using personal API keys to bypass organizational governance, creating ungoverned data flows outside any audit trail.

Threat Actor: AI Agent Pipelines

Risk: Automated workflows accessing sensitive records and triggering actions without human oversight or privacy controls.

Mitigation

Privaro's interception layer is applied before all AI interactions — whether initiated by human users or autonomous agents. All requests are subject to the same PII detection, policy enforcement, tokenization, and blockchain-certified audit trail regardless of origin.

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable

3 Architecture Overview

Privaro consists of five tightly integrated components, each with a single responsibility. The design follows a zero-trust model — no component trusts any other without verification.

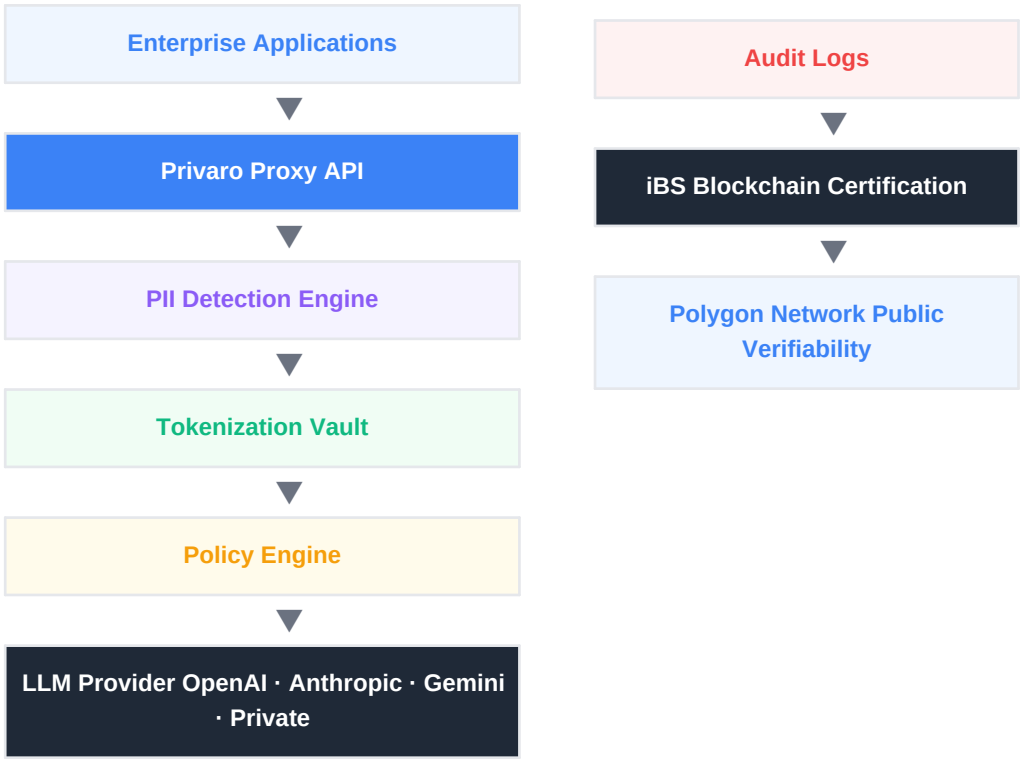
Component Map

Component	Technology	Responsibility
Proxy API	FastAPI (Python 3.11) Railway / Docker	Intercepts prompts, applies policies, manages tokenization flow
PII Detection Engine	Microsoft Presidio spaCy es_core_news_lg	Semantic detection of EU PII entities: DNI, IBAN, names, health records
Tokenization Vault	Supabase + AES-256-GCM Format-preserving encryption	Reversible pseudonymization with per-organization key segmentation
Blockchain Evidence	iBS API (iCommunity Labs) Polygon Network	Immutable SHA-512 certification of every audit event
Frontend Dashboard	React + Supabase Realtime Lovable	Chat interface, policy engine, audit logs, admin panel

3.1 Architecture Diagram

The following diagram illustrates Privaro's position within an enterprise AI pipeline.

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable



Parallel Layer: every request generates an Audit Log entry → certified on Polygon blockchain via iBS

4 Data Flow — Request Lifecycle

The following sequence describes the complete lifecycle of a single enterprise AI request through Privaro. The entire flow completes in approximately 50–100ms. Blockchain certification occurs asynchronously and does not block the response.

① **Enterprise App → Proxy API**
The enterprise application sends the prompt to Privaro's Proxy API (POST /v1/proxy/protect) using an organization-scoped API key (format: prvr_xxxx). TLS 1.3 enforced on all connections.

② **PII Detection**
The engine processes the prompt using Microsoft Presidio + spaCy (es_core_news_lg) for semantic entity recognition, augmented with custom regex for Spanish/EU entities: DNI, NIE, IBAN ES, SIP health records.

③ **Policy Evaluation**
Each detected entity is evaluated against the organization's policy rules — configurable per entity type, LLM provider, and regulatory framework. Actions: tokenise | pseudonymise | anonymise | block.

④ **Tokenization & Audit**
PII entities are replaced with format-preserving tokens ([NM-0001], [BK-0001]). The protected prompt, original hash (SHA-256), and detection metadata are written to audit_logs with ibs_status='pending'.

⑤ **LLM Forward**
The protected prompt — containing no real PII — is forwarded to the configured LLM provider. The provider's response is optionally re-identified for authorized DPO/admin users before being returned.

⑥ **Blockchain Certification (Async)**
Asynchronously, the Proxy API posts a SHA-512 hash of the audit payload to the iBS API. iBS certifies the hash on Polygon and sends a webhook (evidence.certified) updating ibs_status='certified' with the transaction hash.

Critical Design Decision: Async Blockchain

iBS does not return the transaction ID synchronously. The Proxy API responds to the client in ~50ms (steps 1-5). Blockchain certification arrives later via webhook (step 6). The ibs_sync_queue table manages automatic retries (max 3) if the webhook fails, alerting the DPO if any log remains uncertified after exhausting retries.

5 Database Schema

Privaro uses Supabase (PostgreSQL 15) with Row Level Security (RLS) enforced on all tables. The org_id field on every table guarantees strict multi-tenant isolation.

Table	Responsibility
organizations	Root of the multi-tenant model. Each enterprise client is one organization. Stores plan, DPO email, data region, and pipeline limits.
users / profiles	Extends Supabase Auth with enterprise roles (admin | dpo | developer | viewer) and org membership. Supports SSO/SAML.
pipelines	Each AI integration is a pipeline. Tracks provider, model, policy set, and aggregate privacy metrics (requests, PII detected, coverage %).
policy_rules	Per-organization privacy rules aligned to specific GDPR articles. Controls which entities are tokenized vs. anonymized vs. blocked per provider.
audit_logs	Immutable record of every PII event. Five ibs_* columns populated by the blockchain webhook. UPDATE restricted to ibs_* columns only via RLS.
pii_detections	Granular record of each detected entity within a request — position, confidence score, detector used, token reference.
tokens_vault	Secure store of reversible tokens. AES-256-GCM encryption with per-organization key segmentation. Access controlled by role.
api_keys	Organization API keys for enterprise application integration. SHA-256 hashed — plaintext never stored anywhere in the system.
ibs_sync_queue	Resilience layer for blockchain certification. Tracks pending certifications and manages retry logic with DPO alerting.
conversations / messages	Chat session persistence. Stores protected content only (tokenized) — original PII never written to the database.

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable

6 Blockchain Evidence Layer

Privaro partners with iBS (iCommunity Blockchain Solutions) to provide cryptographic proof that audit records have not been tampered with. This addresses GDPR Art.5(2) accountability requirement — organizations must be able to demonstrate compliance, not just assert it.

Why Blockchain for Audit Evidence?

A database audit log, no matter how well secured, can theoretically be altered by a database administrator. A blockchain-anchored hash cannot. When Privaro certifies an audit_log entry on Polygon, the SHA-512 fingerprint of that record is permanently and publicly verifiable at checker.icommunitylabs.com — independent of Privaro's own infrastructure.

Certification Details

Property	Value
Hash Algorithm	SHA-512 of the complete audit_log JSON payload, base64-encoded
Blockchain Network	Polygon (EVM-compatible, low-cost, high-throughput, carbon-neutral)
Evidence API	POST https://api.icommunitylabs.com/v2/evidences
Verification URL	https://checker.icommunitylabs.com/check/polygon/0x...
Certification Mode	Asynchronous — webhook evidence.certified callback
Retry Resilience	ibs_sync_queue: max 3 retries, DPO alert on exhaustion
Evidence ID Format	evd_ + 22 alphanumeric characters
TX Hash Example	0x13e0528d72783dee17f0b01f1f3190ed643e357747fe93...

7 Regulatory Compliance

Privaro's detection engine, policy framework, and audit architecture are designed to address the specific requirements of the following regulatory frameworks. Coverage is configurable per organization and per pipeline.

GDPR

EU General Data Protection Regulation

- Art.5(1)(f): Integrity and confidentiality of personal data
- Art.5(2): Accountability — verifiable proof via blockchain certification
- Art.25: Data protection by design and by default
- Art.30: Records of processing activities (audit_logs)
- Art.32: Technical measures — AES-256 encryption and pseudonymization
- Art.9: Special categories — health, biometric, and genetic data detection

HIPAA

Health Insurance Portability and Accountability Act

- PHI detection: patient names, medical record numbers, diagnoses, dates of birth
- Minimum necessary standard enforcement via configurable policy rules
- Audit controls: immutable log of all PHI access via AI with blockchain proof
- Transmission security: TLS 1.3 plus tokenization before any LLM call

PSD 2

Payment Services Directive 2 (EU)

- IBAN, card numbers, and payment references detected and tokenized automatically
- Strong customer authentication data never reaches LLM providers
- Audit trail for all financial data processing events with blockchain certification

FCA

Financial Conduct Authority (UK)

- Consumer Duty: protection of customer financial data in AI-assisted workflows
- Operational resilience: blockchain-certified evidence of privacy controls
- Senior Manager accountability: DPO dashboard with full audit log access

EU AI Act

EU Artificial Intelligence Act

- High-risk AI systems: mandatory fundamental rights impact assessment support
- Transparency: explainable PII protection with per-entity action audit logs
- Data governance: controlled data pipelines with centralized policy enforcement

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable

8 Enterprise Use Cases

Privaro is designed to protect enterprise AI workflows across multiple industries. The following examples illustrate how Privaro operates in real-world scenarios.

■■ Legal Document Review

Law firms increasingly use AI to analyze contracts and legal documents. Privaro ensures that client identities, addresses, and confidential clauses are tokenized before reaching the AI model. Attorneys interact with AI-generated analysis while sensitive data never leaves the tokenization vault.

■ Healthcare Documentation

Clinical notes often contain patient names, identification numbers, and diagnoses. Privaro detects and pseudonymizes Protected Health Information (PHI) before AI processing — ensuring HIPAA and GDPR Art.9 compliance for every interaction with clinical AI tools.

■ Financial Analysis

Financial institutions use AI to analyze customer portfolios, payment histories, and transaction data. Privaro automatically detects IBAN numbers, card references, and financial identifiers, replacing them with reversible tokens before forwarding to any LLM provider.

■ Customer Support Automation

AI assistants analyzing customer emails and support tickets frequently encounter personal data. Privaro ensures all customer identifiers are protected before AI analysis, enabling automation without privacy risk — fully compliant with GDPR Art.6 lawful basis requirements.

■ Enterprise Copilots

Internal copilots integrated with CRM, ERP, or HR systems can safely query enterprise data without exposing sensitive records to external AI providers. Privaro governs every data request — whether initiated by a human user or an autonomous agent workflow.

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable

9 Role & Access Model

Privaro implements a four-tier role model with cumulative permissions. Each role inherits all capabilities of the roles below it. Row Level Security (RLS) in Supabase enforces these boundaries at the database level — not just the application layer.

Role	Typical User	Key Capabilities
viewer	Business user, employee	Use the secure chat interface, view own conversation history, view dashboard metrics
developer	Engineer, integrator	Everything above + create and configure AI pipelines, manage technical pipeline settings
dpo	Data Protection Officer	Everything above + reveal tokenized values (with 2FA), export GDPR compliance reports, manage privacy notifications
admin	IT Admin, CTO	Everything above + manage users and roles, configure LLM provider catalog, API keys, billing, org settings

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable

10 Competitive Landscape

The enterprise AI privacy space is nascent. Most organizations currently address the problem with one of three inadequate approaches: blocking AI entirely, accepting the risk, or building custom internal solutions.

	Privaro	Block AI	Accept Risk	Build In-house	Generic DLP
Transparent to users	✓	–	✓	~	~
Real-time PII detection	✓	–	–	~	✓
Reversible tokenization	✓	–	–	~	–
Blockchain audit proof	✓	–	–	–	–
GDPR Art.5(2) ready	✓	~	–	~	~
Multi-LLM support	✓	–	✓	~	–
EU data residency	✓	✓	–	~	–
Time to deploy	**< 1 day**	< 1 day	0 days	6-18 months	1-3 months
Ongoing maintenance	**Zero**	Zero	High risk	High	Medium

✓ = Fully supported ~ = Partial / requires custom work – = Not supported

11 SLAs & Technical Guarantees

Privaro provides the following technical commitments for production deployments. Exact SLA terms are defined in the enterprise subscription agreement.

Performance

Proxy latency overhead	< 50ms p50, < 100ms p99 (detection + tokenization)
API availability	99.9% uptime SLA (enterprise plan)
Throughput	1,000 requests/minute per pipeline (scalable on request)
Cold start	< 2 seconds (Railway always-on deployment)

Security

Encryption in transit	TLS 1.3 enforced on all connections, no downgrade
Encryption at rest	AES-256-GCM for tokenized values, per-org key segmentation
API key security	SHA-256 hashed — plaintext never stored or logged
Audit log immutability	UPDATE restricted to ibs_* columns only via RLS policy

Data Residency

Default region	EU (Supabase EU West — Frankfurt data center)
LLM provider filtering	Admin can restrict to GDPR-compliant providers only
Data retention	Configurable per org (minimum 365 days recommended for GDPR)
Right to erasure	Token revocation + audit log anonymization on Art.17 request

Blockchain Evidence

Certification time	Typically 30–120 seconds after request (Polygon finality)
Retry resilience	Automatic retry x3 via ibs_sync_queue if webhook fails
Public verifiability	Any certified hash verifiable at checker.icommunitylabs.com
Failure handling	DPO alert + ibs_status='failed' if all retries exhausted

© 2026 iCommunity Labs · privaro.io · Public document — freely shareable

12 Roadmap

Privaro follows a 9-phase technical roadmap. Phases 0–4 are complete and deployed. Phases 5–7 are in active development. Phase 8 defines the long-term platform vision.

✓ **Phase 0–1 — Foundation**
Supabase schema + RLS, Lovable dashboard, FastAPI proxy, /proxy/detect and /protect with regex detection engine

✓ **Phase 2 — NLP Engine**
Microsoft Presidio + spaCy es_core_news_lg — semantic entity detection for Spanish/EU PII including names without keywords

✓ **Phase 3 — Tokens Vault**
AES-256-GCM tokenization with per-org key segmentation, /proxy/restore endpoint with 2FA and full audit trail

✓ **Phase 4 — Policy Engine + Admin**
Full CRUD policy rules with sector presets, LLM provider catalog, user management, API key management, billing

Phase 5 — Blockchain Evidence
iBS integration: POST /evidences, webhook receiver, ibs_sync_queue cron, Polygon checker link in Audit UI

Phase 6 — Full Proxy Forward
/proxy/forward: protect → forward to LLM → optional re-identification → response. GDPR Art.5(2) report generation

Phase 7 — Pilot
First enterprise client onboarding, E2E testing, risk and coverage report, blockchain certification validation

Phase 8 — Enterprise AI Governance Platform
Privaro evolves beyond a privacy proxy into a full enterprise AI governance platform: enterprise-wide AI usage monitoring, AI risk scoring for sensitive data exposure, automated compliance reporting, agent governance and autonomous workflow supervision, cross-model governance across multiple AI providers

13 Architecture Principles

Privaro's architecture is built around five core principles that guide every design decision, from the database schema to the blockchain evidence layer.

1

Interception before inference

Privacy controls are applied at the network layer — before any data reaches an AI model. This is architecturally superior to post-processing or application-level filtering, which can be bypassed.

2

Zero raw PII exposure

No plaintext personal data is ever transmitted to an external AI provider. Tokenization occurs within Privaro's trusted boundary. The original values are stored exclusively in the encrypted Tokenization Vault.

3

Verifiable compliance evidence

Audit records are not merely stored — they are cryptographically certified on a public blockchain. Compliance is demonstrable to regulators, auditors, and data subjects independently of Privaro's own infrastructure.

4

Provider-agnostic AI governance

Privaro enforces the same governance rules regardless of which LLM provider is used. Organizations can switch between OpenAI, Anthropic, Google, or private models without changing their compliance posture.

5

Minimal operational friction

Privaro is transparent to end users — no workflow changes, no new interfaces to learn. The ~50ms proxy latency is imperceptible. Privacy protection should not come at the cost of productivity.

Ready to evaluate Privaro?

Request a technical demo or integration consultation with the iCommunity Labs engineering team.

privaro.io · hello@icommunity.io