

Pryvasee.AI

The Pre-LLM Trust, Safety & Governance Layer

PRODUCT INFORMATION

August 2026

Executive Overview

Pryvasee.AI is a pre-LLM trust, safety, and governance middleware platform for consumers and enterprises. It sits between the user and any large language model, sanitizing sensitive data before it is ever transmitted, and evaluating the trustworthiness of what comes back, turning ad hoc AI use into a governed, auditable interaction.

Rather than monitoring AI activity after the fact, Pryvasee.AI intervenes at the point of interaction: the moment a person or system is about to send a prompt, document, or image to an AI model. This pre-LLM positioning is the platform's defining characteristic and its primary point of differentiation from conventional data-loss-prevention and post-hoc AI monitoring tools.

The Problem

As individuals and organisations adopt AI models such as ChatGPT, Gemini, Grok, and DeepSeek into daily workflows, three risks recur without adequate tooling in place:

- **Data leakage:** personally identifiable, health, and payment information (PII/PHI/PCI) is routinely pasted into prompts, uploaded in documents, or embedded in images, with no interception before it leaves the organisation's control.
- **Unverified output:** AI responses can be ungrounded, inconsistent across models, or outright hallucinated, with no built-in mechanism to assess how trustworthy a given answer is before it is acted upon.
- **No audit trail:** most AI interactions leave no durable, tamper-evident record, making it difficult for regulated organisations to demonstrate governance or investigate incidents after the fact.

The Solution

Pryvasee.AI addresses these risks with a single interception layer that works across any LLM provider, rather than a walled garden tied to one model vendor. Every interaction is screened, routed, scored, and logged before and after it reaches the underlying AI model.

Pryvasee Guard

DATA PROTECTION

Pre-LLM detection of PII, PHI, and PCI data across text, documents, and images, using OCR and content-safety screening. Sensitive content is identified and redacted before a prompt or upload ever reaches an AI model.

Pryvasee Thread

MULTI-LLM

Aggregates and compares responses across multiple leading LLM providers, including OpenAI, Gemini, Grok, and DeepSeek, from a single interface, so outputs can be cross-checked rather than taken on faith from one model.

Pryvasee Trust Engine

RESPONSE SCORING

Produces an AI Integrity Score (0–100) for each response, assessed across seven sub-metrics: groundedness, factual accuracy, hallucination risk, source attribution, contextual coherence, semantic consistency, and bias detection.

Adaptive Privacy Layer

DETECTION FRAMEWORK

A broader detection framework extending Guard's capabilities, incorporating OCR and image analysis to catch sensitive data embedded in screenshots, scans, and other non-text formats.

Family Safety Plan

CONSUMER

A consumer-facing safety product that extends Pryvasee's protection to households, helping families use AI tools without exposing personal or sensitive information.

AI Integrity Score, How Responses Are Assessed

Every response processed through Pryvasee Thread is evaluated by the Trust Engine across seven weighted sub-metrics before being returned to the user:

Sub-metric	What it assesses
Groundedness	Whether the response is supported by verifiable source content rather than the model's own inference
Factual accuracy	Correctness of claims against known or retrievable facts
Hallucination risk	Likelihood that the response contains fabricated or unsupported content
Source attribution	Whether claims can be traced back to an identifiable source
Contextual coherence	Consistency of the response with the surrounding conversation and prompt intent
Semantic consistency	Stability of meaning across repeated or rephrased queries

Bias detection	Presence of skew or one-sidedness in framing or conclusions
----------------	---

Technology & Architecture

Pryvasee.AI is built natively on Microsoft Azure as a containerized, cloud-native service. Every layer of the platform, from network ingress to data storage, runs within a single governed Azure tenant, rather than being assembled across disconnected third-party infrastructure.

Infrastructure stack

Layer	Technology	Purpose
Compute & orchestration	Azure Kubernetes Service (AKS)	Containerized, horizontally scalable service deployment
Edge & security	Azure Front Door with Web Application Firewall	Global request routing, DDoS mitigation, and layer-7 threat filtering
Identity	Microsoft Entra External ID	Customer identity and access management, with B2B identity federation for enterprise customers
Application data	Azure SQL Database	Core transactional data storage
Caching	Redis	Low-latency session and response caching
Secrets management	Azure Key Vault	Centralized encryption key and credential management
Deployment	Terraform, Flux, Azure DevOps	Version-controlled, GitOps-based infrastructure and release pipelines
Observability	Prometheus, Grafana, OpenSearch	Real-time monitoring, alerting, and log analysis across the platform

Detection engine

Pryvasee Guard's detection layer combines OCR with content-safety screening to identify sensitive data across text, documents, and images before it reaches an AI model. Detection covers PII, PHI, and PCI data by default, and the underlying framework is extensible to organisation-defined sensitive data categories for enterprise deployments.

Multi-model integration

Pryvasee Thread connects to multiple leading LLM providers, including OpenAI, Gemini, Grok, and DeepSeek, through a single unified interface. Requests are routed and responses retrieved without permanently tying a customer's data or workflow to any one model vendor.

Trust and audit layer

Every interaction is written to a tamper-evident audit trail built on Azure SQL Ledger technology, giving regulated organisations a cryptographically verifiable record for compliance reporting and incident investigation.

Azure-native security posture

Because the platform runs entirely within Microsoft's Azure ecosystem, using Managed Identity and a single governance surface, customer data does not need to cross cloud boundaries between separate providers. This simplifies security review and procurement for enterprise buyers already standardised on Microsoft infrastructure.

Platform & Availability

Pryvasee.AI is available as a web application and on iOS and Android, with the MVP live since July 2026. The platform is currently pre-revenue, with early registered users and an active enterprise proof-of-concept pipeline. A Microsoft Azure Marketplace listing is in the publishing and certification process under the Microsoft AI Cloud Partner Program (ISV).

Contact

Pryvasee AI Ltd | Company No. 16366859 | 71–75 Shelton Street, Covent Garden, London WC2H 9JQ, United Kingdom

Email: corporate@pryvasee.ai | Web: pryvasee.ai