STRICTLY CONFIDENTIAL - DO NOT DISTRIBUTE

Enterprise Security & Architecture

Trust, Data Privacy, and Secure AI Deployment

As artificial intelligence transforms customer engagement, Chief Information Security Officers (CISOs) and IT leaders face a critical mandate: deploying generative AI without compromising proprietary data or violating global privacy frameworks.

This technical brief outlines how your agency safely utilizes Large Language Models (LLMs) through a secure, isolated architecture designed specifically for enterprise environments.

1. Introduction: The Secure Use of LLMs

The primary security concern regarding AI adoption is data leakage—the fear that proprietary business data or Personally Identifiable Information (PII) might be ingested into a public model and regurgitated to unauthorized users.

Our platform mitigates this risk by fundamentally decoupling your business knowledge from the foundational language models. We do not use "fine-tuning" on public models; instead, we utilize an advanced architectural framework known as **Retrieval-Augmented Generation (RAG) and PageIndex.**

> **Core Principle:** Your data is used exclusively to provide context for your specific user's query. It is NEVER used to train, improve, or fine-tune public baseline models (such as OpenAI's GPT or Anthropic's Claude).

2. Data Handling & Privacy Mechanics

To ensure total data sovereignty, our platform enforces strict boundaries on how information is stored, processed, and transmitted:

- **Data Isolation (Siloing):** When you upload PDFs, documentation, or sync your website, that data is vectorized and stored in a secure, tenant-isolated vector database. Your data is cryptographically separated from all other clients.

- **Ephemeral Processing:** When a user asks a question, our system retrieves only the relevant paragraphs from your vector database. This fragment is sent securely to the LLM solely to format the response. The LLM retains zero memory of this transaction after the response is generated.
- **PII Redaction Engine:** (Optional Module) Automatically detects and masks sensitive data—such as credit card numbers or Social Security Numbers—before any text is processed by the generative engine.
- **Encryption at Rest & In Transit:** All data is secured using AES-256 encryption at rest and TLS 1.2+ for all data in transit.

3. High-Level Architecture

The diagram below illustrates the strict one-way flow of information, demonstrating how user queries interact with your isolated Knowledge Base without compromising data integrity.

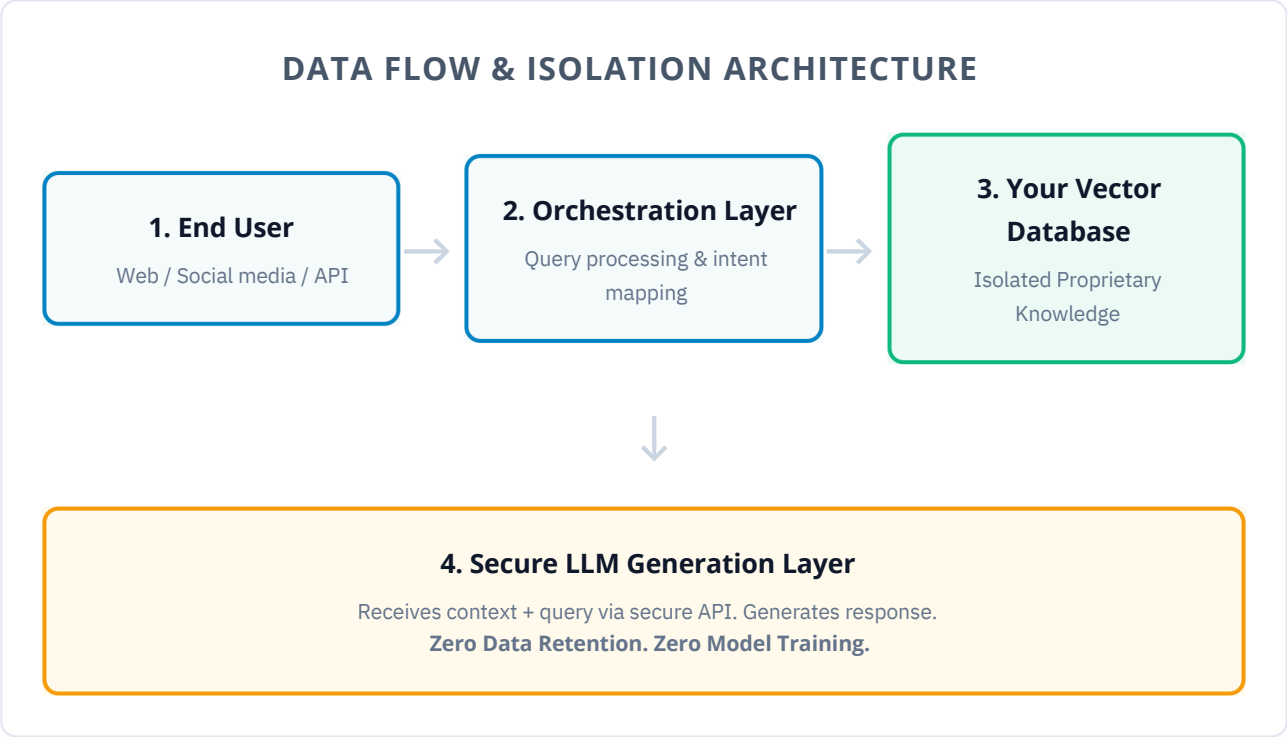

4. Compliance & Certifications

We build our infrastructure to meet the stringent requirements of heavily regulated industries, including healthcare, finance, and legal services.

Secure	**Penetration Testing** Platform adheres to Trust Services Criteria regarding security, availability, and confidentiality. We conduct regular penetration testing and vulnerability assessment.
GDPR	**General Data Protection Regulation (GDPR)** We are fully compliant with EU and UK data protection laws. Our platform supports localized data residency (EU servers available), granular consent tracking, and automated "Right to be Forgotten" (data deletion) requests through our dashboard API.
HIPAA	