OmniIndex Boudica

Powerful, Private, and Precise AI for Your Business

Move Beyond Hallucination

OmniIndex

White Paper

OmniIndex

Why Boudica?
The AI Trust Gap.

The adoption of Generative AI (GenAI) in the enterprise is hampered by fundamental security flaws, creating an "AI Trust Gap." Especially when it comes to those working with the most sensitive and regulated data.

This is because traditional Large Language Models (LLMs) pose three unacceptable risks for organizations handling confidential or regulated data (PII, IP, financial records):

Data Leakage
Hallucinations
Business Blindness

The Boudica Chat engine, leveraging OmniIndex's patented expertise, is engineered to eliminate these vulnerabilities by operating under a Zero Data Trust framework. It transforms data security from a compliance barrier into a competitive asset, enabling confident, secure AI-driven intelligence without hallucination.

BOUDICA

OmniIndex

Standard AI	
Data Leakage	PII/Proprietary data exposed to external models and third parties.
Hallucinations	False, contradictory, or biased answers leading to misinformed and costly decisions.
Business Blindness	Inability to securely access or analyze live, highly sensitive, or siloed company data.

Boudica AI	
Complete Control & Security	Data never leaves your environment and is never used for external training.
Verifiable Accuracy	Answers derived exclusively from your authenticated, encrypted data.
Full Context & Real-Time Analytics	Natively connects to all your live data sources while processing remains encrypted.

OmniIndex

Boudica AI: How It Works

Boudica Chat operates under OmniIndex's "never decrypt" philosophy, ensuring data is shielded from exposure throughout its entire lifecycle.

OmniIndex

The Security Cycle (Zero Data Trust)

Boudica implements a five-step security cycle that guarantees data integrity and privacy:

Encrypt Data (FHE): Data is secured using patented Fully Homomorphic Encryption (FHE) technology.

AI Analysis & Reports: The AI analyzes the encrypted data *without decryption* to securely create reports and GenAI outputs. This eliminates the vulnerability window traditional systems expose during processing.

No Permanent Storage: Data is not stored permanently within the chat environment and is immediately deleted after analysis, eliminating persistent third-party access risk.

No External Training: The AI is not trained on your data and cannot decrypt your encrypted data, preventing data exfiltration or absorption into the model.

Tamperproof Audit Trail: An immutable audit trail enables authorized administrators to check data access and use, providing necessary accountability for compliance frameworks.

The Boudica Chat interface is built on a minimal attack surface: it uses pure HTML, CSS, and JavaScript with no third-party libraries and strictly no application or third-party cookies. All client-side processing is minimal, and the core security logic is handled server-side.

OmniIndex

Moving Beyond Hallucination

Hallucinations (the generation of factually incorrect or inconsistent information) pose a critical cybersecurity risk, potentially causing vulnerabilities to be missed, audit logs and compliance checks to be misinterpreted, or business critical reports to be fatally flawed.

Boudica eliminates this risk through its unique SLM-Ontology Architecture:

Small Language Model (SLM) Architecture: Unlike massive LLMs trained on heterogeneous, uncontrolled data pools, Boudica is an SLM trained on algorithms and strict rules of conduct. It requires minimal data for initial training (focusing on rules) and operates exclusively on small, controlled, client-specific datasets. This focused approach inherently reduces bias and inconsistency.

Server-Side Context Management: Hallucination prevention is managed server-side to prevent client-side "context mash-up" or accidental user misunderstandings, ensuring only the most verifiable and probable answer is delivered based on the client's authoritative encrypted data.

Controlled & Confined Results: Boudica will never give an answer if there is not one. Instead, it will return an error or request for a different prompt or data source. This ensures no answer is ever forced or made up for the sake of giving one.

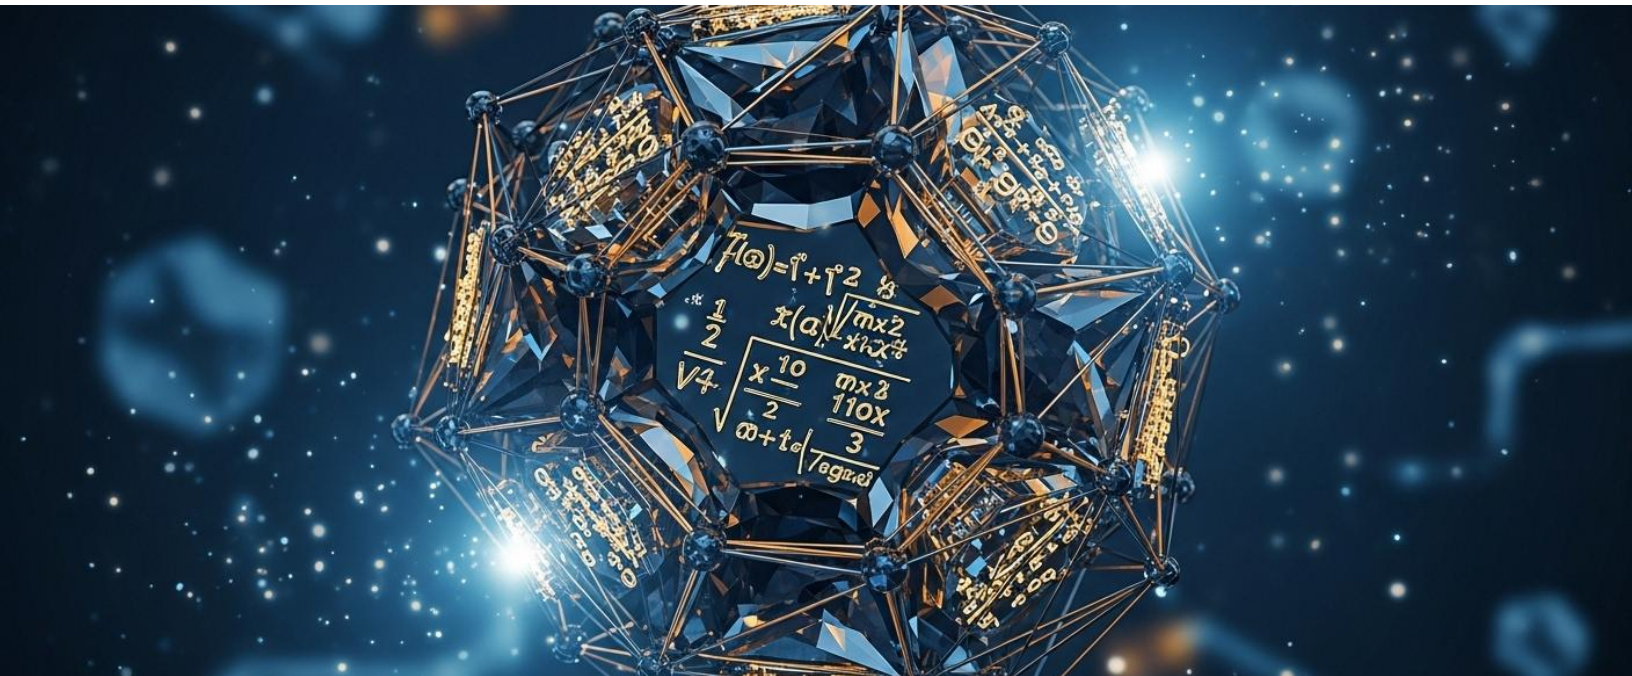

OmniIndex

Seamless and Secure Data Integration

Boudica Chat is engineered for enterprise adoption, enabling seamless, secure connectivity to existing business data sources without requiring massive, risky data migration projects through the Chat's native Bots.

These are read-only connections which enable authorized users to search and analyze their external content without risk of exposure or data leakage.

**OmniIndex are developing more connectors continually based on customer feedback.*

Box/Dropbox (File Storage)

Documents are subject to Additional Privacy Masking policies defined by JSON configuration (e.g., JavaScript supported regular expressions) before analysis.

HubSpot (CRM/Marketing Data)

Connector configurations restrict the dataset to specific, encrypted tables and limit the scope of the analysis. PII fields are automatically redacted.

Enterprise Databases (e.g., Postgres, Oracle)

The connection defines the database structure, enabling Boudica to create a query that acts only on the specific encrypted data block requested by the user.

OmniIndex

Secure LLM Firewall

The adoption of LLMs exposes organizations to a new, non-traditional attack surface.

The Boudica AI Engine functions as a sophisticated LLM Firewall, sitting as an intelligent proxy that inspects, analyzes, and sanitizes all interactions between users and the underlying language models. Allowing users to get the benefits of GenAI from an LLM without any of the exposure.

No sensitive data is shared with the LLM with interactions controlled by Boudica's unique ability to handle fully encrypted data. No files or user data is passed to the LLM and hallucinations are eliminated by Boudica.

Financial Fraud Analytics: Run real-time queries across encrypted customer transaction logs to detect fraud. Benefit: Maximizes detection speed and coverage without ever exposing sensitive PII to the AI, ensuring global compliance.

Confidential Legal Summarization: Instantly summarize complex, proprietary case files and contracts using natural language. Benefit: Provides accurate, verifiable intelligence derived only from encrypted documents, eliminating misleading legal advice from hallucinations.

Proactive Threat Hunting: Query encrypted server log files and access data using a natural language interface. Benefit: Delivers immediate, actionable threat intelligence on anomalous behavior, drastically accelerating incident response and protecting corporate infrastructure data.

The future of enterprise AI demands absolute security and verifiable accuracy. Boudica Chat delivers on both fronts.

By pioneering the use of Fully Homomorphic Encryption in a conversational AI interface, OmniIndex has removed the operational paralysis caused by the AI Trust Gap.

Boudica provides users with instant, trustworthy intelligence derived solely from their protected data, mitigating compliance risk, protecting intellectual property, and accelerating accurate decision-making. Adopting Boudica Chat is not merely a security measure; it is a strategic step towards seizing a competitive advantage in the new era of secure, private Generative AI.

OmniIndex

www.omniindex.io +1 (650) 297-4682

BOUDICA