Architecting OmniIndex PGBC

PGBC resolves the conflict between data security and utility. This technical introduction explores how our PostgreSQL fork enables analysis on data that remains fully encrypted and immune to attack.

OmniIndex

White Paper

The Headline

PGBC is an award-winning inoculated data platform created by OmniIndex. It is a web3 postgres fork enabling users to keep data encrypted at all times – even while searched & analyzed.

It is powered by 14 patents which cover areas of secure data management, AI & homomorphic encryption.

What's it for?

PGBC provides enhanced security and compliance for data management, as well as powerful AI productivity through its native SLM and ability to work with fully encrypted data all of the time.

PGBC is available as a DBaaS and also as a self-hosted platform, and it also powers OmniIndex's Dropblock file store.

Context:
What is Postgres?

PostgreSQL is an open-source relational database management system (RDBMS) that stores data in tables with defined relationships. This allows for efficient storage, retrieval, and analysis of complex data.

The system enforces data integrity through foreign keys and reduces redundancy by referencing existing data. Additionally, PostgreSQL offers robust security features like access control, data encryption, and auditing to safeguard sensitive information.

PostgreSQL has a large and varied user base, including industry giants and small businesses alike. Stack Overflow's developer survey further highlights this popularity, with nearly half (45%) of over 90,000 respondents reporting PostgreSQL use compared to 41% for MySQL.

What's more, because it is open source it is primed for innovation with developers able to build on the established strengths of the database to add their own advancements and bring out their own versions.

It is because of these secure and scalable foundations that OmniIndex chose to use it as the foundation for their own data platform: PGBC (The OmniIndex Postgres Blockchain).

OmniIndex

The OmniIndex PGBC Platform: *Inoculated. Intelligent. Yours.*

PGBC leverages standard PostgreSQL data structures but revolutionises the underlying storage and computational layers.

1. Immutable, Decentralized Storage: Data is not stored in mutable files but is committed to a decentralized chain as immutable blocks. This design prevents modification of existing records, offering inherent protection from data corruption and ransomware attacks.

2. Advanced Native Capabilities: PGBC integrates an SLM (Small Language Model) for built-in secure data management with automated key derivation & zer-trust verification. t also enables privacy-preserving analytics, allowing real-time search and computation on data that remains fully encrypted.

The following table provides an at-a-glance comparison:

Functionality	PostgreSQL	PGBC
Read/Write	Yes	Yes
Modify Data	Yes	No (on the block data)
Multi-Node Replication	No	Yes
Multi-Node Cluster Support	No	Yes
Peer to Peer Clustering	No	Yes
AI Support Out of the box	No	Yes
ML Support out of the box	No	Yes
Extensibility	Yes	Yes
Full SQL Support	Yes	Yes
Record Level Encryption Without Administrator Access	No	Yes
Search on Encrypted Data	No	Yes

Key Technical Features of PGBC

PGBC integrates AI-driven analytics with a blockchain foundation that makes data immutable and tamper-proof. This allows for advanced computations and insights to be generated directly on fully encrypted information, ensuring data remains secure & useful.

Blockchain Storage

PGBC is a Web3 data store where all data is recorded on a blockchain, yet managed with the simplicity of SQL.

This unique integration allows users to interact directly with the chain through familiar commands. For example, a new set of records can be cryptographically sealed and added to the chain using a command like: CREATE BLOCK new_transactions (user_id INT, sale_amount DECIMAL);

The security and privacy benefits of this model over legacy databases revolve around two core blockchain principles: data immutability and network decentralization.

1. Immutable Data, Guaranteed by Cryptography:

Every block of data is linked to the previous using a cryptographic hash. This means that once data is written, it is computationally impossible to modify or delete it without breaking the entire chain. This immutability provides a definitive defence against common threats like data corruption or ransomware, as an attacker has no mechanism to overwrite or encrypt existing, validated records.

2. Decentralized and Resilient Storage:

Instead of residing on a single server, the PGBC ledger is replicated across a network of nodes, eliminating any single point of failure. An attack on one node has no impact on the integrity of the overall network.

Furthermore, if a node is compromised or fails, its data can be instantly and reliably recovered from other peers in the network, ensuring zero data loss and continuous availability.

Homomorphic Encryption (FHE)

OmniIndex's patented encryption enables users to search and perform computations on fully encrypted structured & unstructured data in real-time.

Our FHE utilizes a unique architecture based on the principles of AES-256 to achieve real-time computation on encrypted data, a significant departure from traditional, slower models.

This means data does not have to be decrypted and left vulnerable to exposure through accidental or deliberate actions to be useful.

For example, Future-X Education in Nigeria is using PGBC to keep the PII of students protected and private while still being able to gain insights into educational outcomes.

It is also a core technology in Dropblock's threat and compliance intelligence for secure file storage. This is because it enables authorized users to analyze their fully encrypted log files to identify vulnerabilities and breaches, as well as to predict potential attacks.

For example, an AI risk assessment of our in-house logs holding over 1.5 million encrypted records pulling back 200,000 rows can be done in under 400 milliseconds.

> “The problem with encrypted data is that you must decrypt it in order to work with it. By doing so, it's vulnerable to the very things you were trying to protect it from by encrypting it.
>
> There is a powerful solution to this scenario: homomorphic encryption.”

FORBES

Native AI (Boudica)

Boudica is an SLM (Small Language Model) AI engine empowering users to gain insights on their encrypted data without any data ever being exposed or shared externally. This private AI engine has a number of direct benefits for users.

1: Privacy

Boudica is a private and native AI solution, meaning your data and queries are never shared externally or accessed by any third party.

This design guarantees true data sovereignty, as your information is processed entirely within your designated network or geographical boundary while encrypted. Furthermore, you retain complete authority and control through granular zero-trust access controls.

2: Accuracy

SLMs like Boudica only work on small pools of controlled data that they are supplied with and are therefore far less likely to contain biases and inaccuracies than LLMs which learn from huge pools of varied data. For instance, a financial institution using Boudica with a financial ontology will receive answers based solely on its own verified data, not external.

Boudica's use of multiple separate thesaurus models and its probability matrices also ensure that only the optimum response is given to a user.

3: Efficient & Optimized

SLMs do not require extensive training on huge pools of data in order to be used and can be adapted to offer specialist services in different languages and areas simply by changing their ontologies. They also require less powerful hardware to run.

Zero-Trust Access

Zero-Trust is a security framework built on a single, powerful principle: never trust, always verify. It dismisses the outdated idea of a secure network perimeter and assumes threats can exist anywhere.

To combat this, our native AI engine, Boudica, continuously monitors all activity, forcing every user, device, and application to prove its identity and authorisation in real-time before every interaction.

OmniIndex Zero-Trust includes:

Granular Access Control:

Users are granted the absolute minimum "least privilege" access required. Our patented Homomorphic Encryption (FHE) takes this a step further, enabling even administrators to perform analytics and manage the system without ever gaining the ability to view the raw, unencrypted data.

Constant Encryption:

Unlike traditional systems, data remains encrypted at all times: at rest, in transit, and crucially, during use. Each interaction automatically derives a unique, single-use cryptographic key, ensuring every query and transaction is independently secured. All analytics and AI functions are performed directly on this encrypted data, eliminating the primary vulnerability where most data breaches occur.

Immutable Audit Trail:

Every action—from login attempts to AI-driven queries—is recorded as a permanent, unalterable transaction on an immutable blockchain ledger. This provides a completely transparent and tamper-proof audit trail, which our AI continuously analyses to detect anomalous patterns and potential threats in real-time.

In an era where data is an organization's greatest asset and its biggest vulnerability, the PGBC platform was architected to resolve this fundamental conflict.

By unifying an immutable blockchain ledger , patented homomorphic encryption, a native AI engine, and a comprehensive Zero-Trust framework, it provides a single, cohesive solution.

The result is the definitive platform for the next generation of secure applications, where data remains tamper-proof, fully encrypted during analysis, and governed by intelligent, verifiable access controls.

OmniIndex's PGBC is the only commercial solution that makes this possible.

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