



Success Story

Real-Time Intelligence at Scale: How Evodian Powers BlaaS with the VAST DataBase



The ability to turn raw data into immediate action is the ultimate competitive advantage. As a managed service provider specializing in data management, analytics, and intelligent automation, Evodian integrates data from diverse business areas such as accounting, logistics, and retail systems to unlock these actionable insights.

Evodian found that the traditional cloud data warehouse model had become a bottleneck, with up to 12-hour latencies and layers of technical complexity that hindered operational agility. To solve this, they brought their data infrastructure inhouse and consolidated their entire Business Intelligence as a Service (BlaaS) stack on the VAST AI Operating System. By eliminating the friction between storage and compute, Evodian now powers real-time decision-making with million-row-per-second ingestion and a unified data lakehouse that scales as fast as their customers' ambitions.

For Aaron Pritchard, Technical Team Lead at Evodian, the mission was to build a data stack that was not only “best-in-class” but also sustainable for long-term growth. To achieve this, Evodian turned to the VAST AI OS to serve as the foundation of their modern data lakehouse.

The Challenge: The Performance and Cost Limits of the Cloud Data Warehouse

Before adopting VAST, Evodian relied on a traditional cloud data warehouse architecture. However, as the firm grew, the limitations of this model became a barrier to operational excellence.

The Latency Gap

The primary pain point was data freshness. Under the previous ETL-heavy model, the technical team saw data latency. “Our BI tools would only be able to see the data as it was 12 hours ago,” says Pritchard. “We want our data to be as current as possible. If you can get low latency on your data, you can make it operational, not just analytical.”



*It is really hard to
imagine a world
without VAST. It
enables us to do pretty
much everything we're
doing.*

Aaron Pritchard,
Technical Team Lead.



Architectural Friction

Operating in a private cloud environment while relying on a remote cloud data warehouse created a disconnect. Evodian's compute and consumption were local, yet their data was being sent to a different data center ten hours away just to be stored. This imposed an inherent bottleneck and complicated data governance responsibilities.

The Cost-Performance Paradox

While the total figures for their cloud bills were manageable, the performance was not proportional to the investment. Evodian found themselves in a cycle of debating efficacy versus expense. "As soon as I'm debating the quality of a solution versus the cost, that is an alarm bell for me," Pritchard explains. "I always want to ask, 'What is the best solution I can provide?' not 'What is the cheapest?'"

The Solution: A Unified Data Lakehouse on VAST

Evodian built what is now one of the most advanced modern data stacks on the market. By integrating dbt Labs, ThoughtSpot, and VAST Data, Evodian has created a seamless pipeline that moves from raw data to insights without the legacy overhead of traditional providers. This transition allowed them to shift from a traditional ETL model to a more agile ELT (Extract, Load, Transform) workflow, utilizing the VAST AI OS as the central repository for their data, consolidating structured data in the VAST DataBase and unstructured data in the VAST DataStore.

The BlaaS Stack: From Ingestion to Insight

- **Data Ingestion:** Evodian utilizes a variety of protocols, including VAST's many SDKs, Kafka-compatible APIs and S3 endpoints, to ingest data from ERPs, CRMs, and external sources like the Bureau of Statistics and meteorological providers.
- **Data Processing and Transformation:** The firm uses dbt in combination with streaming technologies to transform raw data directly within the VAST DataBase, ensuring high performance and data integrity. This rapid execution is powered by VAST's unique write path, which instantaneously commits dbt schema updates and transactions to persistent memory (SCM) before intelligently migrating the data into an exabyte-scale columnar flash layout optimized for immediate analytical queries.
- **Analytics with Open Query Engines:** Trino acts as the high-speed query engine, serving as the intermediary between the VAST DataBase and the analytics layer.
- **Self-Service Analytics:** At the head of the stack is ThoughtSpot, which provides a search-driven BI experience. This allows end-users, from factory floor managers to purchasing teams, to perform ad-hoc queries without needing specialized data analysis skills.



When we looked at the leadership and the vision behind VAST, we were confident that it would become something special. Even in the early stages, it was clear that the team was building ahead of the curve. We chose to align ourselves with that trajectory, knowing that as VAST grew, our capabilities would grow with them.

Aaron Pritchard, Technical Team Lead, Evodian.

The Results: Operational Simplicity and Massive Throughput

VAST has simplified Evodian's end-to-end data pipeline, allowing a lean team to manage complex, multi-tenant environments without the operational overhead of maintaining disparate systems like standalone Kafka clusters. "Every additional component in our stack adds complexity. With each release from VAST, we realize there's another thing we can consolidate," says Pritchard.

1

Breakthrough Performance

By leveraging VAST's unified architecture, Evodian has eliminated the throughput bottlenecks of their previous architecture.

We were usually bottlenecked in the past, getting maybe 10,000 rows per second. With an in-house solution on VAST, we are now ingesting over one million rows per second.

Aaron Pritchard, Technical Team Lead, Evodian.

This staggering leap from 10,000 to over 1,000,000 rows per second is made possible by the unique VAST Disaggregated Shared-Everything (DASE) architecture. Unlike traditional data warehouses built on legacy "shared-nothing" architectures, which rely on complex clustering that causes debilitating data serialization, hotspots, and noisy-neighbor traffic as data scales, DASE shatters these limitations. By completely decoupling stateless compute logic from a stateful, globally accessible pool of flash media over an ultra-low latency NVMe fabric, VAST allows every single compute node to see the entire storage estate. This eliminates the intra-cluster communication overhead, allowing Evodian to scale performance and capacity completely independently while easily maintaining a smooth, parallel distributed path for million-row-per-second ingestion.

2

Operational BI

With near real-time data, BI has moved from a passive reporting tool to an active operational driver. For example, purchasing teams can now see sales trends instantly, allowing them to anticipate stock needs and adjust habits dynamically rather than waiting 12 hours for a report.

3

Cost Governance and Scale

By keeping data within their private cloud, Evodian has eliminated the unpredictable ingress and egress costs of public cloud providers. This transparency allows them to offer their customers a more stable and secure managed service.

An Extension of the Team: VAST Support

The relationship between Evodian and VAST extends far beyond the technology, functioning more as a strategic partnership than a vendor-customer transaction. Aaron Pritchard views the VAST support team as a transparent, high-touch extension of his own technical department. By utilizing integrated Slack channels for real-time communication and direct access to R&D, Evodian is able to troubleshoot complex architectural questions without the friction of traditional support tiers. This collaborative environment ensures that the platform is not just maintained, but continuously optimized as Evodian adopts new features like the VAST DataEngine.

"I love the VAST support team. They feel like an extension of our own technical team. Honestly, I couldn't ask for more than that."

Aaron Pritchard, Technical Team Lead, Evodian

Future Outlook: IoT and the VAST DataEngine

Looking ahead, Evodian is preparing to integrate high-volume IoT data into their platform, particularly for customers in the manufacturing vertical. Utilizing the VAST Event Broker, they plan to use custom sensors and RFID scanners to provide real-time visibility into factory floor machine processes.

The technical framework behind this upcoming architecture relies on the clientless design of the VAST Event Broker. Natively compatible with Kafka APIs, it ingests streaming telemetry directly into VAST DataBase tables without requiring external Kafka clusters, ZooKeeper instances, or brittle stream-processing layers. Because events are natively indexed into structured tables upon arrival, they avoid "east-west" network serialization traffic entirely, transforming real-time streams into instantly queryable historical data on a single platform.

The team is also architecting new solutions around the VAST DataEngine. One upcoming use case involves the automated processing of computer-aided designs; using the DataEngine to instantly trigger transformation pipelines integrated with ERP systems. Another would see the DataEngine run machine learning models across incoming proof-of-delivery media, pinpointing and extracting useful information without delay.



VAST is not just storage to us; it's an ecosystem. It is the foundation that allows us to solve our customers' problems with a truly dynamic, best-in-class stack.

Aaron Pritchard, Technical Team Lead, Evodian.