

Vector Analytics Database

Faster Analytics Across Your Data Center and Clouds

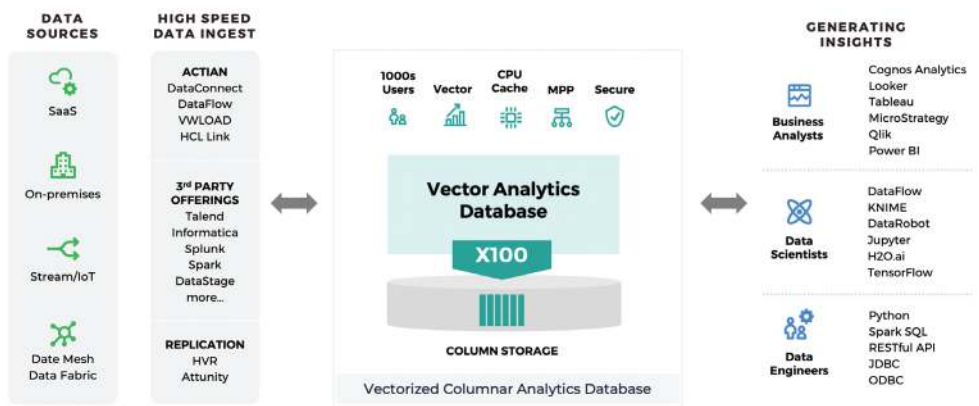
KEY BENEFITS

- Run complex queries against billions of records on commodity hardware in seconds
- Leverage vector processing (SIMD) to maximize performance of complex analytic queries.
- Increase SQL performance at least 100% over traditional tools
- Execute updates without performance penalty
- Exploit dedicated CPU core and caches running 100x faster than RAM
- Scan data faster using self-indexed storage
- Use native Spark-powered direct queries

FLEXIBLE DEPLOYMENT

- Linux
- Windows
- Google
- AWS
- Azure
- Hadoop

The Vector Analytics Database makes it easy for business users to get the latest data to power their analytics with superior performance and no additional latency during updates.



Processing gigabytes of compressed data per second!

Figure 1. Works in any enterprise, on any data, for any use case

Consistently deliver fast analytics to more users

- **Execute complex queries in seconds** not hours
- **Access operational and streamed data** faster
- **Apply updates from operational systems** with no impact to query performance
- **Support 1000s of active users** to increase the return from your data investment
- **Accelerate machine learning and artificial intelligence** for deeper learning and faster insights
- **Operate continuously** with Vector’s 99.999% availability

Realize record-breaking analytic performance

You can achieve extreme performance on commodity hardware with little to no performance tuning. An advanced columnar implementation minimizes I/O while retrieving data from disk. Vectorized compute leverages CPU Single Instruction Multiple Data (SIMD) and processes data in the L1/L2 CPU cache instead of RAM, leading to significantly faster performance.

FEATURES

- Highly scalable SMP or MPP architecture
- Full ACID compliance
- Zero-penalty real-time updates
- CPU cache optimization
- Column-based storage and execution
- Query result caching
- SQL-2016 compliant
- User Defined Function support for Python, SQL and JavaScript
- Open, industry standard data access, including Spark, ODBC, JDBC, and .NET
- On-premises and BYOL cloud deployment

Deliver on-premises and hybrid cloud for maximum flexibility

Deploy as an on-premises solution on Windows and Linux and as a private or managed cloud on Google Cloud, AWS, and Azure. This enables organizations to realize the true potential of hybrid cloud by bringing compute power to the place where their data resides. You can leverage the same database engine, physical data model, ETL/ELT tools, and BI tools across clouds.

Take advantage of patented features and best practices

Column-based storage: reduce I/O to relevant columns, optimize data compression and deliver better cache performance with the flexibility to store indexed columns in the same block.

Live data updates: maintain real-time access as compressed data stored on disk is frequently updated., resulting in I/O and CPU savings and shorter execution time.

Data compression: provides multiple options to save storage space.

Parallel execution: maximize concurrency while enabling load prioritization using adaptive algorithms.

Rapidly deploy with built-in integration for BI and analytics

Integration with tools including DataFlow for fast data loading and DataConnect with over 200 connectors and templates make it to easy to source data at scale. You can load structured and semi-structured JSON data, including event-based messages and streaming data without coding.

An open API allows business analysts, data scientists and data engineers to use their favorite tools for visualization, reporting, analytics, ML and AI through ODBC, JDBC, .NET and Python. Examples include Looker, Power BI, Tableau, and MicroStrategy. The Vector database delivers:

- **Simplified, futureproof architecture:** Rely on the same patented vectorized database engine in the cloud and on-premises; work with a single data model, consistent ETL integration, and one technology to learn.
- **Stronger compliance and security:** Protect sensitive data with encryption at rest and in transit with the ability to re-key the database with new encryption keys. Dynamic data masking and column-level de-identification provide a safe way to share data across stakeholders .
- **Phased, non-disruptive migrations:** Immediately move suitable workloads to the cloud while those that should remain on-premises run on amortized infrastructure for lower TCO.