

Actian Vector Analytic Database

High performance database for the operational data warehouse

Key Benefits

- Run complex, ad hoc queries against billions of records in seconds
- Process hundreds of records in a single CPU instruction cycle with vector processing
- Execute updates without any performance penalty
- Get consistent query results even if the data changes
- Exploit dedicated CPU core and caches running 100x faster than RAM
- Scan data faster using self-indexed blocks

Flexible Deployment

- Linux
- Windows
- Hadoop
- AWS
- Azure

[Data sheet](#)

Actian Vector delivers on the promise of in-the-moment analytics with the industry's fastest analytics database. Actian Vector makes analytics more accessible to business users, freeing them from the common limitations of traditional data warehouses. Achieve extreme performance on commodity hardware or cloud-based platforms, with little or no database tuning.

Financial services, retail, telecommunications, media and many other enterprise users need to make decisions based on fresh data and the flexibility to explore their data beyond the beaten path.

Deliver fast analytics for your Operational Data Warehouse

- **Execute queries in seconds** not hours
- **Analyze larger datasets faster** whether operational or streamed
- **Apply updates from operational systems** with no impact to query performance
- **Support more concurrent users** to increase the return from your data investment
- **Iterate more quickly** - more responsive ad hoc queries without tuning

Record-breaking analytic performance

Actian Vector leverages vector processing technology to unlock hidden performance features in your existing hardware. The results below compare the sums of query execution times for all 22 TPC Benchmark™H (TPC-H) queries using popular SQL on Hadoop solutions. This workload benchmarks decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions.



Key features

Vectorized query execution: exploits Single Instruction, Multiple Data (SIMD) capabilities in commodity Intel x86 architecture CPUs, enabling processing of hundreds or thousands of data values using a single instruction.

SMP and MPP architecture provides exceptional vertical scalability on SMP systems by exploiting multiple CPUs and multiple cores. Hadoop clusters provide scale-out to thousands of users, nodes, and petabytes of data, with greater redundancy and system-wide data protection.

Full ACID compliance: performs data updates with multi-version read consistency, maintaining transaction integrity.

Zero-penalty real-time data updates: enable in the moment computing using patented Positional Delta Trees (PDTs) for incremental small inserts, updates and deletes without impacting query performance.

CPU cache optimization: uses dedicated CPU cores and caches as execution memory to run queries 100x faster than from RAM, delivering significantly greater throughput than conventional in-memory approaches.

CPU optimizations: include hardware-accelerated string-based operations, for accelerating selections on strings using wildcard matching, aggregations on string-based values, and joins or sorts using string keys.

Column-based storage: reduce I/O to relevant columns and provide the opportunity for greater data compression and enable storage indexes to maximize efficiency.

Data compression: provides multiple options to maximize compression, from 4-6x to greater than 10x for Hadoop storage.

Storage indexes: provide automatic min-max indices to enable fast block skipping on reads and eliminate the need for an explicit data partitioning strategy.

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

Spark powered direct query access: provides direct access to Hadoop data files stored in Parquet, ORC, or other standard formats allowing users to realize significant performance benefits without converting to the Vector file format first.

Extensive SQL support: uses standard ANSI SQL with advanced analytics, including cubing, grouping, and window functions.

Deploy on-premises and in the cloud: on single servers or in Hadoop clusters scaling to hundreds of nodes. Cloud deployment supports perpetual and subscription licensing, pay as you go (hourly and annual plans) or as a service, managed by Actian, on your choice of cloud service provider.



www.actian.com | Toll Free +1 888 446 4737 | Tel +1 650 587 5500
2300 Geng Rd., Suite 150, Palo Alto, CA 94303



© 2018 Actian Corporation. Actian is a trademark of Actian Corporation and its subsidiaries. All other trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. (DS52-0918_Actian Vector ODS)