

MemSQL Software

The Cloud-Native Operational Database
Built for Speed, Scale, and SQL

Overview

The value of data decreases sharply over time. To take advantage of opportunities and threats effectively, businesses have to invest in solutions which provide in-the-moment context to aid better decision-making. They also need to do this at the speed and scale the business demands while taming the complexity of a diverse data infrastructure landscape across multiple clouds and on-premise environments. Our mission at MemSQL is to deliver The Database of Now™ which provides the speed, scale, and SQL in a cloud-native solution which simplifies data infrastructure by providing a converged data platform that is optimized for real-time applications in hybrid and multi-cloud environments.

MemSQL is a distributed, highly-scalable relational SQL database that can run anywhere. MemSQL can handle both OLTP and OLAP workloads in a single system, which fits with the direction of new applications to combine transactional and analytical requirements. MemSQL can ingest millions of events per second with ACID transactions while simultaneously analyzing billions of rows of data with standard SQL. Besides all the popular structured data types like strings, date/time, numbers, binary, text, and so on, MemSQL also has native support for JSON, Full-text search, and geospatial types. MemSQL JSON support simplifies the developer's job when processing semistructured data, while maintaining the benefits of SQL and the relational model.

MemSQL is designed to deliver maximum performance for transactional and analytical workloads with familiar relational models. MemSQL offers both multi-cloud and hybrid options, ranging from a database-as-a-service, to Kubernetes-based hybrid and private deployments, to traditional installations on VMs or commodity hardware, in the public cloud or on premises. MemSQL SingleStore is a breakthrough in database architecture, allowing transactional and analytical workloads to be processed using a single table type. With MemSQL SingleStore, and the new system of record improvements, our vision of “one database to rule them all” begins to be realized.

MemSQL is ideal for real-time applications that require fast data ingestion, low-latency queries and elastic scaling with familiar relational SQL. MemSQL can deliver 10x performance at one third the cost of traditional databases. It can handle converged workloads with rowstores containing tens of terabytes executing over 10,000 complex queries per second, and columnstore data sets in the multi-petabyte range, while supporting aggregations of over a trillion rows per second. Additional innovations which enable MemSQL's performance characteristics are lock-free skip lists, SIMD vectorization, highly-efficient data compression, and high-performance distributed joins. MemSQL lets you achieve ultra-fast query response across both live and historical data using familiar ANSI SQL.

Benefits

MemSQL enables ultra-fast operational analytics and real-time insights to achieve faster, more informed decisions, improved customer experiences, and operations.

Latency-Free Analytics: MemSQL lets you achieve ultra fast query response with high concurrency across both live and historical data using familiar ANSI SQL


Ultra-fast Event-to-Insight Performance: Deliver against the toughest service level agreements using parallel, distributed lock-free ingestion and real-time query processing

Scale Limitlessly: Elastic scale-out architecture with distributed massively parallel data processing delivers consistent, predictable response under high ingest and user concurrency

Ease of Use and Flexibility: MemSQL SingleStore™ brings simplicity and ease to your data architecture by allowing OLTP and OLAP workloads to be processed on operational data using a single table type

Drop-in Compatibility: Plug-in directly with existing technologies and skills with support for standard SQL, BI and distributed technologies like Amazon S3, Spark, Kafka and Hadoop

Powerful Programmability: Supports simple, fast, and powerful in-database programming via MPSQL and an extensive set of built-in data types and functions



Technical Specifications

Universal requirements: Each MemSQL node requires a host machine with an x86_64 CPU with at least four CPU cores and 8GB of RAM available per node. When provisioning your host machines, the minimum Linux kernel version required is 3.10 or later.

Recommended platforms: RHEL/CentOS 6 or 7 (version 7 is preferred), Debian 8 or 9 (version 9 is preferred)

Hardware recommendation for optimal performance: CPU: 8 vCPU per host machine, Memory: At least 4GB per core and 32GB minimum per leaf node, Storage: Provide a storage capacity for each node with at least three times the capacity of main memory, and SSD storage is recommended for columnstore workloads.

Network Hardware Recommendation: 10Gigabit ethernet switch or better recommended for cluster node interconnection.

Features

MemSQL Pipelines: Built-in parallel data ingestion technology natively ingests high-throughput real-time data from external sources such as Apache Kafka, Amazon S3, Azure Blob, Filesystem, Google Cloud Storage and HDFS data source.

MemSQL SingleStore: SingleStore architecture allows you to support large-scale Online Transaction Processing (OLTP) and Hybrid Transactional and Analytical Processing (HTAP) at a lower total cost of ownership (TCO). It is a continuing evolution of the columnstore, supporting transactional workloads that would have traditionally used the rowstore.

Compiled, Vectorized Query Execution: Built-in distributed query optimizer evenly divides the processing workload to maximize the efficiency of CPU usage. Query plans are compiled to machine code and cached to expedite subsequent executions

Multi-Version Concurrency Control (MVCC) and Lock-Free Data Structures: With these technologies, data remains highly accessible, even amidst a high volume of concurrent reads and writes.

Distributed Ingest, Bulk or Streaming, with Concurrent Non-Blocking Reads: MemSQL offers a lock-free architecture which is based on the skip list index, that efficiently processes transactions and updates without locking or blocking concurrent reads, resulting in delivering the capability to perform bulk and/or streaming ingestion online, simultaneously with query workload.

Highly Compatible: MemSQL is wire protocol compatible with MySQL making it instantly accessible from any BI tool such as Tableau, PowerBI or Looker and with widely-available bindings to popular programming languages such as Golang, Rust, Python, NodeJS, R, Java, and C++.

Flexible Storage: Memory-optimized rowstore and an on-disk columnstore to handle both highly concurrent transactional and analytical workloads

MemSQL Studio: Visual user interface tool that allows you to easily monitor, debug and interact with all of your MemSQL Clusters.

Enterprise Security: Ensures military-grade security with RBAC, encryption, auditing, password policy management, and strict mode to isolate data from administrators.

Extensible: Supports in-database programming via the MPSQL language, which allows definition of user-defined functions, stored procedures (SPs), table-valued functions, and user-defined aggregates. Language features include embedded SQL, dynamic SQL, arrays, records, a full suite of control structures, exception handling, and recursion. SPs and functions are compiled to machine code for speed. An extensive set of built-in functions, JSON functions, and AI functions (cosine similarity and Euclidean distance), can be used in SQL statements, SPs, and functions.

Use Cases

Operational Analytics: MemSQL delivers the fastest and most scalable reporting and analytics across all of your operational data; including streaming, real-time, and historical data. Typical scenarios include real-time applications for retail inventory analytics, A/B testing, game play analysis, threat detection, and streaming media quality analytics.

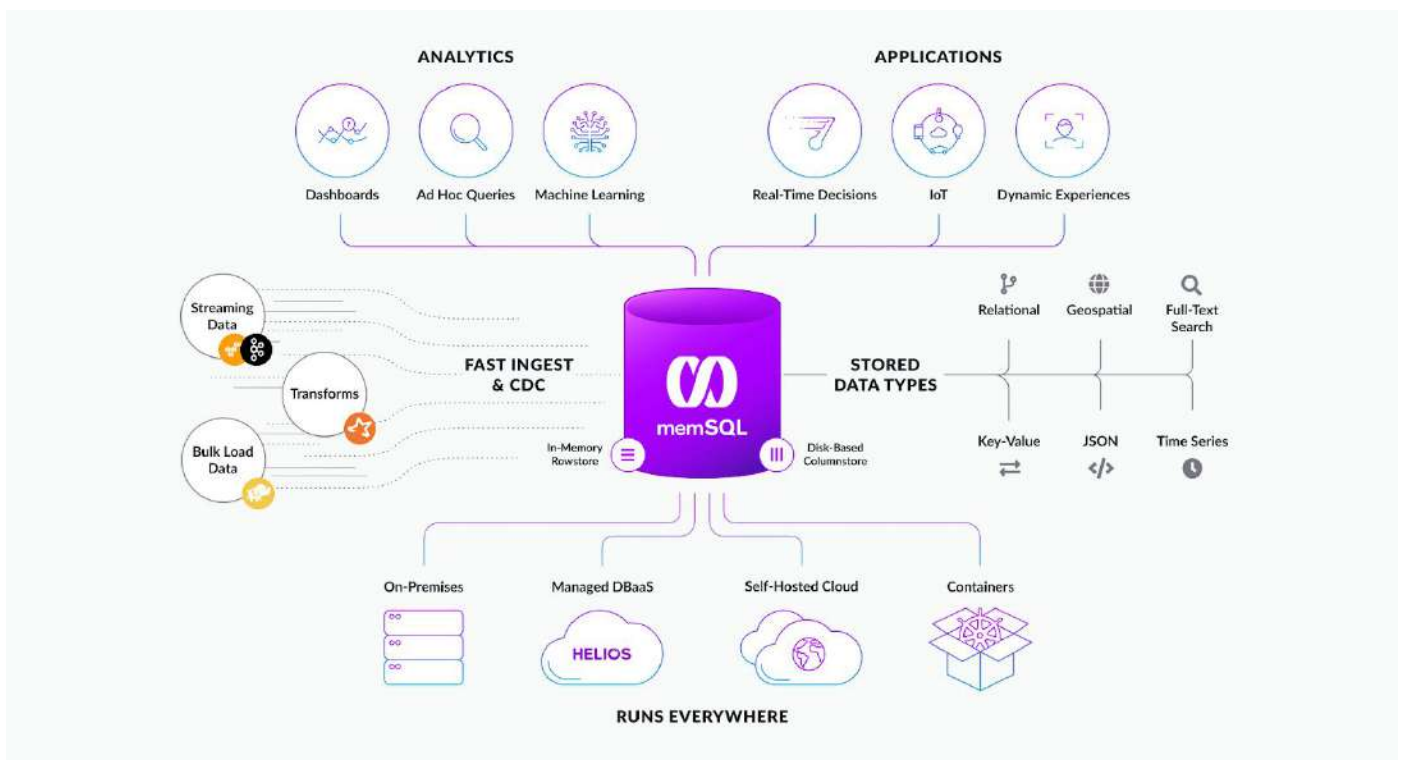
Operationalizing ML and AI Apps: MemSQL can transform your business with scalable ML/AI applications with an ultra fast ingest and query platform that enables real-time model scoring on both streaming and historical data, and vector similarity testing for image matching and other deep learning applications. Typical uses are real-time applications for fraud detection.

Monitor and Detect in Real-Time: MemSQL gives organizations the ability to monitor and detect anomalies in rapidly changing data through its innovative real-time ingestion, database and analytics platform. Typical scenarios include real-time applications for predictive analytics for energy demand-response, and location analytics.

Accelerating Legacy Data Platforms: MemSQL can accelerate legacy systems with scalable rapid data ingestion and fast queries on large data sets, by leveraging the simplicity of SQL. Typical scenarios include modernizing data lakes and data warehouses built on Hadoop, object stores & legacy analytic databases by bringing low-latency queries to the data layer.



MemSQL Architecture



Get Started with MemSQL for Free Today at memsql.com/software/