MemSQL Helios

A Fully-Managed, On-Demand, & Elastic Cloud Database

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 SQL database that can run anywhere. MemSQL's database-as-a-service offering is called *MemSQL Helios*. MemSQL Helios gives you the full capabilities of MemSQL without the operational overhead and complexity of managing it yourself. MemSQL Helios provides a resilient database with cloud-agnostic deployment support on AWS and Google Cloud Platform (with support for Azure and others to come). MemSQL Helios can handle both OLTP and OLAP workloads in a single system as cloud-native HTAP, which fits with the direction of new applications to combine transactional and analytical requirements.

MemSQL Helios 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 relational data types, MemSQL Helios also supports semi-structured data with native support forJSON and full-text search with efficient indexing for each. With the support for geospatial, time-series, and key-value use cases, MemSQL Helios provides an easy-to-use, efficient converged database which reduces complexity.

MemSQL database lets you achieve ultra fast query response across both live and historical data using familiar ANSI SQL. With MemSQL Helios, cluster provisioning, cluster management, deployment, upgrades, alerting, and troubleshooting are all handled by MemSQL. This greatly reduces operational expenses, by shifting the database administration (DBA) tasks needed to operate your database from your organization to MemSQL. Just as importantly, MemSQL Helios is offered at a price point dramatically lower than traditional database vendors, while our ultra-efficient query engine means that operational costs for MemSQL also tend to be lower than the proprietary offerings from the cloud service providers.

MemSQL Helios 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

Benefits

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

Effortless Deployment & Elastic Scale: Get the full capabilities of MemSQL with one-click deployment and easy cloud scalability

Superior TCO vs. Legacy & CSPs: Dramatically cheaper to buy & operate than either legacy data platforms or public cloud databases

Mult-Cloud & Hybrid Flexibility: Avoid cloud lock-in with a database available across public cloud providers and on-prem environments

Other Key Benefits Include:

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

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 Helios lets you achieve ultra-fast query response across both live and historical data using familiar ANSI SQL.

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

Technical Specifications

Minimum Requirements: 2 unit cluster with a total of 16vCPUs, 128 GB Memory, 2TB Storage.
A cluster's size can be increased on-demand.

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

Highly Compatible: MemSQL is wire protocol compatible with MySQL and supports advanced features not in MySQL such as Distributed SQL, Geospatial, JSON, Window Functions, and Time Series Functions.

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.

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.

Enterprise Security: Ensures military-grade security with RBAC and encryption 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.

Uber COMCAST Kellogg's DELL EMC Akamai SAMSUNG cisco pandora monday

---

MemSQL Helios Architecture

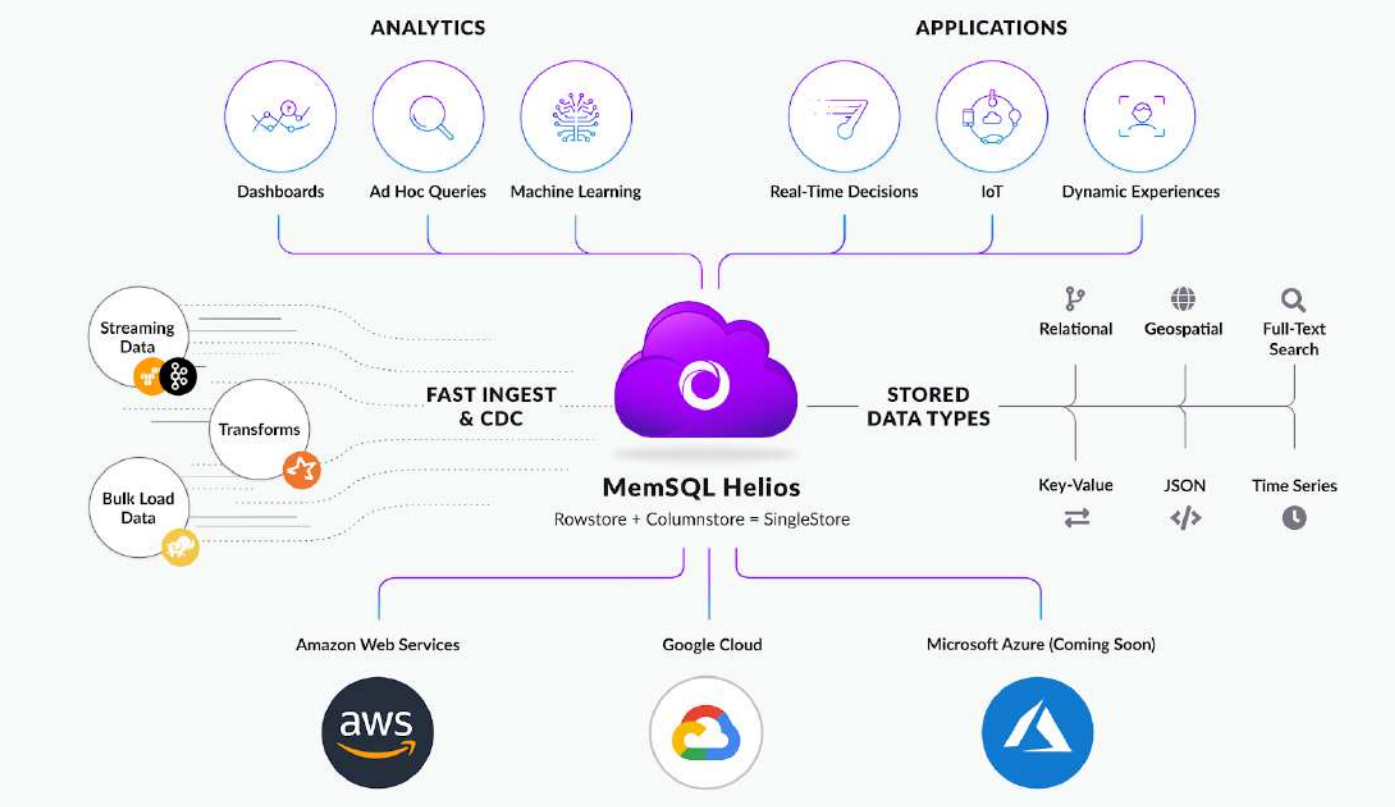

Get Started with MemSQL Helios for Free Today at **memsql.com/helios**