



Cisco UCS and MemSQL: Real-Time Data Warehouse for the Enterprise

Adapt and Learn from Your Data in Real Time



Highlights

Real-Time Performance with Linear Scalability

Cisco Unified Computing System (Cisco UCS) Integrated Infrastructure for Big Data and Analytics, combined with the MemSQL real-time data warehouse, provides a simplified intelligent data infrastructure with high performance and scalability to meet growing real-time business demands.

Integrated Infrastructure Built on Cisco UCS Advantages

Cisco UCS Integrated Infrastructure for Big Data and Analytics is a proven platform for enterprise analytics applications. MemSQL delivers high-speed simultaneous data ingestion and SQL analysis with high levels of concurrency for large-scale application requirements.

Ease of Deployment

Cisco UCS Manager simplifies infrastructure deployment with an automated, policy-based mechanism that helps reduce configuration errors and system downtime.

Optimized Real-Time Analytics

MemSQL is a real-time data warehouse that spans in-memory and disk-based technologies to deliver extremely fast data ingestion and analytics.

Built-In Resilience and Security

Efficiently scale out MemSQL on Cisco UCS hardware to provide highly available and secure analytics that is easy to set up, maintain, and scale. Protect your data against threats with enterprise-grade security encompassing authentication, encryption, and auditing.

Extreme Performance for Real-time Analytics

Real-time analytics on live data, and the ability to analyze petabytes of data to gain insights, are key for any on-demand and digital organization. Acting on insights as they happen can determine the next best customer interaction, identify a new opportunity, or prevent costly expenditures. The explosive growth in digital interactions across customers and devices has put new pressures on traditional data collection and analysis technologies.

The need to be able to ingest millions of events per second has become greater than ever. With always-on devices and sensors proliferating, their data must be ingested and actionable in the moment. The infrastructure behind each Internet of Things (IoT) application will be powered by a real-time database.

There's no question that the growth of digital business affects data architectures. This growth is driven by the expansion of applications, devices, and data. A digital business runs on data to enhance real-time decisions, improve customer experiences, and optimize business operations. To become a responsive data-driven business, organizations must address their current data latency challenges. These challenges are commonly found across three general areas:

- Slow data loading: Loading data into a format that is easy to analyze, moving past batch processes, and receiving data in real time. Most data warehouses load data in batches, which results in stale or latent insights into the business.
- Lengthy query execution: Operational insights must be readily available for in-the-moment decisions. A data warehouse that delivers fast query response can deliver insights whenever the application or users require, ultimately providing a differentiated service or identifying opportunities.
- Low concurrency: Digital business assumes a large-scale use of data across an entire business or customer base. As more users engage and interact with data, the response time for those interactions must maintain a consistent experience. A scalable data warehouse will help ensure that data and user growth will not negatively affect the operational system.

MemSQL on Cisco UCS Integrated Infrastructure for Big Data and Analytics as shown in Fig 1 provides a scalable, real-time data warehouse platform for high-performance applications that require fast, accurate, secure, and always available data. MemSQL can linearly scale millions of events per second while analyzing petabytes of data for insights.

MemSQL supports fast ingestion and concurrent analytics needed for sensor systems, recommendation systems, and similar use cases that deliver instant, actionable insights.

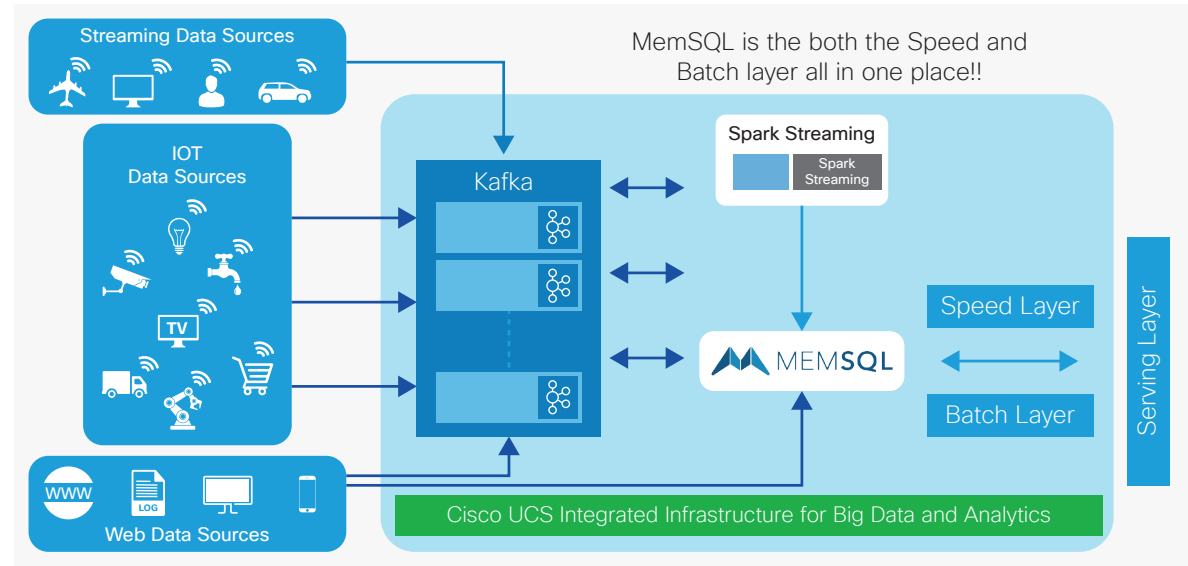
Cisco UCS Integrated Infrastructure for Big Data and Analytics

The Cisco UCS Integrated Infrastructure for Big Data and Analytics provides an end-to-end architecture for processing high volumes of real-time or archived data, both structured and unstructured. At the same time, it quickly and efficiently delivers out-of-the-box performance while scaling from small to very large as business requirements and big data and analytics requirements grow.

Cisco UCS Manager

Cisco UCS Manager provides unified, embedded management of all software and hardware components in the Cisco UCS system. It enables rapid, consistent server configuration using service profiles and automates ongoing system maintenance activities such as firmware updates across the entire cluster as a single operation. Cisco UCS Manager enables fast, accurate configuration of computing, networking, and storage resources.

Figure 1. Solution Architecture of MemSQL with Cisco UCS



It also offers advanced monitoring with options to raise alarms and send notifications about the health of the entire cluster.

Cisco UCS 6200 and 6300 Series Fabric Interconnects

Cisco UCS 6200 Series Fabric Interconnects provide high-bandwidth, low-latency connectivity for servers, with integrated, unified management provided by Cisco UCS Manager for all connected devices. The Cisco UCS 6300 Series Fabric Interconnects are the latest version of this technology. The 6300 Series is a core part of Cisco UCS, providing low-latency, lossless, 10 and 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions with management capabilities for systems deployed in redundant pairs.

Cisco fabric interconnects offer the full active-active redundancy, performance, and exceptional scalability needed to support the large number of nodes that is typical in clusters serving big data applications.

Cisco UCS C-Series Rack Servers

The Cisco UCS C240 M4 Rack Server is an enterprise-class server designed to deliver exceptional performance, expandability, and efficiency for storage and I/O-intensive infrastructure workloads.

The server uses dual Intel® Xeon® processor E5-2600 v4 Series CPUs and supports up to 1.5 TB of main memory and a range of hard-disk drive (HDD) and solid-state disk (SSD) options.

The performance-optimized option supports 24 small-form-factor (SFF) disk drives, and the capacity-optimized option supports 12 large-form-factor (LFF) disk drives. This server can be used with the Cisco UCS Virtual Interface Card (VIC) 1227 for 10-Gbps network connectivity or with the VIC 1387 for 40-Gbps network connectivity, depending on the fabric interconnect that is being used.

MemSQL: A Real-Time Data Warehouse

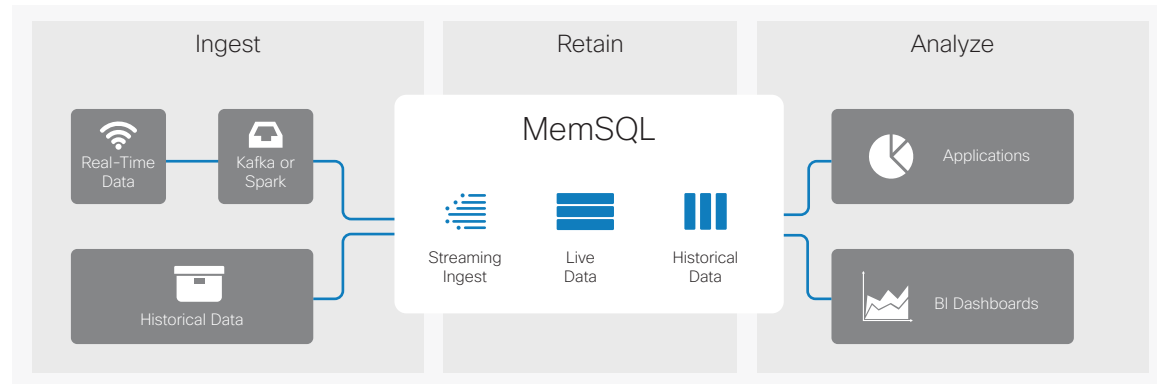
MemSQL is a scalable, real-time data warehouse that ingests data continuously to perform analytics for the front lines of the business. It can ingest and transform millions of data events per day while simultaneously analyzing billions of rows of data using standard SQL. The database can query traditional SQL, JSON, and geospatial data types in real time.

MemSQL enables businesses to process transactions and perform analysis simultaneously in a single operational database with a lock-free data structure. With data loaded into scalable memory, analytics can be executed across extremely large data sets in real time. Having immediate access to both live and historical data enables MemSQL to deliver new opportunities for customer engagement, personalization, and IoT applications, and new analytical models.

Another area where MemSQL excels as shown in Fig 2 is in capturing real-time data at the point of ingestion and fusing this data with other operational data to deliver new applications or customer experiences or to allow for operational analytics for real-time

dashboards and reports. Further, MemSQL accelerates application performance without costly rewrites using ANSI SQL and the MySQL wire protocol.

Figure 2. Scalable Real-Time Data Warehouse with MemSQL



Key capabilities of MemSQL include:

- **Fast data ingestion:** Collect data using common message brokers such as Apache Kafka while maintaining durable, consistent delivery with exactly-once semantics.
- **Ultra-fast analytics:** Query petabytes of data using advanced data compression using disk-optimized tables with up to 10x compression and vectorized queries for fast analytics.
- **Real-time analytics:** Use memory-optimized tables to analyze real-time events.
- **Drop-in compatibility:** Integrate with existing infrastructure using Java Database Connectivity (JDBC) or Open Database Connectivity (ODBC) and ANSI SQL and MySQL compatibility.

- **Geospatial support:** Store, query, and index geographic data types, including polygons and points, to support area, distance, and location analytics in real time.
- **JSON optimized:** Store and query JSON data as a column type to efficiently store and analyze multiattribute objects
- **Fully distributed joins:** Scale out fully distributed joins across any table and column for simple, efficient query access.

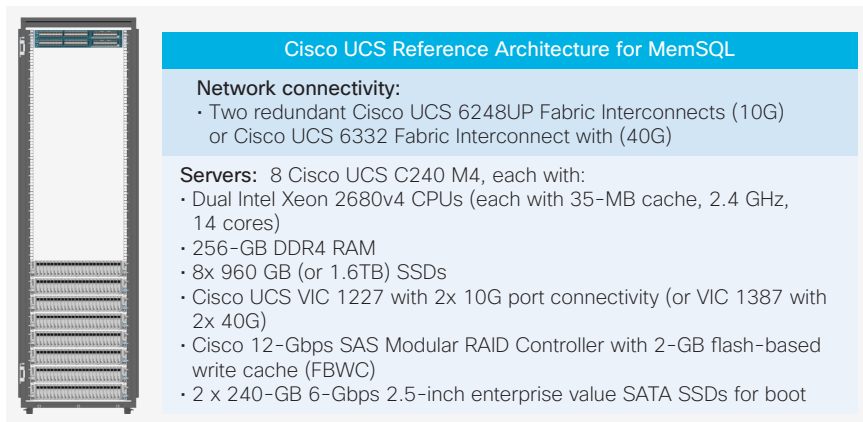
Use cases include:

- Real-time analytics
- Internet of Things
- Personalization and recommendations
- Risk management
- Monitoring and detection
- Customer 360

Reference Architecture

The Cisco UCS Integrated Infrastructure for Big Data and Analytics for MemSQL includes eight or more C240 M4 servers, each with dual Intel Xeon 2680v4 processors, 256-GB RAM, dual 10-Gbps (or 40-Gbps) network connectivity, and eight SSDs. These servers are connected to Cisco UCS 6248 fabric interconnects. Figure 3 shows the reference architecture for MemSQL.

Figure 3. Cisco UCS Reference Architecture for MemSQL



Performance Tests and Results

Several performance benchmarks were used to highlight the strength of MemSQL and Cisco UCS for real-time data warehouses. These tests focus on the most critical components for real-time analytics: ingest, concurrency, and scaling. Both the ingest and query performance were increased to show how MemSQL and Cisco can be dynamically scaled to meet the needs of today's businesses.


Ingest Performance

To achieve real-time analytics, the system must be able to populate the database as quickly as possible with a variety of sources. MemSQL is unique in that it can efficiently ingest data from Kafka streams as well as traditional APIs. The ingest performance is shown below with a simple Upsert.

Upsert performance is commonly used for IoT environments where values are sampled and events are counted over specific time frames.

MemSQL and Cisco UCS were able to ingest over 4 million Upserts per second and achieved near-linear speedup (1.97x) when the cluster size doubled from four to eight nodes (Figure 4).

Figure 4. MemSQL on Cisco UCS – Ingest Performance



The test above was run with a direct API connection to MemSQL.

MemSQL pipelines support a direct connection to Kafka streams. Pipelines manage the Kafka topics and offsets to ensure exactly-once semantics for ingesting real-time data. Essentially, pipelines read from a Kafka stream in micro-batches that are managed entirely by MemSQL.

Concurrent Query Performance with Ingest

This use case is very common in the ad technology industry. Consider a system that has numerous advertisers that populate real-time dashboards to show how effectively current campaigns are performing and to submit ad hoc queries. This must scale to meet the needs of the multiple advertisers and demands of the business.

To show concurrent query performance, we used the [dbbench toolkit](#) created by MemSQL to run and scale concurrency. This toolkit reports the overall “per second” rate of dashboard refreshes and ad hoc report queries. Finally, during the query measurements there was a steady-state ingest of approximately 285,000 rows per second from two Kafka pipelines.

Concurrent Dashboards with Linear Scaling

The dashboard use case plans to support up to 900 dashboard refreshes per second as the application is rolled out over 3 years. The goal is to be able to support this refresh rate while staying well under a 1-second response time. MemSQL showed excellent scaling in order to handle the desired dashboard refresh rate by simply scaling up the cluster from four to eight nodes (Figure 5).

Ad Hoc Campaign Reports

These reports are requested by advertisers on an ad hoc basis. It is critical to be able to increase the throughput rate while maintaining a consistent response time that meets the SLAs. The desired goal is to be able to support more than 16 ad hoc reports per second by year 3 while maintaining a response time of less than 1 second (Figure 6).

Figure 5. Dashboard Report with Linear Scaling


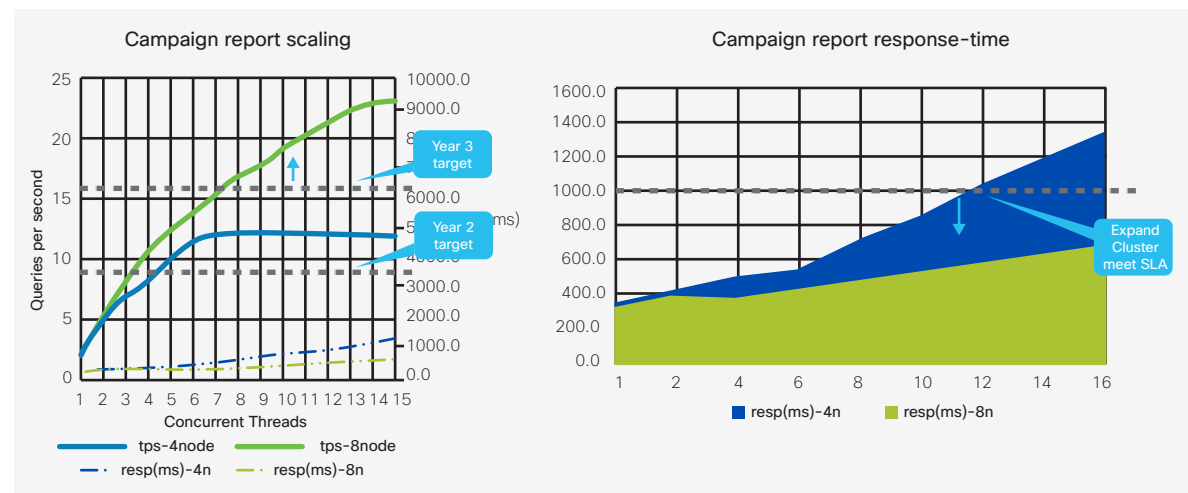


Figure 6. Campaign Report Scaling (left) and Campaign Report Response Time (right)



MemSQL allows for easy expansion of the cluster simply by adding nodes. This increases the concurrency while maintaining the desired SLAs.

Conclusion

The Cisco UCS Integrated Infrastructure for Big Data and Analytics with MemSQL provides a simplified, intelligent infrastructure and a real-time data warehouse with the scalability to meet growing business demands.

For More Information

- To find out more about Cisco UCS big data solutions, see <http://www.cisco.com/go/bigdata>.
- To find out more about MemSQL, see <http://www.memsql.com/>.



Copyright © 2017 MemSQL, Inc. All Rights Reserved.

MEMSQL and the MEMSQL Logo are trademarks, registered trademarks or service marks (Marks) of MemSQL, Inc. All other trademarks and service marks, which may be registered in certain jurisdictions, may be claimed by the holder or holders of such marks.

MemSQL products are protected by patents (issued and pending) in the U.S. and elsewhere as set forth at www.memsql.com/patents.

© 2017 Cisco and/or its affiliates. All rights reserved. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

C22-739001-00 05/17