



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 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 Unified Computing System™ (Cisco UCS) hardware to provide highly available and secure analytics and a solution that is easy to set up, maintain, and scale. Protect your data against threats with enterprise-class security that encompasses authentication, encryption, and auditing.

Extreme performance for Real-time analytics

Real-time analytics on live data and the capability to analyze petabytes of data to gain insights are essential for on-demand, digital operations. The capability to act immediately on insights can help you determine the best next customer interaction, identify a new opportunity, and avoid costly expenditures.

The dramatic growth in digital interactions across customers and devices has created challenges for traditional data collection and analysis technologies. Organizations now must be able to ingest millions of events per second. Always-on devices and sensors are proliferating, and their data must be ingested and actionable immediately. The infrastructure that supports each Internet of Things (IoT) application must be powered by a real-time database.

The growth of digital business affects data architectures. This growth is fueled by the expanding number of applications and devices and increasing amount of data that must be processed. 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 data latency challenges. These challenges commonly include the following:

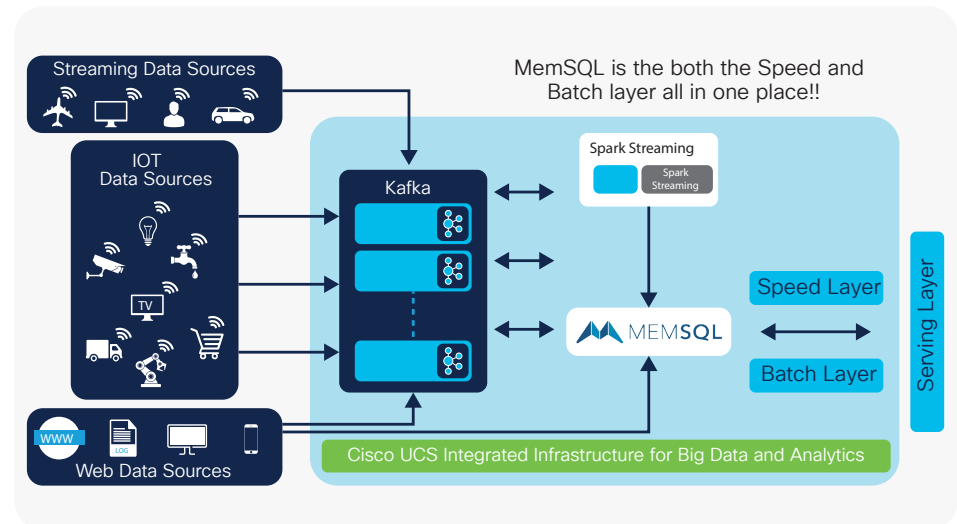
- Slow data loading: Data must be loaded in a format that is easy to analyze, moving quickly through batch processes and receiving data in real time. Most data warehouses load data in batches, which results in stale data and potentially obsolete insights into the business.
- Lengthy query execution: Operational insights must be readily available for in-the-moment decisions. A data warehouse that delivers fast query responses can deliver insights whenever the application or users require them, ultimately providing a differentiated service and identifying opportunities.

- Low concurrency: Digital business assumes a large-scale use of data across the entire business or customer base. As more users engage and interact with data, the response time for those interactions must not degrade, maintaining a consistent user experience. A scalable data warehouse helps ensure that data and user growth does not negatively affect operations.

Cisco UCS® Integrated Infrastructure for Big Data and Analytics with MemSQL provides a scalable, realtime data warehouse platform for high-performance applications that require fast, accurate, secure, and always-available data (Figure 1). MemSQL can linearly scale millions of events per second while analyzing petabytes of data for insights.

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

Figure 1. Cisco UCS Integrated Infrastructure for Big Data and Analytics with MemSQL Solution Architecture



Cisco UCS Integrated Infrastructure for Big Data and Analytics

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 deployments as business requirements and big data and analytics requirements grow.

Organizations today must be sure that the underlying physical infrastructure can be deployed, scaled, and managed in a way that is agile enough to change as workloads and business requirements change. Cisco UCS Integrated Infrastructure for Big Data and Analytics has redefined the potential of the data center with its revolutionary approach to managing computing, network, and storage resources to successfully address the business needs of IT innovation and acceleration.

Cisco UCS 6300 Series fabric interconnects

Cisco UCS 6300 Series Fabric Interconnects provide high-bandwidth, low-latency connectivity for servers, with Cisco UCS Manager providing integrated, unified management for all connected devices. Cisco UCS 6300 Series Fabric Interconnects are a core part of the Cisco Unified Computing System™ (Cisco UCS), providing low-latency, lossless 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions.

Cisco® fabric interconnects offer the full active-active redundancy, performance, and exceptional scalability needed to support the large number of nodes that are typical in clusters serving big data applications. Cisco UCS Manager enables rapid and 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 also offers advanced monitoring with options to raise alarms and send notifications about the health of the entire cluster.

Cisco UCS C240 and C220 M5 Rack Servers

Cisco UCS M5 rack servers are dual-socket, 2-Rack-Unit (2RU) servers offering industry-leading performance and expandability for a wide range of storage and I/O-intensive infrastructure workloads, such as big data, analytics, and collaboration. These servers use the new 2nd Generation Intel Xeon® Processor Scalable Family with up to 28 cores per socket. They support up to 24 Double-Data-Rate 4 (DDR4) Dual in-Line Memory Modules (DIMMs) for improved performance and lower power consumption. The DIMM slots are 3D XPoint ready, supporting nextgeneration nonvolatile memory technology.

Depending on the server type, Cisco UCS rack servers have a range of storage options. The Cisco UCS C240 M5 Rack Server supports up to 26 Small Form-Factor (SFF) 2.5-inch drives (with support for up to 10 Non-Volatile Memory Express [NVMe] PCIe Solid-State Disks [SSDs] on the NVMe-optimized chassis version) or 12 Large-Form-Factor (LFF) 3.5-inch drives plus 2 rear hot-swappable SFF drives with a Cisco 12-Gbps SAS Modular RAID Controller. The Cisco UCS C220 M5 Rack Server supports up to 10 SFF 2.5-inch drives (with support for up to 10 NVMe PCIe SSDs on the NVMe-optimized chassis version) In addition, all the servers have two modular M.2 cards that you can use for boot. A modular LAN-Onmotherboard (mLOM) slot supports dual 40 Gigabit Ethernet network connectivity with the Cisco UCS Virtual Interface Card (VIC) 1497.

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, JavaScript Object Notation (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 performed 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, IoT applications, and new analytical models.

MemSQL also excels at capturing real-time data at the point of ingestion and fusing this data with other operational data to deliver new applications and customer experiences and allow operational analytics for real-time dashboards and reports (Figure 2). In addition, using ANSI SQL and the MySQL wire protocol, MemSQL accelerates application performance without costly rewrites.

The main capabilities of MemSQL include:

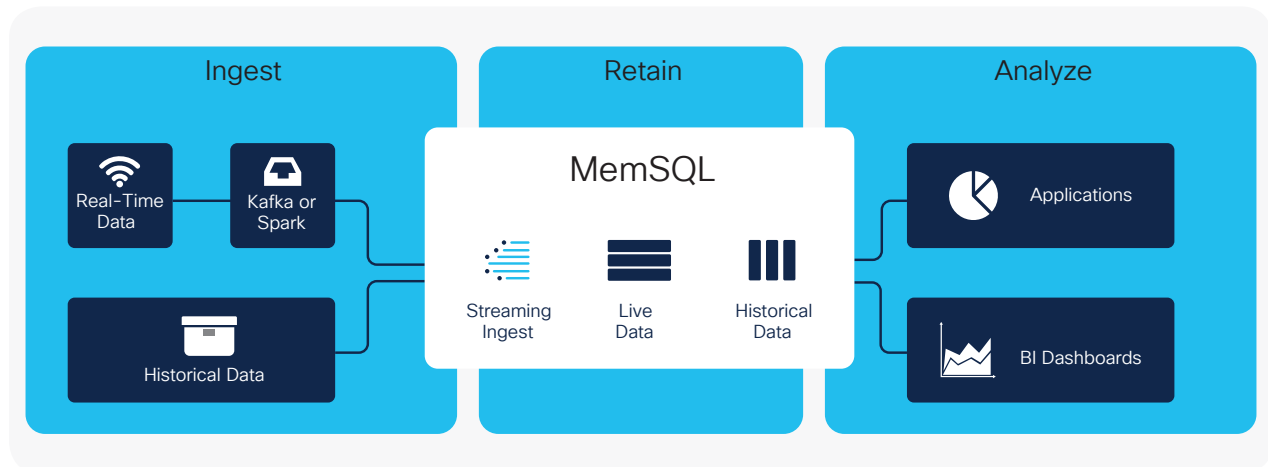
- **Fast data ingestion:** Collect data using common message brokers such as Apache Kafka while maintaining durable, consistent delivery with exactly-once semantics.
- **Ultra-fast analytics:** Query petabytes of data using advanced data compression using disk-optimized tables with up to 10x compression 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 optimization:** Store and query JSON data as a column type to efficiently store and analyze objects with multiple attributes.
- **Fully distributed joins:** Scale out fully distributed join operations across any table and column for simple, efficient query access.

Use cases include:

- Real-time analytics
- IoT
- Personalization and recommendations
- Risk management
- Monitoring and detection
- 360° customer view

Figure 2. Scalable real-time data warehouse with MemSQL Business intelligence Dashboards



Reference architecture

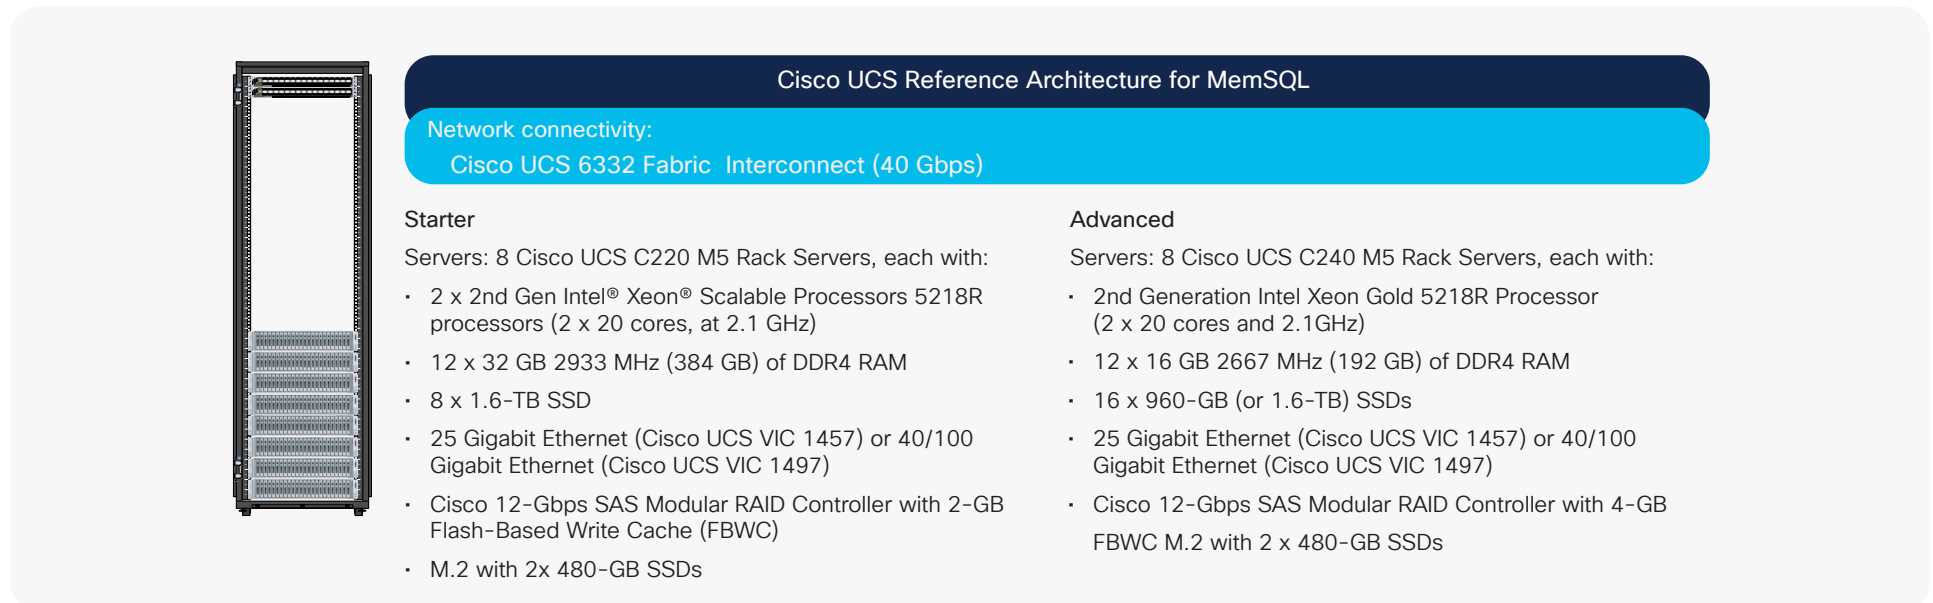
Cisco UCS Integrated Infrastructure for Big Data and Analytics for MemSQL includes eight or more Cisco UCS C240 M5 (or C220 M5) servers, each with dual Intel Xeon Processor Scalable Family 6132 CPUs (2 x 14 cores and 2.6 GHz), 384 GB of RAM, dual 40-Gbps network connectivity, and 8 (or 16) SSDs. These servers are connected to Cisco UCS 6332 Fabric Interconnects.

Figure 3 shows the 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: ingestion, concurrency, and scalability. Both the ingest and query performance were increased to show how MemSQL and Cisco can be scaled dynamically to meet the needs of today's businesses.

Figure 3. Cisco UCS Reference Architecture for MemSQL



Ingest performance

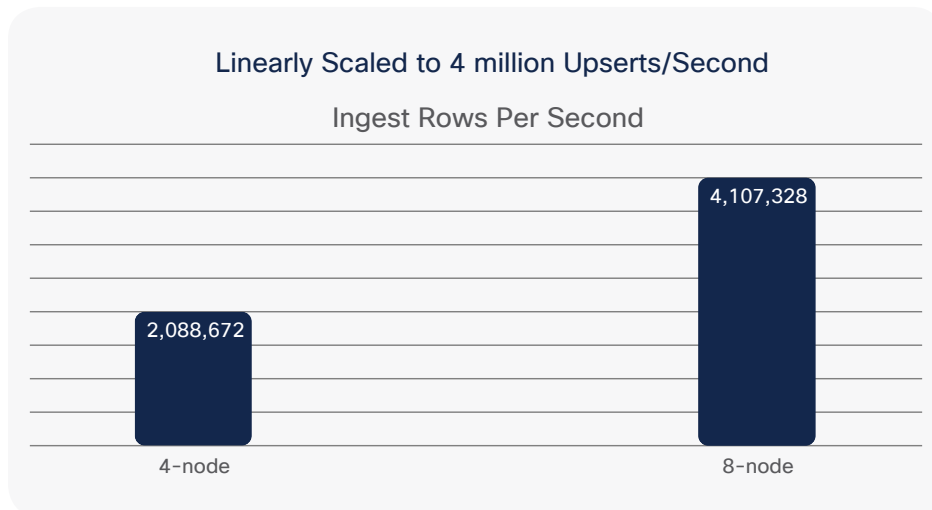
To achieve real-time analytics, the system must be able to populate the database as quickly as possible with a variety of sources. MemSQL is unique in that it can efficiently ingest data from Kafka streams as well as traditional APIs. The ingest performance is demonstrated here with a simple upsert operation. Upsert performance is commonly used for IoT environments in which values are sampled and events are counted over specific time frames.

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

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

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

Figure 4. MemSQL on Cisco UCS: Ingest Performance



Concurrent query performance with ingest

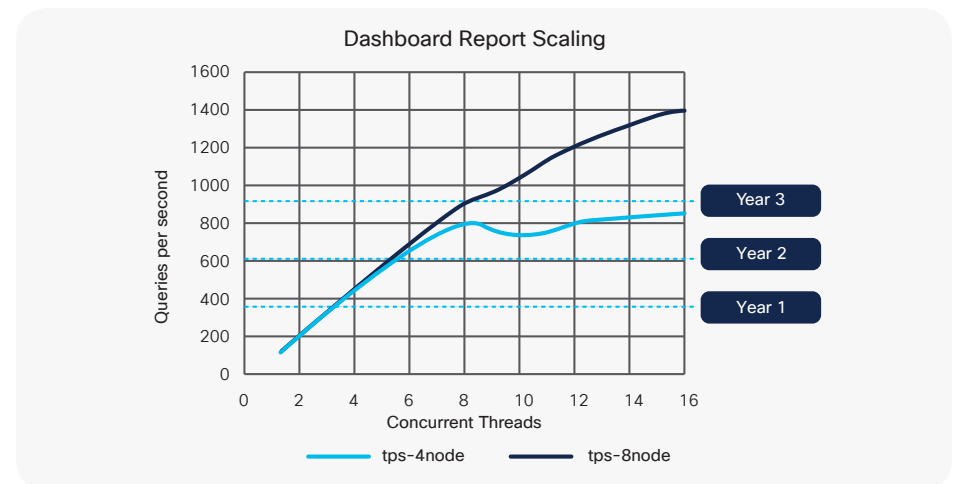
The concurrent query use case is common in the advertising technology industry. For instance, consider a system that has numerous advertisers that populate real-time dashboards to see how effectively current campaigns are performing and that accept impromptu queries. This system must scale to meet the needs of the multiple advertisers and business demands.

To demonstrate 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 impromptu report queries. During the query measurement, there was a steady-state ingestion of approximately 285,000 rows per second from two Kafka pipelines.

Concurrent dashboards with linear scaling

The goal for the system is to support up to 900 dashboard refreshes per second as the application is rolled out over three years. The goal is to be able to support this refresh rate while staying well under a 1-second response time. MemSQL showed excellent scalability, handling the desired dashboard refresh rate by simply scaling up the cluster from four nodes to eight nodes (Figure 5).

Figure 5. Dashboard Report with Linear Scaling



Conclusion

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 <https://www.cisco.com/go/bigdata>.
- To find out more about MemSQL, see <http://www.memsql.com/>.



© 2020 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)

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

Impromptu campaign reports

Campaign reports are requested by advertisers on an impromptu basis. Systems must be able to handle an increase in the throughput rate while maintaining a consistent response time that meets Service-Level Agreements (SLAs). The desired goal is to be able to support more than 16 impromptu reports per second by year three while maintaining a response time of less than 1 second (Figure 6).

MemSQL allows easy expansion of the cluster through the addition of nodes. This simple approach increases concurrency while maintaining the desired SLAs.

Figure 6. Campaign Report Scaling (Left) and Campaign Report Response Time (Right)

