

# Building Scalable Real-Time Analytics



## Table of Contents

Introduction .....	3
The Need for Real Time .....	3
The Limitations of Legacy Data Architectures .....	4
Legacy Database Architectures Fall Short – Here’s Why .....	4
SingleStoreDB: The Powerful, Scalable Modern Databases .....	5
Scalable Real Time in Real Life .....	7
Conclusion .....	8

## Introduction

From hailing a cab with Uber, to streaming music on Spotify or checking your portfolio, in today's fast-paced, data-intensive world, real-time applications are everywhere and continuing to gain more traction.

You may want to provide targeted offers and recommendations to your customers in real time. Or maybe you need to improve your cybersecurity posture with immediate data ingestion and analytics on security event streams. Perhaps you are trying to power real-time trading platforms and crypto-exchanges with time-series data and ultra-fast analytics. Or, you're looking to provide real-time portfolio analytics and wealth management to your top-tier clients.

**Can your data architecture handle these challenges?**

## The Need for Real Time

Real-time decision making is crucial for companies of all sizes. Those that continue to rely on stale data and analytics, served up on slow apps and dashboards with limited interactivity and responsiveness, have a built-in disadvantage. Because when business leaders lack the latest data, poor decisions are made and market opportunities are missed.

And the bad news doesn't end there. Slow, unresponsive or stale data in your applications also contributes to poor customer and end-user experiences – which can negatively impact sales and lead to customer and partner churn.

Fast data is the next major competitive differentiator in business. Your organization's ability to access data about every aspect of its business (e.g., customers, transactions, supply chains, asset performance, etc.) in real time and transform it into immediate, meaningful insights to act upon is an immensely powerful competitive advantage. Organizations able to implement the needed data infrastructure and operationalize it behaviorally will be the winners of tomorrow.

That's why the need to move to real time is growing more urgent every day.

**But, are you ready for real time?**

## The Limitations of Legacy Data Architectures

If you're still relying on legacy data engines and architectures, you're probably not ready for the real-time revolution. At least, not easily.

Why? Because legacy platforms and databases can't handle high-intensity, fast-moving and fast-changing data streams. They weren't designed for parallel streaming ingestion, super-low latency queries and analytics on streaming data as it lands – and simply can't keep up.

## Legacy Database Architectures Fall Short – Here's Why

The ability to access real-time information and analytics doesn't happen through sheer force of will. Real-time access comes from a deliberate foundation – a data architecture specifically designed to drive instant analytics.

As modern businesses demand speed, scale and agility, traditional architectures that utilize an online transactional processing (OLTP) database together with an analytical engine or a data warehouse are unable to handle high throughput streaming ingestion, super low-latency analytics and concurrency needs that enterprises demand. A new approach to powering modern analytical data applications requires real-time insights or fast analytics on dynamic data for complex queries, all within sub-seconds.

### Key Challenges with Legacy Architectures



#### Streaming Ingestion

Inability to ingest, process and analyze fast-moving streaming data necessary to power real-time analytics



#### Batch Processing

ETL batch windows create delays in getting analytics to users and applications



#### Low-Latency Query Performance

Designed for pre-aggregated data sets, and not fast query performance or ad-hoc analytics



#### High Concurrency

Not optimized to handle large numbers of concurrent users or queries, unable to easily scale data applications or analytical dashboards

And what good is **having** data if you're not able to **use** it to improve customer experiences, deliver custom recommendations and ultimately, drive sales? The frustrating reality is that – even with custom workarounds that might enable faster data ingest – neither a transactional database nor a cached datastore will perform well for analytics. You're stuck managing yet another database to handle your analytics, which brings with it the latency, data degradation and frustration that come from porting information from one type of database to the other. And, in the end, you're acting on old insights that do little, if anything, to help you drive your business forward.

## SingleStoreDB: The Powerful, Scalable Modern Database

The good news? There's a better solution.

Modern SaaS applications need to be able to drive fast, interactive analytics on operational data in real time. And SingleStoreDB provides a modern, distributed database that unifies transactional and analytical workloads in a single engine to power real-time applications and analytics – empowering you to break through the performance bottlenecks of legacy architectures with an infinitely scalable, distributed system.

What you need is a responsive database architecture that fits your size and needs, and can deliver results in real time. One that combines data ingest, transactions and analytics in a single, continually updated, ultra-fast database. One that enables your users to enjoy a tailored experience with responses in the blink of an eye.

This database must be:

- **Relational:** Built on a relational foundation with all of the sophisticated query capabilities it imparts, in industry-standard ANSI structured query language (SQL).
- **Distributed:** Conceived, designed and built as a distributed system that does not require external middleware to segment, operate and scale.
- **Cloud-native:** Purpose-built to run on cloud infrastructure for elastic scaling of compute and storage in both hybrid and multi-cloud environments.
- **Multi-model:** Native support for relational and non-relational data types and data models.


In other words, if you want to empower scalable, real-time decision making, you need SingleStoreDB.

SingleStoreDB is a real-time, distributed SQL database designed specifically to handle today's business requirements at the scale you need. It offers the best of both worlds: the reliability and consistency found in relational databases and the full scalability that was formerly only available with NoSQL. Plus, SingleStoreDB simplifies your database architecture by unifying data ingest, transactions (OLTP) and analytics (OLAP) in a single, powerful data engine.

With SingleStoreDB, you can take advantage of:

- **Parallel high throughput streaming ingestion:** SingleStoreDB Pipelines is a built-in parallel data ingestion technology that enables users to natively ingest high-throughput, real-time data from external sources such as Apache Kafka, Amazon S3, Azure Blob, Filesystem, Google Cloud Storage and HDFS data source. Check out our blog on [How to Load 100 Billion Rows of Data Ultrafast](#).
- **Exceedingly fast analytics:** SingleStoreDB can process and analyze **more than one trillion database rows per second**, as established in this blog – [Shattering the Trillion-Rows-Per-Second Barrier With SingleStore](#). This enables low-latency, single-digit millisecond responses while scaling to handle thousands of concurrent users.

# Powering real-time applications with SingleStoreDB



- **Three-tier storage architecture:** Our unique tiered storage architecture spreads data across three tiers – in-memory row-store (Tier-1), disk-based columnstore (Tier-2) and Cloud Object Storage (Tier-3) – to maximize performance for the data needed most. Hot data is kept in memory, cooler data on local disks and cold data is kept in blob storage. This can be done without impacting query latency. All this coordination happens in the backend without any manual intervention, delivering ultra-fast performance at scale, while optimizing for costs.
- **Superior price-performance:** SingleStoreDB is designed for distributed environments, scaling out to meet demand with no external middleware. Separation of storage and compute allows for cost savings, as well as improved performance and elasticity. And it delivers a minimum of 10 times the performance of legacy databases. Moreover, it costs less than half of Snowflake-MySQL stack and is 2.5 times cheaper than Redshift-PostgreSQL, according to a [recent GigaOm Benchmark Report and Study](#).
- **Familiar and flexible SQL:** SingleStoreDB features familiar SQL tooling and is MySQL wire protocol compatible, facilitating quick adoption and go-live production. Powerful, industry-standard relational SQL queries are supported natively.

We go beyond other solutions to unify data types and workloads, including:

- **Multi-model data:** SingleStoreDB unifies all data types in a single multi-model database, supporting relational, document, key-value, time-series, geospatial, full text and streaming data that you can utilize to build your real-time application.
- **Multiple workloads:** SingleStoreDB unifies traditional rowstore and columnstore into a single table type, called [Universal Storage](#). This design breakthrough allows transactions and analytics – processing tasks that were previously mutually exclusive – to be executed simultaneously in the same table without any data movement or replication.

## Scalable Real Time in Real Life

Sounds too good to be true? It really isn't. And we have real-life data to prove it.

Some of the best known and most respected companies in the world rely on SingleStoreDB to deliver speedy ingest and lightning-fast analytics at the scale they need to meet their customers' expectations.



Take Siemens, a technology company helping customers build more resource-efficient factories, resilient supply chains, smarter buildings and grids, cleaner and more comfortable transportation, and more advanced healthcare. When the company first built their Siemens Pulse Analytics data dashboard to empower customers to access insights, they quickly found that they weren't able to serve up results fast enough. Painfully slow results led to frustrated customers. Not a good situation.

After a few false starts with other solutions, a Google search [led Siemens to SingleStoreDB](#). And they've never looked back. Siemens Pulse Analytics is now a hybrid containerized app combining SingleStoreDB, Kubernetes and AWS, and they're delivering impressive results for their customers, including:

- A **10-100X improvement** in data integration, which used to take 30 minutes.
- Streamlined processes – **from 10 steps down to just one**.
- The ability to **slice and dice data in real time** to determine how valid a correlation is.
- The ability to **query billions of rows and petabytes of data** in real time to accelerate time to insight. SingleStoreDB executes queries in **100 milliseconds or less**.
- **Autoscaling** that accommodates spikes in usage – from as few as 500 to as many as 100,000 concurrent users.

Super-fast results, streamlined processes and flexibility leading to actionable insights for customers. What's not to love?

## Conclusion

It's true that legacy databases played an important role in growing the business ecosystem to what it is today. But today's environment operates at an intense pace that legacy architectures simply cannot match. And businesses that embrace the outdated model of managing multiple databases for transactions and analytics won't be able to deliver the real-time, custom experience their customers have come to expect.

Modern businesses need a modern architecture — one that is purpose-built to handle data-intensive applications and combines streaming data ingestion, transactions and real-time analytics in a single, continually updated, ultra-fast database.

Modern businesses need SingleStoreDB.

The most successful companies in the world rely on SingleStoreDB to power their real-time, scalable data transactions and analytics. Shouldn't you?

## Ready to see what SingleStoreDB can do for you?

Visit [www.singlestore.com/cloud-trial/](https://www.singlestore.com/cloud-trial/) to sign up for a free trial of SingleStoreDB Cloud — delivering the full capability of a real-time distributed SQL database without the operational overhead and complexity of managing it yourself.