

Real-time insights has become a powerful differentiator for companies across many industries.

Transforming CX with embedded real-time analytics



During Black Friday in 2024, Stripe processed more than \$31 billion in transactions, with processing rates peaking at 137,000 transactions per minute, the highest in the company's history. The financial-services firm had to **analyze every transaction in real time** to prevent nearly 21 million fraud attempts that could have siphoned more than \$910 million from its merchant customers.

Yet, fraud protection is only one reason that Stripe embraced real-time data analytics. Evaluating trends in massive data flows is essential for the company's services, such as allowing businesses to bill based on usage and monitor orders and inventory. In fact, many of Stripe's services would not be possible without real-time analytics, says Avinash Bhat, head of data infrastructure at Stripe. "We have certain products that require real-time analytics, like usage-based billing and fraud detection," he says. "Without our real-time analytics, we would not have a few of our products and that's why it's super important."

Stripe is not alone. In today's digital world, data analysis is increasingly delivered directly to business customers and individual users, allowing real-time, continuous insights to shape user experiences. Ride-hailing apps calculate prices and estimate times of arrival (ETAs) in near-real

Key takeaways

- 1 Businesses can benefit significantly from real-time analytics. Companies that analyze and make use of the right data in near real time will see advantages, while businesses that lack real-time visibility into their operations face greater challenges to customer engagement and differentiation.
- 2 Real-time analytics can be measured using three metrics: latency, freshness, and concurrency. Queries should be quickly processed, new data should represent the current reality, and the infrastructure should be able to serve a large customer base.
- 3 Advances in data infrastructure and purpose-built software have significantly changed the economics of real-time analytics. Capabilities once considered cost-prohibitive can now deliver greater insight at a lower operational cost than non-optimized infrastructure.

time. Financial platforms deliver real-time cash-flow analysis. Customers expect and reward data-driven services that reflect what is happening now.

In fact, having the capability to collect and analyze data in real time correlates with companies' ability to grow. Business leaders that scored company in the top quartile for real-time operations saw 50% higher revenue growth and net margins, compared to companies placed in the bottom quartile, **according to a survey** conducted by the MIT Center for Information Systems Research (CISR) and Insight Partners. The top companies focused on automated processes and fast decision-making at all levels, relying on easily accessible data services updated in real time.

Companies that wait on data are putting themselves in a bind, says Kishore Gopalakrishna, co-founder and CEO of StarTree, a real-time data-analytics technology provider.

"The basis of real-time analytics is – when the value of the data is very high – we want to capitalize on it instead of waiting and doing batch analytics," he says. "Getting access to the data a day, or even hours, later is sometimes actually too late."

Business leaders that score
their company in the top quartile
for real-time operations saw

50%

higher revenue growth
and net margins

Source: Compiled by MIT Technology Review
Insights based on data from MIT Center
for Information Systems Research and
Insight Partners

The perception of real-time analytics is shifting. Once criticized as costly and of questionable value – often lacking clear connections between fresh insights and better decisions – it's now proving its worth through external data products that customers actively rely on. At the same time, purpose-built real-time analytics databases like Apache Pinot have eliminated the cost and complexity of retrofitting legacy systems, allowing real-time to not just be viable, but essential.

The benefits of real-time data analytics

Real-time data is also critical for Uber, and most companies have some functions that are equally reliant on real-time data, says Yupeng Fu, principal software engineer at ride-hailing service.

“For any business, there could be events, there could be changes happening in real time, and they have to react in a way to solve those problems, one-by-one, to further grow the business,” he says. “I really believe having this real-time capability is a game changer, or we can at least rethink many business problems in the context of real time capabilities.”

With companies increasingly adopting AI, the supply of real-time data and the capability to analyze that data will become essential. Machine learning and AI-based systems are fueled by high-quality, fresh data, and analytics and AI are the way the data is made useful. Successful companies will industrialize their data analytics, not only treating transactional data – but also the analysis of that data – as core data products.

For Stripe, real-time analytics creates value by enabling merchants to monitor cash flow trends, detect spikes in sales or refunds, track churn risk in subscription metrics, and surface anomalies as they happen – empowering

businesses to make smarter decisions and respond instantly to customer behavior or operational issues. Stripe has made those real-time capabilities part of its service guarantee, delivering the most up-to-date data to the customer's dashboard within 30 seconds.

While many businesses consider up-to-date data as a “nice-to-have” feature, soon it becomes indispensable, Bhat says. “Once people get used to the fresh fast data, it's like you can't go back,” he says. “Not all the reports need to be that fresh for all use cases, but once they get used to having that visibility in a particular workflow within their company, then they demand it.”

Understanding challenges

Yet, most companies face challenges in deploying these capabilities, because the infrastructure has traditionally been complex, large-scale, and expensive. Delivering fast responses to queries on the latest data, while minimizing computational overhead and serving millions of users simultaneously, is a significant software design and deployment challenge.

Companies need to provide fast responses to queries (known as query latency), quickly update the database from real-time source data (known as freshness), and serve millions of users (known as concurrency).

The trick is performing well across all these metrics at a very large scale, Gopalakrishna says. While general database systems can handle small-scale deployments, non-optimized infrastructure will quickly run into significant problems. “That's the challenge that most companies face: Existing batch systems are not built to solve for this scale, and when they try to retrofit, they try to do micro-batching – running the batch cycles faster and faster – they end up actually making it worse

“Once people get used to the fresh fast data, it's like you can't go back. Not all the reports need to be that fresh for all use cases, but once they get used to having that visibility in a particular workflow within their company, then they demand it.”

Avinash Bhat, Head of Data Infrastructure, Stripe

because it's very inefficient," Gopalakrishna explains. "And it isn't giving you what you really want, which is both the freshness of the data and the latency that is needed for the users."

Once a company decides to adopt real-time analytics at a large scale, the technology infrastructure needs to evolve, says Fu. "When scale is small, many things can actually work," he says. "When we started, there was only one PostgreSQL database for the entire infrastructure. The scaling part meant we had to adapt, and it became another forcing function for us to take into account with the solution."

Failing to redesign the data acquisition, processing, and analysis infrastructure can lead to a poor-performing and expensive system that will not support agile business operations, Fu says.

The last mile of open source real-time analytics

Historically, real-time analytics was considered prohibitively expensive and only used for the most mission-critical use cases. Many organizations attempted to retrofit traditional batch databases to support real-time workloads, leading to expensive and fragile architectures. A trio of open source technologies – Apache Kafka, Flink, and Pinot – built with real-time performance in mind have quickly become a common toolset.

Uber leverages open source technologies like Apache Kafka and Apache Pinot to build its real-time analytics infrastructure. **The ride-sharing service sees** a 64% reduction in page load times (from 14 seconds to under 5 seconds), reduced ingestion lag to under 10 milliseconds, and a decline in data loss and query timeouts.

The fact that Uber's service provider extensively uses an open source software (OSS) foundation for their technology is critical, says Fu. According to StarTree data, **Uber saw** a 70% reduction in infrastructure costs when the company modernized its infrastructure with a real-time focused solution, saving more than \$2 million per year, an 80% reduction in CPU cores, and a 66% reduction in data footprint.


"We can get many features for free from the community and also the technology is validated at scale by other companies, so OSS really improves the ROI of the engineering investment from our perspective," he explains. And there are additional benefits, "It also helps our hiring and the branding as well – we can get the talent more easily because this software is used and developed by other people, so it also makes our knowledge transfer to a new hire much, much easier."

Like Uber, Stripe's streaming data pipeline similarly includes open source components, such as Apache Kafka for storage, Flink for processing, and Pinot for analytics, aiming for greater self-service and cost effectiveness.



"The basis of real-time analytics is – when the value of the data is very high – we want to capitalize on it instead of waiting and doing batch analytics. Getting access to the data a day, or even hours, later is sometimes actually too late."

Kishore Gopalakrishna, Co-Founder and CEO, StarTree



With companies increasingly adopting AI, the supply of real-time data and the capability to analyze that data will become essential.

Stripe's data infrastructure establishes three critical capabilities. The first stage is a mechanism that captures any event that occurs on its platform, known as stream storage, and adds the event to its data lake. The second stage, known as stream processing, optimizes queries and the structure of the data to make data analytics more efficient. The final stage, known as stream analytics, allows users to query the pre-processed data in an efficient way.

"You need these three pieces to make real-time analytics work," outlines Bhat. "If you just use the analytics by itself, it can get very expensive, and you'll likely drop events coming into stream table, which will lead to missing data, and that doesn't work."

From edge case to core infrastructure

By adopting purpose-built real-time infrastructure, businesses gain the flexibility to consume and act on data in ways that drive continuous innovation. While some open source projects may support individual real-time use cases, they weren't necessarily designed for multi-tenancy or high concurrency, making them difficult to scale as real-time demands grow.

Real-time analytics is rapidly transitioning from a specialized capability adopted by digital-native enterprises to a foundational element of modern data infrastructure. Emerging real-time data sources are now central to competitive strategy in industries far beyond Silicon Valley. This shift is not solely about operational efficiency – it directly impacts the quality, personalization, and responsiveness of the customer experience.

As data volumes grow and latency-sensitive use cases proliferate, organizations across sectors are re-evaluating

An open source trio

Real-time analytics requires efficient collection, processing, and storage of real-time data, delivering low latency, high concurrency, and the ability to update the datastore in (near) real time. Three open source projects are often used to efficiently perform different stages of the pipeline.

Apache Kafka

A distributed event-streaming platform used for real-time movement of event data — whether from user interactions, IoT devices, or financial systems — by acting as a stable, high-throughput backbone for streaming data.

Apache Flink

A real-time data-processing engine that can perform transformations on data streams with low latency and high throughput, useful for event-driven applications or a component in a data pipeline.

Apache Pinot

A low-latency, high concurrency datastore for online analytical processing (OLAP). With its ability to handle hundreds of thousands of concurrent queries, it is well-suited for delivering customer and agent facing analytics.

Together, they create a complete real-time insights platform.

- Kafka streams high-throughput event data.
- Flink enriches and transforms the data stream.
- Pinot is where the analytics happen.


This trio can help enterprises evolve raw events into actionable insights in real time, at scale, and without the inefficiencies of traditional architectures.

how they process and operationalize information. Telecommunications providers rely on real-time network telemetry to maintain service quality, resolve issues proactively, and reduce customer churn. Retailers use real-time inventory and behavioral data to dynamically adjust pricing, improve product availability, and deliver personalized offers – enhancing both conversion and loyalty. Financial services firms use real-time analytics to detect fraud as it occurs, deliver instantaneous alerts, and provide tailored financial insights – raising both trust and perceived value.

In each case, the ability to act on live data is becoming integral to delivering highly responsive, differentiated customer experiences. Real-time analytics is no longer confined to the operational backend; it is increasingly the engine of customer-facing innovation.

Planning for success

Rather than treating real-time analytics as a niche solution for a handful of high-priority applications, companies should embrace a platform approach, embedding real-time analytics as a foundational layer within their broader data architectures, says Gopalakrishna. Part of this is reviewing how other innovative businesses are using real-time data to better manage operations and deliver new features to customers.



"Initially, keep things small, and don't try to boil the ocean in the beginning," advises Gopalakrishna. "Don't try to say, I'm going to turn everything into real-time analytics – you still need to keep in mind how your infrastructure actually scales."

"For any business, there could be events, there could be changes happening in real time, and they have to react in a way to solve those problems, one-by-one, to further grow the business."

Yupeng Fu, Principal Software Engineer, Uber

“Transforming CX with embedded real-time analytics” is an executive briefing paper by MIT Technology Review Insights. Laurel Ruma was the editor of this report, and Nicola Crepaldi was the publisher. MIT Technology Review Insights has independently collected and reported on all findings contained in this paper. We would like to thank the sponsor, StarTree, as well as the following experts for their time and insights:

Avinash Bhat, Head of Data Infrastructure, Stripe

Kishore Gopalakrishna, Co-Founder and CEO, StarTree

Yupeng Fu, Principal Software Engineer, Uber

About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of MIT Technology Review, the world’s longest-running technology magazine, backed by the world’s foremost technology institution – producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the U.S. and abroad and publishes a wide variety of content, including articles, reports, infographics, videos, and podcasts. This content was researched, designed, and written entirely by human writers, editors, analysts, and illustrators. This includes the writing of surveys and collection of data for surveys. AI tools that may have been used were limited to secondary production processes that passed thorough human review.

From the sponsor

At **StarTree**, we turn high-volume data streams – clickstreams, logs, vector embeddings, IoT signals – into instant insights. Built on Apache Pinot™, StarTree Cloud is a fully managed service that delivers sub-second analytics at scale for customer-facing apps, AI agents, and operational intelligence – while lowering costs compared to traditional architectures, an uncommon feat when increasing speed. Companies like Citi, Stripe, DoorDash, Together AI, and Wix use StarTree to serve customers and AI agents with fresh, instant insights at scale. We work with AWS, Google Cloud, Microsoft, and Confluent to help teams get the most out of their real-time data.

star✶**tree**

Illustrations

Cover art by Adobe Stock and spot art assembled by Chandra Tallman Design with art from The Noun Project and Adobe Stock.

While every effort has been taken to verify the accuracy of this information, MIT Technology Review Insights cannot accept any responsibility or liability for reliance on any person in this report or any of the information, opinions, or conclusions set out in this report.

© Copyright MIT Technology Review Insights, 2025. All rights reserved.



MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com