

5 Key Reasons to Choose Apache Pinot For User Facing Analytics

As organizations recognize the significant business benefits of user-facing analytics, they're increasingly turning to Apache Pinot as the best-in-class platform for high speed analytics at massive scale — for five key reasons.

User-facing analytics — also known as customer-facing analytics — has recently emerged as an essential value driver for successful brands. Capable of powering unprecedented growth and engagement, user-facing analytics nevertheless presents an enormous technical challenge to organizations, namely: *how to deliver interactive analytics on the freshest data to potentially millions of users*. Apache Pinot has become the leading solution to this problem.

Consider the ask: To deliver user-facing analytics, we need ultra-low (i.e., millisecond) latency at extremely high throughputs (commonly tens of 1000s of queries per second). We need to be able to query huge datasets that include historical data as well as data ingested in real-time, and we also need flexible indexing approaches in order to handle complex queries and maintain performance at scale. Finally, we need all of this to be economical, and simple enough for both data scientists and lines of business to use. Across each of these requirements, Apache Pinot stands out.



1. Blazing-fast query performance and massive throughput

In head-to-head comparisons, Apache Pinot demonstrates a speed and throughput that far exceeds that of its competitors. Currently supporting hundreds of 1000s of queries per second at millisecond latencies for companies such as LinkedIn, Uber, and Stripe, Pinot's performance is unrivaled in the industry.



2. Hyper-efficiency delivers the freshest data — at scale

Pinot's unparalleled efficiency ensures that users get actionable insights on both raw and aggregated data, with ingest support from batch sources like S3, HDFS, ADSL, and GCS, as well as streaming sources such as Kafka or Kinesis. And Pinot maintains real-time access to all that data even at *petabyte* scale.



3. Transformative indexing drives flexibility and speed

To deliver user-facing analytics in real-time, Pinot features the most powerful set of indexes in the industry — including inverted, sorted, range, JSON, text, geospatial, and star tree indexes. These are key to minimizing the work required for any given query and enabling industry-leading speed and scalability.



Apache Pinot helped deliver insights and uncover data blindspots with high accuracy in real-time... we were able to implement user-facing analytics on real-time data residing in Apache Pinot and discover opportunities to serve our customers better in real-time.



— **Saritha Ivaturi, Guitar Center | VP of Data Platform & Engineering**



4. Lower costs from better design

Pinot's advanced indexing and excellent compression (via columnar storage) reduce compute and storage requirements. Paired with support for multi-tenancy and tiered storage, the result is a highly performant and flexible database that saves you time and money.



5. One simple solution, multiple use cases

A single, production-ready platform for user-facing analytics, anomaly detection, and operational analytics, Pinot leverages cluster management, automation, observability, and quality of service across your multi-cloud to streamline delivery of user-facing analytics and accelerate your time to value.



Pinot enables us to execute sub-second, petabyte-scale aggregation queries over fresh financial events in our internal ledger. We chose Pinot because of its rich feature set and scalability, which has enabled better performance than our previous solution — at a lower cost.

— Peter Bakkum, Stripe | Engineering Manager

Talk to us about powering real-time analytics in your organization.