

Apache Druid®

The Analytics Database When Real-Time Matters

For decades, analytics have been confined to static executive dashboards and reports. Now, leading companies like Atlassian, Citrix, and Salesforce are going beyond traditional BI and realizing a whole new world of analytics use cases powered by developer-built applications. Developers are using Apache Druid to build real-time analytics applications that require sub-second queries at scale and under load on streaming and batch data.

Why Apache Druid?

Apache Druid® is the open-source, real-time analytics database used by developers at 1000s of leading organizations to power real-time analytics applications. These applications are used for operational visibility at scale, customer-facing analytics, rapid data exploration, and real-time decisioning. Unique capabilities include:

- **Sub-second at scale.** Druid's unique architecture delivers sub-second response times for OLAP queries on high-cardinality and high-dimensional data sets with billions to trillions of rows without pre-defining or caching queries in advance.
- **High concurrency at the lowest cost.** Druid features a highly efficient and massively parallel architecture that supports 100s to 100,000s of queries per second at consistent performance and uses less infrastructure than other databases.
- **Real-time and historical insights.** Druid was built natively for streaming data and features a connector-free integration for Apache Kafka® and AWS Kinesis. It supports massive scale event-based ingestion with exactly-once semantics and can query both real-time and historical data via a single SELECT.

“Druid is our choice for any application where somebody on the other end is waiting for a response. If you want super fast, low latency, that’s when we recommend Druid.”

Parth Brahmabhatt, Senior Software Engineer, Netflix

