

The Database for Modern Analytics Applications Built from Apache Druid

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 applications. Their developers are building analytics applications that enable interactive data experiences from streaming data and deliver real-time insights to both internal and external users. They are turning to Apache Druid with Imply to power these modern analytics applications.

Why Apache Druid?

Apache Druid is the open-source, real-time analytics database used by developers at 1000s of leading organizations to power modern analytics applications. These applications deliver operational visibility at scale, rapid drill-down exploration, real-time recommendations, and insights externally to customers.

Developers turn to Apache Druid for its unique ability to enable interactive analytics at any scale, high concurrency at the best value, and insights on streaming and batch data. Its hyper-efficient architecture delivers sub-second response on billions to trillions of rows for 100s to 1000s of concurrent users with near-infinite scale.

How is Imply, Built from Apache Druid, Unique?

Imply, powered by Apache Druid, is the right choice when powering an analytics application at any scale, for any number of users, and across streaming and batch data. Unique capabilities include:

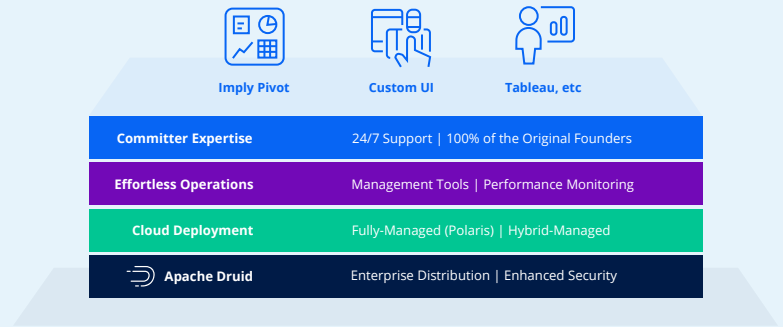
Sub-second at scale. From TBs to PBs and 100s to 1000s of concurrent queries, Druid's unique distributed architecture delivers consistent sub-second query response times - and does it without breaking the bank. It combines the performance of shared-nothing with the elasticity of shared-everything.

True stream ingestion. Druid was built for streaming data with native integration with Apache Kafka and AWS Kinesis. It supports massive scale ingestion of millions of events per second, query-on-arrival to analyze events instantly, and guaranteed exactly-once consistency to ensure data quality.

Non-stop reliability. For always-on applications, Druid is designed to never go down and never lose data. Its architecture is built for high availability and no data loss for streams via continuous backup and automatic recovery and rebalancing.

ImPLY Completes Druid

ImPLY delivers the complete developer experience for Apache Druid. Founded by its original creators, ImPLY adds to the speed and scale of the database with committer-driven expertise, effortless operations, and flexible deployment to meet developers' application requirements with ease.



Quick Comparisons



Cost efficient for infrequent queries and long-running reports, not for interactivity at scale



Easy in a single node cluster, challenging to manage at scale and potential data loss



General-purpose, not designed for analytics at scale



Good for transactions and simple aggregations, not for complex queries



Built for full-text search, performance limitations for analytics and reliability at scale

Key Use Cases

Observability

Drill-down visibility into the health and performance of each layer of a complex network or microservices application..

Clickstream Analytics

Harnessing clickstream data to optimize and monetize user experience from the first mouse click to the last.

External Analytics

Real-time operational and usage insights as standalone and embedded applications for external customers.

Security Analytics

Search, detection and investigation in real-time to quickly find anomalies to reduce risk and exposure.

IoT/Telemetry

Sensor data analysis for a wide range of use cases including predictive maintenance and condition-based alerts.