


Analyze event streams

Druid provides fast analytical queries, at high concurrency, on event-driven data. Druid can instantaneously ingest streaming data and provide sub-second queries to power interactive UIs.



Utilize reimagined architecture




Druid is a new type of database that combines ideas from [OLAP/analytic databases](#), [timeseries databases](#), and [search systems](#) to enable new use cases in streaming architectures.

Build event-driven data stacks

Druid integrates natively with message buses (Kafka, AWS Kinesis, etc) and data lakes (HDFS, AWS S3, etc). Druid works especially well as a query layer for stream hubs and stream processors.



Deploy anywhere



Druid can be deployed in any 'NIX environment on commodity hardware, both in the cloud and on premise. Druid is cloud-native: scaling up and down is as simple as adding and removing processes.