



Google Cloud
Partner



Ask Bigger with SQream Blue

The only data lakehouse for which no data set is too large & no pipeline is too complex

Digital-first businesses face bottlenecks and limitations when it comes to querying their rapidly growing datasets

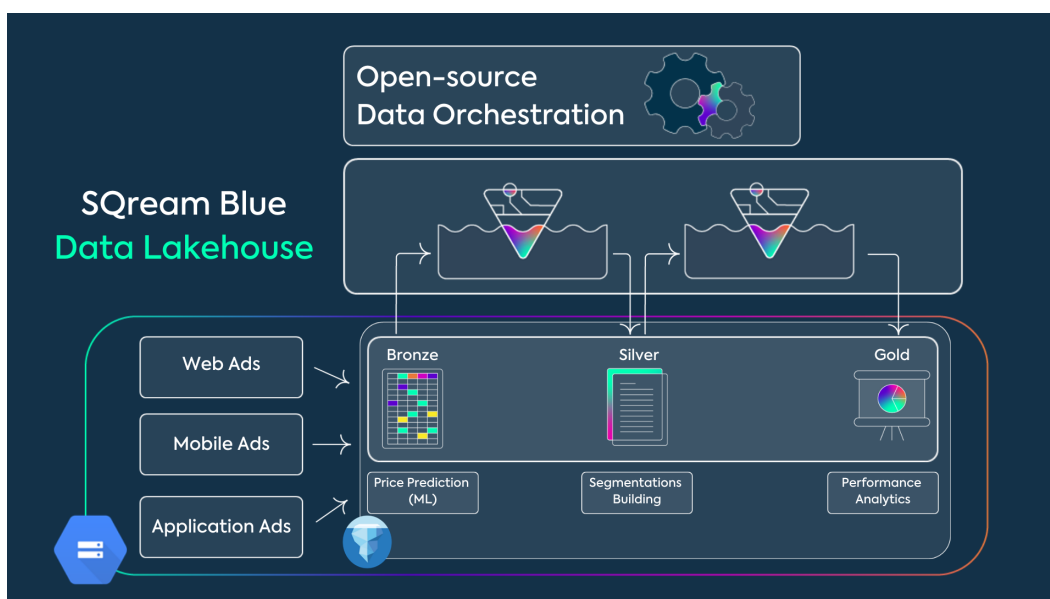
SQreamBlue is a cloud-native fully-managed data lakehouse built for fast, reliable, and cost-effective data usage utilizing a patented GPU optimization engine. The platform enables easy data preparation and transformation from and to the data lake, for faster analytics.

The SQream Blue advantage

- Unmatched scale**
Petabyte-scale data pipelines
- Lowest Price**
at 50% of other offerings
- Speed**
2x faster for data preparation

SQream Blue for Data preparation

Transform raw open-standard data into 'analytics-ready' for BI or ML. With a Medallion Architecture the data will never leave your cloud storage, and will be ready for analytics anytime. Orchestrate tasks like denormalization, pre-aggregation, feature generation, enrichment and validation at the speed of the GPU.



SQream Blue can be also used to analyze data stored in open-standard formats (ORC, Avro, Parquet, JSON) on cloud storage (data lake) with your favorite BI tool connected to Blue's GPU processing engine.

Solution highlights

GPU processing engine

Blue's performance leans on patented GPU-acceleration, synchronizing all available resources (CPU, GPU, RAM) and using the brute force of the GPU for the most complex analytical tasks. Blue uses the GPU to achieve parallel data processing. By splitting large tasks into smaller processes, SQream distributes operations between multiple GPU cores, while allowing admins to balance parallelism and concurrency according to their business needs

Architecture

SQream Blue doesn't require ingestion or data movement and relies on direct access to data in open-standard formats. Through the entire data preparation cycle, all data remains at the customer's low-cost cloud storage, maintaining privacy and ownership at best, while preserving a single source of truth and eliminating the need for data duplication.

Connectivity

Blue easily integrates with common open-source workflow management and orchestration tools (Apache Airflow, Dgaster, Prefect), along with support for industry-standard ODBC, JDBC, and Python connectors. Moreover, Blue's cluster management has a REST API.

Columnar optimization

Blue's processing engine utilizes Apache Parquet's column-oriented structure and metadata by saving unnecessary data read.

Trusted by global leaders



About SQream

SQream makes it possible to ask bigger questions of extremely large, complicated data sets. With no off-limits questions, analysts get unprecedented new insights at exceptional speed. For too long, high costs and complexity have caused big data projects to fail at an alarming rate. To succeed, companies need a shortcut for asking bigger data questions to make better decisions. With SQream, you can finally "dig deeper, go faster, and reach anywhere" so you're able to ask those bigger questions – on prem, in the cloud, or at the edge. This is why organizations ranging from fast-growth startups to Fortune 100s all rely on SQream. Ask bigger.™

To learn more, visit sqream.com or follow us on twitter @sqreamtech.