

Ask Bigger with SQreamDB

The only SQL database for which no data set is too large and no query is too complex

Introducing SQreamDB

SQreamDB is an SQL database that empowers organizations to query and perform analytics on complex, petabyte-scale datasets, and gain time-sensitive insights faster than ever and at 1/10 of the cost.

With SQreamDB you can ingest and prepare data, or instantly create ad-hoc queries using it as a query engine directly on your data lake. It enables you to get answers to key business questions with accuracy and at any scale.

SQreamDB is unique in that it doesn't take the typical approach of distributed data processing which causes cluster latency and bottlenecks.

Instead, it takes a proprietary approach of leveraging the parallel data processing of the GPU. Thus, by splitting large tasks into smaller processes, SQreamDB distributes operations between multiple GPU and CPU cores.

Therefore you can dig deeper into your data to uncover more meaningful and impactful insights, as well as to go faster and get these insights much sooner and far more cost effectively.

SQreamDB use cases



Data warehouse

Store and manage enterprise data, so decision-makers, business analysts, and data engineers and scientists can analyze it and gain valuable insights.



Query engine

Analyze data sets from any source, and in any data or table format, on top of existing analytical and storage solutions, without any data duplication required.



Data preparation

Transform raw data through denormalization, pre-aggregation, feature generation, and cleaning so it can be ready for BI and ML processing, and enriched as a part of the ETL/ELT pipeline.

Solution highlights

Query

Processing is performed with the ANSI-SQL syntax, where queries can run through the built-in SQream Acceleration Studio or through a third-party BI tool.

Architecture

Compute and storage are completely separated, without any processing performed in memory. Multiple compute units run in parallel, storing and retrieving data from a single or multiple storage sources, delivering greater flexibility, and allowing easy scalability.


Connectivity

Solution integrates into existing ecosystems and supports industry-standard ODBC and JDBC connectors, as well as Python and C#, .Net, Java, and others.

Compression

All ingested data is automatically compressed, at a 5:1 ratio, on average.

The SQreamDB advantage



Unmatched scale

querying up to
100 petabytes




Lowest TCO

at **1/3** of
other offerings



Complex queries

Handling
infinite complexity



Minimal hardware

1/10 less



Speed

5 x faster
to business insights



Smallest carbon footprint

1/10 less

Trusted by global leaders



About SQream

SQream makes it possible asking bigger questions on their data to make better decisions. 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. This is why organizations ranging from fast-growth startups to Fortune 100s all rely on SQream. Ask bigger.

For more information go to www.sqream.com.