

# Ask Bigger with SQreamDB

The only data analytics platform 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 one-tenth of the cost.

With SQreamDB you can ingest and prepare data, or instantly create ad-hoc queries using the capabilities of its query engine. It enables you to get answers to key business and operational 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 so 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 structured and semi-structured enterprise-scale data, so decision-makers, business analysts, and data engineers and scientists can analyze the data and gain valuable insights.



### Query engine

Analyze data sets from any source, any technology, 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, cleaning, and dataset combining, so it can be ready for BI and ML processing, and enriched as a part of the ETL/ELT pipeline.

## The SQreamDB advantage

### Unmatched scale

querying up to  
**100 petabytes**

### Complex queries

**Handling  
infinite complexity**

### Speed

**5 x faster**  
to business insights

### Lowest TCO

**at 1/3** of  
other offerings

### Minimal hardware

**at 1/10 less**

### Smallest carbon footprint

**at 1/10 less**

## 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 for 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.

## 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.

For more information go to [www.sqream.com](http://www.sqream.com).