

INSIGHTS AT SCALE WITH GPU ACCELERATED ANALYTICS AND ALL-FLASH DATA LAKES

CHALLENGE

We live in a data-driven world - where organizations that can quickly gain insight from massive data sets have a clear competitive advantage. The data warehouse is now a critical asset that gains value with each data point. But along with this growth comes challenges of scale that can lock away that value and slow innovation.

Most data warehouses were designed before the big data explosion, and struggle to support modern workloads. To make do, many companies are cutting down on their data pipelines, severely limiting the productivity of data professionals. There is frequently a trade-off between coarser data granularity and faster querying, making future fine-grained drill down analysis impossible. These restrictions compromise the quality and accuracy of the analytic results and end up costing businesses millions in missed insights and lost profits.

SOLUTION: THE NEW APPROACH TO DATA EXPLORATION

SQream DB harnesses the power of VAST Universal Storage and NVIDIA GPUs to perform fast, flexible, and cost-efficient analysis of massive datasets.

This powerful trio brings game-changing acceleration, with SQL query performance improvements up to 150% compared to GPU-equipped off-the-shelf x86 servers, all while reducing cost and system complexity.

The solution empowers data professionals to broaden their queries, breeze through petabytes of raw data, save time on data-preparation, and extract previously unobtainable insights in less time and at lower costs.

BENEFITS

FASTER TIME TO INSIGHTS

Up to 5X faster loading time and up to 3X query time vs CPU-based servers

ALL-FLASH STORAGE PERFORMANCE, ARCHIVE ECONOMICS

Data efficiencies drive down costs at scale

No tiering, no silos, all data at all-flash speed

TURNKEY & SIMPLE

Enterprise support and simplicity

Focus on your data not your infrastructure with appliance simplicity

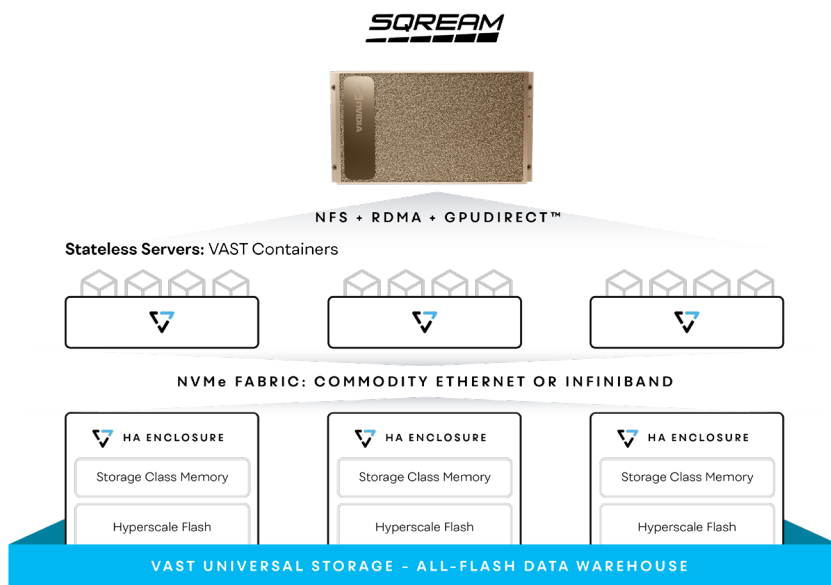
SCALE PERFORMANCE AND CAPACITY

Scale compute and storage independently

Linear performance with unlimited scalability

SHRINK YOUR DATA CENTER

Supercomputer capabilities in a small package for lower TCO



SQREAM ON VAST:

Proven Scale and Performance for modern big data workloads

SQream and VAST built a sample architecture for testing with the industry standard Transaction Processing Performance Council's Express Big Bench (TPCx-BB). The results demonstrated not only impressive performance for each workload but critically that performance scaled as the data set grew.

On the SQream + VAST architecture, the field test derived from the TPCx-BB benchmark results are dramatic. Starting with 30TB, followed by a massive 500TB load SQream + VAST achieved ingestion at an impressive 2.5GB/s. Complex query performance clocked in at an astonishing 14GB/s Read and 5GB/s Write. The combination of SQream + VAST seamlessly scales performance as data size increases to handle any big data challenge with linear scalability.

"In extensive lab validation, VAST achieved impressive Read and Write speeds with no signs of performance impact as the data set grew to massive scale."

– SQream

SQREAMDB

SQreamDB is empowering data consumers to Ask Bigger questions of their ever-growing datasets using the high throughput capabilities of the GPU.

SQreamDB was built from the ground up as a fully relational SQL database with extensive parallelism capabilities (not in-memory). SQreamDB completely disaggregates storage, compute, and metadata for lightspeed performance. As a columnar database, SQreamDB is optimized for heavy ad-hoc OLAP tasks, allowing every enterprise to Ask Bigger and gain more from its existing big data.

VAST DATA

VAST has reimaged every aspect of storage system design to break all the tradeoffs of DAS and shared-nothing architectures, delivering a system that provides unprecedented price, performance, and exabyte scale. The result is a highly scalable and affordable all-flash, file and object platform that allows you to run petabyte and exabyte scale analysis at less than half of the cost of traditional all-flash solutions while being many times faster. VAST DATA's unique Disaggregated and Shared-Everything (DASE) architecture with support for GPU direct access has the performance to match SQream's data processing capabilities and the scale for massive data sets.