SQREAM

Petabyte Scaling Using SQreamDB

Introduction

Analyzing petabytes of data can provide enterprises with valuable insights into complex problems, and help to make data-driven decisions that can have a significant impact on business outcome. Yet, despite the rapid development of modern data stack and cloud computing, analyzing the full scope of data at a petabyte-scale is very challenging. Most organizations find it too expensive or too slow and wearisome, so they compromise and base their analytics on a fraction of the data.

If your organization processes petabyte-scale data, analyzing it in full scope would benefit you in several ways:

- Greater accuracy of the insights gained from analysis, as more data points are taken into account.
- Deeper understanding of patterns and trends that may not be visible in smaller data samples.
- Improved predictive modeling for machine learning (ML) projects, as bigger and more quality datasets can make more accurate predictions.

In this technical whitepaper, we will explain how we implemented the cutting-edge technology of SQreamDB at one of the biggest electronics enterprises in Asia, allowing it to analyze data at a petabyte-scale, save tremendous costs on analytics, and increase yield.

Modern technology presents amazing opportunities

Modern technological advancements like innovative, high-throughput hardware acceleration present amazing opportunities for businesses to modernize the way they access their data. Legacy solutions were not designed to take advantage of high-throughput computing, such as multicore processors, scale-out storage solutions and hardware-accelerated algorithms, and may have limitations when scaling to meet the demands of modern data centers.

However, by using techniques like General Purpose Graphics Processing Unit (GPGPU) computing, it is possible to achieve more advanced processing capabilities with less hardware. GPGPU utilizes the massive parallel processing power of Graphics Processing Units (GPUs) to accelerate data processing tasks, allowing businesses to achieve faster and more efficient data analysis.

By embracing these modern advancements, businesses can gain access to more powerful data processing capabilities while reducing the need for costly hardware investments.

Petabyte-scale analytics: is it worth the investment?

Many enterprises have implemented unstructured and ungoverned data lakes, often built around the Hadoop ecosystem, as an alternative to traditional data warehouses that are limited in their data processing capabilities. However, the huge Hadoop clusters required to support these data lakes have resulted in a high total cost of ownership (TCO) and poor performance.

While these data lakes are flexible and can store large amounts of semi-structured and unstructured data, they require intense data preparation due to the "schema-on-read" approach. This need for data preparation has shifted autonomy away from data analysts and scientists, who are now forced to navigate difficult-to-use tools and jump through hoops to access the data they need.

The Hadoop ecosystem has ultimately failed to deliver on its promise of providing flexible and interactive access to data, and has instead discouraged data professionals by hiding data behind programming APIs and inadequate tools. As a result, businesses must consider the high TCO of maintaining large Hadoop clusters and evaluate whether the benefits of using these data lakes outweigh the costs.

General system architecture

SQreamDB is an enterprise-grade GPU-accelerated database for SQL analytics, designed from scratch to harness the raw brute-force power and high throughput capabilities of the GPU. With MPP-on-chip capabilities and a fully relational SQL database, SQream is designed to empower data consumers and support constantly-growing data. It is not an in-memory database or an SQL query engine for Hadoop or Postgres.

SQream accelerates data processes throughout the entire workload, from data preparation to insights, providing the best cost-performance compared to other available solutions. In contrast to traditional data warehouses and NoSQL data platforms, which tightly couple storage and computing together on the same physical nodes, SQreamDB decouples computing and storage, enabling scale-out of specific components to achieve linear scale.

Scaling these traditional systems for additional concurrent users, more complex queries, or additional data storage can be challenging and lead to performance limitations. SQream's decoupling approach provides greater flexibility and improved scalability. The solution is practical and built to support the needs of modern enterprises, empowering them with faster data processing and more advanced analytics capabilities.

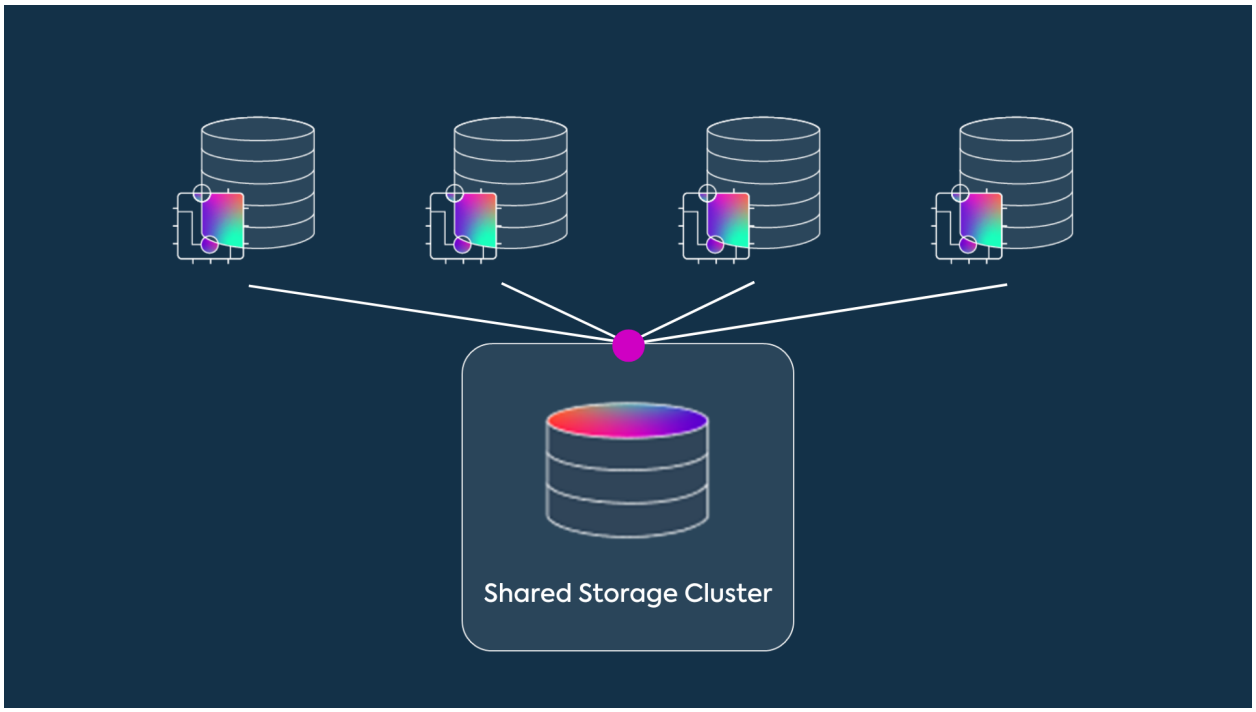

About the customer

A prominent electronics manufacturer in Asia was having difficulties with managing its data infrastructure. To address this issue and improve their business performance metrics, the manufacturer decided to embark on a project to centralize their data and create an all-encompassing artificial intelligence (AI) platform that included machine learning (ML) capabilities. The aim of this project was to identify and categorize production faults and defects as early as possible, using advanced anomaly detection algorithms, and take advantage ofthis insights to maintain and adjust the production floor machines accordingly. The ultimate goal was to optimize the production process and enhance efficiency.

The following chapter outlines the technical challenges that needed to be resolved and the reasons for previous difficulties. Additionally, it highlights how SQreamDB provided technical solutions that facilitated the creation of an AI platform.

Technical challenges and technical solutions

MES and MIS data loading

Challenge

Daily batches of Manufacturing Execution System (MES) and Management Information System (MIS) data loading **took over two days**. This extended loading time impacted the overall efficiency of the anomaly detection process, as the predictions **were not relevant by the time they were produced**. Additionally, the large volume of data loaded daily made it difficult to identify and investigate any potential anomalies or faults.

Solution

At the time, the manufacturer was using the Greenplum massively parallel processing (MPP) database management system over the Hadoop platform, which provided storage and parallel data processing capabilities. To reduce the time it took to load MES and MIS data, SQream recommended that the manufacturer improve the parallel processing capabilities by migrating to SQreamDB's centralized data warehouse management system. This would not only improve loading time but eventually would lead to a **lower Total Cost of Ownership** (TCO).

The SQreamDB system architecture introduced two operations critical to the efficiency of data loading: storing data in logical chunks and splitting the data ingestion into multiple sub-processes. Logical chunk storage, which can be thought of as micro-partitioning, **enables GPU core processes to run independently**. This is due to the fact that each chunk encompasses data of a specific column and logical unit (or sub-unit, such as hour, day, or geolocation), and each GPU worker is capable of performing ingestion and retrieval of data from a specific chunk. Splitting ingestion into multiple sub-processes supports the independence of GPU core processes since each sub-process may be assigned to a different GPU worker.

SQreamDB's ETL management system enabled the manufacturer to ingest **over 30TB a day** by using merely four servers equipped with 20 GPUs. SQream's centralized storage architecture also offered an advantage over Hadoop when executing SQL statements involving large JOIN clauses and aggregations, as the Hadoop-distributed HDFS systems required many compute nodes to locally access each piece of data. Hadoop's approach to reading and writing data from disk could result in costly operations, particularly when dealing with terabytes or petabytes of data, and it was not designed for in-memory calculations, leading to significant processing overhead. SQreamDB uses both RAM and GPU RAM to optimize the execution of statements, and additionally, its efficient query processing for hundreds of columns was

superior to Hadoop, where the number of nodes involved in query execution increased with the number of columns queried.

MES and MIS analysis capabilities

Challenge

The manufacturer faced challenges due to the lack of a unified analytical platform that could efficiently integrate data from their Manufacturing Execution Systems (MES) and Management Information Systems (MIS). This resulted in a fragmented approach that made it difficult for them to comprehensively analyze data and make informed decisions to optimize their operations.

The proliferation of disconnected data silos further complicated their data management, hindering their ability to access and refine meaningful insights. Without a centralized analysis environment for both MES and MIS data, the business could not gain the necessary operational insights to identify areas for improvement and make data-driven decisions to optimize production processes.

Solution

SQream has partnered with Weka and HITACHI Vantara to provide centralized analysis capabilities for MES and MIS data. ETL processes are consolidated into a centralized data lake, created using HITACHI Content Storage for File (HCSF), and managed by the Weka File System (WekaFS). The data lake is divided into two tiers: a high-performance block storage tier for structured data and an object storage tier for consistency and parallel access to each storage node. SQreamDB's shared storage architecture enables each server or instance to access the same data, share resources, and be centrally managed by Weka's single mounting point. HCSF optimizes and moves data chunks between tiers, and unlike the previous Hadoop-based system, Hitachi now saves a single copy of the data, reducing storage costs. This partnership enables customers to take advantage of a proven linear scale storage solution and, when combined with SQreamDB's GPU technology, optimize storage bandwidth and reduce costs.

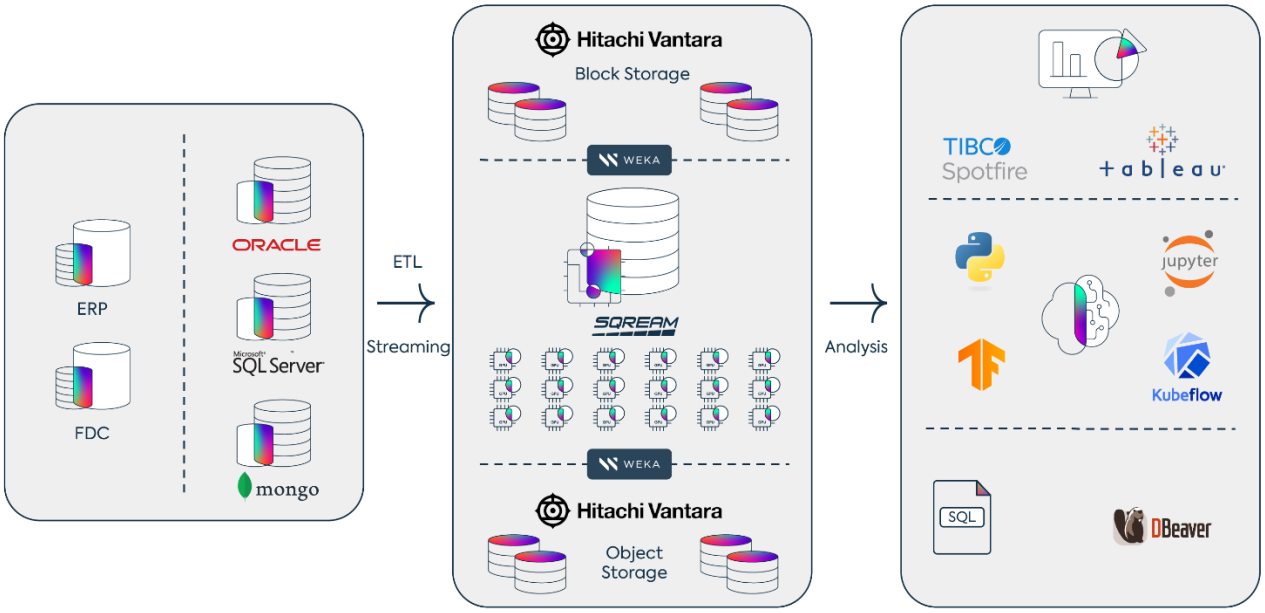

Historical data analysis

Challenge

The manufacturer needed to analyze a massive amount of historical data, with a scale of petabytes. They needed to analyze approximately five years' worth of data, which amounted to a dataset of approximately 12PB.

Solution

SQream's architecture leverages the power of GPUs for efficient data analysis without the need for intermediate steps. By utilizing the raw computing power of GPUs, data can be loaded and analyzed immediately. SQream also incorporates features such as fine-grained, low-overhead Zone Maps and Data Skipping to facilitate efficient analysis of large datasets. Unlike traditional data warehouses that rely on a fixed set of resources, SQream's flexible solution can allocate resources dynamically to handle varying workloads. This includes utilizing available CPU, GPU, RAM, and storage resources, with a focus on maintaining a balance between CPU and GPU operations for optimal performance. The SQream interface layer's statement compiler is designed to analyze each query and determine which parts should run on the CPU or GPU, with some queries running solely on the CPU to avoid unnecessary overhead.

High concurrency abilities

Challenge

The manufacturer needed to develop a system that could handle a high concurrency level for at least 100 data analysts and scientists.

Solution

SQreamDB's centralized data warehouse enables petabyte-scale data management in a single analytical environment, eliminating the need for a disparate data mart architecture and enabling simultaneous access to the same data by data analysts and scientists. To meet the manufacturer's need for high concurrency of over 100 data analysts and scientists, SQream enabled the creation of Resource Pools that can share GPUs between multiple processes executing different statements, with workload management assigning groups of SQream workers to different kinds of tasks based on the required resources. By efficiently managing the workload of 94 workers running on a 20 GPU compute architecture and split into 10 different services, SQreamDB achieves maximum utilization of hardware resources, allowing different users with different analytical tasks to access free resources and concurrency.

Case study achievements

Analytic achievements

- Automating the manual querying process resulted in 280 daily reports generated 8X times faster.
- Centralizing all structured and unstructured data on a single data lake environment for management and analytics.
- Preparing and preprocessing the data for the ML training pipeline - 50X times faster.

Business achievements

- Migrating from Hadoop to SQreamDB saved $25.5M for data collection and loading phases.
- Automatic reporting and shortening time for ad-hoc queries saved $3.5M per year.
- Better precision and prediction using the established ML platform improved the factory yield from below 50% to 90% in 8 months.

One more word about the Importance of ML data preparation

Machine Learning (ML) data preparation involves transforming raw data into a format suitable for machine learning algorithms to analyze and model. The quality of the data preparation process directly affects the accuracy and effectiveness of machine learning models. With the ever-growing complexity of ML models, the need for ML Operations (MLOps) has become essential in deploying successful models into production and continuously improving them with new data. SQream as a fast SQL analytical database for petabyte-scale can be useful for data preparation, distributed model training, and parallel inference due to its ability to perform many small processes across its GPU cores.

Transforming raw data into a "clean" format that can be used to train ML algorithms includes data preprocessing (cleaning, integration, transformation, and reduction) and data wrangling (filtering, grouping, enhancing features, and enhancing accuracy). Estimates regarding the time taken to prepare data for ML practitioners range from 40%—80% of the entire MLOps process. SQream's strength is its speed and efficiency in data preparation on large-scale datasets, especially for analytical use cases that require multiple JOIN clauses and aggregations.

Summary

This whitepaper describes how SQream enabled a leading manufacturer to overcome the challenges of data analytics and machine learning at a petabyte-scale while replacing its legacy data infrastructure. SQreamDB was implemented as the centralized data warehouse for the enterprise, delivering outstanding improvement in data loading and reporting times, reducing TCO of infrastructure, and boosting its yield up to 90%.

About SQream

SQream makes it possible to Ask Bigger questions of extremely large and complicated datasets. 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 or in the cloud. This is why organizations ranging from fast-growth startups to Fortune 100s all rely on SQream. Ask Bigger.™

To learn more, follow SQream on LinkedIn, Twitter, Facebook, or visit www.sqream.com