

Database records are generally comprised of data that is organized in a consistent manner. This means that the operations to query through this data can be disassembled into simple remedial processing steps.

Multicore CPUs are not sufficient to scale to billions of records. CPUs were designed for complex branch prediction and pipeline processing, and they accomplish this task well. However, this type of processing is not necessary when mining database entries.

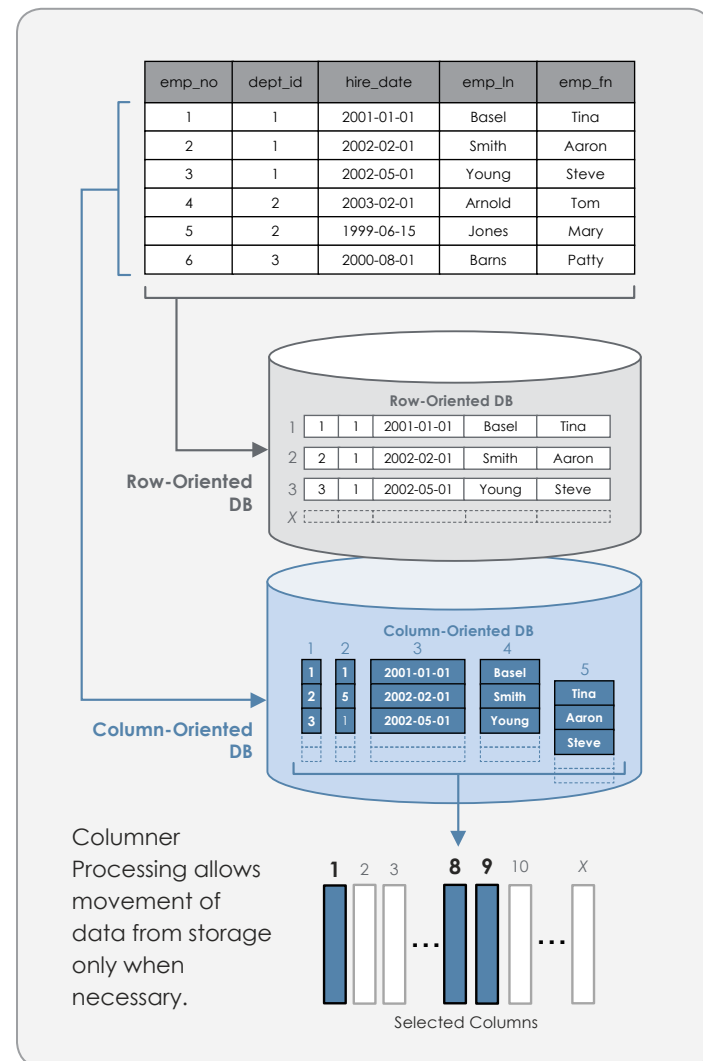
GPUs have thousands of cores that are specifically designed to perform remedial redundant processing. This type of processing is perfectly suited for database querying, which can involve processing over billions of records.

Technica was inspired by GPU processing and believes combining these results with other novel software design approaches lays out the path for efficient database processing that can be performed affordably using commodity hardware.

DESIGNED APPROACH

COLUMN-ORIENTED DESIGN


SQadron is a column-oriented database, designed for Online Analytics Processing workloads. Whereas a row-oriented design may be optimal for transaction processing, SQadron handles analytical queries more efficiently by targeting only the required data categories.



OPTIMIZATION TECHNIQUES


DYNAMIC COMPRESSION

Compression algorithms save disk space and reduce I/O time required to transfer data from internal hard drives to the CPU for processing. SQadron utilizes compression techniques that preserve order so that the data can be queried without decompressing.




DATA SKIPPING

Data skipping ignores large blocks of data that are not needed for the query. This further minimizes memory access and boosts read performance.




GPU ACCELERATION

SQadron leverages the massive parallel processing capability of the GPU for columnar database operations. Thousands of GPU cores boost query performance by executing tasks simultaneously, performing parallel compression, and operating directly on compressed data. Queries are further accelerated with a multi-GPU solution.



TIME SERIES DATA


Data points measured over time can be used to create models of reality. SQadron's analytic functions help analysts identify trends, cycles, and anomalies, and predict opportunities and failures BEFORE they occur.



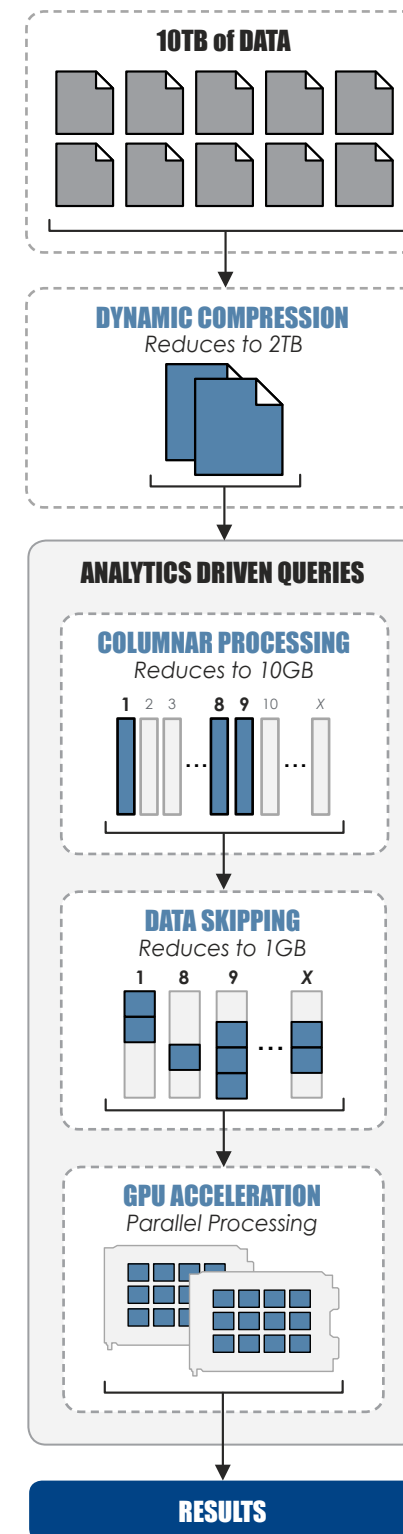
IN-DATABASE ANALYTICS

There are benefits to performing analytics within the database:

- Eliminating the need to extract the data.
- Saves time and cost, removing the need for a dedicated analytics server.
- Increases security by lessening the transfer of sensitive data.

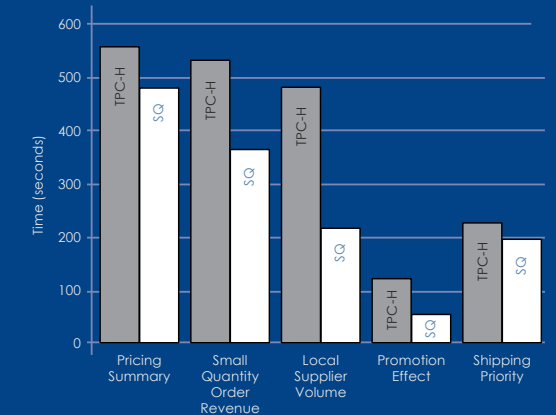


SQadron provides an in-database analytics library to allow for statistical and machine learning analysis beyond basic database query analysis. K-means clustering classifies a given query result into a specified number of clusters.



PERFORMANCE RESULTS

SQadron query execution times were validated against the Transactions Processing Performance Council TPC-H benchmark. The goal is to exceed performance for all queries.



Typical Competing System

SQadron's performance is achieved on a cost-efficient system using commodity hardware components. The typical competing system is a CPU-based in-memory database appliance, and sacrifices cost-efficiency to produce comparable results.

SQadron


- GPU: NVIDIA Tesla K-40c-12GB
- CPU: 6 Cores, 3.4 GHz
- RAM: 64 GB RAM, DDR3 1600 MHz
- Storage: 24TB RAID HDDs

Competitor

- CPU: 56 Intel Xeon ES-2680 v2 2.8GHz
- RAM: 1.5 TB RAM
- Storage: 50TB HDDs, 1.2TB SSDs



SQuadron™



MISSION

The goal of SQuadron is to provide terabyte scale database processing at an affordable price. SQuadron uses cutting edge software techniques and the parallel power of GPU processing to manage and query data efficiently. By utilizing affordable commodity hardware, SQuadron provides cost-effective solutions in support of Big Data analytics.



PROCESSING BIG DATA ON A BUDGET



SQuadron

Technica provides professional services, products, and innovative technology solutions to the Federal Government. We specialize in network operations and infrastructure; cyber defense and security; government application integration; systems engineering and training; and product research, deployment planning, and support.


- IT Solutions since 1991
- CMMI Level 2, DEV and SVC v1.3
- Certified PMP Resources
- Certified ITIL v3 Expertise

Technica®

Technica Dulles, VA (HQ)
22970 Indian Creek Drive, Suite 500
Dulles, VA 20166

703.662.2000

technicacorp.com



Technica®