

SOLUTION BRIEF

GPU Accelerated Insights with Hitachi Content Software for File & SQream

Hitachi Content Software for File & SQream joint solution provides the fastest time to market and time to insights with the best economics.

Enabling Fastest Time to Market and Time to Insights

Unprecedented data growth straining enterprise resources


Organizations worldwide are facing the challenge of managing an exponentially increasing variety and volume of high velocity data. While the term “Big Data” may be going out of fashion, the data formerly known as “Big” is not. What is changing are customer expectations for the platforms used to handle these massive datasets. Enterprises require solutions that deliver real time insights without consuming half the datacenter, or a massive cloud provider bill at the end of each month. A solution that scales linearly, allowing them to start small while planning big, to grow from terabyte scale to petabyte scale without introducing disparate systems and storage. Hitachi and SQream have partnered to deliver exactly such a solution. An analytics and data platform that merges the fastest file system in the world with the fastest GPU accelerated SQL analytic engine, and exabyte scale object storage to allow organizations to Ask Bigger.

A data platform for the modern enterprise

Hitachi Content Software for File is the fastest-growing data platform for modern enterprise workloads such as AI/ML and high-performance data analytics (HPDA). A data platform that enables enterprises to harness the power of their data to drive innovation, explore new market opportunities, and bring products to markets faster. Hitachi Content Software for File leverages cutting-edge cloud, compute, storage, and fast networking technologies to unleash the value of your data up to exabyte scale with fast access to data when needed. Our patented architecture eliminates the need for data copies across your entire workflow, reducing operational complexity, enhancing pipeline efficiency, and increasing GPU utilization.

SQreamDB is an SQL based analytics platform designed for accelerating petabyte-scale time-sensitive business insights. From ingestion to query, SQreamDB allows customers to perform complex analytics on a petabyte-scale of data, faster than any other solution. Accelerated by GPU, SQreamDB scales linearly on-premises or on private cloud to minimize time to insight and reduce TCO with less hardware. And when it comes to Big Data, no other SQL based platform can handle the variety of structured, semi-structured, and unstructured data as well as SQream.

Solution Benefits



Best of Object and File in a Single High-Performance SQL Analytics Platform:

Hitachi's platform integrated with SQream accelerates any ultra-high performance or high-capacity data analytics application including AI and ML workloads at massive scale.




Data Lake Acceleration:

Realize increased analytics performance while saving on infrastructure costs. The solution can supplement or replace your existing data lake/warehouse.



Speed up Data Ingestion:

Ingestion of terabytes to hundreds of petabytes of raw data, cutting prep time



GPU Accelerated Analytics:


Combining SQream GPU accelerated processing with Hitachi storage, the solution eliminates data access latencies that otherwise result in idle GPUs.

How it works

The foundation of the solution is the ultra-modern Hitachi Content Software for File parallel file system, a storage platform that delivers faster than local storage performance by leveraging the latest advancements in storage and networking technologies. The patented Hitachi Content Software for File software bypasses slow kernel based device drivers to directly manage networking and storage devices and eliminate the latencies that result in CPUs and GPUs waiting for IO. This means that more work gets done in less time with fewer resources. Already the fastest GPU accelerated SQL platform, when placed on the Hitachi Content Software for File foundation SQream really does what the name implies! Combined with Hitachi's world class object storage, the solution seamlessly provides the ultimate in scalability, enterprise grade data protection, and data lifecycle management at a much lower cost when compared to traditional Big Data platforms.

Shared High Performance Storage for Every Stage with Zero Unnecessary Copies

Edge / Ingest / ETL / Validate / Curate / Train / Analyze / Infer / Archive



Hitachi Content Software for File with SQream Integration

Notable Features

Effortlessly and linearly scalable at every level

- Gone are the days of throwing compute at a storage problem or vice versa. The combined solution enables easy scaling of only the constrained resources.
- Whether you're scaling up compute, storage, or object, you can expect linear or sub-linear performance gains.

Zero-copy Architecture

- With concurrent multi-protocol support for NFS, SMB, POSIX, Object/S3, NVIDIA® GPUDirect® Storage (GDS), and Kubernetes CSI customers can eliminate data silos by consolidating on a massive pool of shared Hitachi Content Software for File storage.

Built-In Data Protection

- The Snap2Object feature now supports storing snapshots on-premises or in the cloud or both for even more protection thereby eliminating the need for additional backup or Disaster Recovery (DR) software costs.

Resilience by Design

- Hitachi Content Software for File delivers resiliency at scale, with the data and metadata distributed across storage nodes, customers benefit from increased data protection without compromising scalability, capacity, or performance.

Enterprise Ready Features

- Hitachi Content Software for File offers the industry's most comprehensive end-to-end encryption across data at-rest and in-flight using XTS-AES 512 bit keys.

Enabling AI Driven Anomaly Detection for Global Manufacturer

Hitachi Vantara has created unique solutions to address the challenge of growing data within the enterprise, and the difficulty organizations face in analyzing the full scope of this data to achieve accurate business insights. Hitachi Content Software for File integrated with SQream's data acceleration platform is currently helping a major global manufacturer enable AI driven anomaly detection. Using the integrated solution, the manufacturer was able to ingest and continually analyze a multi peta-scale database composed of manufacturing machine sensor events, ingested to thousands of tables. This led to a significant improvement in Overall Equipment Efficiency.

“The joint solution allows customers to accelerate their data analysis while increasing data volumes. Run more queries and process ETLs faster and on larger data sets. Train your AI/ML model faster and data extraction is done faster. Allow your users to facilitate their work by applying ad-hoc analytics. The combination allows better interoperability between the storage and software while your data is turned into business value with unprecedented cost-effectiveness.”

Deborah Leff
CRO
SQream

ABOUT HITACHI VANTARA

Hitachi Vantara, a wholly-owned subsidiary of Hitachi Ltd., delivers the intelligent data platforms, infrastructure systems, and digital expertise that supports more than 80% of the fortune 100. To learn how Hitachi Vantara turns businesses from data-rich to data-driven through agile digital processes, products, and experiences, visit hitachivantara.com.

Learn More →

Hitachi Content Software for File

Blazing fast distributed file system for high-performance computing, AI, ML and analytics workloads.

Hitachi Vantara



Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
hitachivantara.com | community.hitachivantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
hitachivantara.com/contact

© Hitachi Vantara LLC 2024. All Rights Reserved. HITACHI and Lumada are trademarks or registered trademarks of Hitachi, Ltd. All other trademarks, service marks and company names are properties of their respective owners.

HV-JM-HCSF-SQream-GPU-Insights-29-Feb-2024