

SQream and Weka: Accelerating Critical Insight at Any Scale



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. SQreamDB works on-premises or on private cloud allowing linear scalability, accelerated by GPUs, minimized time to insight, reduced TCO and fewer hardware.

WEKA delivers a high performance Data Platform for analytics and AI/ML that leverages cutting-edge cloud, compute, storage, and fast networking technologies. Organizations using WEKA 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.

To win with data, you've got to Ask Bigger

Enterprises win with data. To win with data, you've got to Ask Bigger. You will always have these use cases that can't tolerate latency, those models otherwise you would just lose relevance.

Those use cases would often require the following to meet different needs:

- Operating with data in the tens, hundreds of terabytes or even petabyte-scale.
- Dealing with most complex and demanding queries.
- Costly investment in a never-ending amount of compute and disparate storage.
- Spending more time on ML model building and testing, and not on data preparation.

Until now, these use cases were off-limits. For this exact reason, SQream and WEKA partnered to help you address these challenges. Offering an easily scaled analytics data platform for enterprises that want to Ask Bigger, take advantage of their big data and use it for better business results. SQreamDB with WEKA Data Platform scales linearly, while maintaining as much as five times faster performance than the competitors for much less hardware and TCO.

SQreamDB and WEKA Data Platform

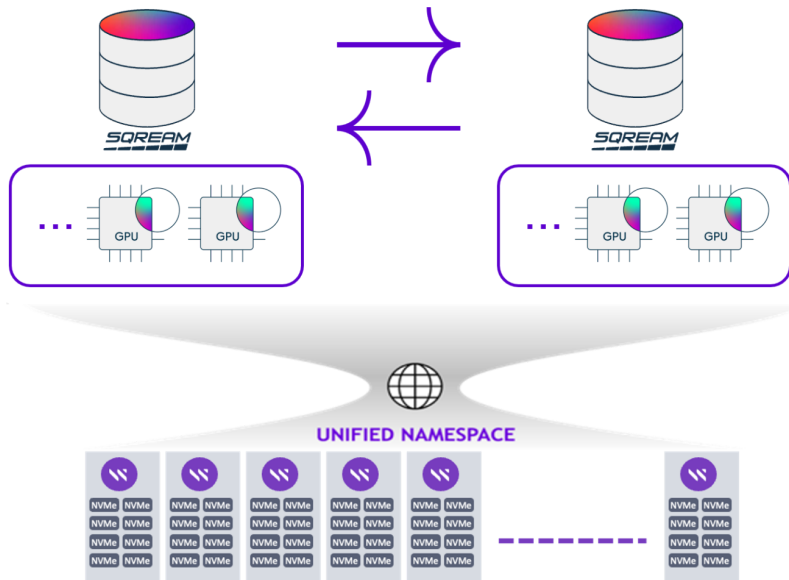


Figure 1: SQreamDB Compute Nodes accessing PBs of data from WEKA Data Platform

SQreamDB runs on off-the-shelf x86 compute servers, accelerated by state of the art NVIDIA GPUs. These systems are connected to WEKA Data Platform over high speed Ethernet or InfiniBand networking, accessing data over a single mount point.

SQreamDB and WEKA Data Platform is already deployed in leading enterprise customers environments and can be delivered on-prem as a private cloud through partners such as Dell, HPE, or Hitachi Vantara. Additionally, customers can leverage GPU, compute and storage resources in the managed clouds to run the end-to-end stack either as a hybrid cloud across on-prem and hosted clouds or as 100% hosted cloud.

Blazing Fast Performance for lowest cost

With SQreamDB working upon WEKA Data Platform, customers can prepare and analyze their data at light speed of less than 60 minutes when running for example *20 complex queries* on *30TB* and reduced TCO. Other solutions would perform poorly and cost more.

Environment	# GPU	Ingestion	Ingestion Rate (GB/s)	Ingest (TB/H per GPU)	Total query time	Total Time to Insight (TTTI)
WEKA + Intel Xeon SP Servers	2x NVIDIA A100	1:54:32	4.3	1.9	0:54:17	2:48:49

The above example is a field test derived from the TPCx-BB benchmark on a 30TB dataset in SQream Labs in July 2022. The testing was done with servers with Intel® Xeon® Scalable Processors and WEKA software version 3.12. SQream and WEKA can run a performance and functional proof of concept for qualified customers using the latest software and supported hardware.

Next steps

For organizations that base their business on data insights and need to make time-sensitive decisions, WEKA and SQream is the perfect combination for big data analytics. We are excited to offer a complementary technical consultation looking at your current and future needs and propose a proof of content. To learn more on how to deploy SQream on WEKA Data Platform infrastructure or run a proof of concept, contact sqream-solutions@weka.io or info@sqream.com or ask your WEKA authorized partner.

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.

For more information go to www.sqream.com.

