


High-Performance Analytics, Simplified



Dremio & VAST Bring Speed and Simplicity to Your Hadoop Workloads

What's next for your aging Hadoop infrastructure? Hadoop ecosystems, which were once at the forefront of revolutionizing data analytics, are no longer able to meet the demands of modern enterprises due to poor performance, costs and complexity. High-latency reporting has given way to instant insights. Data governance and lineage are now fundamental requirements rather than afterthoughts. Hadoop cluster management is complex and requires special staffing in order to maintain operations. Management complexity is further aggravated by the need to maintain governance and security frameworks that provide limited function. Tasks need to be run to gather and transform data from disparate sources, duplicating data into silos that quickly get out of sync. The list goes on. Dremio and VAST are modernizing these environments, by eliminating silos, improving performance and increasing security while reducing costs and simplifying nearly every aspect of management.

VAST and Dremio pair to deliver a simple and elegant analytics solution combining the performance, scalability and simplicity of a VAST all-flash data lake with Dremio's powerful SQL query engine and semantic layer. Dremio gives you a unified access layer over your data lake and all of your ancillary data sources in the cloud or on-premises, complete with comprehensive governance, data lineage, and security capabilities.

Why VAST + Dremio?



Faster Time To Insight

VAST's massively parallel architecture delivers concurrency and real-time access to billions of files and objects for performance at any scale.



Built for Scale

Elastically scale storage and performance as you add more users, query engines, or simply have more data.



Best in Class Price-performance

Get all-flash performance for all data operations and searches at 1/2 the cost of traditional flash solutions.



Superior Data-Reduction

Get up to 1.5:1 additional data reduction for pre-compressed Parquet files and 3.5:1 for un-compressed CSV files

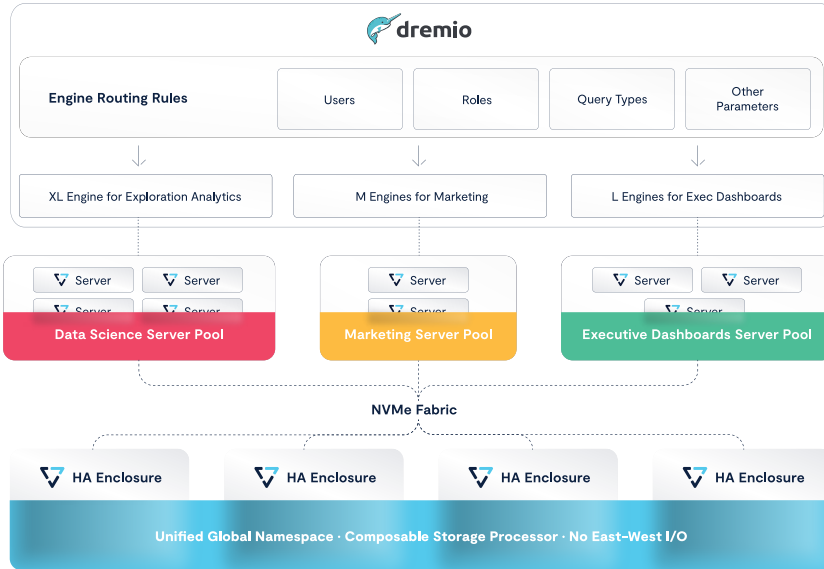


No Resource Contention

Dedicated QoS for Dremio's query engine ensures demanding analytics jobs get consistent performance

The VAST Data Platform

VAST's Similarity-Based Data Reduction delivers advanced levels of data reduction where traditional deduplication techniques fail. VAST's data reduction algorithms reduce data in similar blocks down to minimum sized delta blocks by compressing similar blocks together. The result is a best of all worlds approach that brings the effective cost of flash in line with HDD-based systems for backup scenarios without compromising on performance or resiliency.



Consolidate analytics workloads on the same namespace

Eliminate noisy neighbors with granular QoS engine

No noisy neighbor. Allocate containers per engine for QoS


Eliminate islands of infrastructure

One Universal Platform For All Your Data

The VAST Data Platform combines exabyte levels of scalability with multi-tenant quality of service capabilities in order to consolidate all data and all applications onto one scale-out tier of affordable flash. Everything is simple when you no longer need to move data across silos of infrastructure.

Trusted by Industry Leaders

VAST is trusted by leading companies around the world to leverage a single radically affordable all-flash innovative platform.



About VAST

VAST Data is the data platform software company for the AI era. Accelerating time-to-insight for workload-intensive applications, the VAST Data Platform delivers scalable performance, radically simple data management and enhanced productivity. Launched in 2019, VAST is the fastest-growing data infrastructure company in history.

Looking to simplify your data center and unlock insights from all of your data?

Contact us at hello@vastdata.com

About Dremio

Dremio is the open and easy data lakehouse, providing self-service SQL analytics, data warehouse performance and functionality, and data lake flexibility across all of your data. Hundreds of organizations, including 3 of the Fortune 5, use Dremio to deliver mission-critical BI on the data lake. As the original creator of Apache Arrow, Dremio is on a mission to reinvent SQL for data lakes and meet customers where they are in their cloud journey. Dremio was founded in 2015 and is headquartered in Santa Clara, CA.



[Scan Here to Read](#) more about VAST + Dremio on our website