

EXECUTIVE BRIEF

Dremio

SQL Lakehouse Platform


Enterprises today need data architectures that can rapidly adapt to the zettabytes of data created by modern-day users and applications. Not only is the data growing exponentially, but user requests for data are also growing rapidly. Everybody wants access to the data and wants to query it themselves. Executives and IT leaders want to empower their business users to have timely and self-service access to organizational data, so that they can support new initiatives and evolving business needs.

As cloud data lake storage becomes increasingly cheap, enterprises are dumping near-infinite amounts of data in their storage. Data architects and Chief Data Officers (CDOs) know the pain and time-consuming process required to set up and maintain complex ETL/ELT processes to copy multiple subsets of that data into a data warehouse for different use cases and reporting purposes. And to make things worse, additional copies of data are frequently created for personalized needs, aggregated analysis, specific use cases and dashboards. With this magnitude of data explosion, the bigger question is where and how organizations will store the data and how they will access this data without making multiple copies of that data and at what price.


CDOs and data architects constantly worry about data drift and lack of data governance and security as they deal with the complexity of such data architectures and multiple data copies. Data consumers cannot access the data themselves and end up relying on data teams, who in turn get overwhelmed with weeks or months of data request backlogs. This environment is best described as the “Pyramid of Doom”:

According to a March 2021 IDC study, in 2020, 64.2 ZB of data was created or replicated and the amount of digital data created over the next five years will be greater than twice the amount of data created since the advent of digital storage.

Pyramid of Doom



A No Copy Architecture



So what should enterprises do to save themselves from the Pyramid of Doom? As more and more organizations migrate to the cloud, or aspire to optimize their existing investments in the cloud, decision-makers are trying to break free from the intrinsic dependencies on traditional data warehouse vendors and liberate their organizations and budgets from the lock-in and economic extortion of proprietary data vendors. Leaders are looking for solutions that provide them the flexibility, cost-effectiveness and scalability of data lakes as well as the performance and data management capabilities that traditional data warehouses have provided. Chief Technology Officers (CTOs) and CDOs want their organizations and architectures to be flexible, open, scalable and ready to change quickly to adopt new technologies and innovations as they become available in the market.

A December 2020 McKinsey Research article, warns CIOs and CDOs that one of the core misbeliefs is that it is always better to deploy vendor-based technologies. McKinsey advises that, in reality, open-source solutions can be cheaper and more resilient.

The Dremio SQL Lakehouse Platform

Dremio is a high-performance SQL lakehouse platform built on an open data architecture. Dremio's SQL lakehouse platform simplifies data engineering and eliminates the need to copy and move data to proprietary data warehouses or create cubes, aggregation tables and BI extracts, providing flexibility and control for data architects and data engineers and self-service for data consumers. Dremio enables high-performing BI dashboards and lightning-fast interactive analytics directly on data lake storage through seamless integrations with Tableau, Power BI and other BI tools, using multiple ground-breaking Dremio technologies and innovations.

Unlike traditional data warehouse vendors, which require companies to load data into their vendor-owned storage as a prerequisite to using their SaaS offering, Dremio Cloud allows companies to keep the data in their own cloud account. This not only improves security and data governance, but also eliminates vendor lock-in, and allows companies to be present and future-ready and use other best-of-breed compute engines on that data.

Dremio Cloud, a frictionless, infinitely scalable service, enables high-performance SQL workloads directly on cloud data storage, eliminating the cost and complexity of copying and moving data. Dremio Cloud's elastic multi-engine architecture scales infinitely and helps manage workloads with the highest concurrency and lowest latency.



The Dremio Advantage

With Dremio, shorten time to insight from days or weeks to minutes with a simplified, no-copy data architecture, high-performing BI dashboards and interactive analytics directly on the SQL lakehouse:

- The **no-copy architecture** eliminates the long backlog of requests and empowers data scientists, data analysts and data consumers through self-service access to business-ready data in real time.
- The industry's only vertically integrated **semantic layer**, with shared business logic, data access, KPIs and security, provides a single, consistent view of the data for all teams and tools and applies security and data governance across the company.
- This shared logic and security eliminates business logic duplication across multiple BI tools and **prevents multiple data copies** at different levels — no more cubes, aggregation tables or extracts — and recurring data drift.
- **Seamless** single sign-on (SSO) **integrations with BI tools** provide lightning-fast queries directly on the cloud data lake with sub-second responses from any BI or dashboard.
- Dremio's **high-performance execution engine**, powered by Apache Arrow, speeds queries by 100x and eliminates >90% of compute costs.
- With Dremio Cloud, your data and data processing **stays in your cloud account**. Within your account, data is encrypted at rest and in transit. Secure your data further with support for customer-managed encryption keys and dynamic data masking.
- You can deploy Dremio anywhere! Run Dremio Software in your own environment, or choose Dremio Cloud, a fully-managed SQL lakehouse platform built on an open architecture.

Get Started today with Dremio Software or Dremio Cloud.

ABOUT DREMIO

Dremio is a SQL Lakehouse Platform company enabling organizations to leverage open data architectures. Dremio's SQL Lakehouse Platform simplifies data engineering and eliminates the need to copy and move data to proprietary data warehouses or create cubes, aggregation tables and BI extracts, providing flexibility and control for data architects and data engineers, and self-service for data consumers. Dremio Cloud, a frictionless, infinitely scalable service, enables high performance SQL workloads directly on cloud storage, eliminating the cost and complexity of copying and moving data. Dremio Cloud reimagines the traditional data lake by combining the best of traditional data warehouses and data lakes into a SQL lakehouse, while removing the limitations of traditional data warehouses resulting from closed data architectures.

Dremio and the Narwhal logo are registered trademarks or trademarks of Dremio, Inc. in the United States and other countries. Other brand names mentioned herein are for identification purposes only and may be trademarks of their respective holder(s). © 2021 Dremio, Inc. All rights reserved.