

## SOLUTION BRIEF

# Accelerate Time to Insight with FlashBlade for Dremio

Transform traditional data lakes and data warehouses into a modern unified analytics platform with Pure Storage® FlashBlade® for Dremio.

As organizations become smarter, extracting value from their rapidly growing data assets becomes imperative. Traditional data warehouses can be rigid, siloed, and struggle to keep pace with increasing data needs. These data warehouses can also lock data into third-party formats. Data lakes offer open data formats as well as the ability to store unstructured and semi-structured data. However, legacy Hadoop-based data lakes combine compute and storage, making them complex to scale. They are also slow to query using traditional business intelligence or SQL-based tools. Further, these legacy solutions can limit the ability to leverage container-based architecture to achieve agility and elasticity.

## Pure FlashBlade for Dremio

FlashBlade for Dremio helps you create a unified analytics platform that brings the advantages of data lakes and data warehouses together in a single solution. With Dremio's data lake query engine, you can directly query data stored on Pure FlashBlade. Whether you store data as files, in an open columnar format such as Apache Parquet, or an Apache Hive data lake, Dremio simplifies running SQL queries without making copies of your data.

Pure FlashBlade's simplicity and consolidation allow for data growth without compromising on performance. The joint solution offers the following benefits:

- **Lightning-fast Queries:** Dremio delivers 4x to 1000x faster performance when compared to traditional data lake engines by using a single high-speed consolidated data lake and data warehouse.



### Accelerate All Queries

Get all-flash performance for all data operations and searches.



### Operational Simplicity

Scale compute and storage independently. Add blades and rebalance data faster.



### Improved Agility

From business analytics to data science, simplify handling of use cases and offer them as a service.

- Scalability and Elasticity:** When combined with FlashBlade, Dremio's elastic engines provide elastic, on-demand scale, thanks to completely disaggregated cloud-native architecture. Engines can automatically start and scale to full size, then scale back down and stop according to query traffic. Add storage capacity non-disruptively and independent of compute simply by adding blades to a Pure FlashBlade. FlashBlade delivers linear, predictable performance while scaling to deliver tens of billions of files and objects. You can consolidate storage for other parts of your data pipeline (Apache Spark, Kafka, Elastic, etc.) on a single storage platform.
- A Flexible and Open Data Platform:** FlashBlade for Dremio enables you to store structured data as well as unstructured data in various formats and query them using different protocols (NFS or S3 protocols) in a single unified environment. This enables you to store data in open industry-standard formats, such as Parquet, or connect to a Hive metastore, making it easy to migrate from legacy and cloud data lakes.

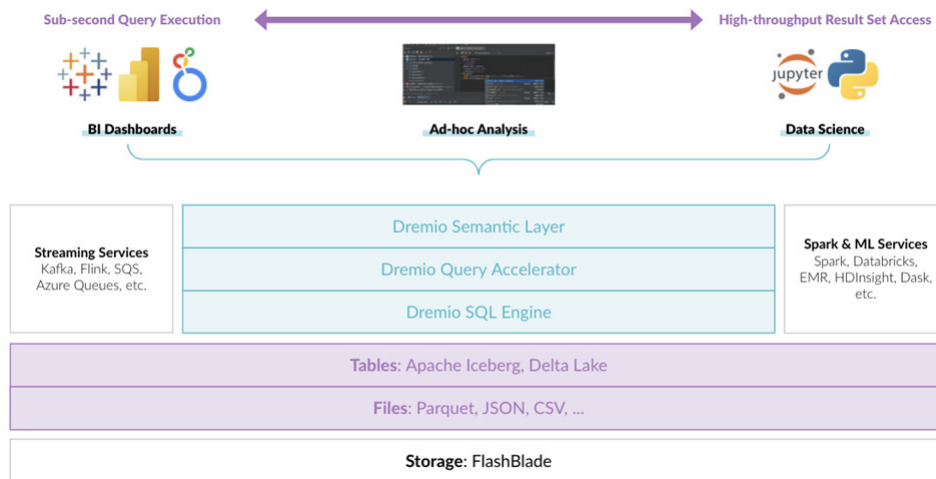


Figure 1. Pure FlashBlade for Dremio application solution

## About Dremio Data Lake Query Engine

Dremio accelerates time-to-insight by providing lightning-fast queries and eliminating the need to copy and move data from your FlashBlade data lake to proprietary data warehouses or create cubes, aggregation tables, and/or BI extracts. A self-service semantic layer provides flexibility and control for data architects and self-service for data analysts. Think faster transactions and more immersive customer experiences.

## Additional Resources

- Learn more about [deploying Dremio on FlashBlade](#).
- Explore what FlashBlade can do [for your organization](#).
- Take advantage of a free [FlashBlade test drive](#).
- Learn more about [Dremio](#) and take a free [Dremio test drive](#).

purestorage.com

800.379.PURE

