LANL's 2024 Efforts in Enabling Near-Data Analysis and Computational Storage

Qing Zheng, Los Alamos National Lab

qzheng@lanl.gov

LA-UR-24-32990

Scientific Data Analysis

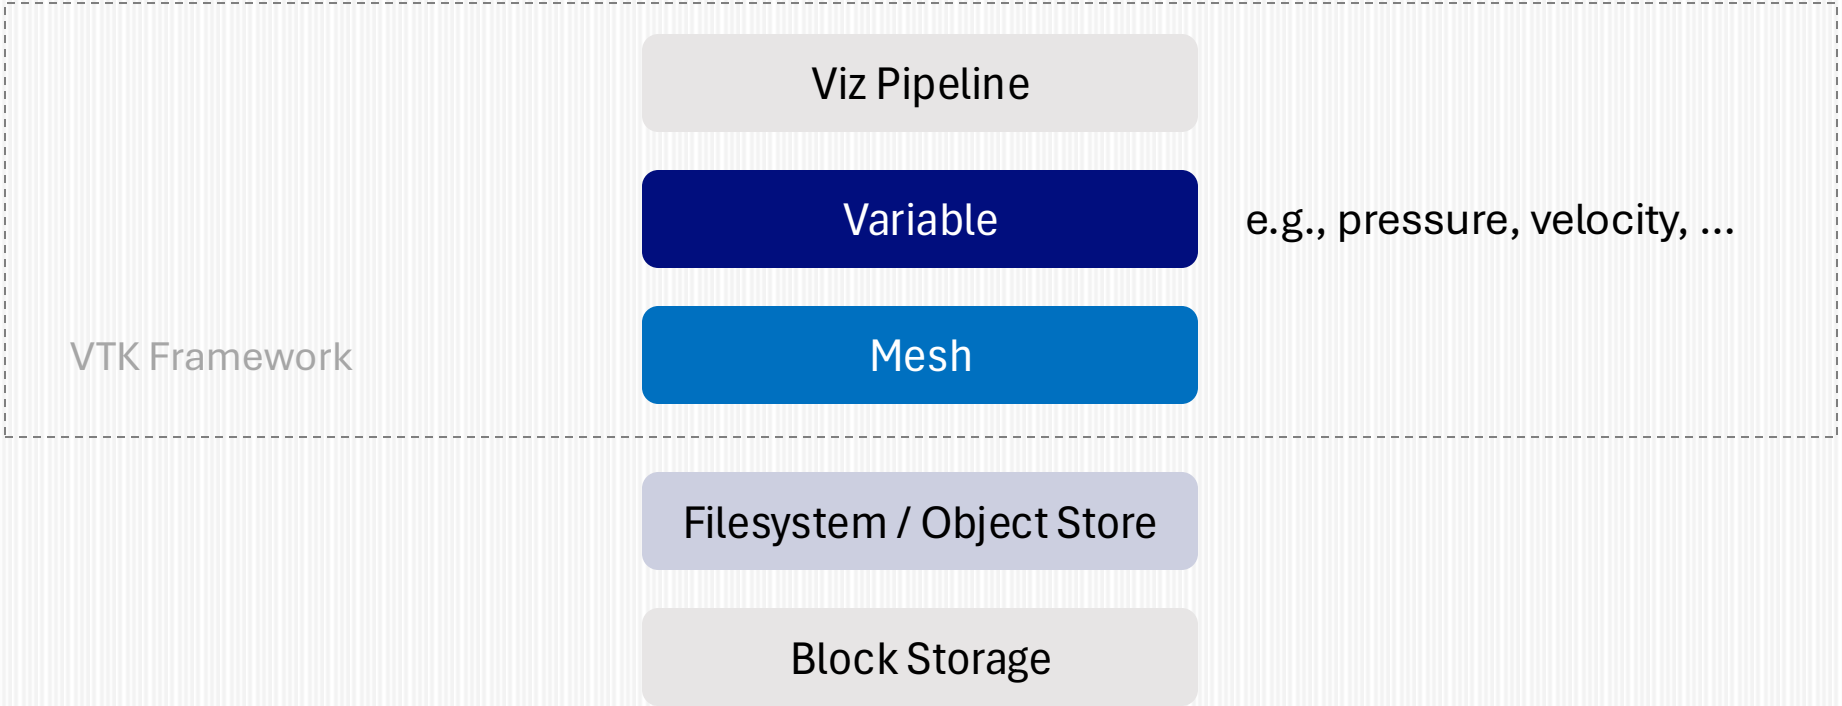

A Relational Perspective

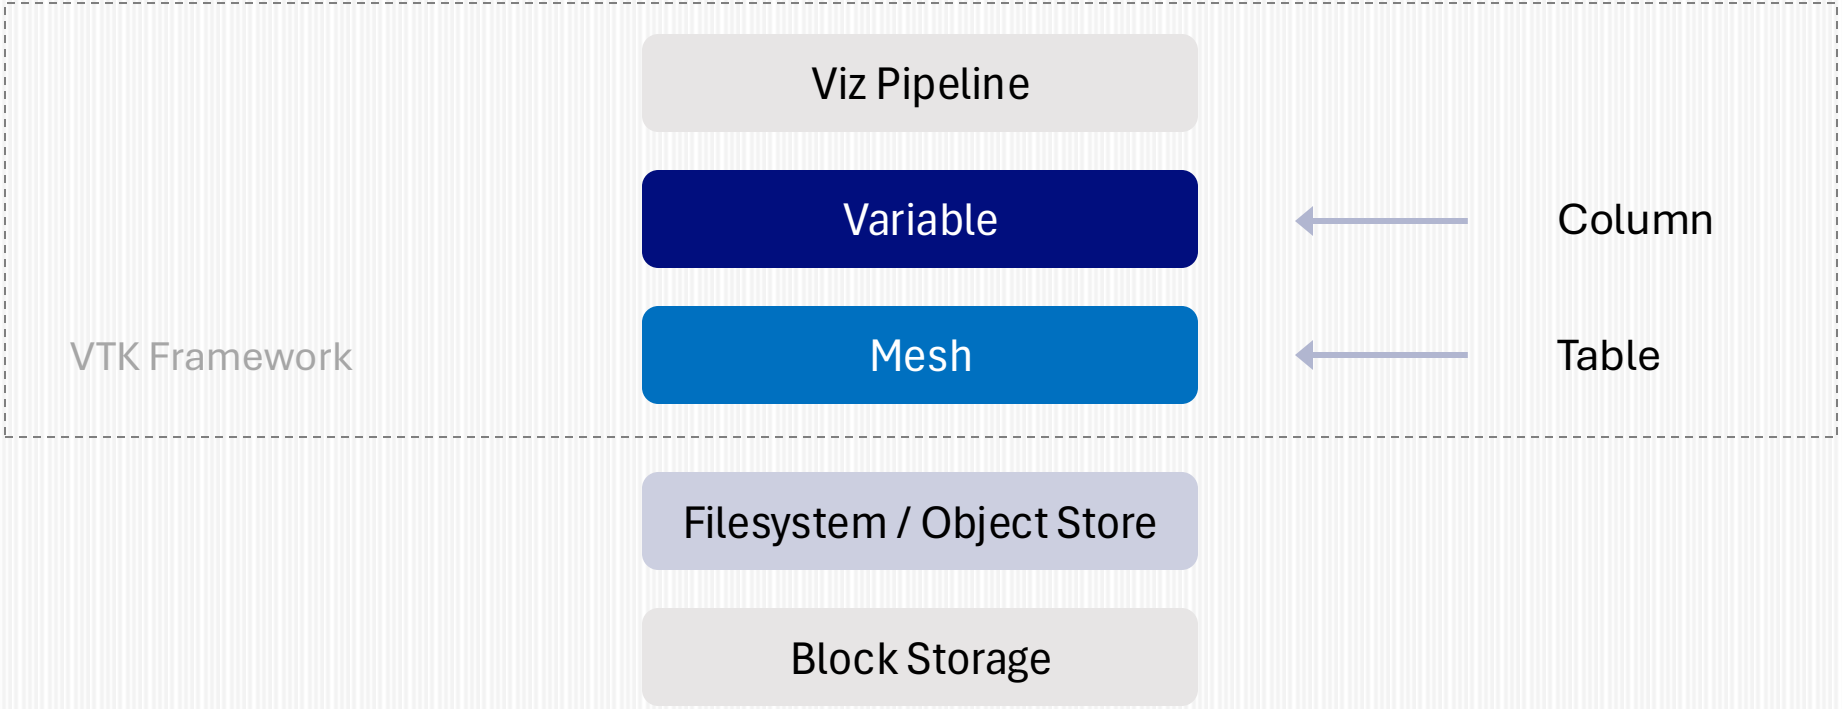

Converging Worlds: HPC and Big Data Analytics

HPC analytics resemble the rest of the world: data is big, moving data is expensive, and analysis often targets a small data subset

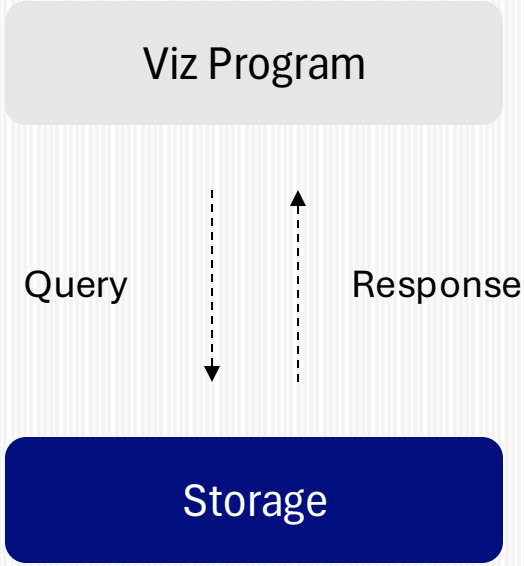

Requires techniques to control and reduce data movement

Computational Storage

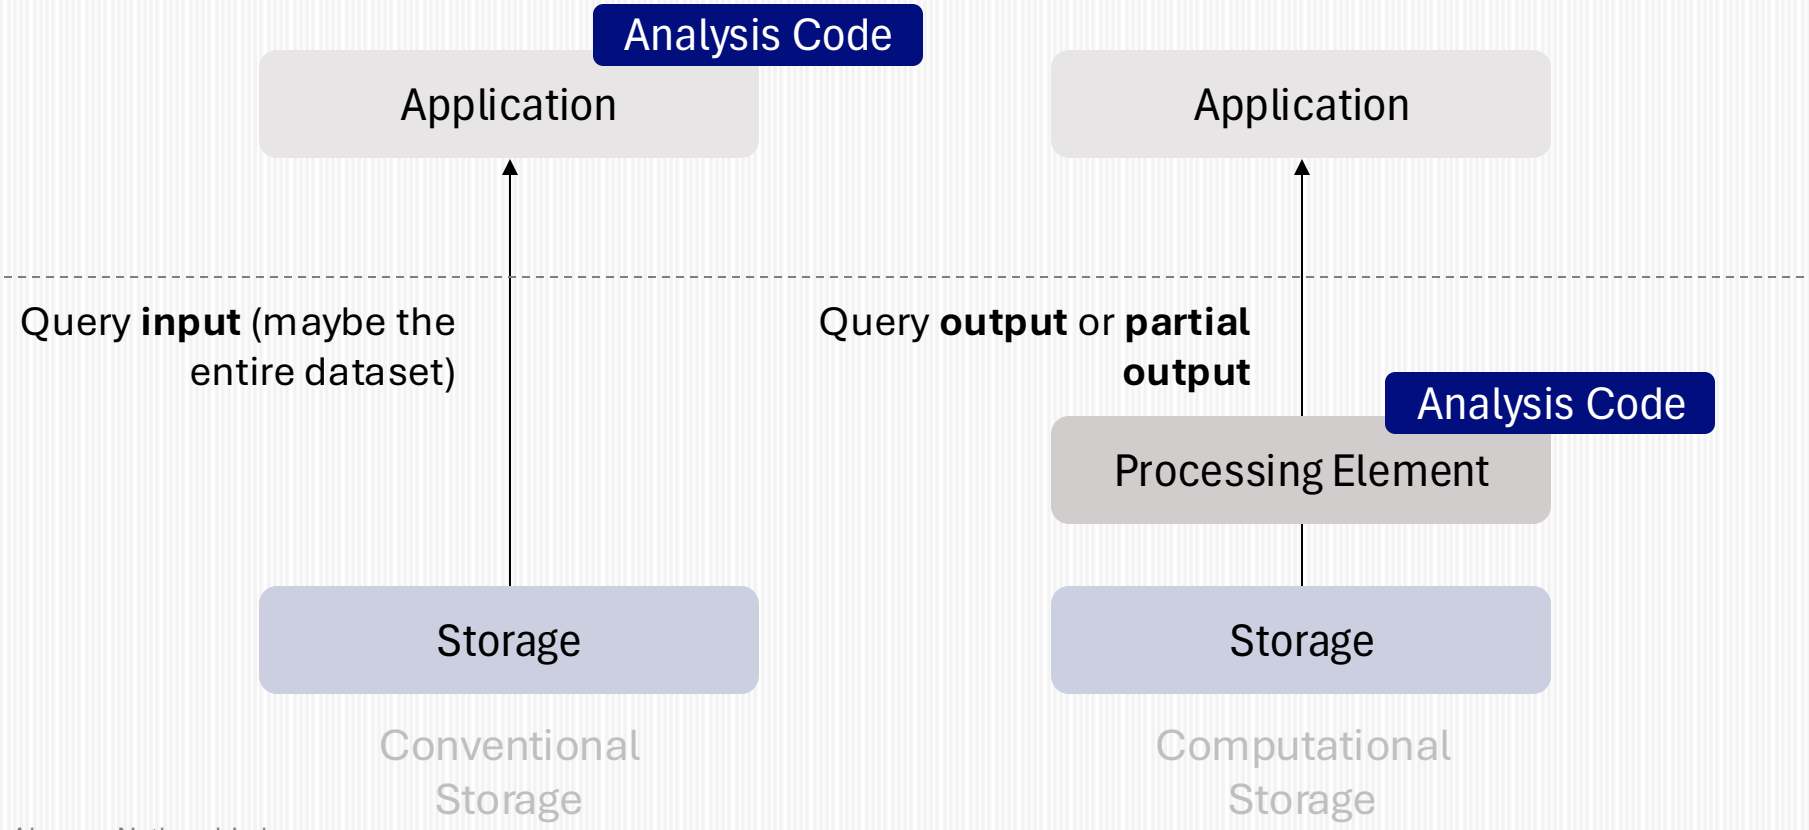

An Example

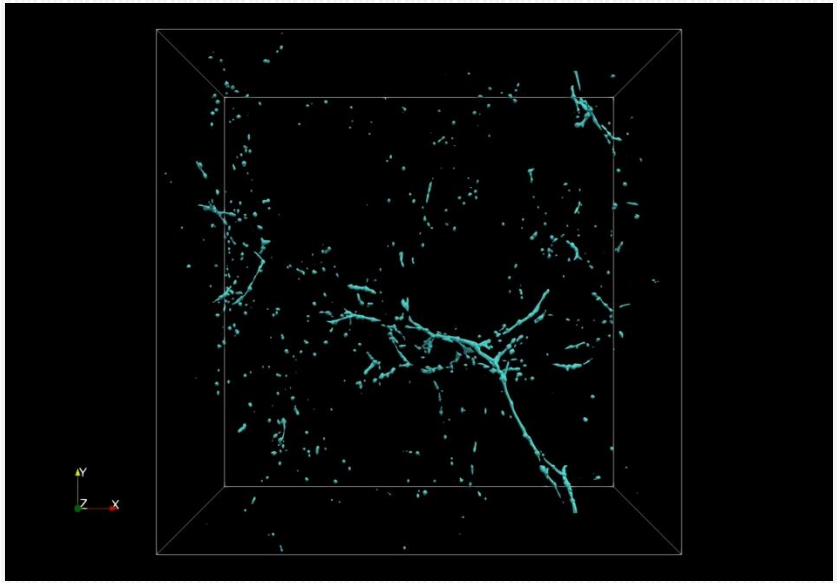

Data selectivity:
0.06%

A contour over baryon density for regions of candidate halos in a cosmological hydrodynamics simulation

A Multi-Year Journey

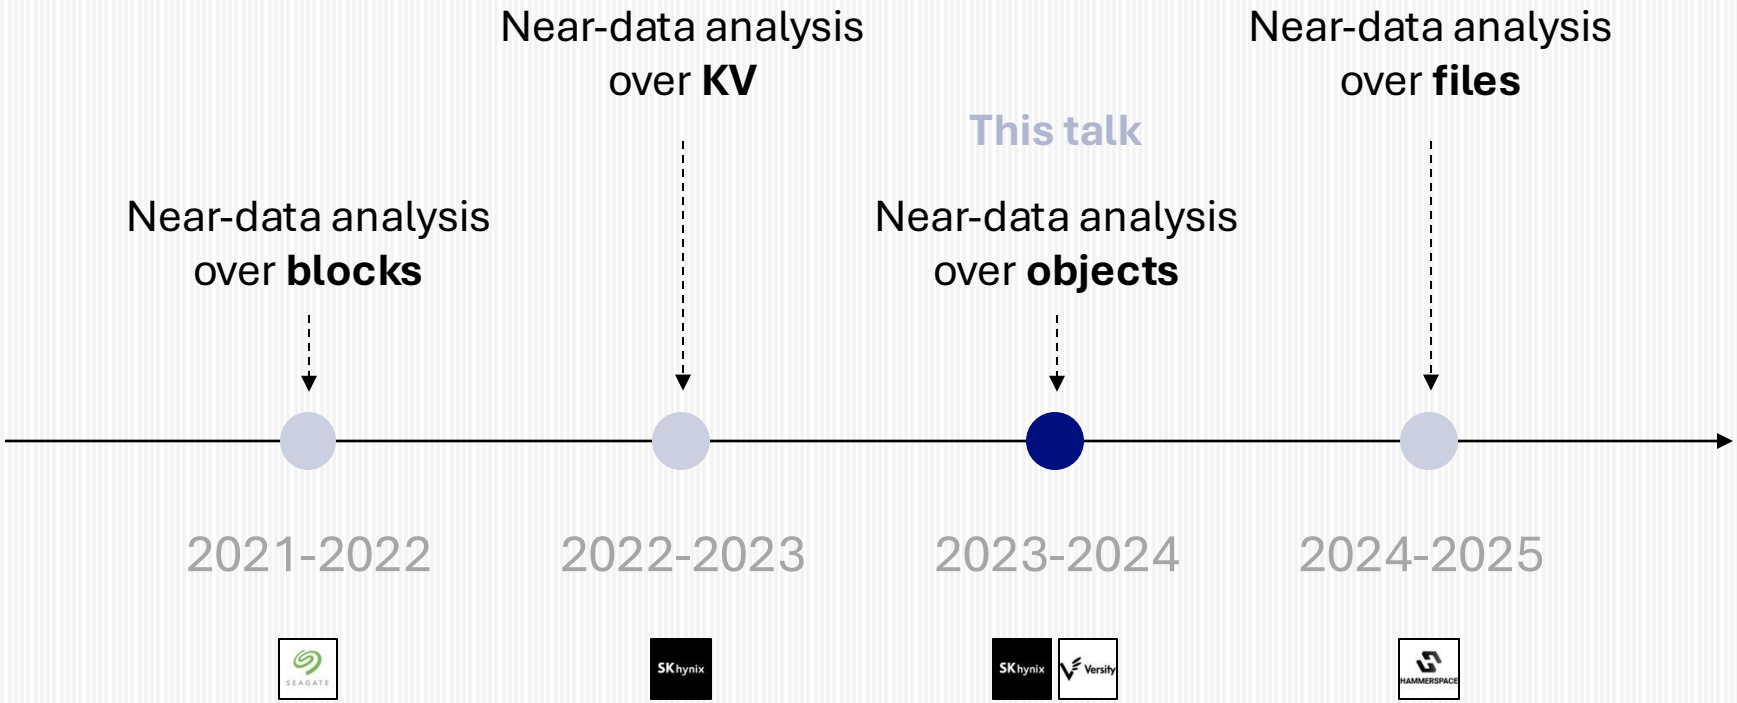

A Quick Look At It

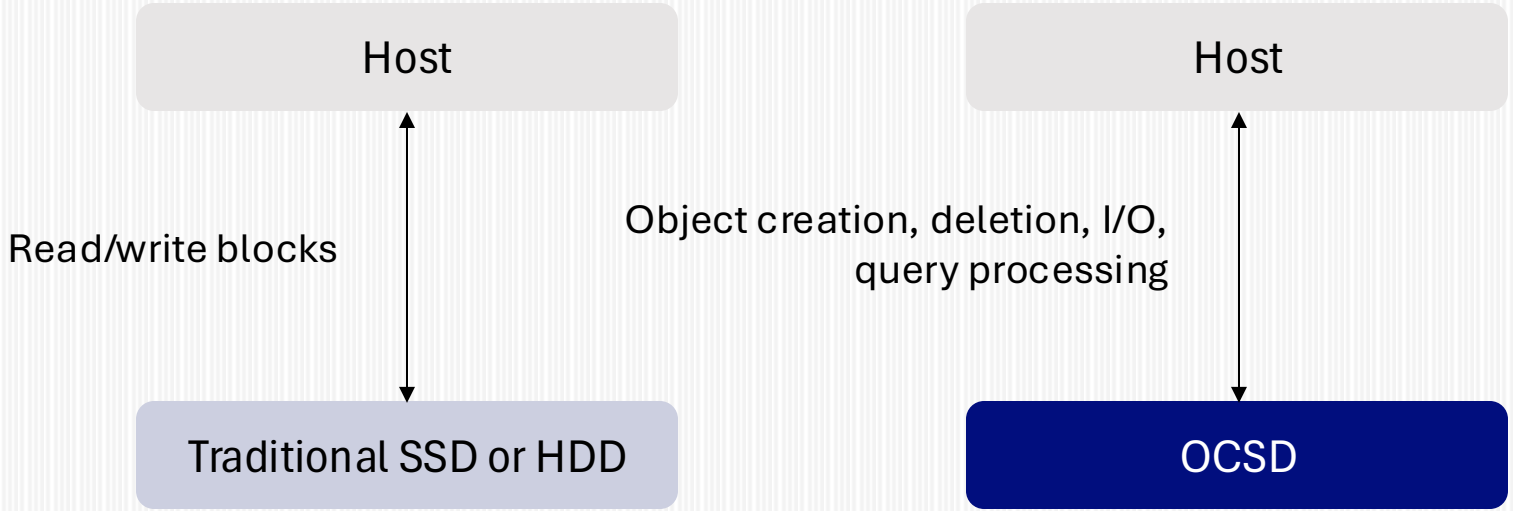

We introduce an Object-based Computational Storage Device (OCSD) concept

Device-Level Data Awareness

Analysis path (***file*** based)

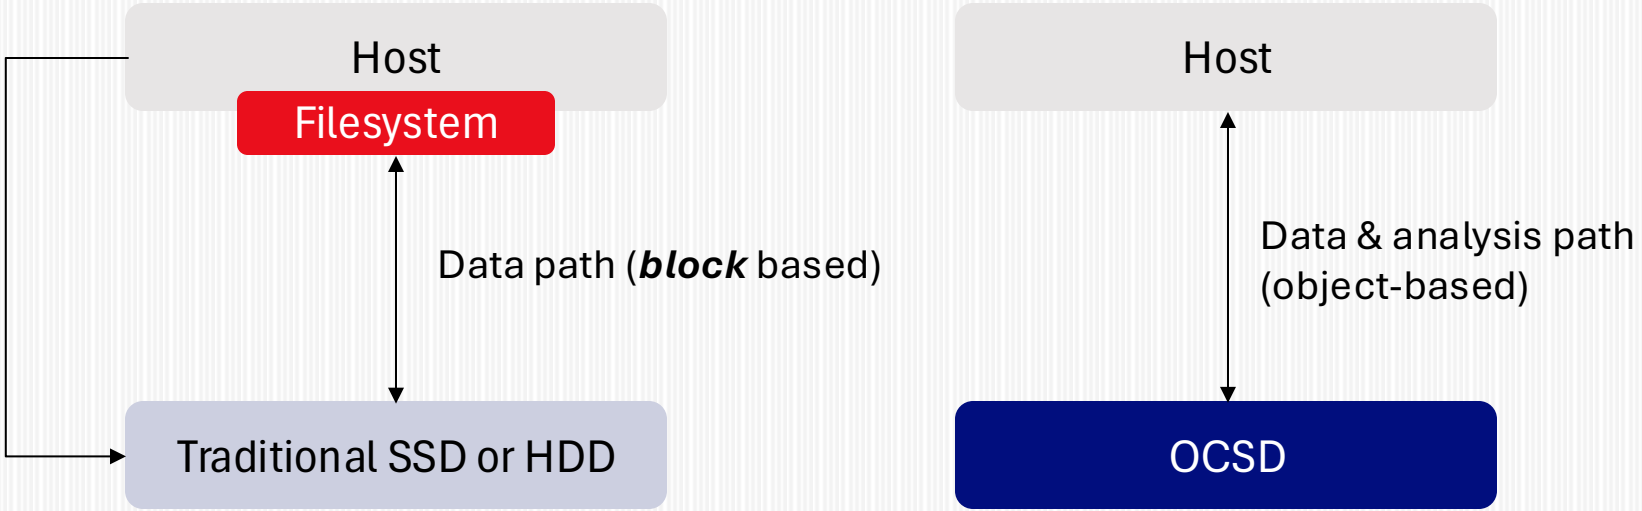

Being object-based enables devices to understand and process data natively,
without requiring a file-to-block translation

Becoming Part of a Big Community

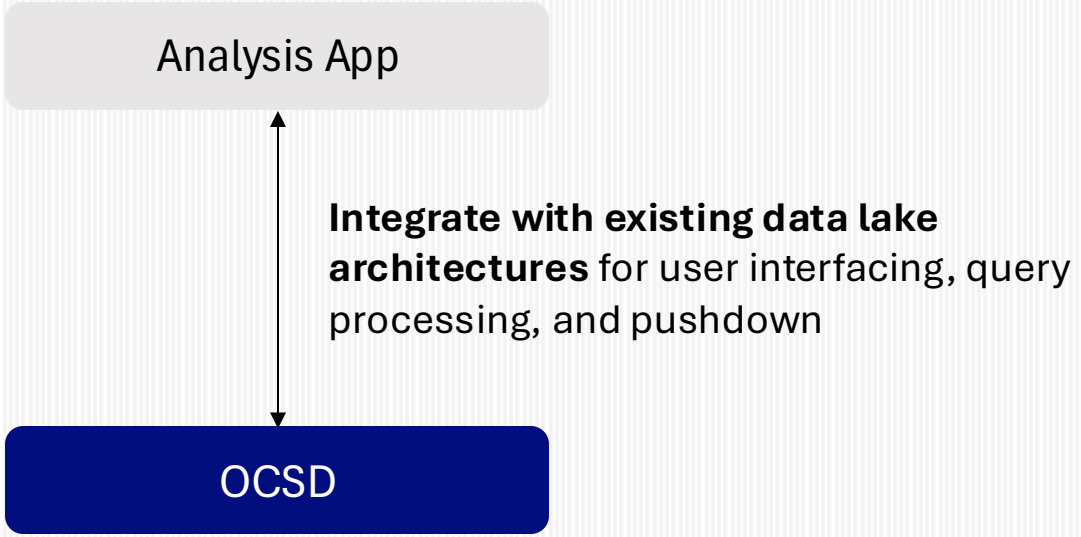

Using off-the-shelf software avoids reinventing the wheel and allows us to build something more powerful

Embracing the State-of-the-Art

❶ **The client** sends queries as SQL statements

❷ **An open SQL engine** (e.g., ***Presto***) for distributed query planning, scheduling, and result aggregation

❸ **An open catalog** (e.g., ***Apache Hive metastore***) for table metadata management

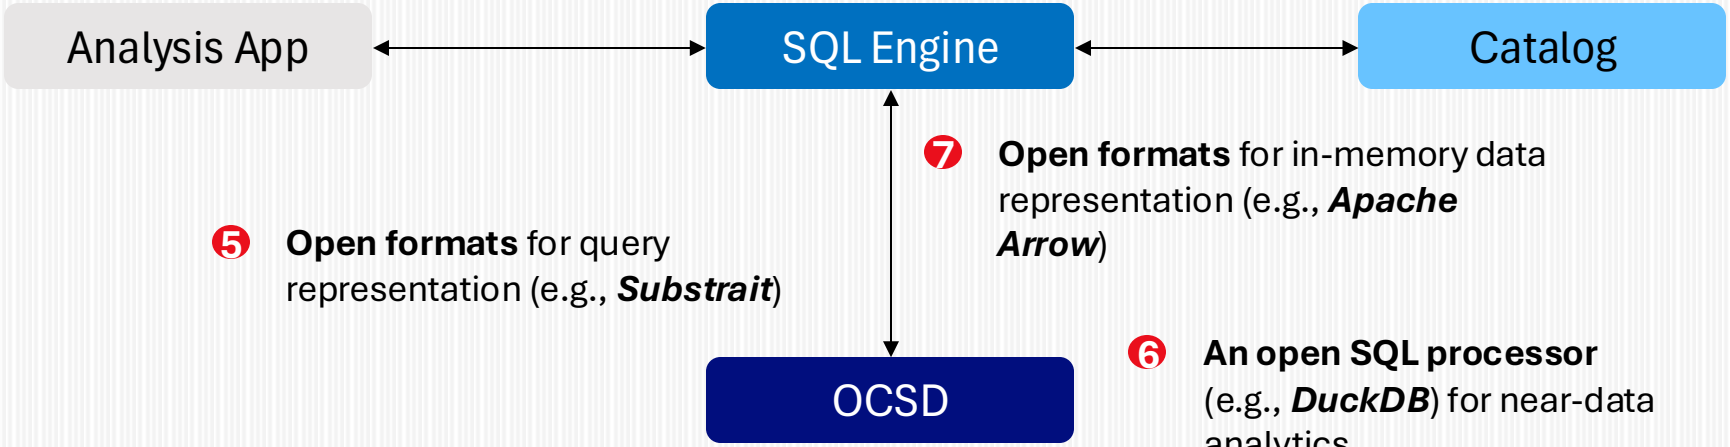

❹ **OCSD** acts both as a relational data source and as a query processor

❺ **Open formats** for query representation (e.g., ***Substrait***)

❻ **An open SQL processor** (e.g., ***DuckDB***) for near-data analytics

❼ **Open formats** for in-memory data representation (e.g., ***Apache Arrow***)

FMS'24 and SC'24 Co-Demonstration with SK

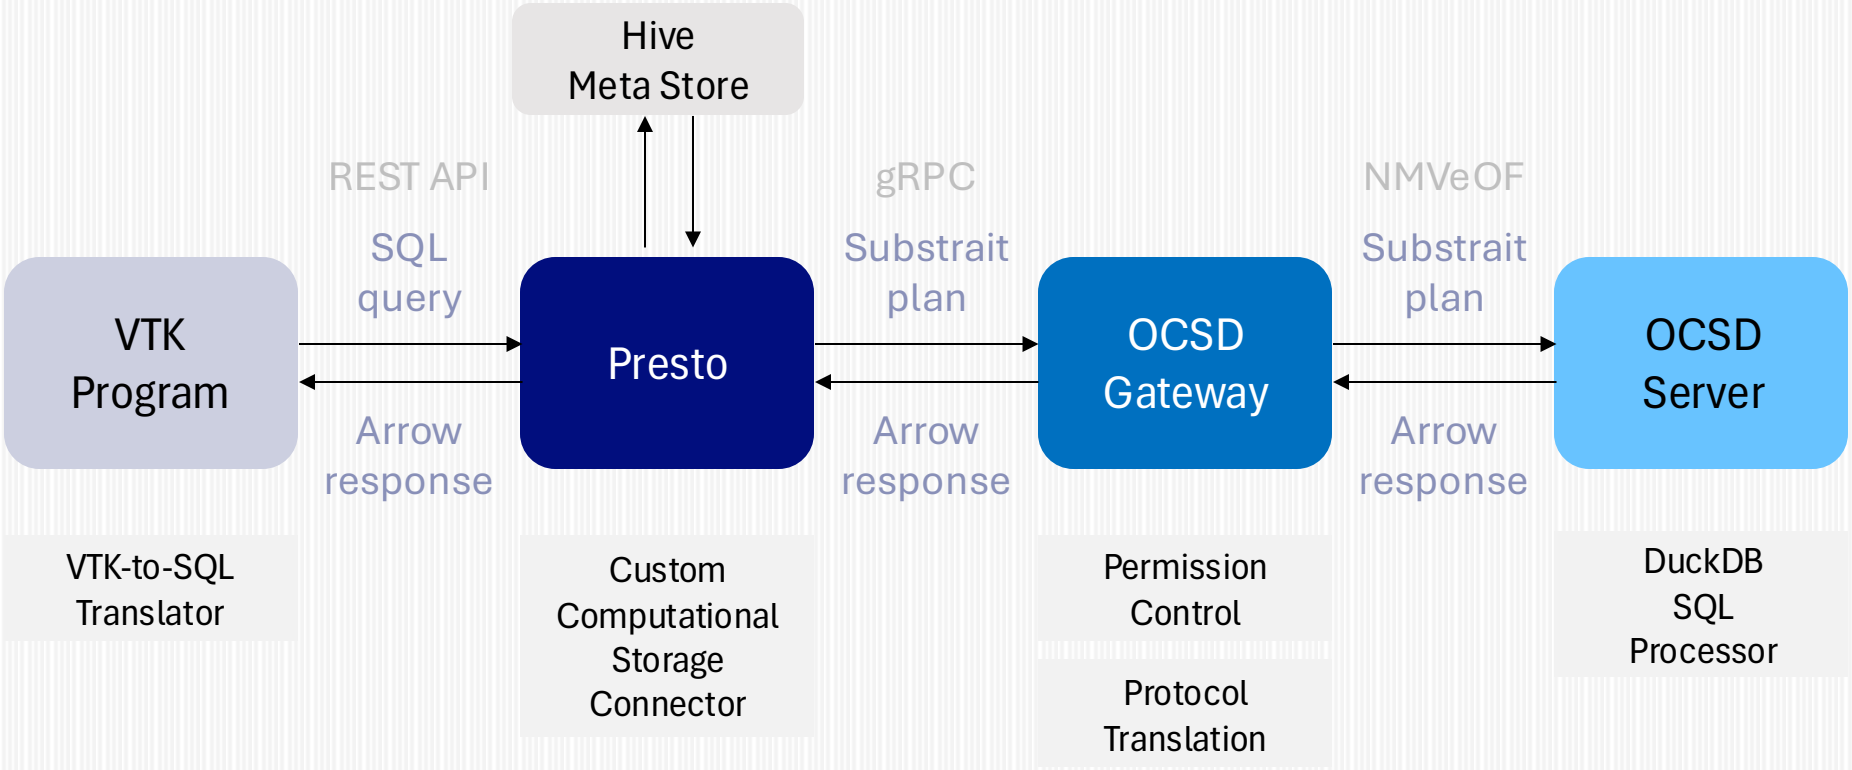

End-to-End Results

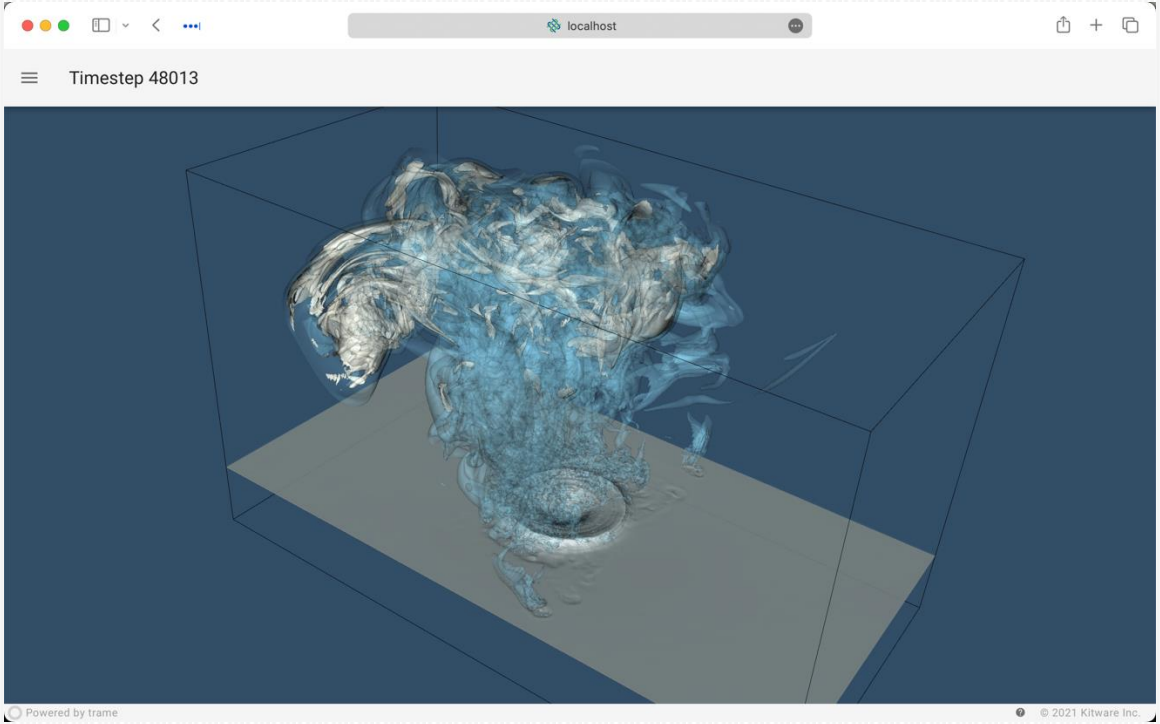

Future Work

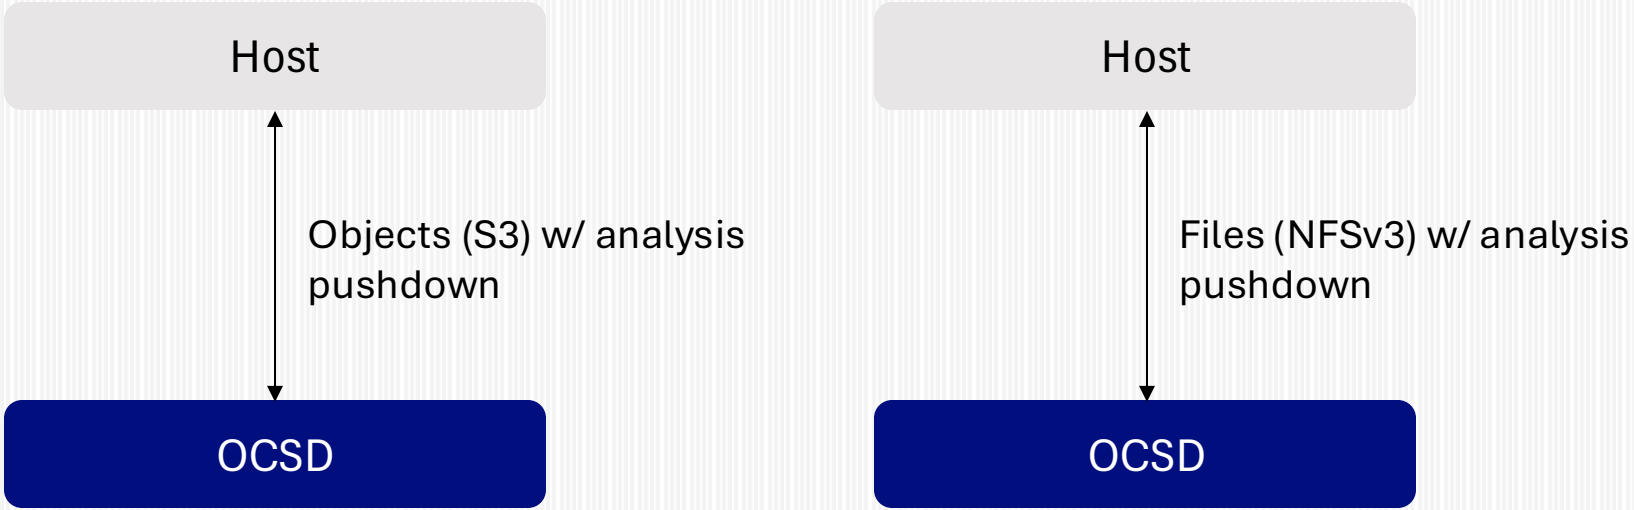

With features from Hammerspace for secure file layout lookup and secure direct local data retrieval

Conclusion

We will continue exploring near-data analysis, expecting deployment in next next generation platforms

Having the right abstraction matters (from block to objects to files)

Leveraging and integrating with community tools as HPC converges with broader industry trends

Los Alamos

NATIONAL LABORATORY