

Vertical query optimization using Data-aware CSD for data analytics


PRESENTER
Jungki Noh




Motivation

- Challenges of Data Analytics System
 - Data analytics in HPC often read more data from storage-nodes to compute-nodes than necessary
 - As the data volume grows rapidly each year, the amount of data being moved also increase steadily
 - Large-scale data movement exacerbates the “Computing” and “I/O” overhead of compute-nodes
 - The unnecessary data movement for analytics poses a challenge to performance improvement

End-to-End Execution Time Breakdown in Google



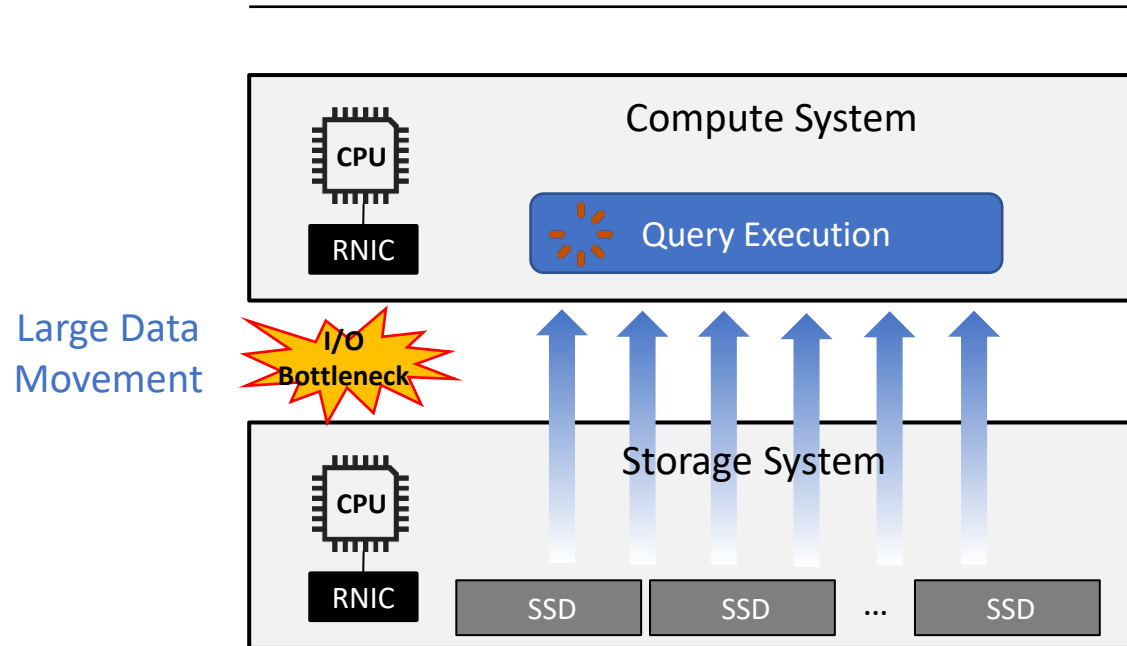
Rapidly Growing Data Volume (2010~2025)



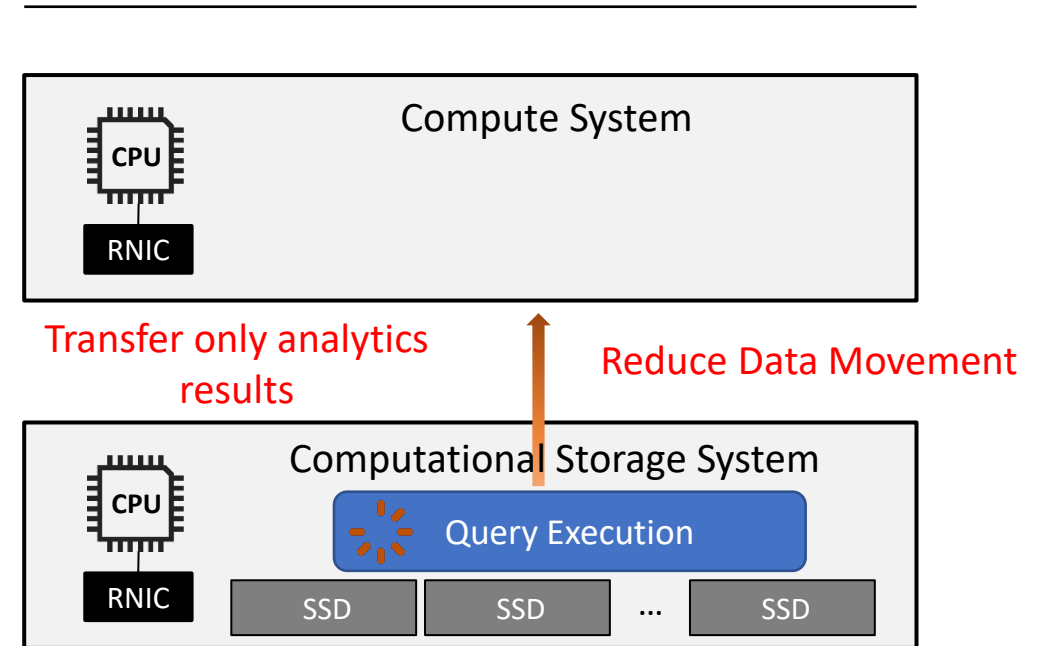
Our Approach : Computational Storage

- Leverage computational storage to minimize data movement by executing queries directly near the data

Conventional Data Analytics



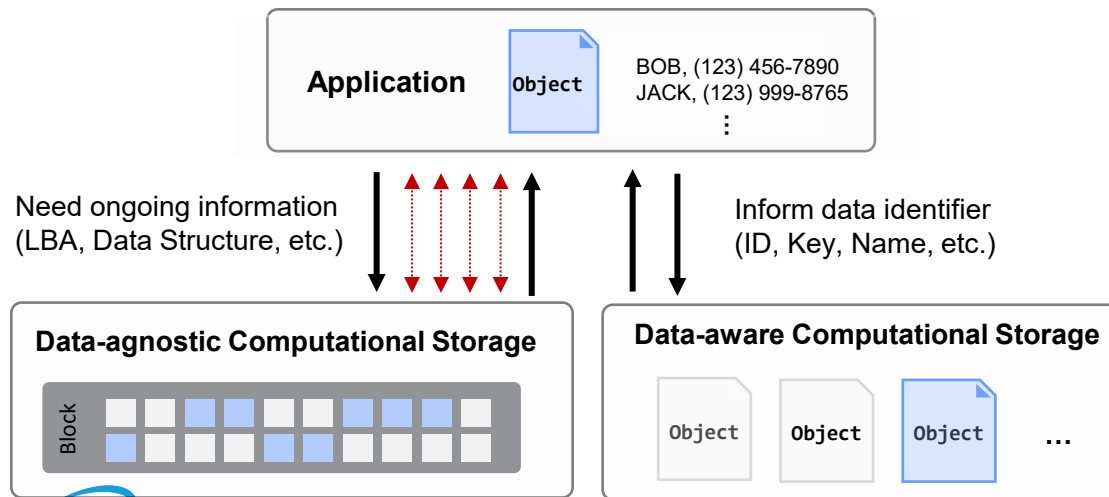
Data Analytics based Computational Storage



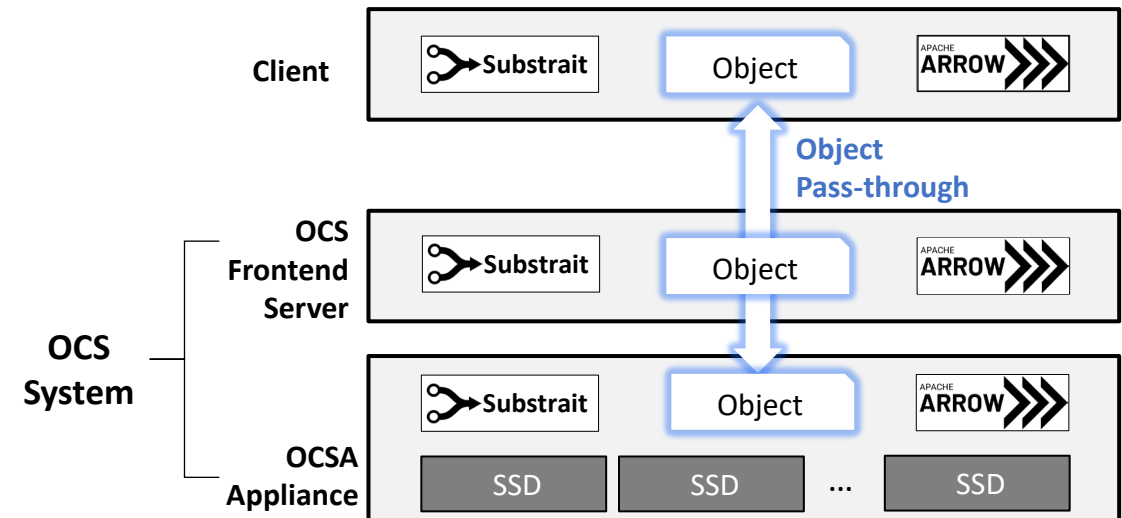
Our Solution : Data-aware Computational Storage

- Data-aware computational storage is optimized for both near-data processing and stored data analytics
- Our Data-aware computational storage operates on Object and provides consistent computational viewpoint everywhere
- We use Apache Arrow and Substrait to enable consistent data analytics processing everywhere

Data-agnostic vs Data-aware

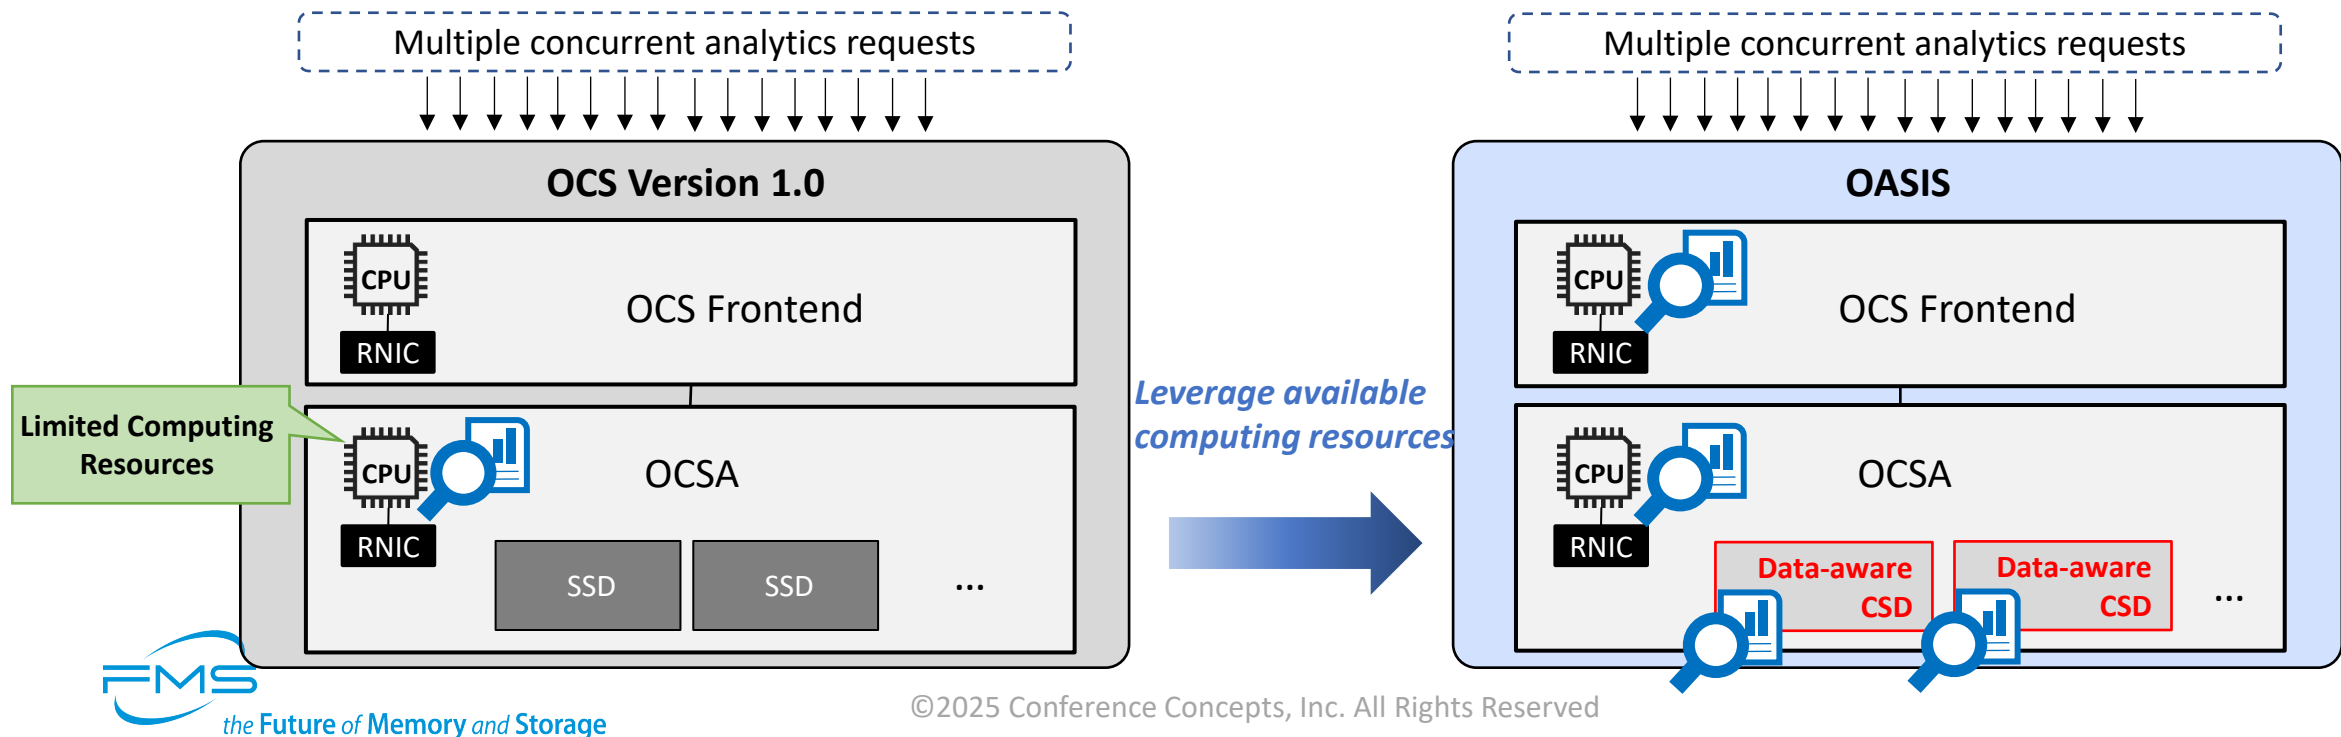


Object-based Computational Storage(OCS) System



Beyond OCS System

- OASIS: Object-based Analytics Storage for Intelligent SQL Query Offloading
 - Integrate Data-aware CSDs with built-in analytics capabilities into the OCS system
 - Optimizing data analytics by leveraging the available computing resources across the entire system



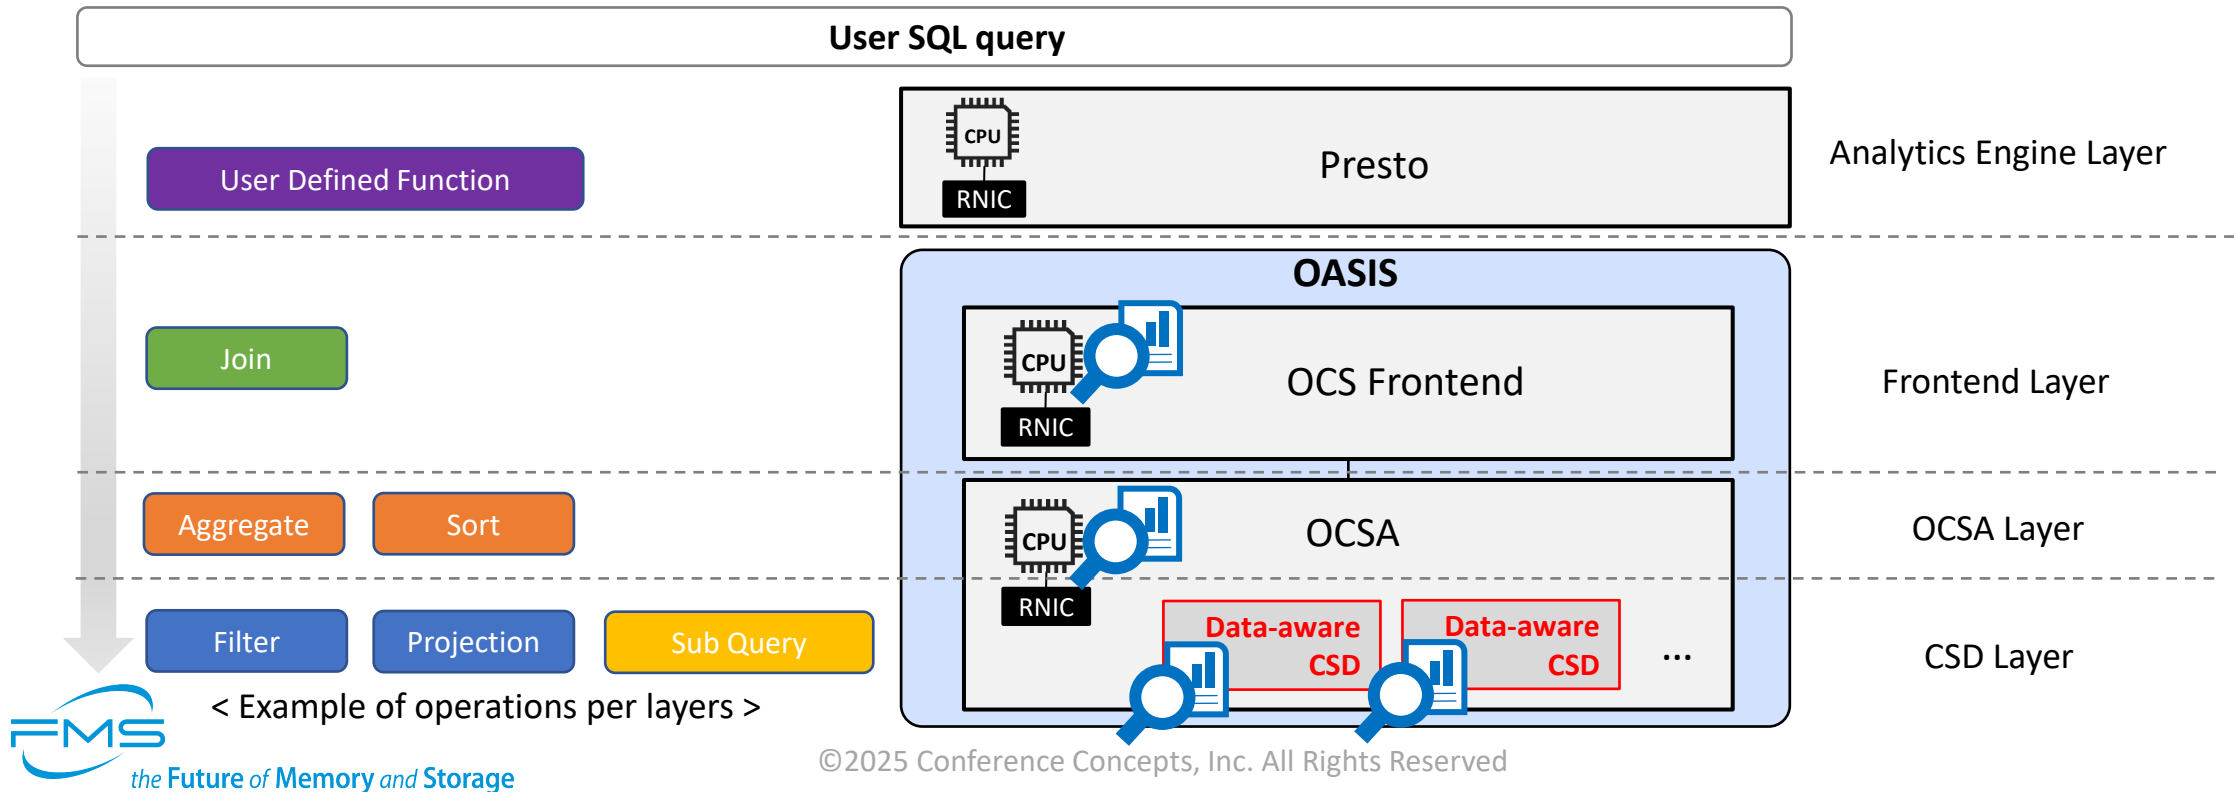
Key Feature of OASIS : Data-aware CSD

- Characteristics of Data Awareness
 - Enable data understanding and analysis by utilizing object-level I/O and management, independent of application assistance
- Analytics Capabilities
 - Support for the ability to interpret and analyze Substrait Query Plans, producing results in Arrow format
 - At present, data-aware CSD supports only the Parquet format



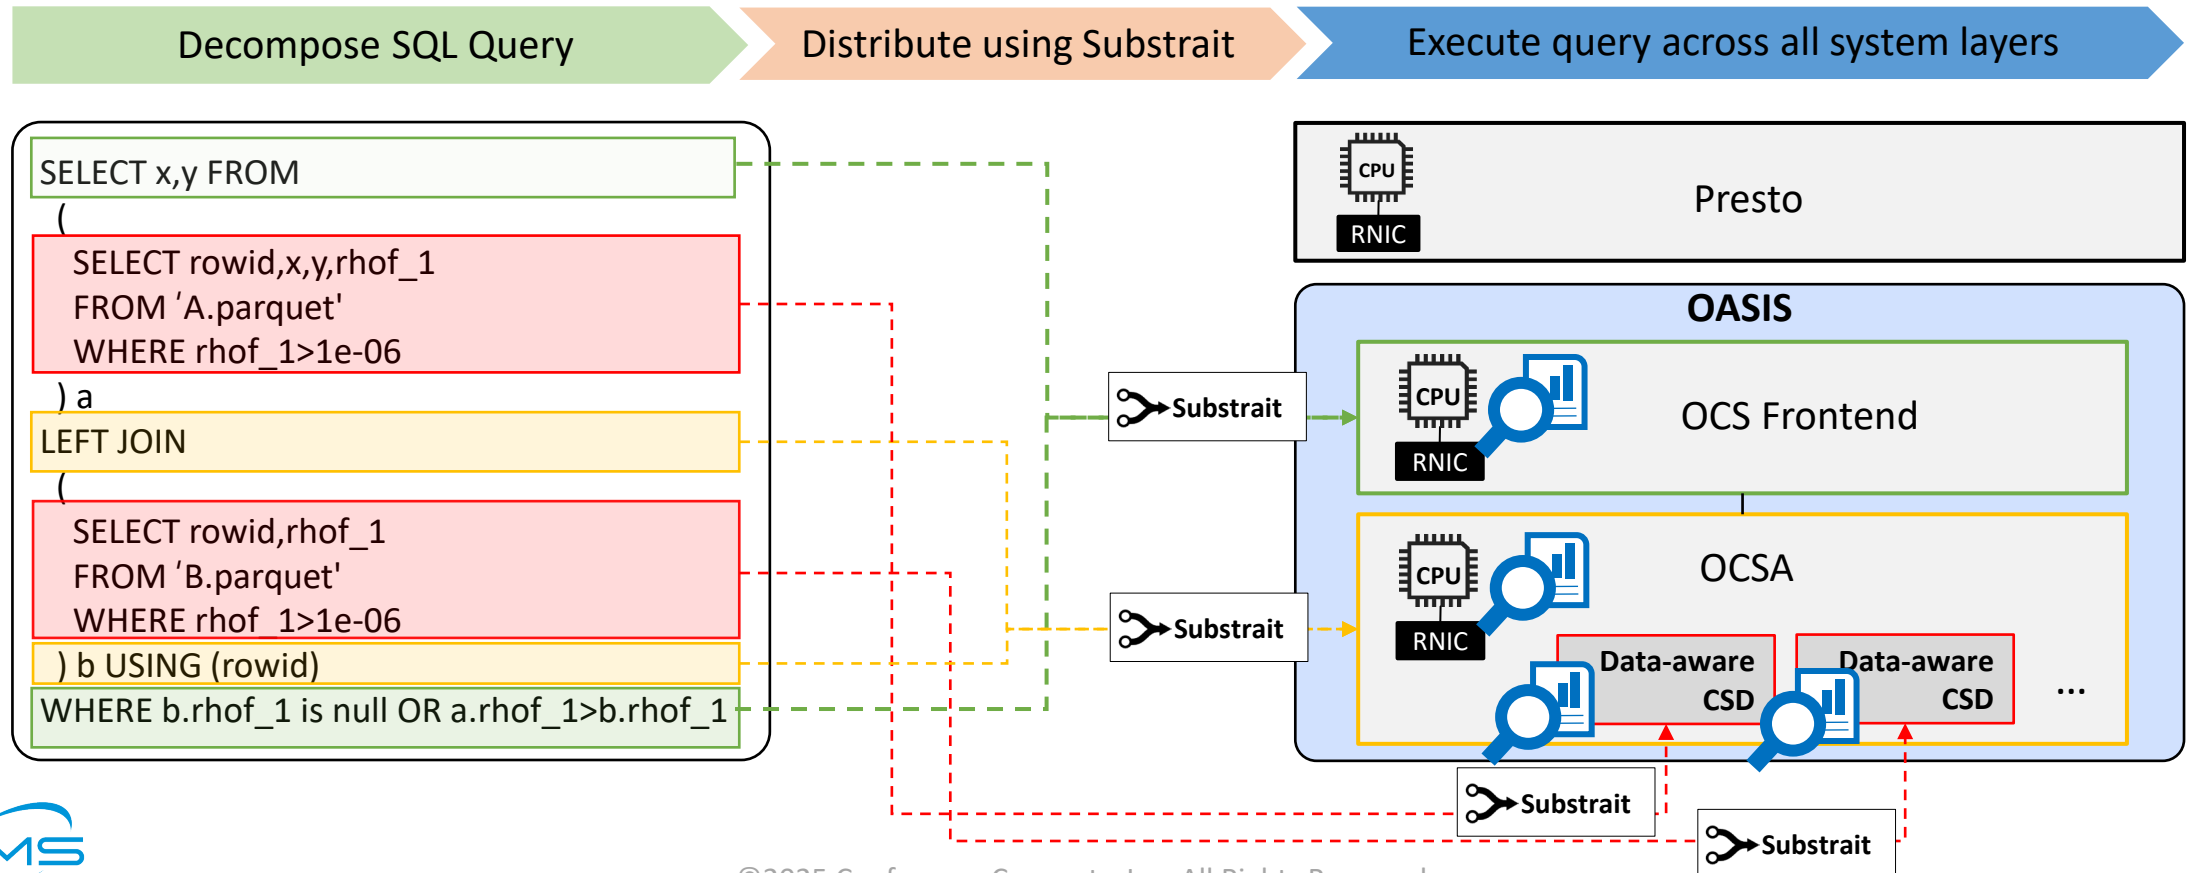
Key Feature of OASIS : Vertical Query Optimization

- User query is decomposed and executed across multiple system layers through vertical query optimization
- Our goal is to construct a parallel analytics pipeline that utilizes compute resources across all system layers



Example of Vertical Query Optimization

- An example of applying Vertical Query Optimization using the LANL HPC Workload




Evaluation Results


- Enhance data analytics performance about 4.1 x times
- Remove data movement up to 99.9 %

Vorticity-Driven Lateral Spread Wildfire


- Real-world scientific simulation data
- Simulation Data of Lateral Diffusion by Vortex (91 Parquet files, each 340 MB)



Average Query Execution Time



Average Data Movement per Query



Summary

- In HPC environments, data analytics often involve transferring more data from storage nodes to compute nodes than necessary, resulting in excessive data movement
- This inefficiency limits performance, and our solution is computational storage - Enhancing near-data processing by integrating Data-aware CSDs into the Object-based Computational Storage System(OASIS)
- To enable the broad deployment of computational storage requires standardization not only at the SSD level, but also between application and the storage software stack, as well as across compute nodes and storage nodes

Booth #207

Meet the future of memory.
Just steps from the entrance.

Innovation starts here,
Literally.

SK hynix

