OCP FUTURE TECHNOLOGIES SYMPOSIUM

OCP Global Summit

November 8, 2021 | San Jose, CA

Accelerating Data Analytics with Computational Storage Drives

Keith McKay, Ning Zheng, Yang Liu, Tong Zhang

What is a computational storage drive?

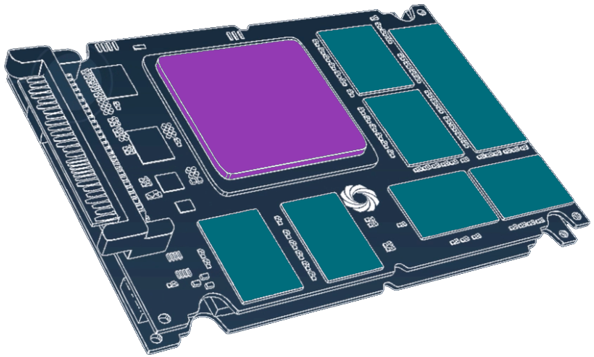

Enterprise PCIe SSD + Compute Offload Engines

Computational Storage Drive or “CSD”

Why use CSDs?

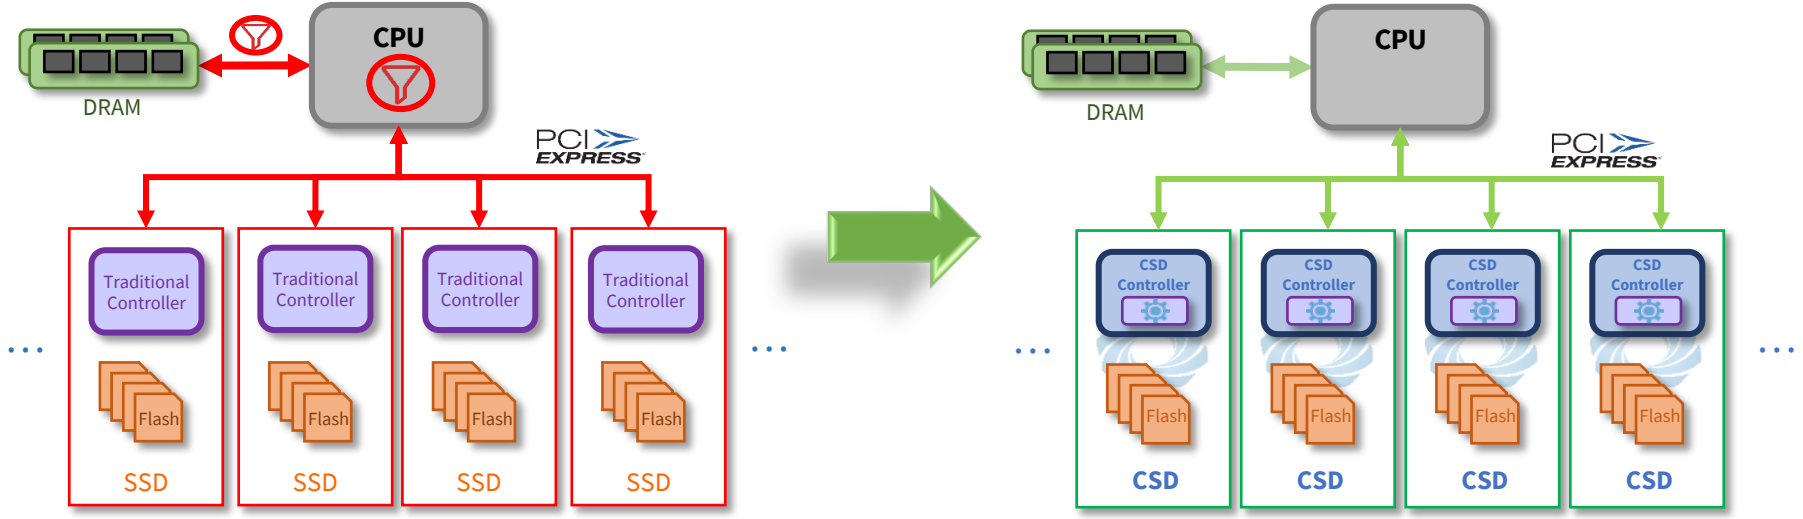

CPU Driven Architecture → Data Driven Architecture

Why use CSDs (Part II)?

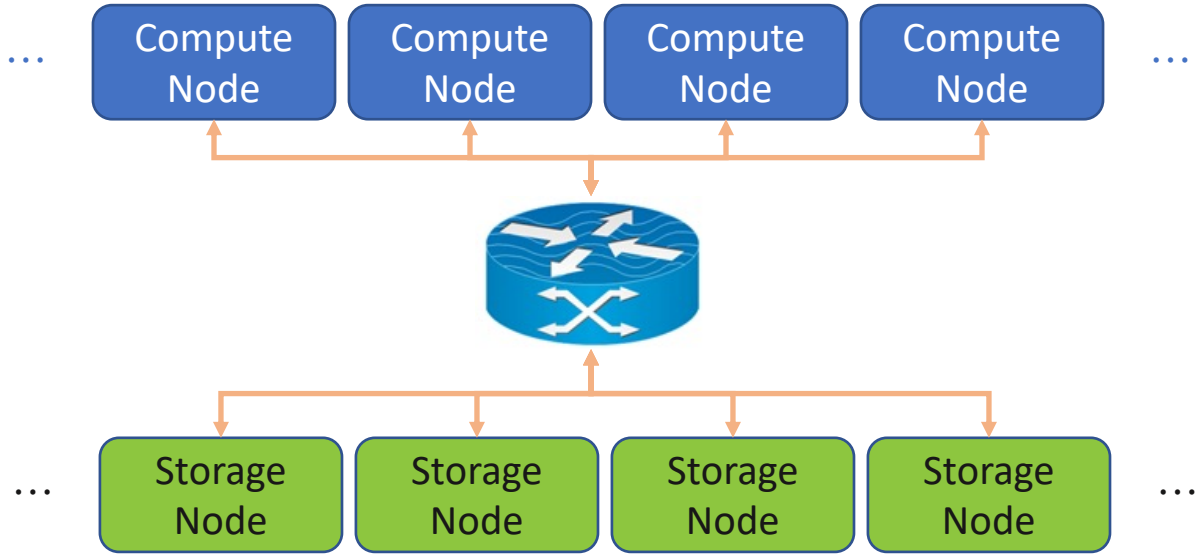

Compute-Storage Disaggregation → Data Movement Challenges

CSDs + Data Analytics

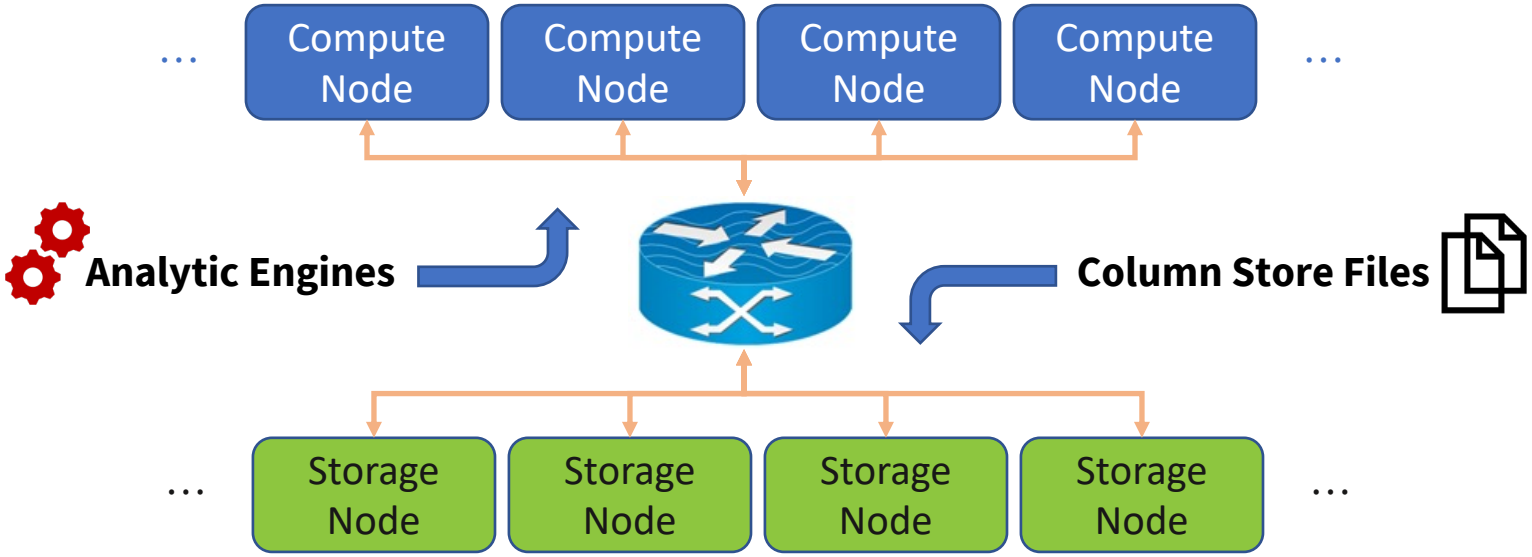

What about scan pushdown?

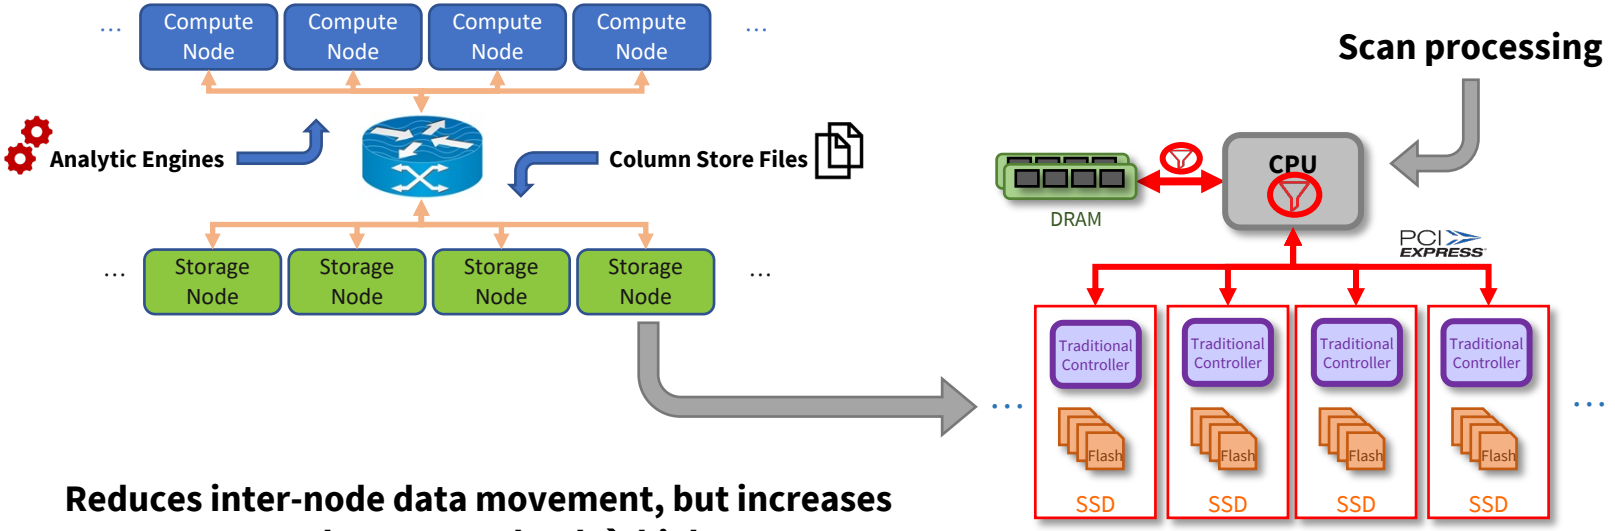

Reduces inter-node data movement, but increases storage node compute load → higher cost

What about scan pushdown with CSDs?

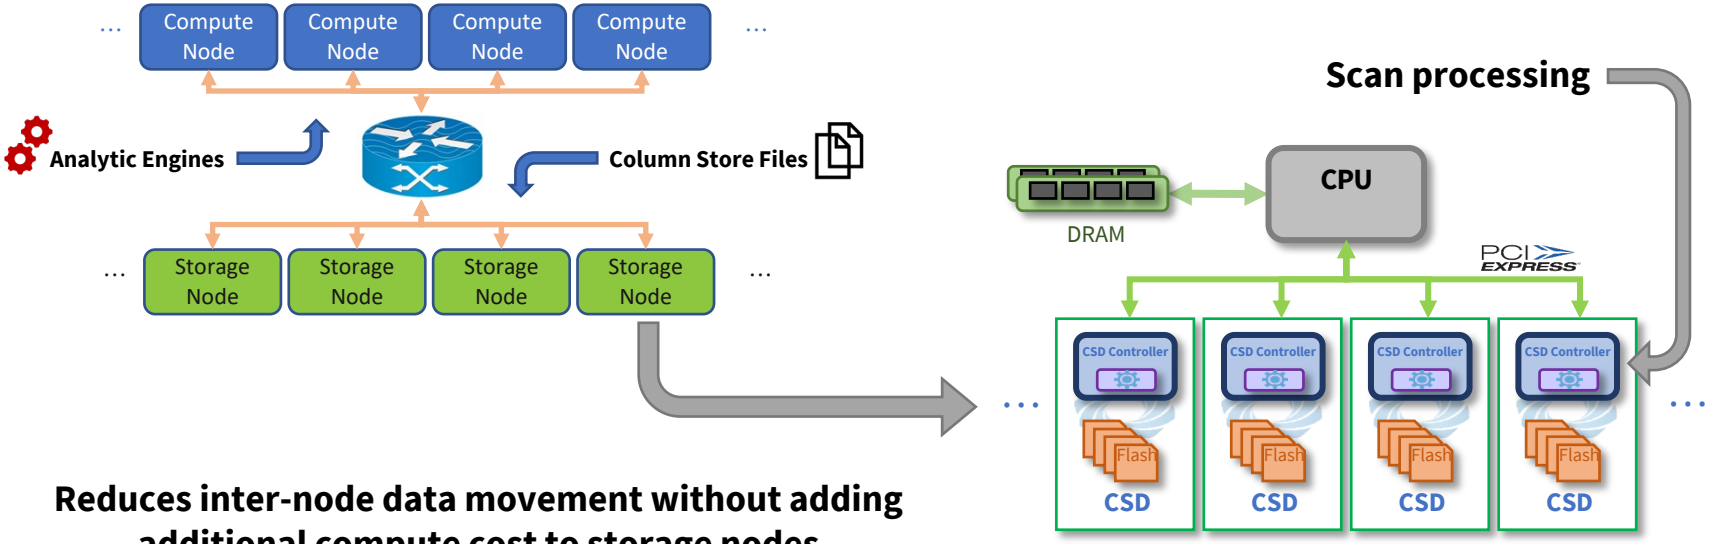

Reduces inter-node data movement without adding additional compute cost to storage nodes

But there is a challenge…

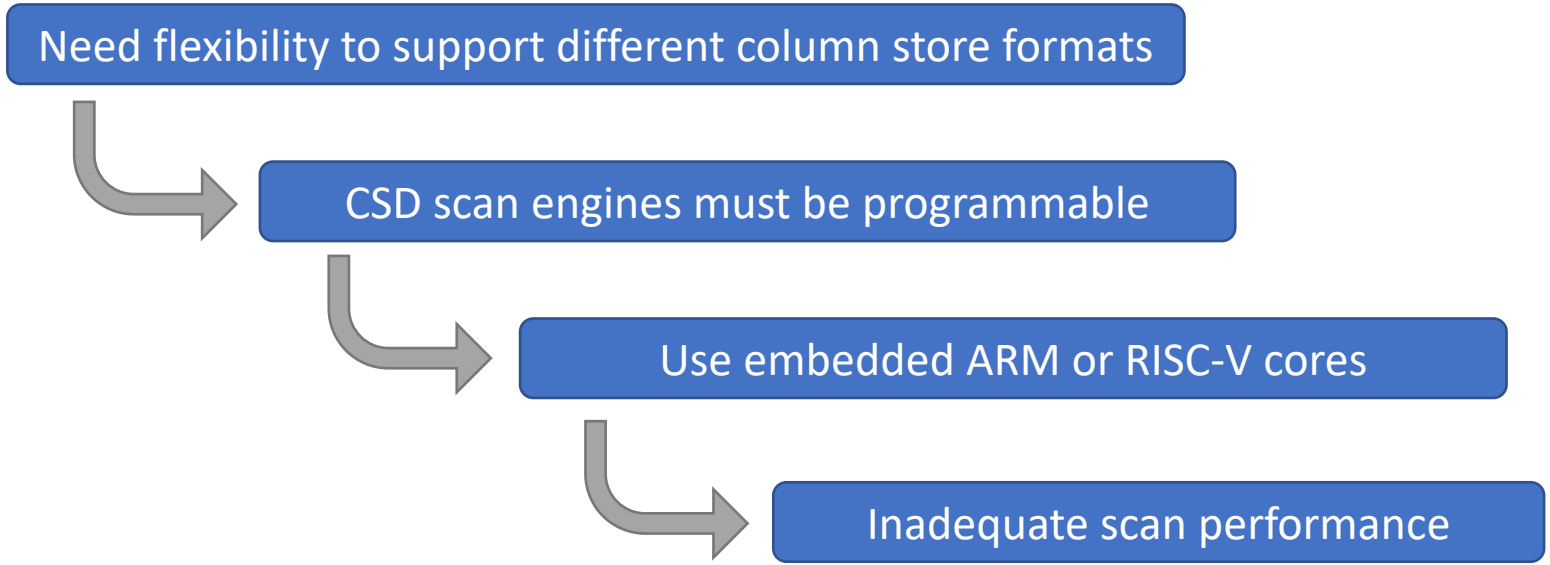

… and a solution

Need prevent embedded cores from becoming a bottleneck

- **Only Option**: Alleviate embedded cores from touching every single data item during scan
  - Enhance the column-store file format to enable data skipping during scan
    - Embed fine-grain metadata into column-store file to facilitate data skipping

Data skipping with fine-grain metadata

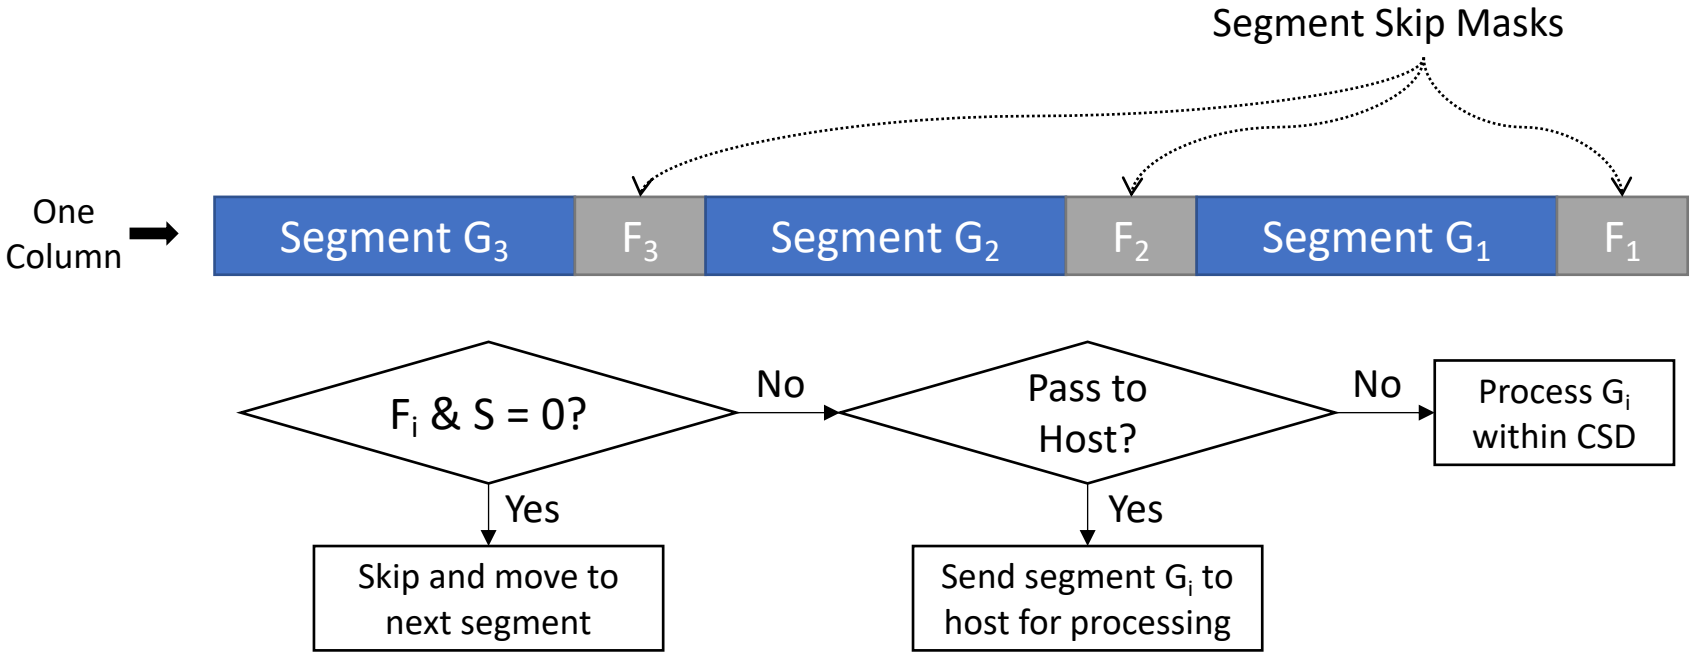

Test Results

- 8x ARM Cortex A53 @ 1.4GHz
- Dataset: TPC-H lineitem Table
- Analytics Engine: ClickHouse
- Assumptions:
  - Pipelined data fetching, decompression, and scan inside CSD.
  - Pipelined intra-CSD processing and host processing.

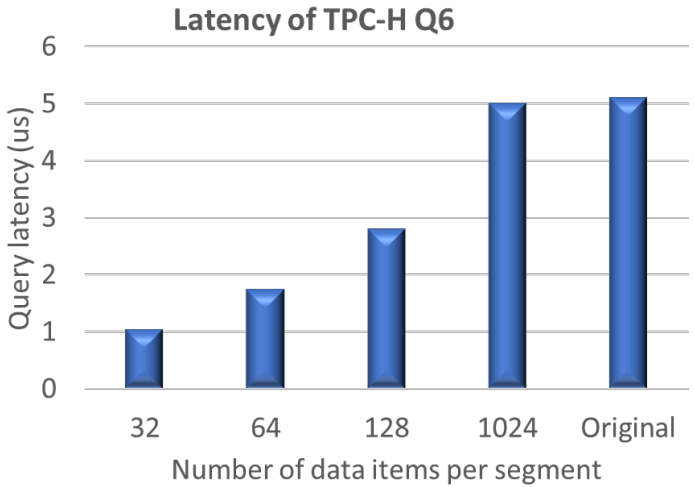

Summary:

- CSD off-load with added metadata reduced query latency 5x compared to host processing alone
- Achieved scan throughput >1GB/s per ARM core
- Granularity of metadata is a key performance tuning factor


OCP
GLOBAL
TOTAL
SOLUTIONS

OCP FUTURE TECHNOLOGIES SYMPOSIUM

2021 OCP Global Summit | November 8, 2021, San Jose, CA