


Compute Made For Analytics

Setting a new standard for query
processing speed and scale

Krishna Maheshwari
Chief Product Officer


Data Processing and I/O speed



Sources: Intel/Samsung and TPC-H data (at 10 TB)

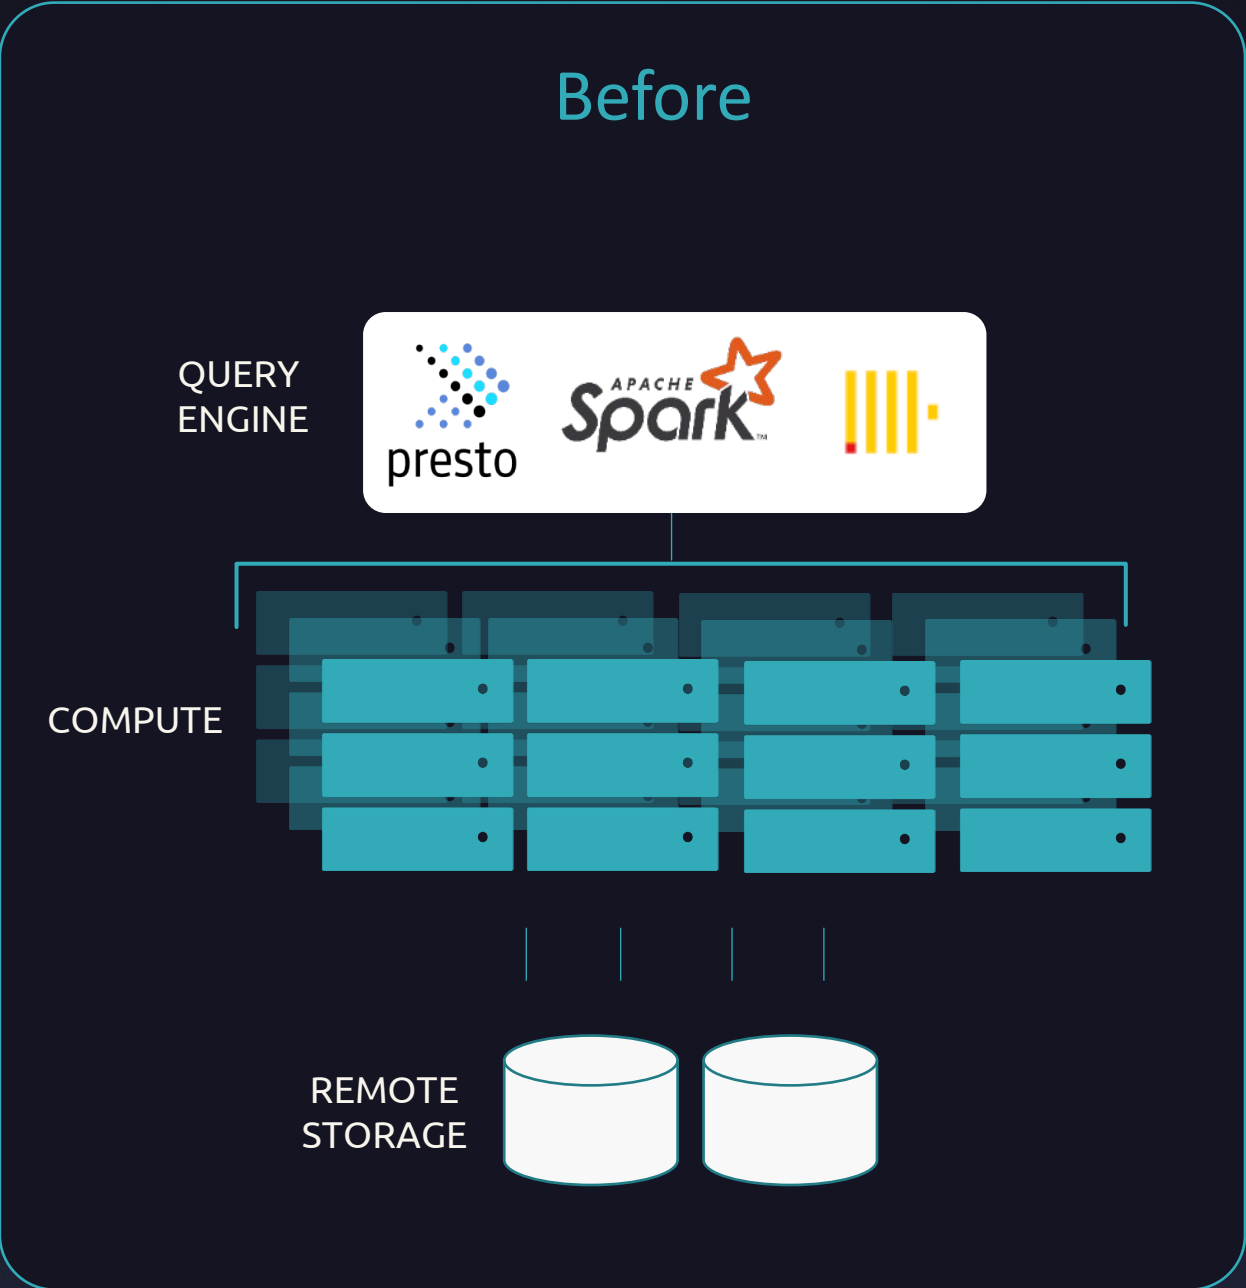


Accelerated Compute, unparalleled performance, reduced footprint, and improved energy efficiency



Case studies at top internet companies

Use cases: A/B testing & Data Science Pipeline

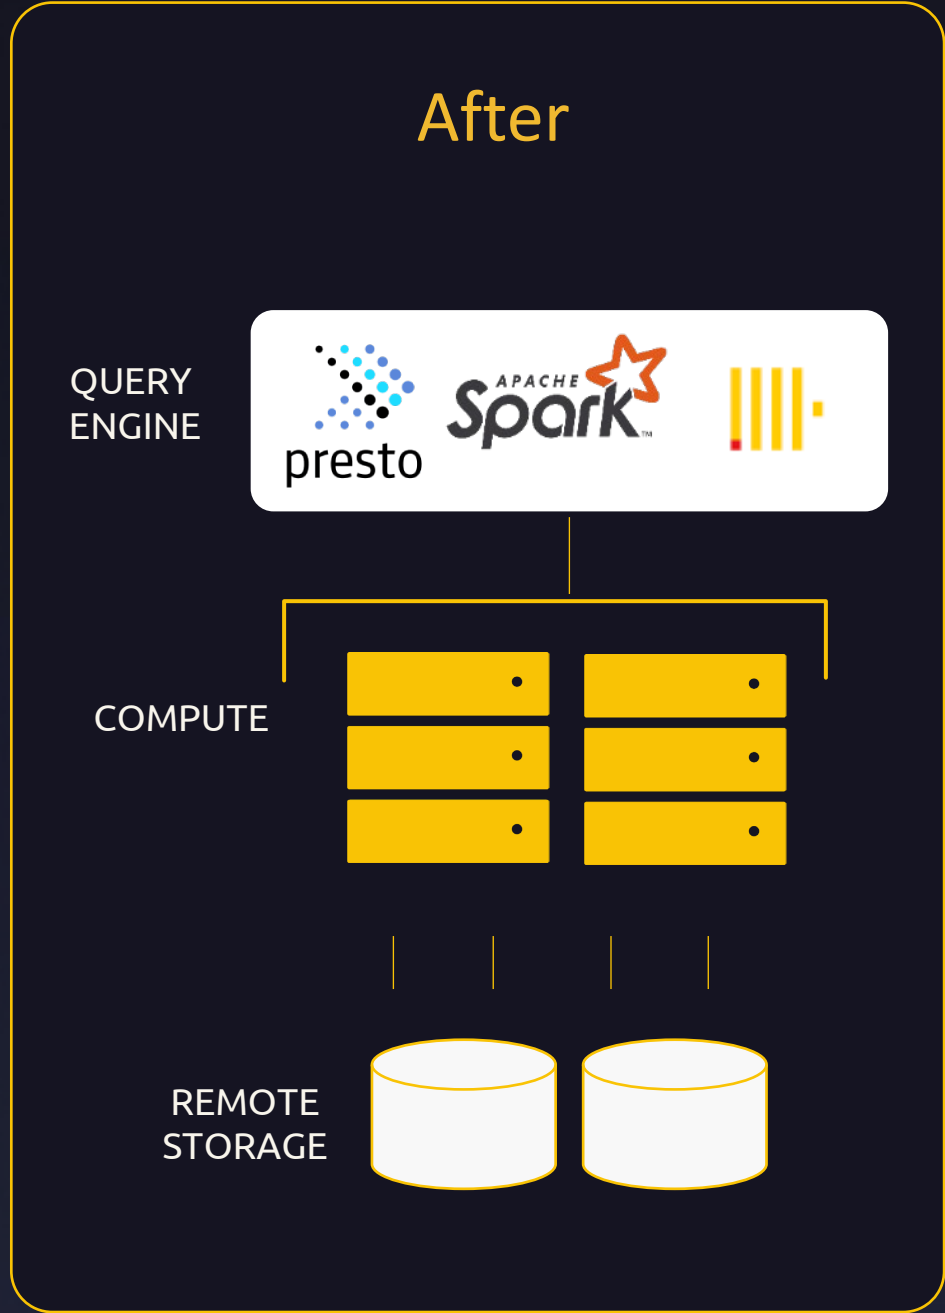


✓ 80% SSD IO Utilization

✓ 21x Data Scan Rate (QDSR)

✓ 15x Faster performance

✓ 70% Reduced TCO



Data Analytics Demo on NeuroBlade

Speed Up (Spark Duration / HEQS Duration)

8.13

Data Reduction (Result Size / Scanned Size)

97%

▼ KIOXIA / NeuroBlade Analytics Demo

KIOXIA CM7 PCIe 5.0 NVMe SSD



Software Only Query Duration (seconds)

123.12

Accelerated Query Duration (seconds)

15.14

Scan Rate (GBps)



Scan Data Size (GB)



Result Data Size (GB)




Scanned Size (GB)

2400

Results Size (GB)

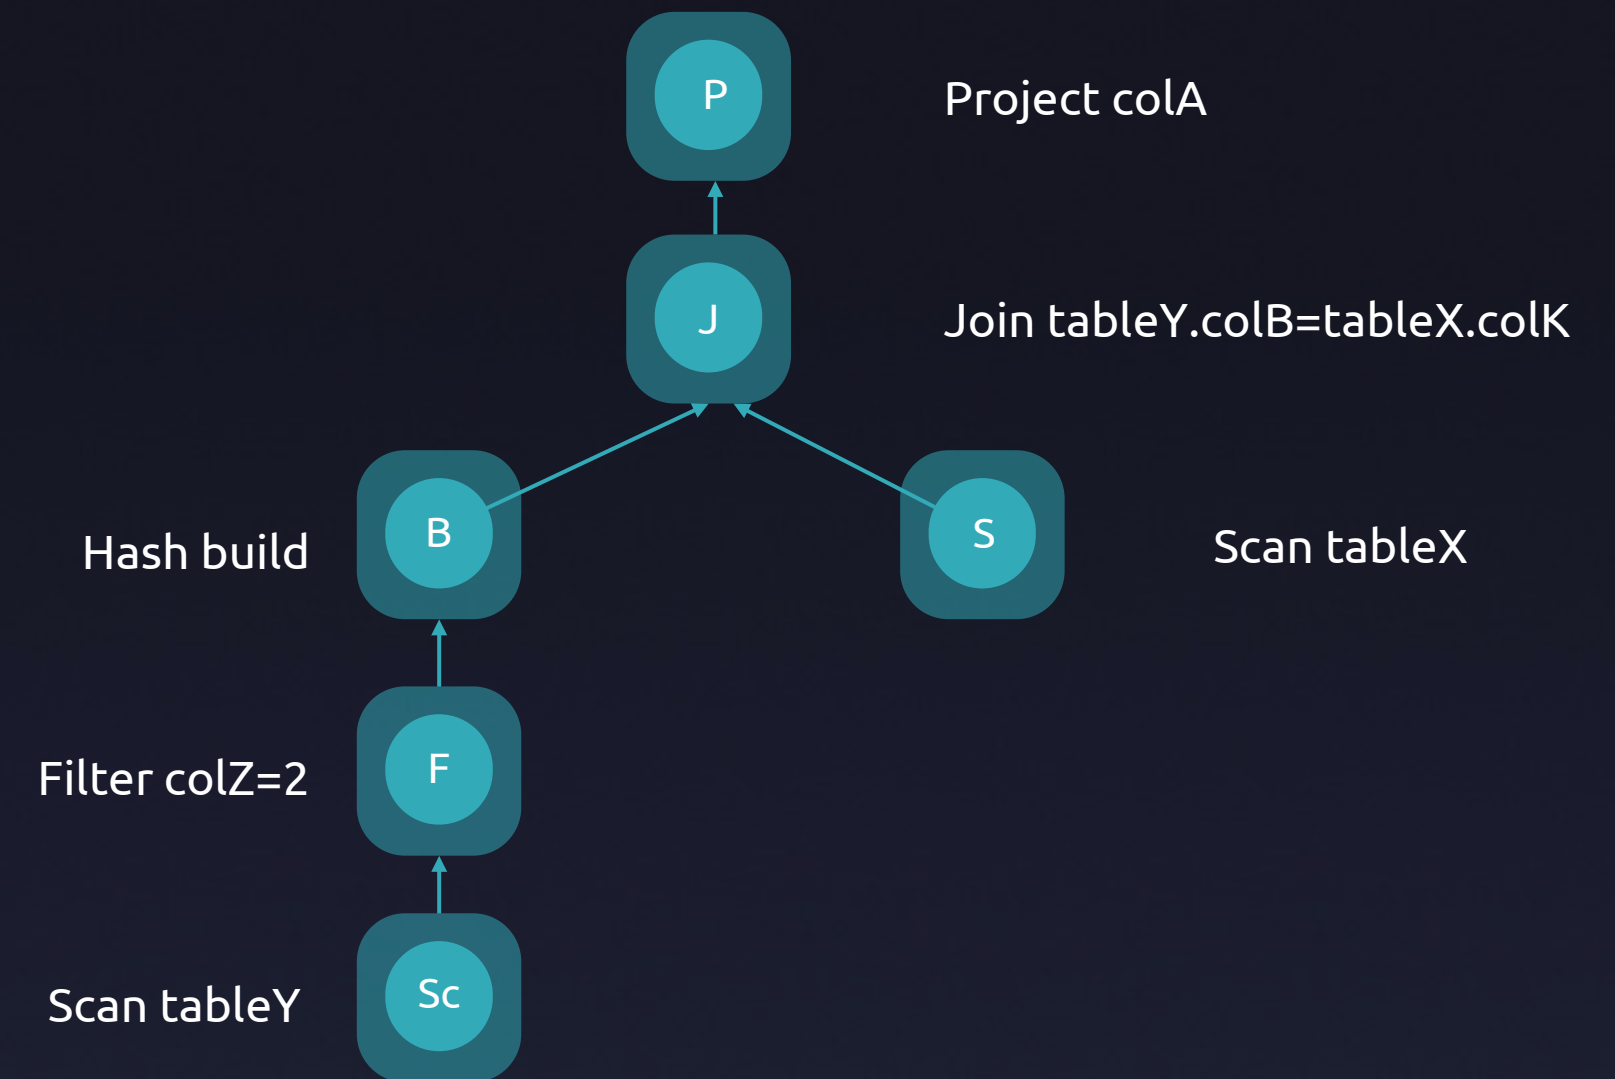
72

NeuroBlade Hardware Enhanced Query System (HEQS)



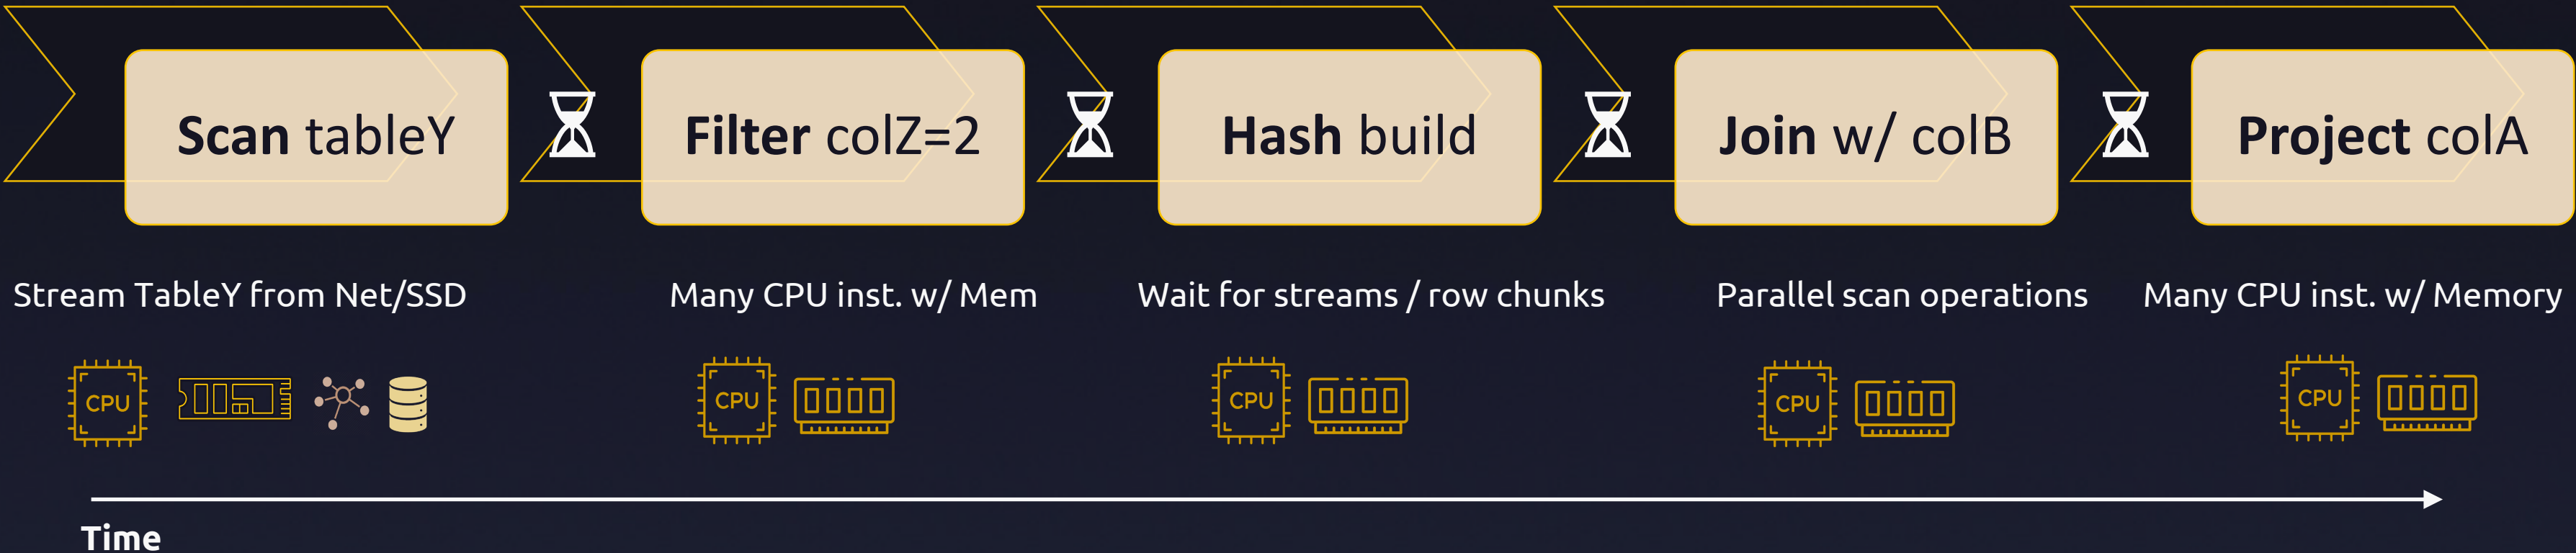
SQL physical plan architecture

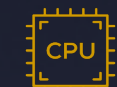





```
SELECT
    colA
FROM
    tableY
WHERE
    colZ=2
JOIN tableX WHERE tableY.colB = tableX.colK
```



Modern CPU architecture provides limited parallelism

Sub-operation parallelism, but *most operations are sequential*



-  Many CPU instructions (= 1 operator) 
-  Network traffic creates variability 
-  Paging creates addition latency 

HW architecture that allows processing directly from drives



SPU processes SQL operation at 32 GB/s



No CPU bottleneck



Direct HW support for Parquet format



Zero copy in-Memory



Increasing density 40U compute → 6U w/ NeuroBlade



**NeuroBlade provides a holistic
approach to acceleration**

**Come to Kioxia Booth #307 to
learn more**



Thank You