

How Close Are We to GPU Data Processing?

Toward Disaggregated Data Systems in the AI Hardware Era

Jigao Luo¹

2026

¹Systems Group, Technische Universität Darmstadt





TECHNISCHE
UNIVERSITÄT
DARMSTADT

SYSTEMS

Intro: Status Quo in Research & Industry

Hardware to reduce the tax

- Hardware evolving very quickly
 - New CPU architectures
 - New ISAs
 - New interconnects
 - **New accelerators**
 - Specialized processors
- How much of this hardware helps data processing?
- How much of this hardware has been developed for data processing?





Raghu Ramakrishnan

Technical Fellow, CTO for Data


**In 3 to 5 years, data analytics
will run on GPUs and FPGAs**

A New Golden Era for Data Management - Gustavo Alonso

A New Golden Era for Data Management, 2026, Gustavo Alonso: slides

Industry Trend: Microsoft

Coming soon



Best Industry Paper award
SIGMOD 2026

GPU accelerated Fabric DW

7x faster query performance
for reporting and apps


Optimized for high-concurrency,
AI-driven analytical workloads

aka.ms/GPUAcceleratedFabricDW


Microsoft Build 2026 Opening Keynote, Satya Nadella: youtu.be/FFMm454fxNA

Industry Trend: Microsoft

Coming soon



Best Industry Paper award
SIGMOD 2026



GPU accelerated Fabric DW

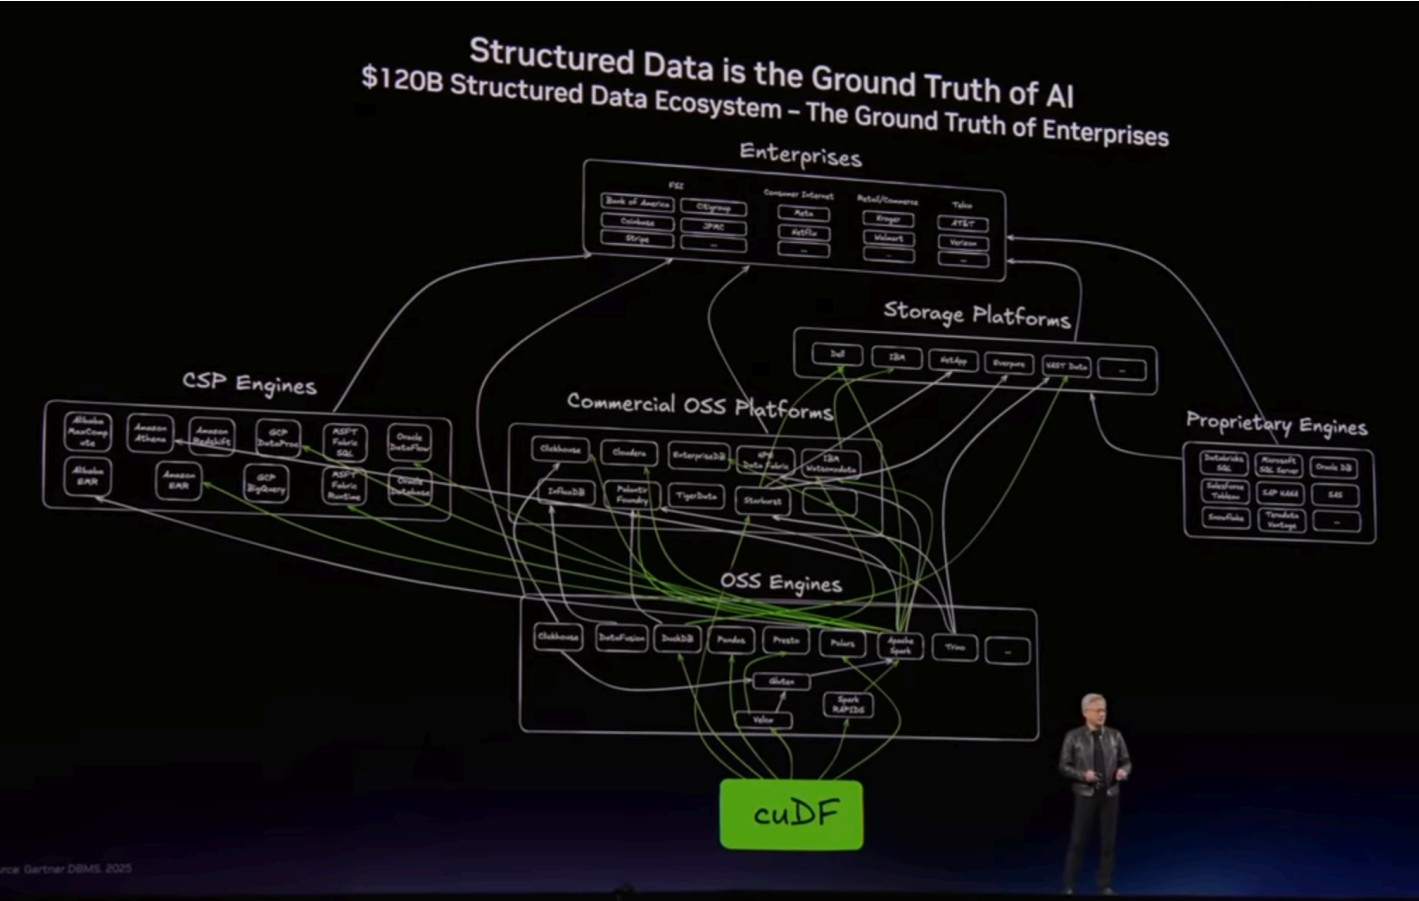
7x faster query performance
for reporting and apps

Optimized for high-concurrency,
AI-driven analytical workloads

aka.ms/GPUAcceleratedFabricDW

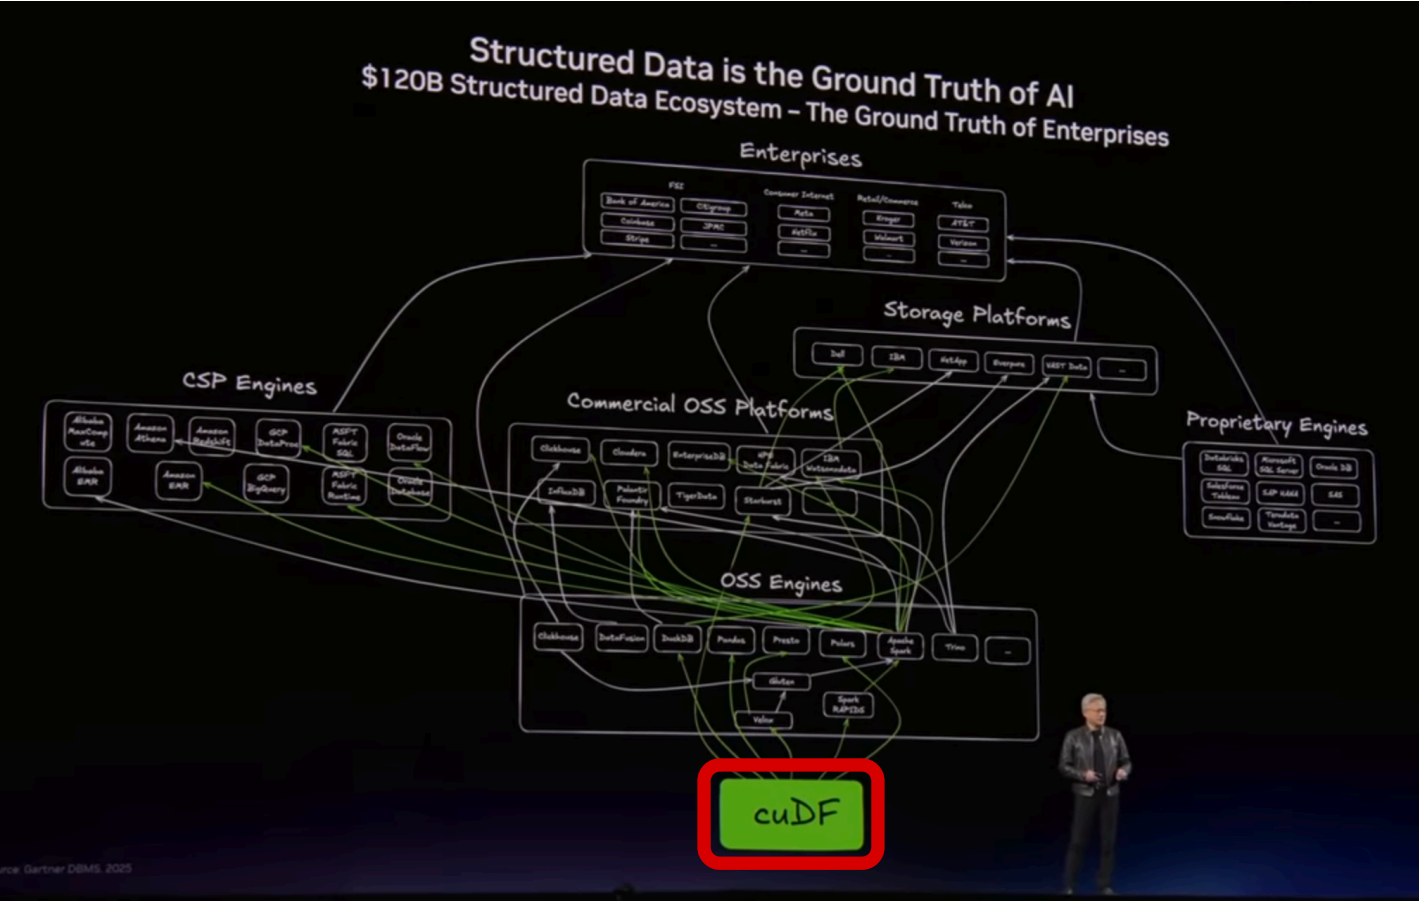
Microsoft Build 2026 Opening Keynote, Satya Nadella: youtu.be/FFMm454fxNA

Industry Trend: NVIDIA



NVIDIA GTC 2026 Full Keynote, Jensen Huang: youtu.be/jlviHI7fqyc

Industry Trend: NVIDIA



NVIDIA GTC 2026 Full Keynote, Jensen Huang: youtu.be/jlviHI7fqyc

GPU Data Processing: Research & Industry

Microsoft Fabric



Microsoft



SiriusDB



DuckDB



WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON

Velox cuDF



Meta



Spark-RAPIDS



Presto



presto



GPU Data Processing: Research & Industry

Microsoft Fabric



Microsoft



SiriusDB



DuckDB



WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON

Velox cuDF



Meta



Spark-RAPIDS



Presto



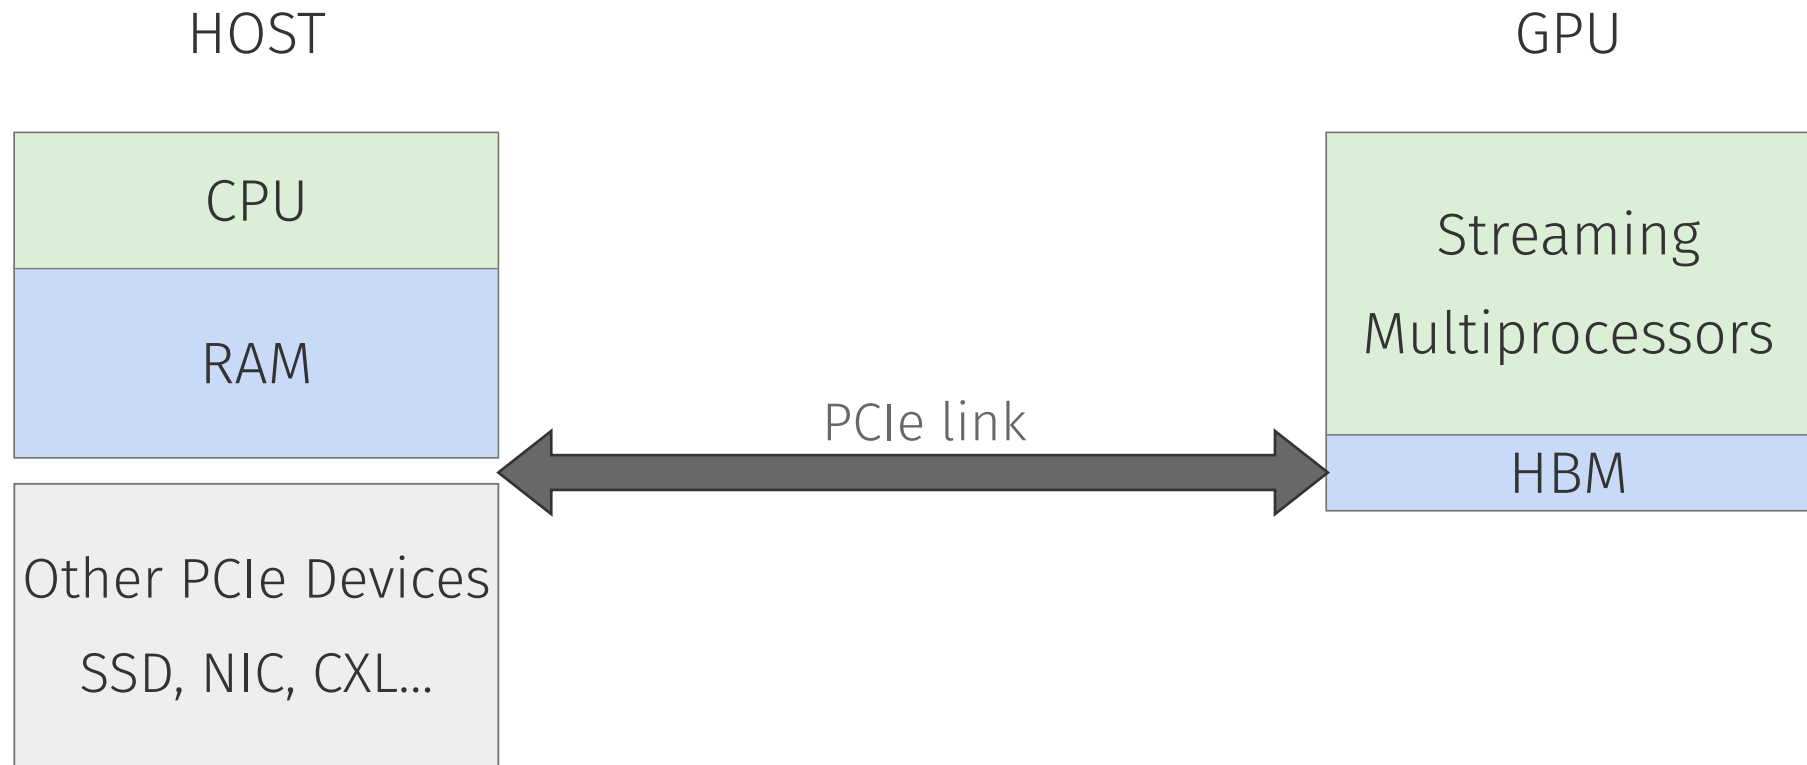
presto



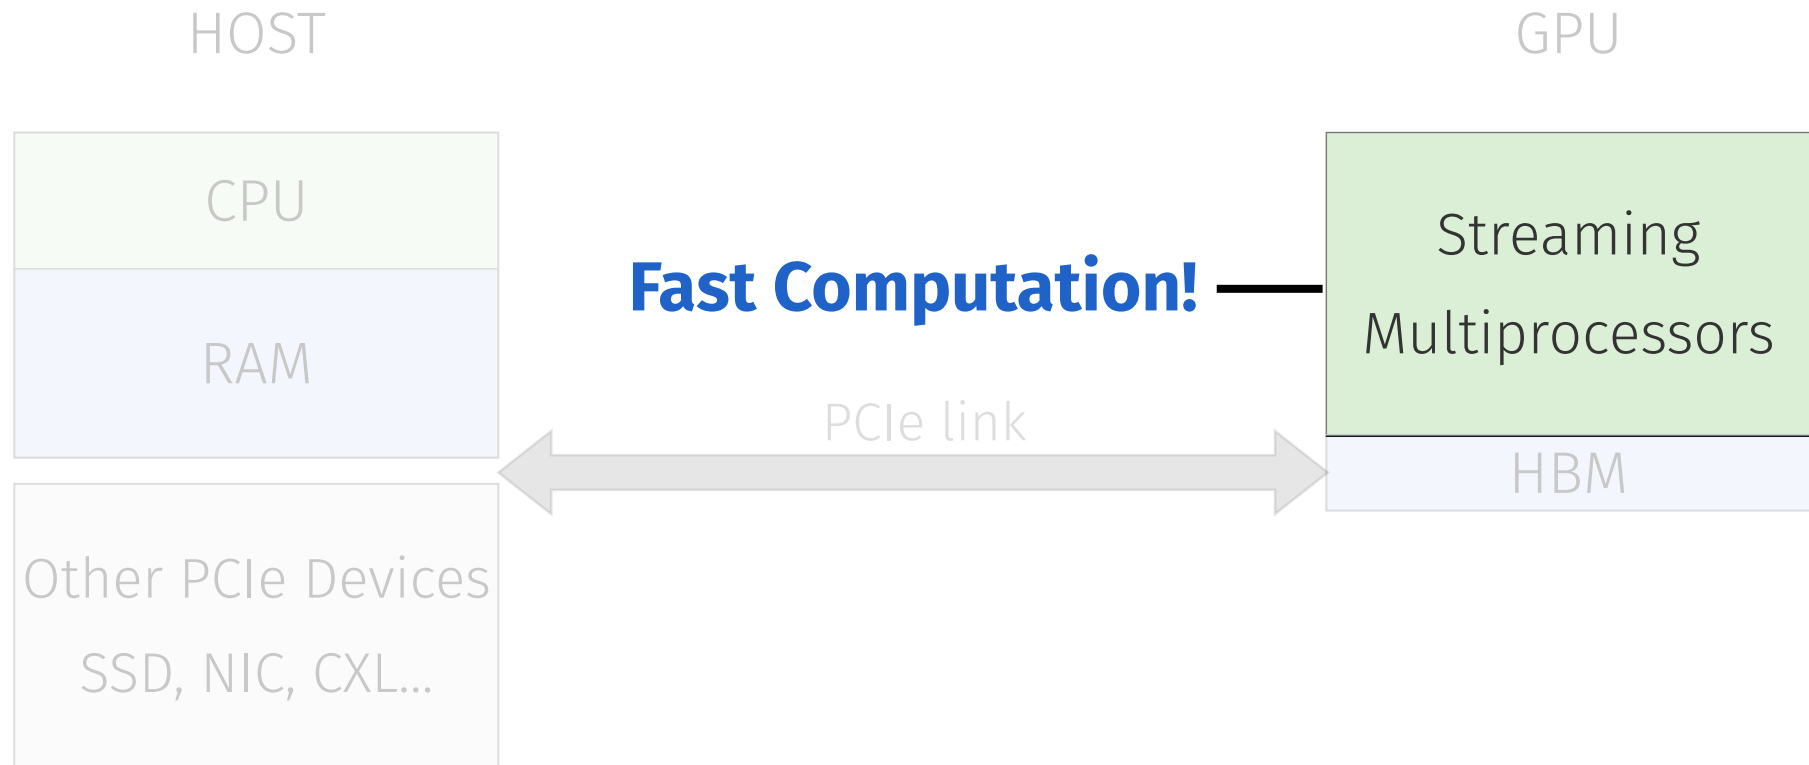
*: More will come with the help of Apache Gluten  Gluten

Background: GPU Data Processing

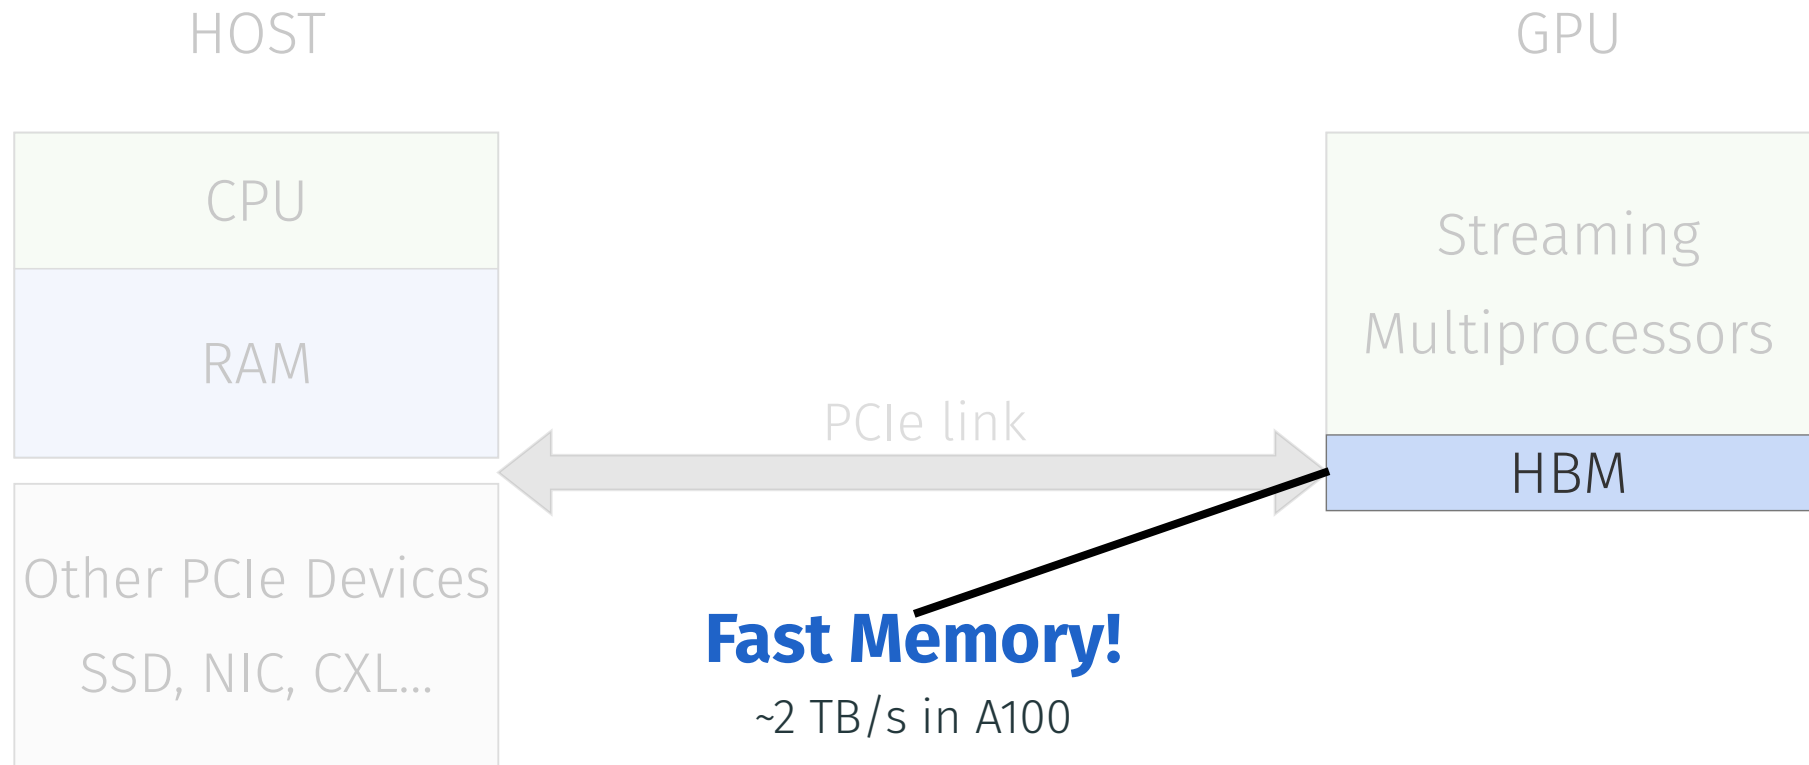
GPU System Primitives



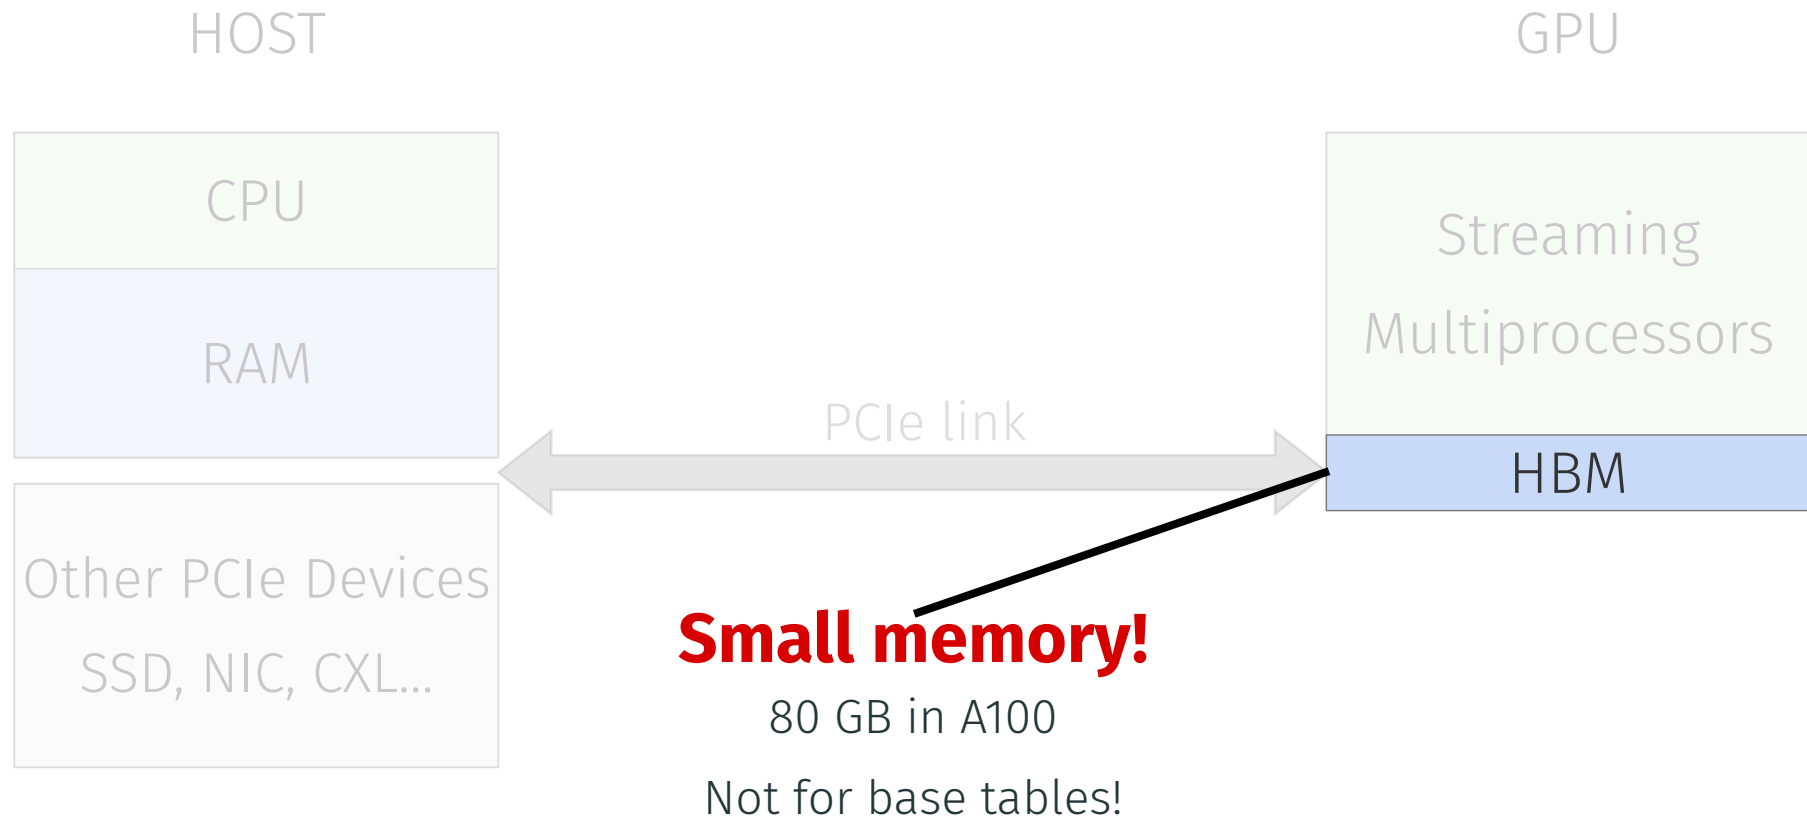
GPU System Primitives



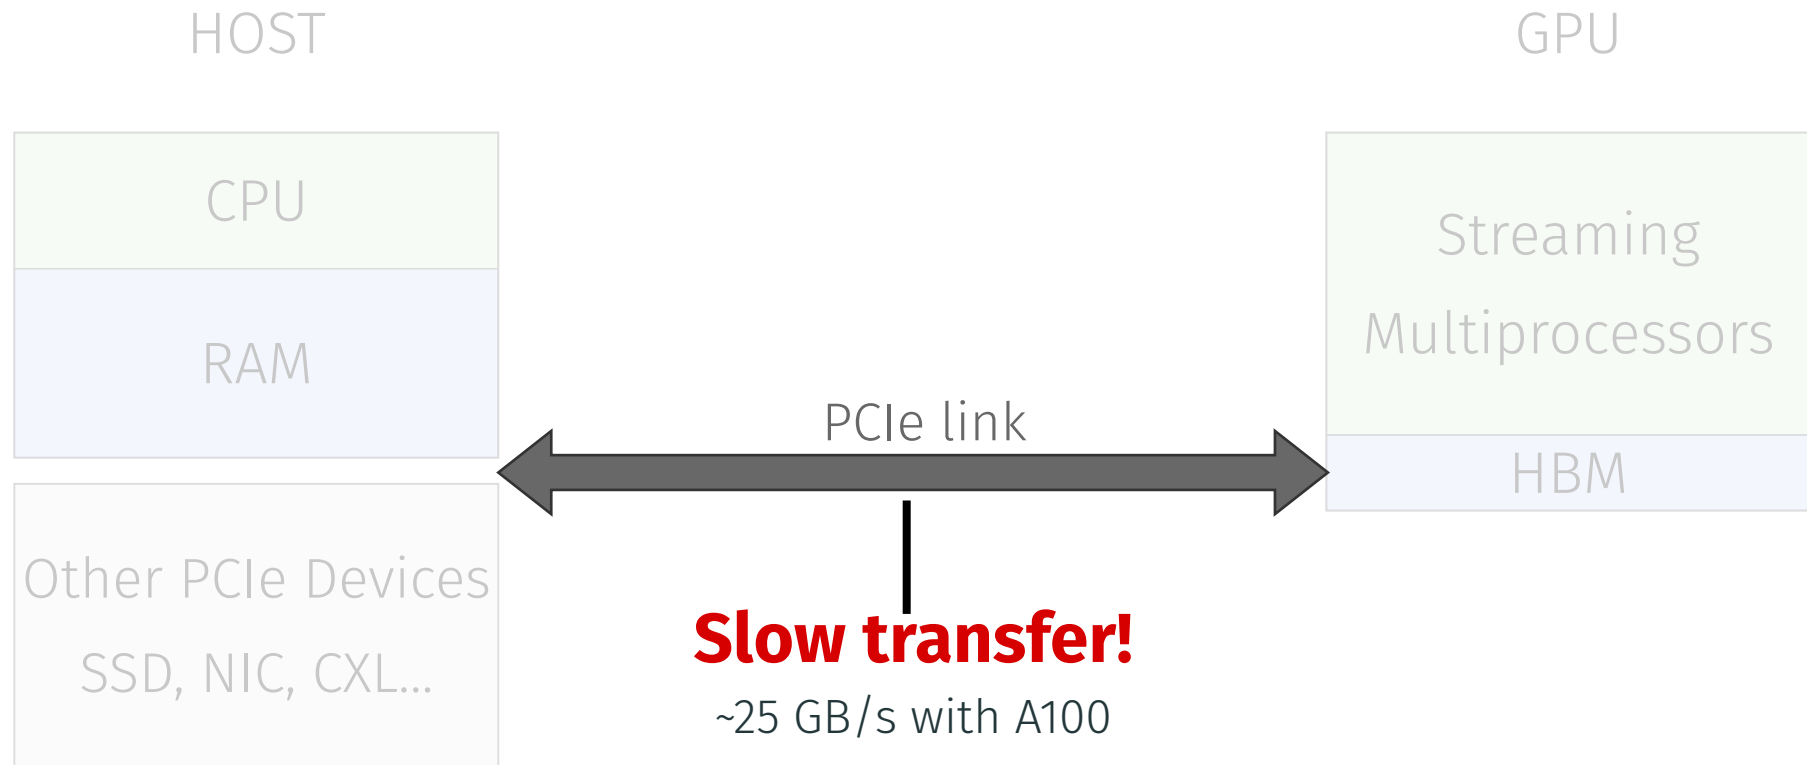
GPU System Primitives



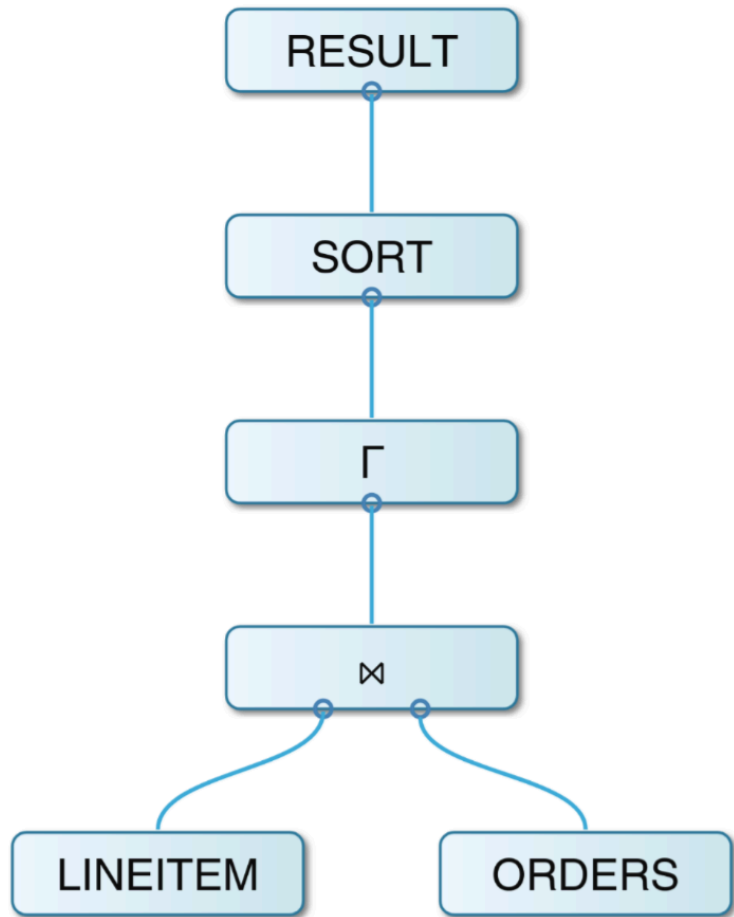
GPU System Primitives: Problems & Challenges




GPU System Primitives: Problems & Challenges



GPU Data Systems




GPU Data Systems



1. join kernel
2. groupby kernel
3. sort kernel

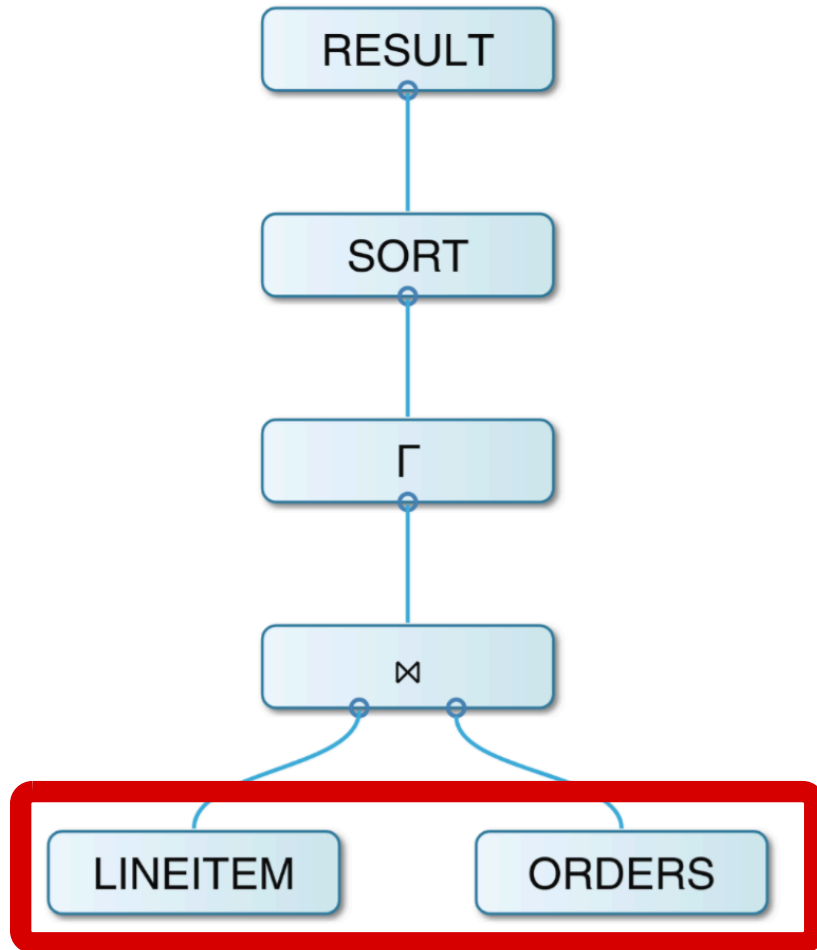
GPU Data Systems



1. join kernel
2. groupby kernel
3. sort kernel

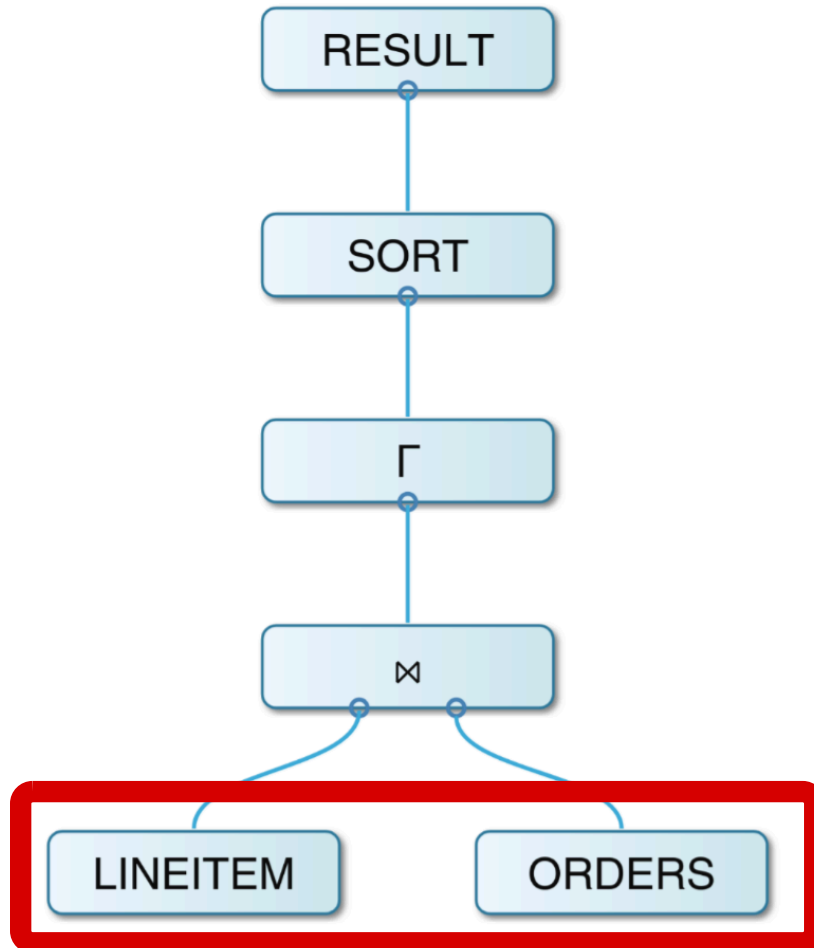
Not in this talk

GPU Data Systems



Where to put base tables?

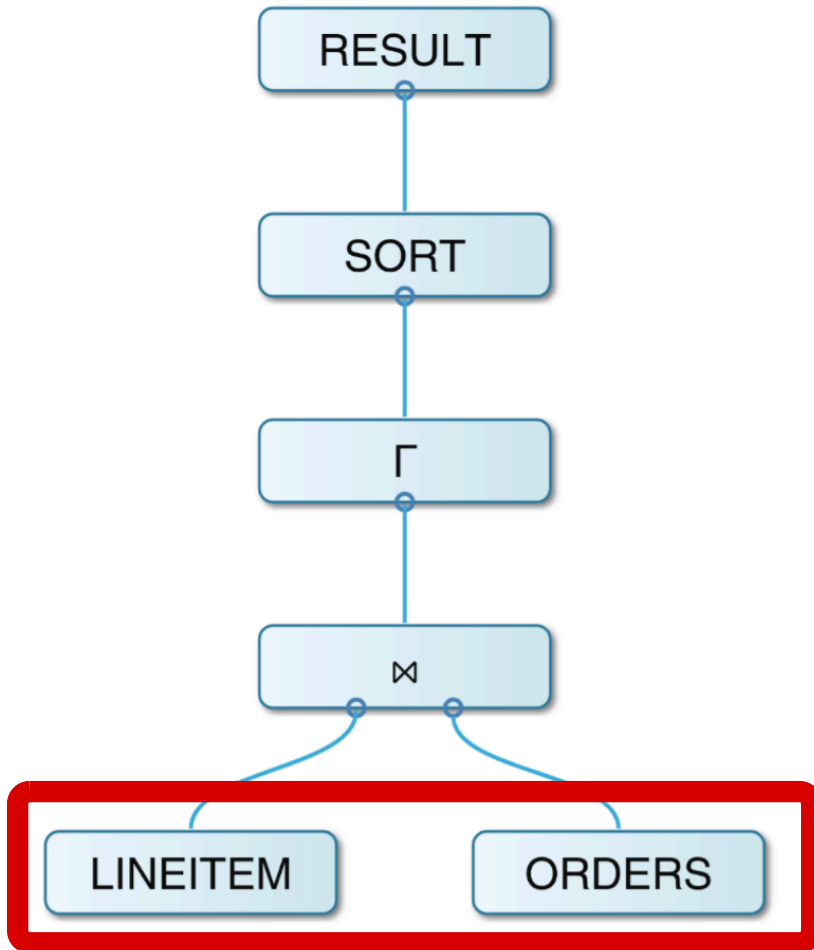
GPU Data Systems



Where to put base tables?

- GPU Memory: HBM


GPU Data Systems



Where to put base tables?

- GPU Memory: HBM
- HBM of neighbour GPUs: NVLink

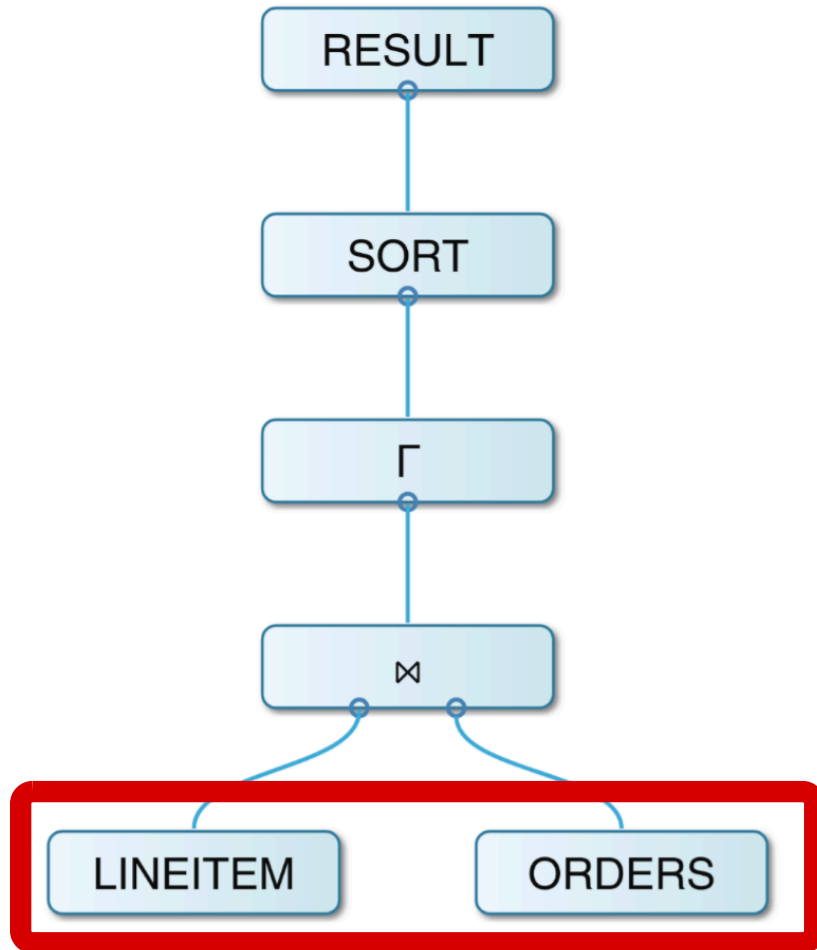
GPU Data Systems



Where to put base tables?

- GPU Memory: HBM
- HBM of neighbour GPUs: NVLink
- CPU Memory

GPU Data Systems



Where to put base tables?


- GPU Memory: HBM
- HBM of neighbour GPUs: NVLink
- CPU Memory
- SSDs
- NICs: remote storage

Common Technique 1: Compression


Raw data

large transfer

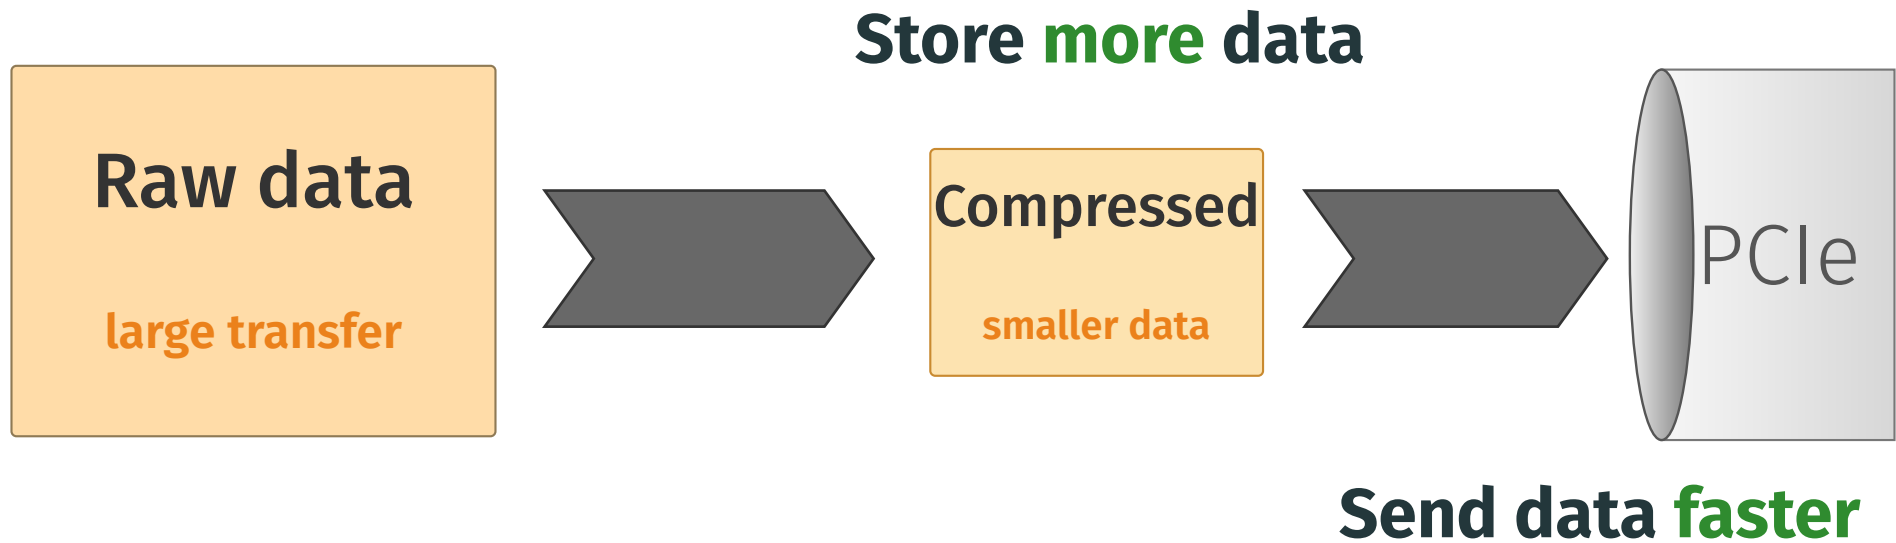
Common Technique 1: Compression



Common Technique 1: Compression




Common Technique 1: Compression




Common Technique 2: Overlapping

Blocking



Common Technique 2: Overlapping

Blocking




Common Technique 2: Overlapping

Blocking



Overlapped




Common Technique 2: Overlapping

Blocking





Overlapped




My Research


My First Meeting with Carsten




My First Meeting with Carsten



Do you know GSL?




My First Meeting with Carsten




Do you know GSL?

No.




My First Meeting with Carsten




Do you know GSL?

Do you know Matteo Interlandi?

No.



My First Meeting with Carsten




Do you know GSL?


Do you know Matteo Interlandi?

No.

No.



My First Meeting with Carsten




Do you know GSL?

No.


Do you know Matteo Interlandi?

No.

Do you know the TQP paper
from Matteo and GSL?



My First Meeting with Carsten



Do you know GSL?


Do you know Matteo Interlandi?

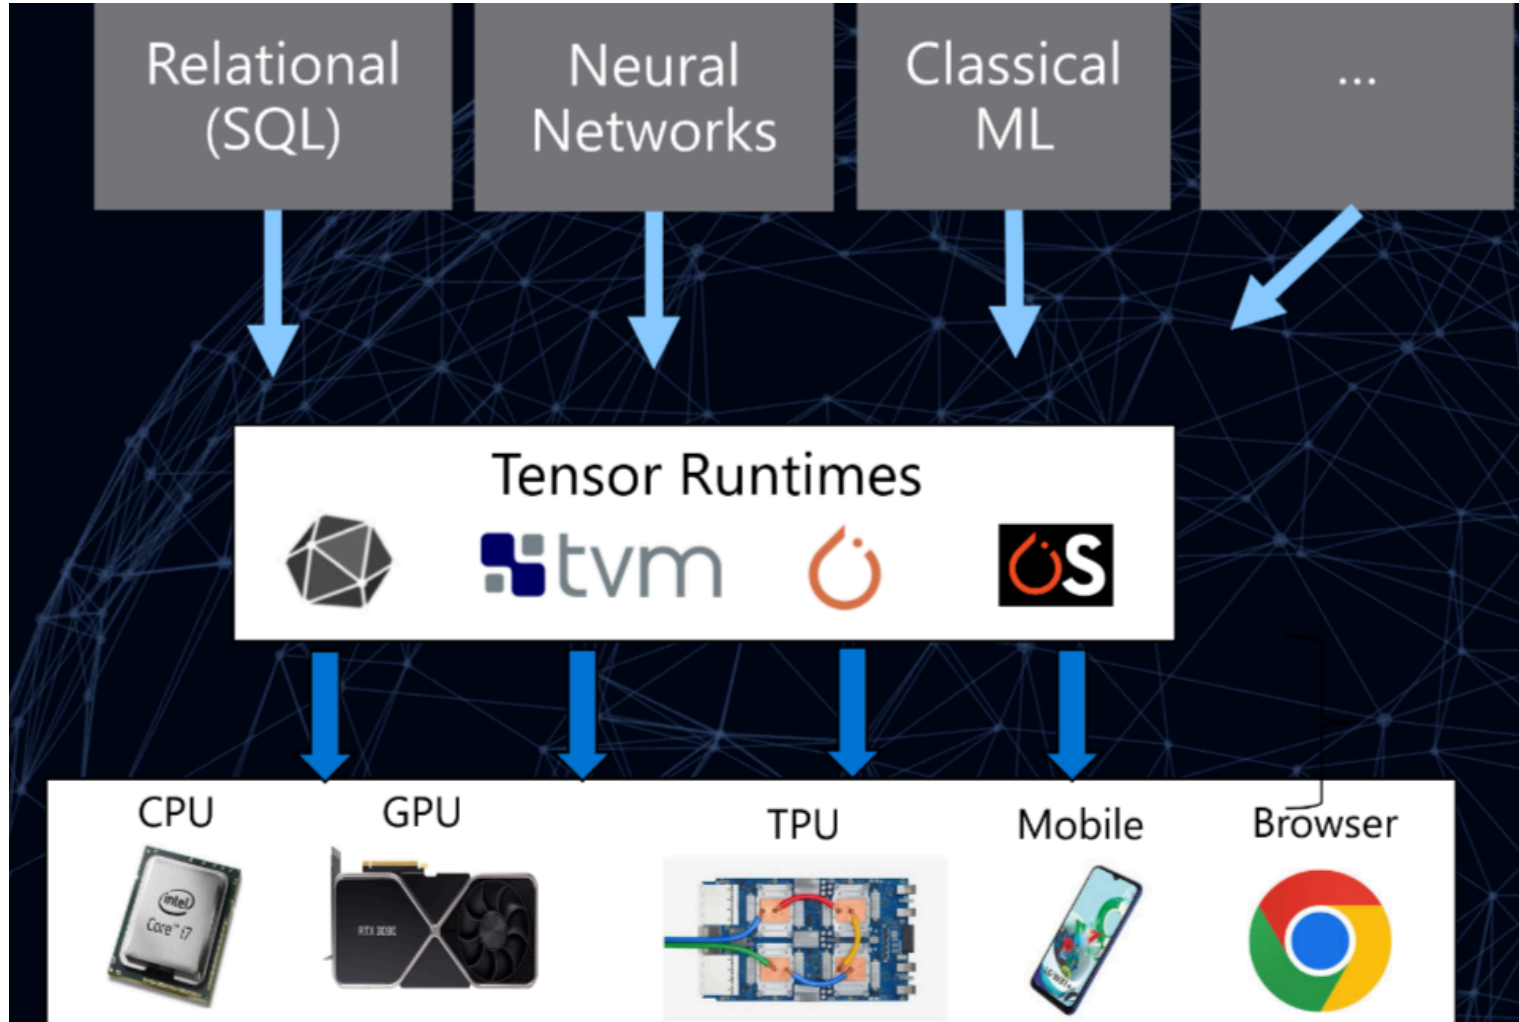
Do you know the TQP paper
from Matteo and GSL?

No.

No.

No.







Matteo Interlandi  • 1st

Principal Scientist Manager at Microsoft GSL

17h • 

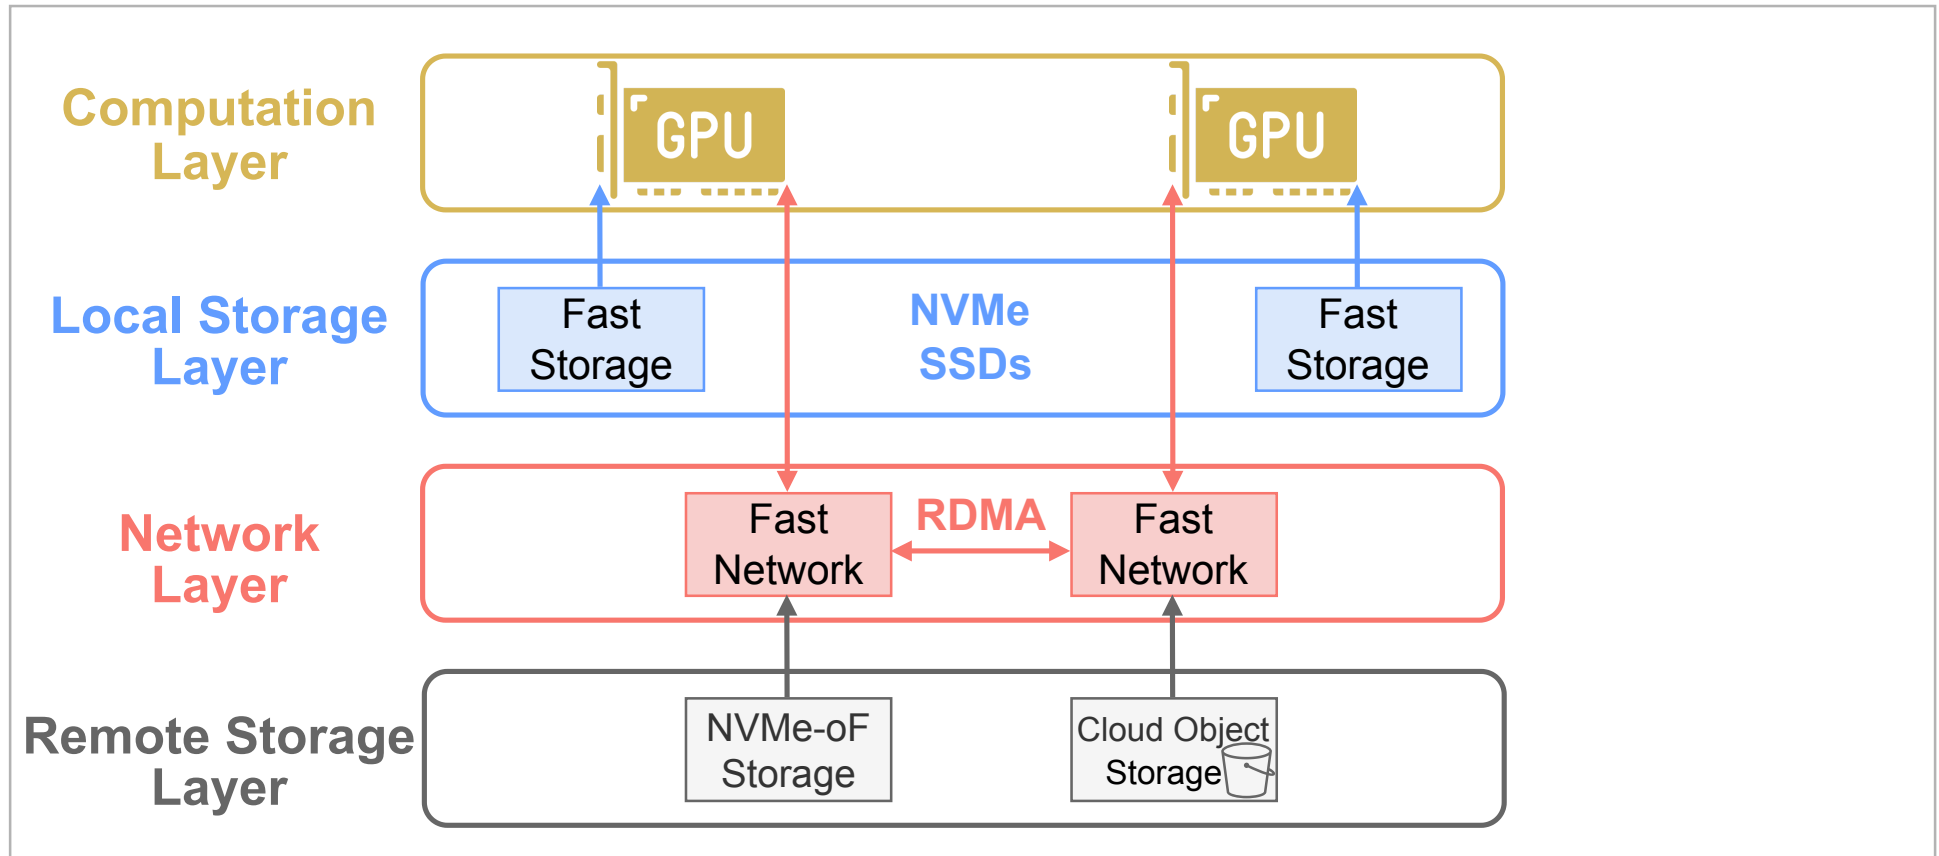


Five years ago, we made a slightly crazy **bet**: what if SQL could run on the same stack that powers AI?

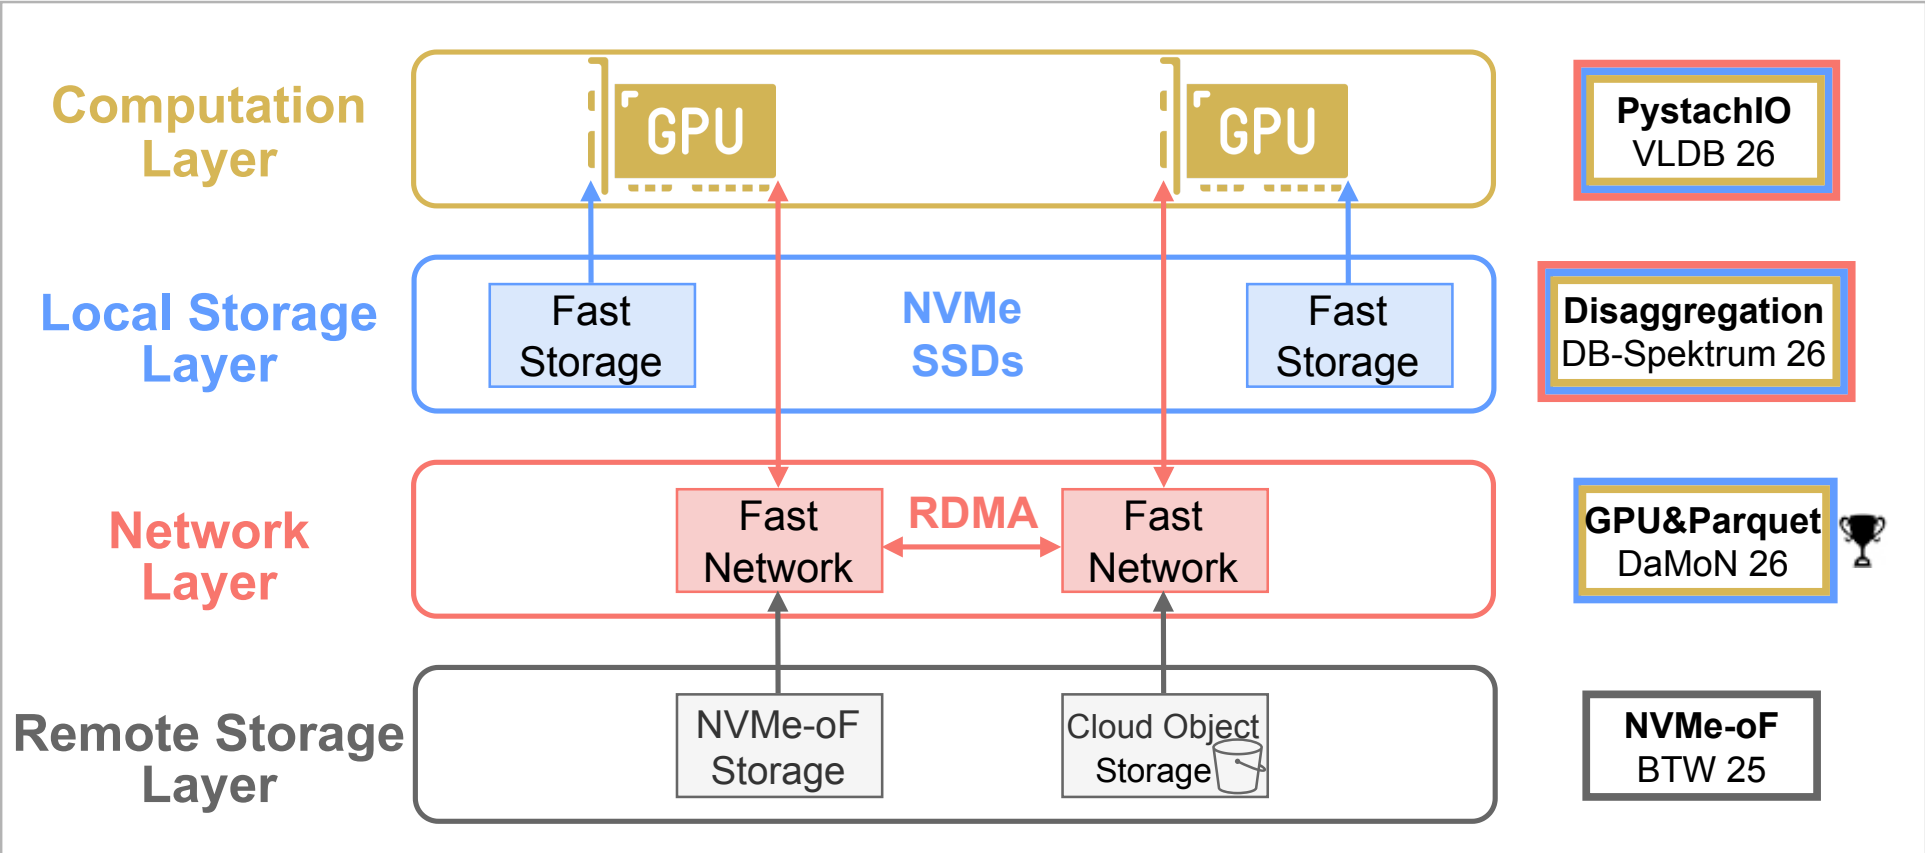
This week, that bet came full circle.

It began 5+ years ago as a research prototype we called the Tensor Query Processor (TQP). The idea was simple: express relational operators on top of the tensor API of PyTorch, so data systems could ride the entire AI wave (GPUs and beyond) for free.

My First Meeting with Carsten: A Roadmap

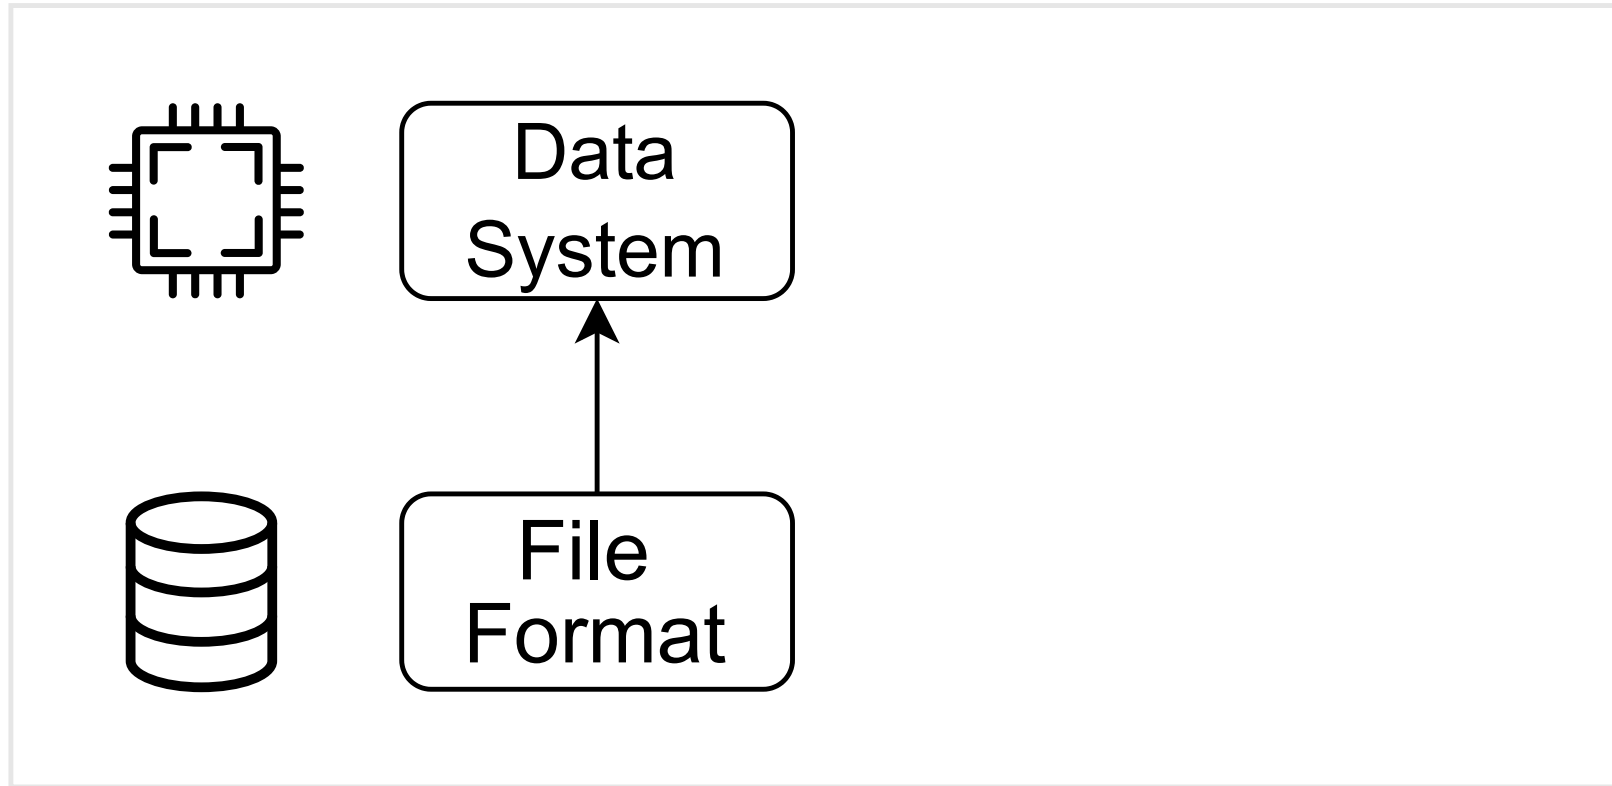


My Research Status

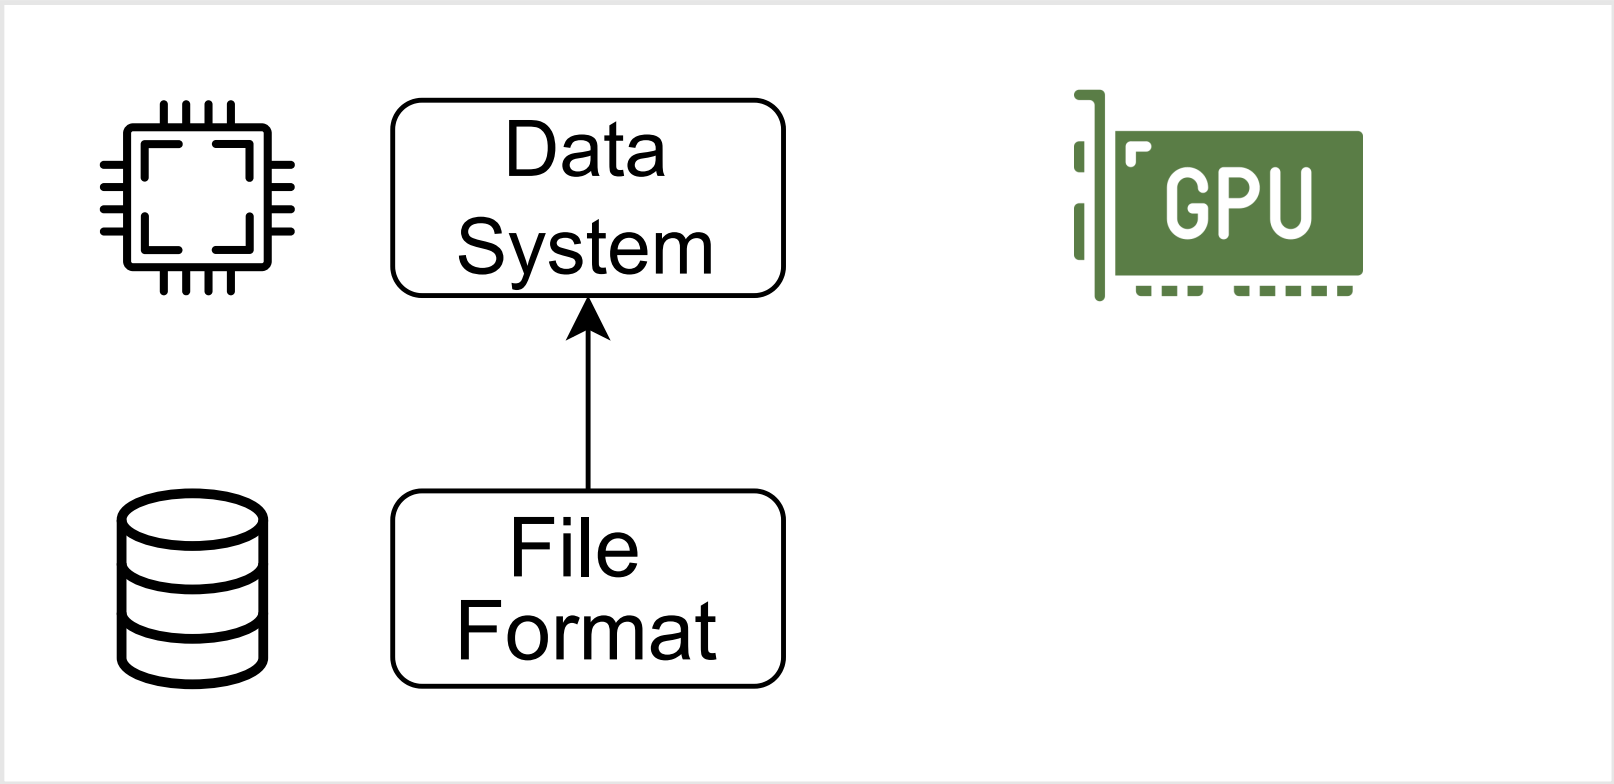


1. GPU & Parquet, DaMoN 2026


Data Processing: GPUs & File Formats



Data Processing: GPUs & File Formats



Data Processing: GPUs & File Formats




Problem: GPU Bottleneck With Parquet

Is Parquet Bad for GPUs?

Problem: GPU Bottleneck With Parquet

Is Parquet Bad for GPUs?



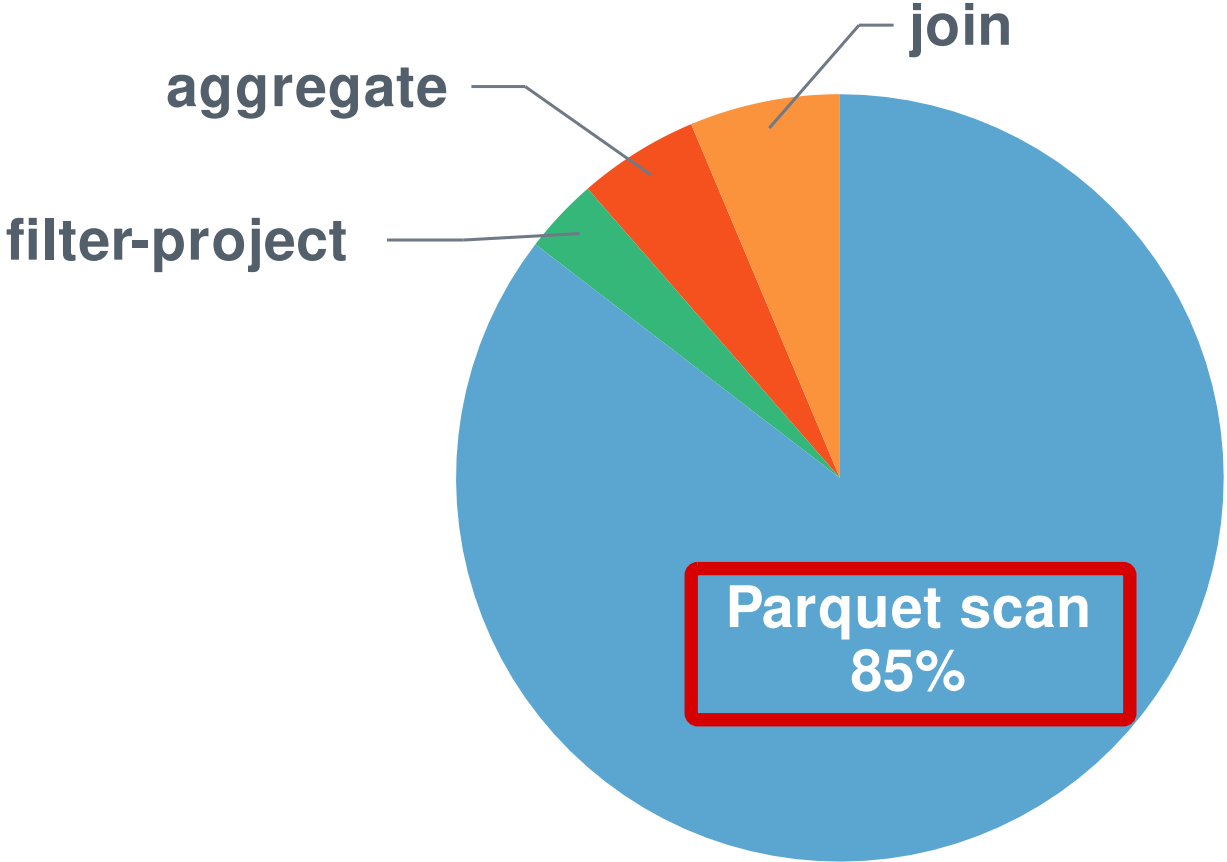
GPU TPC-H SF100 runtime breakdown.

Source: Accelerating Velox with RAPIDS cuDF, Velox-cuDF Team, 2025

Problem: GPU Bottleneck With Parquet

Is Parquet Bad for GPUs?

- **85%** in Parquet Scan



GPU TPC-H SF100 runtime breakdown.


Source: Accelerating Velox with RAPIDS cuDF, Velox-cuDF Team, 2025

Problem: GPU Bottleneck With Parquet

Is Parquet REALLY Bad for GPUs?

Is Parquet Bad for GPUs?

- 85% in Parquet Scan



GPU TPC-H SF100 runtime breakdown.


Source: Accelerating Velox with RAPIDS cuDF, Velox-cuDF Team, 2025

Problem: GPU Bottleneck With Parquet

Is Parquet REALLY Bad for GPUs?

Is Parquet Bad for GPUs?

- 85% in Parquet Scan




Do GPUs REALLY Need New File Formats?

GPU TPC-H SF100 runtime breakdown.

Source: Accelerating Velox with RAPIDS cuDF, Velox-cuDF Team, 2025

Parquet Configuration Defaults for CPUs

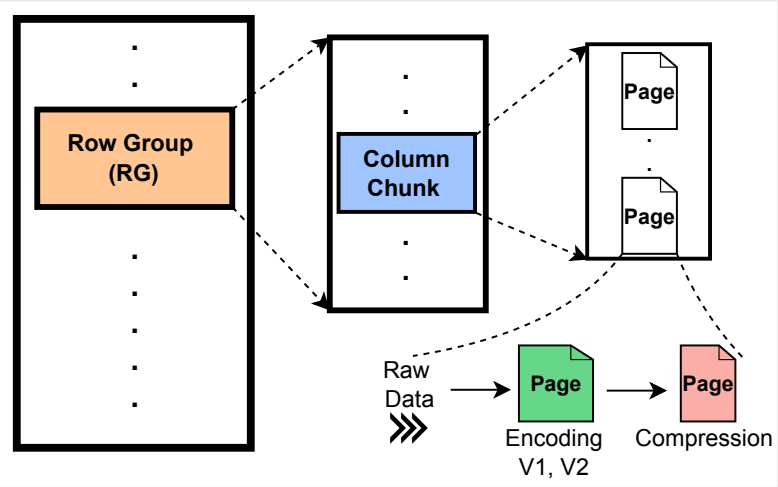
CPU Parquet Config. Defaults in  DuckDB:



Parquet Configuration Defaults for CPUs

CPU Parquet Config. Defaults in  **DuckDB**:

- **Row group**: 122880 rows



Parquet Configuration Defaults for CPUs

CPU Parquet Config. Defaults in  **DuckDB**:

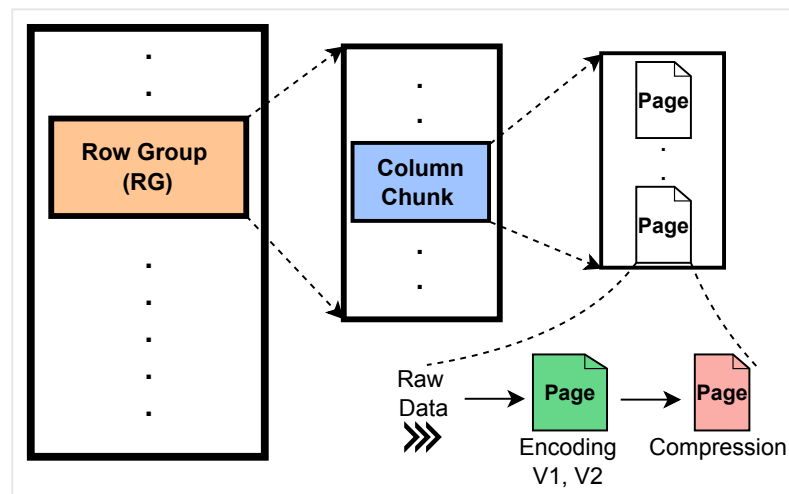
- **Row group**: 122880 rows
- **Column chunk**: 1 page only



Parquet Configuration Defaults for CPUs

CPU Parquet Config. Defaults in  **DuckDB**:


- **Row group**: 122880 rows
- **Column chunk**: 1 page only
- **Encoding**: V1 only



Parquet Configuration Defaults for CPUs

CPU Parquet Config. Defaults in  **DuckDB**:


- **Row group**: 122880 rows
- **Column chunk**: 1 page only
- **Encoding**: V1 only
- **Compression**: used even without size reduction



Issue: CPU Defaults Config. \neq GPU Optimal Config.

CPU Parquet Config. Defaults in  **DuckDB**:


- **Row group**: 122880 rows
- **Column chunk**: 1 page only
- **Encoding**: V1 only
- **Compression**: used even without size reduction



Issue: CPU Defaults Config. \neq GPU Optimal Config.

CPU Parquet Config. Defaults in  **DuckDB**:

- **Row group**: 122880 rows
- **Column chunk**: 1 page only
- **Encoding**: V1 only
- **Compression**: used even without size reduction

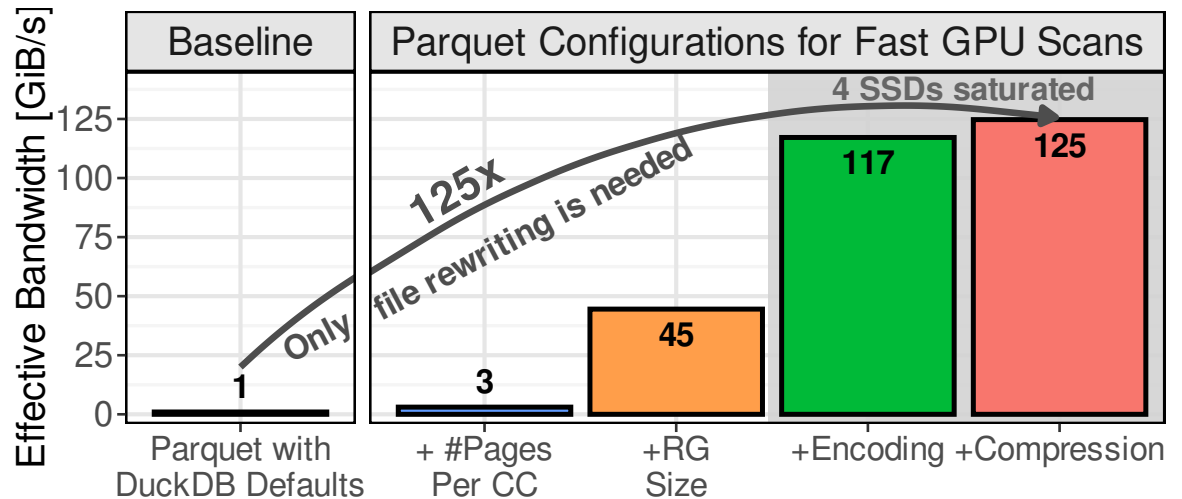


CPU Defaults Config. \neq GPU Optimal Config.

Issue: CPU Defaults Config. \neq GPU Optimal Config.

CPU Parquet Config. Defaults in DuckDB:

- **Row group**: 122880 rows
- **Column chunk**: 1 page only
- **Encoding**: V1 only
- **Compression**: used even without size reduction

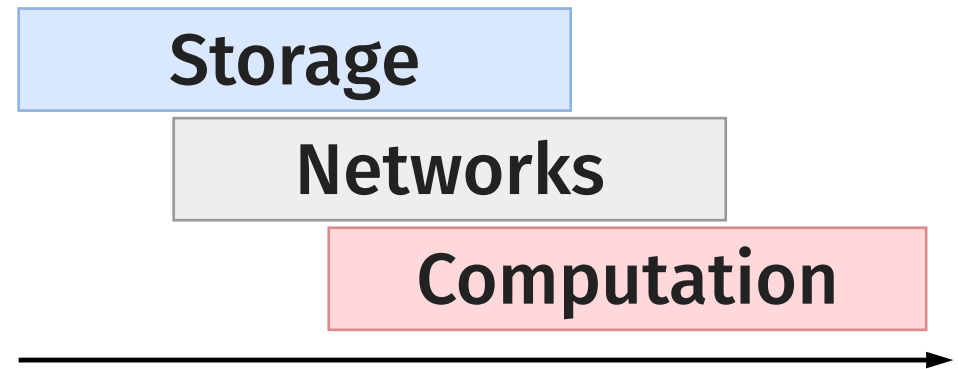


2. PystachIO, VLDB 2026

Overlapping: Storage, Networks, and Computation

- Storage engine reading Parquet
- Network for distributed join
- **PystachIO**: Storage + Network

Overlapped



PystachIO: Fast Storage and Fast Network

- Storage optimized
- Network optimized
- Next?

Storage



Networks & Join

PystachIO: Fast Storage and Fast Network

- Storage optimized
- Network optimized
- Next?

Storage

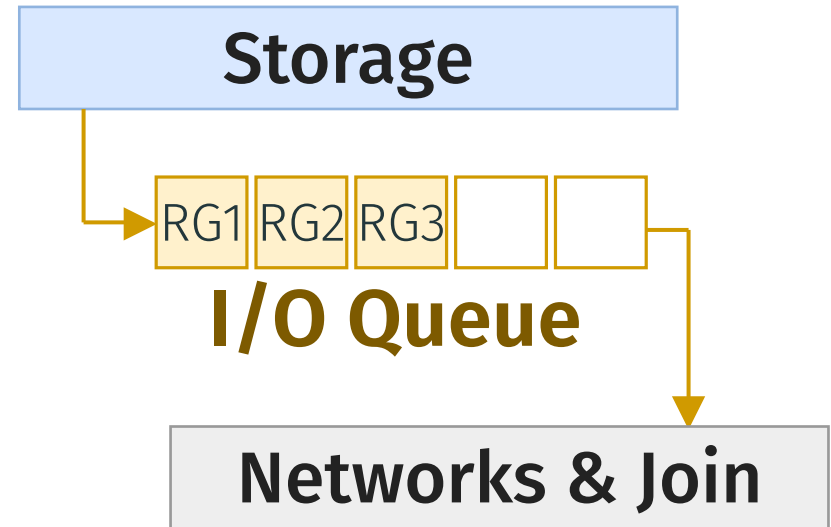
But how to overlap?

?

Networks & Join

PystachIO: I/O Queue

- Storage: enqueueing
- Query engine: dequeuing
- Then network shuffling




Overlapping Network, Storage, and Computation




Overlapping Network, Storage, and Computation




Overlapping Network, Storage, and Computation



Overlapping Network, Storage, and Computation



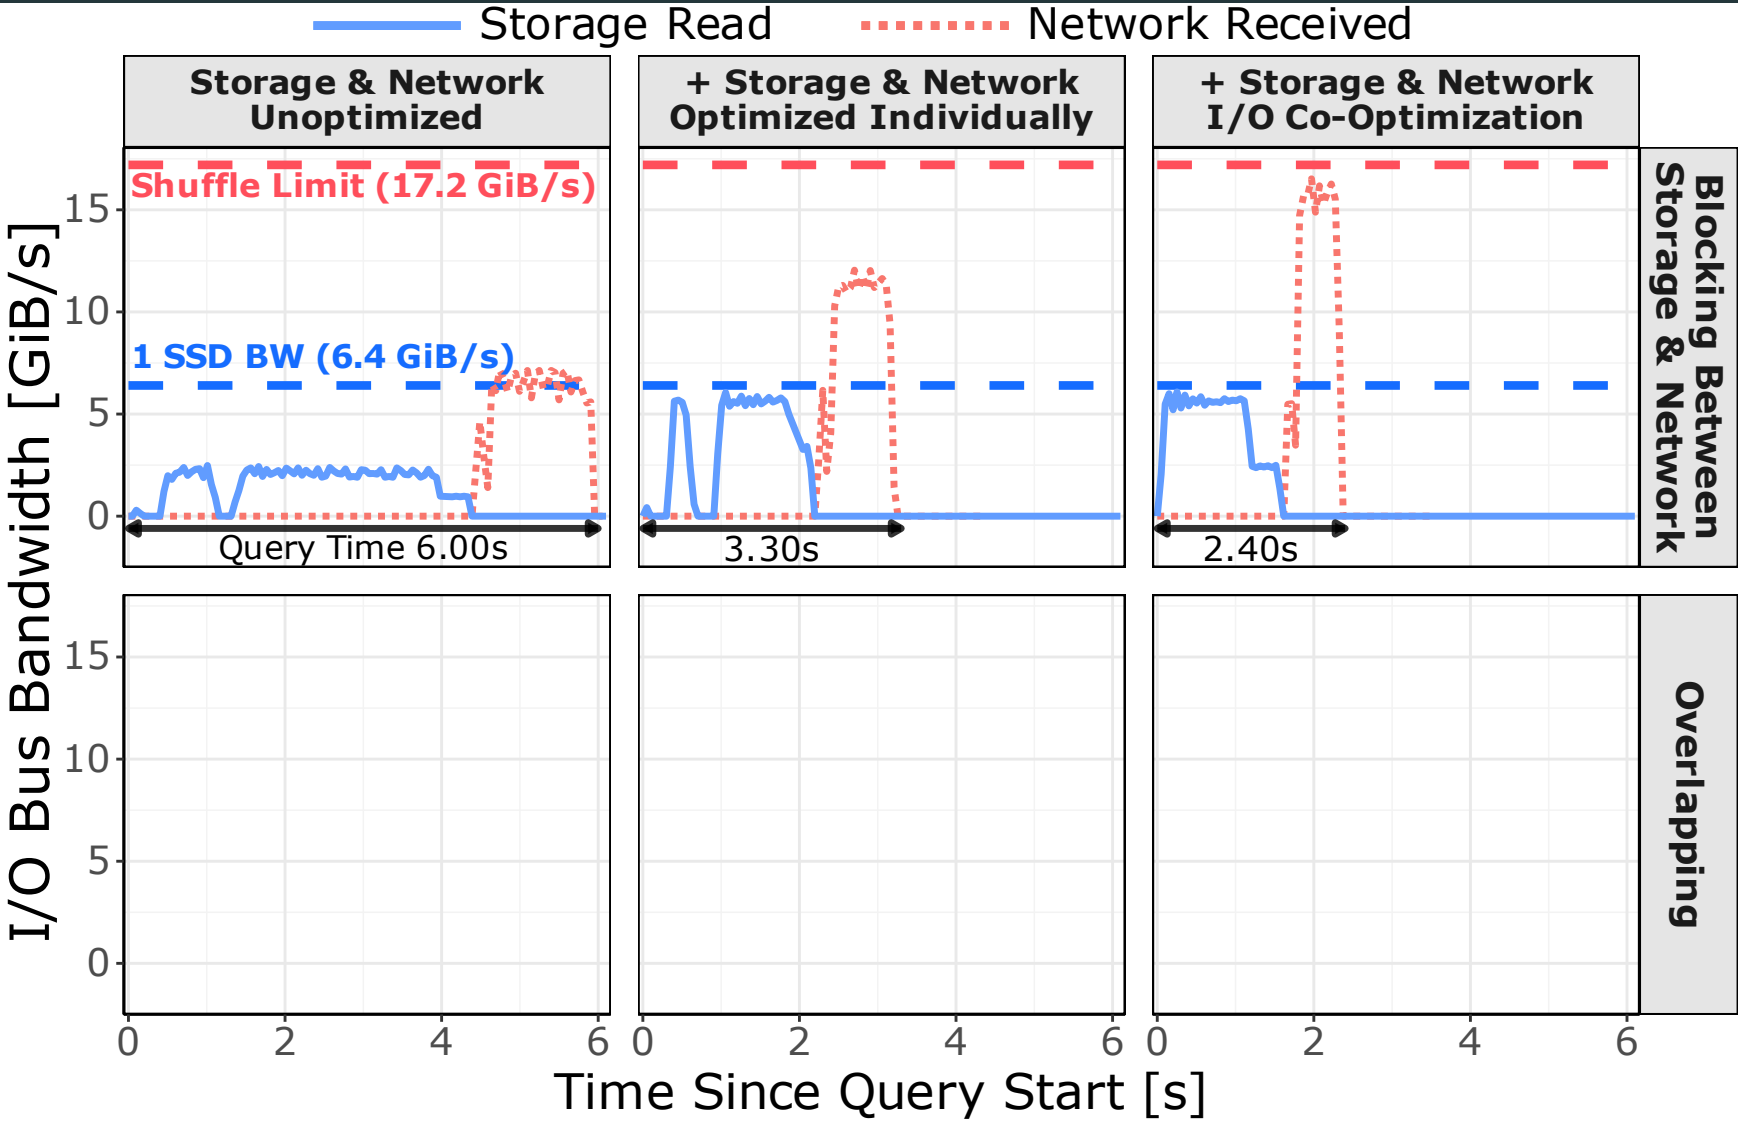
Overlapping Network, Storage, and Computation



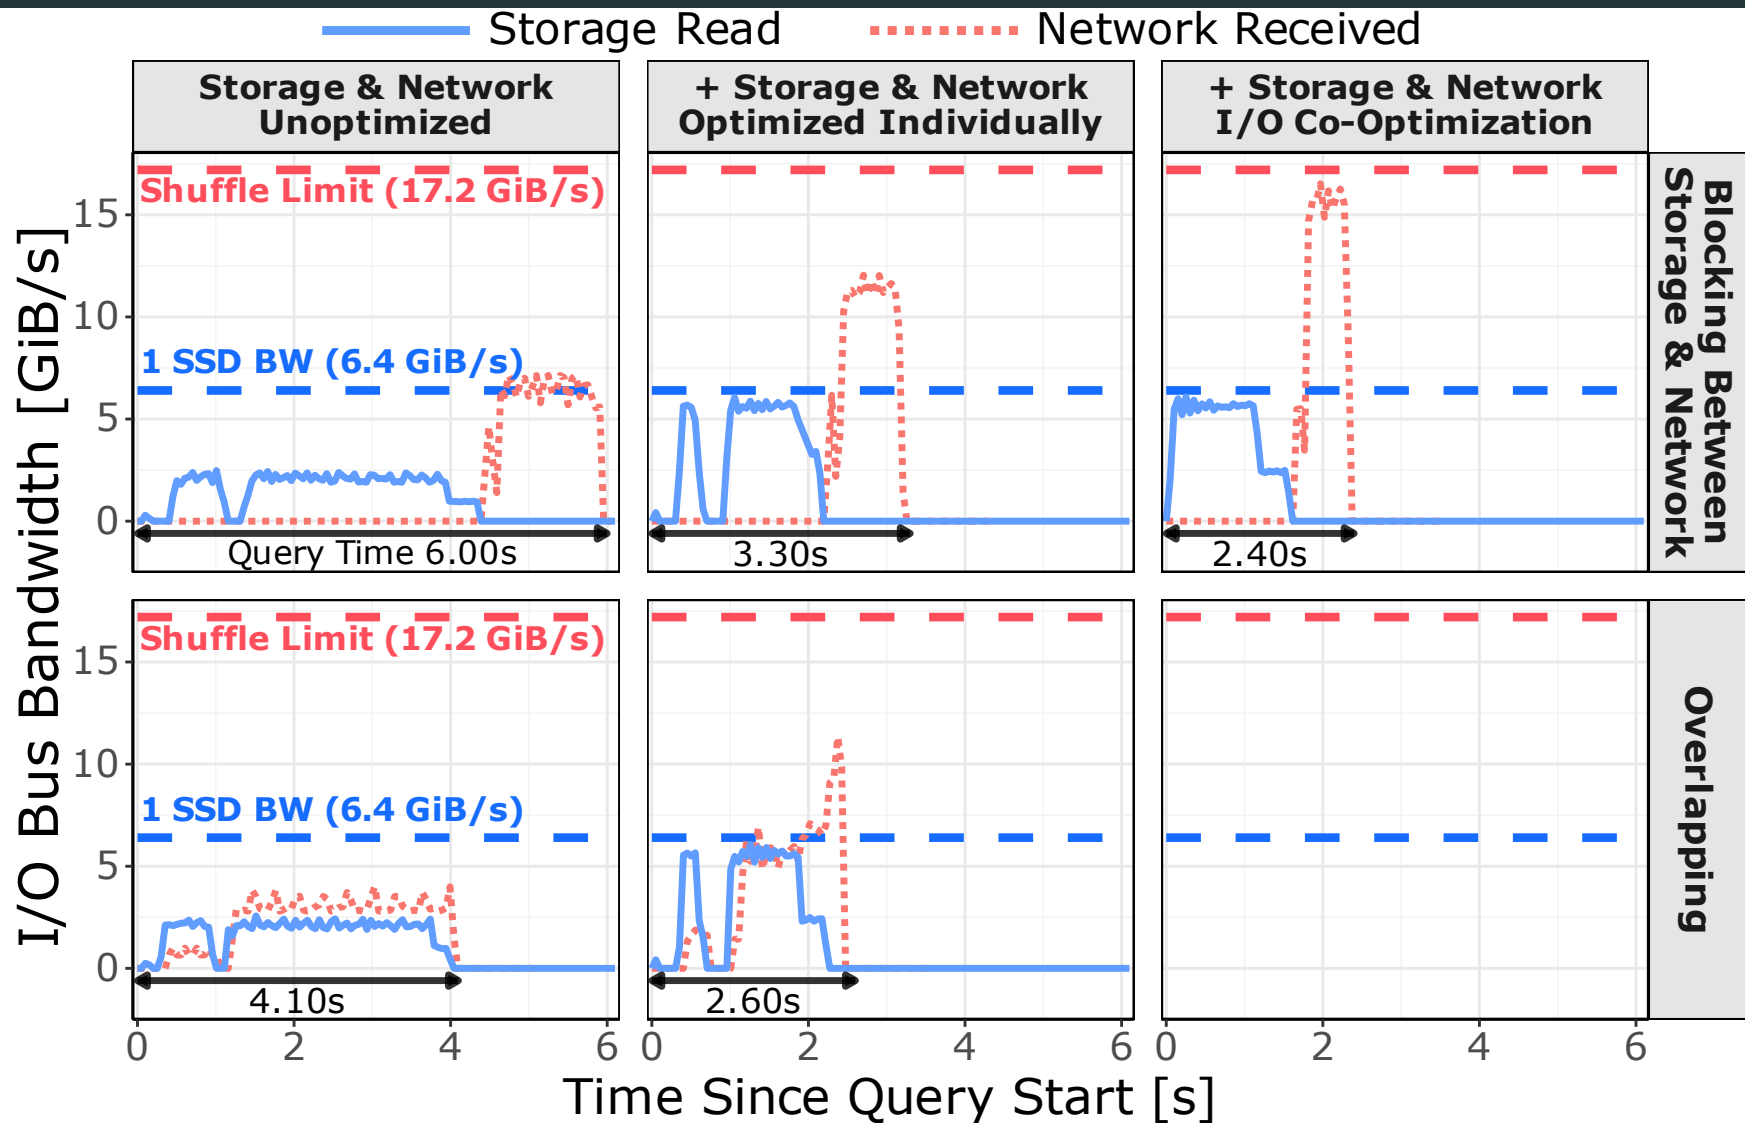
Overlapping Network, Storage, and Computation



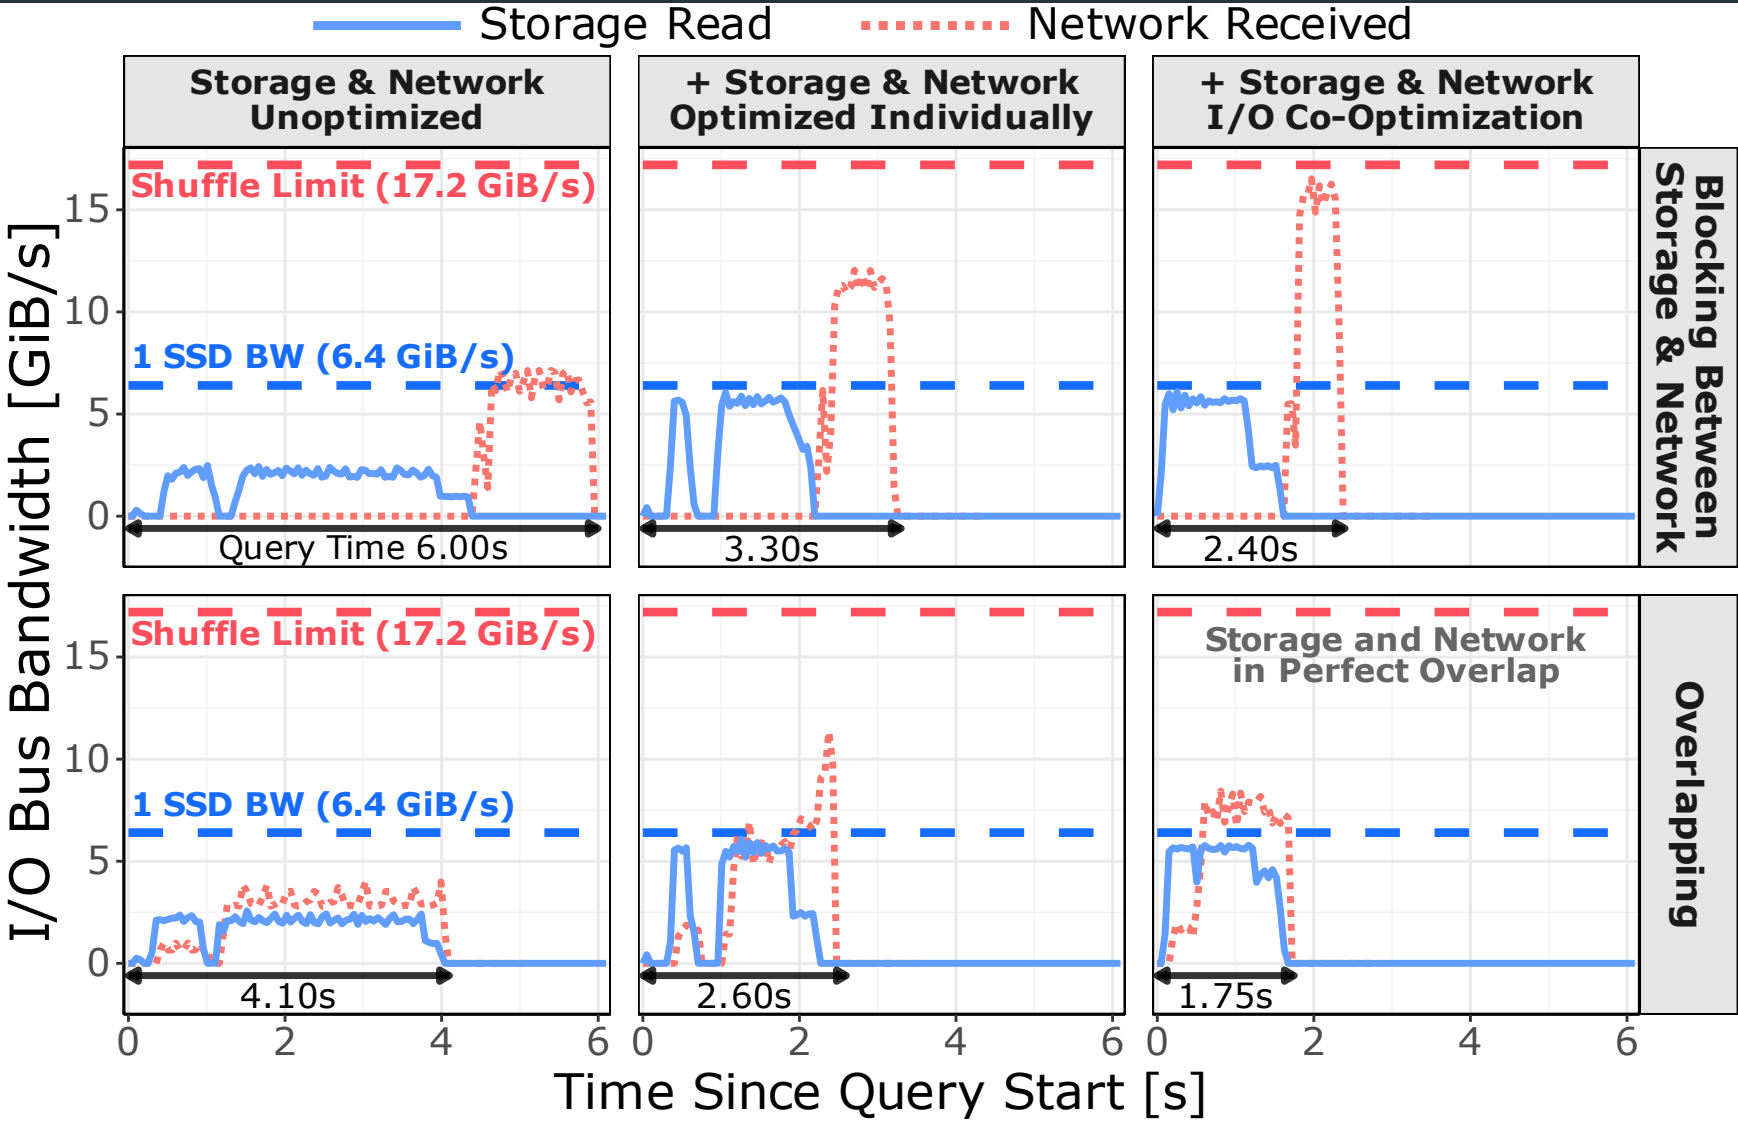
Overlapping Network, Storage, and Computation



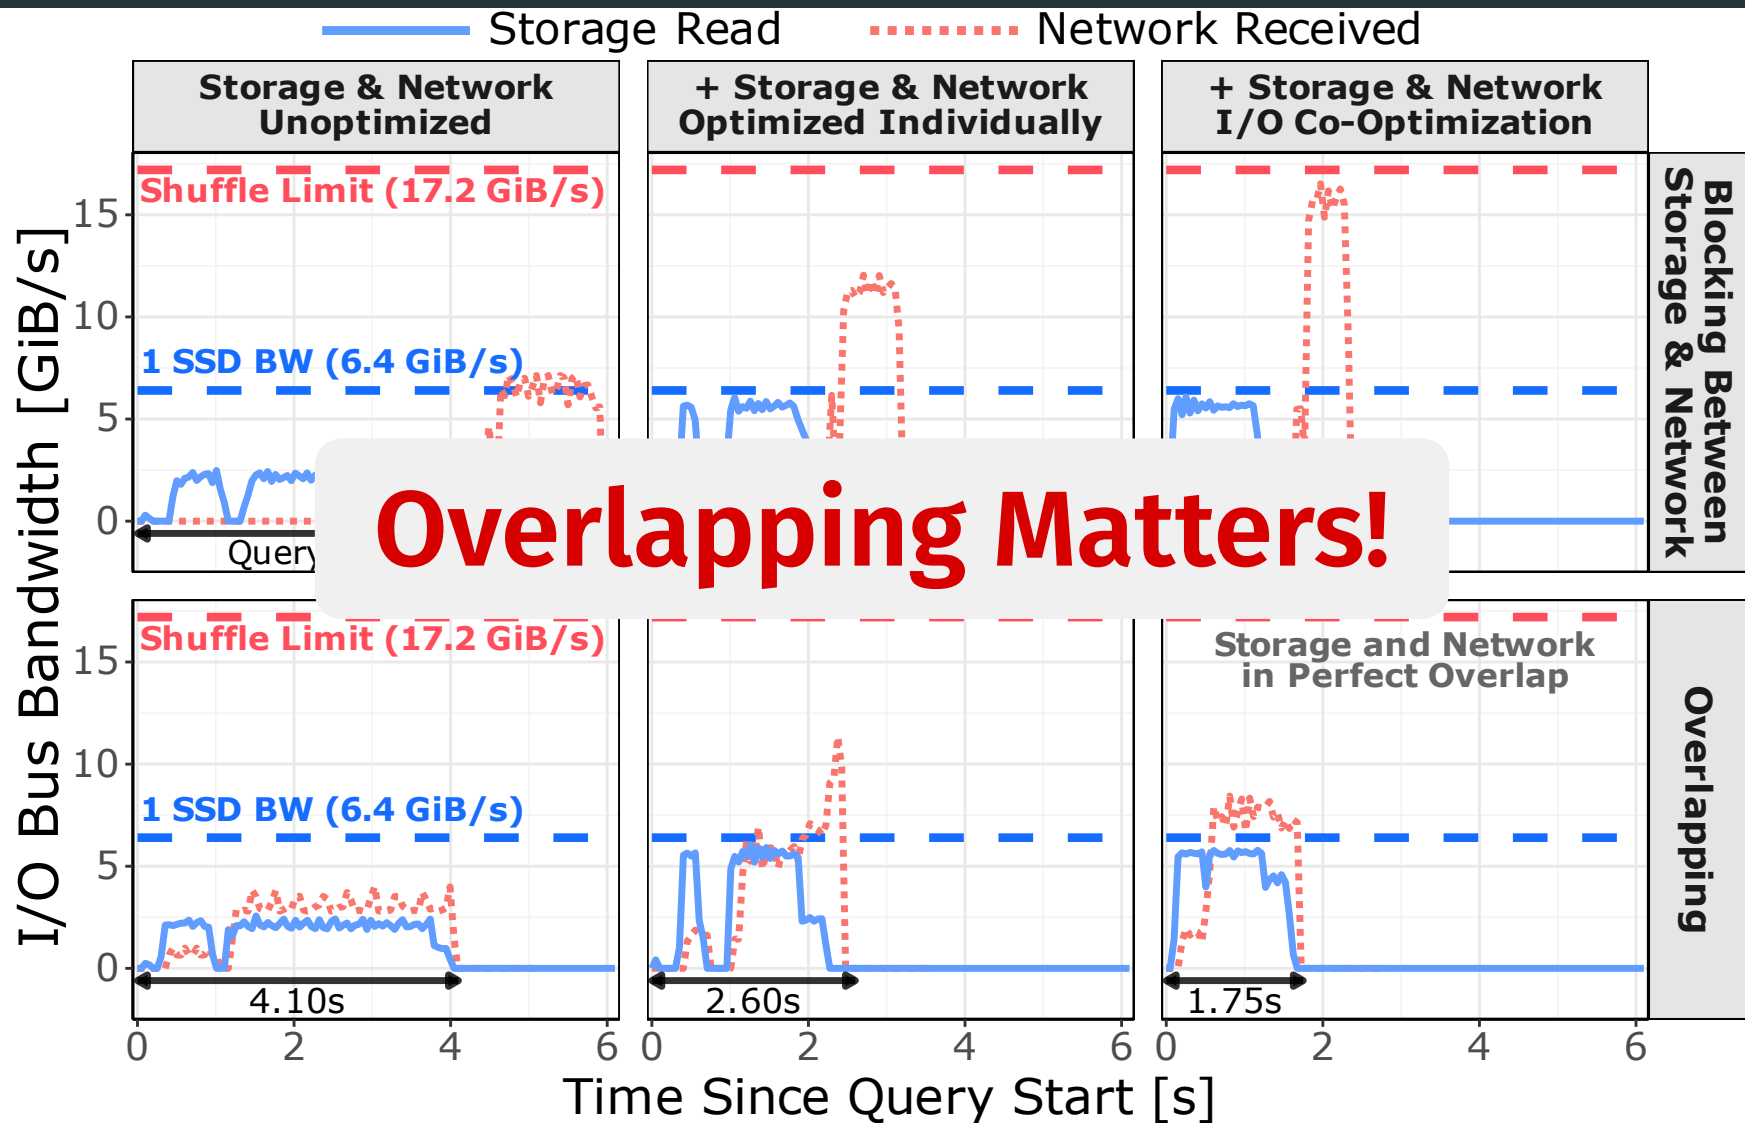
Overlapping Network, Storage, and Computation



Overlapping Network, Storage, and Computation








Overlapping Network, Storage, and Computation



More Than Research

-  **nVIDIA** 

More Than Research

-  **NVIDIA** 
-   **Parquet**  **APACHE DATAFUSION™**

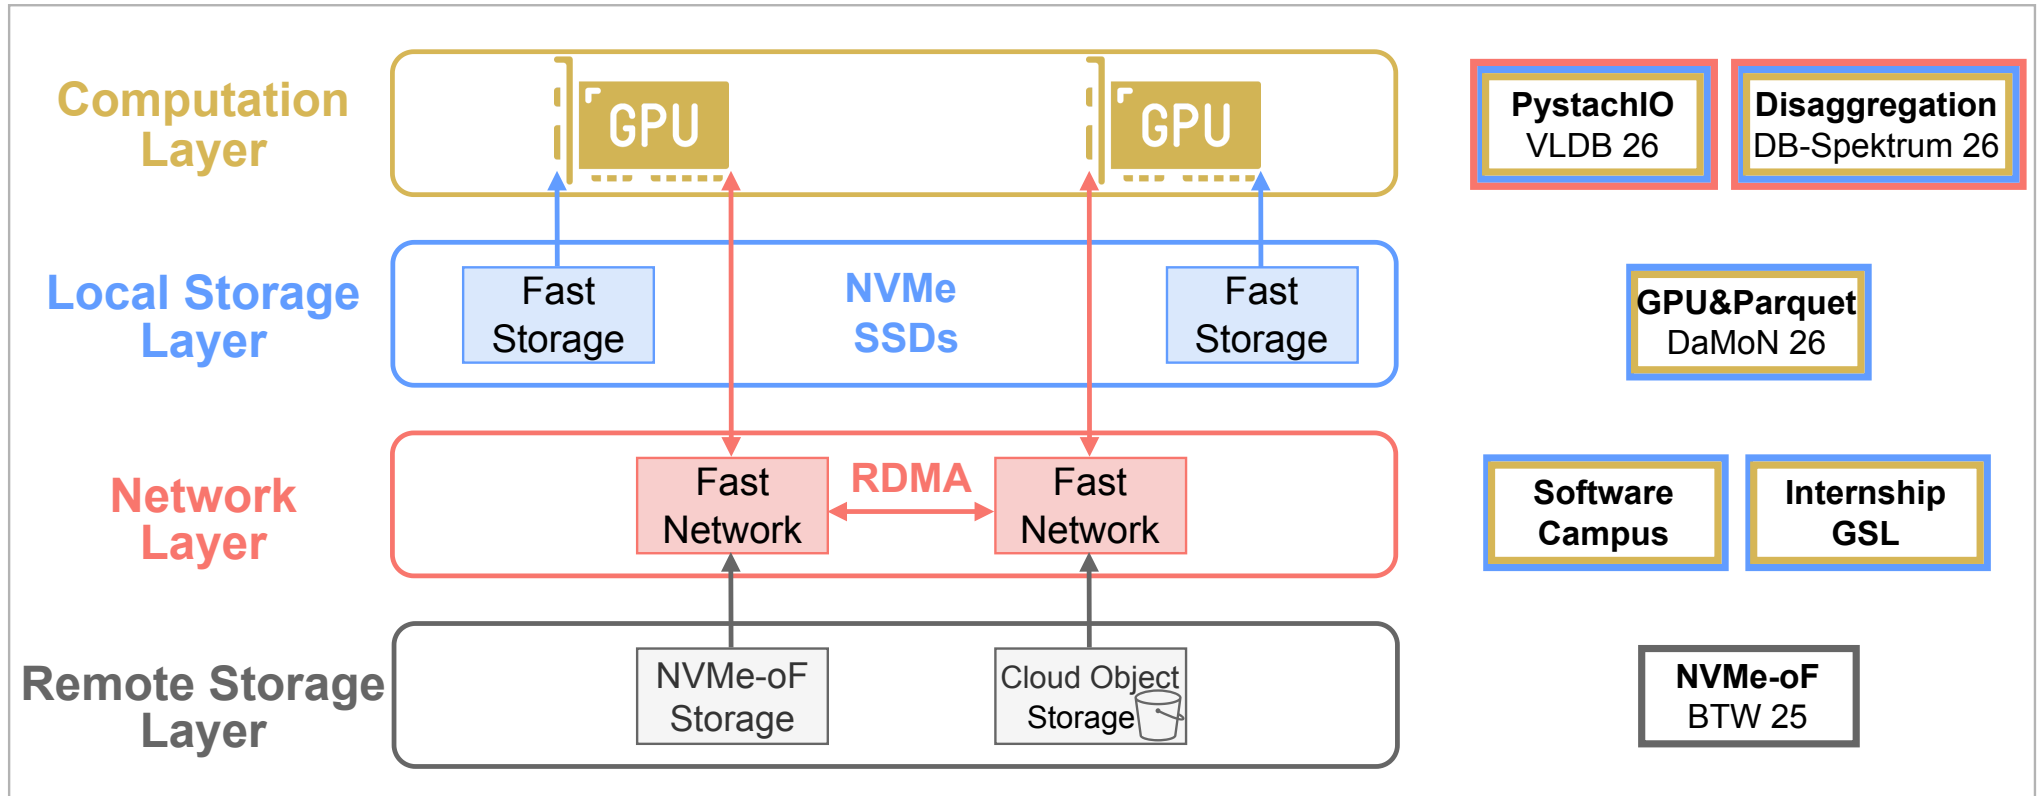
More Than Research

-  **NVIDIA** 
-  **ARROW**  **Parquet**  **APACHE DATAFUSION™**
-  **snowflake**  **software campus** in the next two years

More Than Research

-  **NVIDIA** 
-  **ARROW**  **Parquet**  **APACHE DATAFUSION™**
-  **snowflake**  **software campus** in the next two years
-  **Microsoft**  in the next three months

My Status



- SS 2026: ADMS
- SS 2025: ADMS, Best Exercise 🏆
- SS 2024: ADMS
- WS 2023: SDMS

Thanks For Your Attention!

Acknowledgement

These slides are inspired by the following works. Many thanks to the authors:

- **DBs+GPUs: Where are we and what's next?**, Bowen Wu, Systems Group ETHz
- **G-ALP: Rethinking Light-weight Encodings for GPUs** @ DaMoN25, Sven Hepkema, ex-CWI DB & Systems Group ETHz
- **GPU Databases – The New Modality of Data Analytics**, Xiangyao Yu, University of Wisconsin-Madison
- **SIGMOD 26 Keynote: A New Golden Era for Data Management**, Gustavo Alonso, Systems Group ETHz