


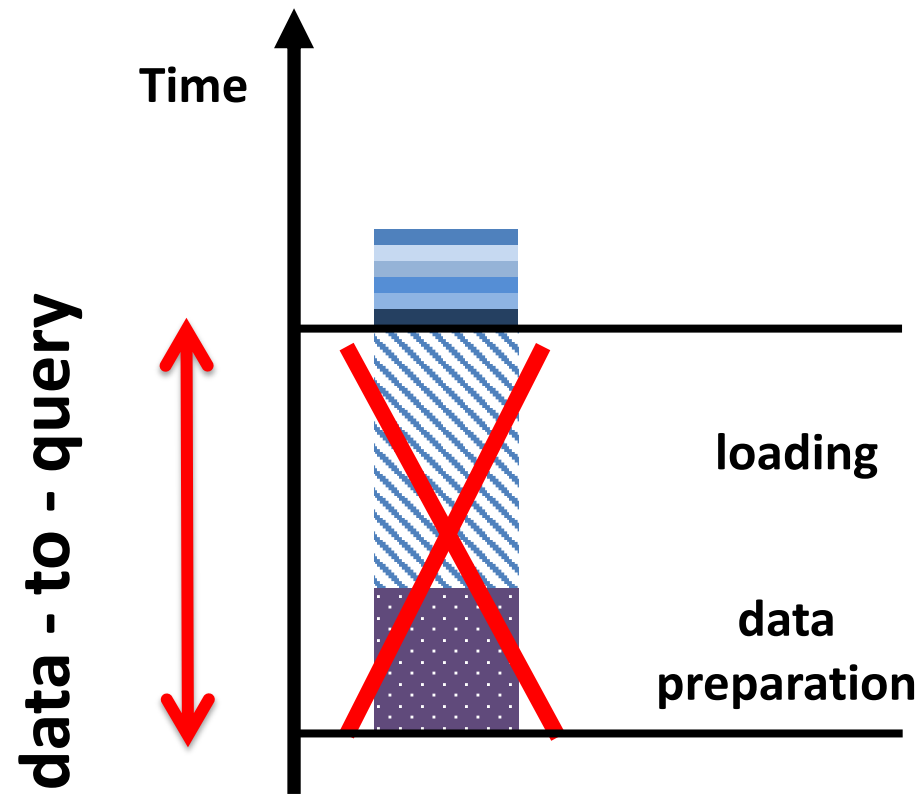
**From
LotsOfCode
to
NoCode**

Anastasia Ailamaki
EPFL



NoDB: minimize data-to-query time

SIGMOD 2022 Test-of-Time award

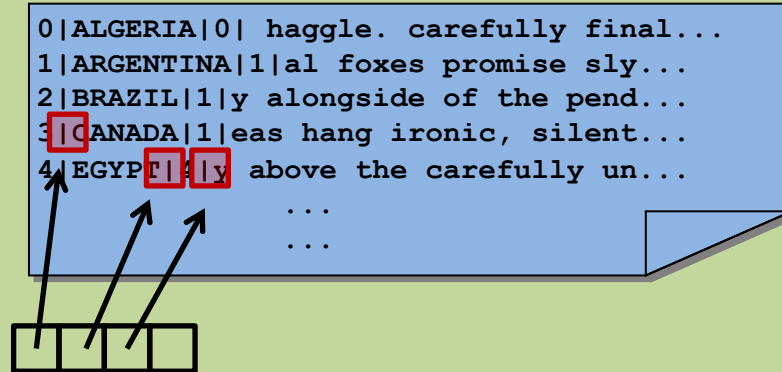


Adapt data on request

NoDB idea:

Decouple functionality from performance

Positional maps



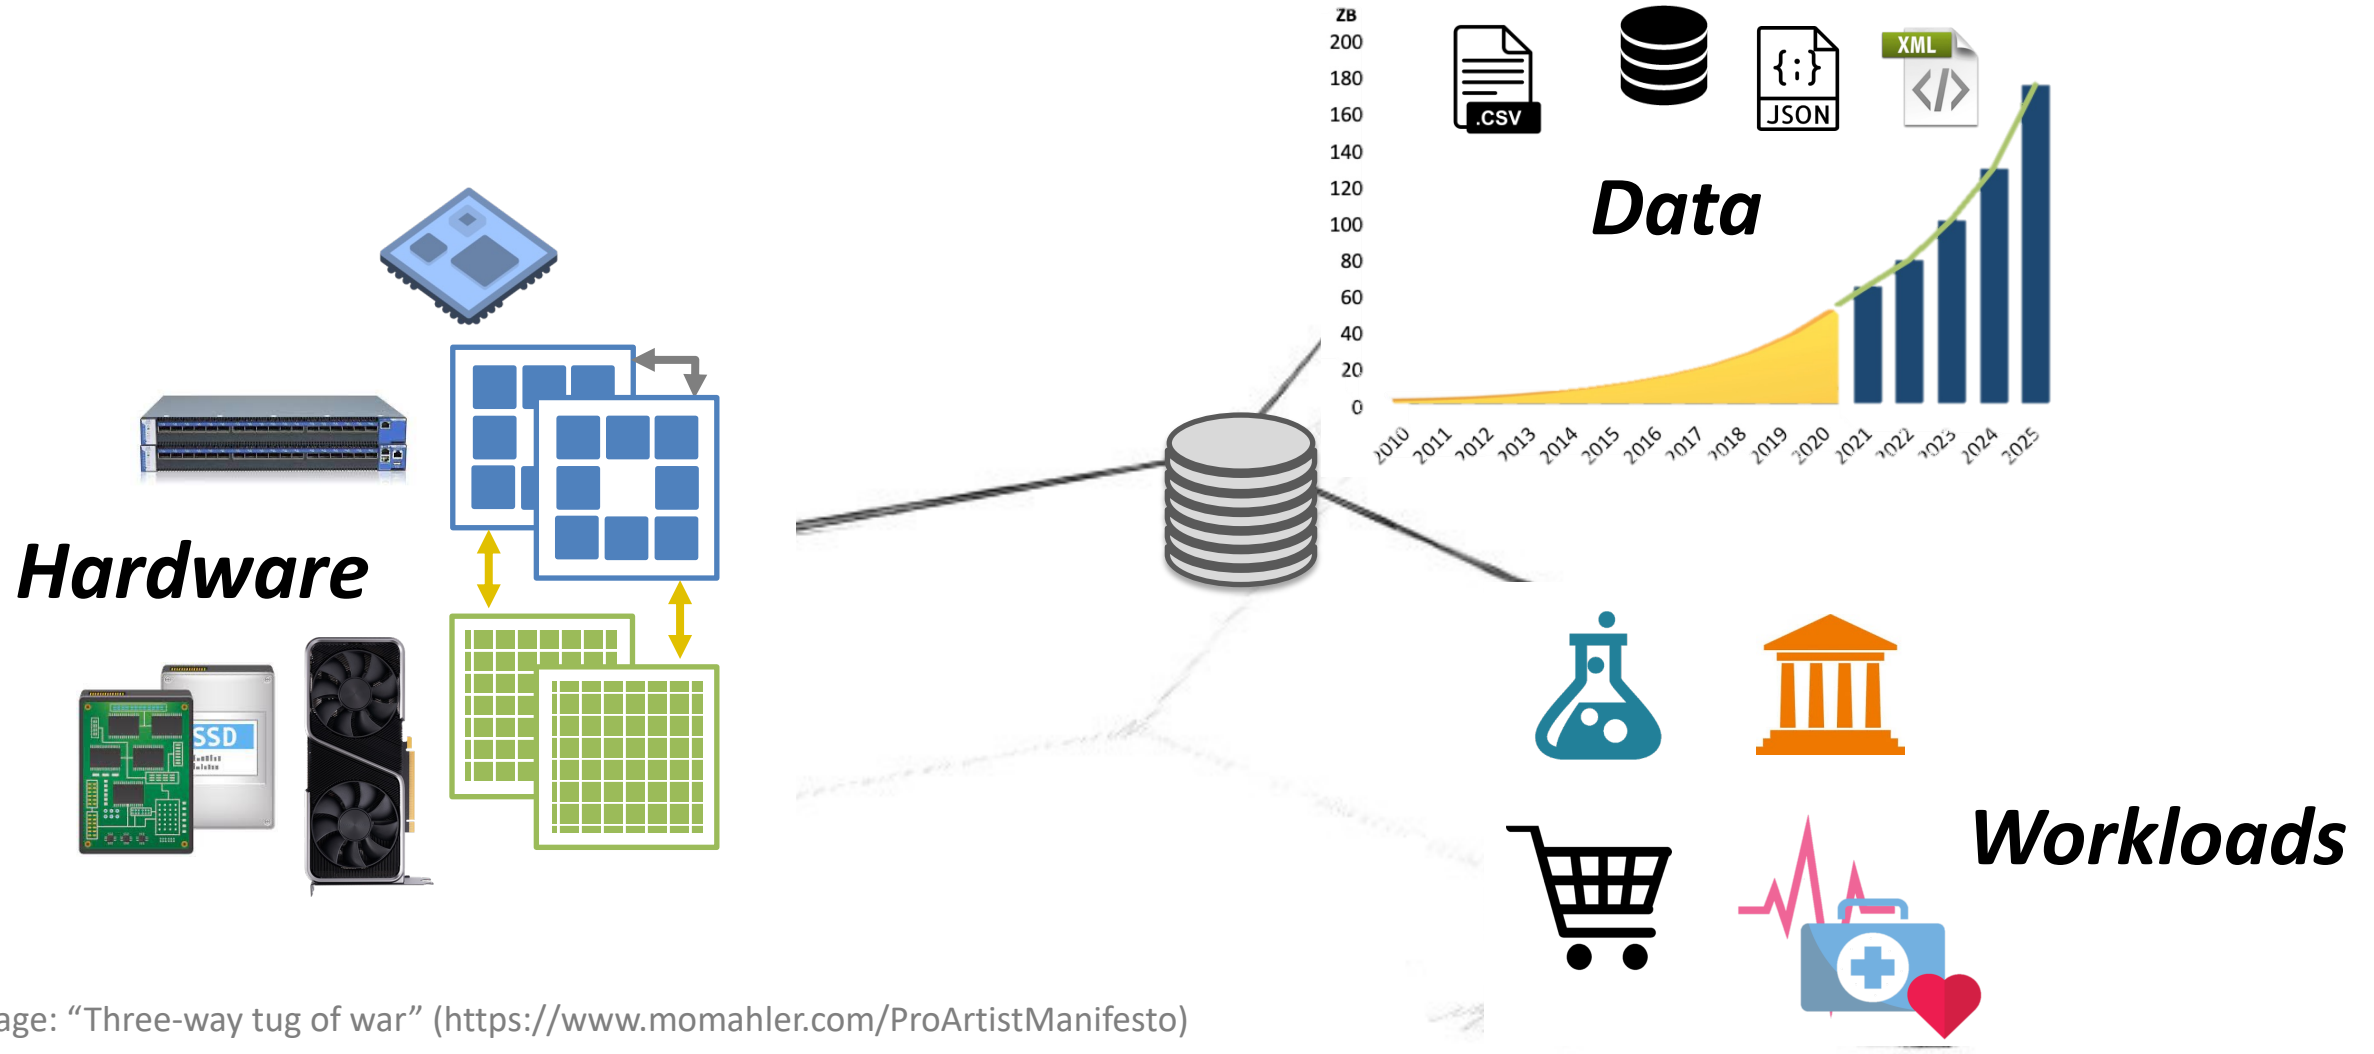
Caching

User does not need to control when, what, how or where data is cached

!= Classical Data Loading

*Data layout
best suited for ...
raw file format
user queries*

An incessantly evolving landscape



Data management faces its most critical challenges

biological disease signatures

coupling

clinical measurements with *validated biomarkers*

Example: Alzheimer's disease



Challenge:

Real-time integration of heterogeneous data

clinical+genetic+imaging data → signature

Patients (CSV)

id	Protein: AACT	Age	Phenotype	...
1	1.4	45	Trauma	...
2	2	55	Chronic Symptoms	...
3	0.2	56

Brain_GrayMatter (Binary)

	0	1	...	n
0	0.45	0.75	...	0.1
1	0.33	0.3	...	0.38
...
m	0.12	0	...	0.47



signature:

age > 50

AND

amygdala.Vol > 0.3

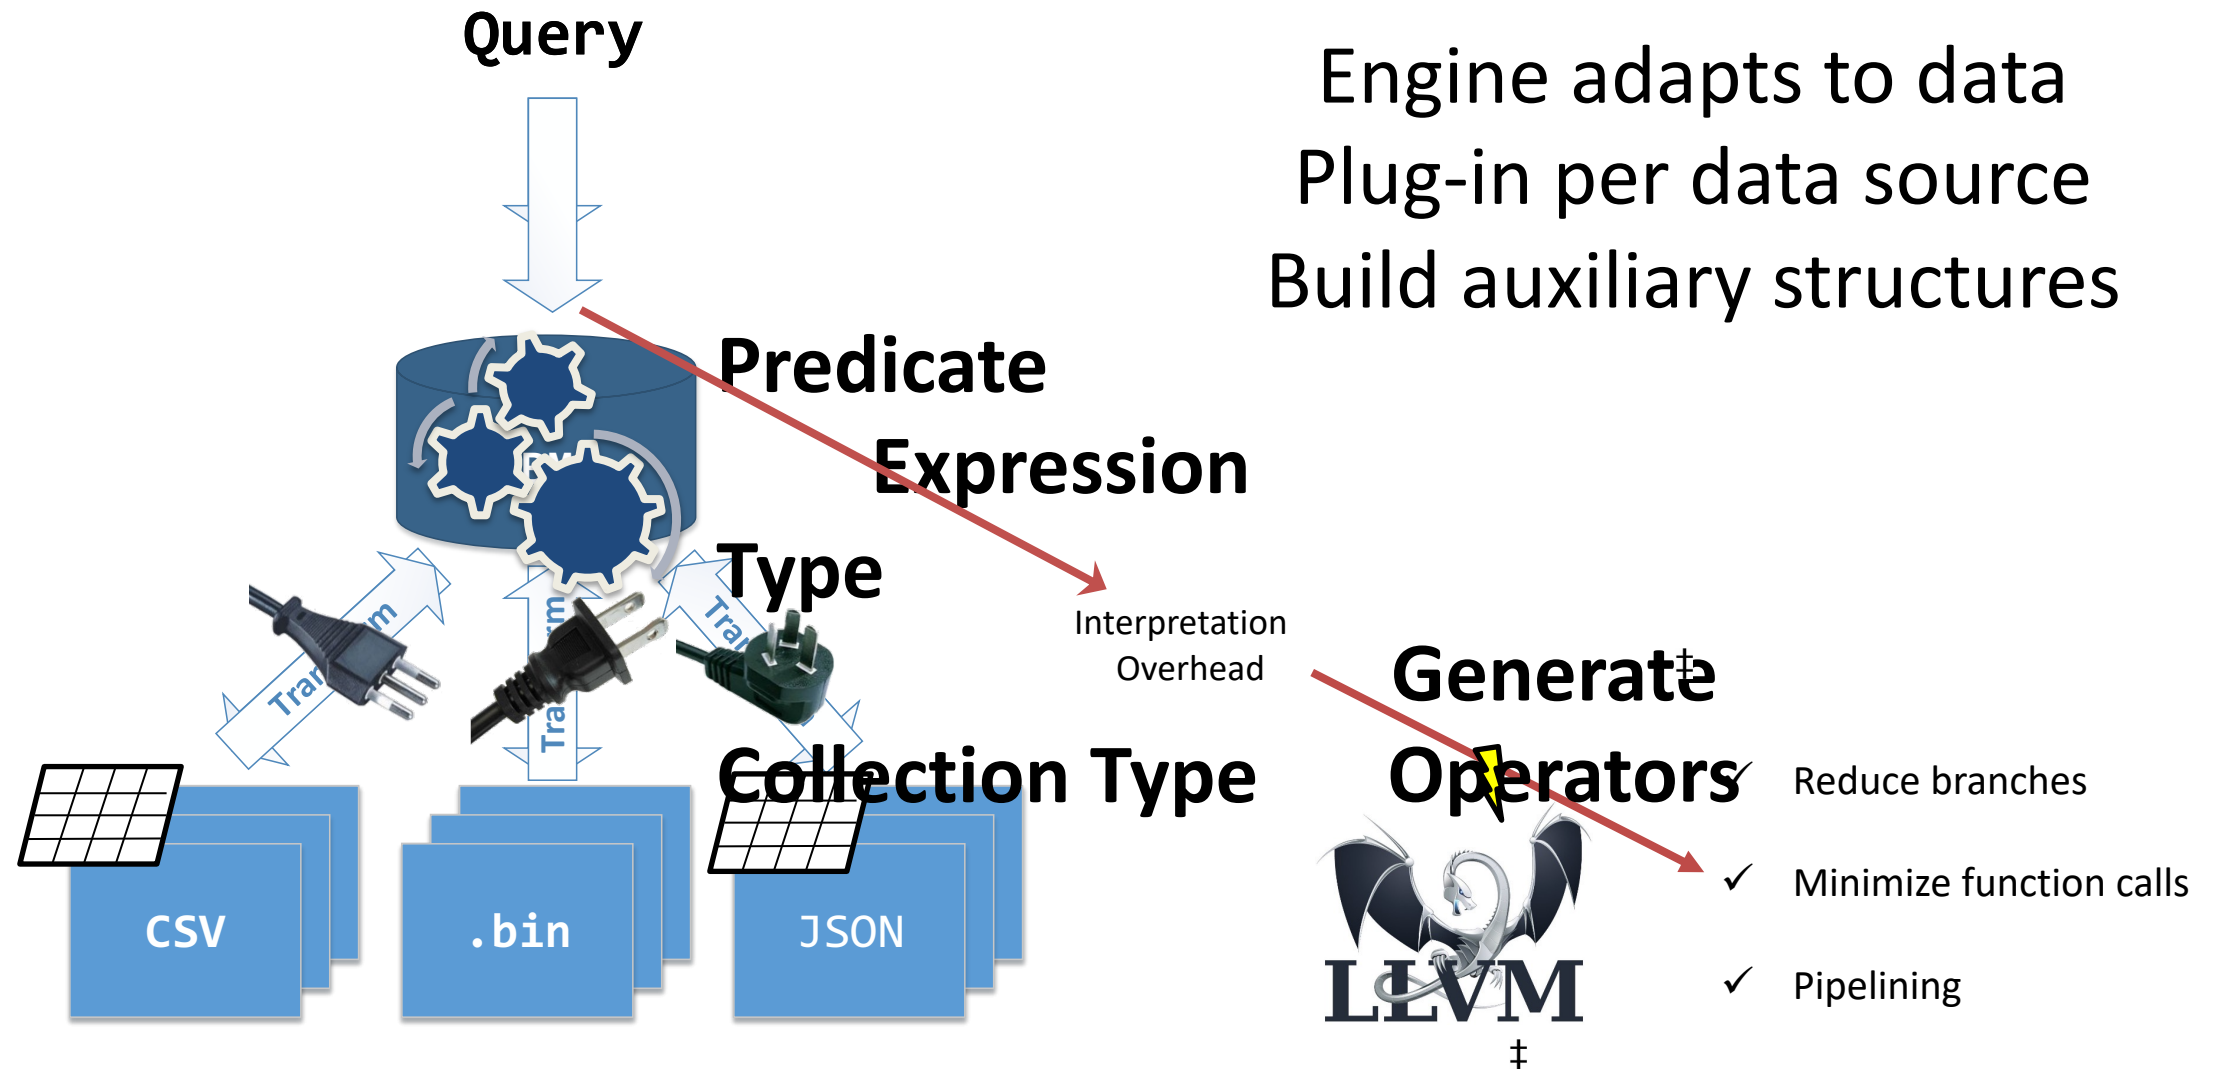
AND

AACT < 1

BrainRegions (JSON)

```
[{"id": 1,
  "amygdala": {"X":15, "Y":20, "Vol": 0.5},
  "hippocampus": {"X":17, "Y":10, "Vol":0.2}},
{"id": 2, ...},
{"id": 3, ...}]
```

From LotsOfCode to NoCode

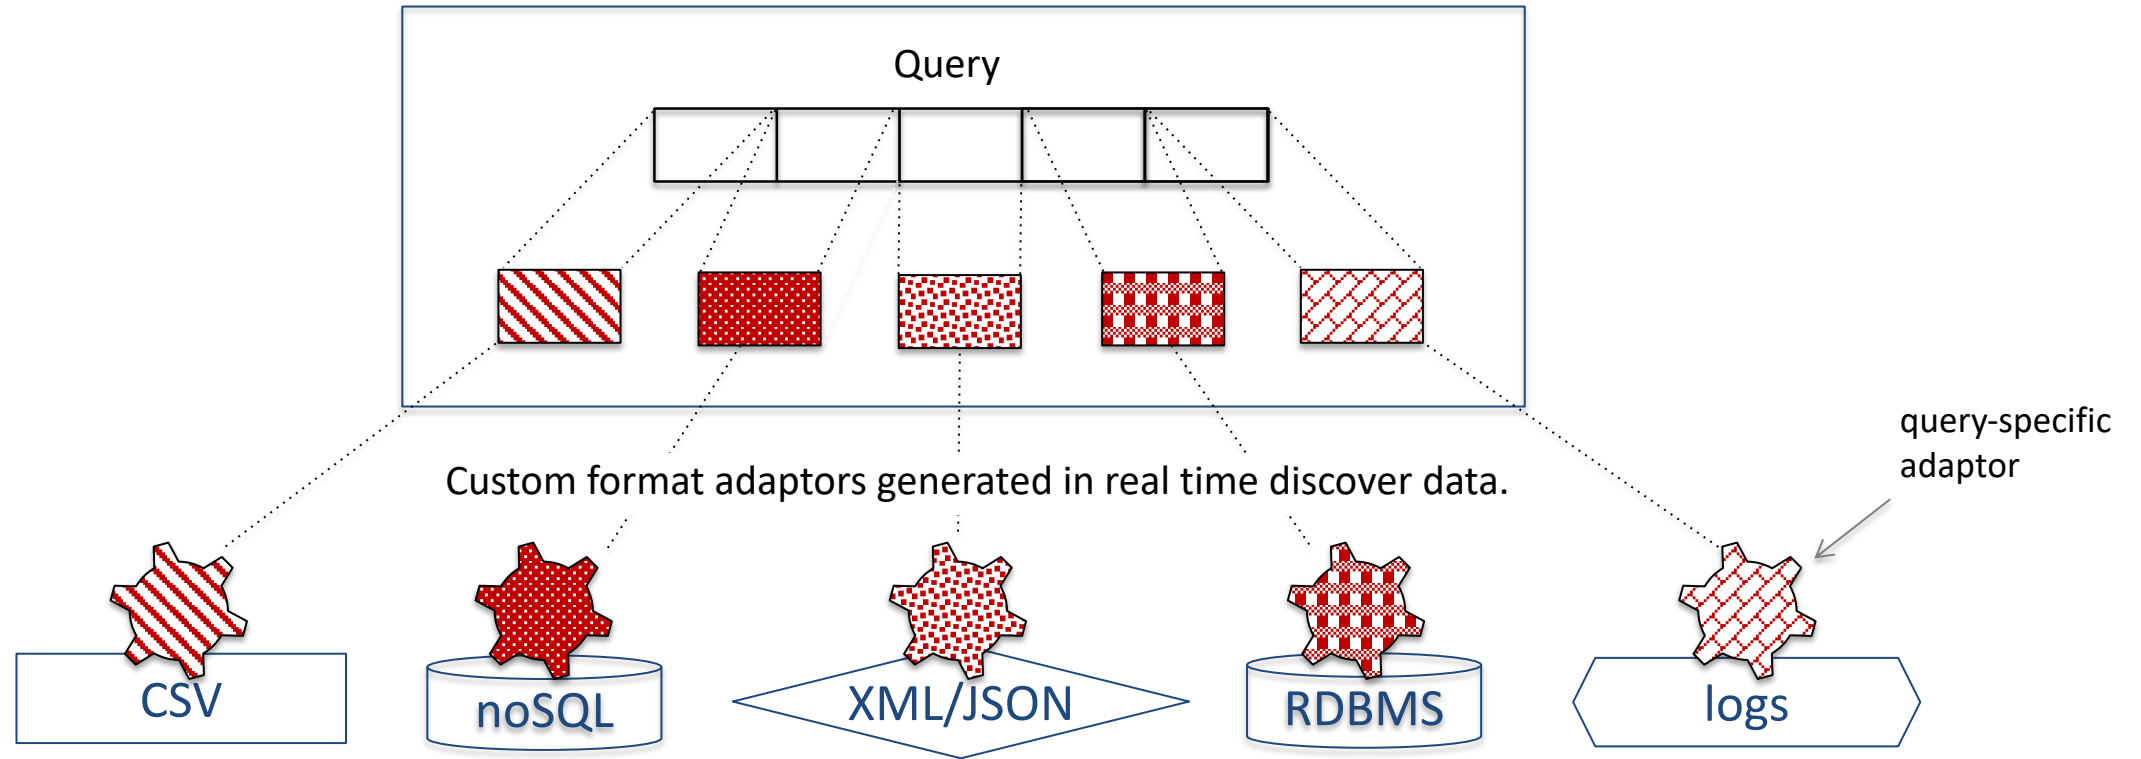


Codegen operators, continuously adapting engine

RAW: a *single* engine for all data

RAW
Just ask.

RAW Query is automatically split up for each data source.





Data is integrated transparently and on-demand.

**Users think of all of their data as a unified database,
without preparation**

Zero-cost (virtual) database

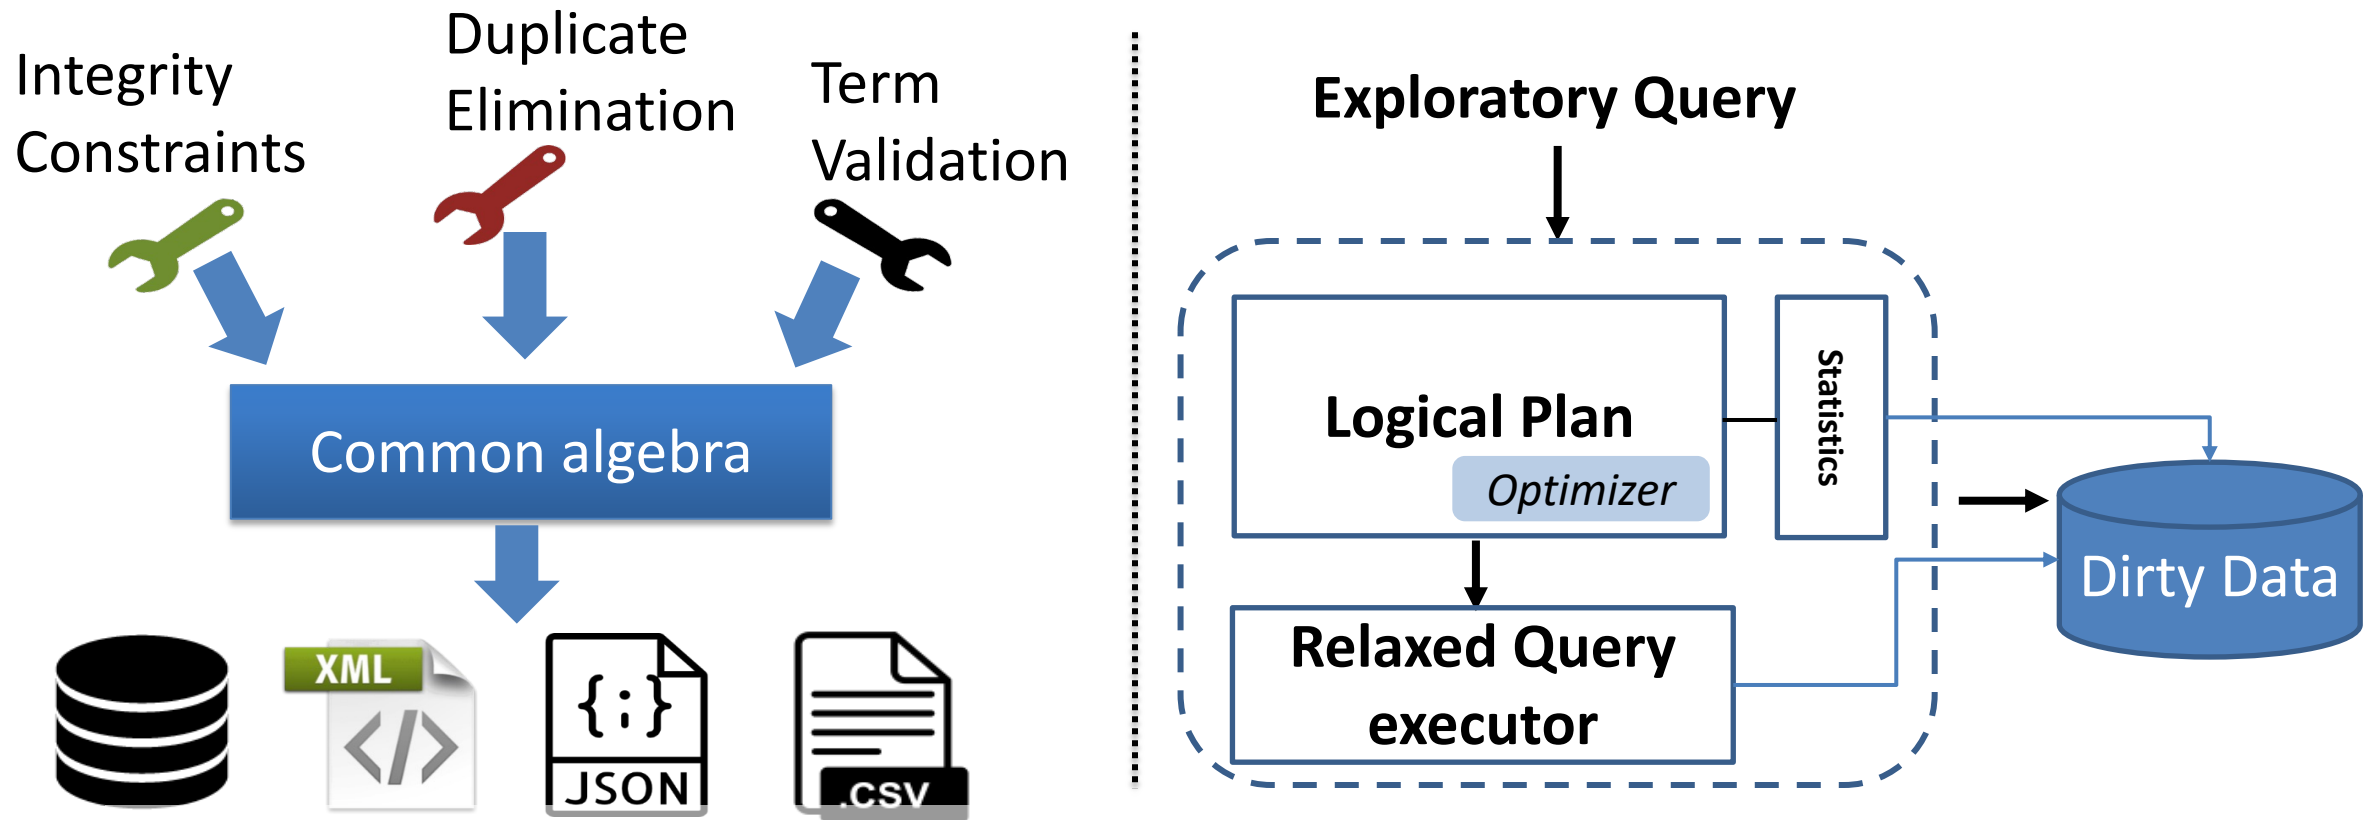
RAW
Just ask.



RAW is now a platform for trusted, live & secure data delivery


Efficient data veracity

Correct ALL errors on ALL data: costly and unnecessary



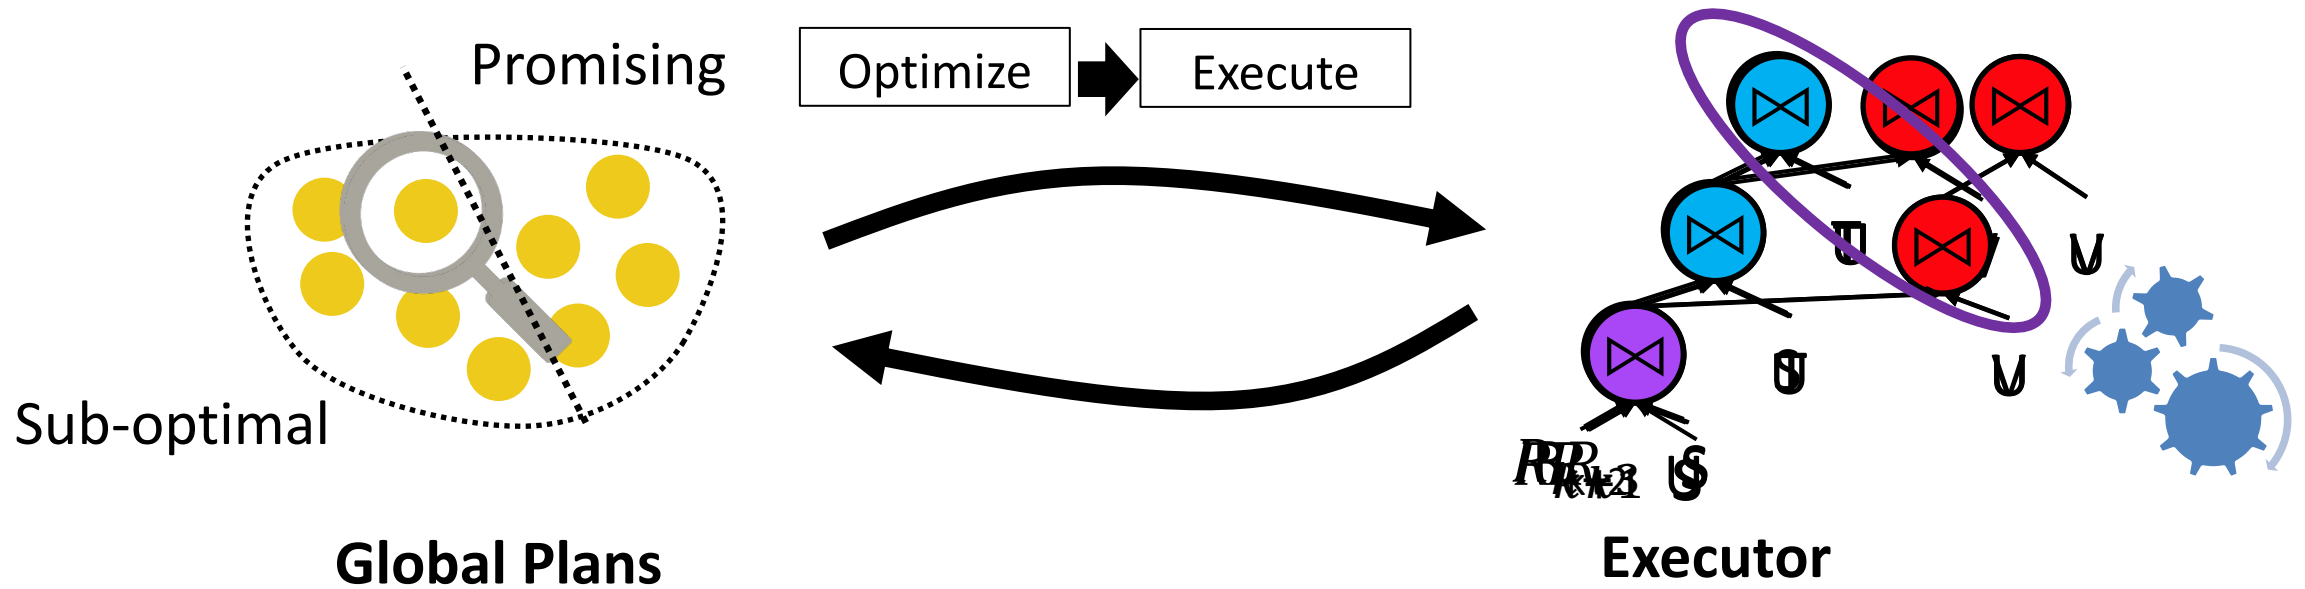
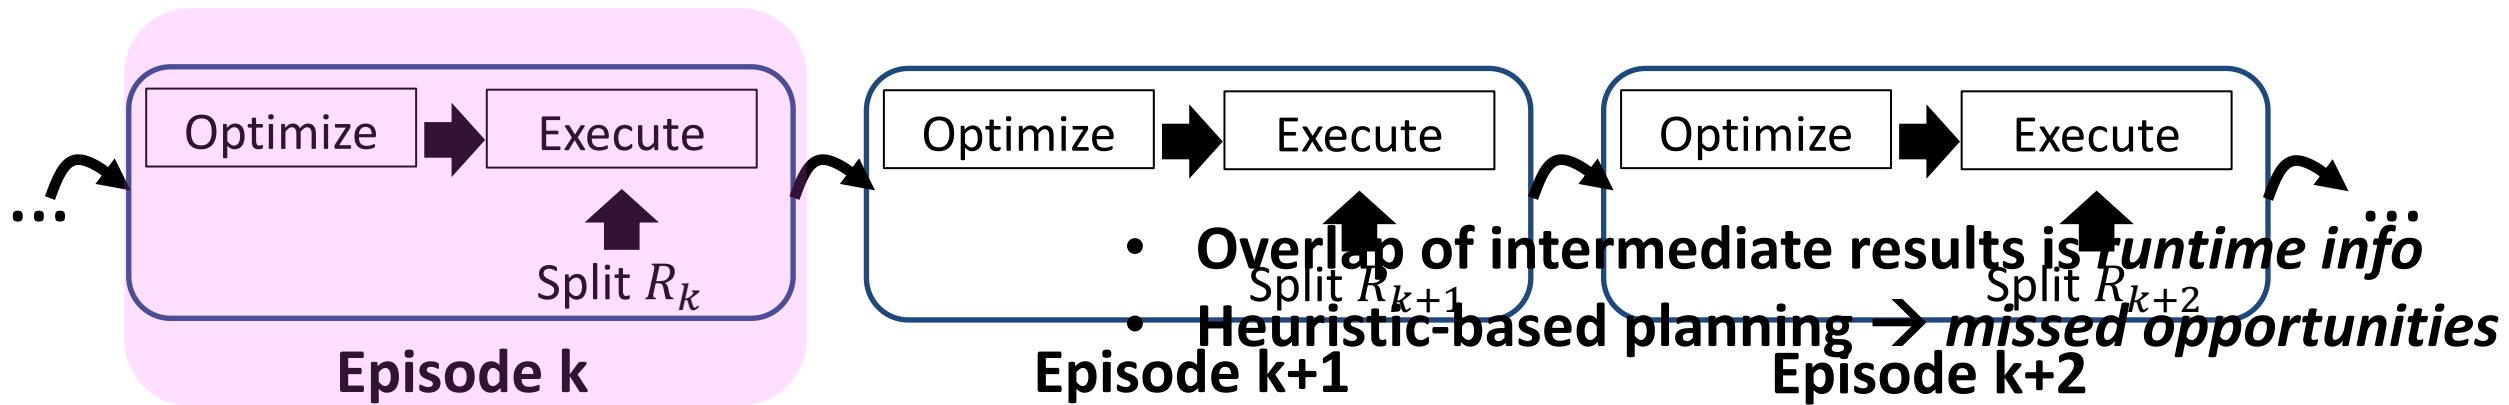
Clean useful data *adaptively* during analysis

Operations vs Analytics





Ever-increasing number of concurrent queries
Data freshness bounded by ETL latency

~10K concurrent tasks



Real-time learning helps reduce work

Hybrid Transactional and Analytical Processing

- Transactions: task-parallel 
 - High rate of short-lived processes
 - Mostly “point accesses” (high data access locality)
- Analytics: data-parallel 
 - Few, but long-running queries



Strong consistency is an invariant

Workload Isolation or Fresh Data?


TRANSACTIONS

ANALYTICS


Fresh Data

Collocated workloads fight for resources


Isolated




Hybrid-Access



Elastic-Compute



Collocated



Interference → better data freshness

No interference → better performance

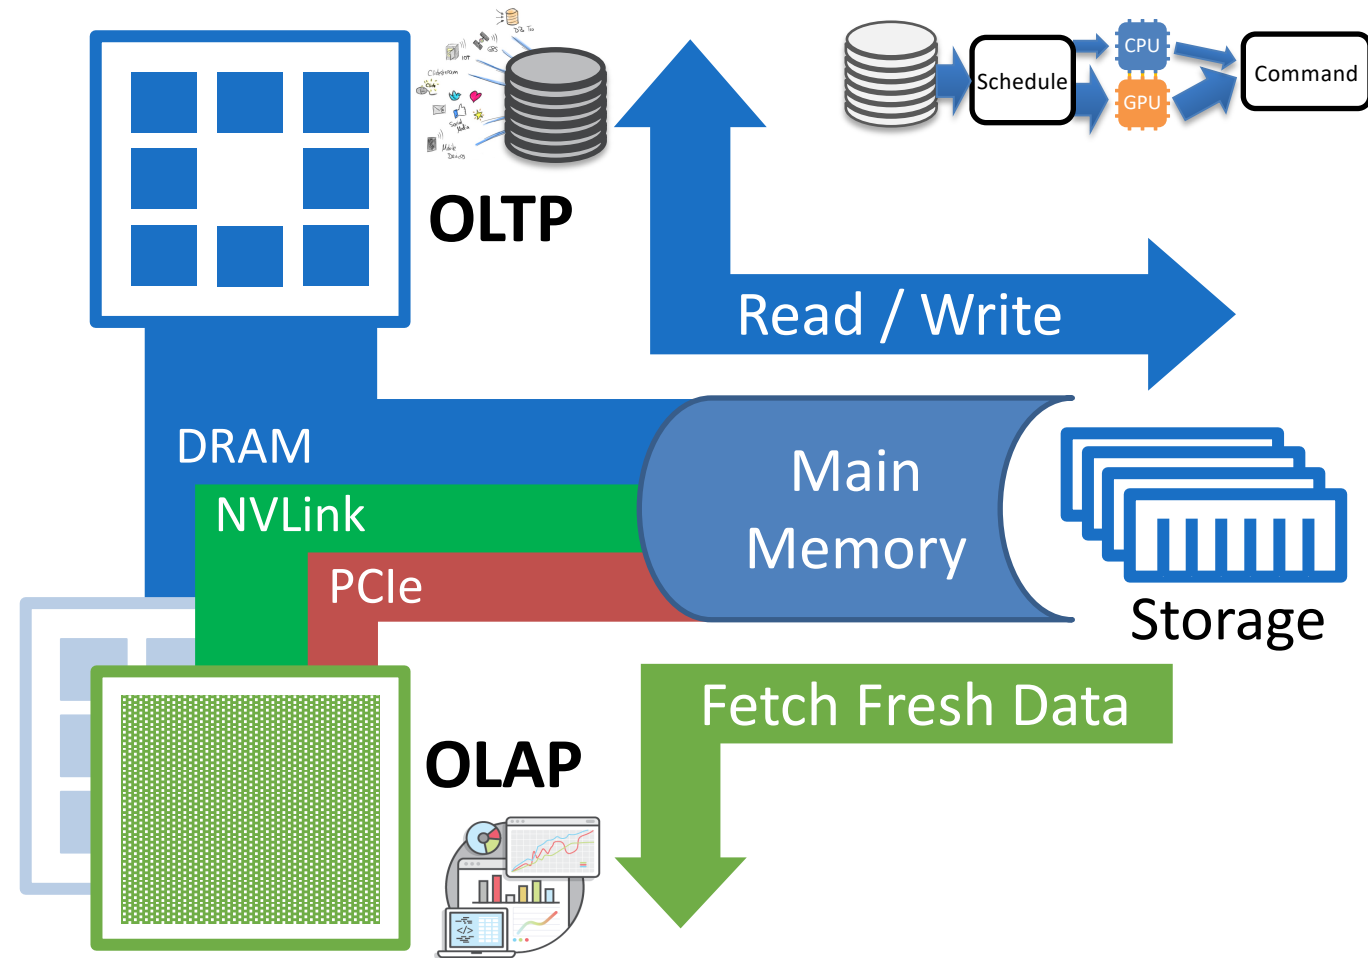
Viable hybrid alternatives

GPU acceleration for HTAP workloads

Transactions store fresh data
on CPU Memory

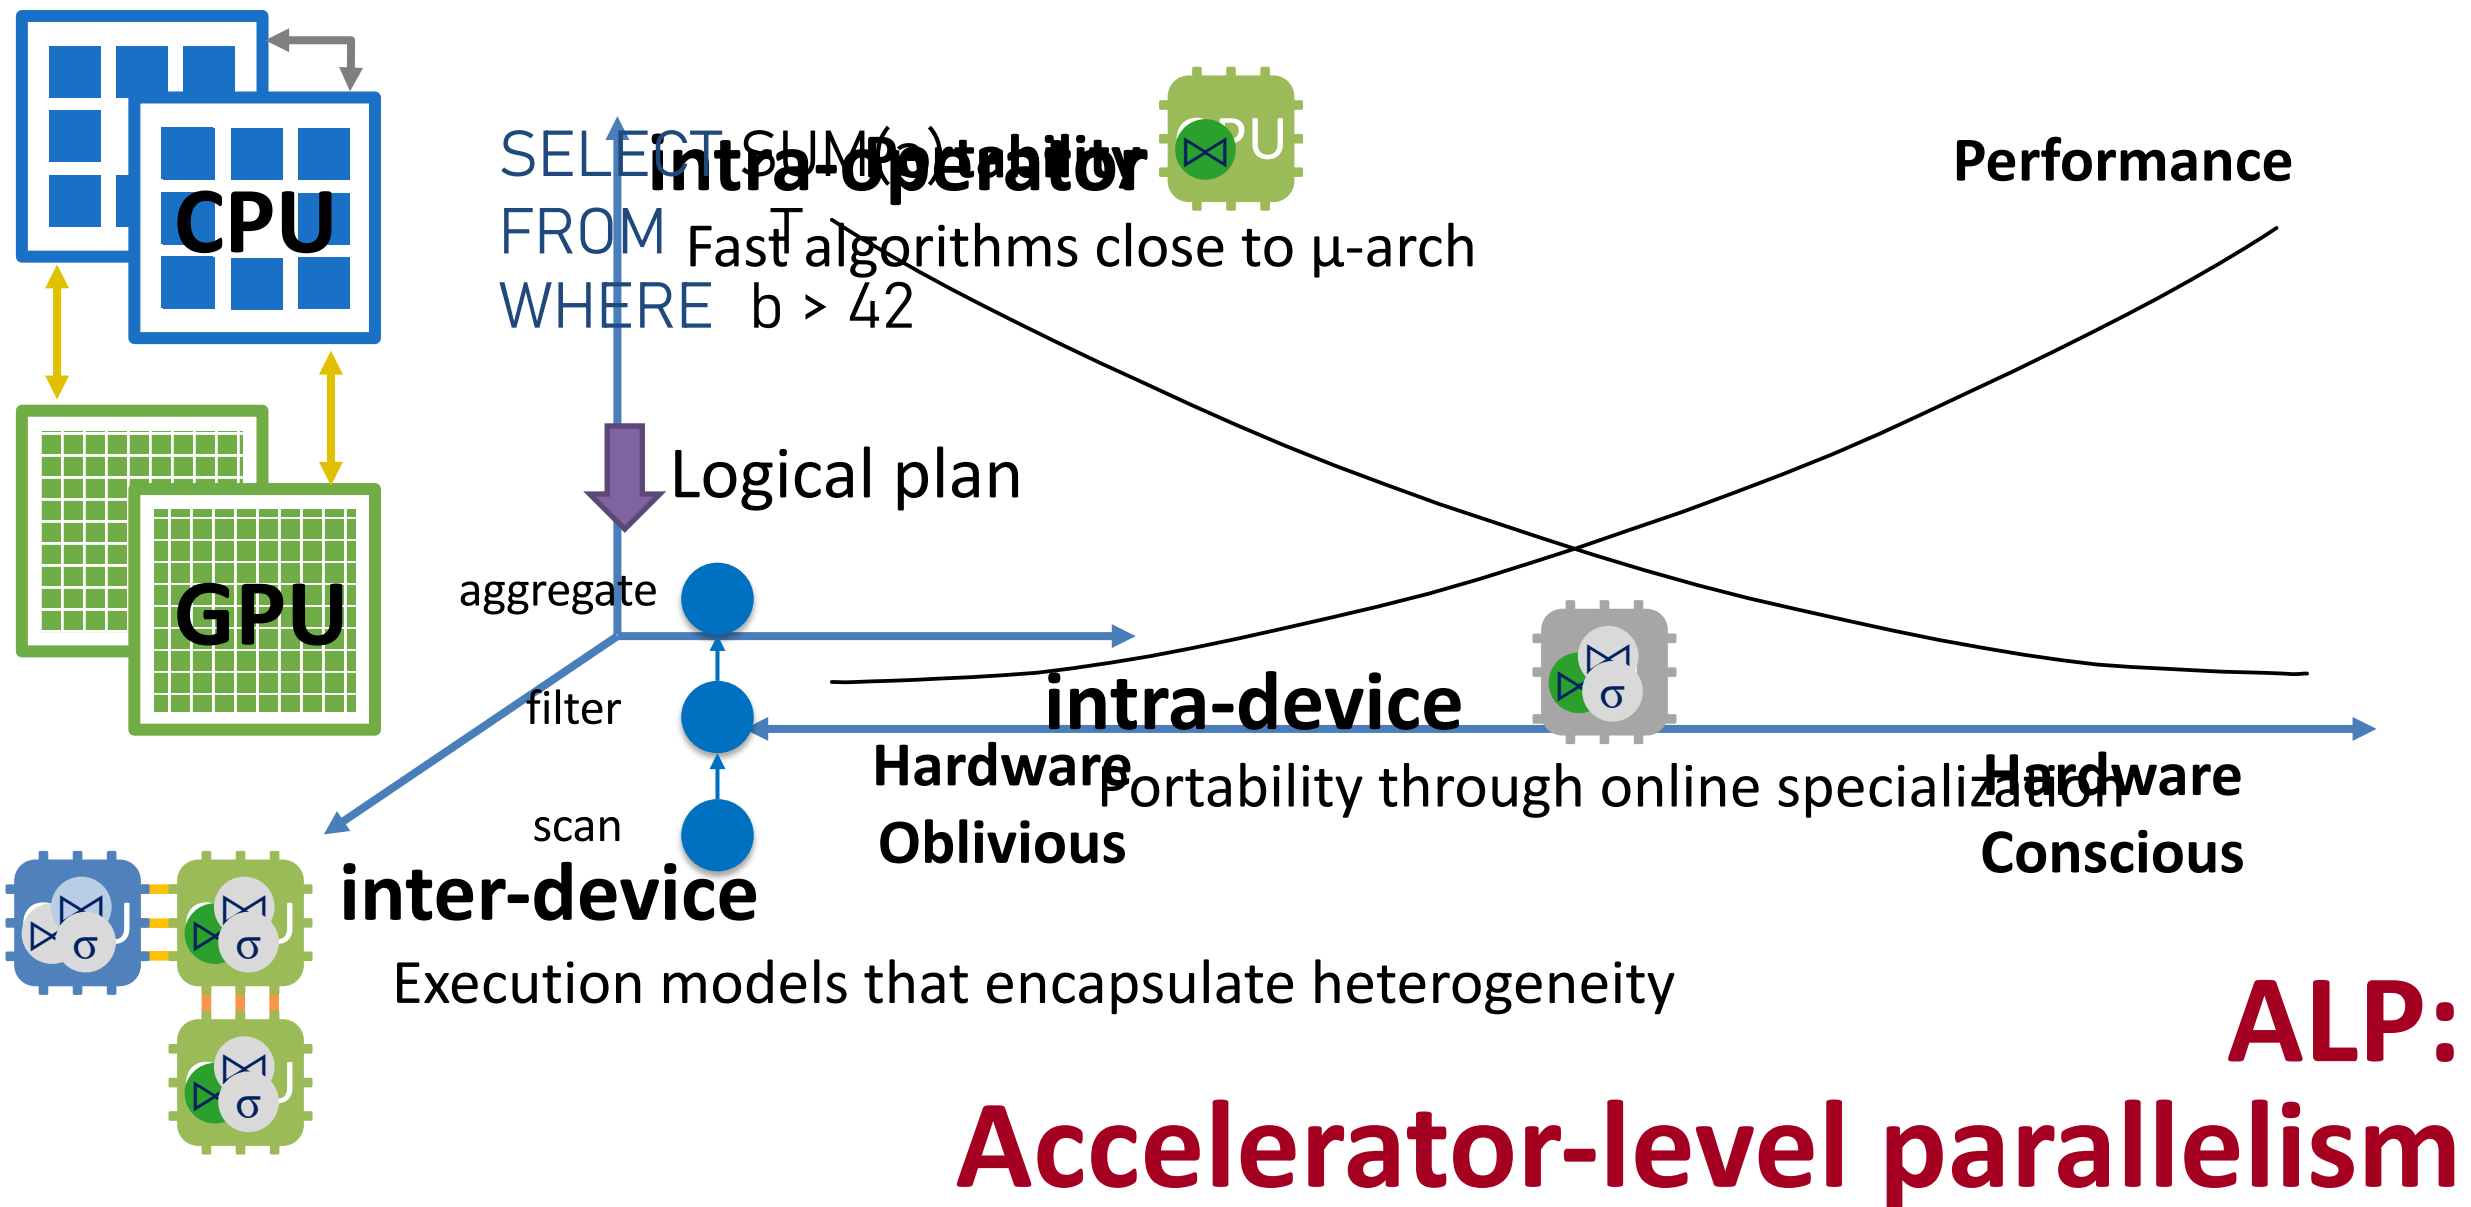
Data access protected by
concurrency control

Analytics access fresh data
through interconnect



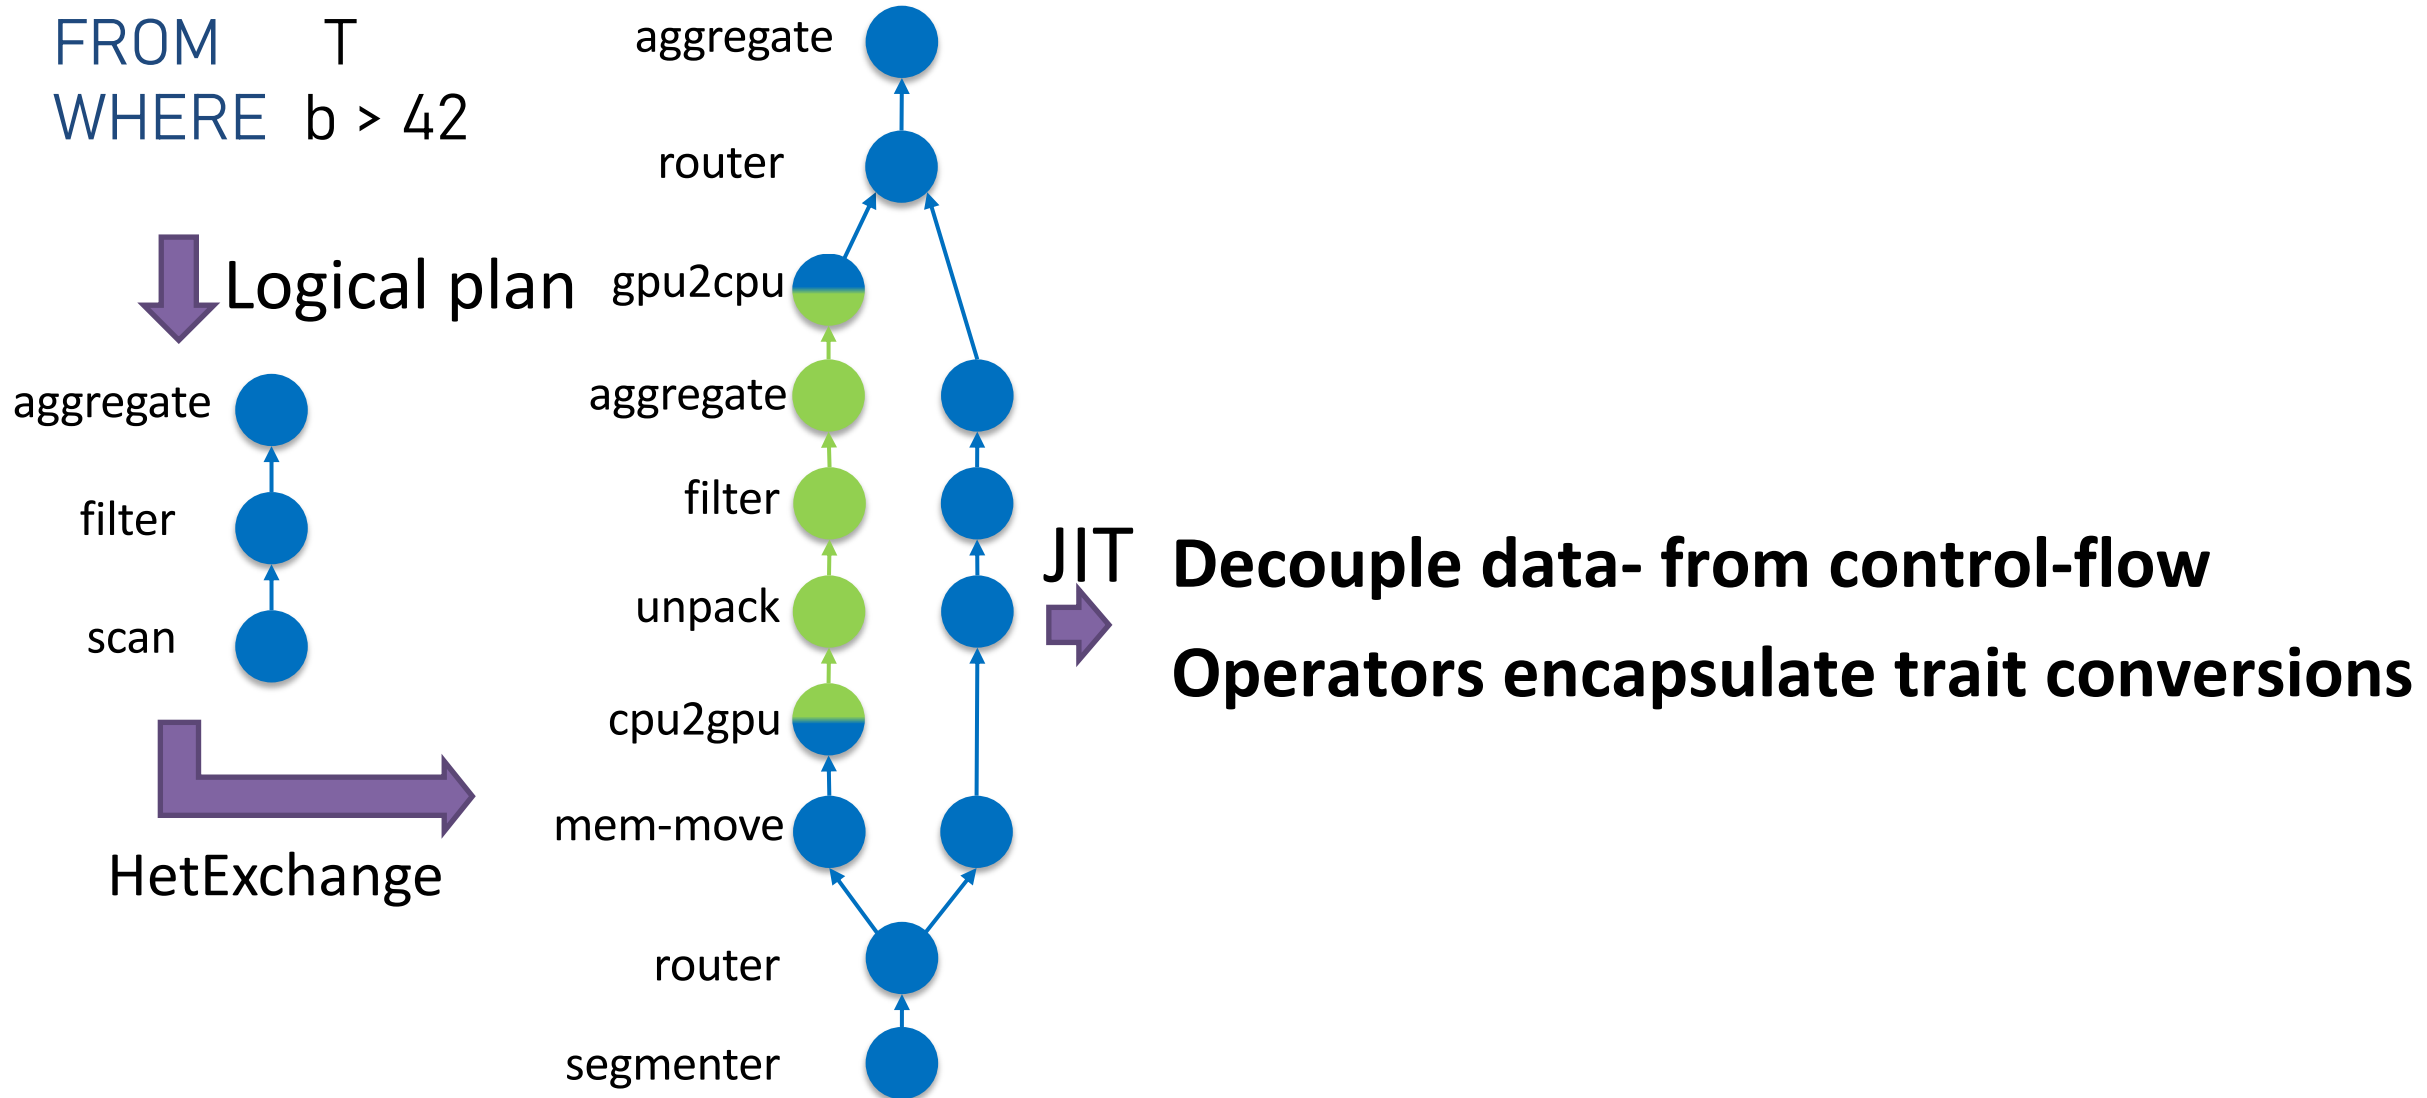
Interconnect speed is critical

Device-conscious processing without regrets

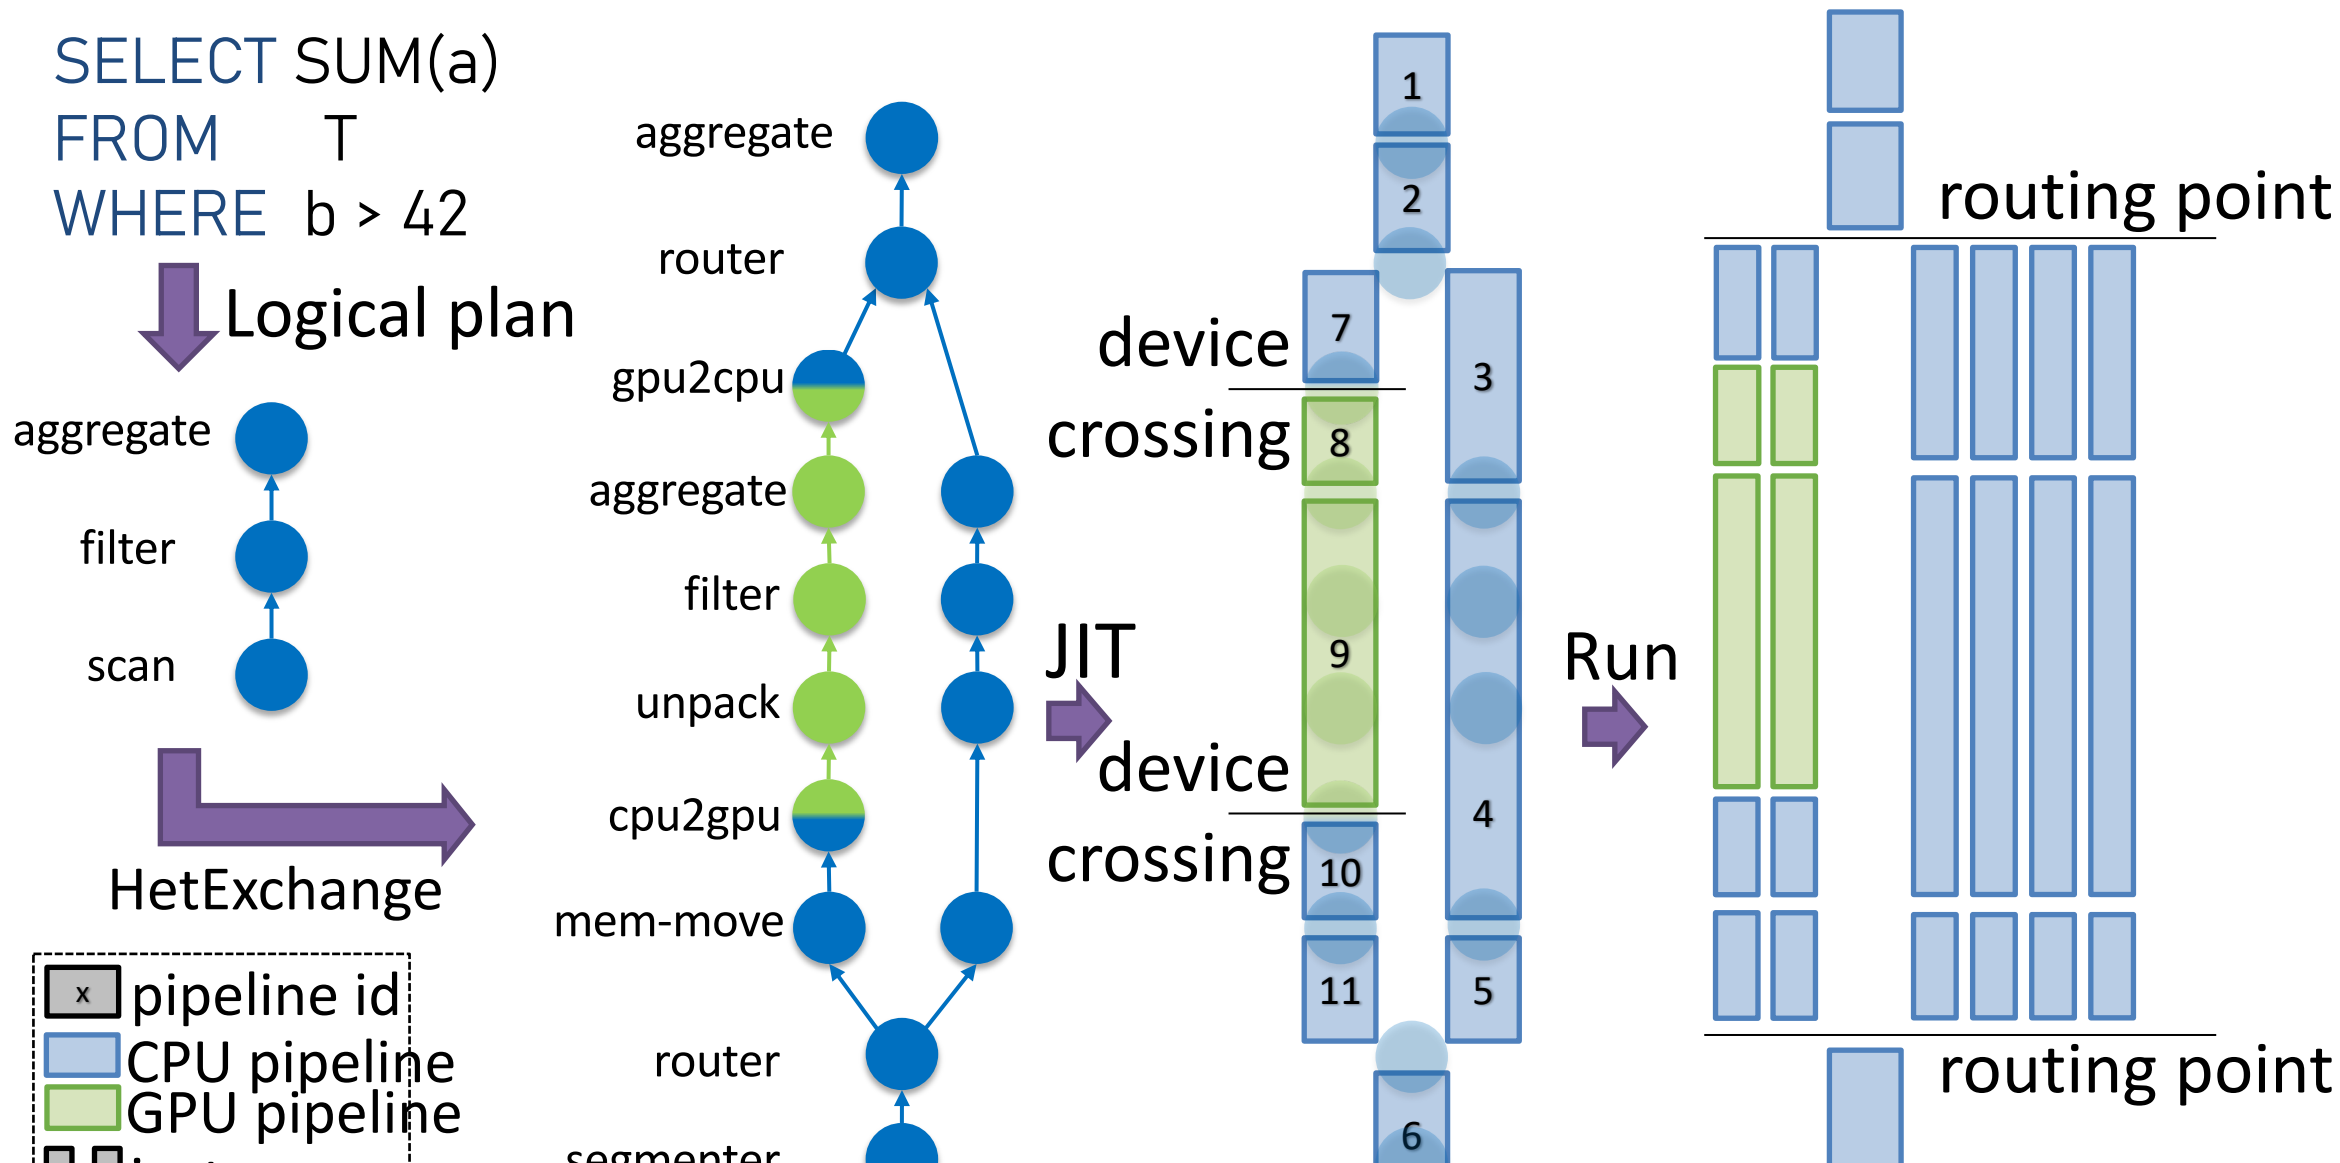


HetExchange: Heterogeneity-aware plans

```
SELECT SUM(a)
FROM T
WHERE b > 42
```




HetExchange in a JITed engine



Maximize inter-device utilization

Rack-scale analytics can use scaleup solutions

- Similar intra-/inter-server interconnect bandwidth
- Local memories and NUMA effects across devices
- CPU-GPU: Capacity-Throughput



Efficient use of heterogeneous interconnected devices

Ever-complex data pipelines

Diverse modern data problems

- IOT, OCR, ML, NLP, Medical, Mathematics etc...



Commercial AI/ML

Database systems catch-up for popular functionality

- Human effort and big delays
- Oblivious to outside workflows



Augmented analytics

Vast resource of libraries

- Authored by domain experts, used by everybody
- Loose library-to-data-sources integration and optimization




Combination of IoT and analytics



Conversational analytics and NLP

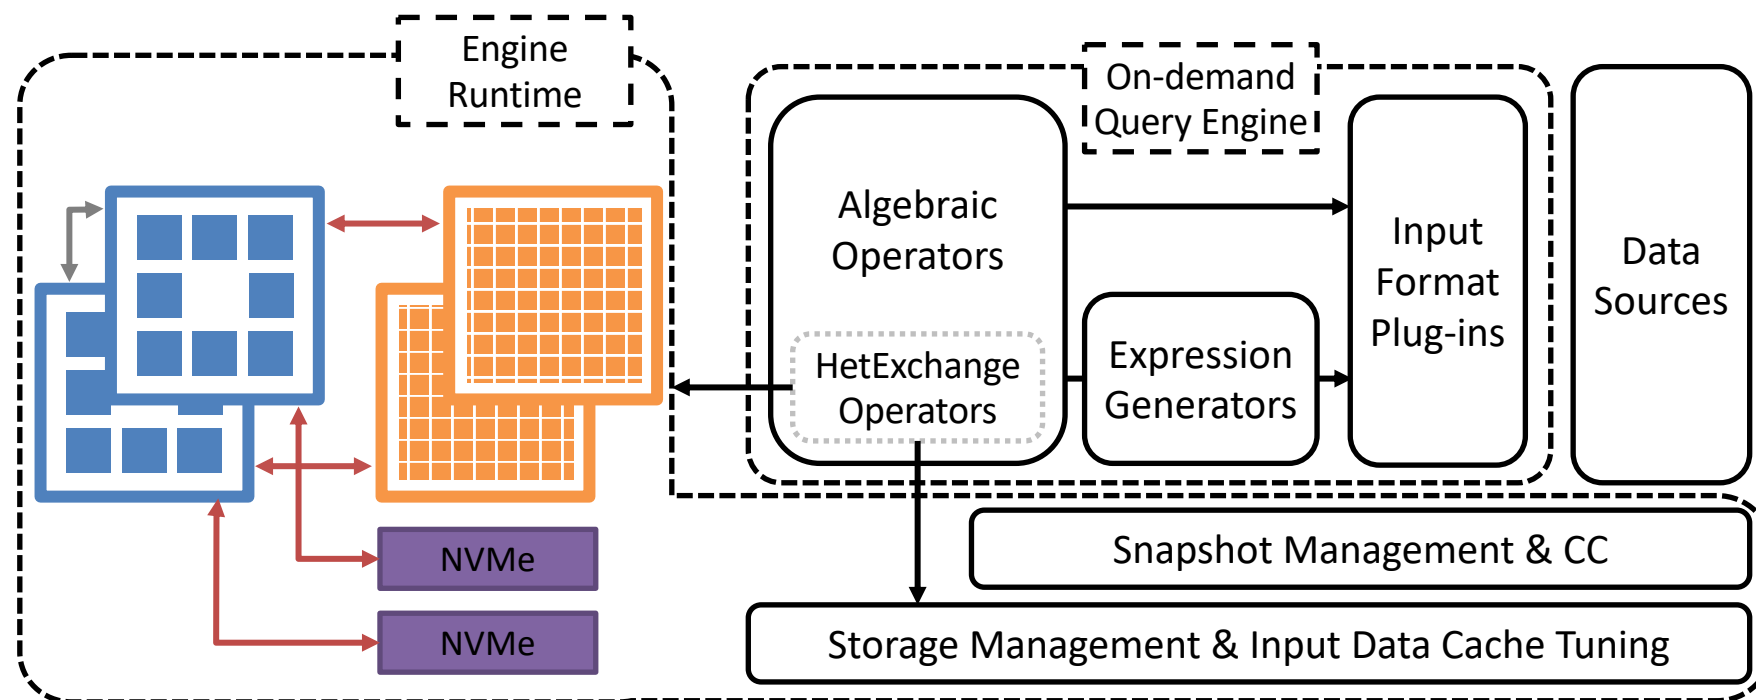
Declarative programming for learning systems

Heterogeneity in every aspect of data pipelines



Not ahead-of-time, JUST-IN-TIME!

Proteus: taming heterogeneity through adaptivity



JIT + no code = fast analytics on fresh data

Intelligent

Real-time

Systems

*Street-smart engines incorporate change as a design principle,
and react to surprises while learning.*

