

Data processing in energy efficient data centers

Gustavo Alonso

Systems Group

Department of Computer Science

ETH Zurich, Switzerland

Most technology forecasting is like a ritual rain dance.

It has no effect whatsoever on the weather that follows
but it makes those who engage in it feel like they are in
control.

Most discussions on the directions of new technology are
directed at improving the dancing, not the weather

Paraphrasing R.L. Ackoff, Wharton Business School

In the real world




≠



In IT




=



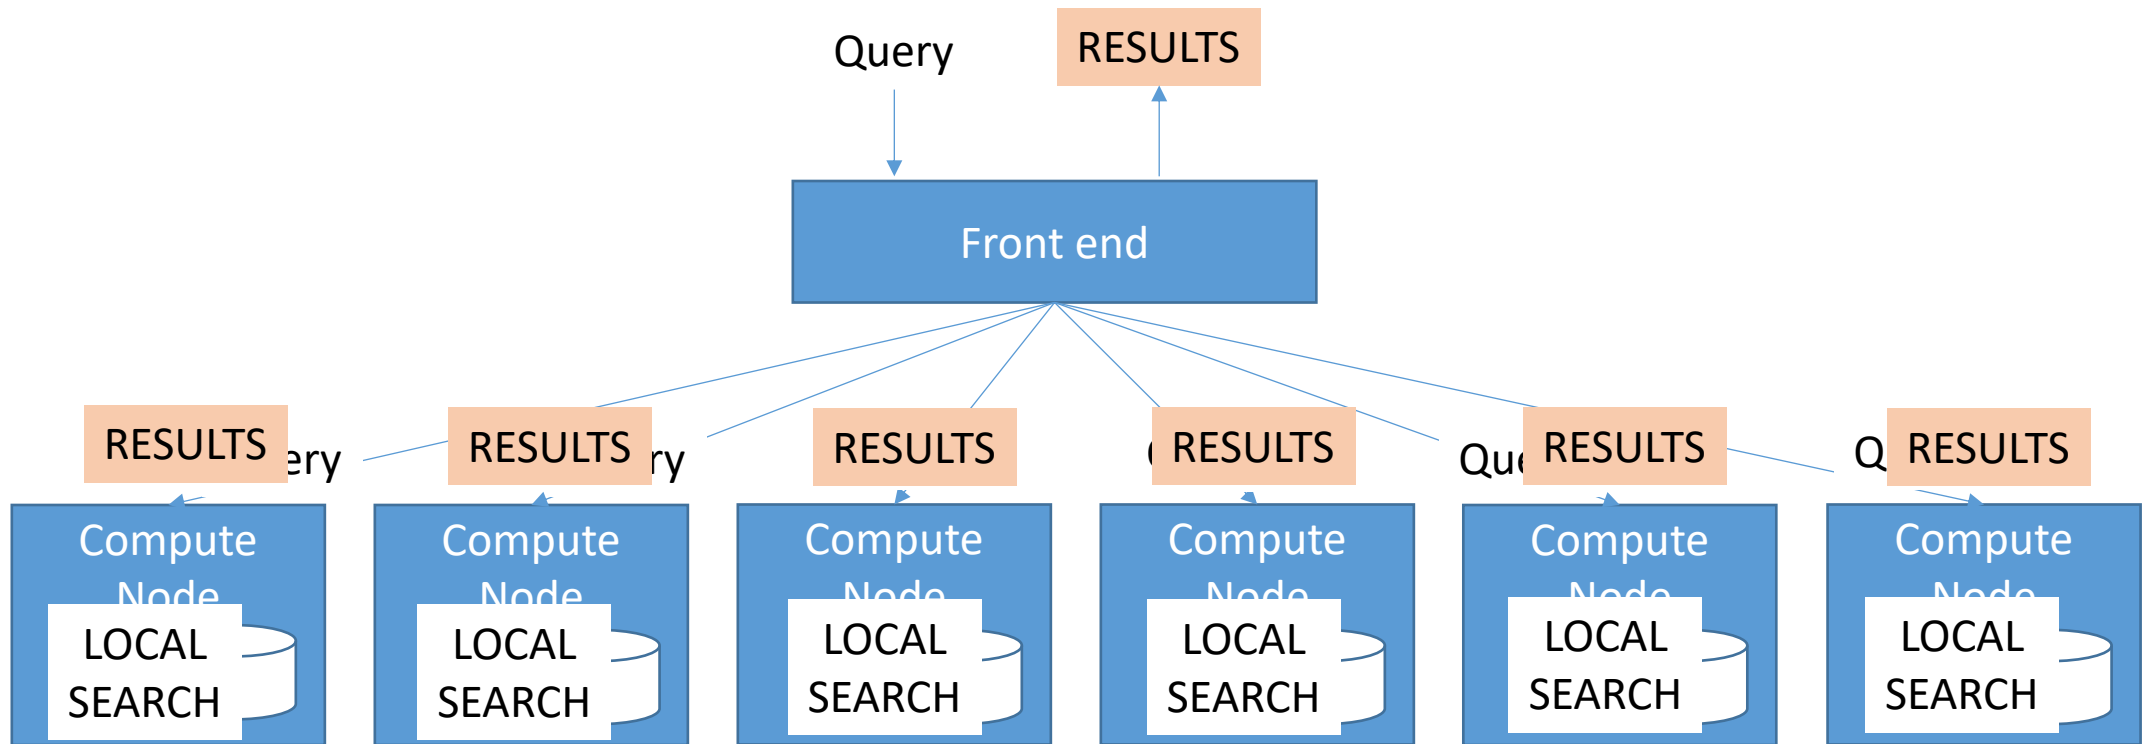
Are our systems efficient?

–


Data processing in the XXI century



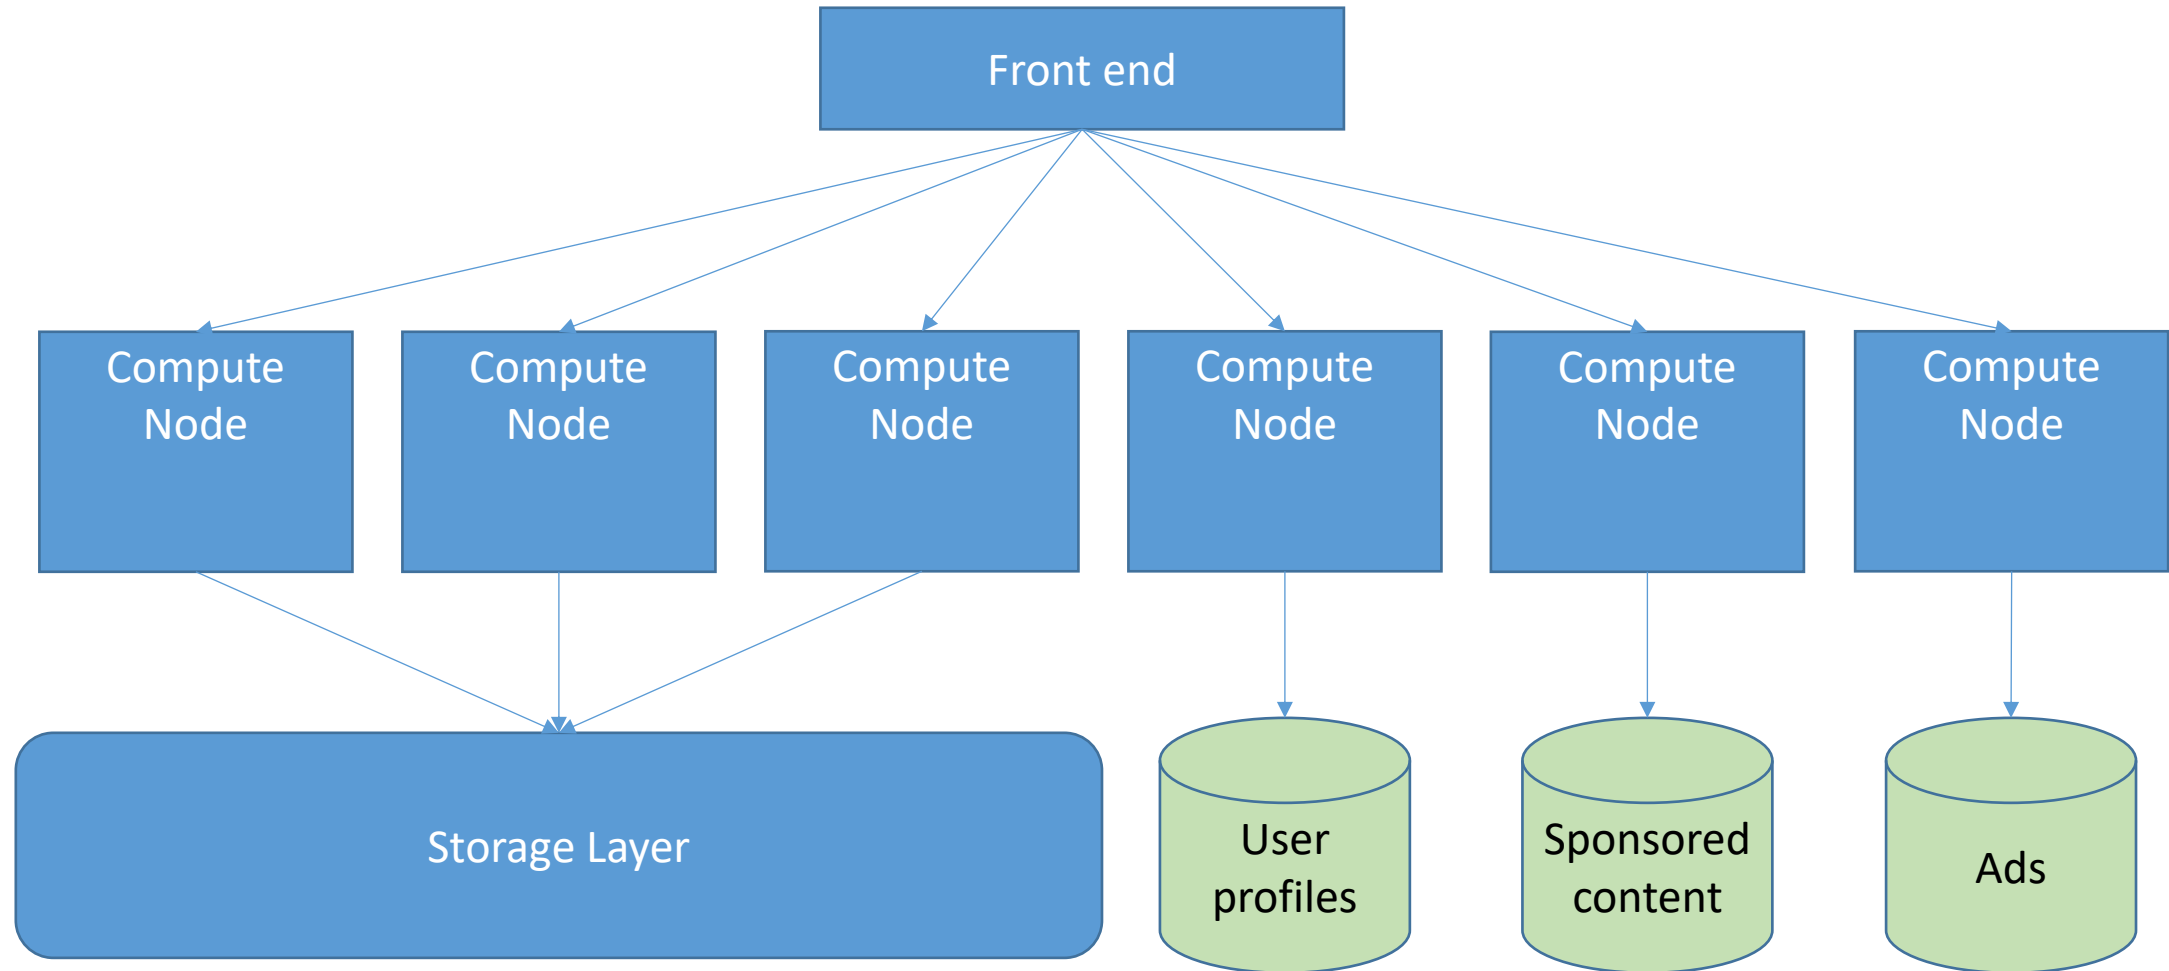
Where do we come from



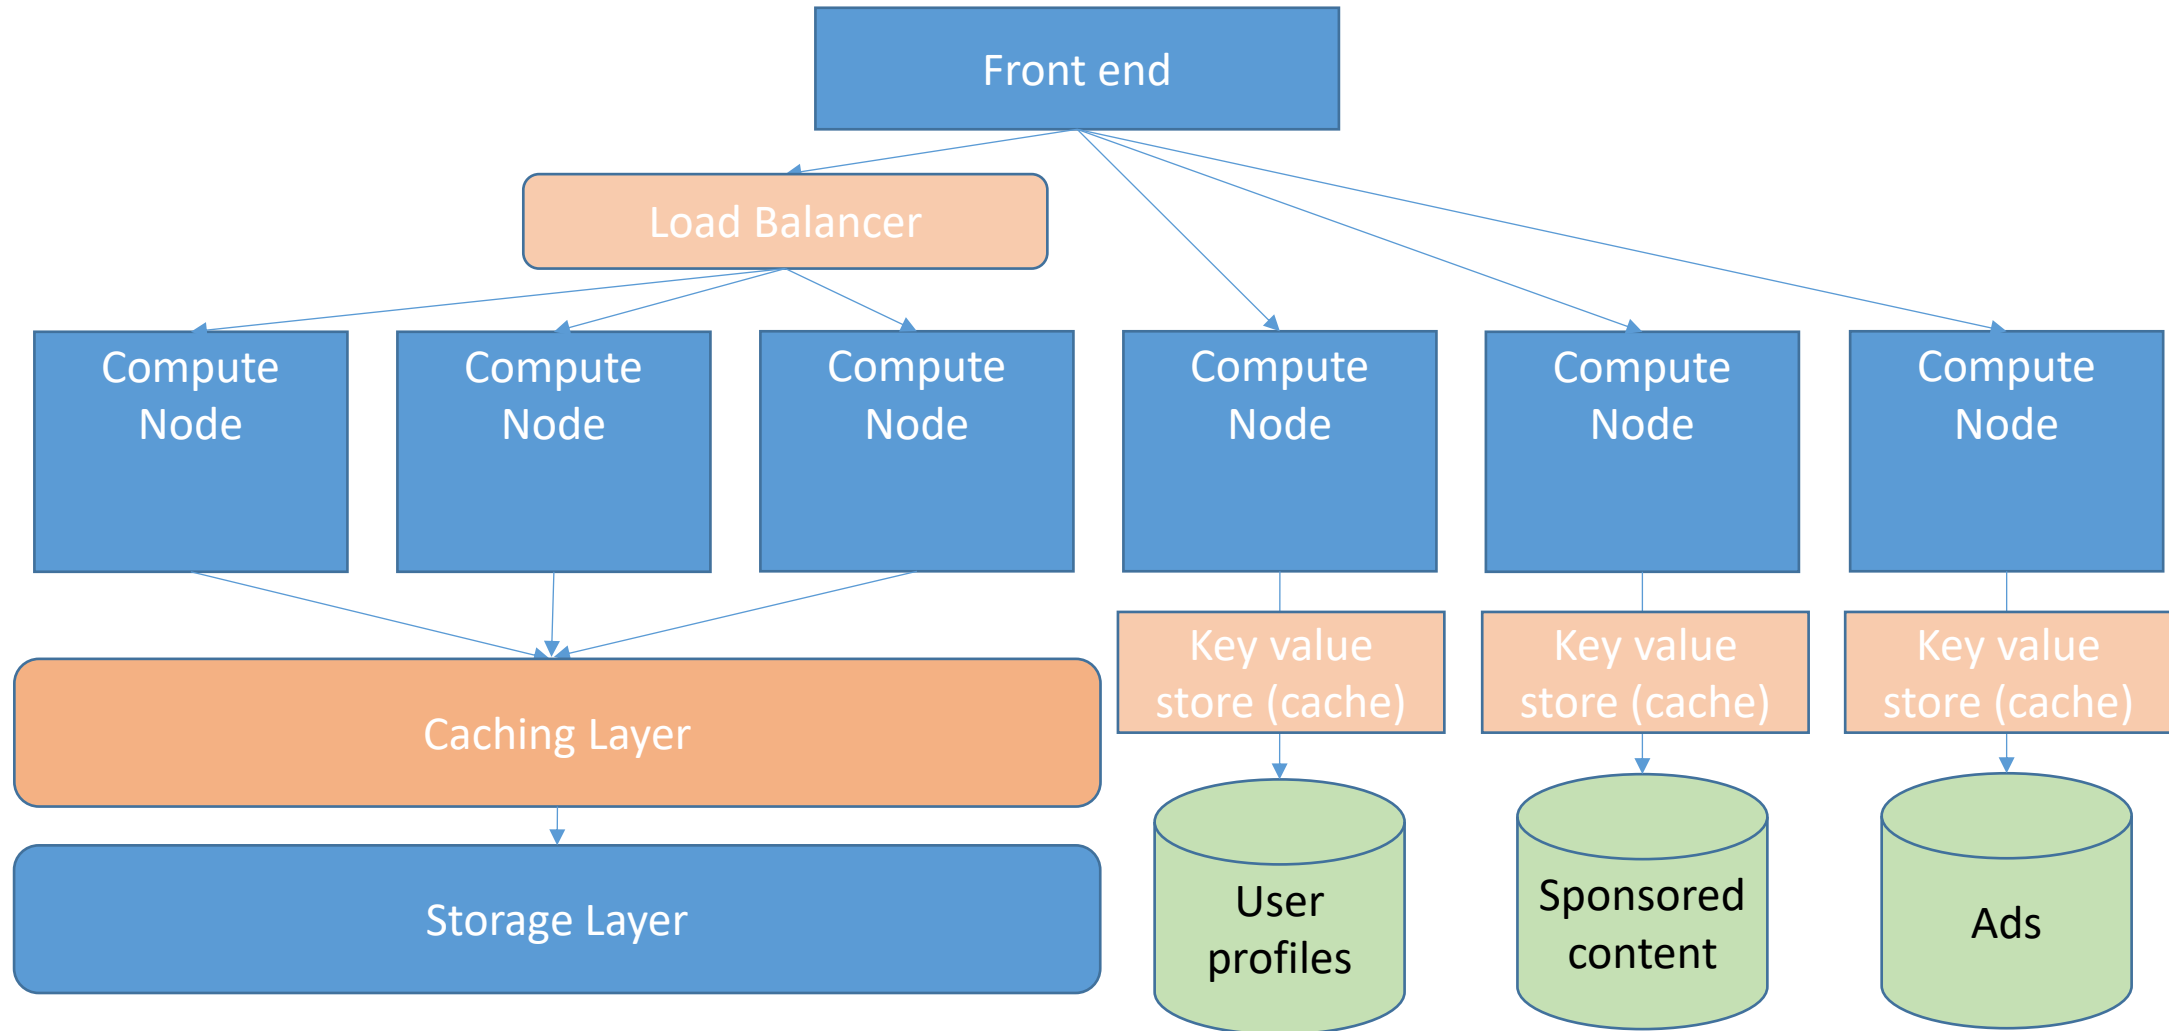
Separation of compute and storage



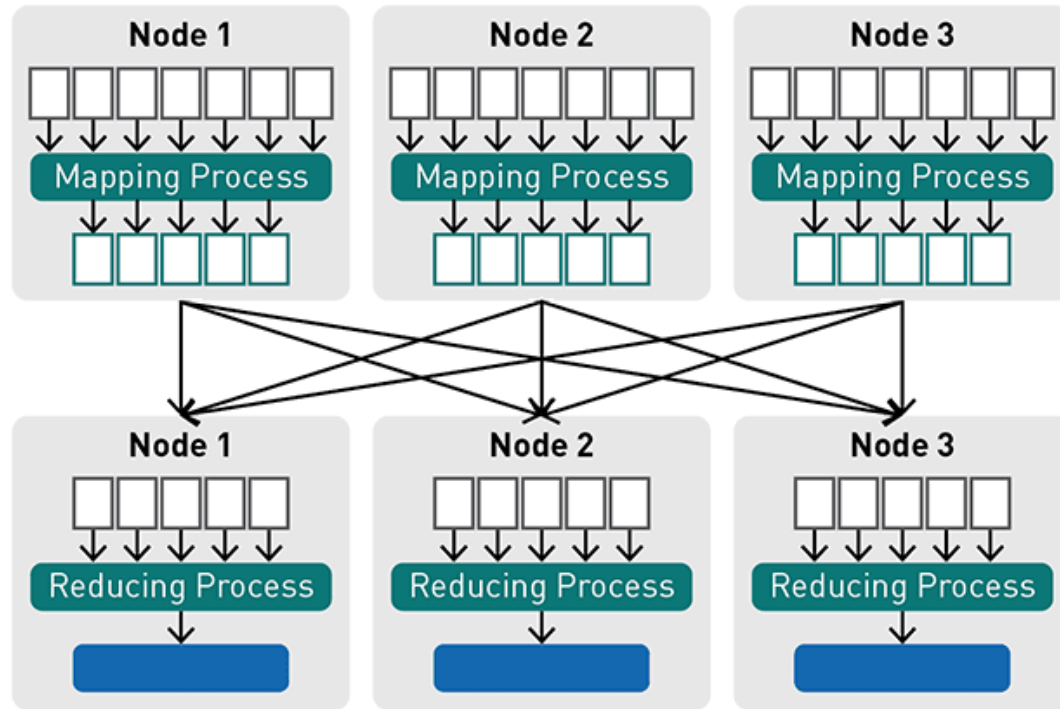
Additional functionality



Compensate for the design problems

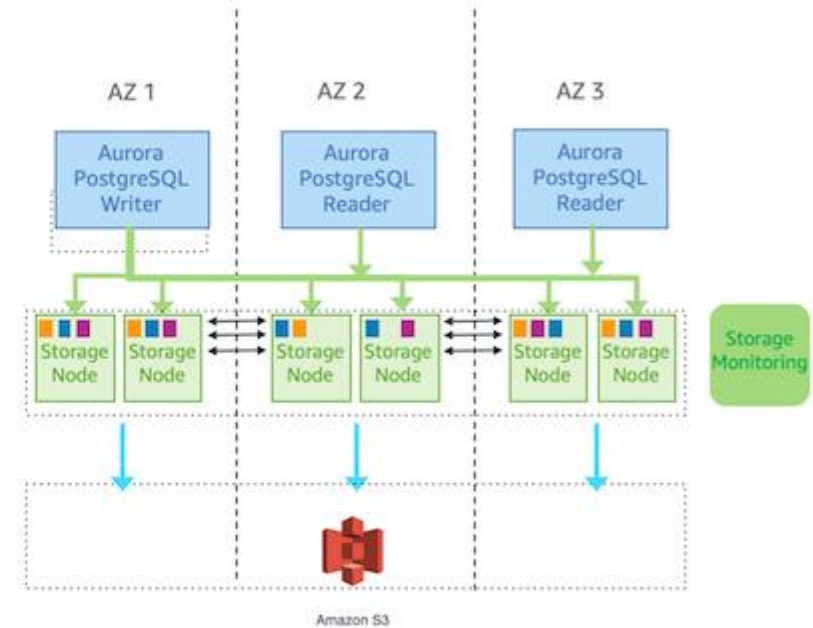


We design around it!




<https://developer.hpe.com/blog/spark-101-what-is-it-what-it-does-and-why-it-matters/>

Amazon Aurora PostgreSQL




The price of distribution



(a) Descendants sorted in increasing order

A Cloud-Scale Characterization of Remote Procedure Calls, SOSPP'23

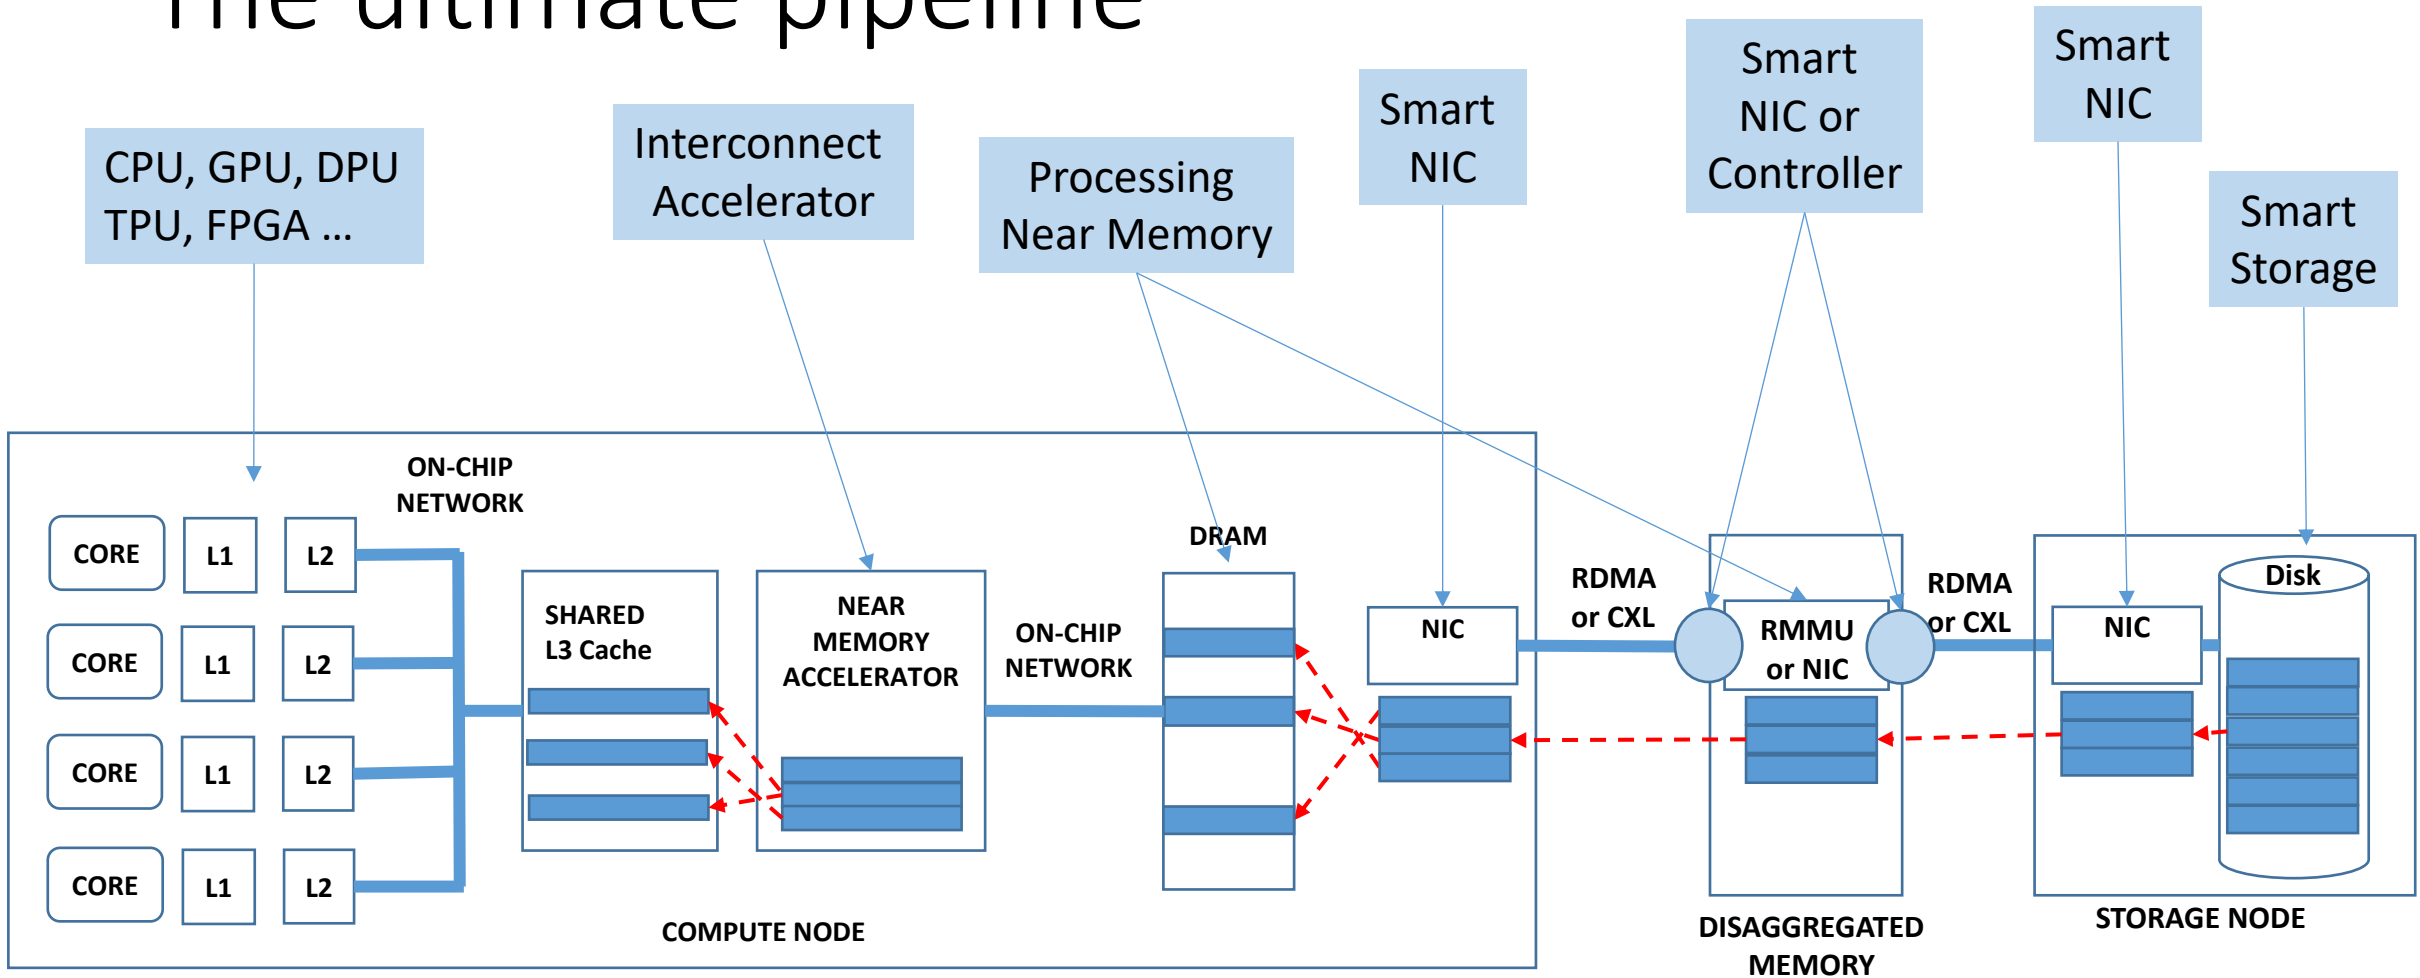


Profiling a warehouse-scale computer, ISCA 2015

Main insight of the talk

**If data has to move,
it has to be
processed along
the way**


The ultimate pipeline



How to make it work

- Has to be part of the infrastructure
 - Homogeneity to simplify software development
- Has to be usable across many use cases
 - Reconfigurable, reprogrammable, multi-tenancy
- Has to be integrated with the rest of the data center
 - Networks, schedulers, management infrastructure, etc.
- Has to be integrated into actual systems
 - Data processing platforms, databases, ML suites, etc.

SLASH (joint work with AMD)




Feasibility argument



Proven gains and implementations

- **Ibex - An Intelligent Storage Engine with Support for Advanced SQL Off-loading.** Louis Woods, Zsolt István, Gustavo Alonso. VLDB 2014
- **Histograms as a side effect of data movement for big data.** Zsolt István, Louis Woods, Gustavo Alonso. SIGMOD 2014
- **Consensus in a Box: Inexpensive Coordination in Hardware.** Zsolt István, David Sidler, Gustavo Alonso, Marko Vukolic. NSDI 2016
- **Caribou: Intelligent Distributed Storage.** Zsolt István, David Sidler, Gustavo Alonso. VLDB 2017
- **Do OS abstractions make sense on FPGAs?** Dario Korolija, Timothy Roscoe, Gustavo Alonso. OSDI 2020
- **StRoM: smart remote memory.** David Sidler, Zeke Wang, Monica Chiosa, Amit Kulkarni, Gustavo Alonso. EuroSys 2020
- **FleetRec: Large-Scale Recommendation Inference on Hybrid GPU-FPGA Clusters.** Wenqi Jiang, Zhenhao He, Shuai Zhang, Kai Zeng, Liang Feng, Jiansong Zhang, Tongxuan Liu, Yong Li, Jingren Zhou, Ce Zhang, Gustavo Alonso. KDD 2021
- **ACCL+: an FPGA-Based Collective Engine for Distributed Applications.** Zhenhao He, Dario Korolija, Yu Zhu, Benjamin Ramhorst, Tristan Laan, Lucian Petrica, Michaela Blott, Gustavo Alonso. OSDI 2024

Enough hardware



Corsica: A project zipline ASIC

Compression without compromise:

- High compression ratio
- Low latency
- Inline encryption, authentication
- High total throughput

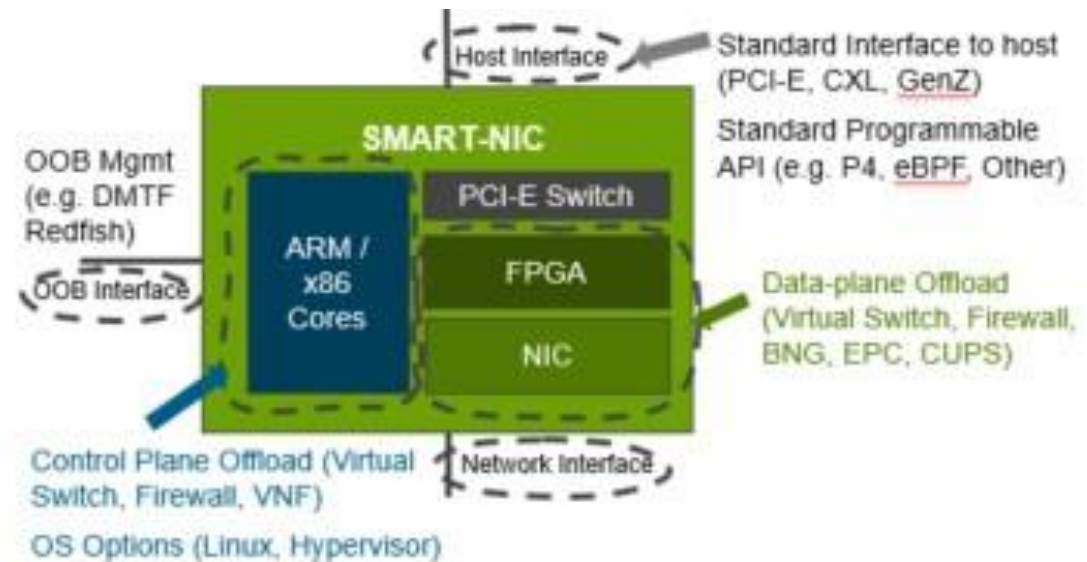
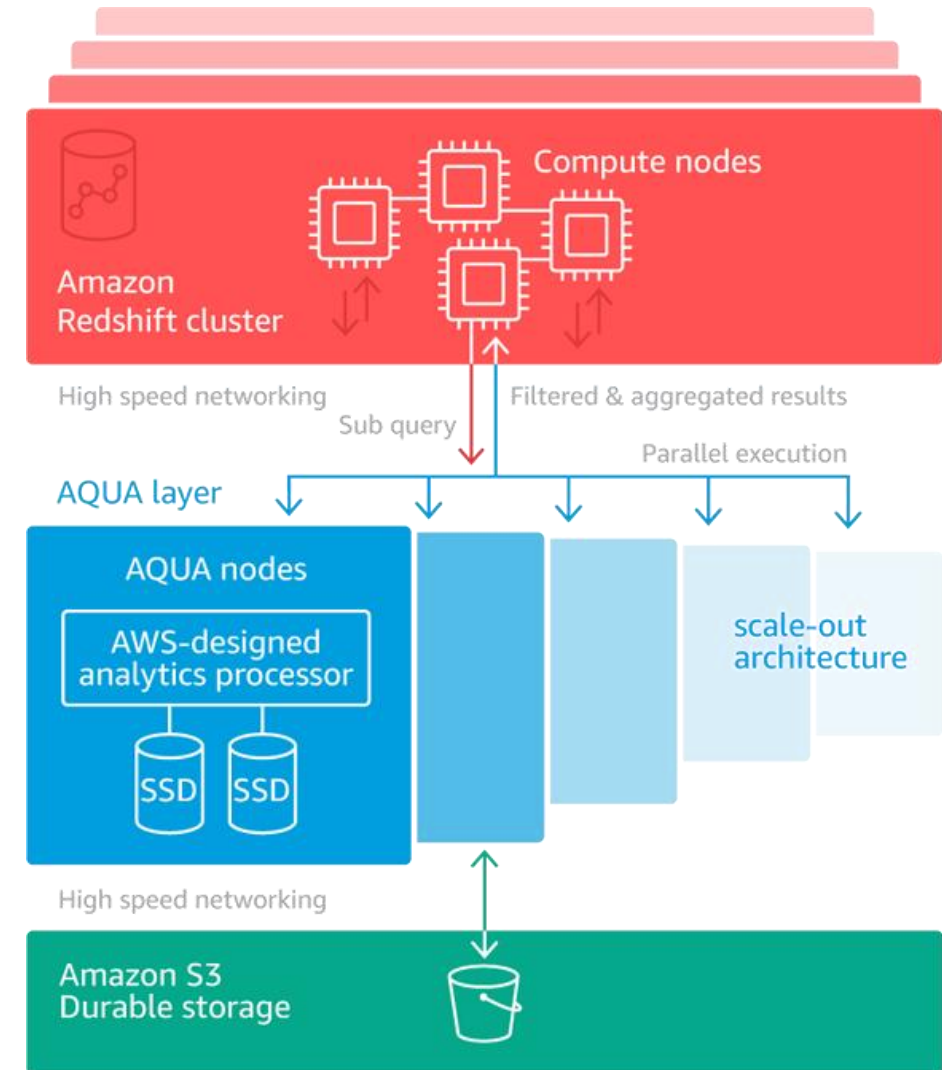
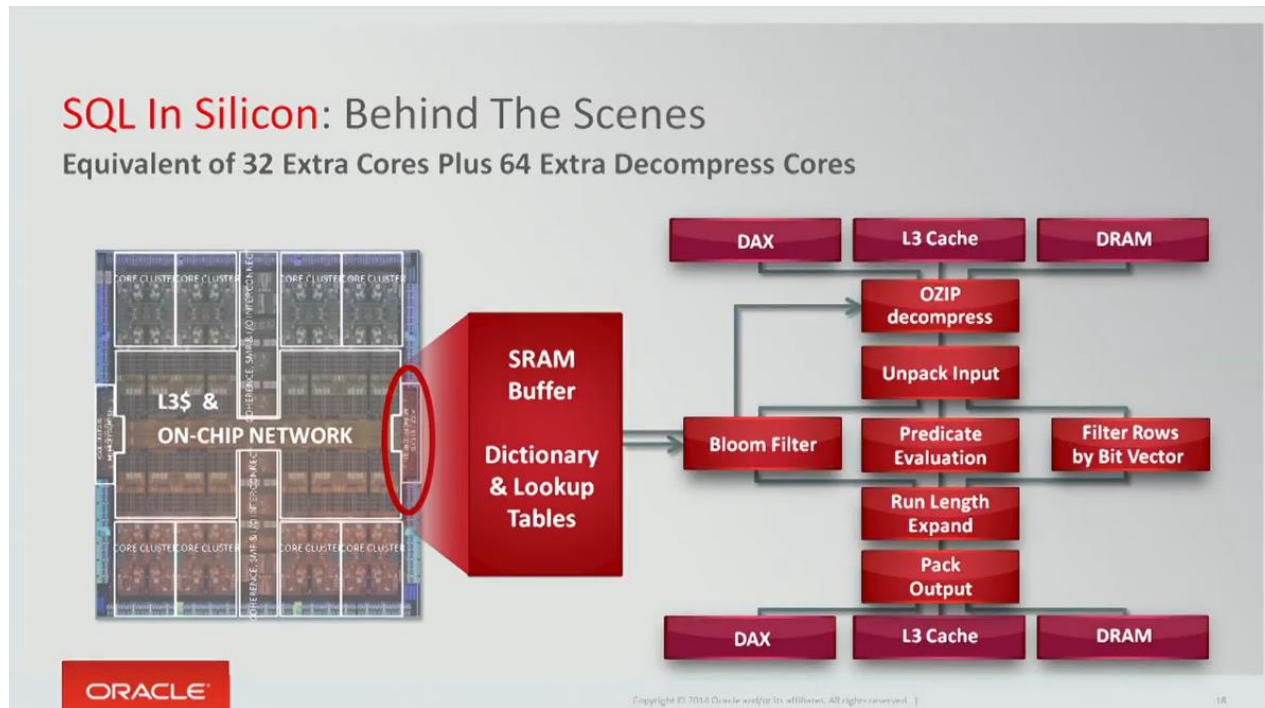


Fig 1. SMART-NIC Architecture

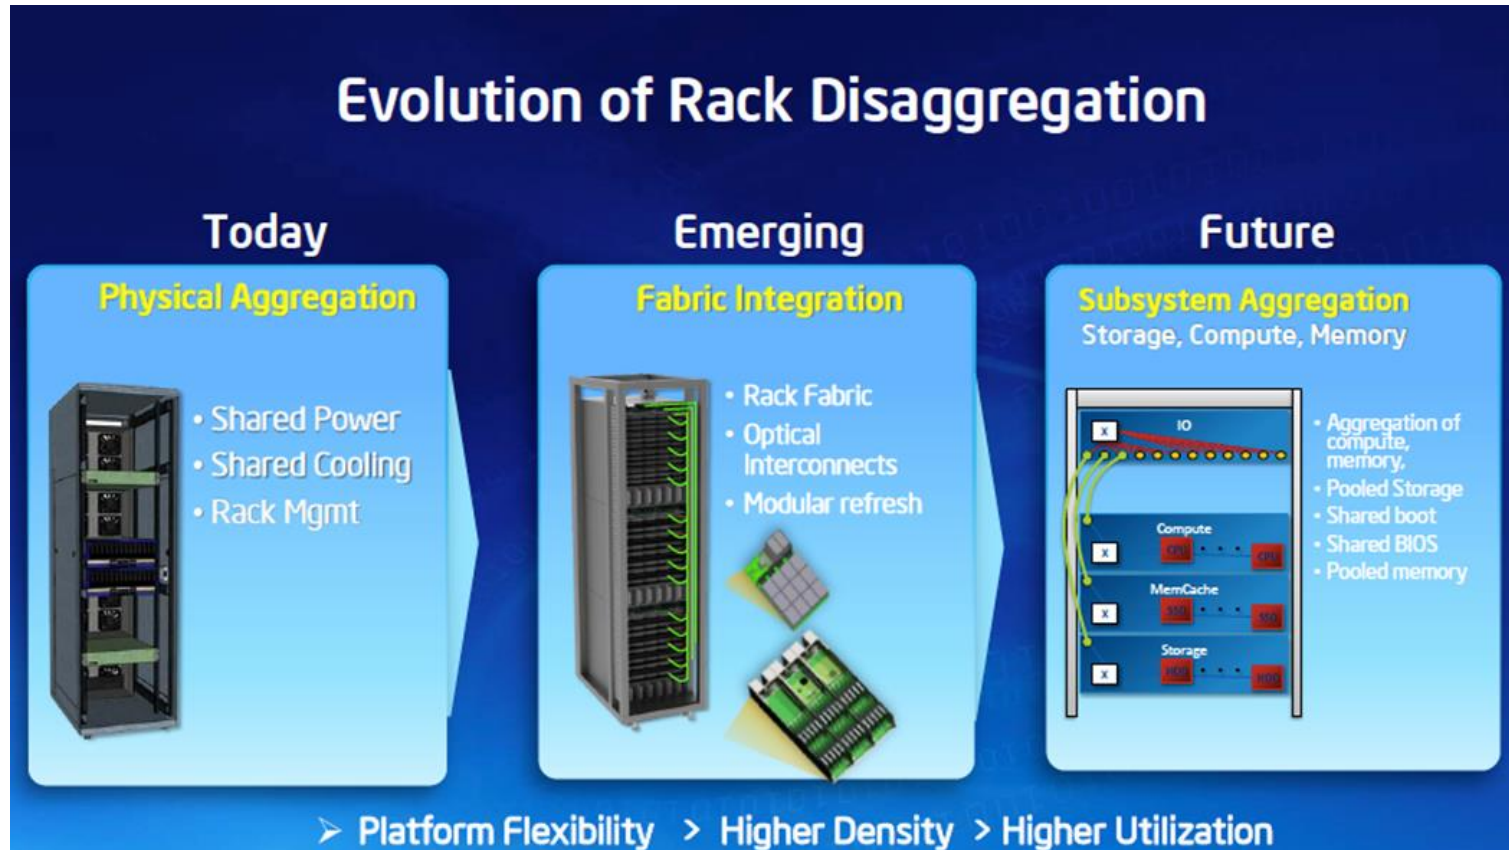
Function shipping

No data shipping



https://pages.awscloud.com/AQUA_Preview.html

Rack scale disaggregation



Conclusions

- FPGAs can play a crucial role here:
 - Exploring architectures
 - Initial deployments for reconfigurability
 - Supporting diversity of solutions
 - Lower energy consumption
- In a heterogeneous, disaggregated architecture, the FPGA can become the center of the system
- But:
 - FPGAs need to become more data center oriented
 - Tools and support need to improve
 - Research is needed on integrating them into real systems