CS 839: Design the Next-Generation Database

Lecture 1: Introduction

Xiangyao Yu

1/21/2020

Who am I?

Xiangyao Yu

- Pronounced like Shiang-Yao Yu.

Assistant Professor in Computer Science

PhD (in computer architecture) and postdoc (in databases) at MIT

Research interests:

- Transaction processing
- New hardware for databases
- Cloud databases

Today’s Agenda

What is this course about?

Course logistics

Class projects

A brief history of database systems

Single-Core, Disk-Based (1970s – 2000s)

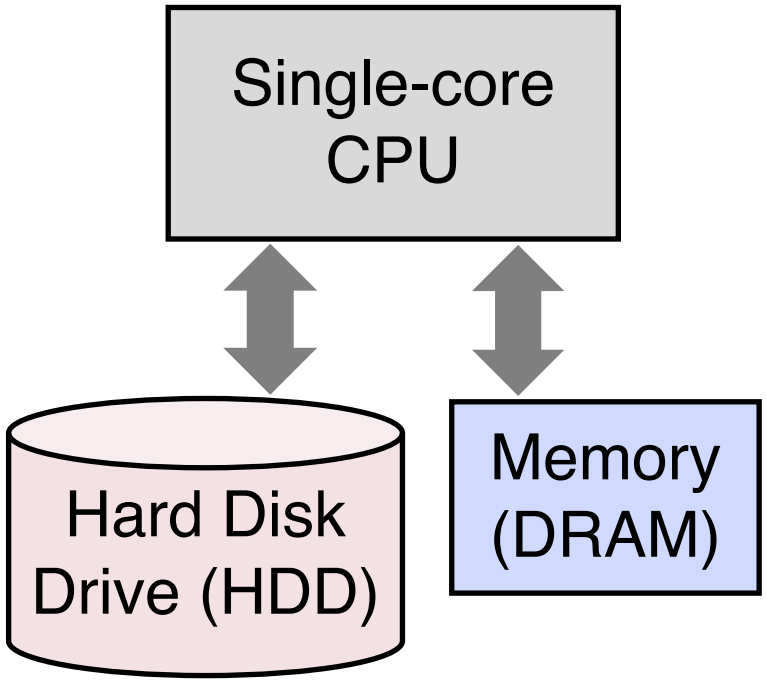

Data stored in HDD

Main memory is a “cache”

Timesharing across users

Distributed, Disk-Based (1980s – 2000s)

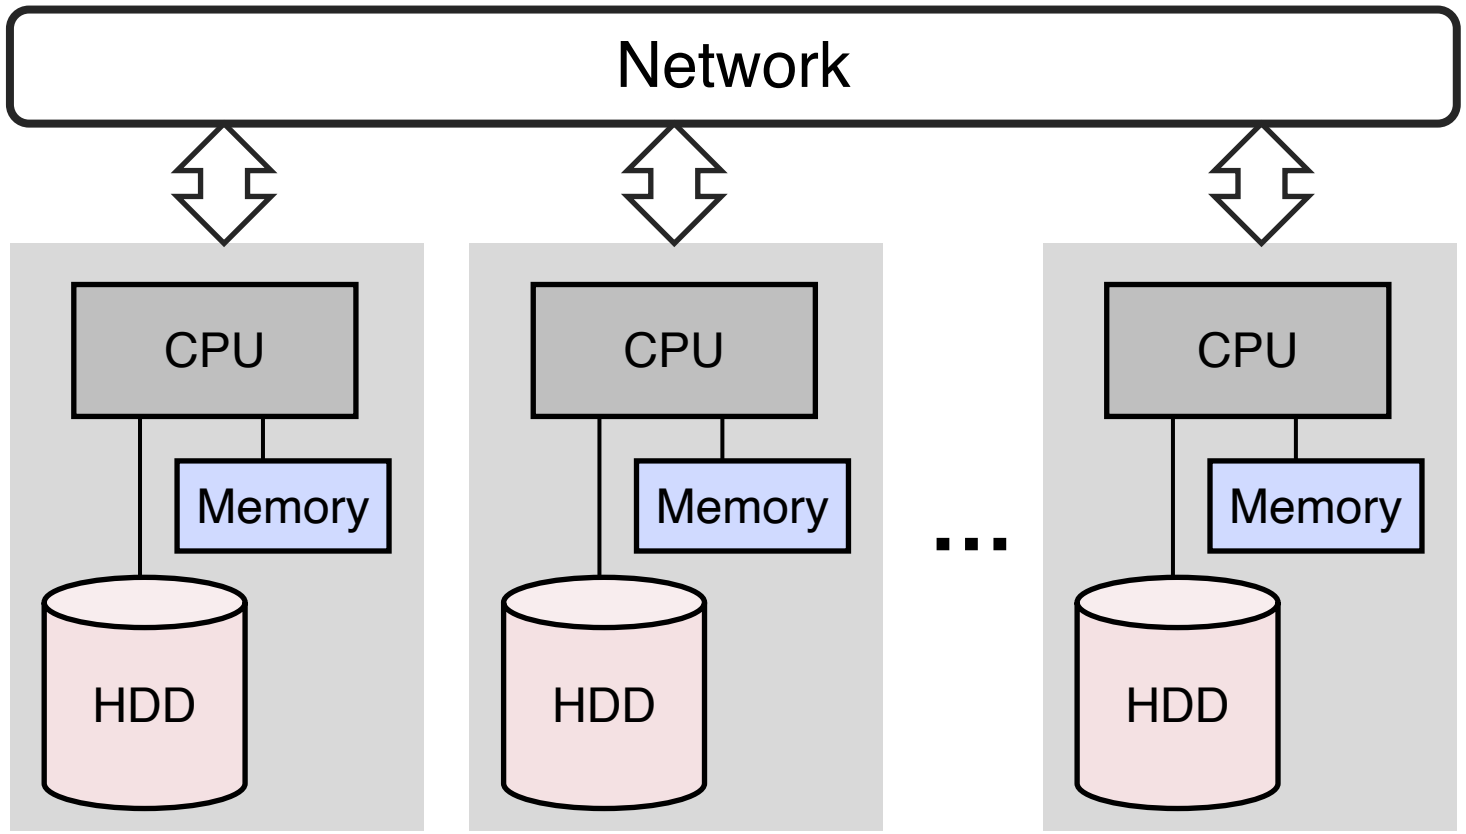

Shared-nothing architecture

Servers communicate over network

Can scale out to thousands of servers

Multicore, In-Memory (2000s – today)

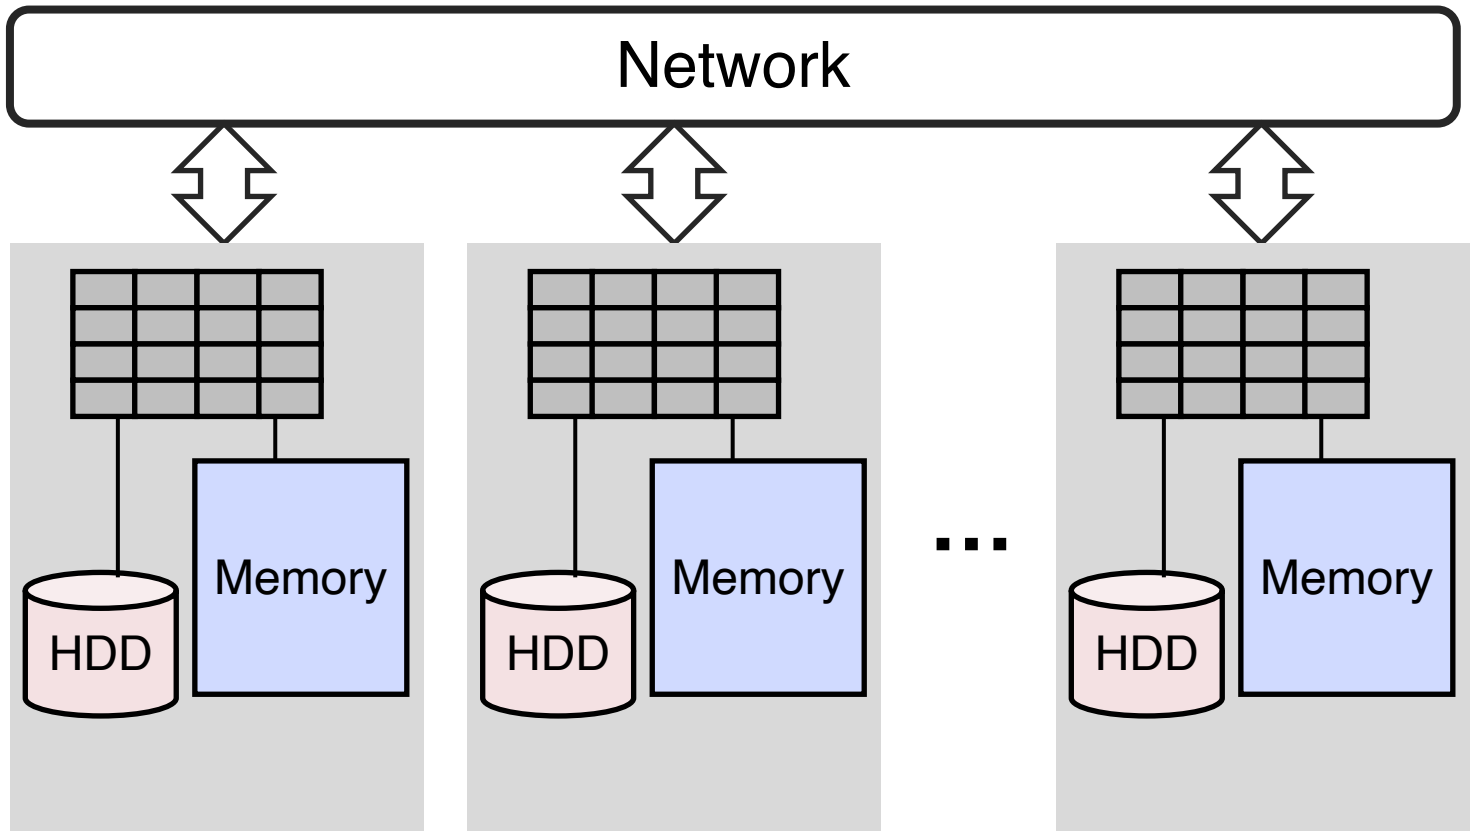

Multicore processors

Data stored in memory

- Memory is cheaper
- Memory capacity increases

What Is Next?

1. New processing units:



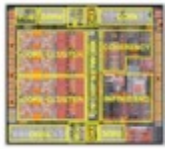

Multicore GPU FPGA Accelerator

2. New memory/storage


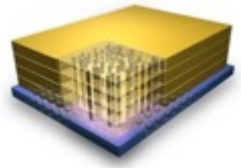

SSD NVM HBM

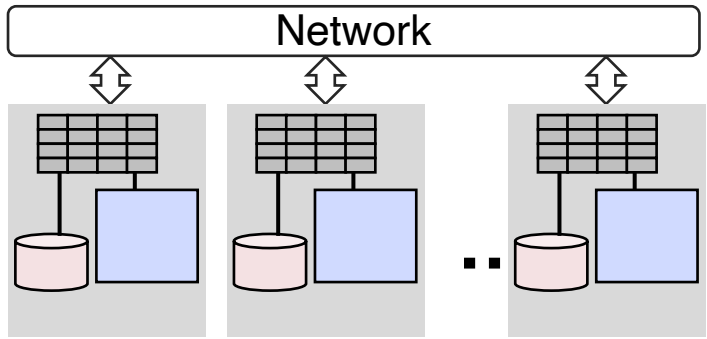

Database system today

RDMA SmartNIC

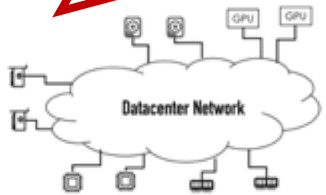

Disaggregation FaaS

3. New network technology

4. Cloud architecture

What Is Next?

1. New processing units:

Multicore


GPU


FPGA

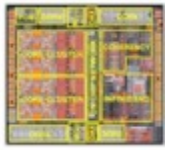
Accelerator

2. New memory/storage

SSD


NVM

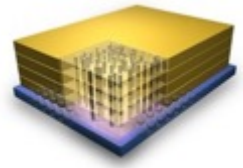
HBM

Next-generation databases have new hardware and system architecture

RDMA


SmartNIC

3. New network technology

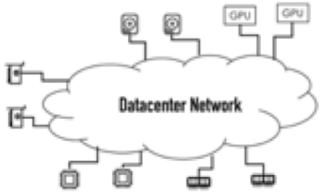

Disaggregation


FaaS

4. Cloud architecture

1. New Processing Units

Multicore

GPU

FPGA, accelerator

1. New Processing Units – Multicore CPU

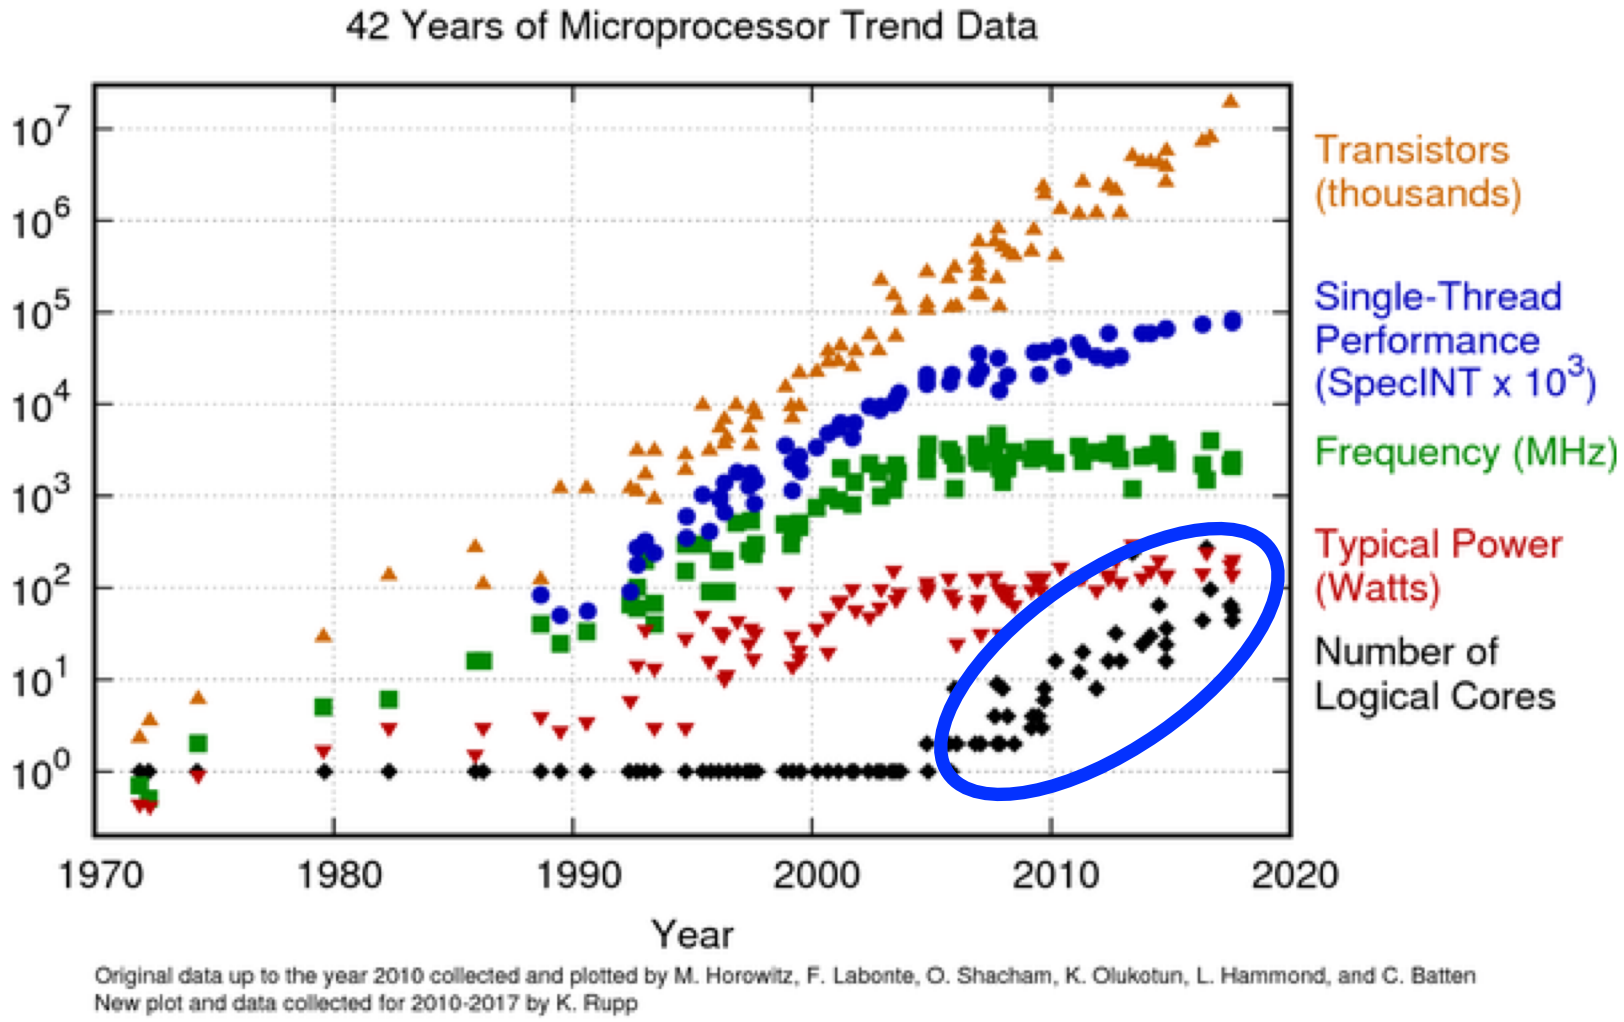

Core count will continue increasing -> scalability challenges

1. New Processing Units – GPU

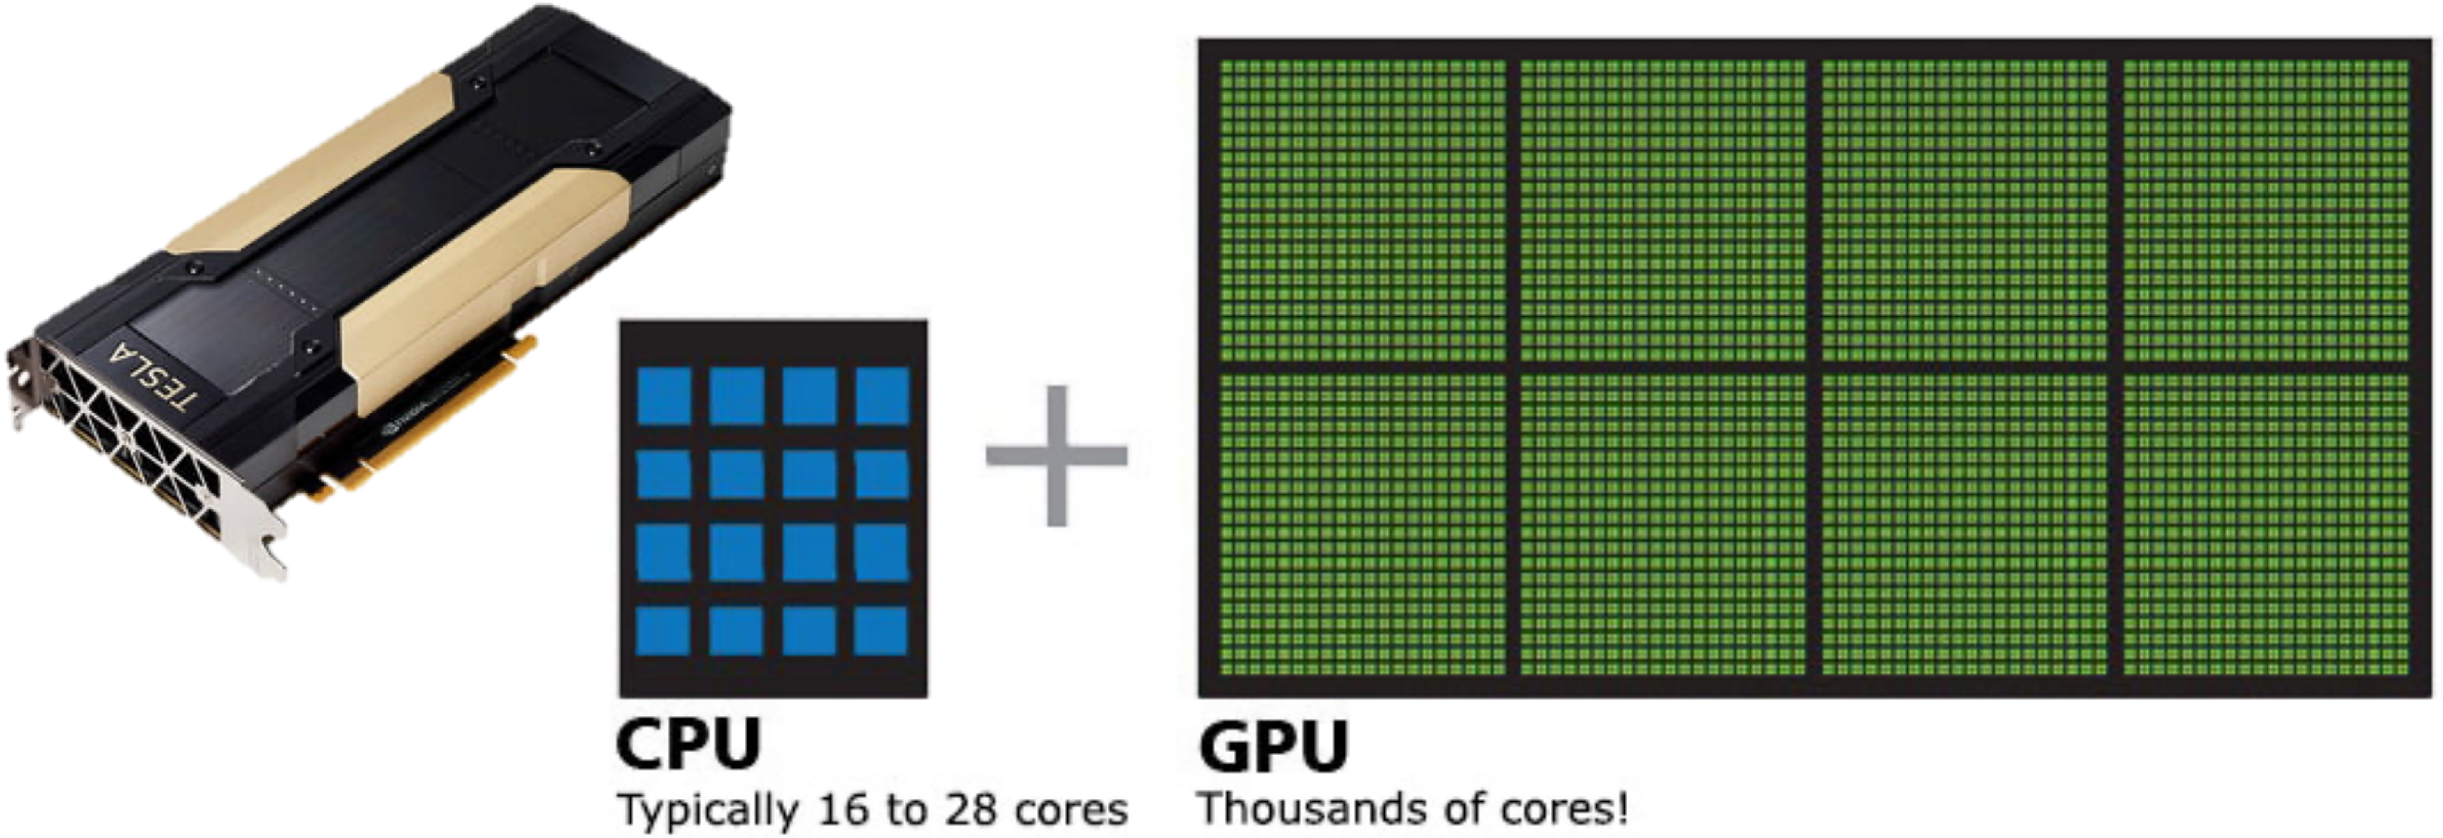

Graphics processing units (GPU) have massive parallelism but limited memory capacity

1. New Processing Units – Accelerators

FPGA

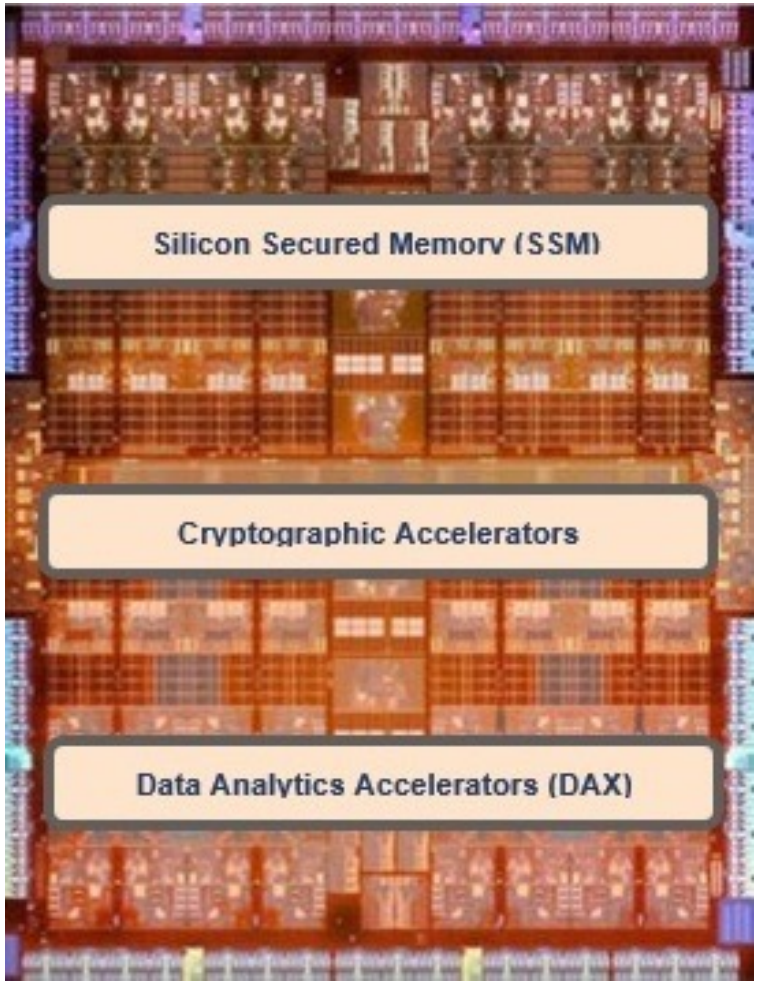

Oracle
software in silicon

Accelerators are effective for compute bound applications

2. New Memory/Storage

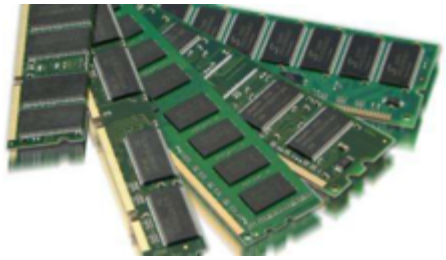

Non-volatile memory (NVM)

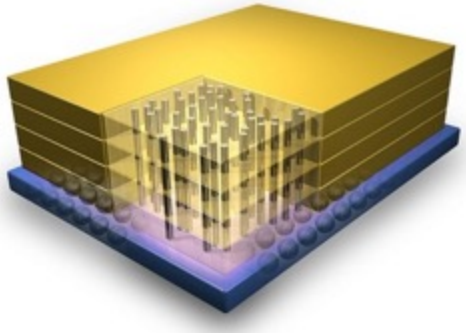

High Bandwidth Memory (HBM)

Process in Memory (PIM) / Smart SSD

2. New Memory/Storage – NVM

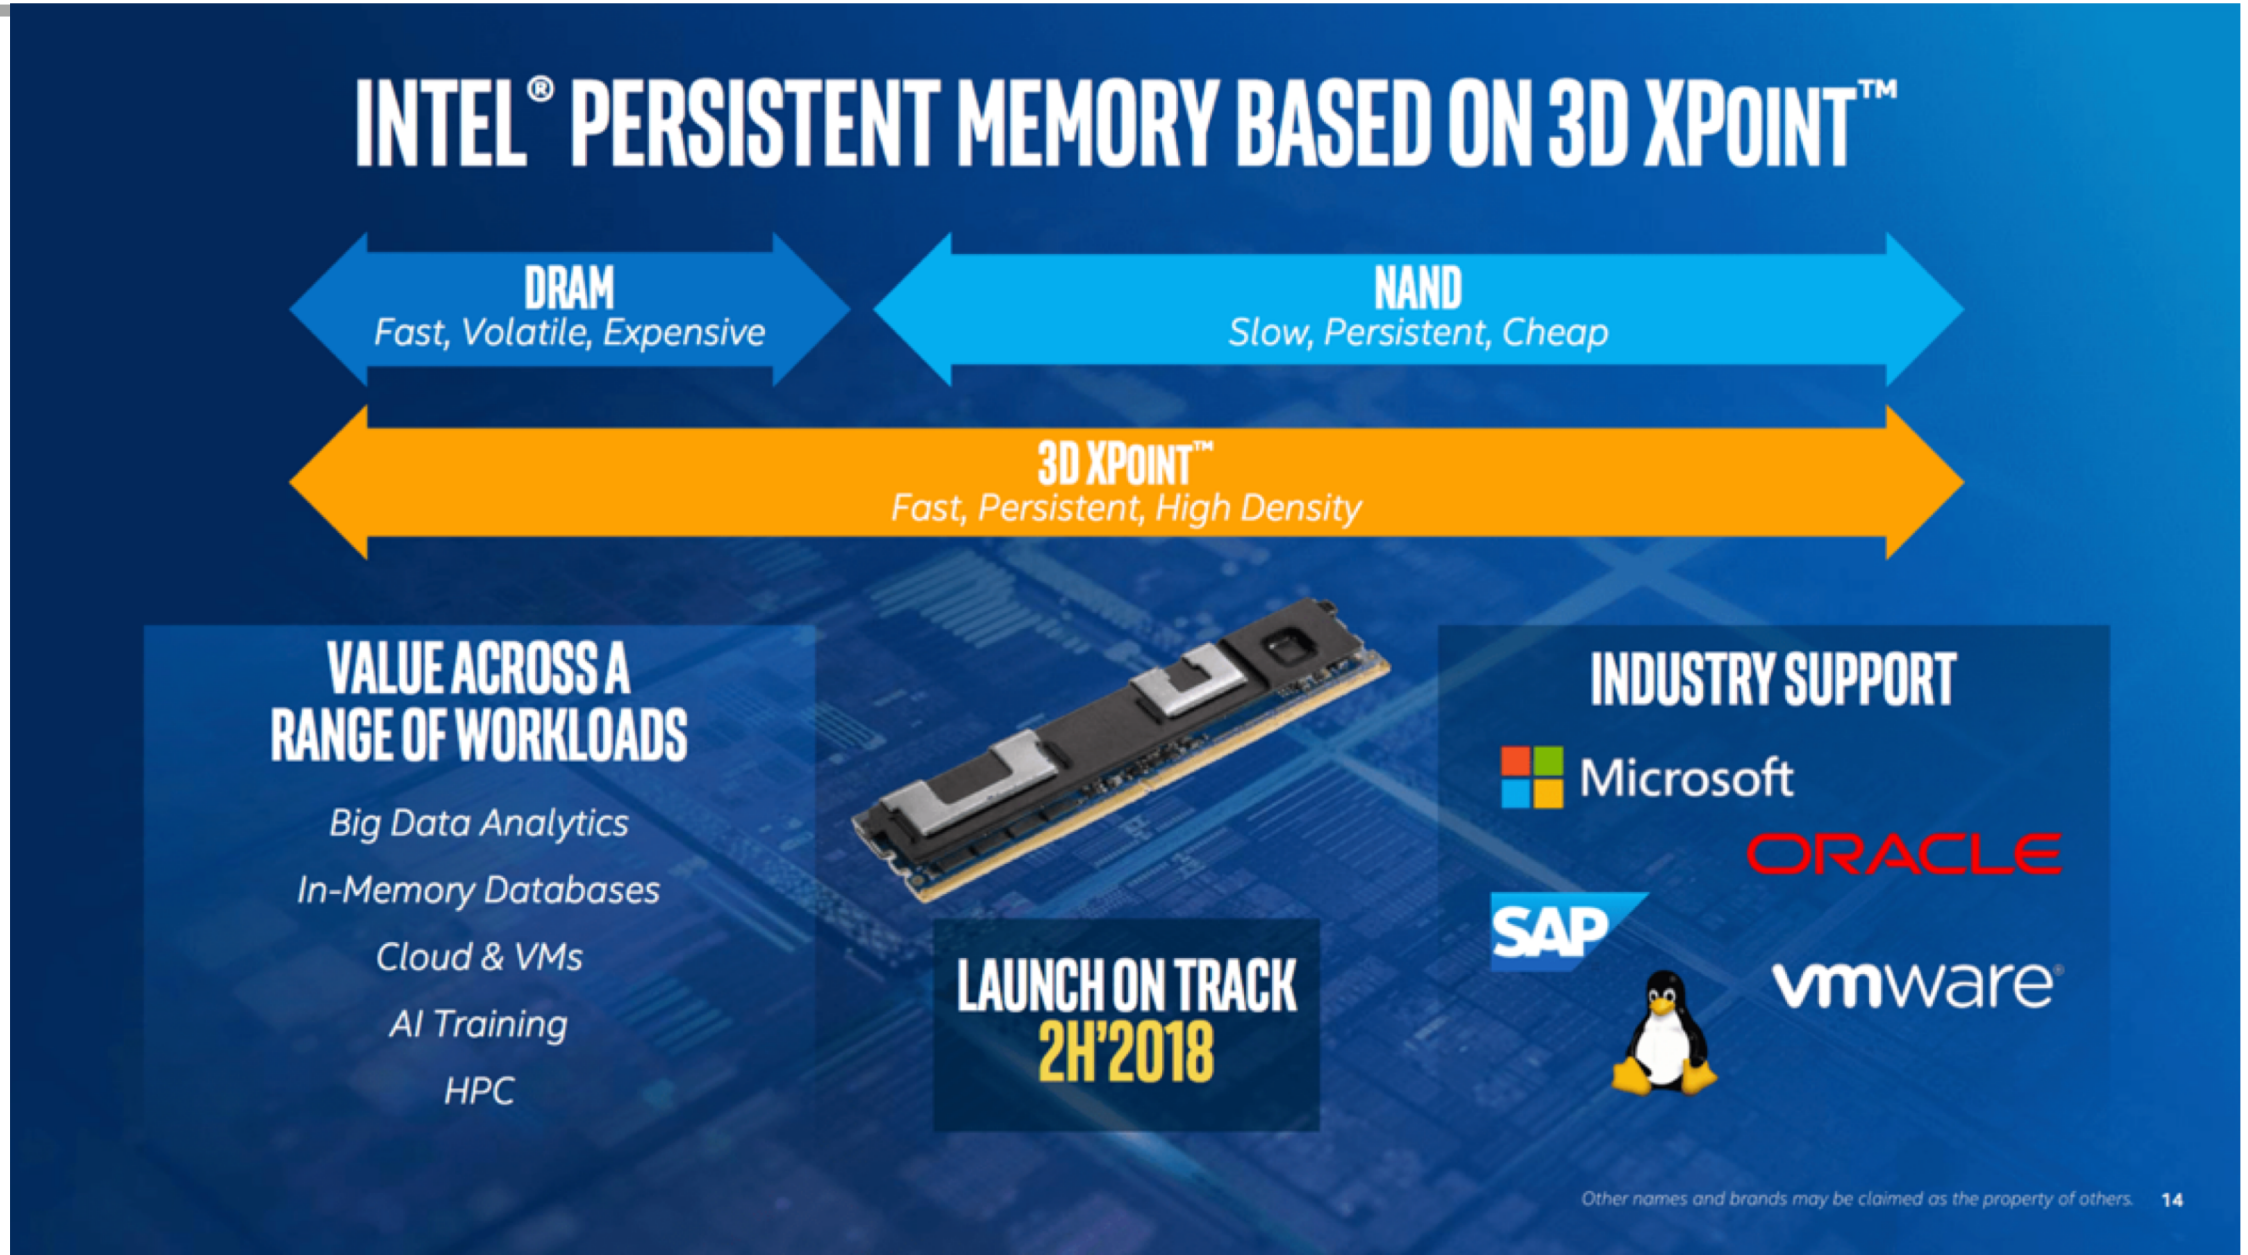

2. New Memory/Storage – HBM

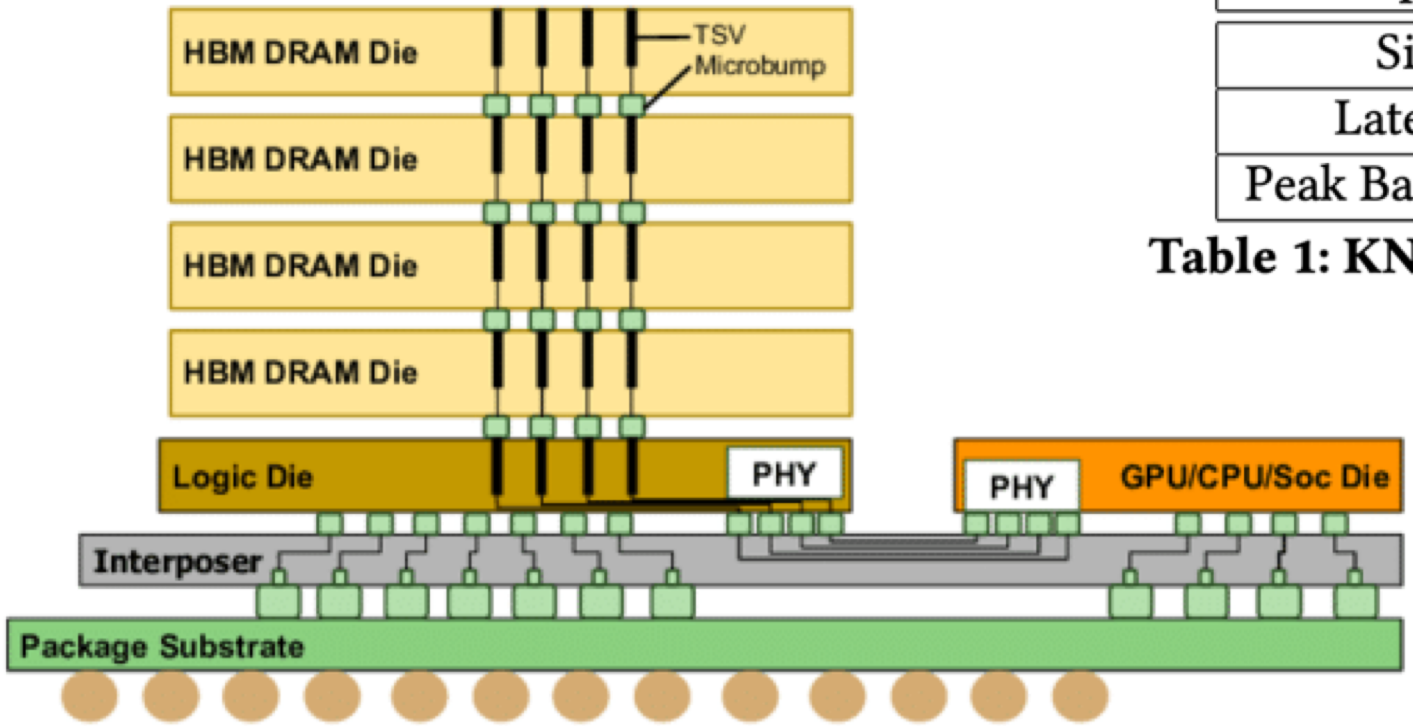

Property	DDR4	MCDRAM
Size	96GB	16GB
Latency	130-145ns	160-180ns
Peak Bandwidth	71GB/s	431GB/s

Table 1: KNL Memory Properties (SNC4 mode)

High bandwidth memory (HBM) has much higher bandwidth than DRAM

2. New Memory/Storage – PIM/SmartSSD

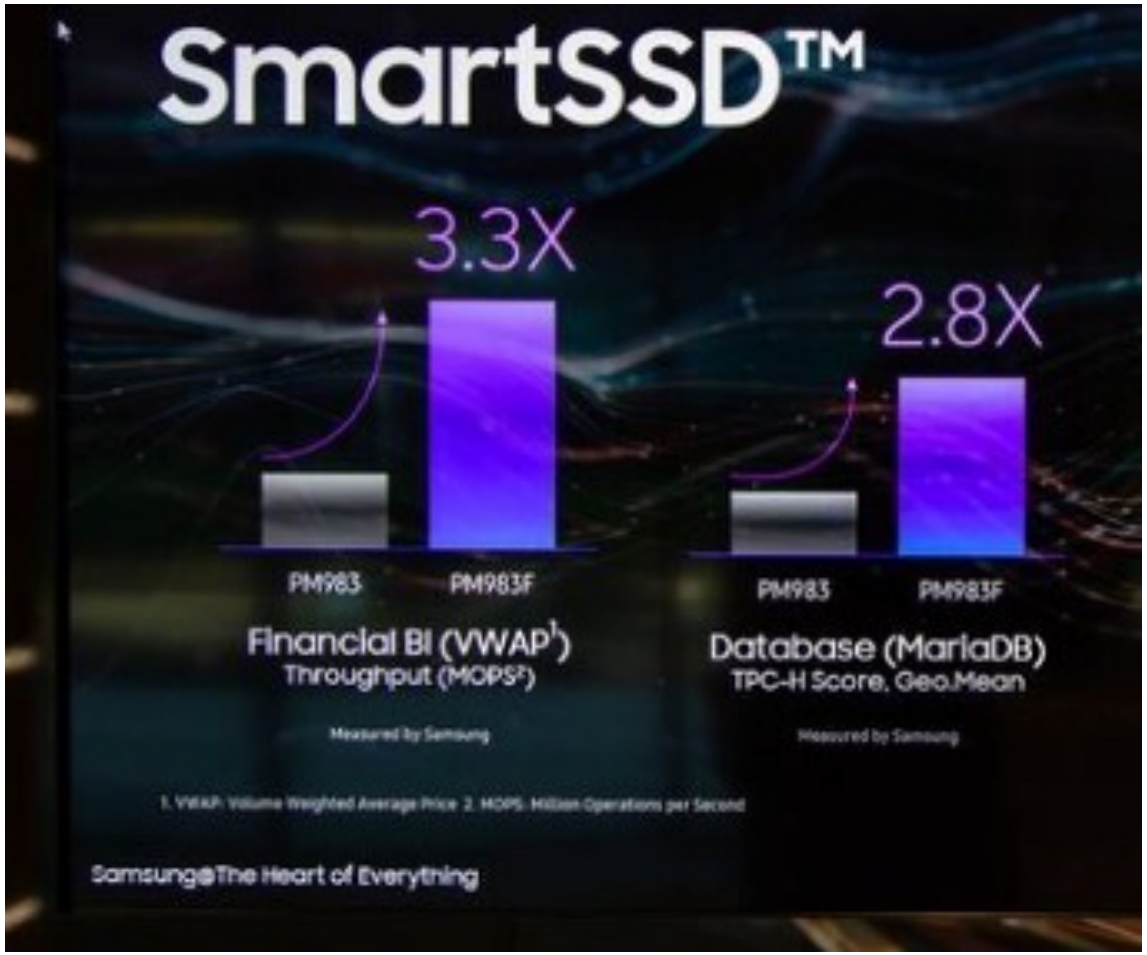

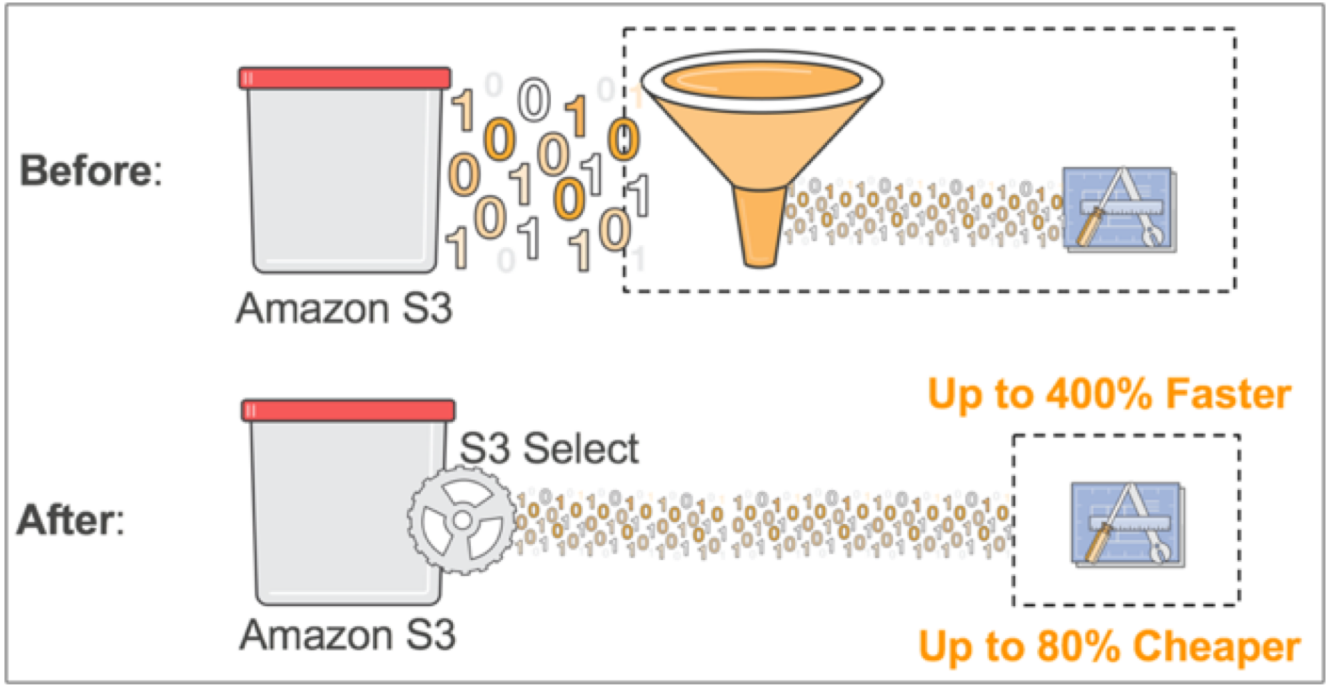

Pushing computation closer to data -> reduces data movement

3. New Network Technology

Remote direct memory access (RDMA)

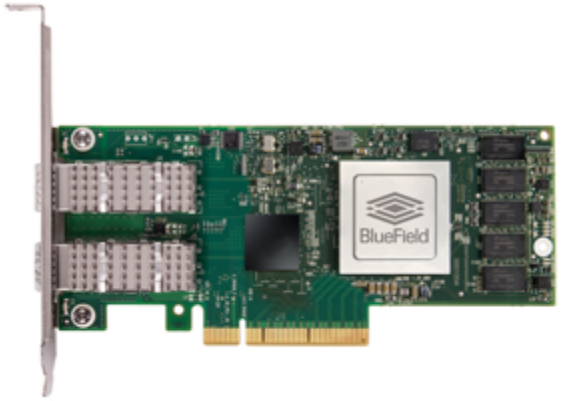

Smart NIC

3. New Network Technology – RDMA

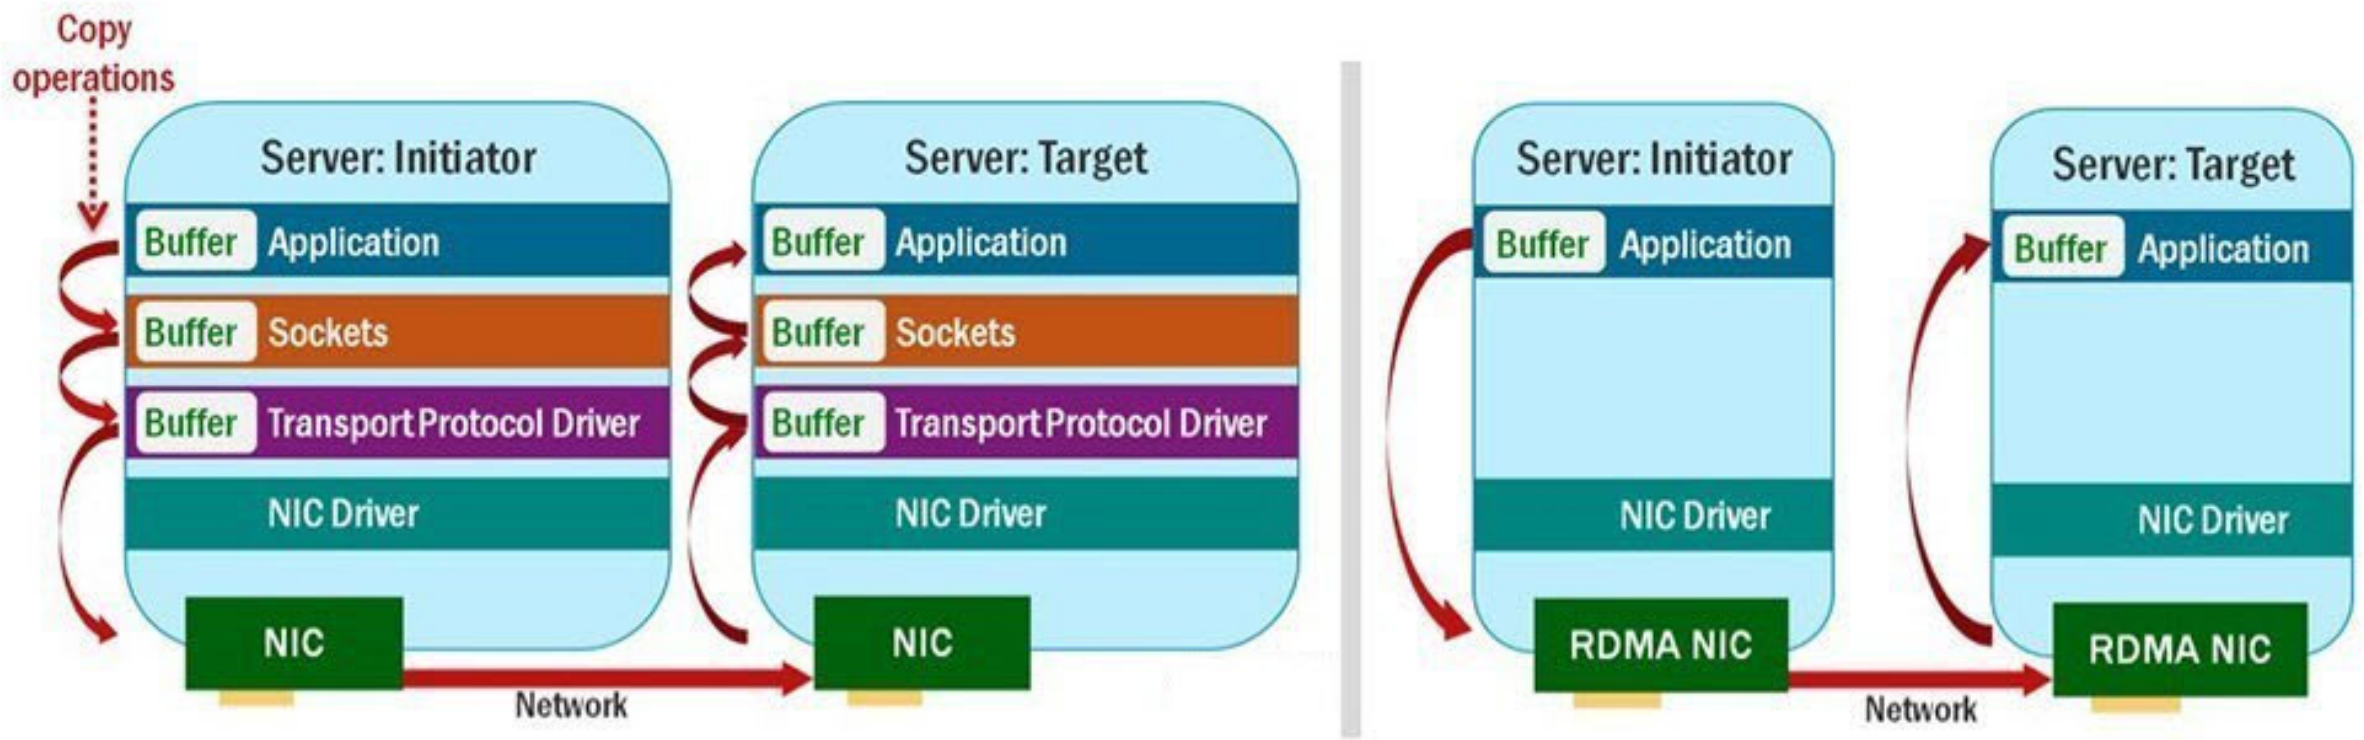

Remote direct memory access (RDMA) networks reduce latency

3. New Network Technology – Smart NIC

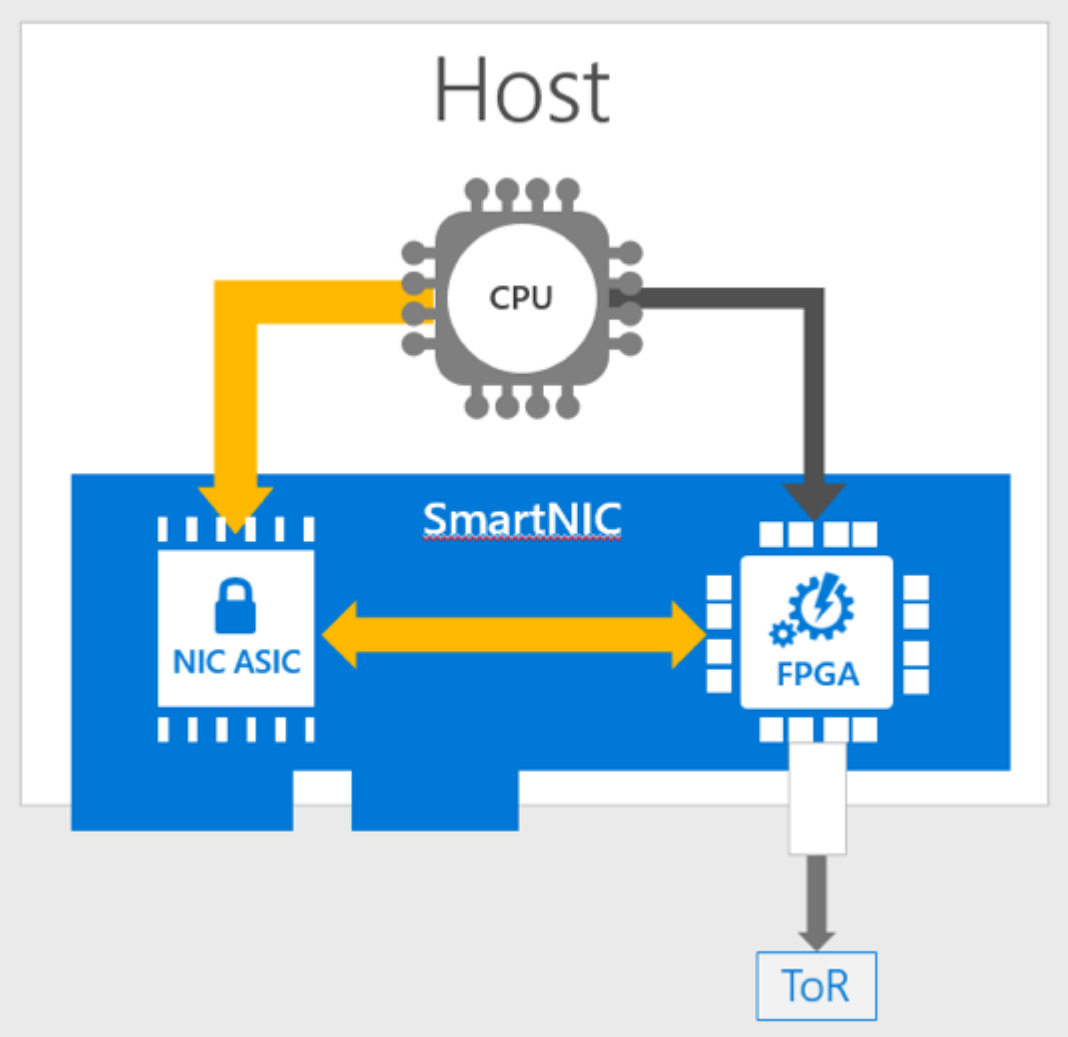

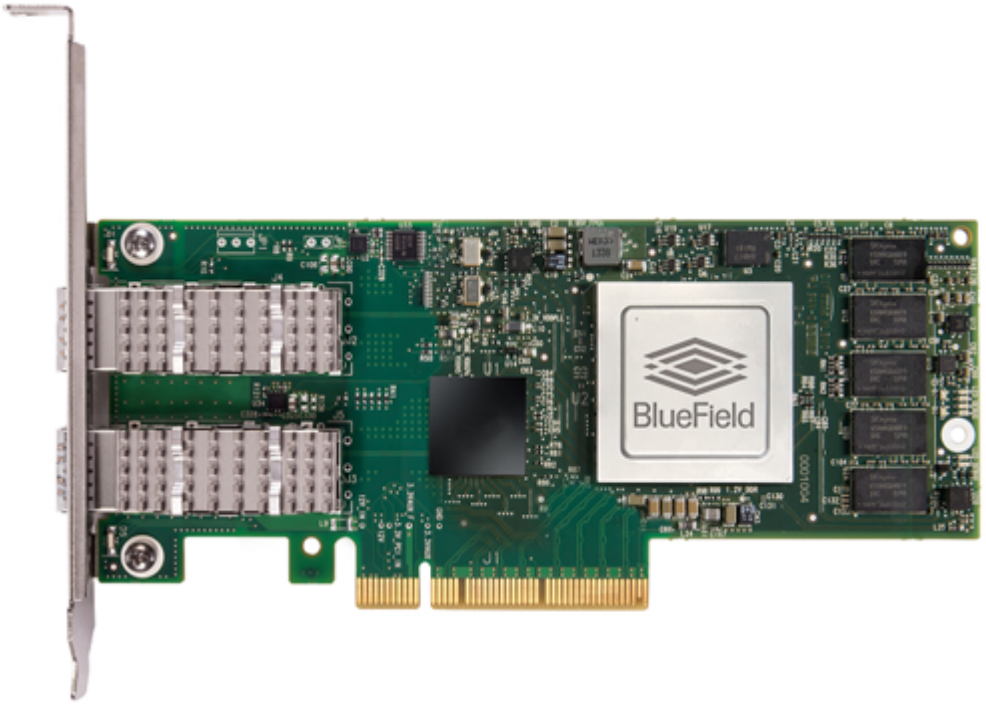

Pushing computation into the network

4. Cloud Architecture

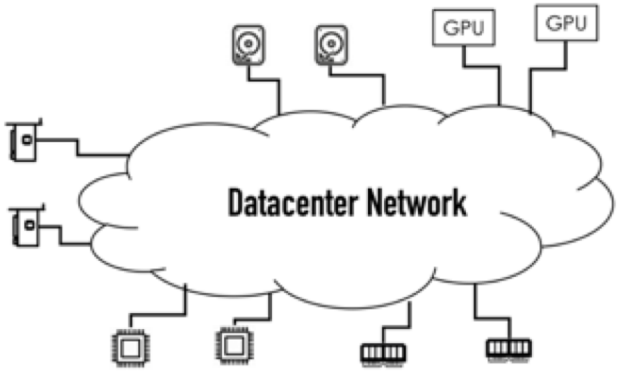

Resource disaggregation

Function-as-a-Service

4. Cloud Architecture – Resource Disaggregation

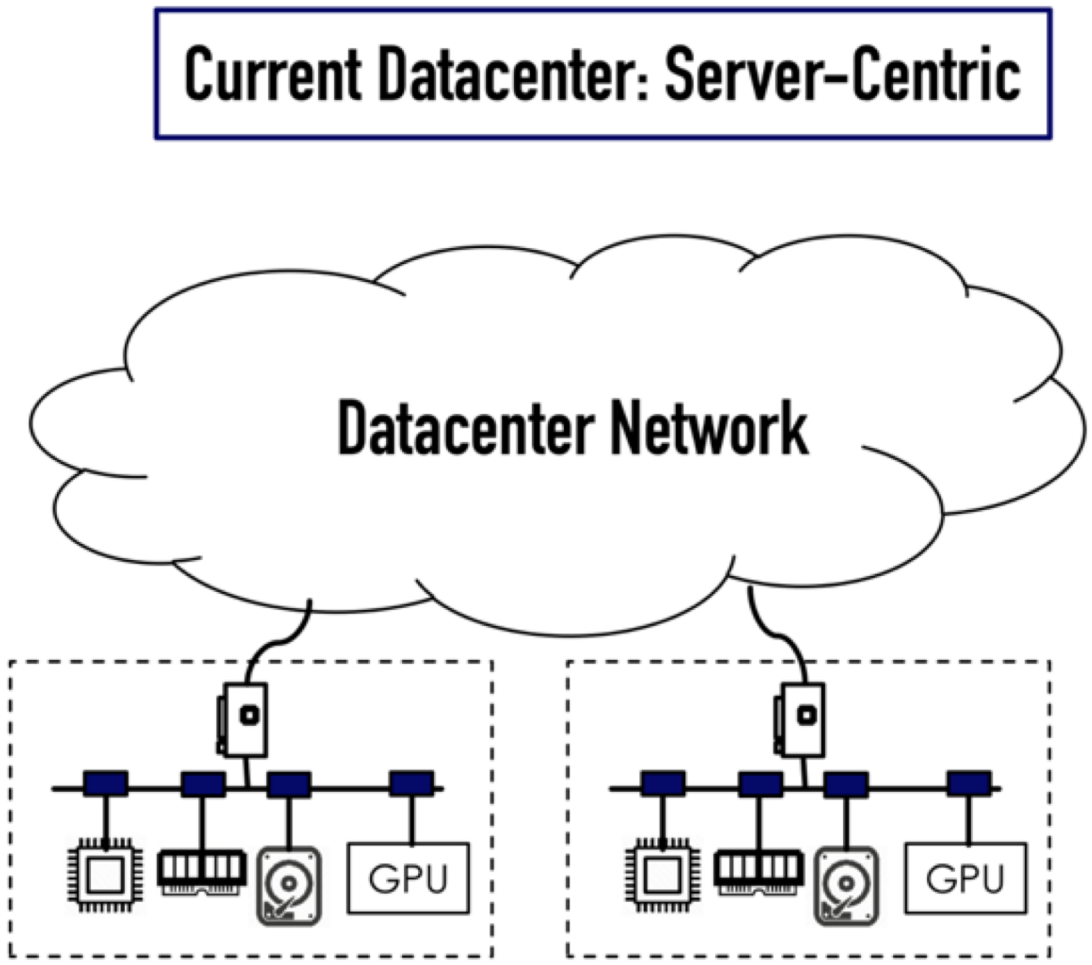

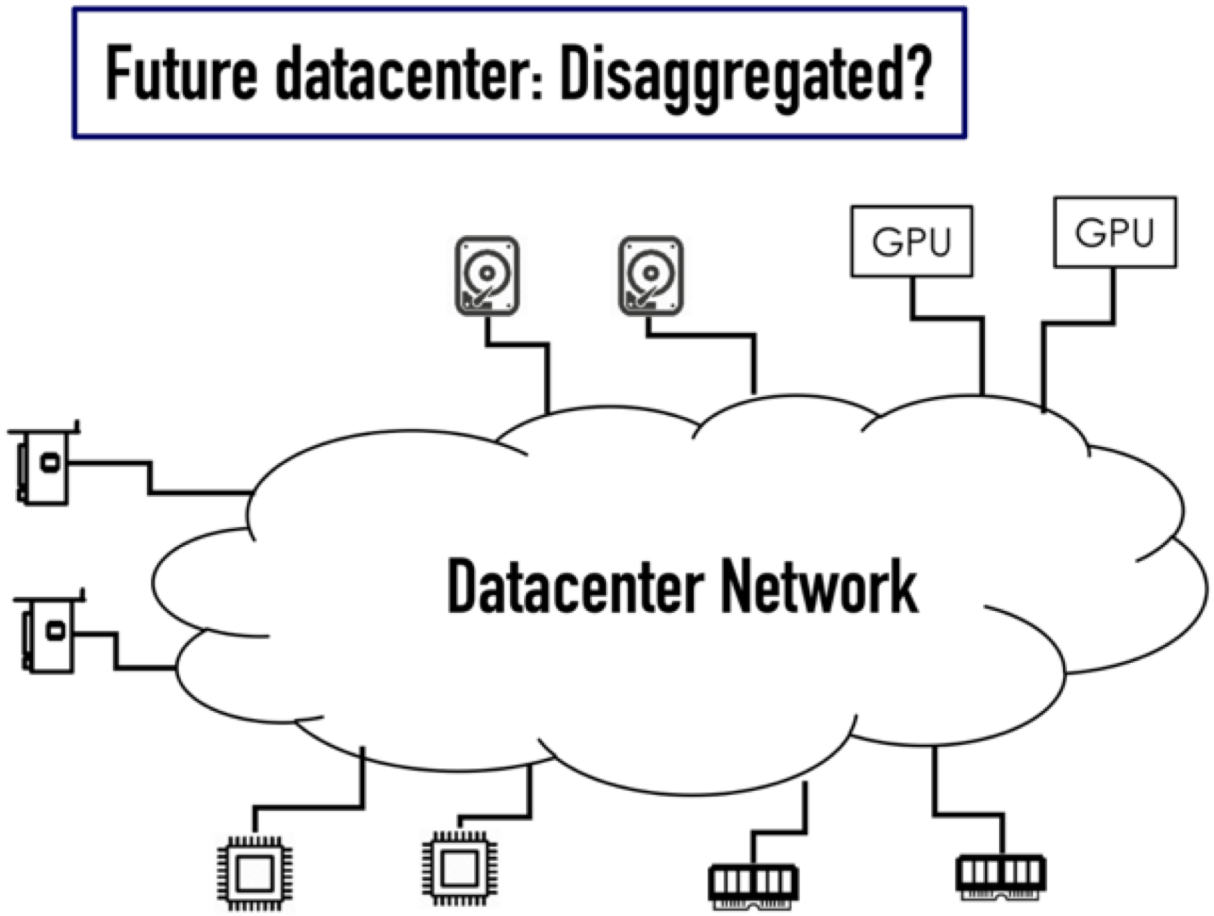

4. Cloud Architecture – FaaS

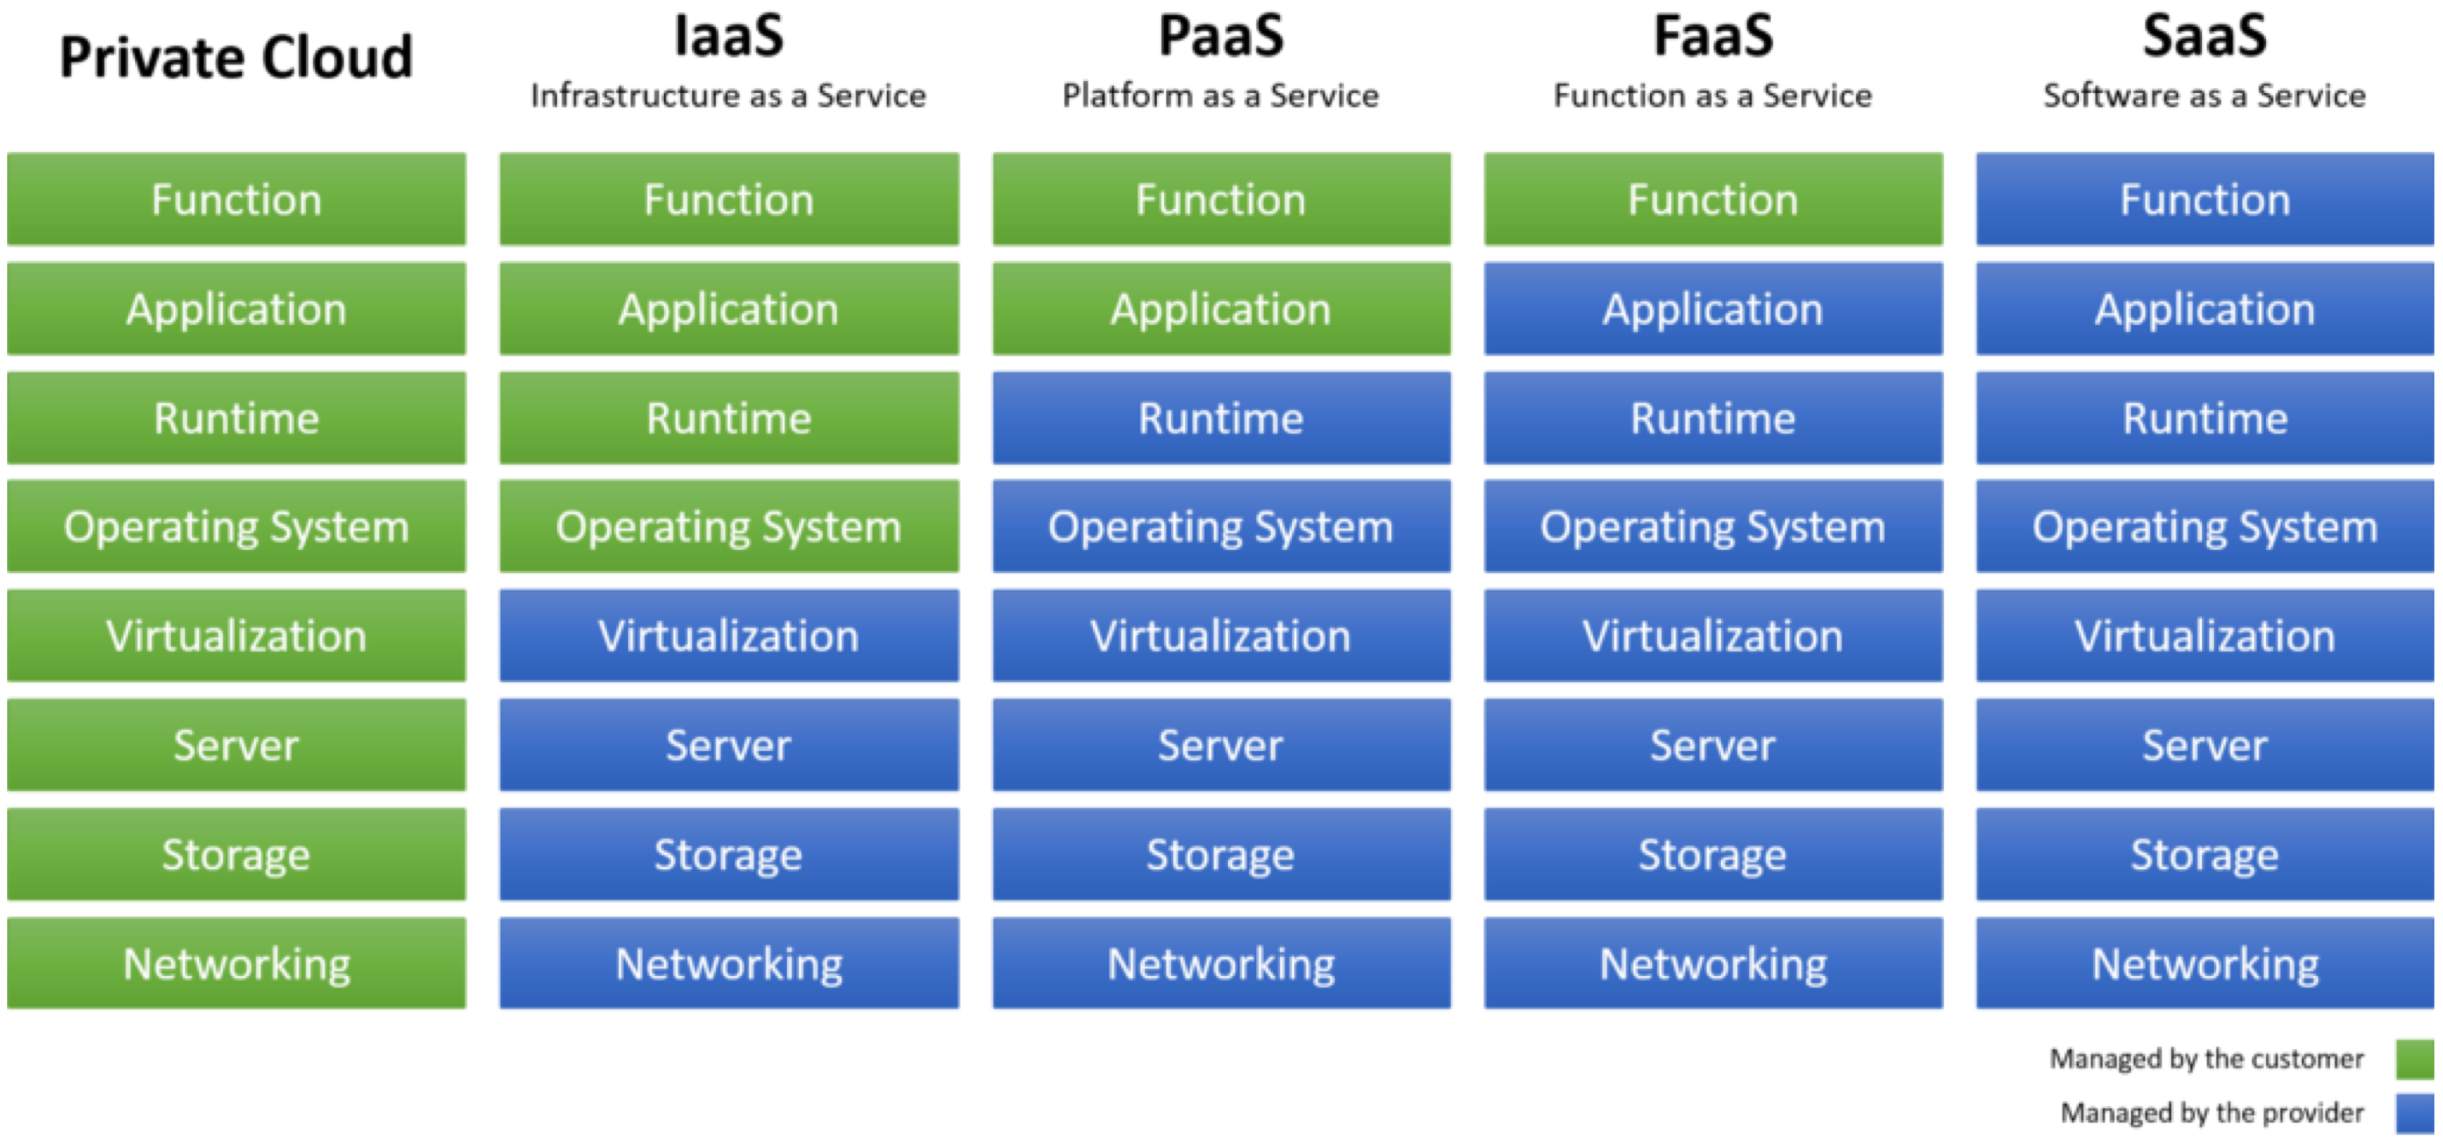

Next-generation databases

1. New processing units:

Multicore


GPU


FPGA

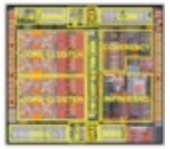
Accelerator

2. New memory/storage

SSD


NVM

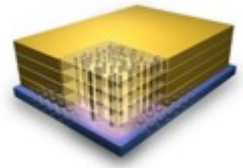
HBM

Next-generation databases have new hardware and system architecture

RDMA


SmartNIC

3. New network technology

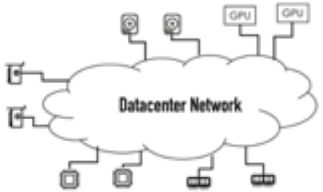

Disaggregation


FaaS

4. Cloud architecture

Goals

If you work on databases:

Take this course to learn future database systems/hardware

If you work on computer architecture:

Take this course to get familiar with an important application

Otherwise:

Take this course to learn both fields

Grading

- Paper review: 20%
- In-class discussion: 20%
- Project proposal: 15%
- Project final report: 30%
- Project presentation: 15%

Lecture Format

Syllabus:

pages.cs.wisc.edu/~yxy/cs839-s20/

Reading: 1 paper per lecture (can skip 3 times)

Upload review to https://wisc-cs839-ngdb20.hotcrp.com before 9am

BONUS: review for optional papers

40 min: Instructor presents the paper

30 min: Group discussion, submit discussion summary

Group Discussion

Discuss the provided topics

- What if we relax assumption X?
- What if metric Y of the hardware improves?
- How does the technique extend to application Z?

Share conclusions with the class

Summarize your discussion and upload to https://wisc-cs839-ngdb20.hotcrp.com

Brainstorm ideas for the course project

Course Project

In **groups of 2—4** students

Option 1: Research project towards top conference paper

Option 2: Survey for a particular area

A list of project ideas will be provided

Encouraged to propose your own ideas

Resources

CloudLab

https://www.cloudlab.us/signup.php?pid=NextGenDB

Chameleon

https://www.chameleoncloud.org

Email me if you need special hardware (e.g., GPU, NVM, RDMA, etc.)

Deadlines

Form groups: **Feb. 27**

Proposal due: **Mar. 10**

Paper submission: **Apr. 23**

Peer review: **Apr. 23 – Apr 30**

Presentation: **Apr 28 & 30**

Camera ready: **May 4**

Before next lecture

[optional] Submit review for

What's Really New with NewSQL?