CS 839: Cloud-Native Database Systems

Lecture 1: Introduction

Xiangyao Yu

09/06/2023

Who am I?

Name: Xiangyao Yu

Assistant professor in Computer Sciences, Database Group

Research interests: **database systems**

- Cloud-native database
- New hardware for database
- Transactions and HTAP

Basic Information

Course website: https://pages.cs.wisc.edu/~yxy/cs839-f23

Instructor: Xiangyao Yu

Office hour:

- Monday 4:00pm–5:00pm CS 4361

Piazza for discussions and questions

- You will be automatically enrolled in piazza once enrolled in the class

Agenda

Cloud database overview

Course logistics

Databases Moving to the Cloud

According to Gartner Report [1]

$39.2 billion, 49% of all DBMS revenue from cloud in 2021

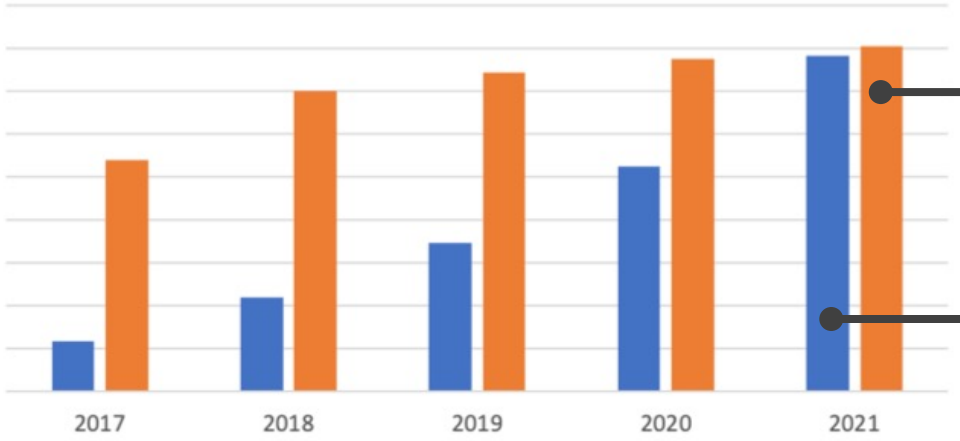

Cloud vs. On-premises Revenue

On-prem database
$40.8B in 2021

Cloud database
$39.2B in 2021

[1] DBMS Market Transformation 2021: The Big Picture, https://blogs.gartner.com/merv-adrian/2022/04/16/dbms-market-transformation-2021-the-big-picture/

Databases Moving to the Cloud

According to Gartner Report [1]

$39.2 billion, 49% of all DBMS revenue from cloud in 2021

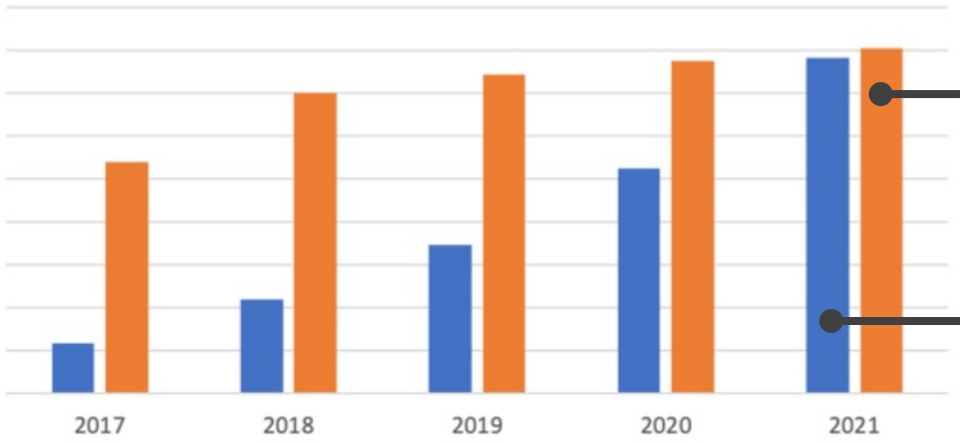

Cloud vs. On-premises Revenue

On-prem database
$40.8B in 2021

Cloud database
$39.2B in 2021

Low Cost

Elasticity

Availability

[1] DBMS Market Transformation 2021: The Big Picture, https://blogs.gartner.com/merv-adrian/2022/04/16/dbms-market-transformation-2021-the-big-picture/

Databases Moving to the Cloud

Transactional DB

Analytical DB

Cloud Computing

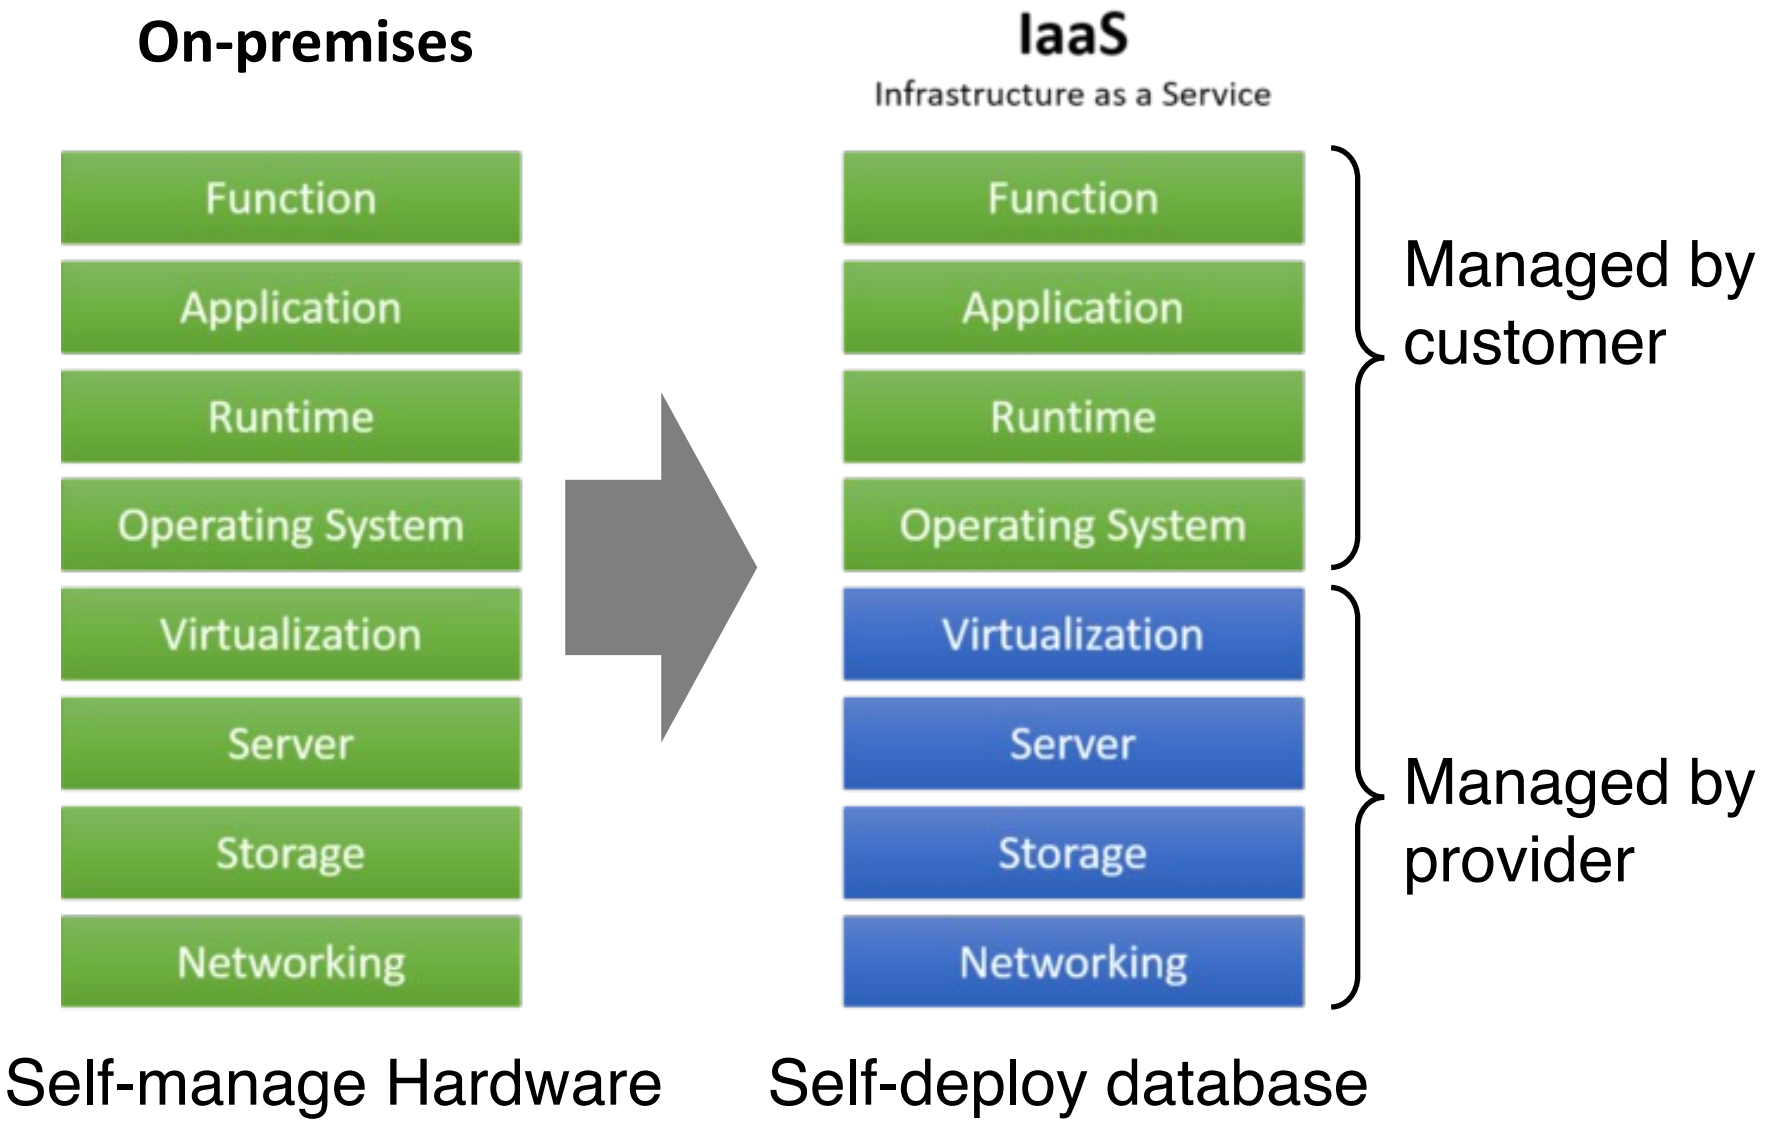

Cloud Computing

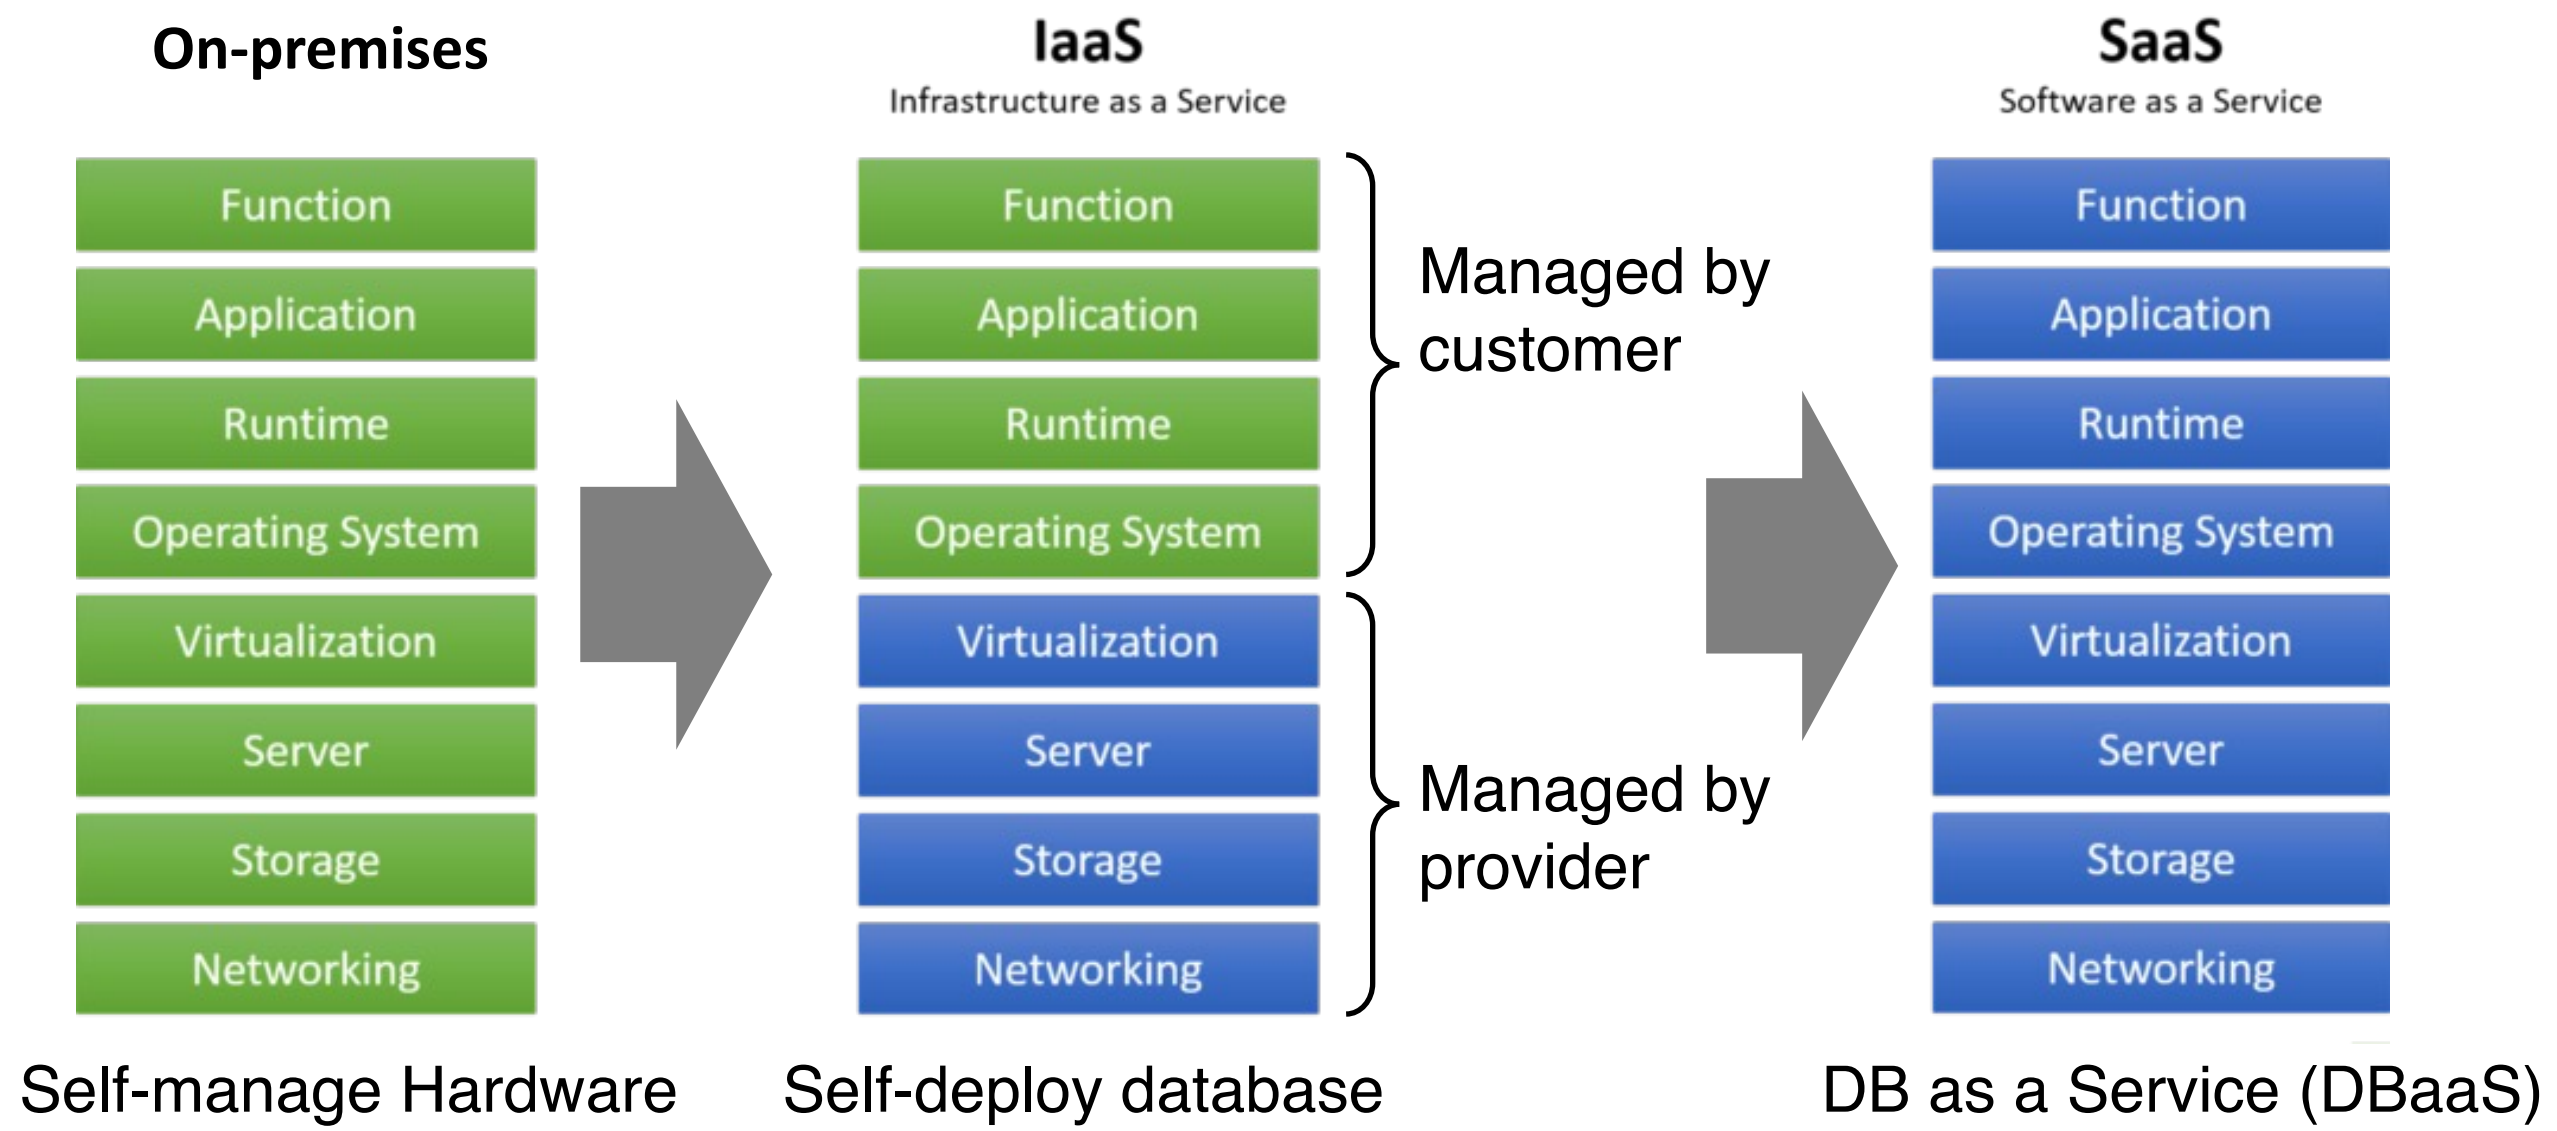

Cloud Databases

Cloud DB 1.0

- On-prem DB architecture in cloud Virtual Machines (VMs)

Cloud Databases

Cloud DB 1.0

- On-prem DB architecture in cloud Virtual Machines (VMs)

Cloud-native databases (Cloud DB 2.0)

- Storage disaggregation architecture
- Serverless
- Autoscaling

Cloud Databases

Cloud DB 1.0

- On-prem DB architecture in cloud Virtual Machines (VMs)

Cloud-native databases (Cloud DB 2.0)

- Storage disaggregation architecture
- Serverless
- Autoscaling

Next-gen cloud database (Cloud DB 3.0)?

- New hardware
- Multi-cloud
- Auto-tuning
- HTAP

New Challenges in Cloud Databases

New Requirements

- Geo-distribution
- High availability
- Low cost
- Elasticity
- Autoscaling

Higher design complexity

Solution: Modularity in distributed system design

Modular Distributed System Design

Conventional distributed system architecture

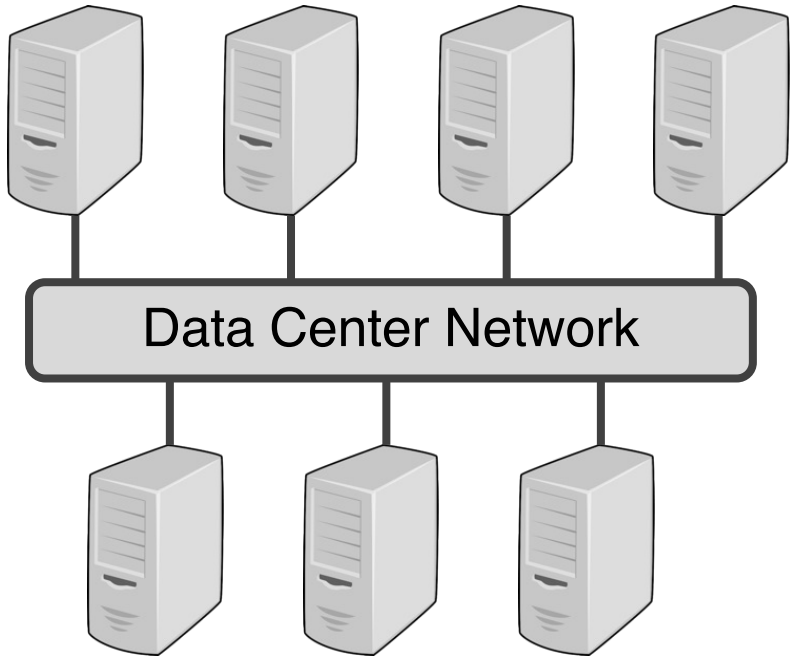

Modular Distributed System Design

Conventional distributed system architecture

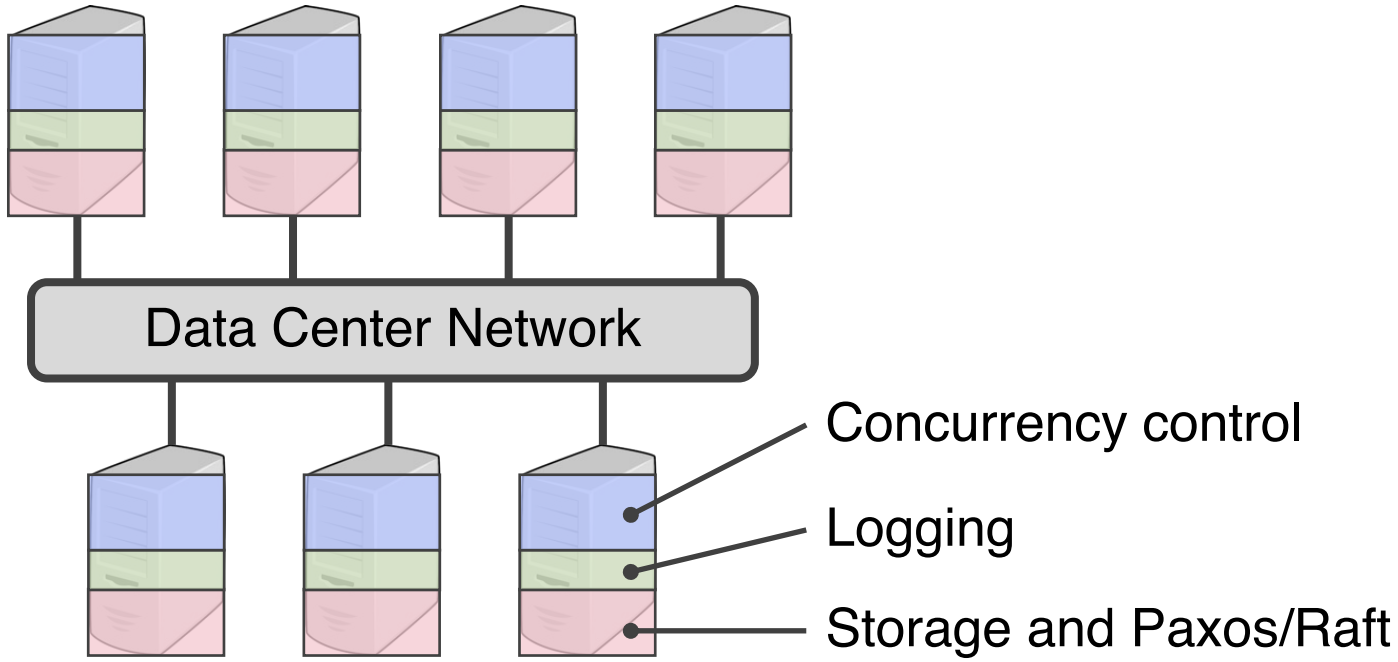

Modular Distributed System Design

Conventional distributed system architecture

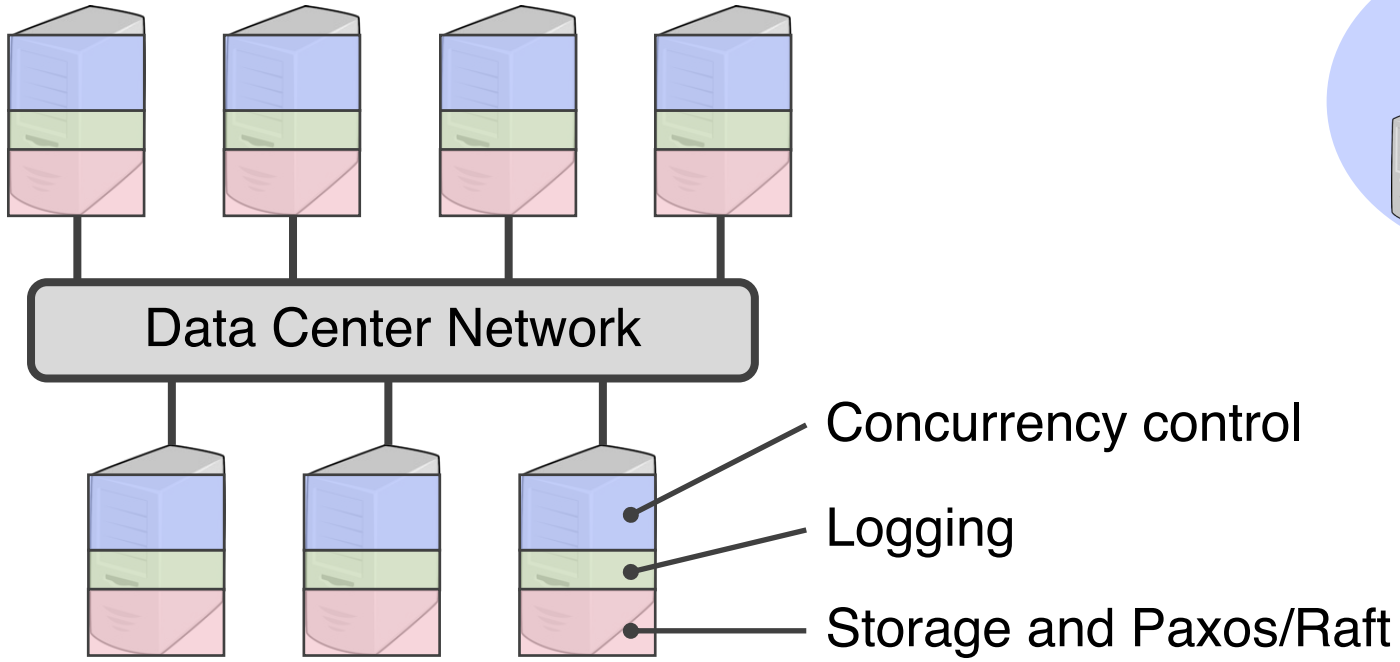

Disaggregated distributed system architecture

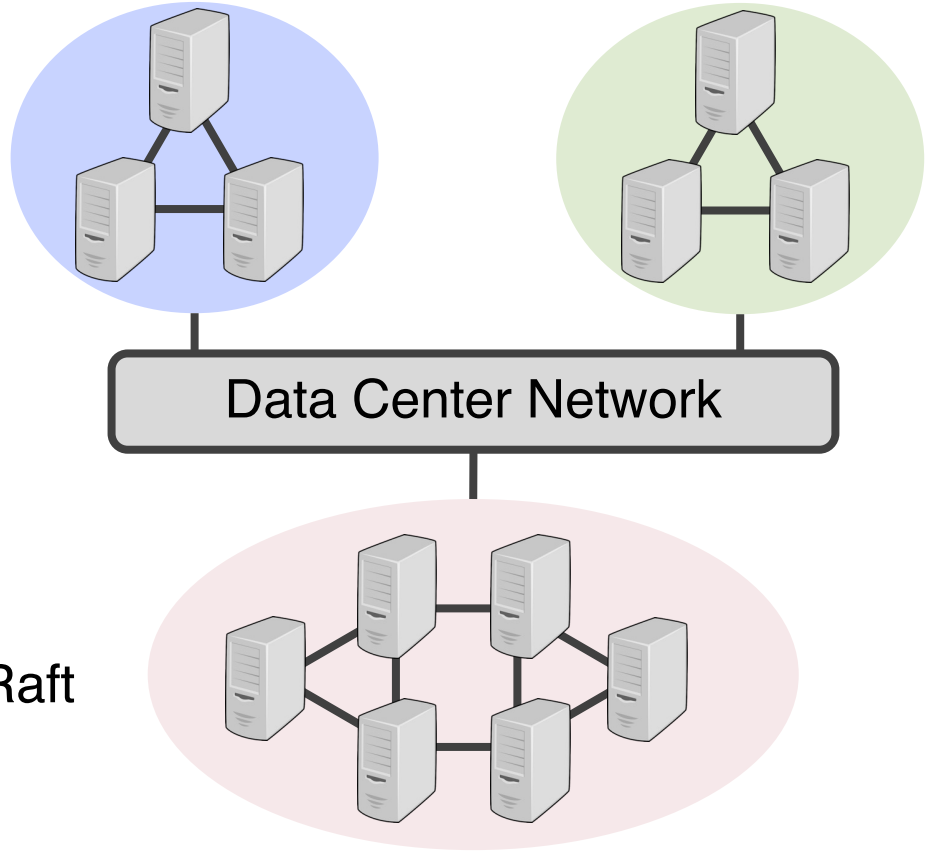

Each service is deployed as a separate distributed cluster

Disaggregated Distributed System

Advantages

- **Scalability**: Services can scale independently
- **Performance and cost**: Services can be custom optimized (e.g., low cost storage service)
- **Separation of concerns**: Services can be independently developed

Disadvantage

- Network can throttle performance

Disaggregated distributed system architecture

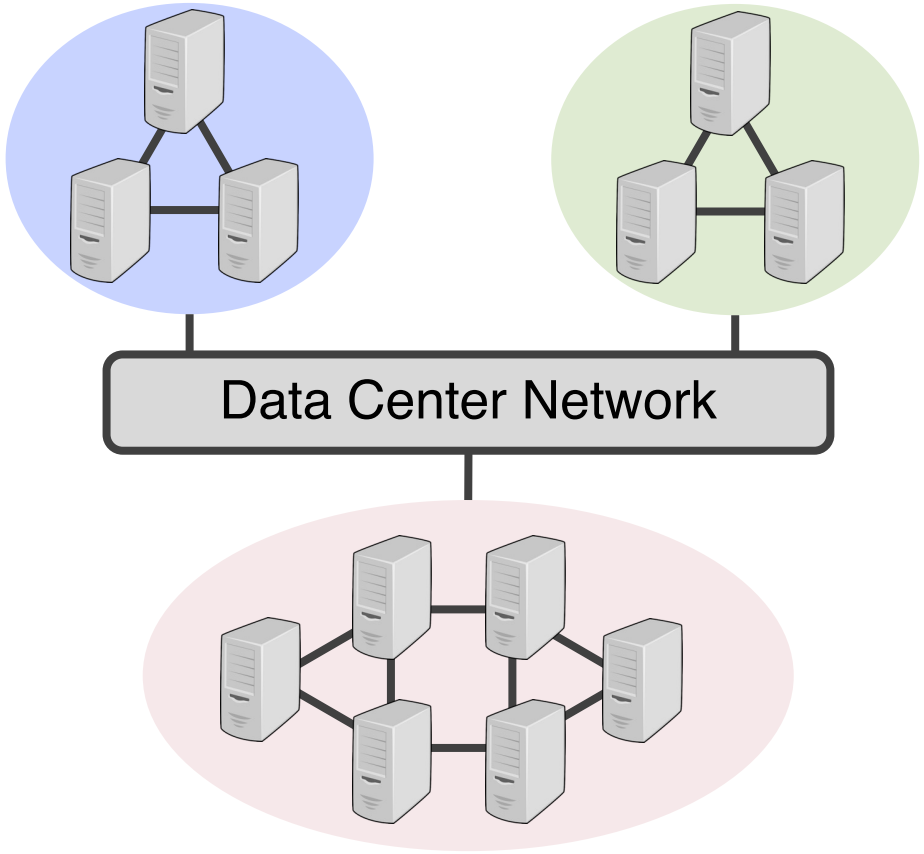

Cloud DB: Storage-Disaggregation

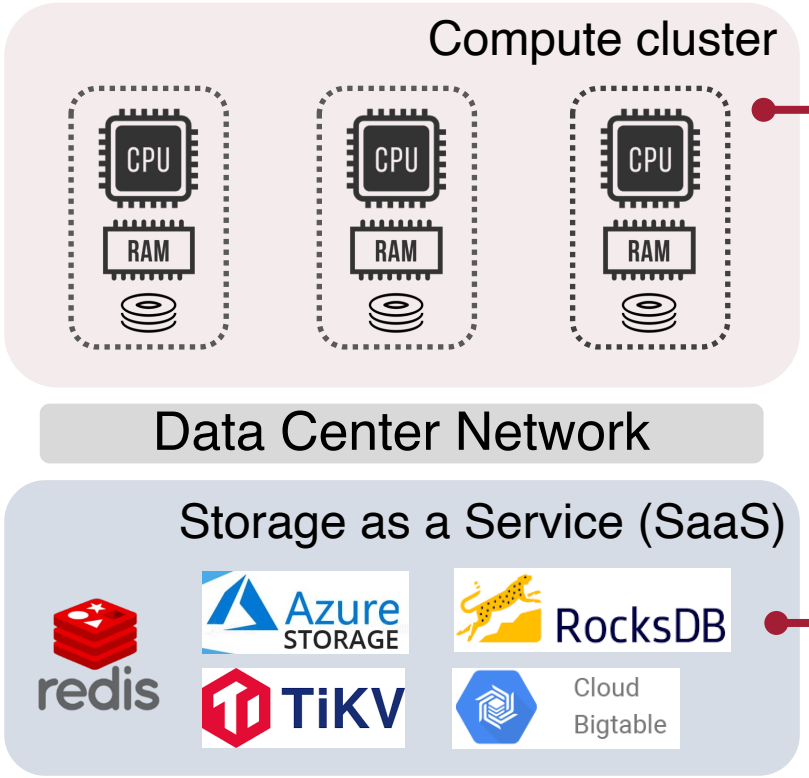

Database logic in computer cluster

Database states (e.g., tables and log) in cloud storage service

Manage computation and storage as separate services

Advantages of Storage-Disaggregation

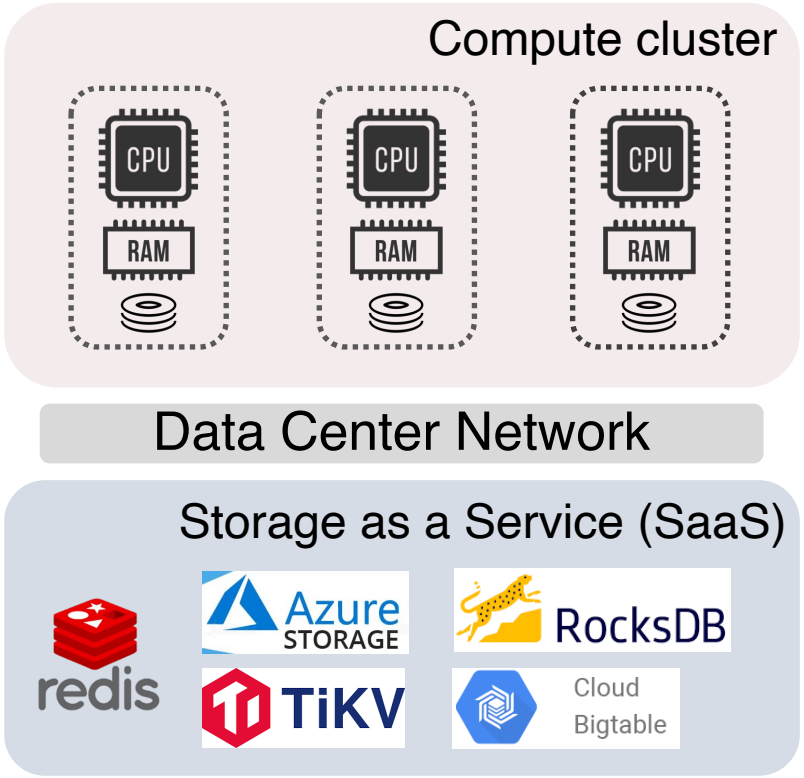

Advantage #1: Elasticity

- Compute and storage resources can scale independently

Advantages of Storage-Disaggregation

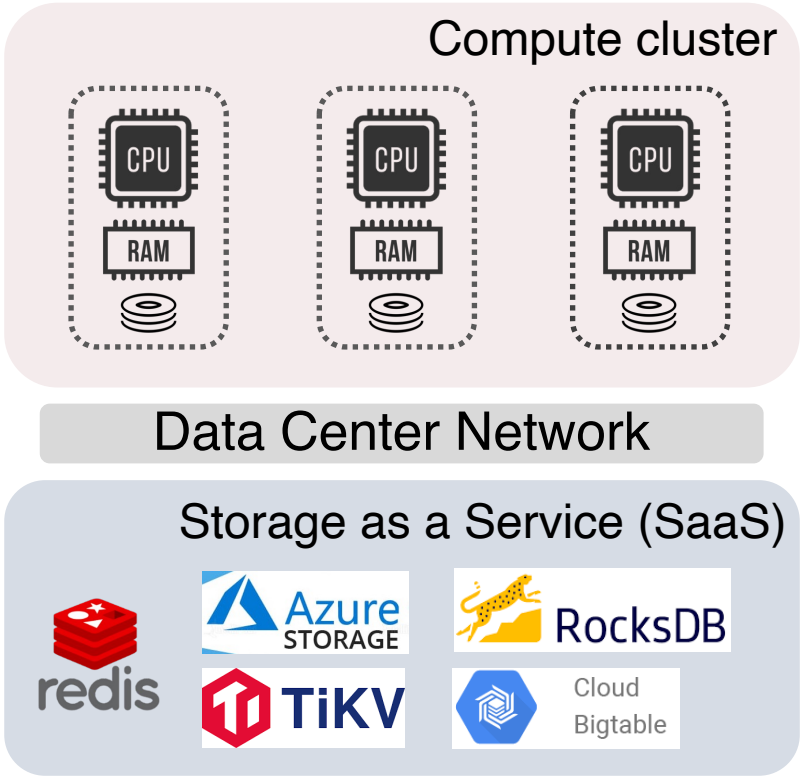

Advantage #1: Elasticity

– Compute and storage resources can scale independently

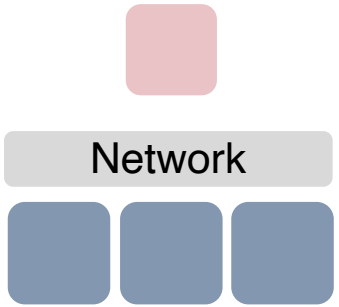

Storage-intensive workload

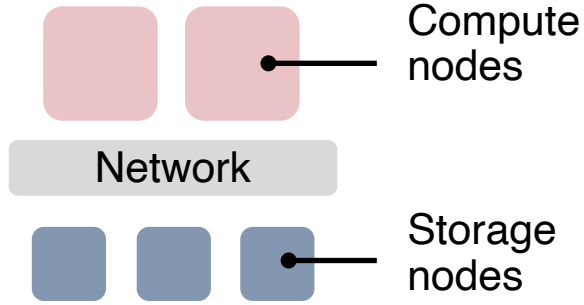

Compute-intensive workload

Advantages of Storage-Disaggregation

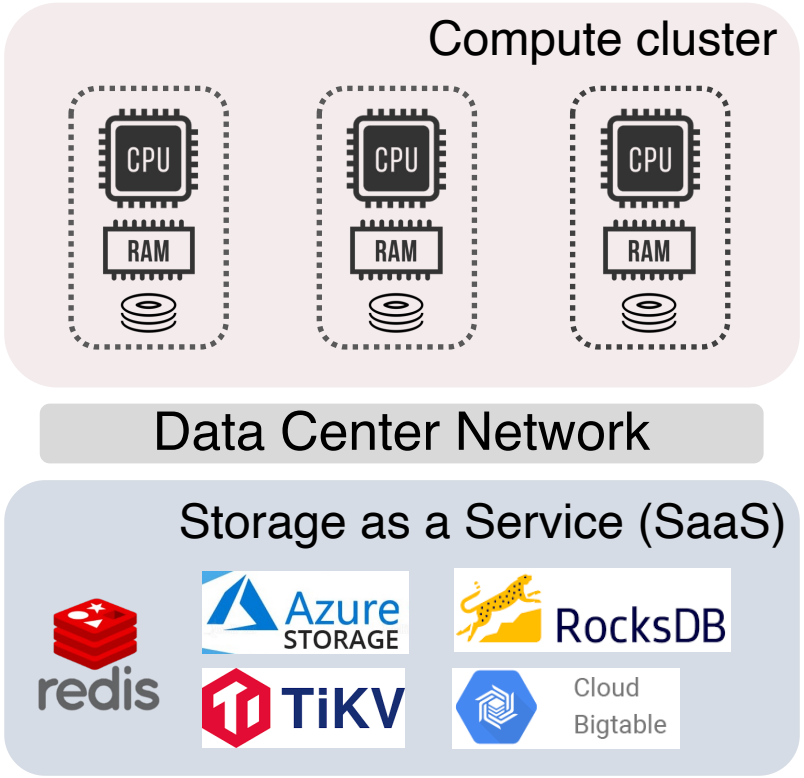

Advantage #2: Low Cost

– Storage service can be much cheaper than compute servers

S3 storage price	**$0.02** per GB per month
16 vCPU Virtual Machine	**$0.5** per hour per VM

Advantages of Storage-Disaggregation

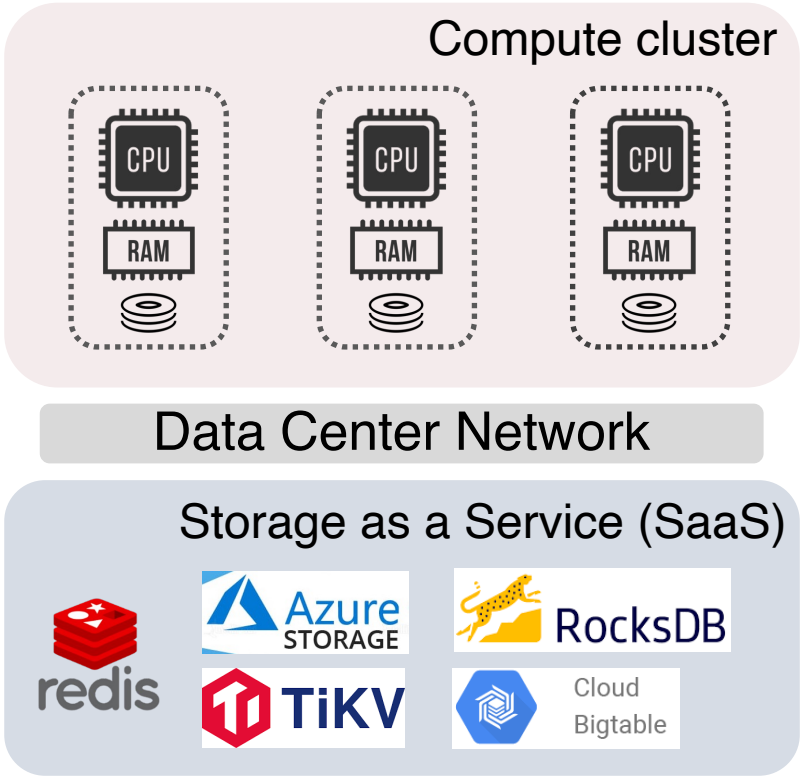

Advantage #2: Low Cost

– Storage service can be much cheaper than compute servers

S3 storage price	**$0.02** per GB per month
16 vCPU Virtual Machine	**$0.5** per hour per VM

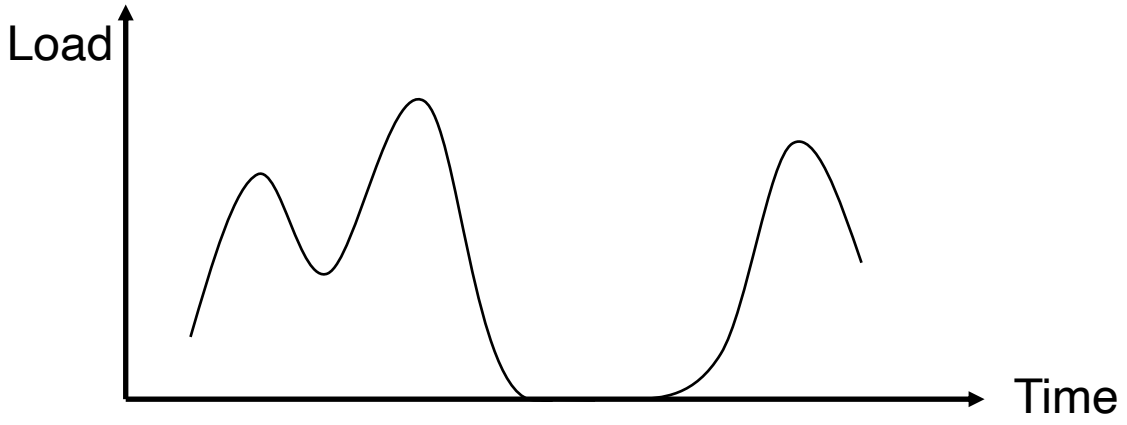

Advantages of Storage-Disaggregation

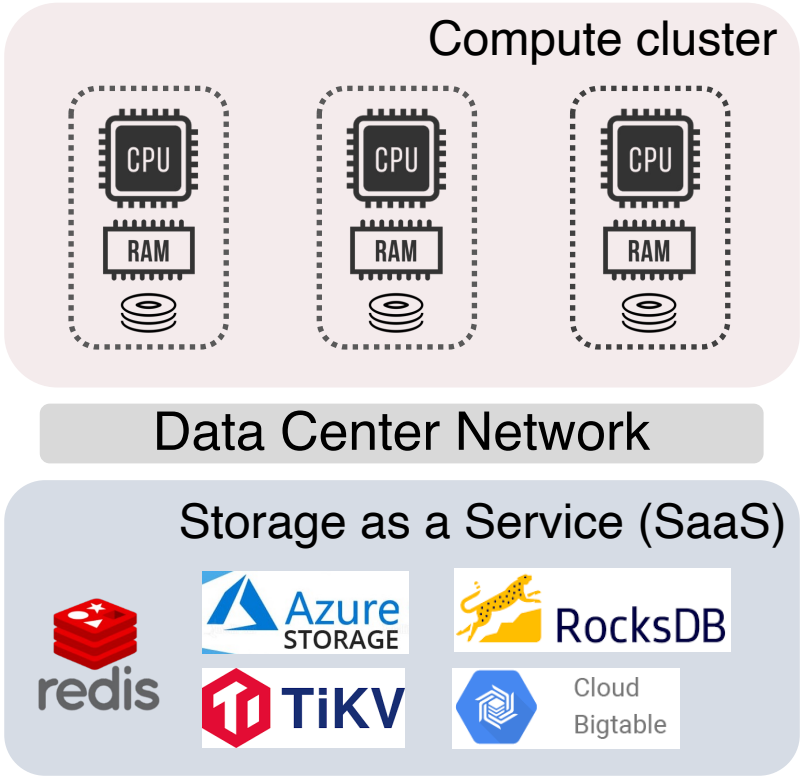

Advantage #2: Low Cost

– Storage service can be much cheaper than compute servers

S3 storage price	**$0.02** per GB per month
16 vCPU Virtual Machine	**$0.5** per hour per VM

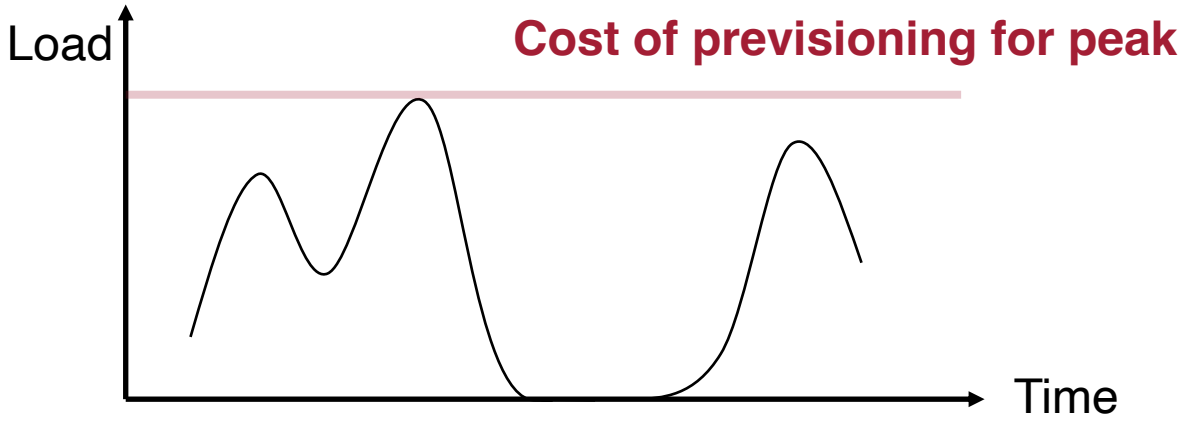

Advantages of Storage-Disaggregation

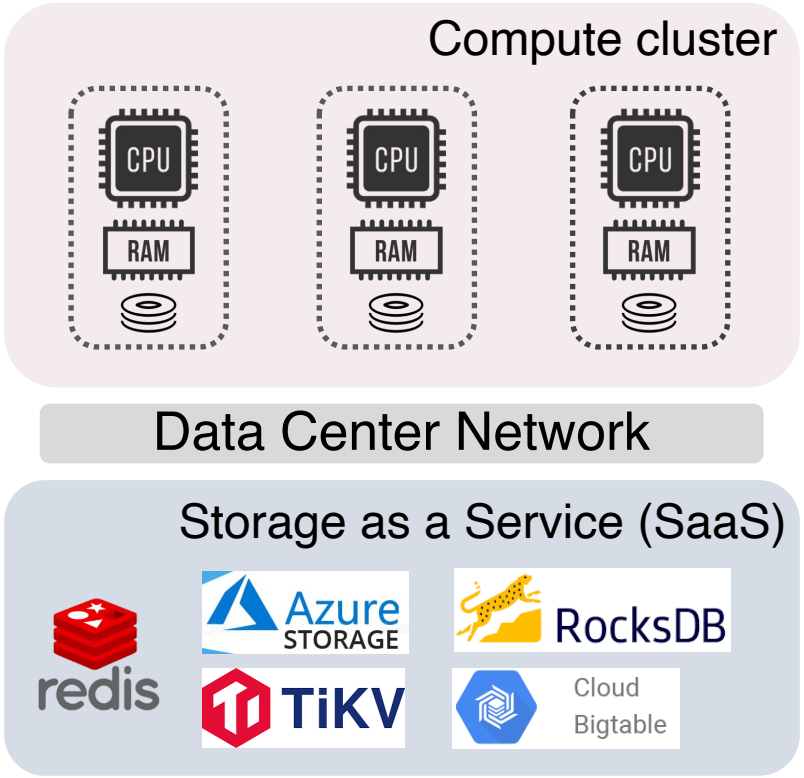

Advantage #2: Low Cost

– Storage service can be much cheaper than compute servers

S3 storage price	**$0.02** per GB per month
16 vCPU Virtual Machine	**$0.5** per hour per VM

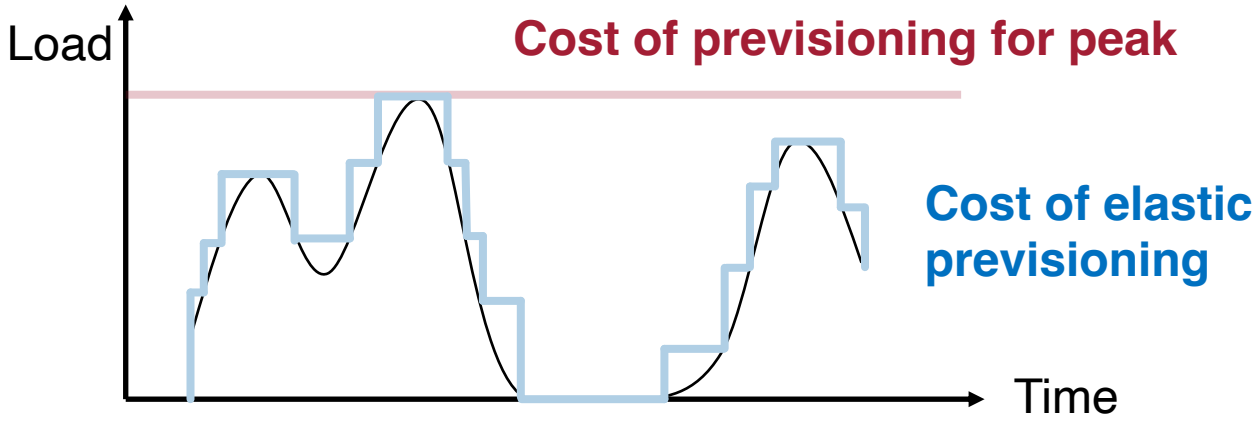

Advantages of Storage-Disaggregation

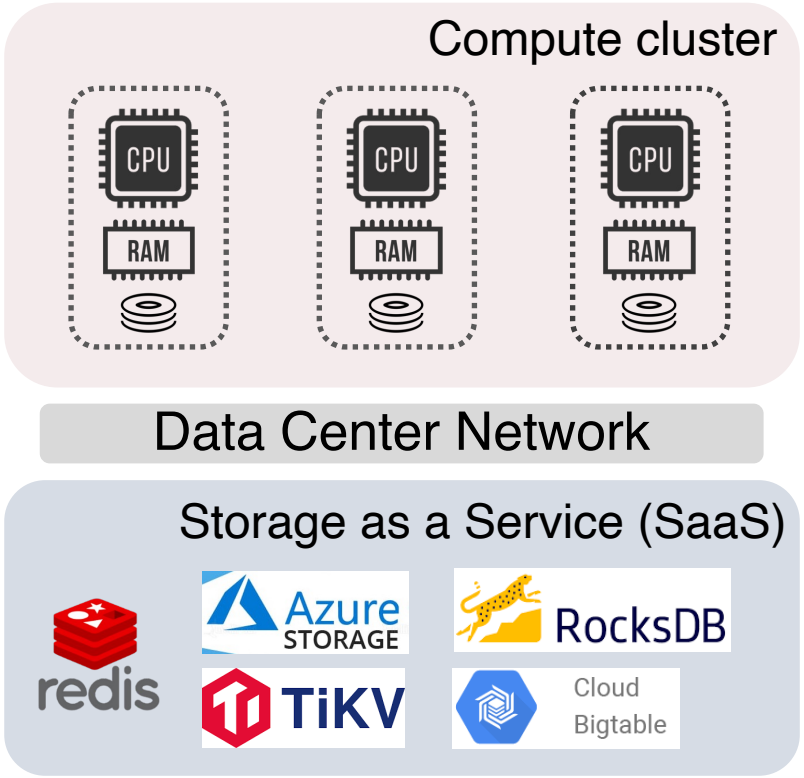

Advantage #2: Low Cost

– Storage service can be much cheaper than compute servers

S3 storage price	**$0.02** per GB per month
16 vCPU Virtual Machine	**$0.5** per hour per VM

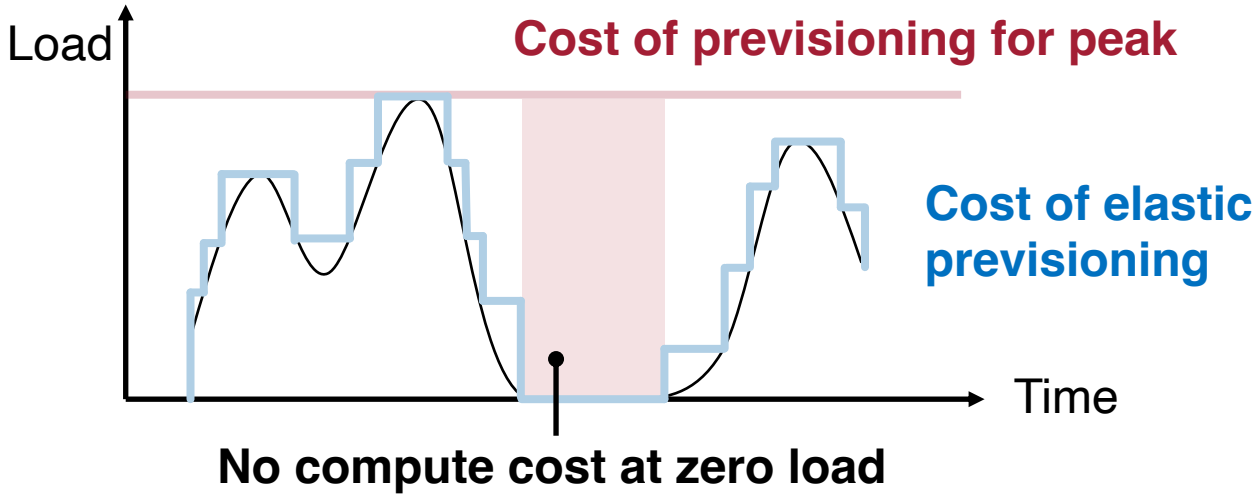

Advantages of Storage-Disaggregation

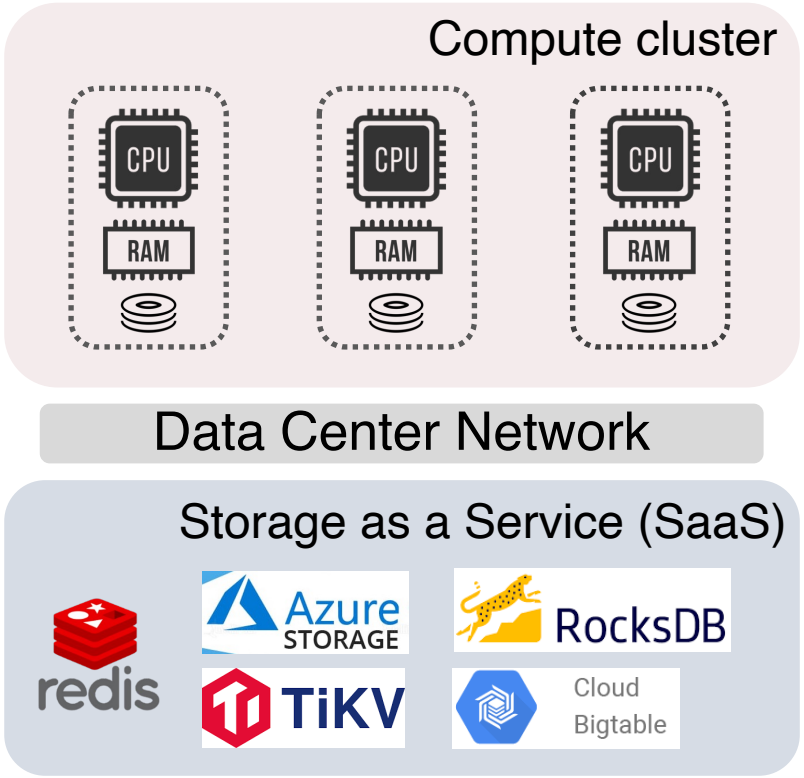

Advantage #3: Availability

- Storage service provides high availability through geo-replication
- Leverage high availability in storage to simplify fault tolerance in database

Separation of concerns between database logic and replication protocols

Advantages of Storage-Disaggregation

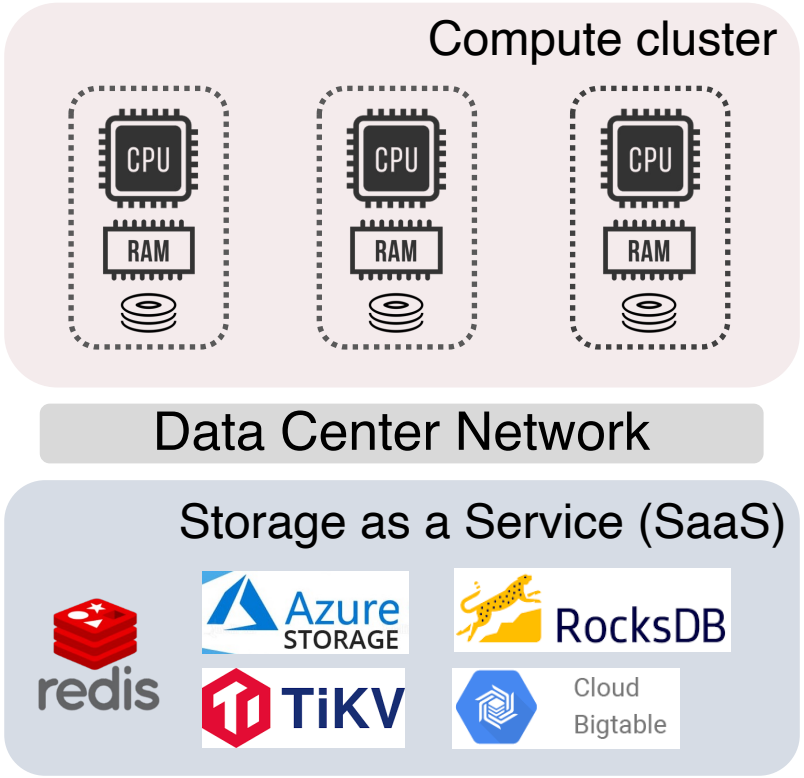

Advantage #1: Elasticity

Advantage #2: Low Cost

Advantage #3: Availability

Storage-disaggregation architecture widely deployed in cloud databases

Amazon Aurora · POLARDB · Microsoft SQL Server · GaussDB · amazon REDSHIFT · Cloud Spanner · TiDB · yugabyteDB · Cockroach DB · Google BigQuery · Azure Cosmos DB · Spark SQL · snowflake · HIVE · presto

Redesign databases in storage-disaggregation architecture

Course Content

Storage disaggregation for analytics
Storage disaggregation for transactions

Course Content

Storage disaggregation for analytics
Storage disaggregation for transactions
Serverless
DBOS
Auto-scaling
Auto-tuning
Multi-cloud

Course Content

Storage disaggregation for analytics
Storage disaggregation for transactions
Serverless
DBOS
Auto-scaling
Auto-tuning
Multi-cloud
HTAP
New hardware
- Memory disaggregation
- GPU
- RDMA
- SmartNIC

Agenda

Cloud database overview

Course logistics

Course Information

Course Website: http://pages.cs.wisc.edu/~yxy/cs839-f23/

Canvas page: https://canvas.wisc.edu/courses/375821

Piazza: https://piazza.com/class/lm5r0a40r6s1bd

- Can be accessed directly through canvas

Prerequisite: CS 564 or equivalent

Lecture Structure

Submit paper review before lecture

- If multiple papers are provided, pick your favorite one

Mixture of presentations from instructor and students

- 2-3 student presentations for most lectures
- A presentation signup sheet will be provided before next lecture

Q/A and in-class group discussion

Summarize group discussion and submit to hotcrp

- Each student needs to write 3-5 summaries during the entire semester.

Grading

Paper review: 25%

Class participation: 25%

Project proposal: 10%

Project presentation: 10%

Project final report: 30%

Paper Review (25%)

Paper reading: pick one paper to read per lecture

– username: cs839 password: dbguru

Paper Review (25%)

Paper reading: pick one paper to read per lecture

- username: cs839 password: dbguru

Upload review: https://tbd.hotcrp.com (must submit before the lecture starts in order to be graded)

- Overall merit
  - Accept or reject?
- Paper summary
  - What main research problem/challenge did the paper address?
  - What is the key contribution of the paper?
- **Comments and questions**
  - Aspects you like or dislike about the paper
  - Questions about that paper that you wish to be discussed in lecture

Paper Review (25%)

Paper reading: pick one paper to read per lecture

- username: cs839 password: dbguru

Upload review: https://tbd.hotcrp.com (must submit before the lecture starts in order to be graded)

- Overall merit
  - Accept or reject?
- Paper summary
  - What main research problem/challenge did the paper address?
  - What is the key contribution of the paper?
- **Comments and questions**
  - Aspects you like or dislike about the paper
  - Questions about that paper that you wish to be discussed in lecture

Grading: You can skip up to 2 reviews without losing points; otherwise 1% of total grade (up to 25%) is deducted for each missing review

Class Participation

2-3 student presentations in most classes

Summary of group discussion

Course Project (50%)

In **groups of 2–4** students

Example project ideas will be provided but you are encouraged to propose your own ideas

- Project ideas for CS 764 2020–2022 are available on the course website
- Example CS 764 projects available on the course website (two papers based on course projects have been accepted to SIGMOD 2022 and SIGMOD 2023)

Course Project (50%)

In **groups of 2–4** students

Example project ideas will be provided but you are encouraged to propose your own ideas

- Project ideas for CS 764 2020–2022 are available on the course website
- Example CS 764 projects available on the course website (two papers based on course projects have been accepted to SIGMOD 2022 and SIGMOD 2023)

Important dates

- Discuss project ideas with instructor: Oct. 9
- Create teams and submit proposal: Oct. 16
- Project meetings with instructor: TBD
- Presentation: Dec. 11 & 13
- Paper submission: Dec. 18

Computation Resources

CloudLab

https://www.cloudlab.us/signup.php?pid=NextGenDB (project name: NextGenDB)

Chameleon

https://www.chameleoncloud.org (project name: ngdb)

Waitlist

If you are enrolled but don’t want to take the class, please drop ASAP so that students on the waitlist can be enrolled

We manage the waitlist first-come-first-serve

Before next lecture

Read the following paper and **submit review**

- Alexandre Verbitski, et al., Amazon Aurora: Design Considerations for High Throughput Cloud-Native Relational Databases. SIGMOD, 2017
- The review website will be ready (hopefully) in a few days