



# CS 839: Topics in Database Management Systems

## Lecture 19: Multi-Cloud

Xiangyao Yu

11/8/2023

# Logistics

---

Mon 11/13. Project meetings (No lecture)

- Meeting (optional) with the instructor to discuss the course project
- Location: CS4361. [Signup sheet](#)

# Group Discussion from Lecture #17

---

Auto-scaling has two aspects: (1) mechanism (e.g., how to add a new process without affecting correctness) and (2) policy (e.g., when should a new process be added or removed). Pick one paper you have read; how does it contribute to both aspects?

What do you see as the biggest challenge to enable an auto-scaling database? You can consider either mechanism or policy in either OLTP or OLAP setting.

# Group Discussion from Lecture #17

---

## Mechanism vs. Policy in Anna

- Mechanism: when adding a new process, update the hash ring. Node sinter-communicate to transfer data
- Policy: spawn new processes when durability constraint is not met or when a node fails.

# Group Discussion from Lecture #17

---

## Challenges to enable an auto-scaling database?

- Scaling against varying transaction rates
- Data consistency across nodes
- Transaction availability
- ACID
- Avoid hot spots

# Discussion Questions

---

Is storage disaggregation still a good architecture in a multi-cloud database? What enhancement would you propose (for OLTP and OLAP) to further improve the performance and cost efficiency for such a database?

What do you think are the new challenges in deploying a multi-cloud or multi-datacenter database?

Submit by **11:59pm CT on Thursday** (11/9). Title starts with “[Discussion L19]”