WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

CS 839: Topics in Database Management Systems

Lecture 16: DBOS

Xiangyao Yu

10/30/2023

Group Discussion from Lecture #15

Today we program the cloud through a service abstraction—developers call APIs of different services (e.g., DB, storage, cloud function, containers, spark, etc.). This approach may not scale as the number of services increase in the future.

- How would developers program the cloud differently from today in ten years?
- What new abstraction are we going to create on top of existing services?
- Does it make sense to build an OS for the entire cloud?

Group Discussion from Lecture #15

How would developers program the cloud in the future?

- OS for lambda. E.g., use ‘fork’ to create new process running in lambda function
- Lambda functions are stateful with fast shared storage
- Program using natural language
- True autoscaling to 0 and to infinity.

Group Discussion from Lecture #15

New abstraction

- Serverless DB + cloud functions
- Programming language for the cloud
- Declarative models
- Multi-cloud abstraction

Group Discussion from Lecture #15

OS for the cloud?

- Extending Linux to the cloud setting.
- A cloud OS may introduce performance overhead
- Functions: Serverless threads, dynamic scheduling, optimized storage abstraction (virtualized storage?), virtual memory, specialized programming language

Discussion Questions

Imagine you are building a cloud OS following the architecture of a cloud database + some cloud runtime (e.g., Lambda function or Kubernetes). Please consider the following questions:

- What can be a good initial killer app for such a system?
- Are there key missing technologies before we can build such a system? If so, what are they?
- Which existing systems/services would you use as building blocks for your initial prototype?

Submit by **11:59pm CT of Tuesday** (10/31). Title starts with "[Discussion L16]"