

# CS15-319 / 15-619

## Cloud Computing

Recitation 10

March 26<sup>th</sup>, 2013

# Announcements

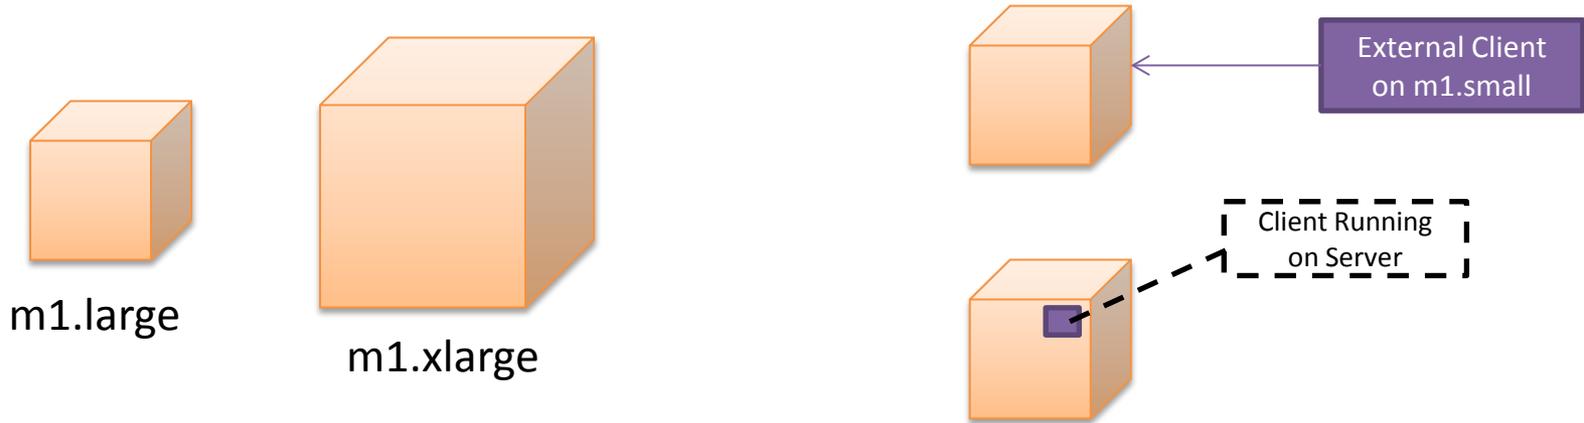
- **Open up S3 location of hand ins:**
  - Give access to your S3 bucket to:
    - public
    - [onlinecloudcomputingcourse@gmail.com](mailto:onlinecloudcomputingcourse@gmail.com)
  - You could lose credit or be penalized otherwise
  - See Piazza Post on how to open up your handin directory
- Encounter a general bug:
  - Post on Piazza
- Encounter a grading bug:
  - Post Privately on Piazza
- Post feedback on OLI

# Project 3 Student Progress

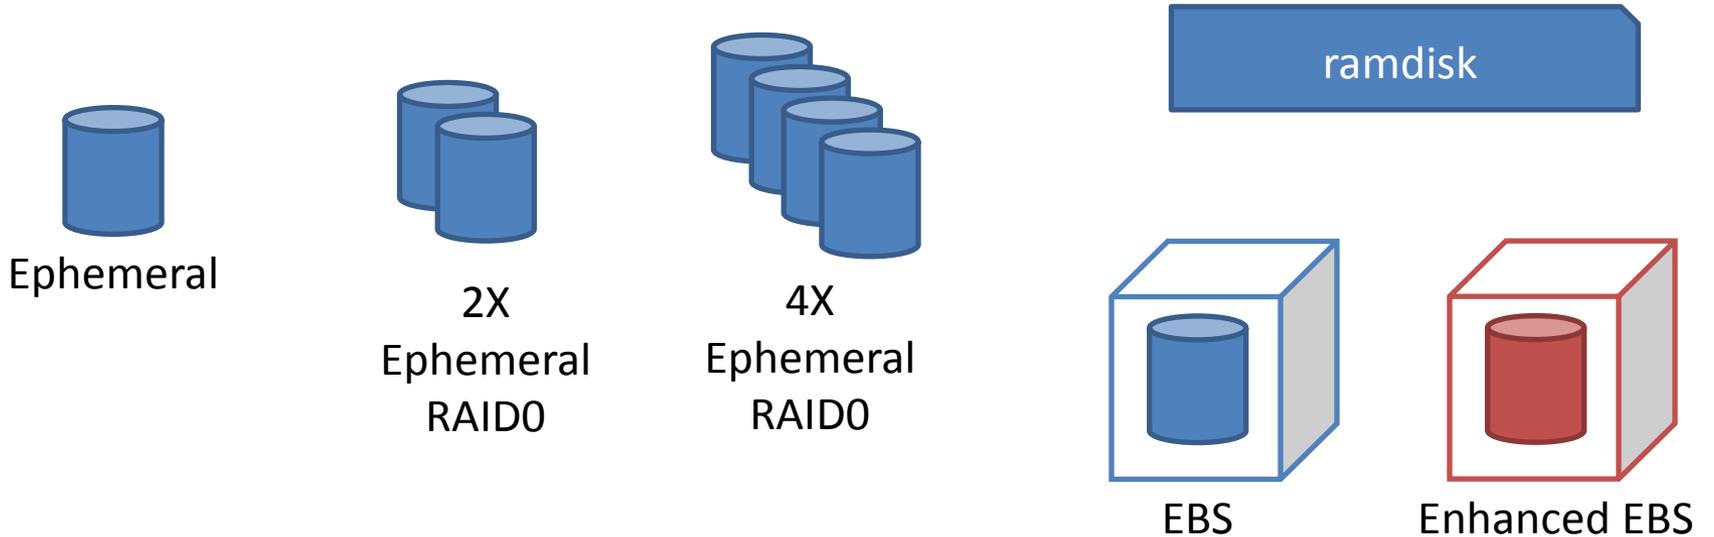
- Part 1: File vs. Database
  - 43/44 Students Completed (98%)
- Part 2: Vertical Scaling
  - 42/44 Students Completed (95%)

# Vertical Scaling

## Instances and Clients



## Storage Options



# Vertical Scaling Results

Scenario	Server Instance Type	Server Storage Type	Run test from	Average TPS for the Class
1	m1.large	ramdisk	the server	499.72
2	m1.large	ramdisk	m1.small client	396.46
3	m1.large	Ephemeral	m1.small client	119.83
4	m1.large	2 Ephemeral RAID0 stripe	m1.small client	188.80
5	m1.large	EBS	m1.small client	375.43
6	m1.large	EBS Optimized	m1.small client	377.12
7	m1.xlarge	ramdisk	the server	1060.19
8	m1.xlarge	ramdisk	m1.small client	449.92
9	m1.xlarge	4 Ephemeral RAID0 stripe	m1.small client	280.66
10	m1.xlarge	EBS	m1.small client	401.33
11	m1.xlarge	EBS Optimized	m1.small client	411.36

- Local Tests (from the server) always faster than Remote Tests (from an m1.small client)
- EBS faster than Ephemeral (Instance Store) and RAID0 Ephemeral
- EBS is almost as fast as ramdisk
- Optimized EBS does not necessarily provide better performance for this application.

# Note on Our Database Experiments

- Synthetic Benchmark
  - No Memory Caching
  - No Index of Tables
  - Stress the disk to see differences between the storage back-ends
- Not an indicator of real world performance!
  - Typically in-memory cache and indexing will give you speeds similar to ramdisk, for any storage backend.
  - Database performance is not purely based on storage performance
    - CPU, Memory, Network and the type of Queries all affect TPS.

# New Modules

- Unit 4 – Cloud Storage:
  - Cloud Storage
    - The Big Picture of Cloud Storage
    - Data and its Abstractions
    - Abstractions in Storage : File Systems
    - Abstractions in Storage : Databases
    - Abstractions in Storage : Objects
    - Block Devices
  - Case Study: Distributed File Systems
    - The Hadoop Distributed File System (HDFS)
    - The Parallel Virtual File System (PVFS)
    - HDFS vs. PVFS

# Project 3 Part 3

- Project 3
  - Part 1
    - Files vs. Databases
  - Part 2
    - Vertical Scaling in Databases
  - Part 3
    - Horizontal Scaling in Databases

# Discussions

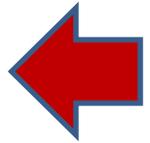
- Your questions...

# Upcoming Deadlines

- Unit 4:

[UNIT 4: Cloud Storage](#)

[Module 14: Case Studies: Distributed File Systems](#)



- Project 3

[UNIT 7: Project 3](#)

[Module 23: Horizontal Scaling in Databases](#)

Horizontal Scaling

[Checkpoint](#)

**Available Now**

Due **3/31/13** 11:59 PM

