15-319 / 15-619 Cloud Computing

Recitation 10

November 3rd and 5th, 2015

Overview

- **Administrative issues**
  Office Hours, Piazza guidelines
- **Last week's reflection**
  Project 3.3, OLI Unit 4, Module 15, Quiz 8
- **This week's schedule**
  - 15619 Project - Query 3 & 4 - November 11th
  - Quiz 9 - November 6th (Unit 4, Module 16 & 17)
  - Project 3.4 - November 8st

Reminders

- Monitor AWS expenses regularly and tag all resource
  - Check your bill (Cost Explorer > filter by tags).
- Piazza Guidelines
  - Please tag your questions appropriately
  - Search for an existing answer first
- Provide clean, modular and well documented code
  - **Large** penalties for not doing so.
- Utilize Office Hours
  - We are here to help (but not to give solutions)
- Use the team AWS account and tag the 15619Project resources carefully.

OLI Unit 4

Review

OLI Unit 4 : Review

- **M14 : Cloud Storage Overview**
- **M15 : Distributed File Systems (HDFS & Ceph)**
  - Topology and design
  - Data distribution strategies
- **M16 : Distributed Database Case Studies**
  - HBase, MongoDB, DynamoDB, Cassandra
- **M17: Cloud Object Storage**

Distributed Databases

- In 2004, Amazon.com began to experience the limits of scale on a traditional web-scale system
- Response was a highly available key-value structured storage system called Dynamo (2007)

Problem	Technique used as solution
Data Sharding	Consistent Hashing
Transient Fault Handling	Sloppy Quorum / Hinted Handoff
Permanent Failure Recovery	Anti-entropy using Merkle trees
Membership and Health Checks	Gossip protocols

- Used in S3, DynamoDB, Cassandra

Distributed Databases

- In 2006, Google published details about their implementation of BigTable
- Designed as a “sparse, distributed multi-dimensional sorted map”
- HBase stores members of “column families” adjacent to each other on the file system - columnar data store

Upcoming Deadlines

Quiz 9 : Unit 4 - Modules 16, 17

Open : 11/06/2015 12:01 AM Pittsburgh

Due : 11/06/2015 11:59 PM Pittsburgh

- Project 3.4 : Social Network with Heterogeneous DBs

  Due : 11/08/2015 11:59 PM Pittsburgh

- 15619Project : Phase 2

  Due : 11/11/2015 4:59 PM Pittsburgh

Project 3

Review

Project 3 Weekly Modules

- P3.1: Files, SQL and NoSQL
- P3.2: Sharding and Replication
- P3.3: Consistency
- P3.4: Social network with heterogeneous backends
- P3.5: Data warehousing and OLAP

Addressing Feedback: P3.1

- Most students familiar with SQL; introduced to HBase
- Complaints about performance variability in:
  - Different availability zones
  - Different experiment times
  - Different instances types

- How are HBase and MapReduce related?

- Lessons learned:
  - Don’t make us generate 100 GB of data

Addressing Feedback: P3.2 & P3.3

P3.2

- Good concurrent programming refresher
- Evaluating different hashing strategies

P3.3

- When/why is causal consistency useful?
- Test cases were too hard to pass / writeup unclear

- Suggestion: ex post facto review of best solutions

Project 3.4 : Introduction

- **Build CinemaChat, a social network about movies:**

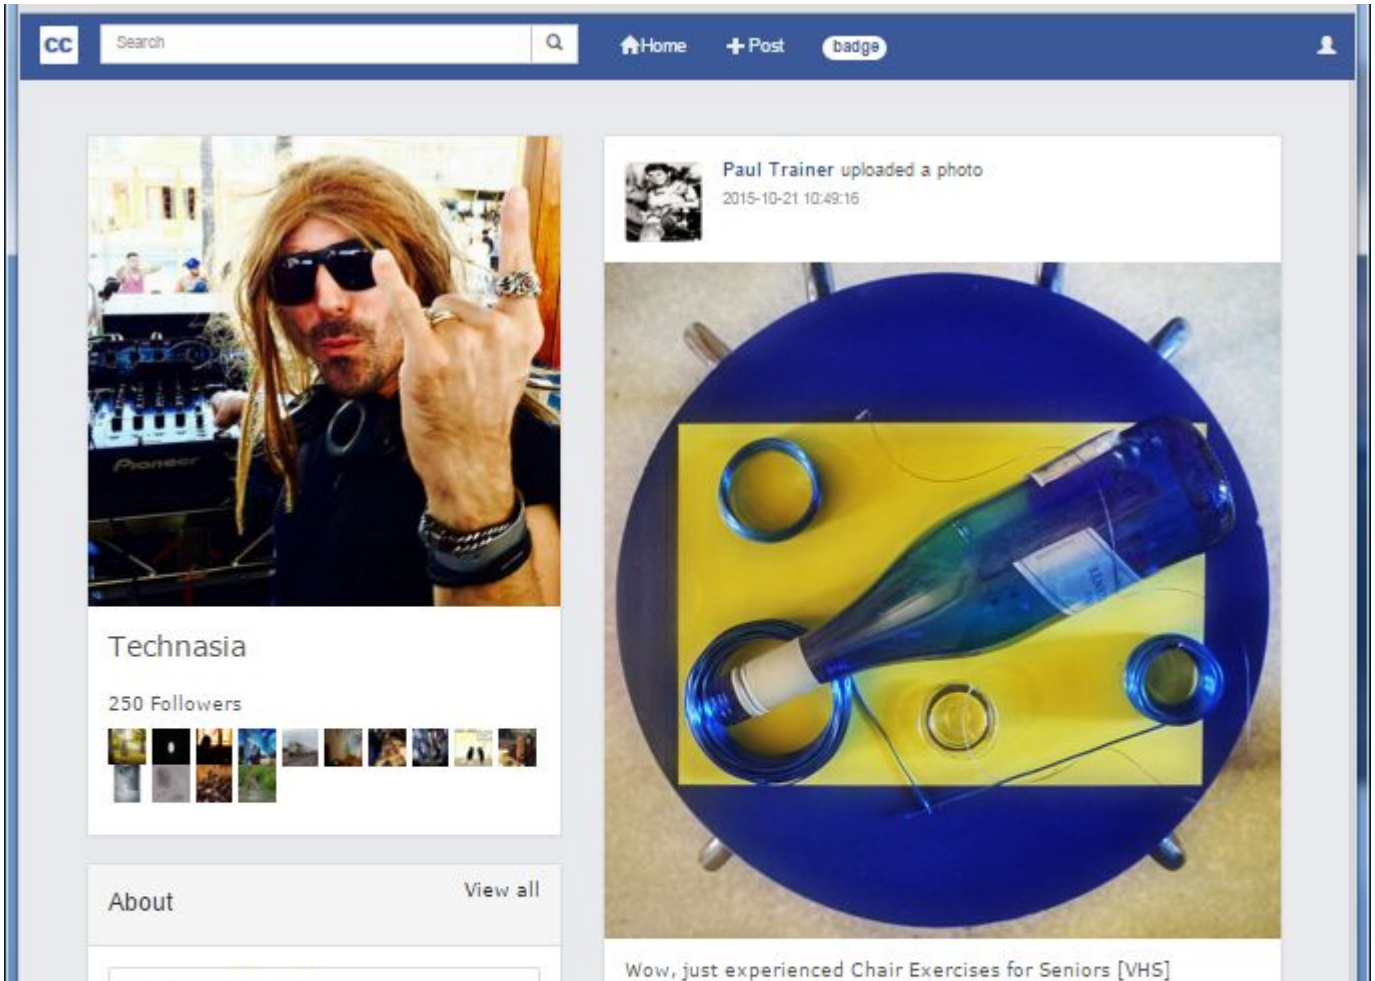

Project 3.4 : Introduction

- **Build CinemaChat, a social network about movies:**

- ✔ Task1: Implementing Basic Login with MySQL
- ✔ Task2: Storing Social Graph using HBase
- ✔ Task3: Build Homepage using DynamoDB
- ✔ Task4: Put Everything Together

Dataset Name	Data Store Used
Login Information	MySQL
User Profile	MySQL
Relation	HBase
Posts	DynamoDB
Profile and Post Images	S3

Project 3.4 : Introduction

- **Build CinemaChat, a social network about movies:**

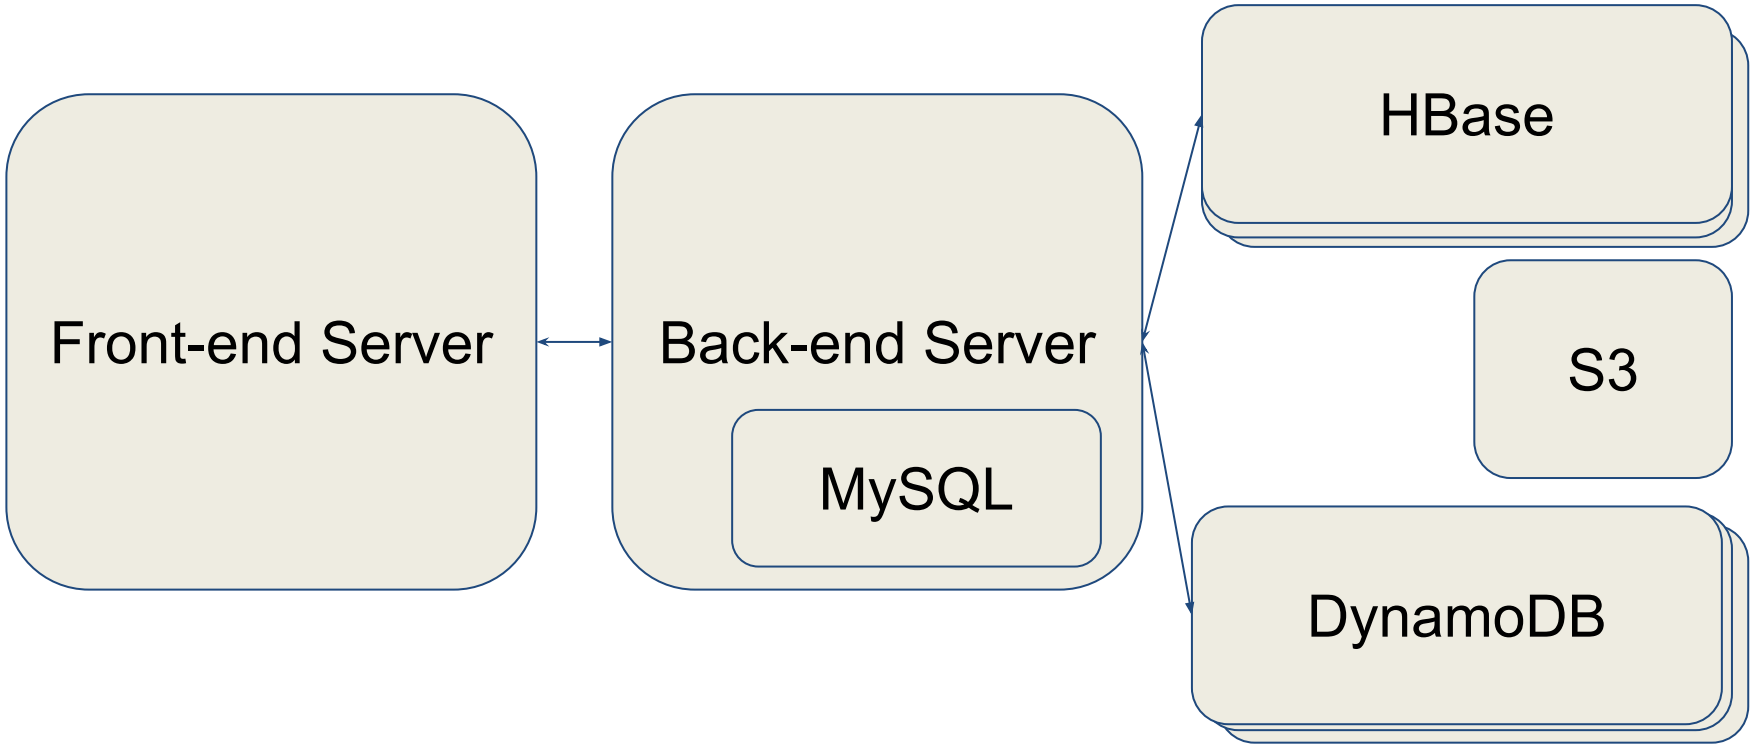

Upcoming Deadlines

- Quiz 9 : Unit 4 - Modules 16, 17

  Open : 11/06/2015 12:01 AM Pittsburgh

  Due : 11/06/2015 11:59 PM Pittsburgh

- Project 3.4 : Social Network with Heterogeneous DBs

  Due : 11/08/2015 11:59 PM Pittsburgh

- 15619Project : Phase 2

  Due : 11/11/2015 4:59 PM Pittsburgh

15619Project

Review

15619Project - Phase 1 : Review

- **Q1 -- Find the best web framework for you**
- **Q2 -- A large key-value store**
- **Feedback:**
  - ETL takes too long
  - Encoding issues
  - HBase optimization is hard

Design Space for the 15619Project

- Understanding the storage hierarchy
- Replication v/s sharding
- Schema design
- Intelligent Load Balancing
- Profiling a request

Upcoming Deadlines

- Quiz 9 : Unit 4 - Modules 16, 17

  Open : 11/06/2015 12:01 AM Pittsburgh

  Due : 11/06/2015 11:59 PM Pittsburgh

- Project 3.4 : Social Network with Heterogeneous DBs

  Due : 11/08/2015 11:59 PM Pittsburgh

- 15619Project : Phase 2

  Due : 11/11/2015 4:59 PM Pittsburgh

Project 3.4 Demo

Social Networking Timeline with Heterogeneous Backends

Architecture

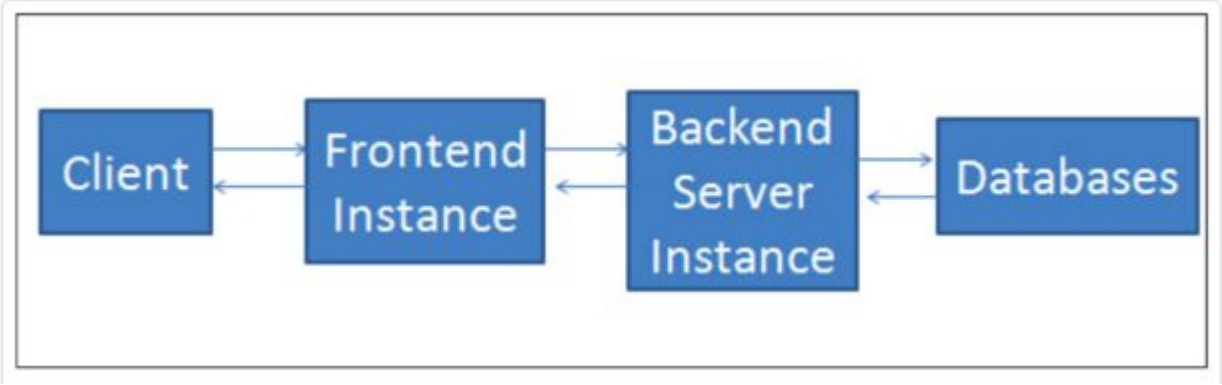

Data Set

1. User Profiles
    1. User Authentication System (such as a Single-Sign-On or SSO) - MySQL
    2. User Info / Profile - MySQL
    3. Action Log
    4. Social Graph of the User: follower, followee, family etc. - HBase
2. User Activity System - All user generated media - DynamoDB
3. Big Data Analytics System
    1. Search System
    2. Recommender System
    3. User Behaviour Analysis

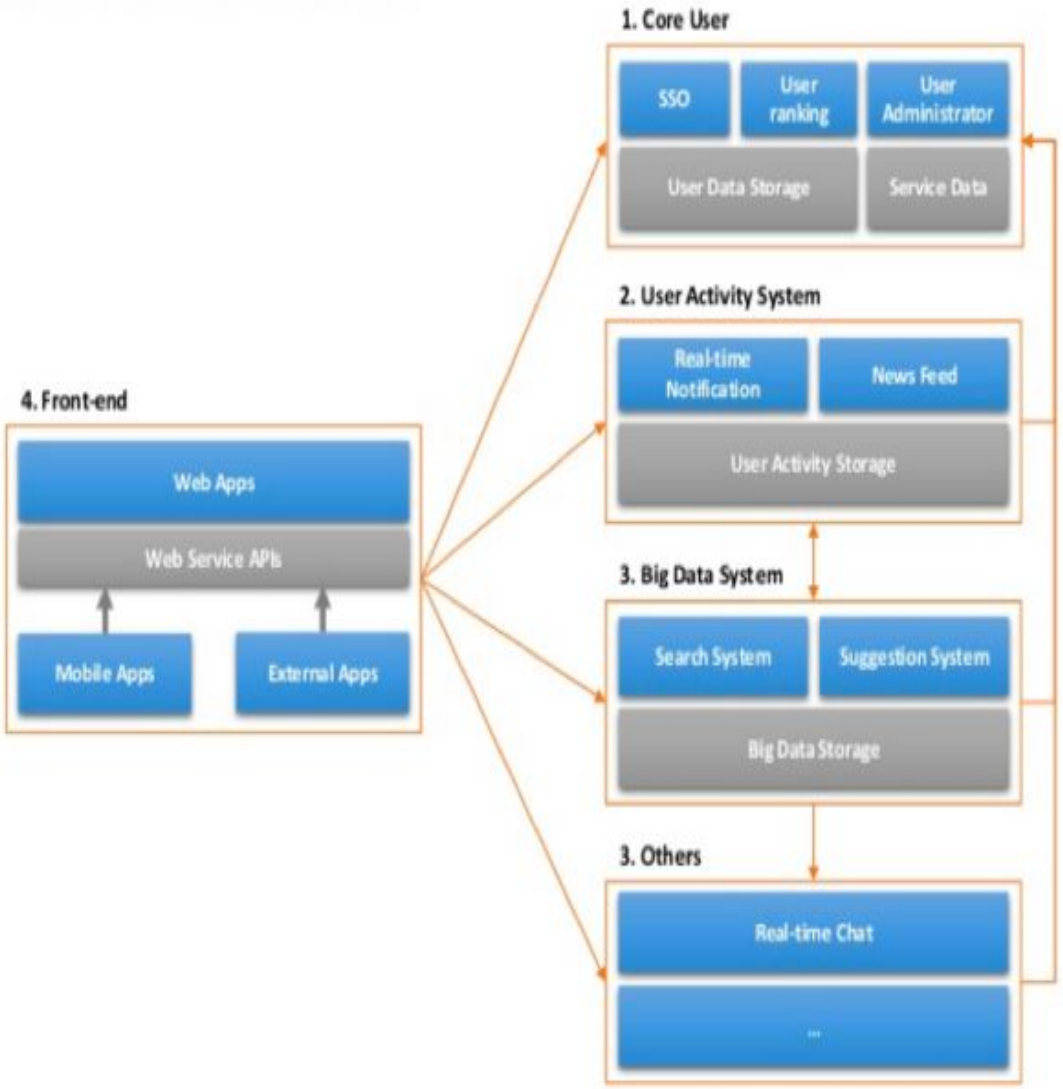

Heterogeneous Backends

Dataset Name	Data Store Used	Description	Location
Login Information	MySQL	[UserID, Password]	/home/ubuntu/users.csv
User Profile	MySQL	[UserID, Name, Profile Image URL]	/home/ubuntu/userinfo.csv
Relation	HBase	[Followee, Follower]	/home/ubuntu/links.csv
Posts	DynamoDB	JSON Format	/home/ubuntu/posts.json
Profile and Post Images	S3	Backend system only need to provide URLs	URLs are in userinfo.csv and posts.json

AWS Technologies Explored

- HBase on Elastic MapReduce
- AWS DynamoDB
- AWS Data Pipeline
- AWS S3

DynamoDB

- Eventual Consistency
- Data Model
  - Tables
    - DynamoDB only requires that a table has a primary key, but does not req

      define all of the attribute names and data types in advance.
  - Items
    - Individual items in a DynamoDB table can have any number of attributes,

      is a limit of 400 KB on the item size. An item size is the sum of lengths of

      names and values (binary and UTF-8 lengths).
  - Attributes
    - Each attribute in an item is a name-value pair. An attribute can be single-valued or multi-

      valued set. For example, a book item can have title and authors attributes

      has one title but can have many authors. The multi-valued attribute is a s

      values are not allowed.

DynamoDB

- Eventual Consistency
- Data Model
  - Primary Key
    - Simple Primary Key
    - The primary key is made of one attribute, a hash attribute. DynamoDB bu
      unordered hash index on this primary key attribute. Each item in the table
      identified by its hash key value.
    - Composite Primary Key
    - The primary key is made of two attributes. The first attribute is the hash a
      second one is the range attribute. DynamoDB builds an unordered hash i
      hash primary key attribute, and a sorted range index on the range primary
      Each item in the table is uniquely identified by the combination of its hash
      attribute values. It is possible for two items to have the same hash attribu
      those two items must have different range attribute values.

TWITTER DATA ANALYTICS: 15619 PROJECT

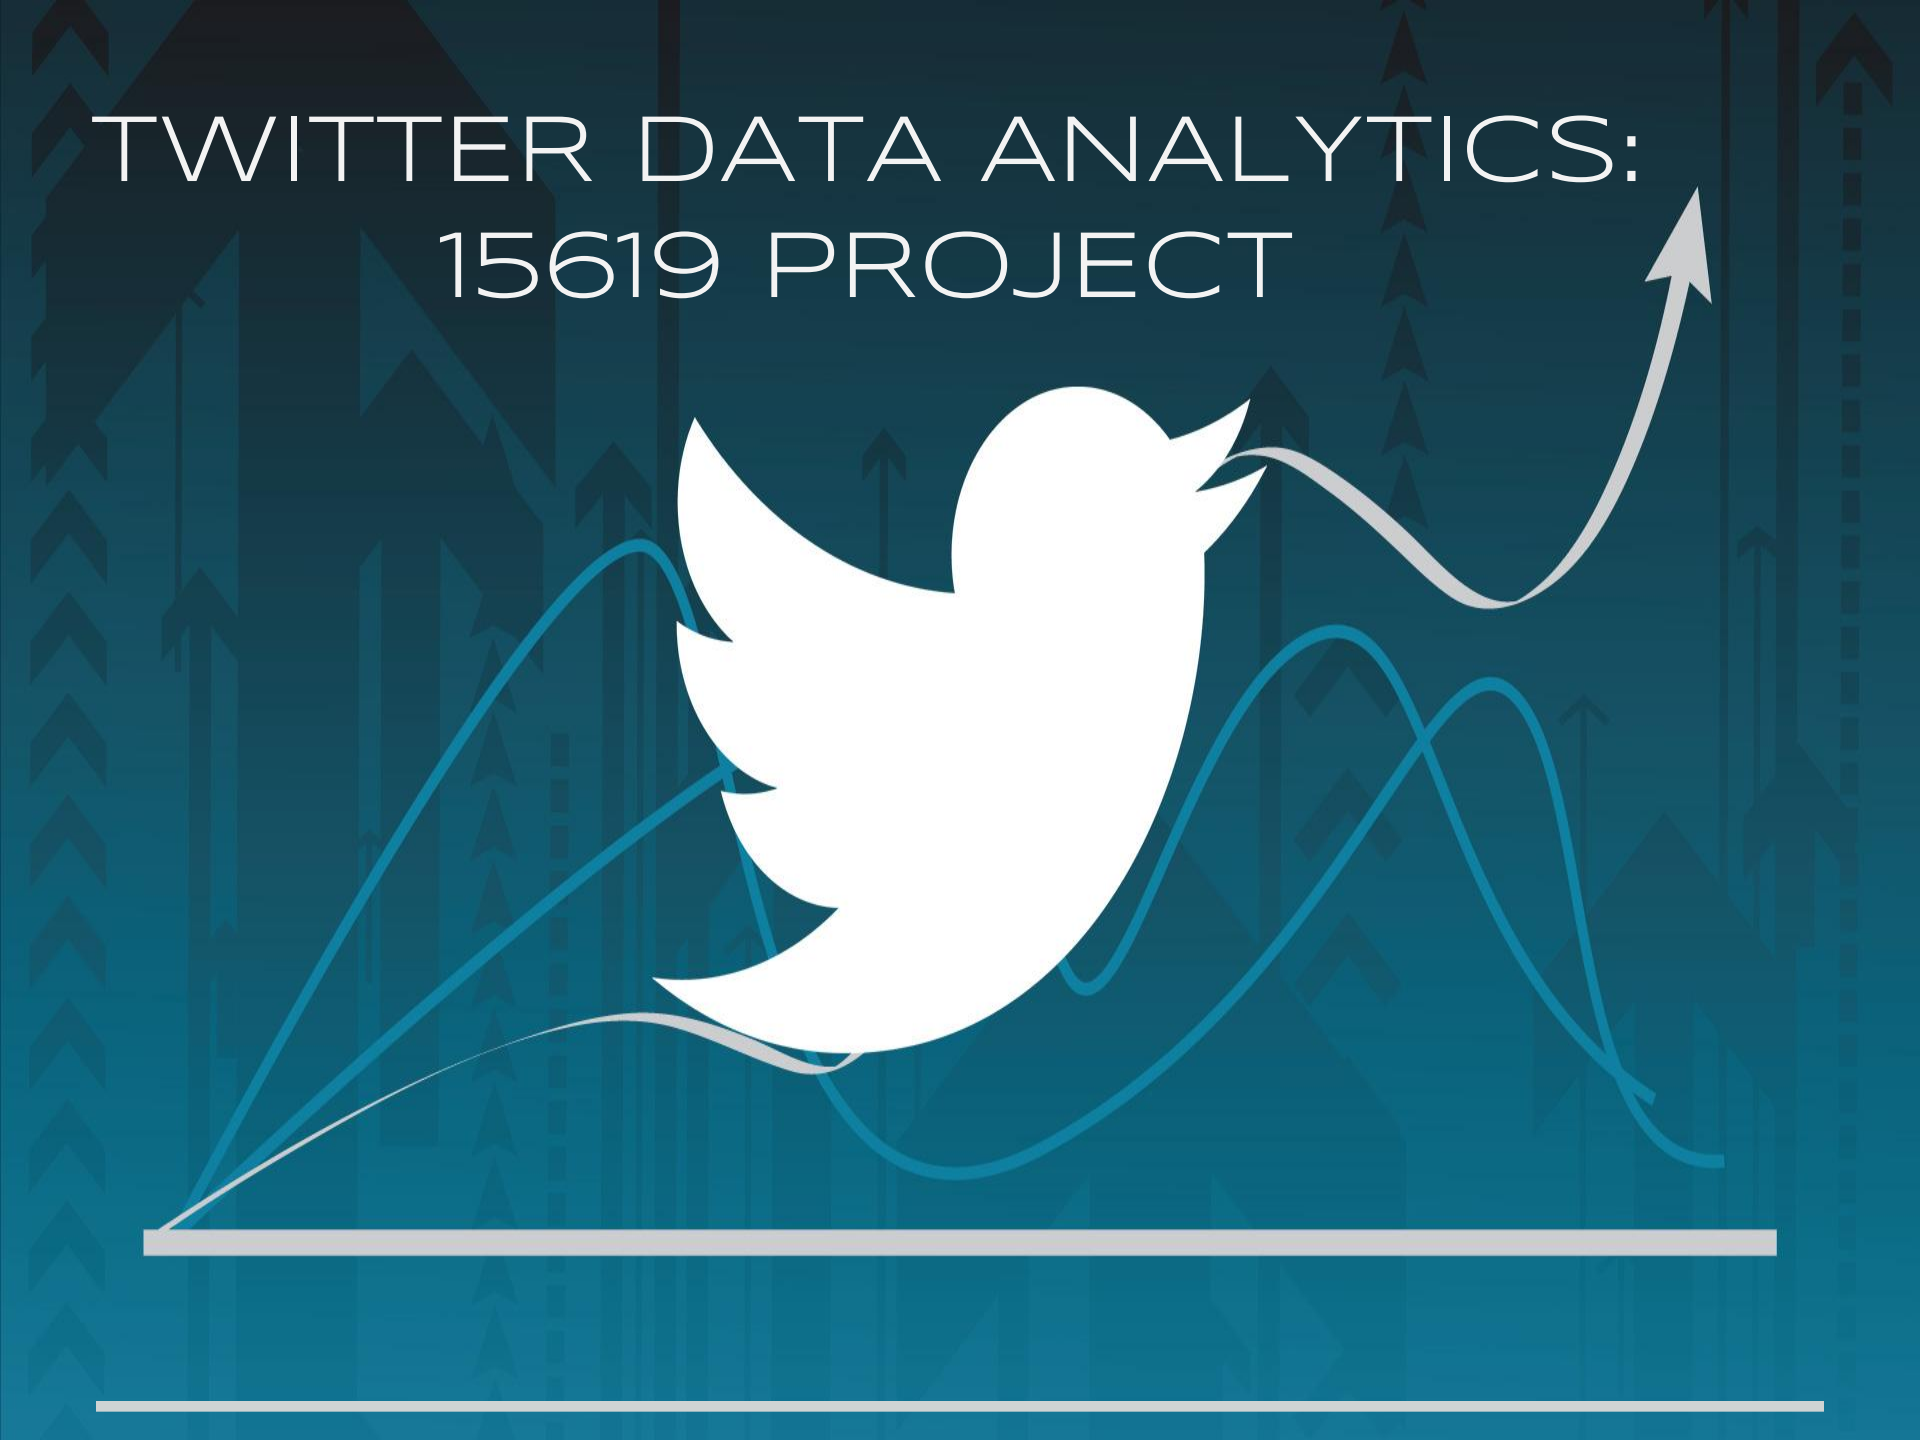

Phase 1 Leaderboard

Rank	Nickname	Time	Total	Q1 Throughput	Q1 Latency	Q1 Error	Q1 Correctness	Q1 Score
1	Yolo	10/21/2015 04:50 -0400	142.05	35511.7	2	0.00	100.00	142.05
2	cmucchackers	10/20/2015 20:43 -0400	141.72	35429.8	2	0.00	100.00	141.72
3	purrito	10/19/2015 06:10 -0400	137.96	34490.4	2	0.00	100.00	137.96
4	UHW	10/21/2015 21:20 -0400	137.45	34361.8	2	0.00	100.00	137.45
5	HUGE	10/18/2015 20:07 -0400	135.85	33962.6	2	0.00	100.00	135.85

Rank	Nickname	Time	Total	Q2 Throughput	Q2 Latency	Q2 Error	Q2 Correctness	Q2 Score
1	theImp	10/29/2015 00:24 -0400	457.85	30680.4	1	0.00	97.00	457.85
2	Directors	10/28/2015 18:43 -0400	240.34	15621.9	3	0.00	100.00	240.34
3	QingfenSpicyPot	10/28/2015 16:54 -0400	233.45	15174.2	2	0.00	100.00	233.45
4	AWSome	10/27/2015 09:15 -0400	210.57	13686.8	3	0.00	100.00	210.57
5	UHW	10/29/2015 02:18 -0400	187	12154.8	3	0.00	100.00	187

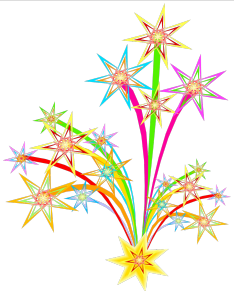

Well done !!!

Congratulations Yolo & theImp

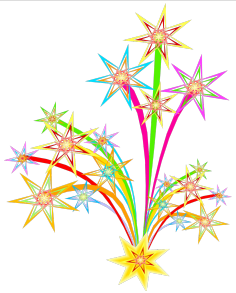

15619 Project Phase 2 Deadlines

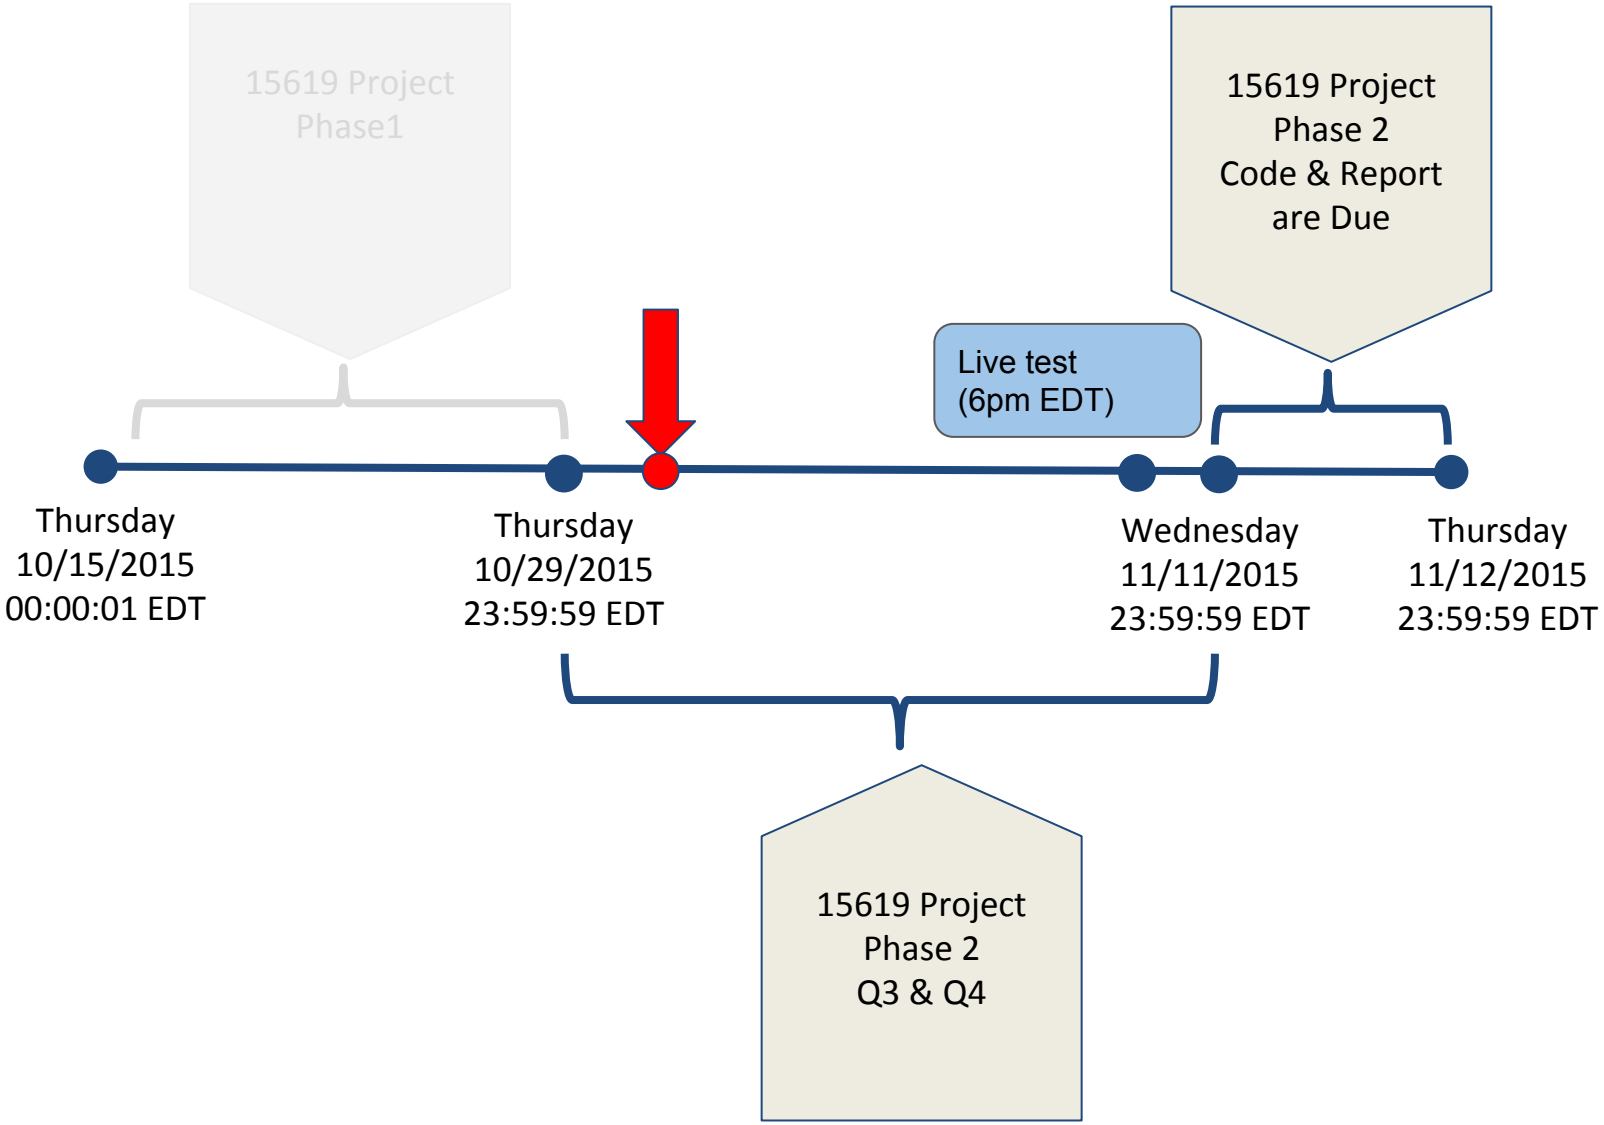

Phase 2

- Two more queries (Q3 and Q4)
  - More ETL
  - Multiple tables and queries

- **Live Test!!!**
  - **Both HBase and MySQL**
    - **Two DNS**
    - **Must get over 50% score on both!**
  - **Includes Mixed-Load**
  - **No more pre-caching of known requests**

Query 3: Spreading Happiness & Sadness

Q. How to compute impact score for each tweet?

score = sentitment_score * (1 + followers_count)

Q. What's "followers_count" and how do I find it?

Read https://dev.twitter.com/docs/platform-objects/tweets

Query 3: Spreading Happiness & Sadness

- Sample Query:

  GET /q3?

  start_date=2012-10-17&end_date=2014-11-30&userid=2334640502

  &n=10

  - n will be less or equal to 10

Query 3: Spreading Happiness & Sadness

- Sample Response:

Positive Tweets

2014-04-18,536,456980867093524480,@AustinMahone Austin please follow me i love you much! ♥ x44
2014-04-07,512,453011017078145024,@ArianaGrande Ariana please follow me i love you much! ♥ x72
2014-04-07,512,453011398722064384,@ArianaGrande Ariana please follow me i love you much! ♥ x86
2014-04-07,512,453012510233624576,@ArianaGrande Ariana please follow me i love you much! ♥ x114
2014-04-07,512,453020030607695873,@ArianaGrande Ariana please follow me i love you much! ♥ x257
2014-04-07,508,453007393199513600,@ArianaGrande Ariana please follow me i love you much! ♥ x8
2014-04-18,270,456984730043711489,@AustinMahone please follow me. I know that it is "impossible" that you see this tweet but if you see it, please!! x45
2014-04-22,270,458751123390615552,@ArianaGrande please follow me. I know that it is "impossible" that you see this tweet but if you see it, please!! x7

Negative Tweets

Query 4: Hashtag Analysis

- Use the hashtag entity
- Sample Query:

  GET /q4?hashtag=01net&n=4

- Sample Response:

  2014-05-16:2:20598769,617996172:#01net: Philips attaque Nintendo et veut faire bannir la Wii U des ventes aux Etats-Unis http://t.co/V55biLpfoJ
  2014-03-28:1:20598769:#01net: Apple pourrait lancer un MacBook Air Retina et un iPad 12'' au second semestre http://t.co/7IRy9QAN2u
  2014-03-29:1:20598769:#01net: TV 4K : faut-il craquer tout de suite ? Nos tests, les prix et ce qui arrive... http://t.co/jRpzOAjtFY
  2014-04-08:1:20598769:#01net: Le rapprochement entre Bouygues Telecom et Free se prépare en douceur http://t.co/d3vipnIYVV

Phase 2 Clarifications

- Censorship? q3: ✓, q4: X
- No required filtering
- Tie breaker?
  - q3:
    - Primary key - **absolute** value of score (descending)
    - Secondary key - tweet id resolve ties by tweet i (ascending)
  - q4:
    - Primary key - count, date with higher count co first
    - Secondary key - date, earliest date comes first

Tips for Phase 2

- Carefully design and check ETL process
- Watch your budget. $60 = phase + Live Test
- Preparing for the live test
  - You are required to submit two URLs, one for the MySQL live test and one for the HBase for live test
  - Budget limited to $1.25/hr for MySQL and HBase web service separately.
  - Caching known requests will not work
  - Need to have all Q1-Q4 running at the same time
  - Don’t expect testing in sequence.
    - Queries will be mixed.

MySQL

- Understand the difference between different storage engines
- Understand the difference between different types of indices
- Remember, replication may not be the best strategy if the tables are big
- Avoid costly operations in front-end
- Make use of connection pools, caching configuration in MySQL
- Consider using a MySQL Cluster

HBase

- Make use of WebUI:
  - Check which tables are getting requests the mosts
  - Use **split**, **move**, **merge** functions to evenly distribute loads
- Make sure your front end is not the bottleneck
- Carefully design schema
- Consider building your own HBase- EMR is expensive and abstracts away details

What's due soon?

- Phase 2 Deadline
  - **Submission by 16:59 ET (Pittsburgh) Wed 11/11**
    - **HBase Live from 6 PM to 9 PM ET**
    - **MySQL Live from 9 PM to midnight ET**
  - Fix Q1 and Q2 if your Phase 1 did not go well
  - New queries Q3 and Q4. Targets 5000 and 8000 rps
  - Get score higher than 50% on both MySQL and HBase
  - Heads up: Phase 2 counts for 30% of 15619Project grade
- Report at the end of Phase 2
  - Make sure you highlight failures and learning
  - If you didn't do well, explain why
  - If you did, explain how
  - Cannot begin to stress how critical this is!!!!