15-319 / 15-619
Cloud Computing

Course Overview and Introduction

August 30 and September 2, 2021

http://www.cs.cmu.edu/~msakr/15619-f21/

Outline

- **What is the course about?**
- What is an online course?
- Administrivia

So What is Cloud Computing?

Evolution of Computing

"Cloud Computing is the transformation of IT from a product to a service"

Evolution of Electricity

Innovation

New Disruptive Technology

Product

Buy and Maintain the Technology

Service

Electric Grid, pay for what you use

A Cloud is …

- Datacenter hardware and software that the vendors use to offer the computing resources and services

Cloud-enabling Data Centers

- Large warehouse scale data centers
- Growing at a rapid rate
- Next is an example from Microsoft Azure
  - Azure US-East2 (Boydton, VA)
  - Azure Expansion 1
  - Azure Expansion 2

Microsoft Azure


US-East2 (Boydton, VA)


US-East2 (Boydton, VA)


US-East2 Hub Expansion

The Cloud

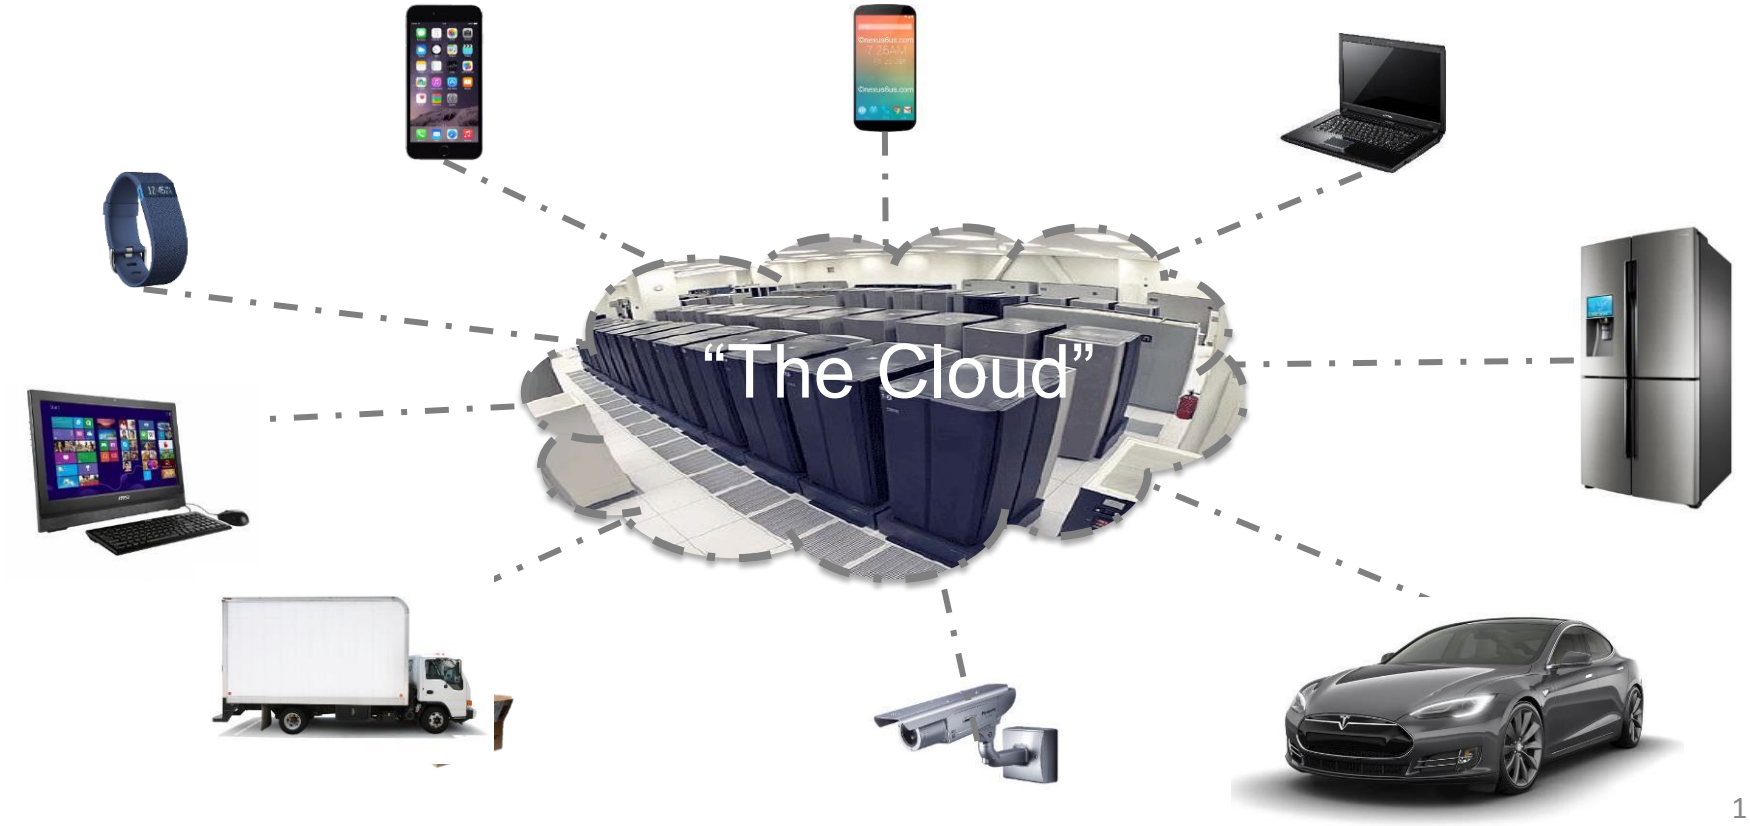

Enabled by Maturing Technologies

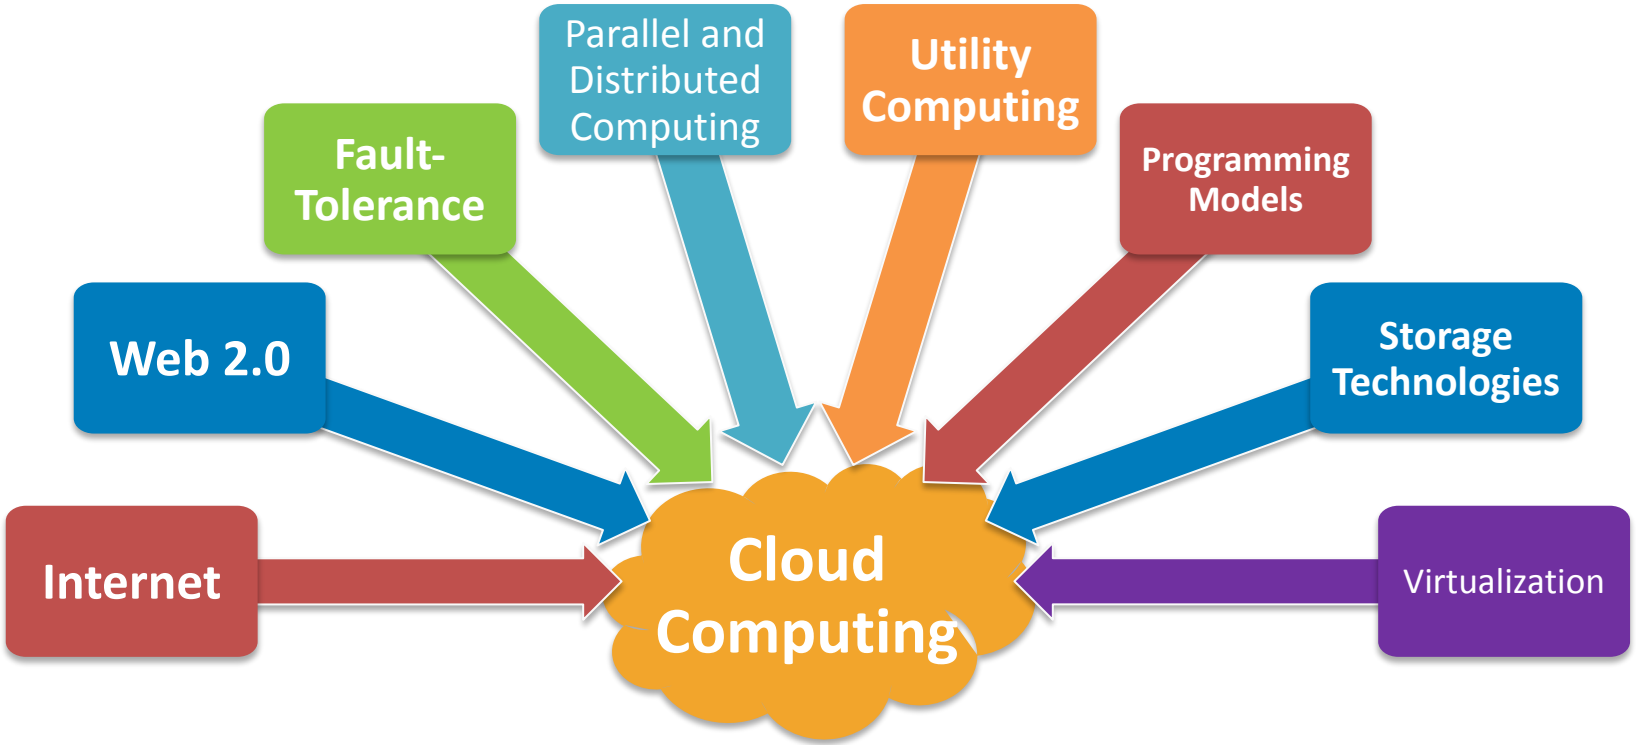

So… how would you transform information technology into a Service?

How to Transform IT to a Service?

- Connectivity
  - For moving data around
- Interactivity
  - Seamless interfaces
- Reliability
  - Failure will affect many
- Performance
  - Should not be slower
- Pay-as-you-Go
  - No upfront fee
- Ease of Programmability
  - Ease of development of complex services
- Manage Big Data
- Efficiency
  - Cost
  - Power
- Scalability & Elasticity
  - Flexible and rapid response to changing user needs

How to Transform IT to a Service?

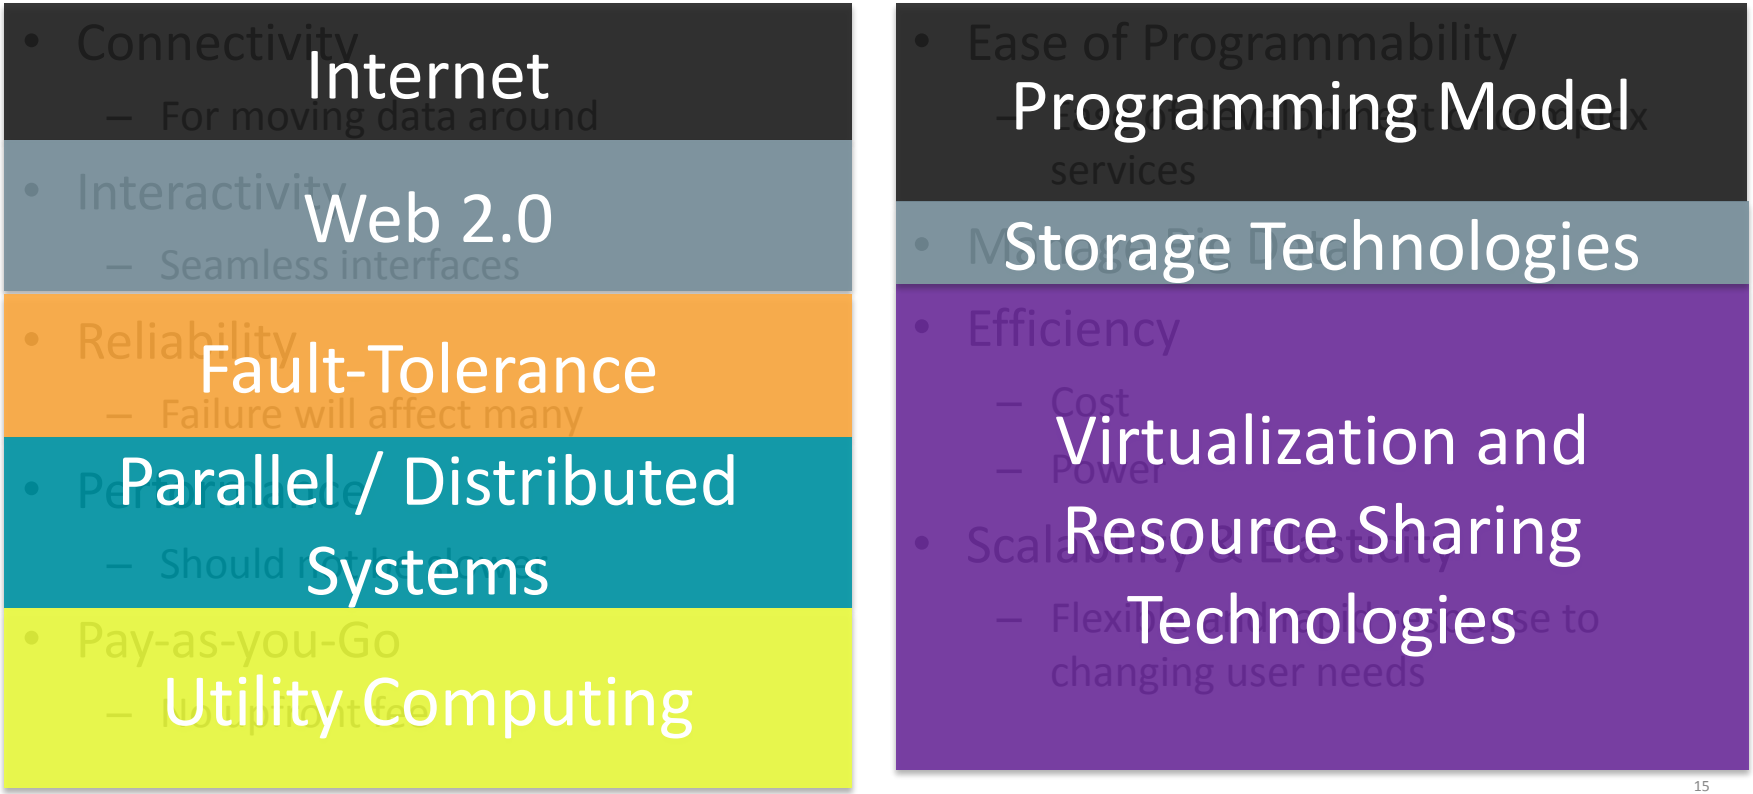

Cloud Building Blocks

Cloud services are available in various forms, corresponding to the layer of abstraction desired by the user

- Software as a Service (**SaaS**)
- Platform as a Service (**PaaS**)
- Infrastructure as a Service (**IaaS**)

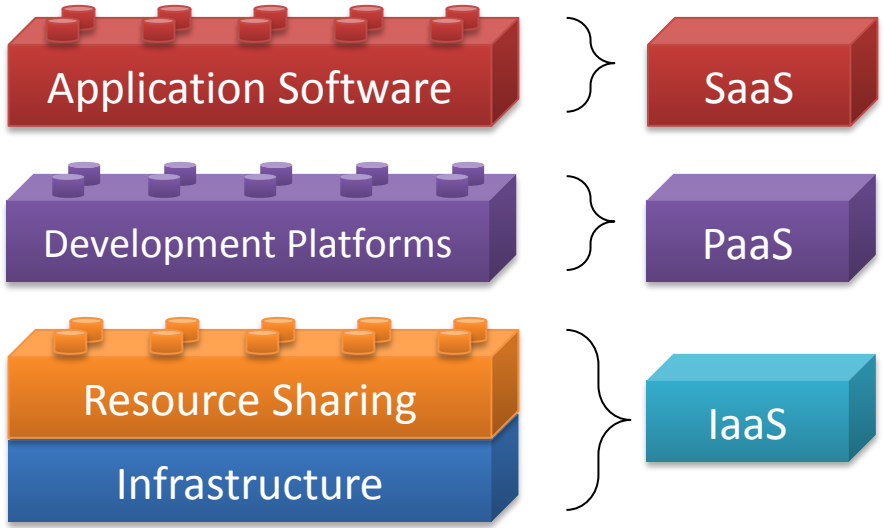

Cloud Computing Stack

- Applications
- Development Platforms
- Elasticity
  - APIs to enable automation,
    Alarms, protocols, triggers, etc…
- Sharing mechanisms
  - Virtualization, Containers, …
- Distributed systems
  - Programming models
  - Storage
- Data centers

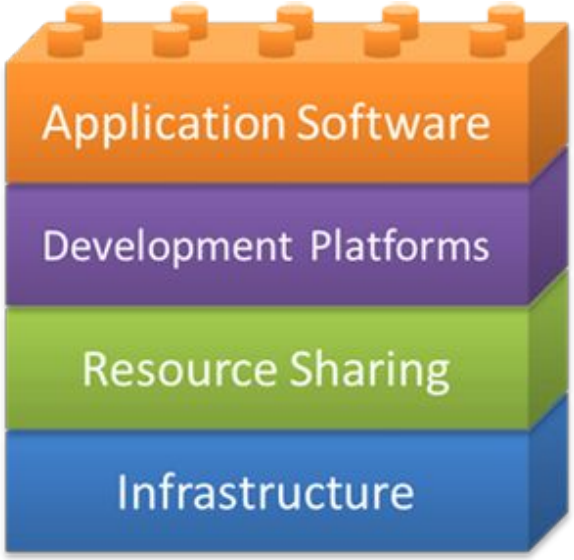

What is this course about?

- Applied aspects of cloud computing
  - Between systems and services

Conceptual content on OLI

1.0 Introduction to Cloud Computing
Service and deployment models, economics and use cases

2.0 Cloud Infrastructure
Components, design considerations and power

3.0 Resource Sharing
CPU, memory and I/O Virtualization, SDN, SDS

4.0 Cloud Storage
Distributed File Systems and Distributed Databases

5.0 Programming Models
MapReduce, Spark, GraphLab, Kafka/Samza

Projects on AWS, Azure, & GCP

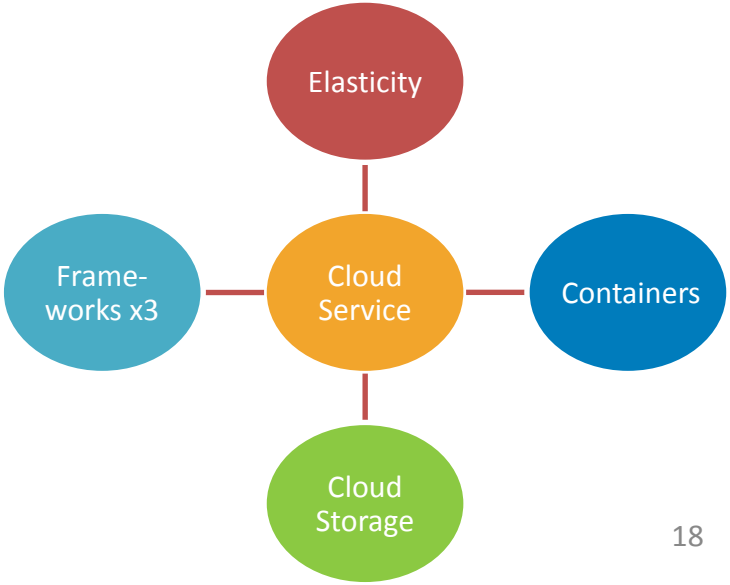

Course Goals

Students gain hands-on experience solving real world problems by completing projects in the areas of cloud **analytics**, **compute and elasticity**, **storage** and **frameworks**, which utilize existing public cloud tools and services. Students are exposed to real-world data scenarios, infrastructure and budgets in order to learn how to:

1. Design, architect, implement, test, deploy, monitor and maintain cloud-based applications;
2. Identify the appropriate tools and architectures to implement a cloud-based design;
3. Analyze the tradeoffs between different tools and cloud offerings to meet real-world constraints;
4. Evaluate performance characteristics of cloud-based services to implement optimizations;
5. [15-619 only] Collaborate with a team on an open-ended project to incrementally realize an optimized end-to-end cloud-based solution.

Conceptual Content on OLI

Unit #	Title	Modules and Description
1	Introduction	Definition and evolution of Cloud Computing Enabling Technologies Service and Deployment Models Popular Cloud Stacks and Use Cases Benefits, Risks, and Challenges of Cloud Computing Economic Models and SLAs Topics in Cloud Security
2	Cloud Infrastructures	Historical Perspective of Data Centers Datacenter Components: IT Equipment and Facilities Design Considerations: Requirements, Power, Efficiency, & Redundancy Power Calculations and PUE Challenges in Cloud Data Centers Cloud Management and Software Deployment Considerations
3	Virtualization	Virtualization (CPU, Memory, I/O) Case Study: Amazon EC2 Software Defined Networks (SDN) Software Defined Storage (SDS)
4	Cloud Storage	Introduction to Storage Systems Cloud Storage Concepts Distributed File Systems (HDFS, Ceph FS) Cloud Databases (HBase, MongoDB, Cassandra, DynamoDB) Cloud Object Storage (Amazon S3, OpenStack Swift, Ceph)
6	Programming Models	Distributed Programming for the Cloud Data-Parallel Analytics with Hadoop MapReduce (YARN) Iterative Data-Parallel Analytics with Apache Spark & GraphLab

Quiz 1, Friday Sep 10, 2021

Projects on AWS/Azure/GCP Clouds

0. AWS/Azure/GCP Account Setup & Data Analytics	• Benchmarking VMs, SSH, Authentication, Billing, Security Groups, Vertical Scaling
1. Scaling, Elasticity and Failure with VMs	• Auto Scaling, Load Balancing, Monitoring
2. Cloud Elasticity with Containers & K8s	• Docker Containers, Kubernetes
3. Cloud Storage	• Standalone MySQL, HBase, Neo4J, MongoDB, Azure Managed SQL DB
4, 5, & 6. Analytics Engines for the Cloud	• Spark, DataBricks, Kafka/Samza, Cloud ML Frameworks
Team Project: A Cloud Native Web Service	• No restrictions on tools or AWS services • Evaluated based on cost and performance

Project Learning Objectives

Compute & Elasticity	1) Design, implement, test, package, deploy and monitor cloud applications using Virtual Machines (VMs), and Containers cloud computing services.
Cloud Storage	2) Explore and experiment with different distributed cloud-storage abstractions and compare their features, capabilities and applicability. 3) Orchestrate, deploy and optimize a unified application that integrates heterogeneous SQL and NoSQL database systems.
Frameworks	4) Design, implement, test and debug applications using interactive, batch and stream processing frameworks and compare their suitability to different problem domains. 5) Illustrate and explain the execution workflow, overhead, fault-tolerance and logical flow of interactive, batch and stream processing frameworks. 6) Train and deploy a machine learning model using a cloud-based framework. 7) Analyze and identify potential sources of bottlenecks in programming frameworks to optimize their performance.
Team Project [15-619 only]	8) Design, build and deploy a performant, reliable, scalable and fault-tolerant cloud native microservice based web service on the cloud within a specified budget. 9) Perform extract, transform and load (ETL) on a large data set. 10) Design schema as well as configure and optimize cloud-based databases to deal with scale and improve the throughput of a web service. 11) Explore methods to identify the potential bottlenecks in a cloud native web service and implement methods to improve system performance.
Overall	12) Practice gathering, cleaning and preparing data for analysis on the cloud. 13) Practice Test-driven Development (TDD) in the software development process. 14) Orchestrate and automate the process of managing and provisioning cloud resources through machine-readable definition files. 15) Make informed decisions about choosing an appropriate cloud tool that will satisfy a set of specified requirements.

P0: Big Data Analytics (ungraded)

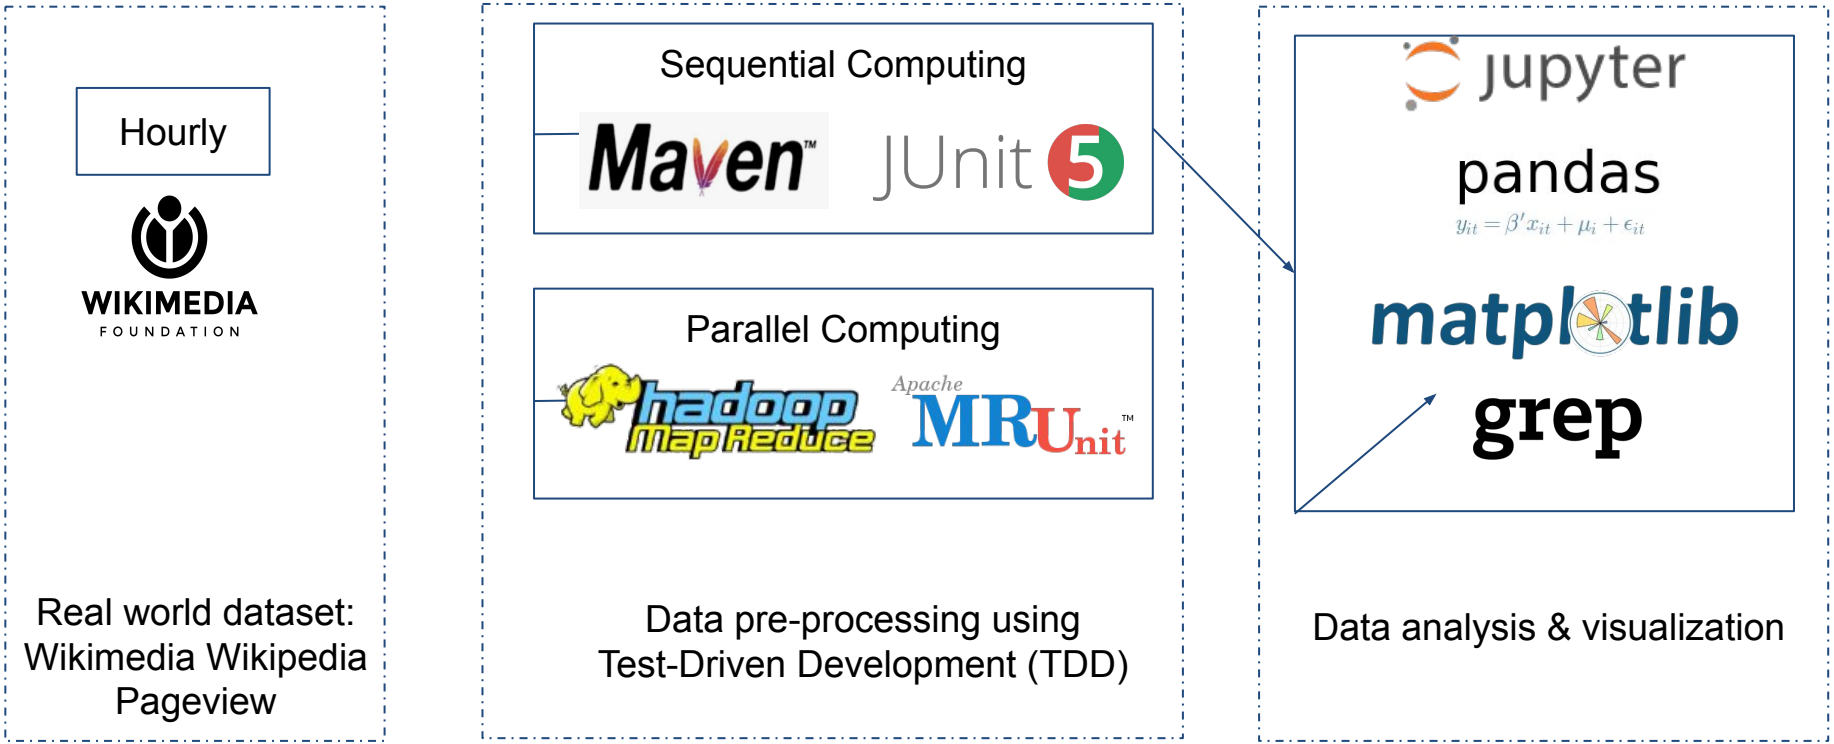

P1: VM Scaling, Elasticity & Failure

- Closed VMs
  - Load Generators
    - Dynamic load
  - Data Center Instances
    - Performance & failure
- Scale out & scale in DCIs to achieve desired RPS within budget
  - Auto-scale groups, elastic load balancers, monitoring, etc.

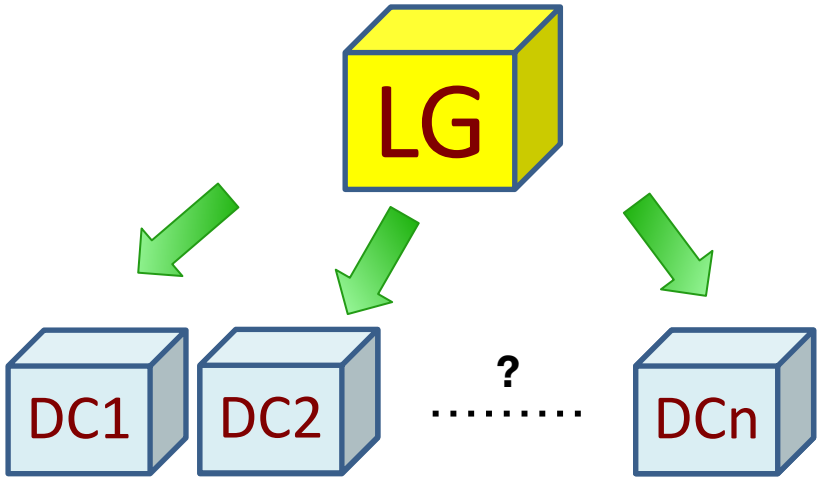

P2: Containers and Kubernetes

- Containerization
- Cluster Management
- Multi-Cloud Deployment

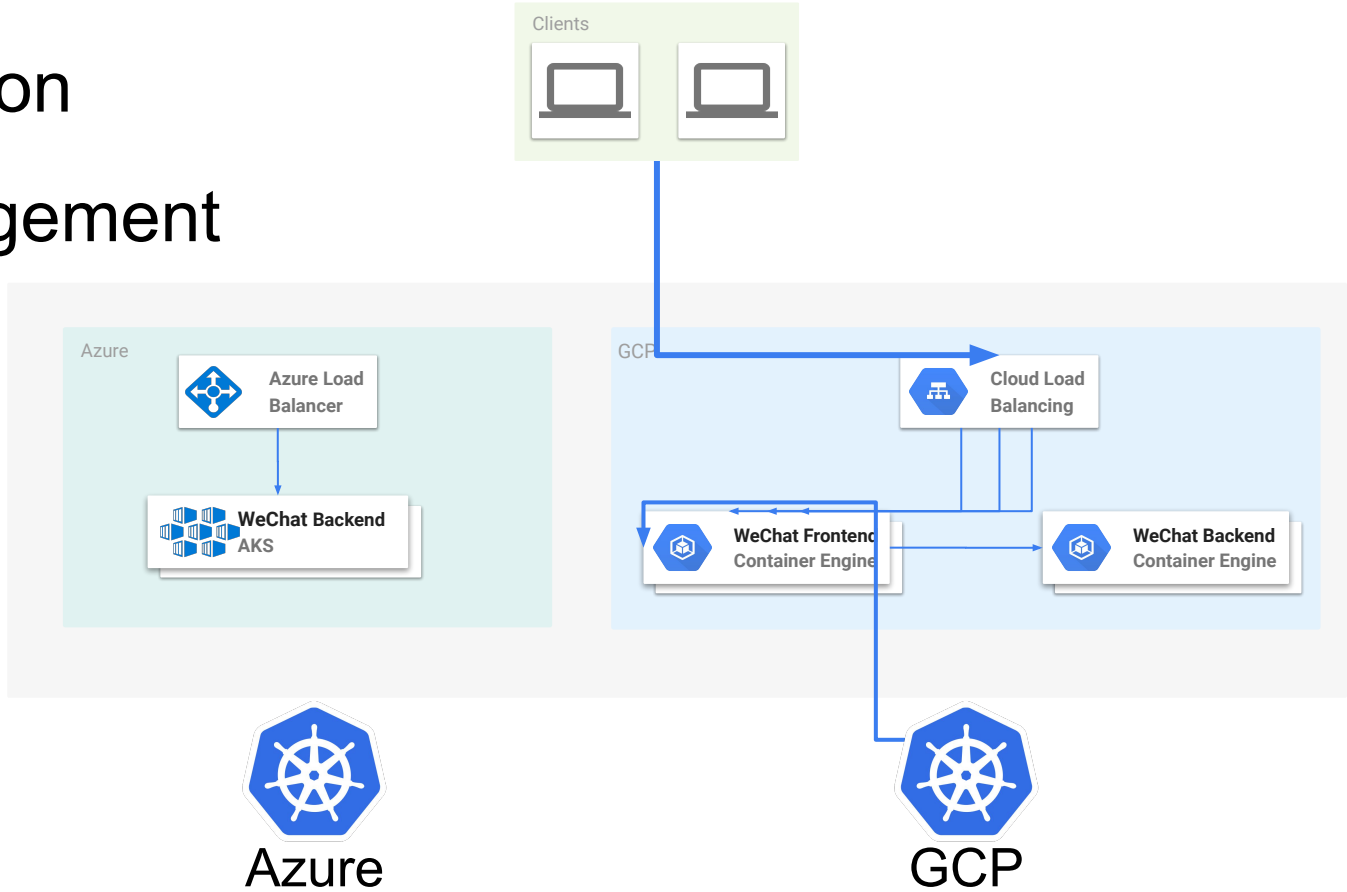

P5: Uber-like Application for NYC

- Stream Processing with Kafka/Samza
  - Stream 1: Car GPS coordinates
  - Stream 2: Customers
- Task:
  - Match customers with cars to minimize travel time & other constraints

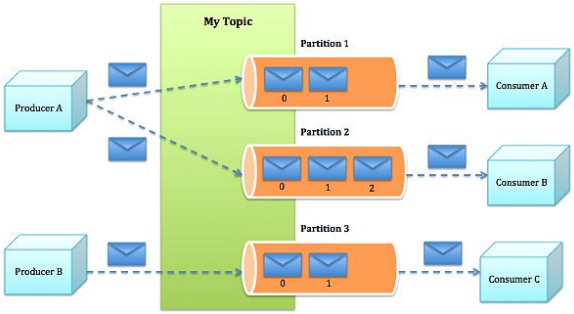

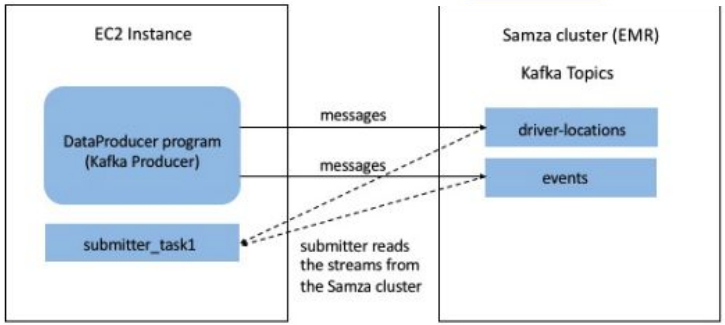

Team Project: Web Service

- Team-based
- 1.2 TB of raw data
- Specified queries
- Constraints
  - Correctness
  - Throughput
  - Budget
  - Time

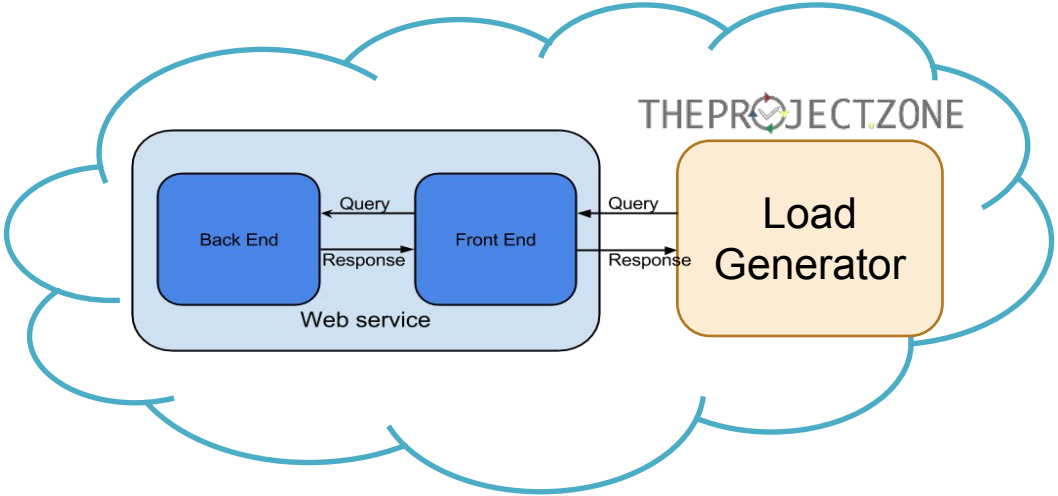

Phases	Duration	Microservices
Phase 1	3 weeks	M1 & M2
Phase 2	2 weeks	M1, M2, & M3
Phase 2 Live Test	6 hours	M1, M2, & M3
Phase 3	2.5 weeks	M1, M2, & M3 (Managed services)
Phase 3 Live Test	6 hours	M1, M2, & M3 (Managed services)

Outline

- What is the course about?
- **What is an online course?**
- Administrivia

Carnegie Mellon Global Course

15-319 - 12 units
15-619 - 15 units

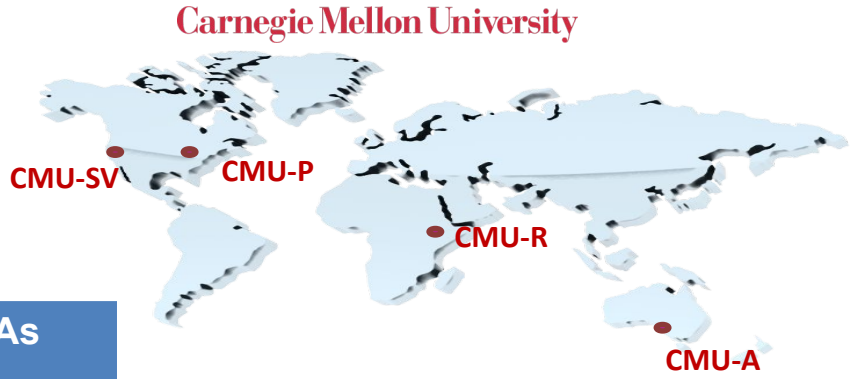

Location	Sections	Students	TAs
CMU Pittsburgh	A & B	87	8
CMU Silicon Valley	C	18	2
CMU Rwanda	D	8	1
CMU Adelaide	E	6	0

Please move to the section for your campus ASAP

Online Course Engagement Model

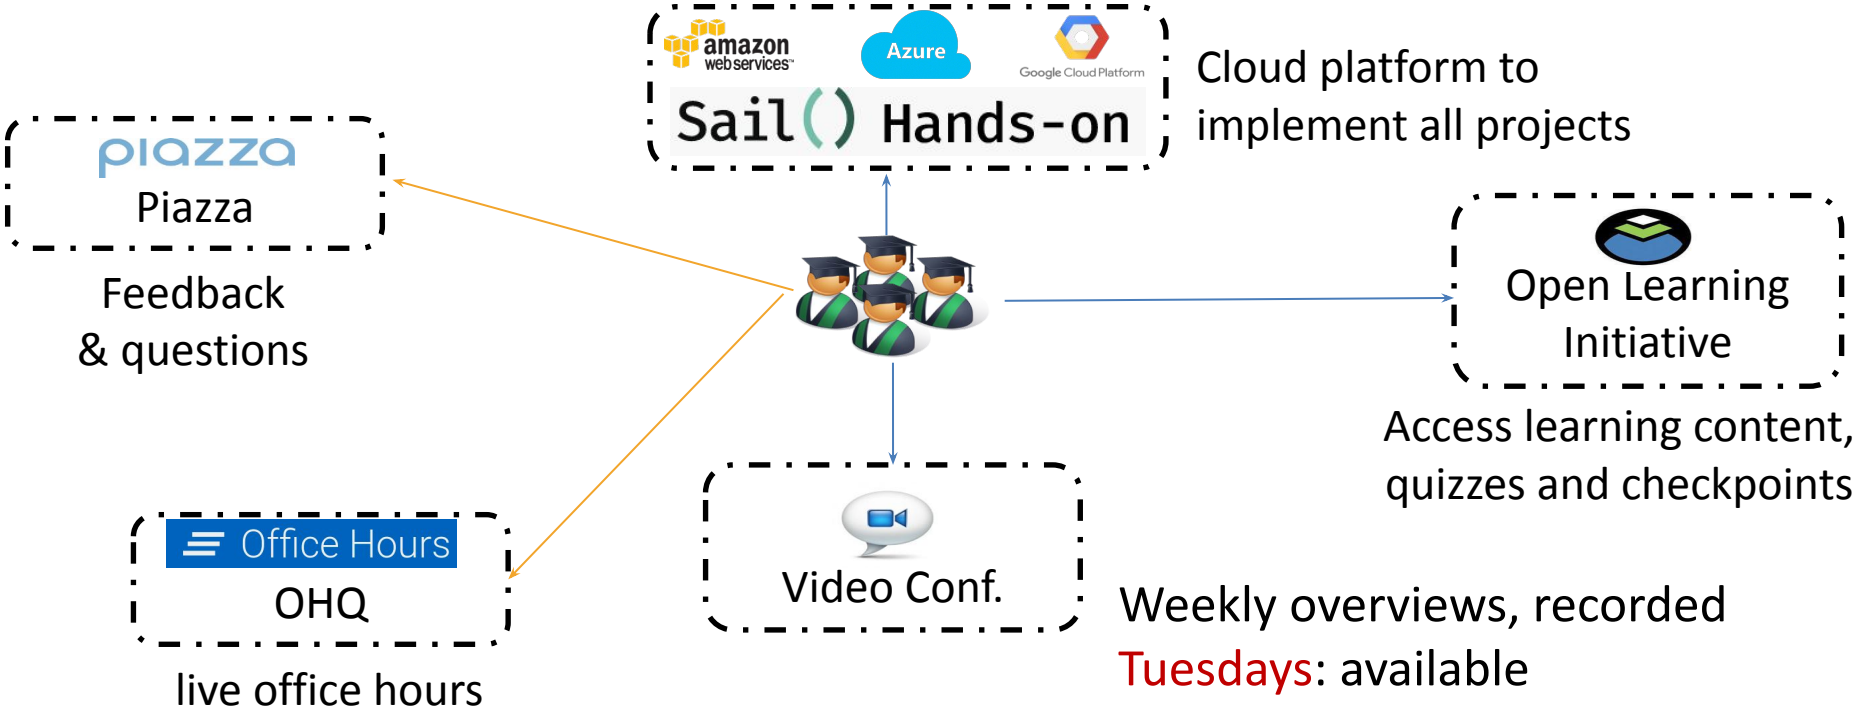

Canvas

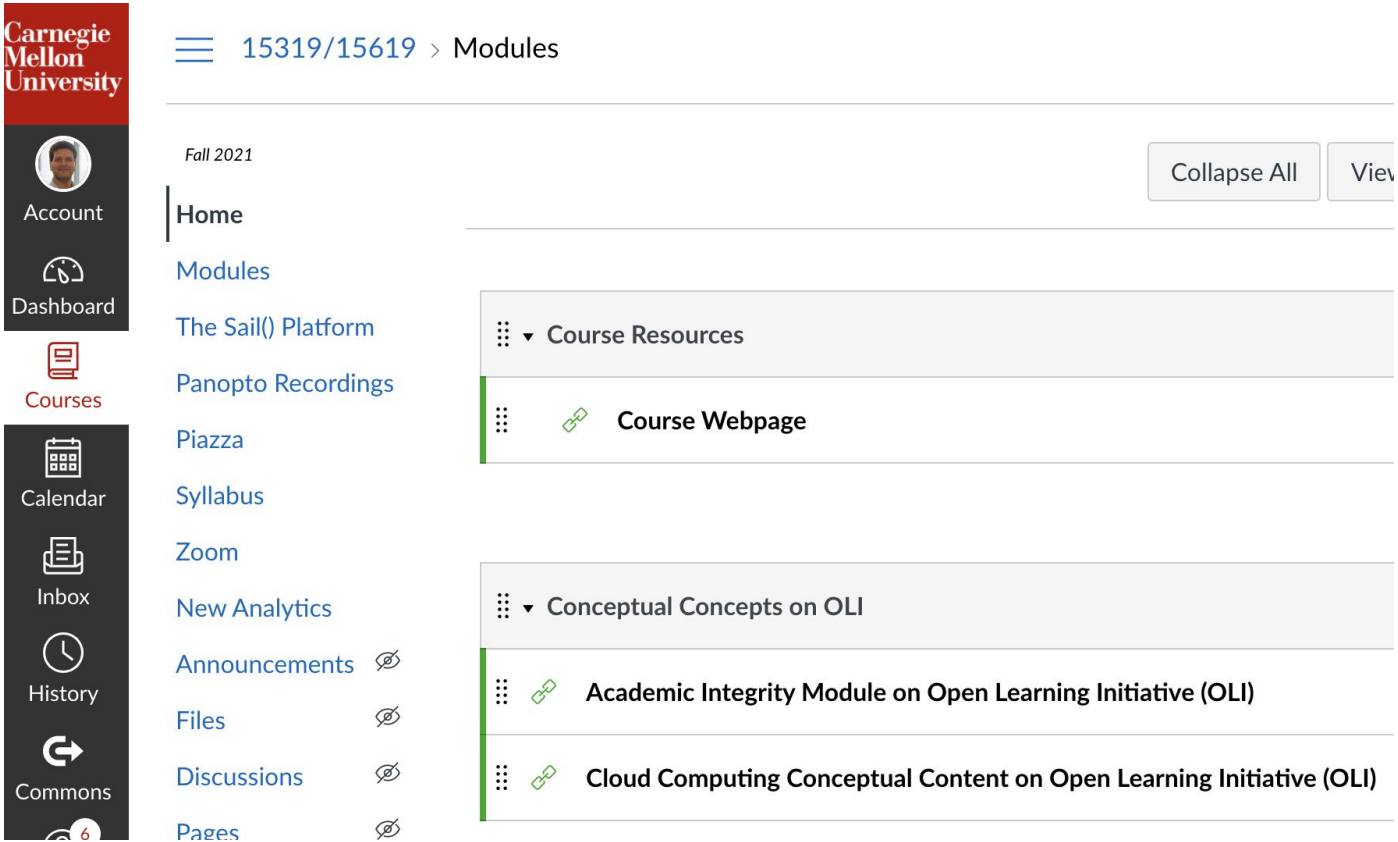
Carnegie Mellon University
Account
Dashboard
Courses
Calendar
Inbox
History
Commons
15319/15619 › Modules
Fall 2021
Home
Modules
The Sail() Platform
Panopto Recordings
Piazza
Syllabus
Zoom
New Analytics
Announcements
Files
Discussions
Pages
Collapse All
Course Resources
Course Webpage
Conceptual Concepts on OLI
Academic Integrity Module on Open Learning Initiative (OLI)
Cloud Computing Conceptual Content on Open Learning Initiative (OLI)

Online Course Content - OLI

Conceptual content is on the Open Learning Initiative:

- Students are automatically registered
- Access to OLI is through Canvas
  - canvas.andrew.cmu.edu
- Provide feedback on OLI
  - Bottom of each page
  - End of each module
- Do not copy or share content

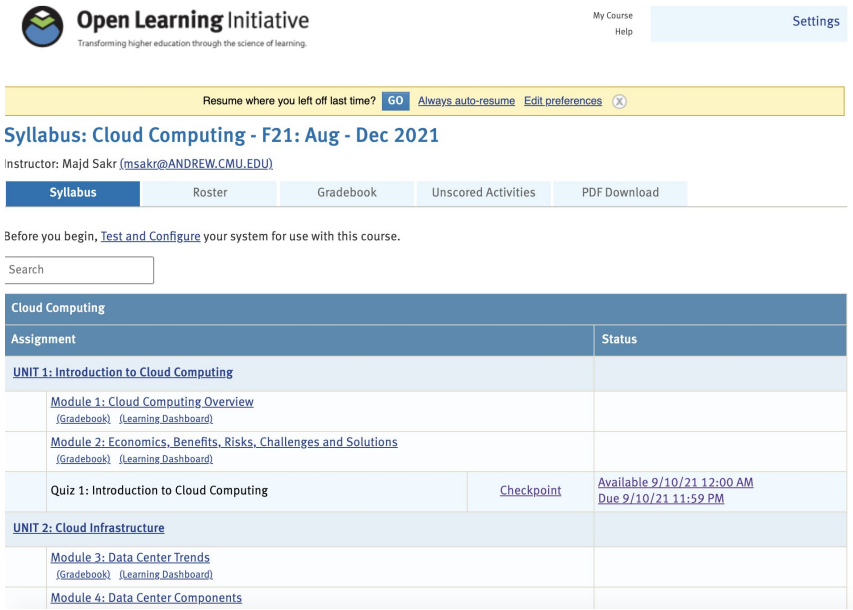

The Sail() Platform

Course projects are on https://projects.sailplatform.org:

- Learn through repetitive attempts and feedback
- Students are automatically registered
- Access through browser
  - Not mobile friendly yet
- Work in progress
  - We will encounter bugs
  - Provide feedback on Piazza
  - Please be patient

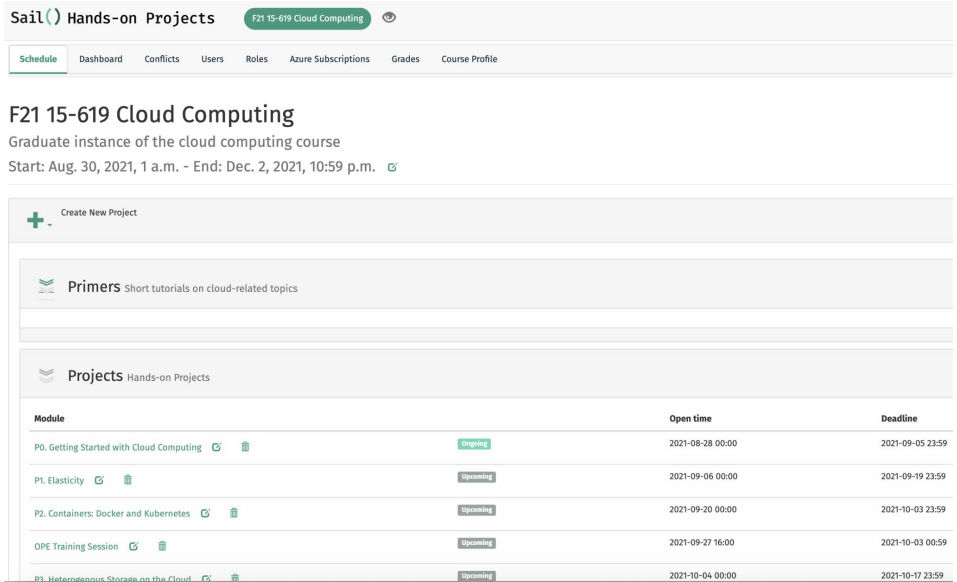

Office Hours

- My Office hours on Zoom (Tue 3pm ET)
- TA Office hours on OHQ+Zoom
  - See Piazza
  - Use OHQ 15-619

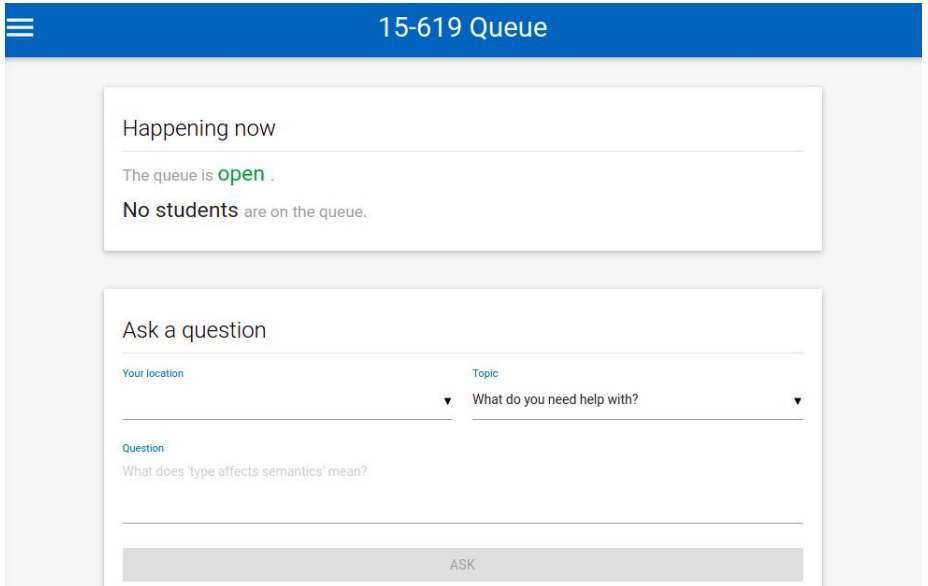

Syllabus

- Updated on webpage
- Provides details on:
  - Course Objectives
  - Learning Outcomes
  - Policies
  - Grading
  - Tentative Schedule

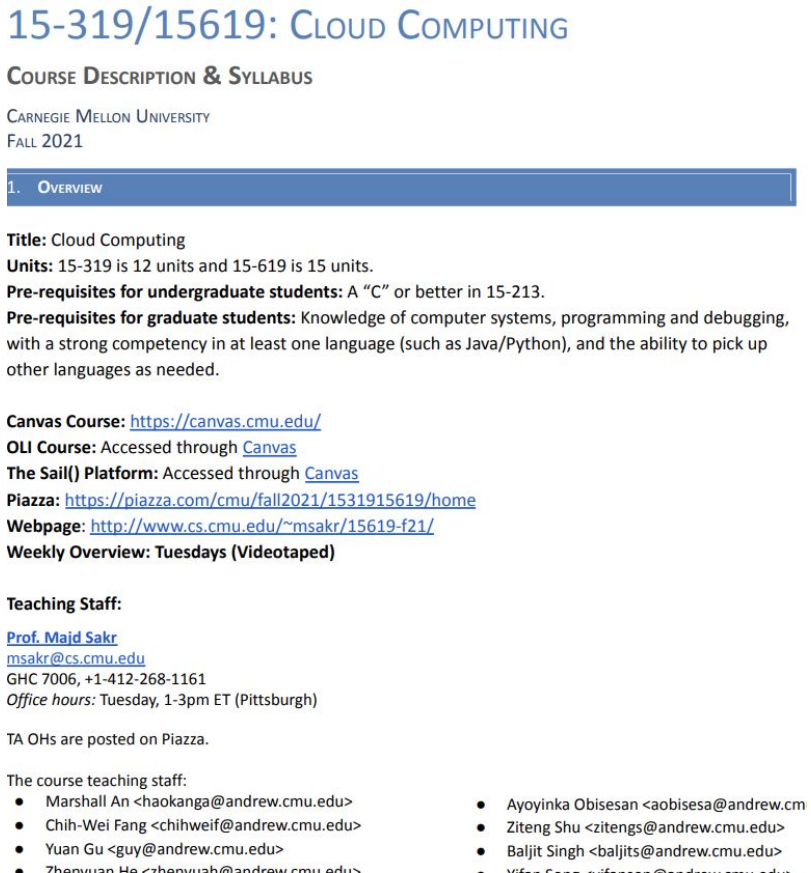

15-319/15619: Cloud Computing

Course Description & Syllabus

Carnegie Mellon University
Fall 2021

1. Overview

Title: Cloud Computing
Units: 15-319 is 12 units and 15-619 is 15 units.
Pre-requisites for undergraduate students: A "C" or better in 15-213.
Pre-requisites for graduate students: Knowledge of computer systems, programming and debugging, with a strong competency in at least one language (such as Java/Python), and the ability to pick up other languages as needed.

Canvas Course: https://canvas.cmu.edu/
OLI Course: Accessed through Canvas
The Sail() Platform: Accessed through Canvas
Piazza: https://piazza.com/cmu/fall2021/1531915619/home
Webpage: http://www.cs.cmu.edu/~msakr/15619-f21/
Weekly Overview: Tuesdays (Videotaped)

Teaching Staff:

Prof. Majd Sakr
msakr@cs.cmu.edu
GHC 7006, +1-412-268-1161
Office hours: Tuesday, 1-3pm ET (Pittsburgh)

TA OHs are posted on Piazza.

The course teaching staff:

- Marshall An <haokanga@andrew.cmu.edu>
- Chih-Wei Fang <chihweif@andrew.cmu.edu>
- Yuan Gu <guy@andrew.cmu.edu>
- Ayoyinka Obisesan <aobisesa@andrew.cm
- Ziteng Shu <zitengs@andrew.cmu.edu>
- Baljit Singh <baljits@andrew.cmu.edu>

Tentative Schedule

- Schedules:
  - Quizzes on OLI
  - Projects on The Sail() Platform
- No extensions

Week	Monday	OLI Content	Individual Projects	Team Project	Quizzes
1	8/30/2021	Unit 1, Module 1, 2	Primers/P0		Q0 (Ac. Integ.)
2	9/6/2021	Unit 1, Module 1, 2	P1 Elasticity		Q1 (Sep 10)
3	9/13/2021	Unit 2, Module 3, 4			Q2 (Sep 17)
4	9/20/2021	Unit 2, Module 5, 6	P2 Containers/K8s		Q3 (Sep 24)
5	9/27/2021	Unit 3, Module 7, 8, 9		Project Out (Sep 27)	Q4 (Oct 01)
6	10/4/2021	Unit 3, Module 10, 11, 12	P3 Storage		Q5 (Oct 08)
7	10/11/2021	Unit 3, Module 13		Phase 1 Due (Oct 17) Phase 2 Out (Oct 18)	Q6 (Oct 15)
8	10/18/2021	Unit 4, Module 14	P4 Batch Proc.		Q7 (Oct 22)
9	10/25/2021	Unit 4, Module 15, 16, 17		Phase 2 Due (Oct 31) Phase 3 Out (Nov 1)	Q8 (Oct 19)
10	11/1/2021	Unit 4, Module 18	P5 Streaming Proc.		Q9 (Nov 05)
11	11/8/2021	Unit 5, Module 19, 20		Phase 3 Due (Nov 14)	Q10 (Nov 12)
12	11/15/2021	Unit 5, Module 21, 22	P6 AI/ML on Cloud		Q11 (Nov 19)
13	11/22/2021	Thanksgiving			
14	11/29/2021				

Grading

Course Elements	#	Weight
Projects	6 + 1	80%
OLI Unit Checkpoint Quizzes	11	20%

- Projects weights
  - 15-319
    - 80 %, 5/6 individual project modules
  - 15-619
    - 60%, 5/6 individual project modules
    - 20%, 1 team project, three phases
- Weekly quizzes (11 in total)
  - 10 out of 11, 2% equal weight

Audit & Pass/Fail option is not available for this course

Outline

- What is the course about?
- What is an online course?
- **Administrivia**

Target Audience

- Technical Majors
- Undergraduate Juniors / Seniors
  - Pre-requisites:
    - 15213 – Introduction to Computer Systems
- Graduate Students
  - Experience:
    - Unix, scripting, python, & java

Course Administration

- Students are automatically registered on OLI through canvas.cmu.edu
- A *single* Piazza course page is created
  - We manually register students to Piazza
- Schedule of units and quizzes is on OLI
  - Content weekly quizzes are due on Fridays
- Schedule of weekly projects is on The Sail() Platform
  - Weekly project modules are due on Sundays

Public Cloud Infrastructure

- Paid Cloud Service
  - billed by the hour/minute
- Start a resource only when you need it
- To explore, use inexpensive instances
- **Terminate all other resources as soon as you are done with them**
- Students will be penalized for over usage
  - We have a fixed budget, do not abuse the resources!
  - Intentional or unintentional abuse ➔ grade penalties
  - Resources need to be tagged, otherwise ➔ penalties

This Week

- Check that you have been enrolled on Canvas and Piazza
- Academic Integrity Module on OLI
  - **Monday, September 6, 2021**
- Become familiar with conceptual content on OLI
  - Start reading Unit 1, Module 1 & Module 2
  - Quiz 1: Unit 1, Module 1 & 2, **Friday, September 10, 2021**
- Create an account on AWS, Azure and GCP (ASAP)
  - Submit your AWS account info using the link provided in the primers on The Sail() Platform
- Projects on The Sail() Platform
  - Primer and P0, due **Sunday, September 5, 2021**

Diverse Technical Preparation

- Students come from diverse backgrounds and technical preparation
  - We offer primers to get you started.
  - If your programming skills are rusty, take the first two weeks to improve.
  - If you don’t think you have the skills required,
    allocate more time each week for the projects.
  - The first couple of weeks are less demanding, take advantage of them.

Perfect Conditions Do Not Exist

- Don’t ask to be trained under perfect conditions
  - We will not provide a sanitized sandbox for you to learn
- You will encounter
  - Badly formed data, inaccurate documentation, intermittent services, insufficient information, etc.
  - Learn how to deal with all these issues
  - Very valuable experience for your career

Time Management is Key

- We are as good as what skills, and hard work we bring to the table.
- Don’t ask for special circumstance due to drama.
  - Find out the source of the drama and make adjustments.
- …

Academic Integrity

It is the responsibility of each student to produce her/his own original academic work.

- Individual work:
  - Weekly Project Modules
  - Unit Checkpoint Quizzes
- Team work:
  - 15-619 Project

Read the university policy on Academic Integrity.

Disciplinary Policies

- First offense:
  - Minimum: worse than not doing the work.
  - Maximum: immediate expulsion.

- Second offense results in expulsion. Always.
  - Previously undiscovered offenses can count as “first offense”!

The Penalties are Severe

- Cheating leads to several students being dismissed from the university every semester

LET IT NOT BE YOU!

Academic Integrity Module on OLI

- Required for all students
- Process
  - Pretest Quiz
    - Please take this without looking at the modules
  - Page 1, Overview
  - Page 2, Policies
  - Page 3, Methods of Prevention
  - Quiz
    - Complete this quiz this week
      - By September 6, 2021

Working within Budgets

- Design is a critical element to success
- Develop a budget for
  - Development
  - Testing
  - Drama
- If funds are left over in the budget, feel free to explore and learn!

Tagging is painful, why the penalty?

- Your boss has a budget and a boss
- The budget is allocated among the team
- Your boss has to keep track of how the resources are being spent in order to
  - Re-allocate budget or ask for more resources
- On the cloud, the only way to keep track is through tagging
  - Learn how to tag correctly, don’t complain about penalties!

Getting Help

- TAs
- Piazza
  - Email does not scale
  - Discussion forum to support each other
- Course Overview
  - Tuesdays (recorded)
    - Will be posted before Tuesday at Noon
- Office Hours
  - Check Piazza for Office Hour schedule
  - Will use OHQ and Zoom [links on piazza]

Teaching Staff

- Majd Sakr
  - msakr@cs.cmu.edu
  - Office Hours
    - Tuesday 3pm ET
    - Zoom link on piazza

Marshall An

Project Scientist @ SCS

Siddharth (Sid) Kandimalla

- Senior Project Scientist @ SCS

Adam Zhang

Data Scientist @ SCS

B.S. Stats ML, 2019

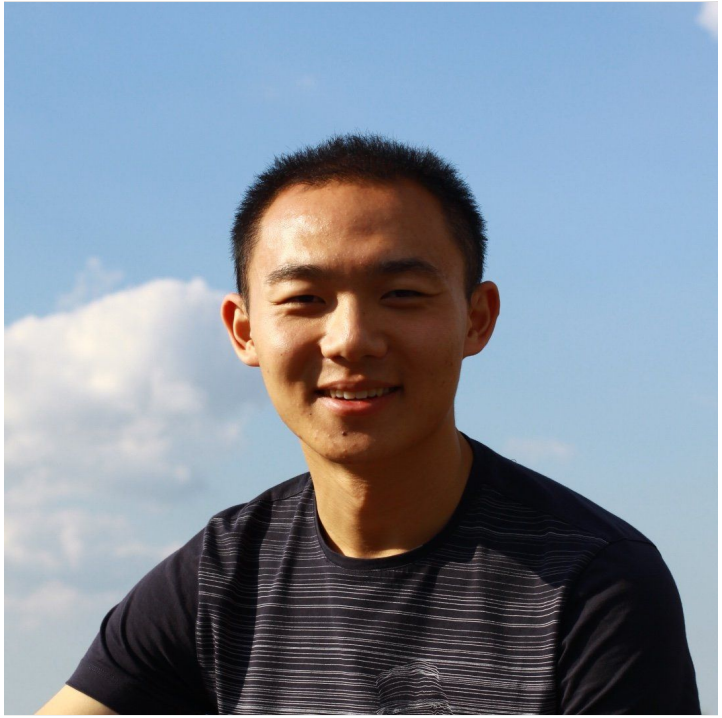

Baljit Singh

Full Stack Engineer @ SCS

Chih-Wei Fang

M.S. in Mobile and IoT Engineering,

Dec. 2021

Yuan Gu

Master of Science in Information Networking, Dec 2021

I like playing basketball and doing workout

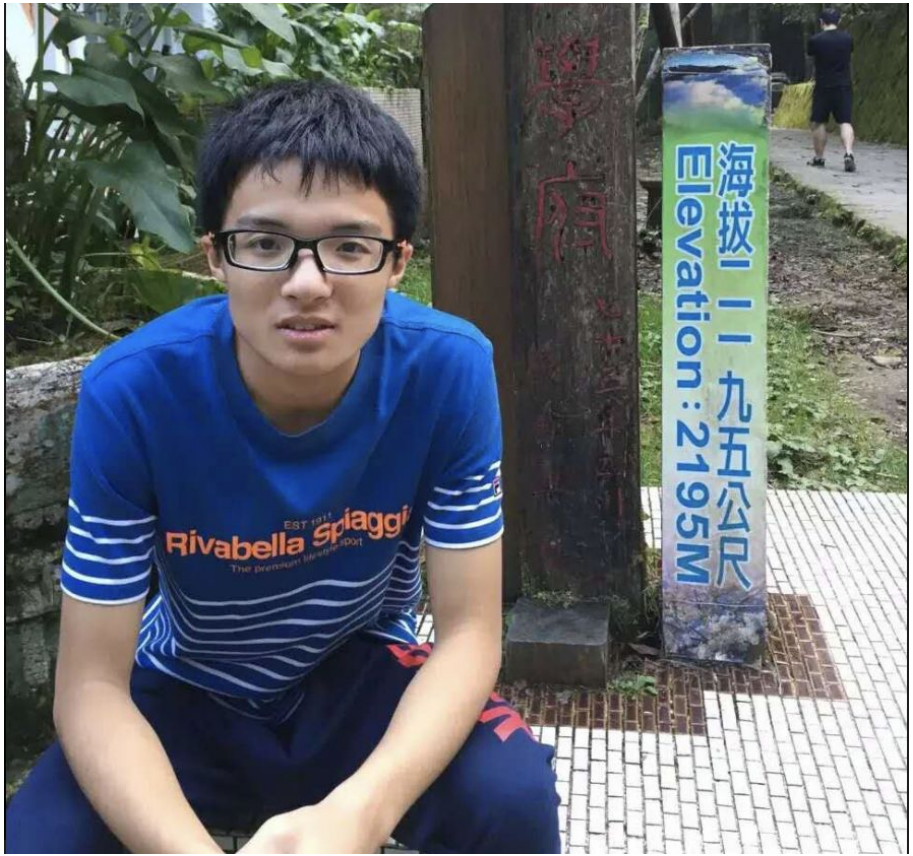

Zhenyuan He

Master of Information Systems Management,

Dec 2021

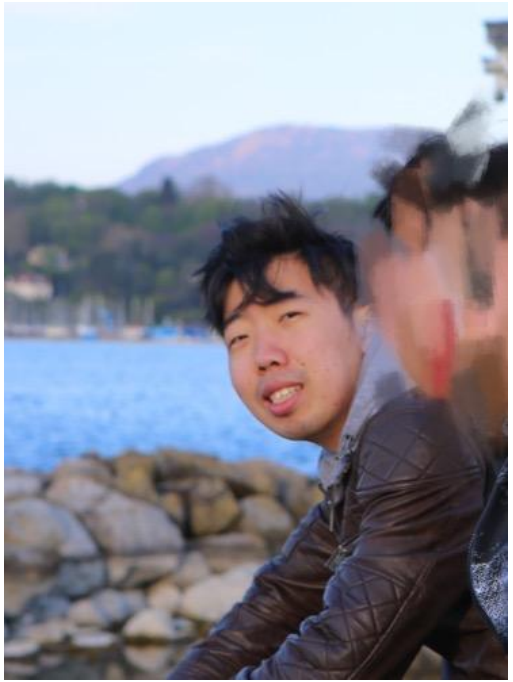

Daria Mashanova

Senior in math and CS

May 2022

Ayoyinka Obisesan

Master of Science in Information Technology, May 2022

Ziteng Shu

M.S. in ECE,

Dec 2022

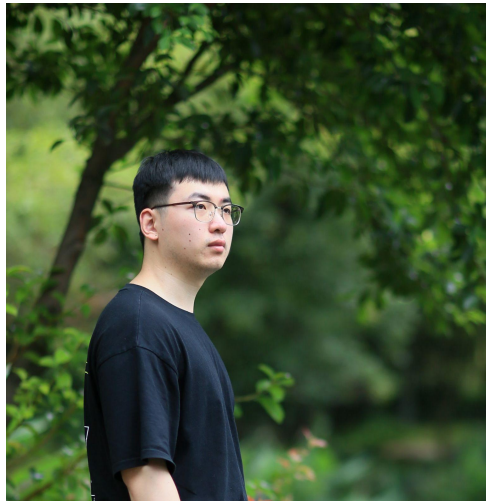

Yifan Song

M.S in Computational Data Science, May 2022

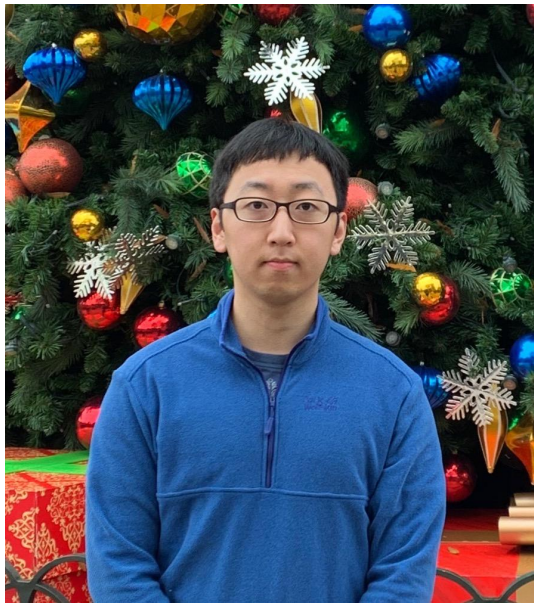

Yao Wang

M.S. in Software Engineering,

Silicon Valley Campus

Dec. 2021

Yuanxin Wang

M.S in Computational Data Science, May 2022

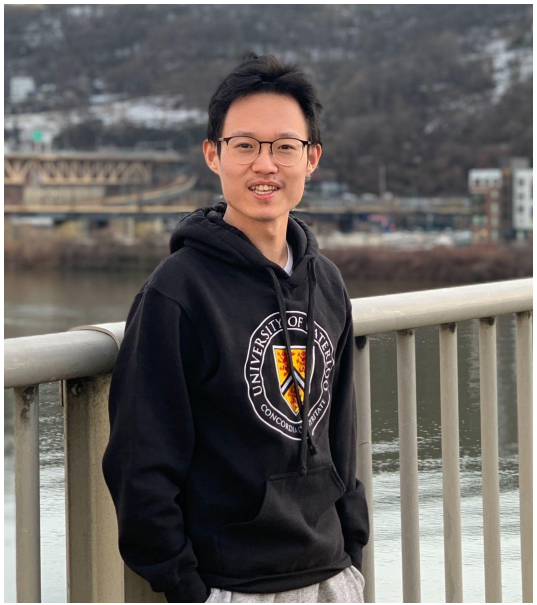

Xuchen Zhang

M.S. in Artificial Intelligence and Innovation

May 2022

I like gaming, skiing

and painting.

Era of Globalization

- Economics
- Communication
- Entertainment
- Sports
- Education
- Compute Services ☺️
  - You're programming the global computer.

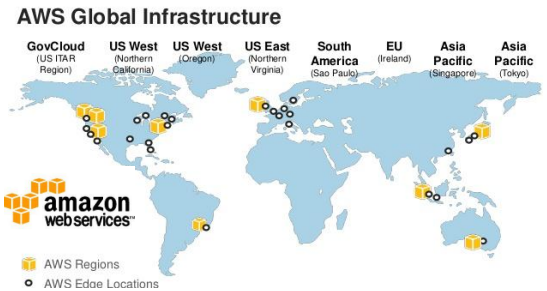

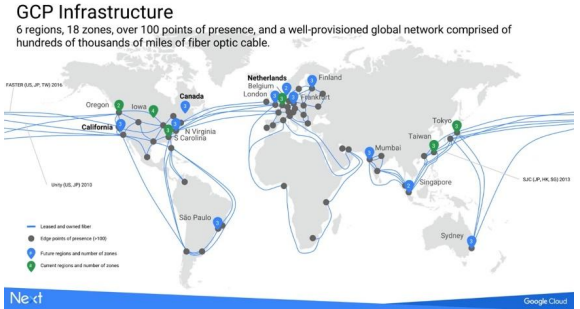

Questions?