

# 15-319/15619: CLOUD COMPUTING

## COURSE DESCRIPTION & SYLLABUS

CARNEGIE MELLON UNIVERSITY

FALL 2018

### 1. OVERVIEW

**Title:** Cloud Computing

**Units:** 15-319 is nine units and 15-619 is 12 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.

**OLI Course:** <http://oli.cmu.edu> (accessed through <https://canvas.cmu.edu>)

**The Project Zone:** <https://TheProject.Zone> (accessed through <https://canvas.cmu.edu>)

**Piazza:** <http://piazza.com/cmu/fall2018/1531915619/home>

**Webpage:** <http://www.cs.cmu.edu/~msakr/15619-f18/>

#### Recitation:

1. Tuesday, 8:00 AM – 8:50 AM, GHC 4307 (Videotaped)
2. Thursday, 4:30 PM – 5:20 PM, GHC 4307 (VC to 118 SV) (First 3 weeks and when needed)

#### Teaching Staff:

**Prof. Majd F. Sakr**

[msakr@cs.cmu.edu](mailto:msakr@cs.cmu.edu)

GHC 7006, +1-412-268-1161

*Office hours:* Tuesday, 3-4pm (Pittsburgh)

**Prof. Seth Goldstein**

[seth@cmu.edu](mailto:seth@cmu.edu)

GHC 7111, +1-412-268-3828

*Office hours:* Thursday, 2-3pm (Pittsburgh)

TAs in Pittsburgh typically hold office hours in GHC 5<sup>th</sup> Floor Teaching Commons. The TA office hours are posted on Piazza:

- Sriharsha Bandaru <[sbandaru@andrew.cmu.edu](mailto:sbandaru@andrew.cmu.edu)>
- Sijia Chen <[sijiac@andrew.cmu.edu](mailto:sijiac@andrew.cmu.edu)>
- Cameron Dashti <[cdashti1@andrew.cmu.edu](mailto:cdashti1@andrew.cmu.edu)>
- Shilun Ding <[shilund@andrew.cmu.edu](mailto:shilund@andrew.cmu.edu)>
- Rong Huang <[rhuang2@andrew.cmu.edu](mailto:rhuang2@andrew.cmu.edu)>
- Kaixiang Li <[kaixianl@andrew.cmu.edu](mailto:kaixianl@andrew.cmu.edu)>
- Ye Li <[yel2@andrew.cmu.edu](mailto:yel2@andrew.cmu.edu)>
- Yanjun Li <[yanjunl2@andrew.cmu.edu](mailto:yanjunl2@andrew.cmu.edu)>
- Zesheng Li <[zeshengl@andrew.cmu.edu](mailto:zeshengl@andrew.cmu.edu)>
- Zizhe Liu <[zizhel@andrew.cmu.edu](mailto:zizhel@andrew.cmu.edu)>
- Bifei Liu <[bifeil@andrew.cmu.edu](mailto:bifeil@andrew.cmu.edu)>
- Clarence Ngoh <[pcn@andrew.cmu.edu](mailto:pcn@andrew.cmu.edu)>
- Xiangyu Song <[xsong1@andrew.cmu.edu](mailto:xsong1@andrew.cmu.edu)>
- Zeyuan Tan <[zeyuant@andrew.cmu.edu](mailto:zeyuant@andrew.cmu.edu)>
- wurong wang <[wurongw@andrew.cmu.edu](mailto:wurongw@andrew.cmu.edu)>
- Zexian Wang <[zexianw@andrew.cmu.edu](mailto:zexianw@andrew.cmu.edu)>
- Runchen Yan <[runcheny@andrew.cmu.edu](mailto:runcheny@andrew.cmu.edu)>
- Wenjie Yang <[wenjiey@andrew.cmu.edu](mailto:wenjiey@andrew.cmu.edu)>
- Wenjie Yang <[wenjiey@andrew.cmu.edu](mailto:wenjiey@andrew.cmu.edu)>
- Tianchen Zhong <[tzhong1@andrew.cmu.edu](mailto:tzhong1@andrew.cmu.edu)>

## 2. COURSE DESCRIPTION

This project-based on-line course focuses on skill building across various aspects of cloud computing. We cover conceptual topics and provide hands-on experience through projects utilizing public cloud infrastructures Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP).

Students will utilize MapReduce, interactive programming using Jupyter Notebooks, and data science libraries to clean, prepare and analyze a large data set. Students will orchestrate the deployment of auto-scaled, load-balanced and fault-tolerant applications using virtual machines (VMs), Docker containers and Kubernetes, as well as serverless computing through Functions as a Service. Students will explore and experiment with different distributed cloud-storage abstractions (distributed file systems and databases) and compare their features, capabilities, applicability and consistency models. In addition, students will develop different analytics applications using batch, iterative and stream processing frameworks. The 15-619 [graduate] students will participate in a team project, which entails designing and implementing a complete web-service solution for querying big data. For the team project, the student teams are evaluated based on the cost and performance of their web service.

Conceptually, the course will introduce this domain and cover the topics of cloud infrastructures, virtualization, software-defined networks and storage, cloud storage, and programming models (analytics frameworks). As an introduction, we will discuss the motivating factors, benefits and challenges of the cloud, as well as service models, service level agreements (SLAs), security, example Cloud Service Providers and use cases. Modern data centers enable many of the economic and technological benefits of the cloud paradigm; hence, we will describe several concepts behind data center design and management and software deployment. Next, we will focus on virtualization as a key cloud technique for offering software, computation and storage services. Within the same theme of virtualization, students will also be introduced to Software Defined Networks and Storage (SDN and SDS). Subsequently, students will learn about different cloud storage concepts including data distribution, durability, consistency and redundancy. We will discuss distributed file systems, NoSQL databases and object storage. Finally, students will learn the details of the MapReduce programming model and gain a broad overview of the Spark, GraphLab programming models as well as message queues (Kafka) and stream processing (Samza).

## 3. COURSE GOALS

In this online course, 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. It is our goal that students will develop the skills needed to become a practitioner or carry out research projects in the domain of cloud computing. Specifically, 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.

Through this process, we aspire for our students to become sophisticated, independent, and resilient problem solvers who are able to overcome challenges and learn.

## 4. LEARNING OUTCOMES

In this project-based course, we have project and conceptual learning objectives.

### 4.1. PROJECT LEARNING OBJECTIVES

The **project learning objectives** (LOs) are common to the undergraduate and graduate offering of the course. The graduate-only (15-619) LOs are designated below. Students will be able to:

Data Analytics	1) Analyze and visualize small and large data sets on the cloud using interactive computing as well as data manipulation and analysis libraries.
Compute & Elasticity	2) Design, implement, test, package, deploy and monitor cloud applications using Virtual Machines (VMs), Containers and Serverless cloud computing services.
Cloud Storage	3) Explore and experiment with different distributed cloud-storage abstractions and compare their features, capabilities and applicability. 4) Orchestrate, deploy and optimize a unified application that integrates heterogeneous SQL and NoSQL database systems. 5) Implement and compare consistency models to recognize the tradeoff between consistency and performance in replicated and distributed cloud storage systems.
Frameworks	6) Design, implement, test and debug applications using interactive, batch and stream processing frameworks and compare their suitability to different problem domains. 7) Illustrate and explain the execution workflow, overhead, fault-tolerance and logical flow of interactive, batch and stream processing frameworks. 8) Train and deploy a machine learning model using a cloud-based framework. 9) Analyze and identify potential sources of bottlenecks in programming frameworks to optimize their performance.
Team Project	10) [15-619 only] Design, build, and deploy a performant, reliable, scalable and fault-tolerant web service on the cloud within a specified budget. 11) [15-619 only] Perform extract, transform and load (ETL) on a large data set. 12) [15-619 only] Design schema as well as configure and optimize cloud-based databases to deal with scale and improve the throughput of a web service. 13) [15-619 only] Explore methods to identify the potential bottlenecks in a cloud-based web service and implement methods to improve system performance.
Overall	14) Practice gathering, cleaning and preparing data for analysis on the cloud. 15) Practice Test-driven Development (TDD) in the software development process. 16) Orchestrate and automate the process of managing and provisioning cloud resources through machine-readable definition files. 17) Make informed decisions about choosing an appropriate cloud tool that will satisfy a set of specified requirements.

### 4.2. CONCEPTUAL LEARNING OBJECTIVES

The **conceptual learning objectives** are below. Students will be able to:

1. Explain the **core concepts** of the cloud computing paradigm: how and why this paradigm shift came about, the characteristics, advantages and challenges brought about by the various models and services in cloud computing.
2. Apply fundamental concepts in **cloud infrastructures** to understand the tradeoffs in power, efficiency and cost, and then study how to leverage and manage single and multiple datacenters to build and deploy cloud applications that are resilient, elastic and cost-efficient.

3. Discuss **system, network and storage virtualization** and outline their role in enabling the cloud computing system model.
4. Describe the overall organization of data and the fundamental concepts of **cloud storage**. Identify the problems of scale and management in big data. Discuss various storage abstractions. Compare and contrast different types of distributed file systems and discuss their design considerations. Compare and contrast different types of databases and discuss their design tradeoffs. Discuss the concepts of cloud object storage. Enumerate the different types of block devices used in data storage.
5. Explain the main execution flow, scheduling and fault tolerance concepts in various **cloud programming models**. Recall and contrast different cloud programming models (MapReduce, Spark, GraphLab, Spark Streaming and Samza).

## 5. COURSE ORGANIZATION

Your participation in the course will involve several forms of activity:

1. Reading the online coursework content for each unit on OLI.
2. Completing the unscored inline activities for each unit (Review activities on OLI).
3. Completing the graded checkpoint weekly quizzes after each unit.
4. Complete projects, which are performed on the cloud and submitted through TheProject.Zone.
5. AssessMe assessments, unscored short quizzes to unlock subsequent project sections.
6. Complete a team project on building a complete web service.

Students should regularly check OLI to see when new content or checkpoint quizzes are made available. Projects and Checkpoint quizzes must be completed by the due dates posted on TheProject.Zone.

## 6. GETTING HELP

Students are encouraged to ask questions about content and projects through **Piazza**, where an online class portal has been created for this course. The course link for Piazza is:

<http://piazza.com/cmu/fall2018/1531915619/home>.

There is video recording of a weekly recitation in Pittsburgh which is made available to all students. The teaching staff will discuss any major questions that have been posted to Piazza or by email. For urgent communication with the teaching staff, it is best to post on Piazza and then send email.

We will use the course website as the basic portal for the class. The course content is entirely on OLI. The project write-ups, submission, scoreboard and grades are on TheProject.Zone. The checkpoint quizzes are on OLI. OLI can be reached through Canvas. Announcements, discussions and questions are posted on Piazza.

## 7. POLICIES

### WORKING ALONE ON PROJECTS

---

Projects that are assigned to single students should be performed individually.

### HANDING IN PROJECTS

---

All assessments are due at 11:59 PM EST (one minute before midnight) on the due dates specified on OLI or TheProject.Zone. All hand-ins are electronic, and use the OLI Checkpoint system and TheProject.Zone.

### APPEALING GRADES

---

After each project module is graded, you have seven calendar days to appeal your grade. All your appeals should be provided by email to Prof. Sakr.

## 8. ASSESSMENT

Inline activities (“Learn by Doing” and “Did I Get This”), which are available in most pages in the OLI course, are simple, non-graded activities to assess your comprehension of the material as you read through the course material. You are advised to complete all of the inline activities before proceeding through to the next page or module. If you missed many of the activities, it is recommended that you review the material again.

There are five units consisting of modules of content on OLI, each week has a Checkpoint Quiz that you must complete before the deadline posted on OLI. Each weekly Checkpoint Quiz will be worth 2% of your total grade. It is your responsibility to ensure that the quiz is submitted prior to the deadline. You will have only a single attempt to complete each Checkpoint Quiz on OLI.

This course includes four individual project themes. Each project theme consists of several project modules. Every week, a project module has to be completed based on the deadlines posted on TheProject.Zone. The write-up required to complete each project module is available on TheProject.Zone. Each module has a submission process that is specific to the project module that is due. It is the students’ responsibility to make sure that all project work is completed and that the project module is submitted prior to the deadline. Students typically have multiple attempts to submit the project module on TheProject.Zone.

15-619 students have to complete a team-based multi-week project in parallel to the weekly Project Modules.

Type	Number	Weight
Content Checkpoint Quizzes	10 out of 11	20%
Projects	10 for 15-319 10 + Team for 15-619	80%
<b>Total Grade</b>		<b>100%</b>

## 9. CHEATING

We urge each student to carefully read the [university policy on academic integrity](#), which outlines the policy on cheating, plagiarism or unauthorized assistance. It is the responsibility of each student to produce her/his own original academic work. Collaboration or assistance on academic work to be graded is not permitted unless explicitly authorized by the course instructor. Each unit checkpoint quiz or project module submitted must be the sole work of the student turning it in. Student work on the cloud is logged, submitted work will be closely monitored by automatic cheat checkers, and students may be asked to explain any suspicious similarities with any piece of code available. The following are guidelines on what collaboration is authorized and what is not:

### WHAT IS CHEATING?

---

1. Sharing code or other electronic files by either copying, retyping, looking at, or supplying a copy of any file. Copying any code from the internet (stackoverflow.com or github or others).
2. Copying answers to any checkpoint quiz from another individual, published or unpublished written sources, and electronic sources.
3. Collaborating with another student or another individual on checkpoint quizzes or project modules.
4. Sharing written work, looking at, copying, or supplying work from another individual, published or unpublished written sources, and electronic sources.
5. Collaboration in team projects is strictly limited to the members of the team.

## WHAT IS NOT CHEATING?

1. Clarifying ambiguities or vague points in class handouts.
2. Helping others use computer systems, networks, compilers, debuggers, profilers, or system facilities.
3. Helping others with high-level design issues.
4. Guiding others through code debugging but not debugging for them.

Cheating in projects will also be strictly monitored and penalized. Be aware of what constitutes cheating (and what does not) while interacting with students. You cannot share or use written code, and other electronic files from students. If you are unsure, ask the teaching staff.

Be sure to store your work in protected directories. The penalty for cheating is severe, and might jeopardize your career – cheating is simply not worth the trouble. By cheating in the course, you are cheating yourself; the worst outcome of cheating is missing an opportunity to learn. In addition, you will be removed from the course with a failing grade. We also place a record of the incident in the student's permanent record.

## 10. CONCEPTUAL TOPICS

The course content will be structured into the following units:

Unit #	Title	Modules and Description
1	<b>Introduction</b>	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	<b>Cloud Infrastructure</b>	Historical Perspective of Data Centers Datacenter Components: IT Equipment and Facilities Design Considerations: Requirements, Power, Efficiency, & Redundancy Power Calculations, PUE and Challenges in Cloud Data Centers Cloud Management and Cloud Software Deployment Considerations
3	<b>Virtualization</b>	Virtualization (CPU, Memory, I/O), Case Study: Amazon EC2 Software Defined Networks (SDN) Software Defined Storage (SDS)
4	<b>Cloud Storage</b>	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	<b>Programming Models</b>	Distributed Programming for the Cloud Data-Parallel Analytics with Hadoop MapReduce (YARN); Iterative Data-Parallel Iterative Analytics (Spark); Graph-Parallel Analytics with GraphLab 2.0 (PowerGraph); Stream Processing (Samza)

## 11. PROJECTS

The programming projects are geared towards providing hands-on experience with various cloud technologies. Students will learn to develop all projects using various public cloud services (AWS, Azure and GCP). For each project, students are expected to work within a specified budget otherwise they risk being penalized.

### 11.1. PROJECT 1: BIG DATA ANALYSIS

Students will setup cloud accounts, work with provisioning management software and launch instances on Amazon EC2. Students will learn the benefits and tradeoffs of running programs in parallel, using AWS EMR or Azure HDInsight or GCP Dataproc, versus sequential on a large dataset within particular cost constraints. Students will utilize Jupyter Notebooks to solve a set of data analytics problems.

### 11.2. PROJECT 2: CLOUD ELASTICITY

In this project, students will learn about cloud elasticity through virtual machines, containers and functions using the AWS, Azure and GCP APIs. Students will be tasked with developing their own elastic services for a dynamically changing load scenario. Students will work with the Load Balancing and Auto Scaling services on AWS to mitigate varying loads on the server. Furthermore, students will deploy a web service using Docker Containers and Kubernetes on multiple cloud platforms. Lastly, students will develop a video indexing service using functions on multiple clouds.

### 11.3. PROJECT 3: CLOUD STORAGE

Using cloud resources and services, students will explore cloud storage technologies to evaluate their capabilities, limitations, and applicability to application domains. Students begin by exploring the limitations of traditional filesystems, and then compare them to relational databases (MySQL) and NoSQL databases (HBase). Next, students build a social network timeline using heterogeneous back-end storage systems. The project will cover several storage systems, including low-latency KV stores, and NoSQL databases, (e.g., Apache HBase, Amazon RDS, MongoDB and others). Finally, students implement strong and eventual consistency for a geo-replicated key-value stores.

### 11.4. PROJECT 4: CLOUD FRAMEWORKS

Students will work on developing applications using the Spark, and Samza and Machine Learning frameworks to experience iterative batch, graph and stream processing. Students will implement a social graph analysis algorithm on Apache Spark. Students will build and deploy a machine learning model using transfer learning on the cloud. Students will learn to analyze streaming data to perform real-time processing of multiple data streams using Apache Kafka/Samza.

### 11.5. 15-619 TEAM PROJECT: TWITTER ANALYTICS WEB SERVICE

Students will work in teams to design and implement a complete web-service that uses the REST interface to respond to queries that require running an analytics job on a large (1.2TB) Twitter data set which is stored in a database (MySQL, HBASE, etc.). Student teams are expected to use different tools and services to achieve build a performing web-service that meets the requirements. The web-services are evaluated through a load generator for a fixed time period (several hours) by measuring the cost of cloud resources used and their system's performance (throughput). There is an upper bound on the budget which could cause students to be disqualified. Students are evaluated based on how their service performs compared to a baseline.

## 12. SCHEDULE

The tentative schedule is as follows (specific deadlines are posted on OLI and TheProject.Zone):

Week	Monday	OLI Content	Individual Projects	Team Project	Quizzes
1	8/27/2018	Unit 1, Module 1, 2	Primers/P0 (Sep 02)		Q0 (Ac. Integ.)
2	9/03/2018	Unit 1, Module 1, 2	P1.1 (Sep 09)		Q1 (Sep 07)
3	9/10/2018	Unit 2, Module 3, 4	P2.1 (Sep 16)		Q2 (Sep 14)
4	9/17/2018	Unit 2, Module 5, 6	P2.2 (Sep 23)		Q3 (Sep 21)
5	9/24/2018	Unit 3, Module 7, 8, 9	P2.3 (Sep 30)		Q4 (Sep 28)
6	10/01/2018	Unit 3, Module 10, 11, 12	P3.1 (Oct 07)		Q5 (Oct 05)
7	10/08/2018	Unit 3, Module 13	P3.2 (Oct 14)	Project Out (Oct 08)	Q6 (Oct 12)
8	10/15/2018	Unit 4, Module 14	P3.3 (Oct 21)		Q7 (Oct 18)
9	10/22/2018	Unit 4, Module 15, 16, 17		Phase 1 Due (Oct 28)	Q8 (Oct 26)
10	10/29/2018	Unit 4, Module 18	P4.1 (Nov 4)		Q9 (Nov 02)
11	11/05/2018	Unit 5, Module 19, 20		Phase 2 Due (Nov 11)	Q10 (Nov 09)
12	11/12/2018	Unit 5, Module 21, 22	P4.2 (Nov 18)		Q11 (Nov 16)
13	11/19/2018	Thanksgiving			
14	11/26/2018			Phase 3 Due (Dec 02)	
15	12/3/2018		P4.3 (Dec 7)		
16	12/10/2018				

## 13. ACCOMMODATIONS FOR STUDENTS WITH DISABILITIES

If you have a disability and have an accommodations letter from the Disability Resources office, I encourage you to discuss your accommodations and needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, I encourage you to contact them at [access@andrew.cmu.edu](mailto:access@andrew.cmu.edu).



## 14. TAKE CARE OF YOURSELF

Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress.

All of us benefit from support during times of struggle. You are not alone. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is often helpful.

If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website at <http://www.cmu.edu/counseling/>. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

If you or someone you know is feeling suicidal or in danger of self-harm, call someone immediately, day or night:

**CaPS:** 412-268-2922

**Resolve Crisis Network:** 888-796-8226

If the situation is life threatening, call the police:

- **On campus:** CMU Police: 412-268-2323
- **Off campus:** 911

If you have questions about this or your coursework, please let me know.