

15-319/15619: CLOUD COMPUTING

COURSE DESCRIPTION & SYLLABUS

CARNEGIE MELLON UNIVERSITY

FALL 2014

1. OVERVIEW

Title: Cloud Computing

Units: 15-319 is 9 units and 15-619 is 12 units.

Pre-requisites for CMU Students: A “C” or better in 15-213.

Pre-requisites for Others: Knowledge of Computer Systems, Java programming.

OLI Course Link: <http://oli.cmu.edu>

Piazza Link: <http://piazza.com/cmu/fall2014/1531915619/home>

Recitation:

1. Tuesday, 8:00 AM – 8:50 AM, GHC 4307 (Videotaped)
2. Thursday, 4:30 PM – 5:20 PM, GHC 4307

Teaching Staff:

Prof. Majd F. Sakr

msakr@cs.cmu.edu

GHC 7006, +1-412-268-1161

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

Pittsburgh:

TAs hold office hours in GHC 4122 and 4126, office hours posted on Piazza:

- Andi Ni
andin@andrew.cmu.edu
- Bin Feng
bfeng@andrew.cmu.edu
- Felipe Faraco
fsoaresf@andrew.cmu.edu
- Jialiang Lin
jialianl@andrew.cmu.edu
- Jiang Xue
jiangx@andrew.cmu.edu
- Jiten Mehta
jitenm@andrew.cmu.edu
- Junqi Wang
junqiw@andrew.cmu.edu
- Kasipan Kanniah
kkanniah@andrew.cmu.edu
- Lina Li
linal@andrew.cmu.edu
- Lu Qu
lqu@andrew.cmu.edu

- Lu Zeng
luzeng@andrew.cmu.edu
- Luning Pan
luningp@andrew.cmu.edu
- Mrigesh Kalvani
mkalvai@andrew.cmu.edu
- Preston Lin
yuyul@andrew.cmu.edu
- Ravi Chandra Bandlamudi Venkata
rbandlam@andrew.cmu.edu
- Suhail Rehman
suhailr@andrew.cmu.edu
- Vivek Munagala
vpm@andrew.cmu.edu
- Xiaokang Zhang
xiaokanz@andrew.cmu.edu
- Yiqi Wu
yiqiw@andrew.cmu.edu
- Yishuang Pang
yishuanp@andrew.cmu.edu
- Yu Wu
yuwu1@andrew.cmu.edu

2. COURSE DESCRIPTION

This on-line course gives students an overview of the field of Cloud Computing, its enabling technologies, main building blocks, and hands-on experience through projects utilizing a public cloud (Amazon Web Services). Cloud computing services are being adopted widely across a variety of organizations and in many domains. Simply, cloud computing is the delivery of computing as a service over a network, whereby distributed resources are rented, rather than owned, by an end user as a utility.

The course will introduce this domain and cover the topics of data centers, virtualization, cloud storage, and programming models. As an introduction, we will discuss the motivating factors, benefits, challenges, and service models. 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. Next, we will focus on virtualization as a key cloud technique for offering software, computation and storage services. We will study how CPU, memory and I/O resources are virtualized, with examples from Xen and VMWare, and present real use cases such as Amazon EC2. Subsequently, students will learn about different cloud storage concepts including data distribution, durability, consistency and redundancy. HDFS, PVFS, HBASE, Cassandra and S3 will be presented as case studies. Students will understand the details of the MapReduce programming model and gain a broad overview of alternative programming models such as Pregel, and GraphLab.

Students will work with Amazon Web Services, use them to rent and provision compute resources and then program and deploy applications that run on these resources. Students will develop and evaluate scaling and load balancing solutions. In addition, students will work with cloud storage systems and learn to develop applications in the MapReduce programming paradigm. The 15-619 students will have to complete an extra project which entails designing and implementing a complete web-service solution for querying big data. For the extra project, the student teams are evaluated based on the cost and performance of their web service.

3. COURSE OBJECTIVES

In this on-line course we plan to give students an overview of the field of Cloud Computing, and an in-depth study into its enabling technologies and main building blocks. Students will gain hands-on experience solving relevant problems through projects that will utilize existing public cloud tools. It is our objective that students will develop the skills needed to become a practitioner or carry out research projects in this domain. Specifically, the course has the following objectives:

Students will learn

- 1) the fundamental ideas behind Cloud Computing, the evolution of the paradigm, its applicability; benefits, as well as current and future challenges;
- 2) the basic ideas and principles in data center design and management;
- 3) different CPU, memory and I/O virtualization techniques that serve in offering software, computation and storage services on the cloud;
- 4) about cloud storage technologies and relevant distributed file systems;
- 5) the variety of programming models and develop working experience in one of them.

4. LEARNING OUTCOMES

The primary learning outcomes of this course are five-fold. 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 the fundamental concepts in **datacenters** to understand the tradeoffs in power, efficiency and cost.
- 3) Discuss **system virtualization** and outline its role in enabling the cloud computing system model.
- 4) Illustrate the fundamental concepts of **cloud storage** and demonstrate their use in storage systems such as Amazon S3 and HDFS.
- 5) Analyze various **cloud programming models** and apply them to solve problems on the cloud.

4.1. BASIC CONCEPTS

This module will provide a broad overview of cloud computing, its history, technology overview, benefits, risks and the economic motivation for it. Upon completion of this module, students will be able to:

- 4.1.1. Explain the concept of “cloud computing”.
- 4.1.2. Briefly recall the recent history of cloud computing, illustrating its evolution.
- 4.1.3. List some of the enabling technologies in cloud computing and discuss their significance.
- 4.1.4. Discuss some of the advantages and disadvantages of the cloud paradigm.
- 4.1.5. Articulate the economic benefits as well as issues/risks of the cloud paradigm for businesses as well as cloud providers.
- 4.1.6. Associate the various layers in the cloud building blocks and differentiate cloud service models.

4.2. DATA CENTERS

This module will provide a historical overview of data centers, along with design considerations. Students will learn to apply methods to evaluate data centers. Upon completion of this module, students will be able to:

- 4.2.1. Describe the evolution of data centers.
- 4.2.2. Outline the architecture of a modern data center in detail.
- 4.2.3. Indicate design considerations and discuss their impact.
- 4.2.4. Demonstrate the ability to calculate various power requirements of a data center.

4.3. CLOUD RESOURCE MANAGEMENT

Students will learn how virtualization can allow software and hardware images (e.g., virtual machines) to run side-by-side on a single cloud data center while provided security, resource and failure isolations. They will understand how virtualization enables clouds to offer software, computation, and storage as services as well as attain agility and elasticity properties. We will discuss resource virtualization in detail and present multiple examples from Xen and VMware. Finally, we will present a real use case such as Amazon EC2. After finishing this unit students will be able to:

- 4.3.1. Identify major reasons for why virtualization is becoming enormously useful, especially on the cloud
- 4.3.2. Explain different isolation types such as fault, resource, and security isolations provided by virtualization and utilized by the cloud

- 4.3.3. Indicate how system complexity can be managed in terms of levels of abstractions and well-defined interfaces, and their applicability to virtualization and the cloud**
- 4.3.4. Define resource sharing as provided by virtualization and discuss how it can be offered in space and time via physical and logical partitioning**
- 4.3.5. Define virtualization and identify different virtual machine types such as process and system virtual machines**
- 4.3.6. Identify conditions for virtualizing CPUs, recognize the difference between full virtualization and paravirtualization, explain emulation as a major technique for CPU virtualization, and examine virtual CPU scheduling in Xen**
- 4.3.7. Outline the difference between classical OS virtual memory and system memory virtualization, explain the multiple levels of page mapping as imposed by memory virtualization, define memory over-commitment and illustrate VMWare memory ballooning as a reclamation technique for memory over-committed virtualized systems**
- 4.3.8. Explain how CPU and I/O devices can communicate with and without virtualization, identify the three main interfaces, system call, device driver and operation level at which I/O virtualization can be carried, and apply I/O virtualization to Xen**

4.4. CLOUD STORAGE

This module will provide a broad overview of storage technologies and concepts of cloud storage. It will also provide a detailed study of Amazon S3, EBS and distributed file systems and databases. Students will be able to:

- 4.4.1. Describe the overall organization of data and storage.**
- 4.4.2. List the various types of data within the data taxonomy and classify different data types within the data taxonomy.**
- 4.4.3. Identify the problems of scale and management in big data. Discuss the various abstractions in storage.**
- 4.4.4. Compare and contrast different types of file systems. Discuss the design considerations of file systems. Compare and contrast Hadoop Distributed File System (HDFS) with Parallel Virtual File System (PVFS).**
- 4.4.5. Compare and contrast different types of databases. Discuss the design tradeoffs involved with databases.**
- 4.4.6. Discuss the concepts of cloud object storage. Enumerate the different types of block devices used in data storage.**

4.5. PROGRAMMING MODELS

Students will be given an overview on a variety of cloud-based programming models. Students will understand the benefits and limitations of each so that they can assess applicability based on the problem domain. Students will gain working experience in one of these programming models. Upon completion of this module students will be able to:

- 4.5.1. Explain the fundamental aspects of parallel and distributed programming models.**
- 4.5.2. Recall the different cloud programming models (MapReduce, Pregel, and GraphLab)**
- 4.5.3. Explain the main concepts in the MapReduce programming model.**

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 inline activities for each unit (“Learn by doing” activities and “Did I get this” Review activities).
3. Completing the graded checkpoint quizzes after each unit.
4. Complete projects which are performed using AWS and submitted through OLI.
5. Complete a team project on building a complete web service.

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

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/fall2014/1531915619/home>.

There is a weekly recitation (either on 8AM on Tuesdays at GHC 4307 or 4:30PM on Thursdays at GHC 4307) in Pittsburgh. The teaching staff will be on hand to highlight and discuss any major questions that have been posted to Piazza or by email. For urgent communication with the teaching staff, it is best to send an email.

We will use the course website as the main information repository about the class. The course content is entirely on OLI. The project write-ups and checkpoint quizzes are on OLI. OLI can be reached through Blackboard. 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 assignments/projects are due at 11:59 PM EST (one minute before midnight) on the due dates specified on OLI or the course web page. All hand-ins are electronic, and use the OLI Checkpoint system or Blackboard.

APPEALING GRADES

After each project phase 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 of content on OLI, each has a Checkpoint Quiz that you must complete before the deadline posted on OLI. Each Checkpoint Quiz will be worth 5% of your total grade. The quizzes require several hours to complete. 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 projects. Each project consists of several Project Modules on OLI. Every week, a Project Module has to be completed based on the deadlines posted on OLI. The write-up required to complete each Project Module is available on OLI. Each module has a Checkpoint Quiz 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 Checkpoint Quiz on OLI is submitted prior to the deadline. Students typically have three attempts to submit the Project Module Checkpoint Quizzes. The attempt with the maximum score is chosen as the grade for that module.

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 (equal weight)	5	25%
Projects (equal weight)	4 (15-319) 5 (15-619)	75%
Total Grade		100%

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 AWS 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 either by copying, retyping, looking at, or supplying a copy of any file. Copying any code from the internet without referencing it.
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 Unit Checkpoint Quizzes or Project Module Checkpoint Quizzes.
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 the computer systems, networks, compilers, debuggers, profilers, or other system facilities.
3. Helping others with high-level design issues.
4. Helping others debug their code.

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. SYLLABUS

The course will be structured into the following units:

Unit #	Title	Modules and Description
1	Introduction	Introduction to Cloud Computing Building Blocks and Service Models in Cloud Computing
2	Data centers	Historical Perspective Datacenter Components Design Considerations Power Calculations
3	Virtualization	Resource Abstraction Resource Sharing Sandboxing Case Study: Amazon EC2
4	Cloud Storage	Introduction to Storage Systems Cloud Storage Concepts Distributed File Systems Cloud Databases Case Study: Amazon Object Storage
6	Programming Models	Introduction to Programming Models Variety of Programming Models Case Studies: MapReduce, Pregel, GraphLab

11. PROJECTS

The programming projects in this course will be geared towards providing hands-on experience with various cloud technologies. Students will learn to work with various Amazon Web Services products including EC2, EMR and S3.

11.1. PROJECT 1: BIG DATA ANALYSIS

Students will work with Amazon AWS and provision their first compute resources. Students will setup AWS accounts, work with provisioning management software and launch instances on Amazon EC2. Students will learn the benefits and tradeoffs of running programs in parallel versus sequential on a large dataset. Students will have to solve a problem using resources provisioned in Amazon within particular cost constraints.

11.2. PROJECT 2: CLOUD ELASTICITY

In this project, students will learn about cloud elasticity. Students will be first tasked with developing their own elastic services for a web server scenario using AWS APIs. Students will then work with the Elastic Load Balancing and Auto Scaling services at AWS to mitigate varying loads on the web server.

11.3. PROJECT 3: CLOUD STORAGE

Students will work on cloud storage technologies to evaluate their capabilities and limitations. Students will be given a workload and will be required to run them on EC2 instances. Student will store datasets on ephemeral, elastic-block-store (EBS) and S3 buckets and run the workload. Student will evaluate vertical and horizontal scaling solutions. Students will deploy DynamoDB, an AWS NoSQL database solution, for a given image dataset. Finally, students will compare DynamoDB to HBASE using the YCSB benchmark.

11.4. PROJECT 4: PROGRAMMING MODELS

In this project, students will work on developing applications using the MapReduce programming model. Students will write their own MapReduce code using Apache Hadoop and provision instances on Amazon EC2 to run them. For this project, students will build their own input text predictor, similar to [Google Instant](#). Students will build the input text predictor from a large text corpus by generating a list of **n-grams**, building a **statistical language model** using the n-grams, and creating a **user interface**.

11.5. 15-619 EXTRA PROJECT: 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 twitter data set which is stored in a database (MySQL, HBASE, etc.). The students' web-services are evaluated through a load generator for a fixed time period (several hours) by measuring the cost of AWS 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):

Date	OLI Content	Quiz	Project	Extra Project
25-Aug-14	Unit 1 –		Primer	
1-Sep-14	Introduction	Unit 1 Checkpoint Quiz	Project 1	
8-Sep-14	Unit 2 –			
15-Sep-14	Datacenters	Unit 2 Checkpoint Quiz	Project 2	
22-Sep-14	Unit 3 – Virtualization			
29-Sep-14				
6-Oct-14		Unit 3 Checkpoint Quiz		
13-Oct-14	Unit 4 – Cloud Storage		Project 3	15-619 Extra Project
20-Mar-14				
27-Oct-14		Unit 4 Checkpoint Quiz		
3-Nov-14	Unit 5 – Programing Models		Project 4	
10-Nov-14				
17-Nov-14				
24-Nov-14				
1-Dec-14		Unit 5 Checkpoint Quiz		