

CS15-319 / 15-619

Cloud Computing

Recitation 1

Course Overview and Introduction

January 15th, 2013

<http://www.qatar.cmu.edu/~msakr/15319-s13/>

جامعة كارنيغي ميلون في قطر
Carnegie Mellon University Qatar

Online Course Content - OLI

Course content is on the Open Learning Initiative
Students should register:

- <http://community.oli.cmu.edu>

Course Objectives

Students will learn:

- the fundamental ideas behind Cloud Computing;
- the basic ideas and principles in data center design and management;
- the resource sharing and virtualization techniques that serve in offering software, computation and storage services on the cloud;
- about cloud storage technologies and relevant distributed file systems;
- the variety of programming models and develop working experience in one of them.

Syllabus

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

15-319/15619: CLOUD COMPUTING

ONLINE LEARNING INITIATIVE
COURSE DESCRIPTION & SYLLABUS
CARNEGIE MELLON UNIVERSITY IN QATAR
SPRING 2013

1. OVERVIEW

Title: Cloud Computing
Units: 15319 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://community.oli.cmu.edu>
OLI Course Key: 15319s13
Piazza Link: <https://piazza.com/class/spring2013/1531915619>
Course Calendar: [Google Calendar Link](#)
Course Mailing List: 15319-s13@lists.qatar.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 4 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 Google App Engine and Amazon EC2. Subsequently, students will learn about different cloud storage concepts including data distribution, durability, consistency and redundancy. HDFS, PVFS and S3 will be presented as examples of underlying distributed file systems. Students will understand the details of the MapReduce programming model and gain a broad overview of alternative programming models such as Pregel, Dryad, Dremel, and GraphLab, among others.

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. In addition, students will work with cloud storage systems and learn to develop applications in the MapReduce programming paradigm.

Target Audience

- CS Majors
- Juniors / Seniors
- Pre-requisites:
 - 15213 – Introduction to Computer Systems


Teaching Staff / Getting Help

Prof. Majd F. Sakr



msakr@qatar.cmu.edu
(1016)

Dr. Mohammad Hammoud



mhhammou@qatar.cmu.edu
(1013)

Suhail Rehman



suhailr@qatar.cmu.edu
(2044)

Jason Boles



jboles@qatar.cmu.edu
(2044)

Course Organization




Course Units



Course Projects



Weekly Recitations



Discussions on Piazza

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
3	Cloud Resource Management	Resource Abstraction Resource Sharing Sandboxing Case Studies :Google Apps, Google App Engine and Amazon EC2
4	Cloud Storage	Introduction to Storage Systems Cloud Storage Concepts Distributed File Systems Cloud Databases Case Study: Amazon Storage
6	Programming Models	Introduction to Programming Models Variety of Programming Models Case Study: MapReduce

Projects

- Four Projects (all on AWS):
 1. Experience with Amazon Web Services
 2. Virtualization
 3. Cloud Storage
 4. Programming Models

Special Note on Amazon EC2



- Paid Cloud Service – billed by the hour.
- Start a resource only when you need them.
- You can keep one micro instance running 24x7
- Terminate all other resources as soon as you are done with them.

Grading

Course Elements	#	Weight
Projects	4	75%
OLI Unit Checkpoint Quizzes	5	25%

Tentative Schedule

Week	OLI Content	Checkpoint Quiz Due Dates	Project
13-Jan-13	Unit 1 - Introduction		AWS Setup
20-Jan-13		Unit 1 Checkpoint Quiz	Project 1
27-Jan-13	Unit 2 - Datacenters		
3-Feb-13		Unit 2 Checkpoint Quiz	
10-Feb-13	Unit 3 - Resource Management		Project 2
17-Feb-13			
24-Feb-13			
3-Mar-13		Unit 3 Checkpoint Quiz	
10-Mar-13	Unit 4 - Cloud Storage		Project 3
17-Mar-13			
24-Mar-13		Unit 4 Checkpoint Quiz	
31-Mar-13	Unit 5 - Programming Models		Project 4
7-Apr-13			
14-Apr-13			
21-Apr-13		Unit 5 Checkpoint Quiz	

15-619 Students

- We are working with Prof. Garth Gibson
 - Project additions
 - Stay tuned for details



Questions?

