CS15-319: Cloud Computing

Lecture 1

Course Overview and Introduction

Prof. Majd F. Sakr

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Our story…

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Our Data Now…

Documents and Media

Personal Data

Emails, Calendars, Contacts, Location Information, etc…

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

We Live in a World of Data…

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

The World of Data

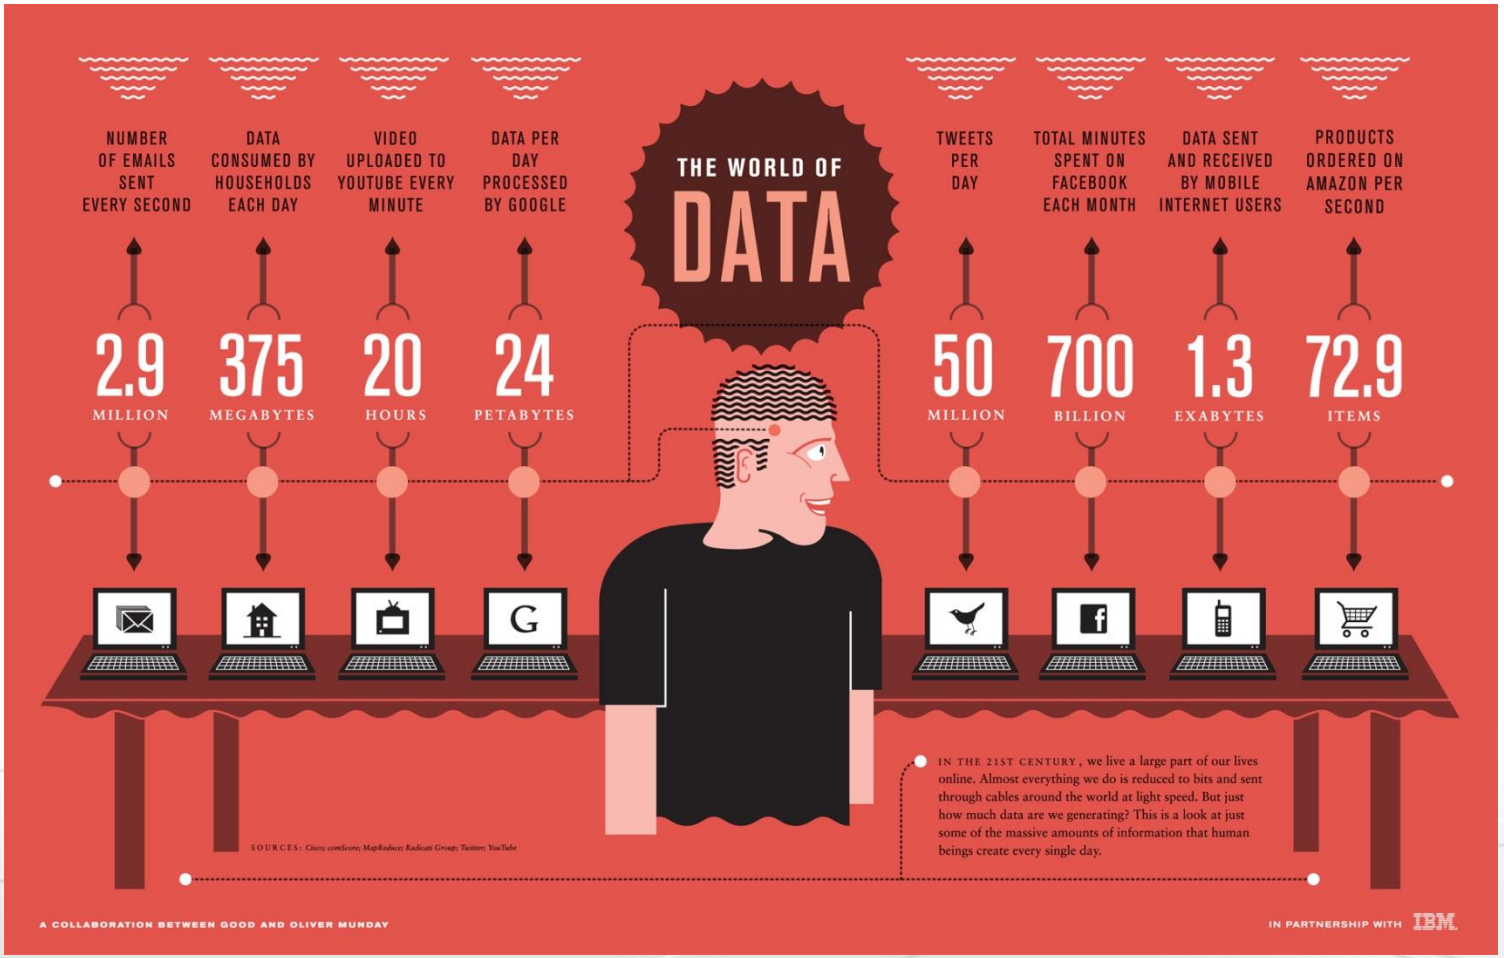

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Big Data

- Big data is defined as large pools of data that can be captured, communicated, aggregated, stored, and analyzed.

- Data continues to grow:

  - In mid-2010, the information universe carried 1.2 zettabytes and 2020 predictions expect nearly 44 times more at 35 zettabytes coming our way.

- Applications are becoming *data-intensive*.

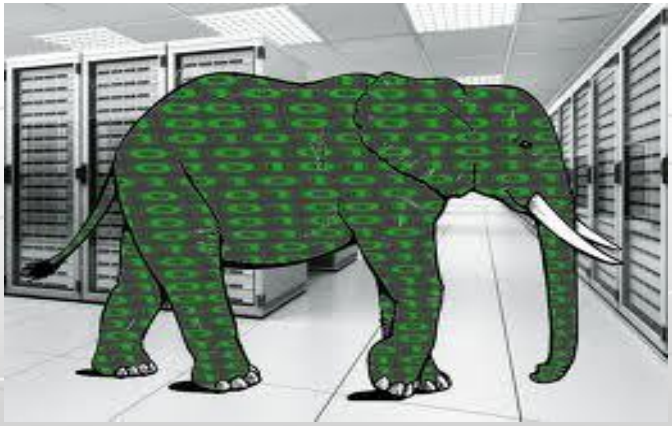

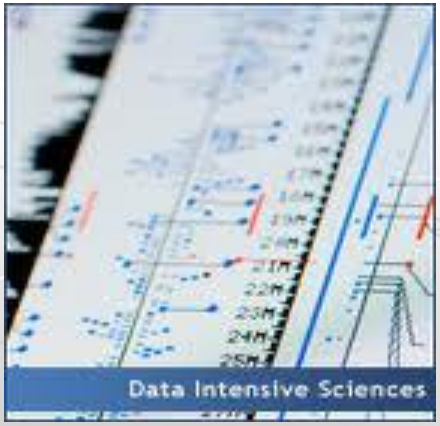

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

What Do We Do With Data?


Store

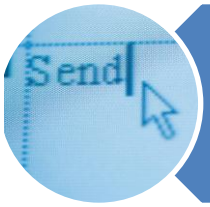

Share


Access


Process


Encrypt


…. and more!

We want to do these seamlessly...

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Using Diverse Interfaces & Devices

Desktops

Mobile Devices

Consumer Electronics

…and even appliances

We also want to access, share and process our data from all of our devices, **anytime, anywhere**!

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

What About the Future?

How will you…

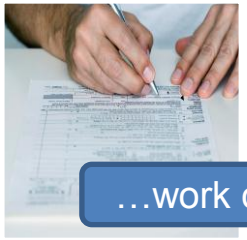

…work on documents?

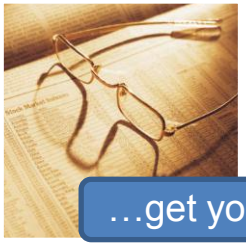

…get your news & info?

…create, access, store and share media?

…navigate?

…communicate with friends and family?

…live in an intelligent home?

…

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

…How Will We Manage Our Data?

Manage it ourselves?	• Personal, but time consuming.
How would you get access to your data wherever you are?	• Would you keep it on your devices? • or would you keep it online?
What if it's managed by someone else?	• and you can get this **"service"** for free or with a subscription?

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Has this Happened Before?

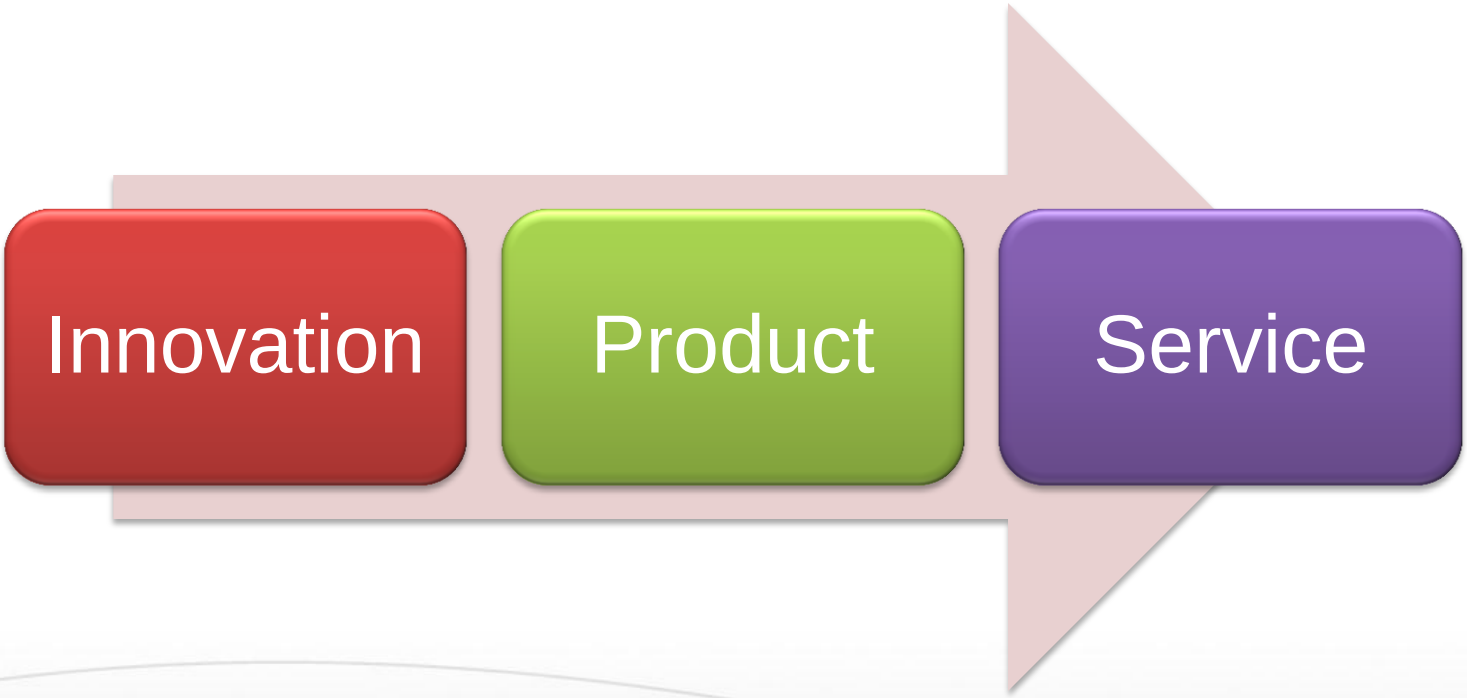

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Think of it this Way …

- Evolution of water Utility

Generate your own utility

Buy it as a product and manage it

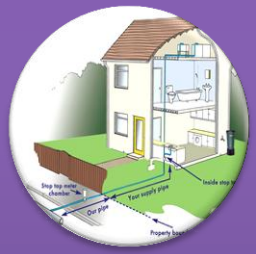

Get a continuous supply of the utility through a dedicated connection

How About Electricity?

- Transformation from a Product to a Service

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

…and Banking?

- Evolution of Banking

No Banks

(Take care of your own money)

Traditional Banking

(Give your money to the bank)

Banking Instruments

(Cheques / Credit Cards)

Internet Banking

(…more services)

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

So What is Cloud Computing?

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Can We Define Cloud Computing?

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

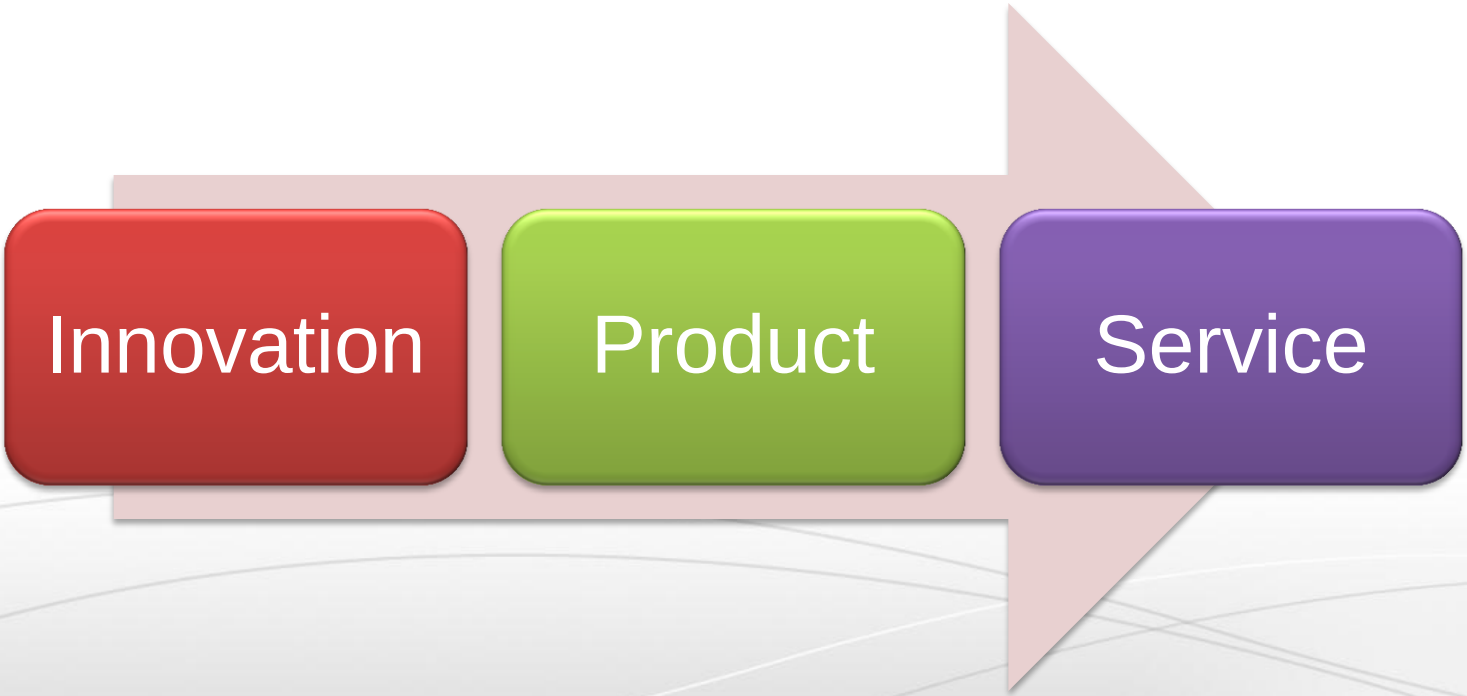

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Cloud Computing

- Transformation of IT from a Product to a Service

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

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

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Requirements to Transform IT to a Service

- Connectivity
  - For moving data around
- Interactivity
  - Seamless interfaces
- Reliability
  - Failure will affect many people, not just one
- Performance
  - It should not be slower or less efficient than what people already have
- Pay-as-you-Go
  - Should not pay an upfront fee for the service
- Ease of Programmability
  - Ease of development of complex services to users
- Manage Large Amounts of Data
  - Big Data
- Efficiency
  - Cost
  - Power
- Scalability & Elasticity
  - Flexible and rapid response to changing user needs

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Requirements to Transform IT to a Service

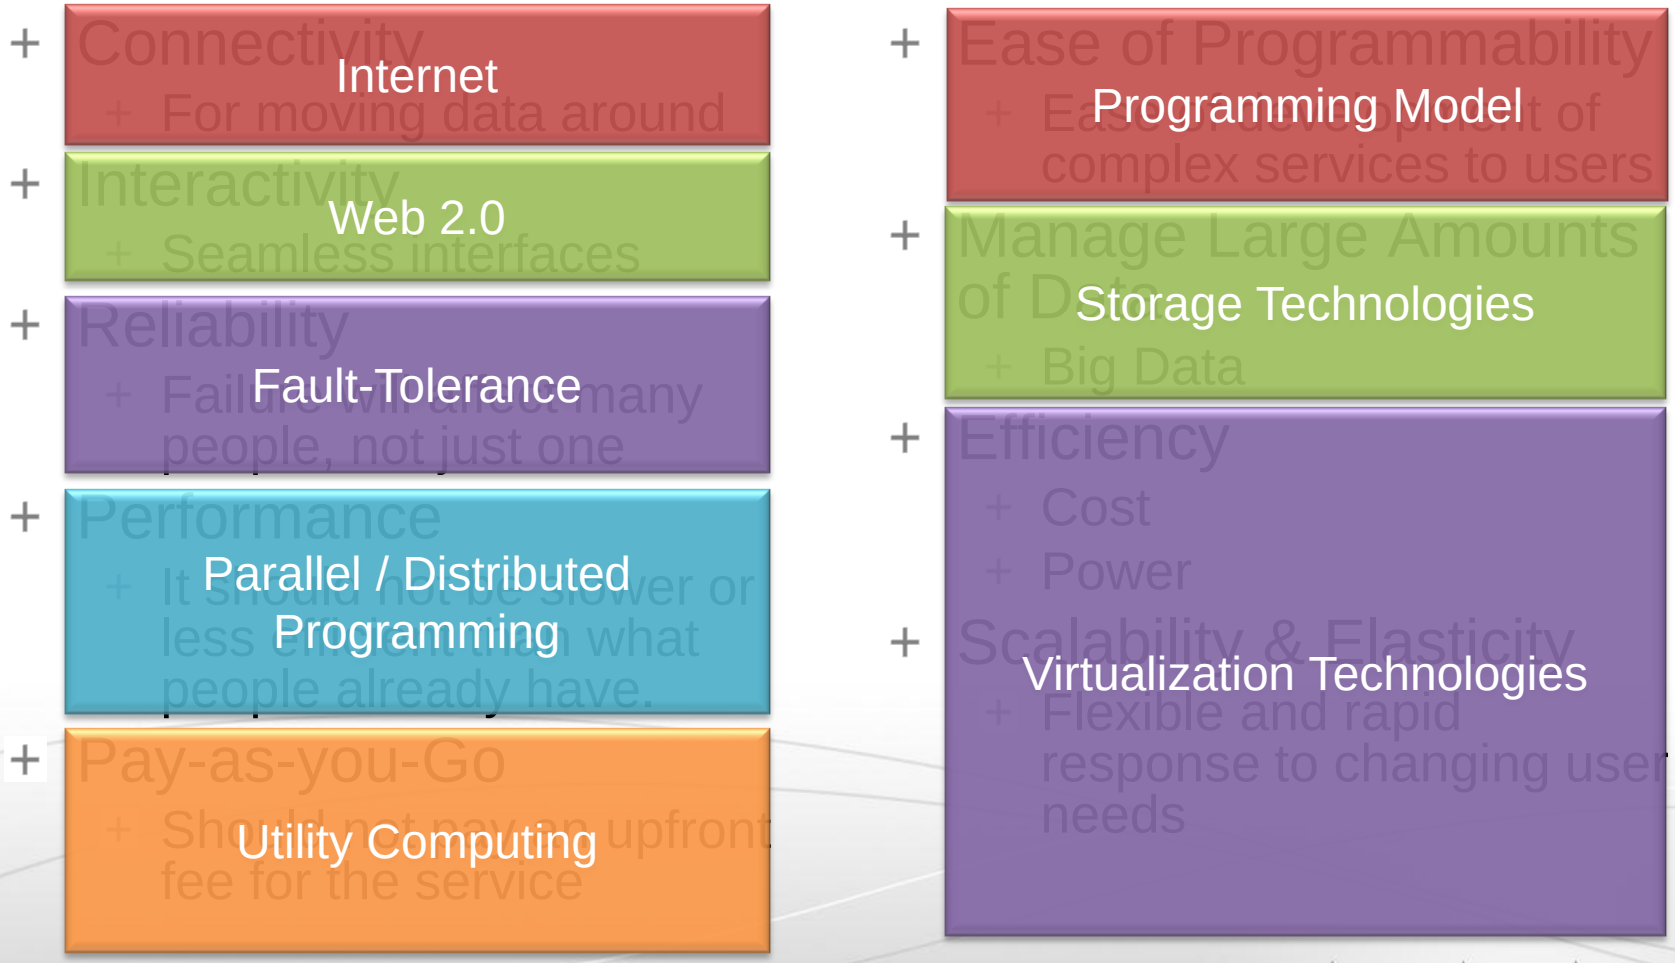

Combine the Enabling Technologies…

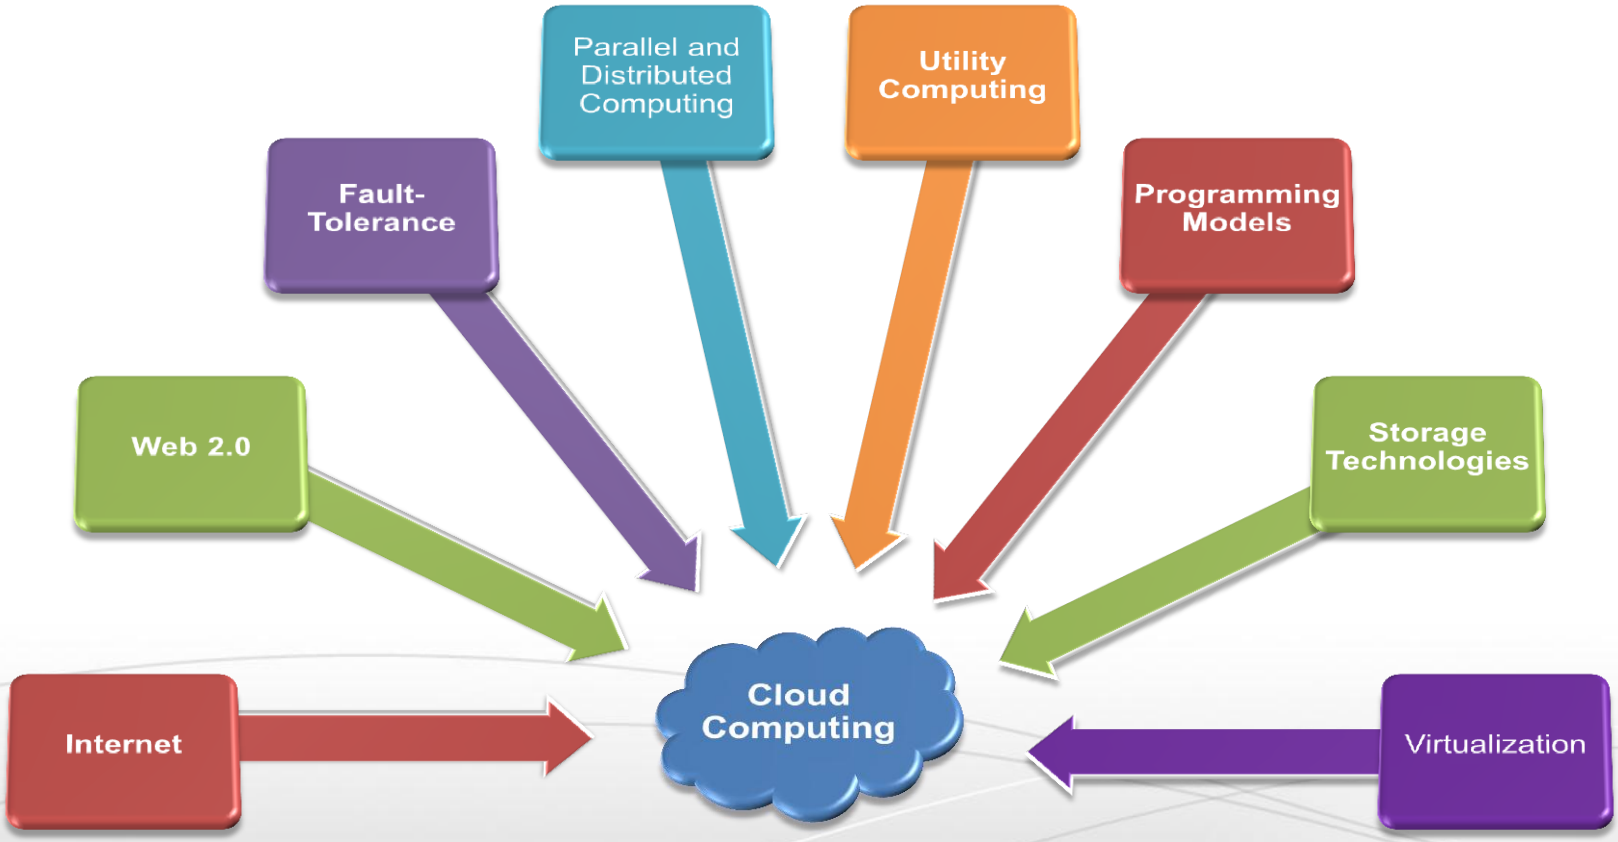

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Cloud Computing

- Think of it as Internet Computing
  - Computation done over the Internet

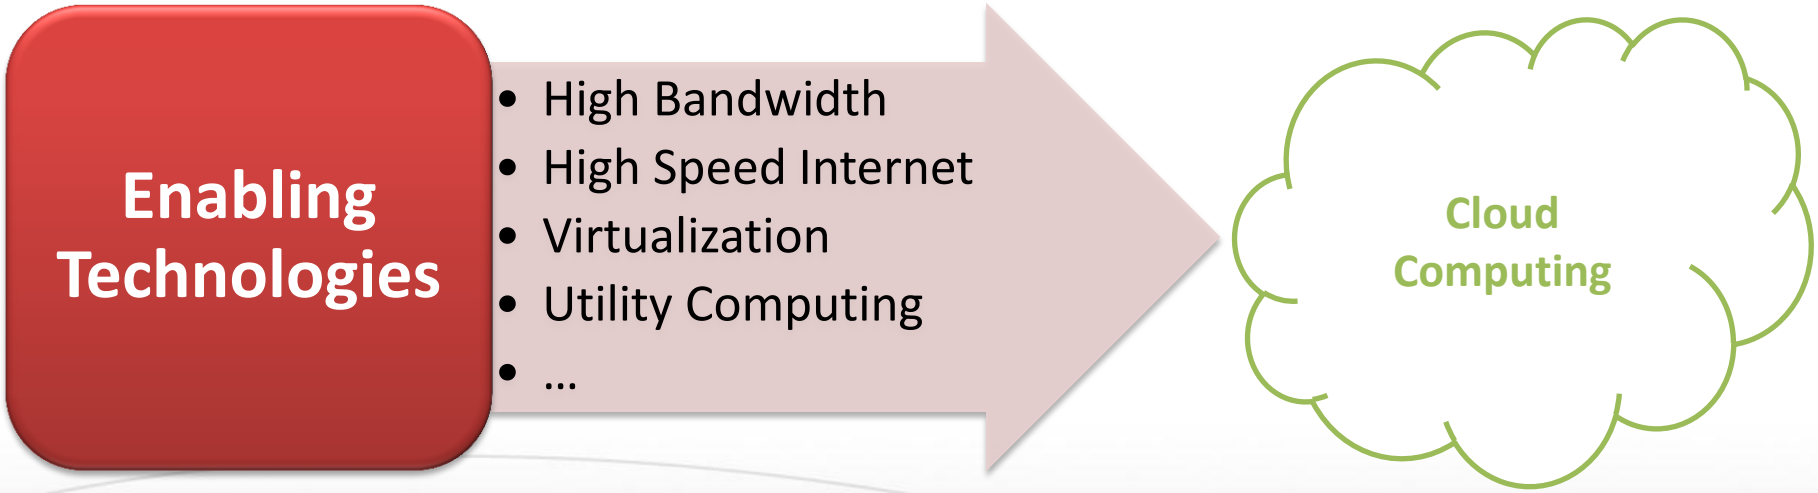

… for a more complete definition!

Cloud Computing is the delivery of computing as a **service** rather than a **product**,

whereby **shared resources, software**, and **information** are provided to computers and other devices,

as a **metered service** over a **network.**

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Why Cloud Computing?

Pay-as-You-Go economic model

- Reduce capital expenditure
- No upfront cost
- Reduced Time to Market

Simplified IT management

- All you need is access to the internet.
- It's the providers responsibility to manage the details.

Scale quickly and effortlessly

- Resources can be rented and released as required
- Software Controlled
- Instant scalablility

Flexible options

- Configure software packages, instance types operating systems.
- Any software platform
- Access from any machine connected to the Internet

Resource Utilization is improved

- Reduce Idle resources by sharing and conolidation
- Better utilization of CPU / Storage and Bandwidth.

Carbon Footprint decreased

- Sharing of resources means less servers, less power and less emissions.

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Applications Enabled by Cloud Computing

High Growth Applications

Startup Businesses

Aperiodic Bursting Applications

Seasonal Businesses

On-Off Applications

Research Computing

Periodic Applications

Changing computational patterns over time

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

High Growth Applications

2001

friendster®

Could not keep up with the growth of their number of users.

vs.

2006

facebook

Growing exponentially

What do you do when your startup gains traction?

Can you grow quick enough?

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

High Growth Applications

Users use it to produce video pieces from their photos, video clips and music.

Animoto's Facebook Plugin doubled traffic to the site every 12 hours for 3 days.

They could scale from 50 servers to 3500 and **go back down using cloud computing services**

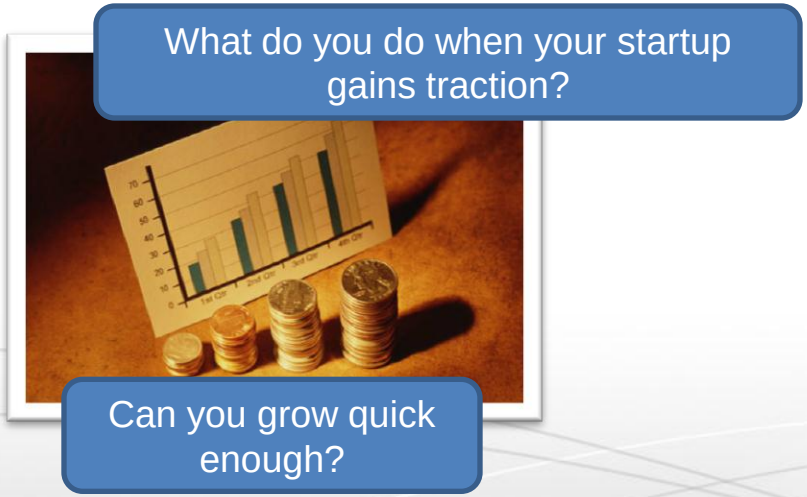

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Aperiodic Bursting Applications

Website went down on 9/11/2001 due to traffic

February 14th – Busiest Day of the Year

US Holiday Season

Website crashed within 10 minutes of the free trouser promotion during Superbowl 2010

Even if you design your website infrastructure to handle peak loads, wont it be idle during other times?

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

On-Off Applications

- Researchers running large-scale scientific simulation using 1000s of computers.

Modern Drug Discovery

- Data-intensive simulation and tests to discover new compounds
- Large compute power required for simulation jobs
- Time to market is crucial

Why not rent computer time to run these simulations?

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Periodic Applications

Sock Market Analysis

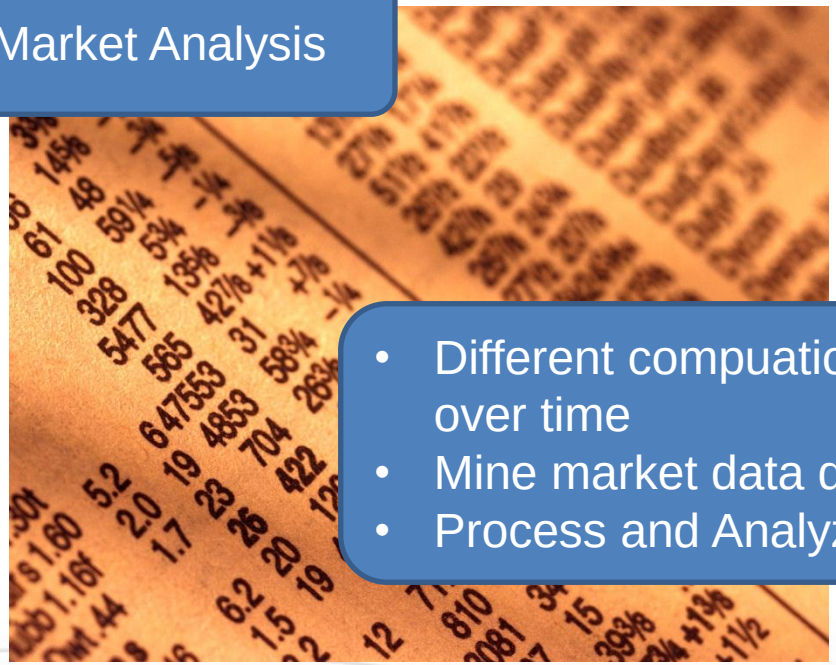

- Different compuational requirements over time
- Mine market data during the day.
- Process and Analyze at night.

Dynamic and Flexible infrastructure can reduce costs and improve performance.

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Technical Challenges

- Programming is tricky but improving
- Tools are continuously evolving
- Moving large data is still expensive
- Security
- Quality of Service
- Green computing
- Internet Dependence

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Non-Technical Challenges

- Vendor Lock-In
- Non-standardized
- Security Risks
- Privacy
- Legal
- Service Level Agreements

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

In 15-319

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Course Objectives

1. Learn the core concepts of cloud computing
2. Hands-on experience in using cloud computing infrastructure
3. Work on a large research project in cloud computing

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Learning Outcomes

- ***Considered:*** a reasonably critical and comprehensive perspective.
- ***Thoughtful:*** Fluent, flexible and efficient perspective.
- ***Masterful:*** a powerful and illuminating perspective.

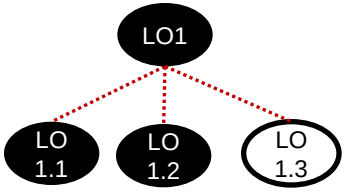

LO1	Explain, apply and identify core concepts and emerging issues in the cloud computing paradigm.
LO1.1	Explain how and why this paradigm came about and the influence of several enabling technologies including programming models (e.g., MapReduce), virtualization (e.g., Xen and Vmware), distributed file systems and cloud storage (e.g., HDFS and HBase), and emerging cloud tools (e.g., Hive)
LO1.2	Examine the process of working on a large research project under the mentorship of a teaching staff member.
LO1.3	Identify some of the emerging cloud research challenges including cloud security, QoS, SLAs, and energy-efficient clouds.

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Syllabus

- Contract!
- All you need to know about the course
  - Course Objectives
  - Learning Outcomes
  - Policies
  - Grading
  - Tentative Schedule

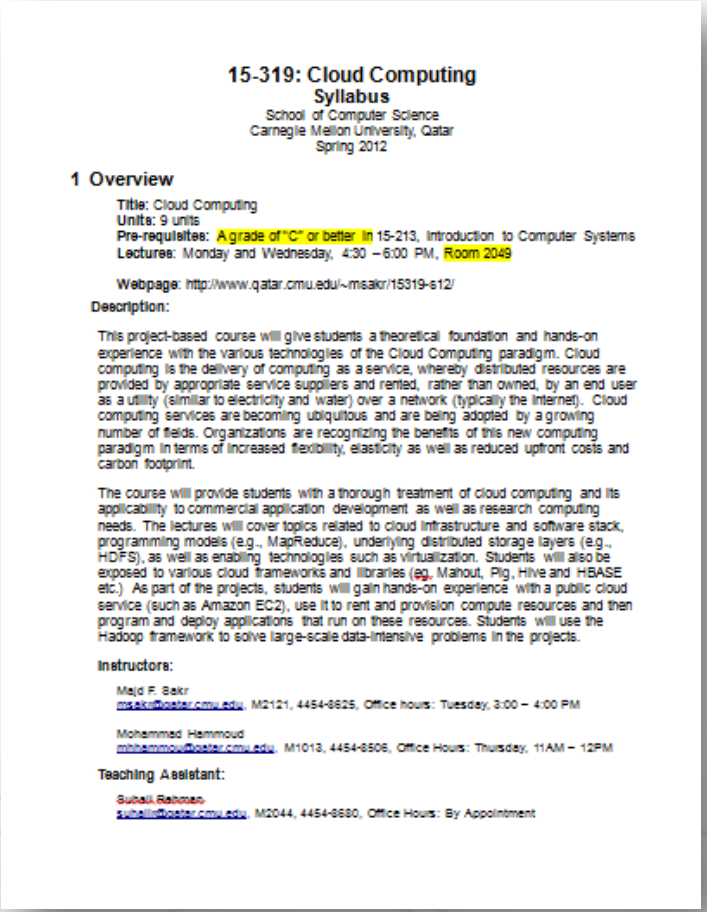

15-319: Cloud Computing
Syllabus
School of Computer Science
Carnegie Mellon University, Qatar
Spring 2012

1 Overview

Title: Cloud Computing
Units: 9 units
Pre-requisites: A grade of "C" or better in 15-213, Introduction to Computer Systems
Lectures: Monday and Wednesday, 4:30 – 6:00 PM, Room 2049

Webpage: http://www.qatar.cmu.edu/~msakr/15319-s12/

Description:

This project-based course will give students a theoretical foundation and hands-on experience with the various technologies of the Cloud Computing paradigm. Cloud computing is the delivery of computing as a service, whereby distributed resources are provided by appropriate service suppliers and rented, rather than owned, by an end user as a utility (similar to electricity and water) over a network (typically the Internet). Cloud computing services are becoming ubiquitous and are being adopted by a growing number of fields. Organizations are recognizing the benefits of this new computing paradigm in terms of increased flexibility, elasticity as well as reduced upfront costs and carbon footprint.

The course will provide students with a thorough treatment of cloud computing and its applicability to commercial application development as well as research computing needs. The lectures will cover topics related to cloud infrastructure and software stack, programming models (e.g., MapReduce), underlying distributed storage layers (e.g., HDFS), as well as enabling technologies such as virtualization. Students will also be exposed to various cloud frameworks and libraries (eg, Mahout, Pig, Hive and HBASE etc.) As part of the projects, students will gain hands-on experience with a public cloud service (such as Amazon EC2), use it to rent and provision compute resources and then program and deploy applications that run on these resources. Students will use the Hadoop framework to solve large-scale data-intensive problems in the projects.

Instructors:

Majd F. Sakr
msakr@qatar.cmu.edu, M2121, 4454-8625, Office hours: Tuesday, 3:00 – 4:00 PM

Mohammad Hammoud
mhhammou@qatar.cmu.edu, M1013, 4454-8506, Office Hours: Thursday, 11AM – 12PM

Teaching Assistant:

Suhail Rehman
suhailr@qatar.cmu.edu, M2044, 4454-8680, Office Hours: By Appointment

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Target Audience

- CS Majors
- Juniors / Seniors
- Pre-requisites:
  - 15-213 – Introduction to Computer Systems

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Instructors / Getting Help

Prof. Majd F. Sakr

msakr@qatar.cmu.edu
(2121)

Dr. Mohammad Hammoud

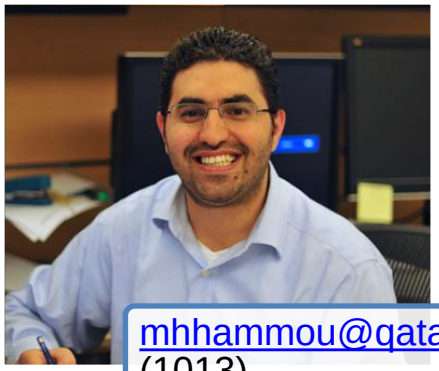

mhhammou@qatar.cmu.edu
(1013)

Dr. Fan Zhang

fanzhan1@qatar.cmu.edu
(1206)

Suhail Rehman

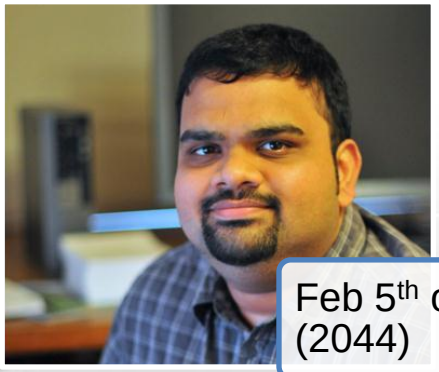

Feb 5th onwards
(2044)

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Course Organization

Lectures

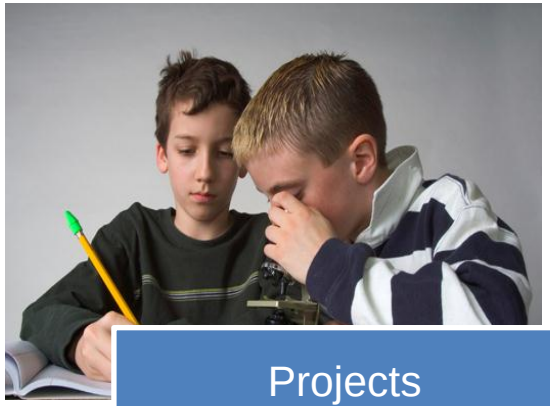

Projects

Project Status Reports and Presentations

In-Class Discussions

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Week	Session	Date	Topic	Teaching Method	Instructor	Projects
1	1	16 Jan	Administrivia and Course Motivation	Lecture	MFS	Project Start
	2	18 Jan	Introduction to Cloud Computing and Cloud Software Stack	Lecture	MFS	
2	3	23 Jan	Course Project and Amazon AWS	Lecture	MFS	
	4	25 Jan	Principles of Parallel Programming	Lecture	MFS	Project Update
3	5	30 Jan	Introduction to Programming Models	Lecture	MFS	
	6	1 Feb	Introduction to MapReduce	Lecture	MHH	Project Update
4	7	6 Feb	MapReduce Algorithms	Lecture	MHH	
	8	8 Feb	Apache Mahout	Guest Lecture	Shannon	Project Update
5	9	13 Feb	Pregel, Dryad, GraphLab	Lecture	MHH	
	10	15 Feb	Major Project Milestone Presentations	Project Presentation		Project Phase I End, Paper V1 Due, Project Phase II Start
6	11	20 Feb	Introduction to Cloud Storage and Distributed File Systems (DFSs)	Lecture	MFS	
	12	22 Feb	Cloud Storage and DFSs– HDFS, PVFS	Lecture	MHH	Project Update
7	13	27 Feb	Cloud Storage and DFSs- BigTable/HBASE	Lecture	MHH	
	14	29 Feb	Apache Pig, Hive, Zookeeper	Lecture		Project Update
8			**Spring Break – No Classes**			
	15	12 Mar	Introduction to Virtualization	Lecture	MHH	
9	16	14 Mar	Resource Virtualization (CPU/Memory/Disk/Network)	Lecture	MHH	Project Update
	17	19 Mar	Virtualization Case Study: Xen	Lecture	MHH	
10	18	21 Mar	Virtualization Case Study: VMware	Lecture	MHH	Project Update
	19	26 Mar	Benchmarking and Application Characterization	Lecture	MHH	
11	20	28 Mar	Major Project Milestone Presentations	Project Presentation		Project Phase II End Project Phase III Start
	21	2 Apr	Cloud Monitoring and Diagnostics	Lecture	MSR	
12	22	4 Apr	Cloud Security	Lecture	Thierry	Project Update
	23	9 Apr	Cloud QoS and SLAs	Lecture	MFS	Poster Due
13	24	11 Apr	Green Cloud	Lecture	MFS	Project Update
	25	16 Apr	How to Write a Scientific Paper	Lecture		
14	26	18 Apr	Guest Lecture	Lecture		Project Update
	27	23 Apr	Final Project Presentations I	Project Presentation		Project End
15	28	25 Apr	Final Project Presentations II	Project Presentation		Project End

Primary Textbook

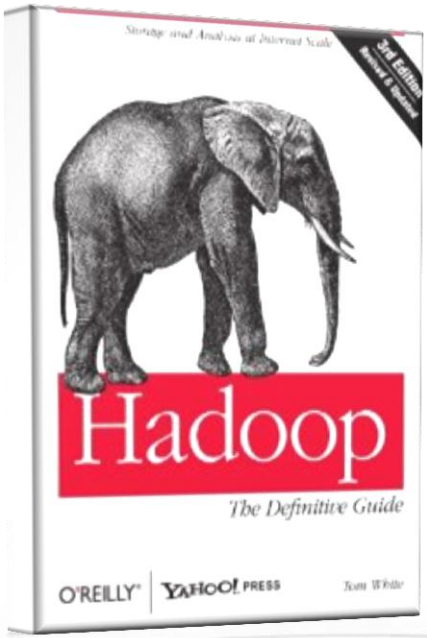

Hadoop: The Definitive Guide (2nd Edition)

Tom White, O'Reilly Media, 2010

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Reference Books

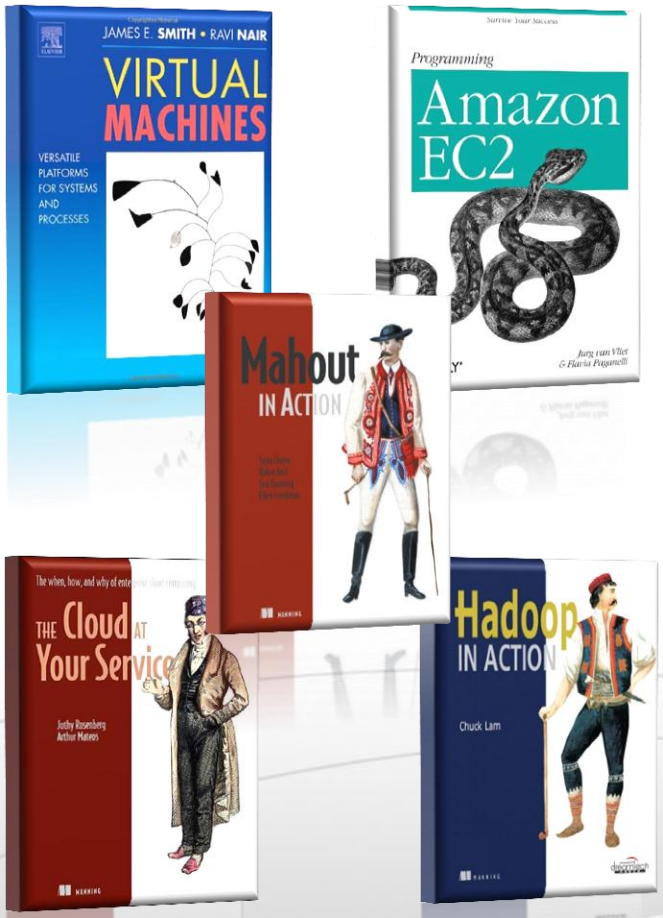

- **Virtual Machines : *Versatile Platforms for Systems and Processes***
  James E. Smith and Ravi Nair, Morgan Kauffman, 2005
- **Programming Amazon EC2**
  Jurg van Vilet and Flavia Paganelli, O'Reilly Media, 2011
- **The Cloud at Your Service**
  Jothy Rosenberg and Arthur Mateos,, Manning Publications, 2010
- **Mahout in Action**
  Sean Owen, Robin Anil, Ted Dunning and Ellen Friedman, Manning Publications, 2011
- **Hadoop in Action**
  Chuck Lam, Manning Publications, 2011

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Projects

- This is a project-based Course
- Select a domain and a mentor
  - NLP, Image Processing, DNA Sequencing
  - Dr. Hammoud, Mr. Rehman, Dr. Zhang
- Three-phase project timeline.
  - Phase 1 – Introduction, Ramp-up, Timeline
  - Phase 2 – Application Implementation
  - Phase 3 – Characterization and Analysis
- You will be using a Public Cloud Service
  - Amazon EC2

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Overall Project Timeline

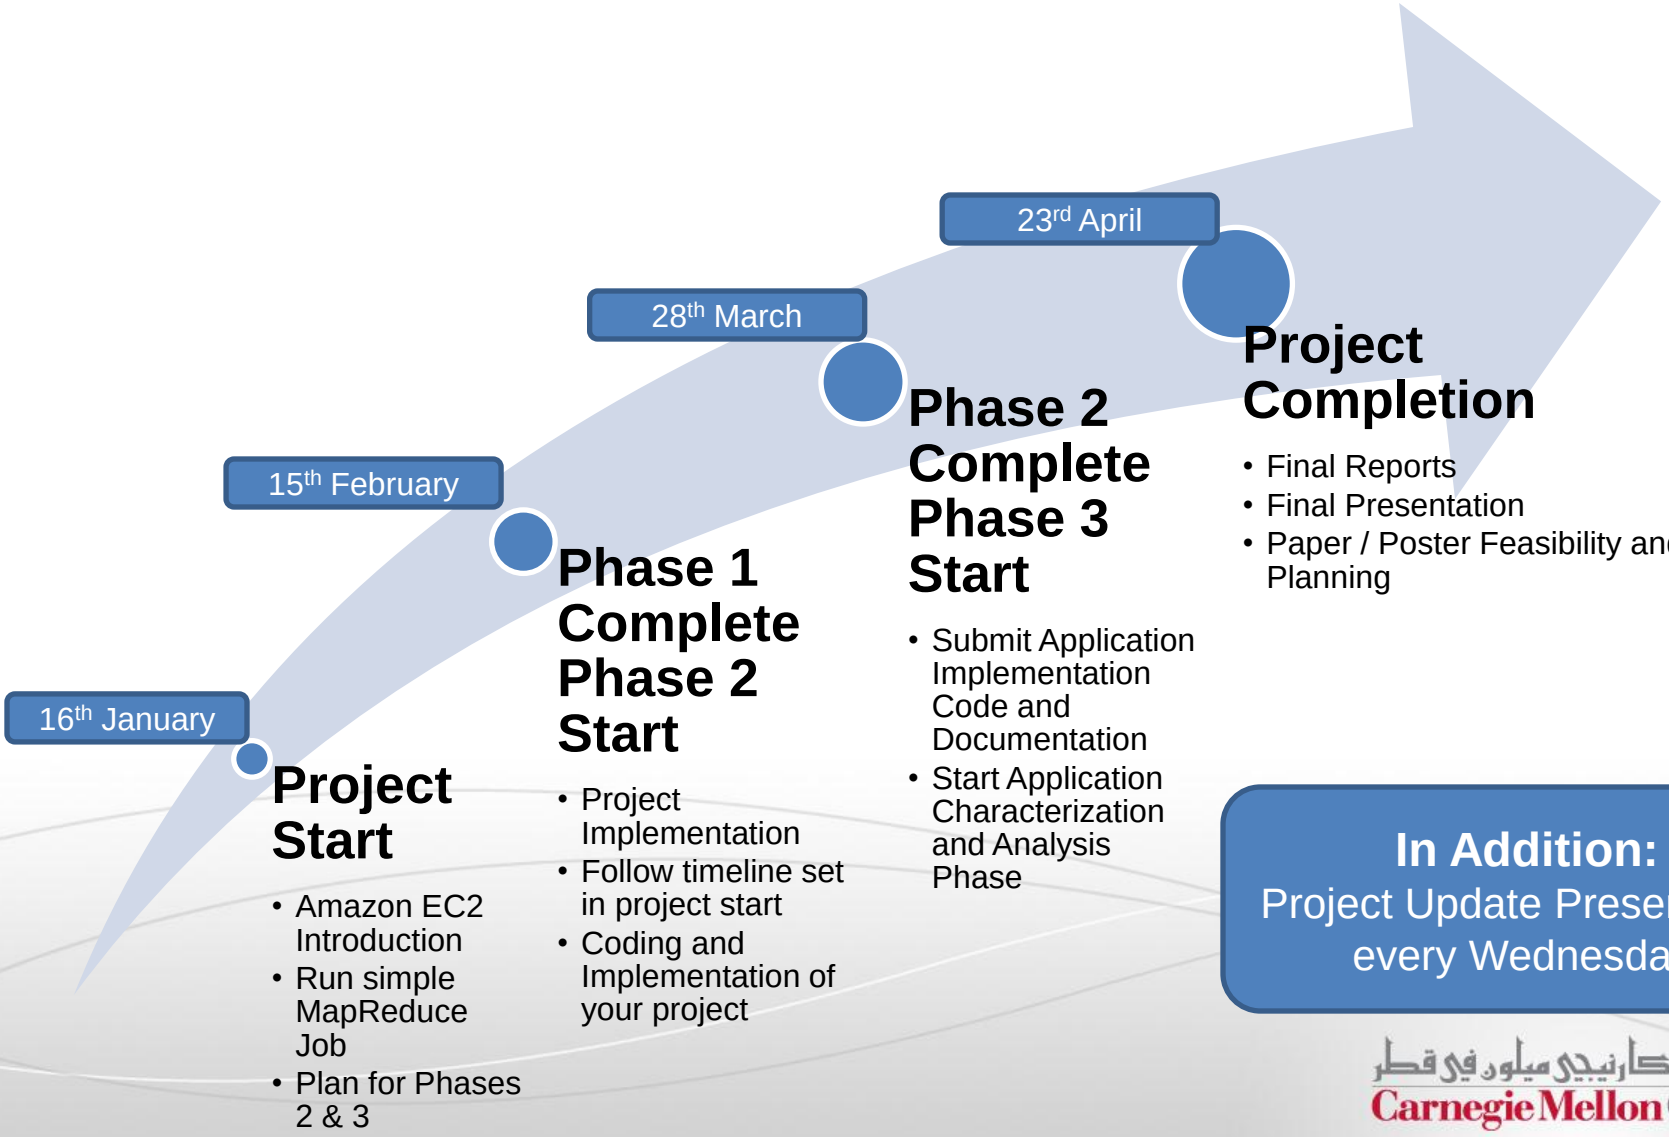

In Addition:
Project Update Presentation every Wednesday.

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Special Note on Amazon EC2

- Paid Cloud Service – you are billed by the hour.
- Start a resource only when you need them.
- Terminate resources as soon as you are done with them.

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Grading

Course Elements	#	Weight
Project Phases I, II & III	3	75%
Project Update Presentations	6	20%
Class Participation and Attendance	28	5%

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar

Questions?

جامعة كارنيجي ميلون في قطر
Carnegie Mellon Qatar