Cloud Computing

ECPE 293A

Introduction

Introductions

- **Introduce yourself**
  - Name
  - Undergrad major
  - Interests
    - Hardware – which area?
    - Software – which area?

Cloud Computing

=

Looking at pictures of
clouds with **computers**

Cloud Computing

≠

Looking at pictures of
clouds with **computers**

Cloud Computing

- **What do you know about cloud computing?**

- **What do you want to know about cloud computing?**

Cloud = The Internet

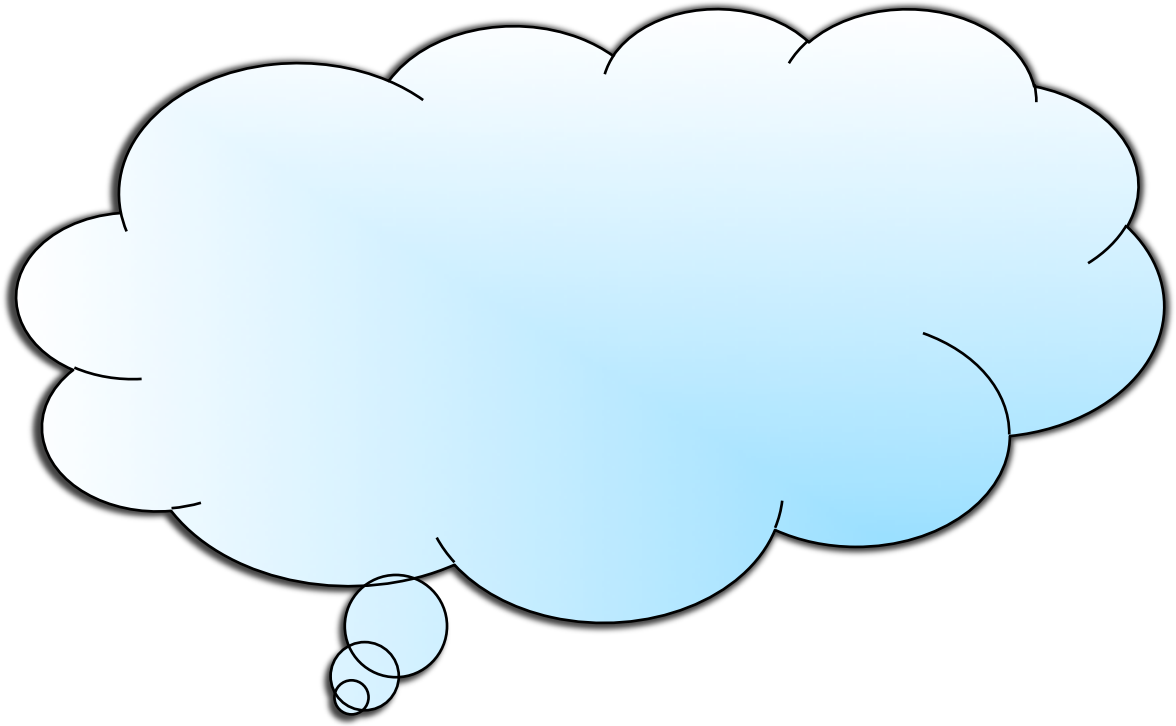

Cloud Computing

- Cloud computing is a computing paradigm where you outsource your data, computation, or both
- Features
  - Resource **scalability** (“infinite”!)
  - On demand / **“just-in-time”** provisioning
  - No upfront cost … **pay-as-you-go**

Cloud Computing

- Economic benefits for customers:
  - No upfront commitment in buying/leasing hardware
  - Can scale usage according to demand
  - Barriers to entry lowered for startups
- Economic benefits for providers:
  - Increased utilization of datacenter resources

Cloud Computing

- Selling "X as a Service"
  - **IaaS**: Infrastructure as a Service
    - Selling virtualized hardware
  - **PaaS**: Platform as a service
    - Selling access to a configurable platform/API
  - **SaaS**: Software as a service
    - Selling software that a user directly interacts with

Cloud Computing Architecture

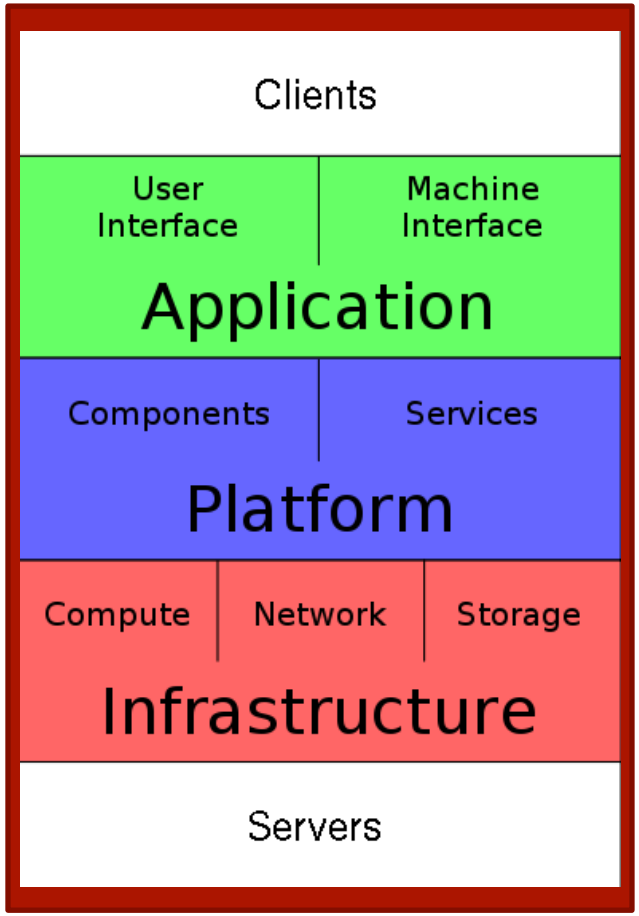

e.g., Web browser

SaaS , e.g., Google Docs

PaaS, e.g., Google AppEngine

IaaS, e.g., Amazon EC2

Course Mechanics

Course Goals

1. Understand the architecture *behind* modern cloud computing infrastructure
   - *"How is the infrastructure built?"*
   - Accomplished by: **reading/discussing research papers**
2. Gain **hands-on experience** with modern cloud computing systems
   - *"How do I use the infrastructure provided?"*
   - Accomplished by: **programming projects**

Course Expectations

- This is not an undergraduate class
- I have higher expectations for you
- **Independence** – You get to select research papers, project topics, and implementation strategies
  - Take advantage of this **flexibility** to focus the course in a specific area that interests you!

Research Papers – Discussion Leader

- During this semester, you will present 1 research paper to the class
- Job of discussion leader
  - Pick a paper (and email me for approval / posting)
  - Read it thoroughly (twice!)
  - Prepare a 20-minute PowerPoint presentation
    - *As if you had done this work and were presenting it at a conference*
    - Do not use existing work on the web! (OK to re-use figures)
  - Lead a group discussion for the remaining class period

Research Papers – Audience Members

- When you are not presenting a paper, you will be an audience member
- Job of audience members
  - ~~Skim the paper 5 minutes before the start of class~~
  - If only 1 person has read the paper, class discussions will be unproductive (and *agonizing*!)

Research Papers – Audience Members

- Read the paper thoroughly
  - **2+ hours** for a good understanding
    - *I'm a speed reader, but these technical papers are **dense!***
    - A little faster as you get up to speed in the field?
- Bring to class (for credit) a 1-2 page document
  - Summary of paper
  - Two strengths
  - Two weaknesses
  - Three questions for the presenter or other audience members

Projects

- Two programming projects
- **Project 1** – Large scale data processing
  - Dataset: CommonCrawl
  - Infrastructure: MapReduce
  - Application: Up to you!
- **Project 2** – Reliability / scalability in the cloud

Honor Code

- You are welcome to talk with your classmates about your projects
  - *Everyone should have a different project focus, anyway…*
- You are welcome use resources online (source code, tutorials, mailing lists, etc…)
  - Must **document and disclose** all of your sources in your project report

Grading

- **Exams** – None!
- **Quizzes** – None!
- **Paper Discussions as *Leader*** – 10%
- **Paper Discussions as *Audience Member*** – 25%
- **Programing Projects (2)** – 65%

Computers

- **Who has a laptop they can bring to class?**
- View research paper PDFs instead of printing them out
- Do in-class tutorials (or do we need a lab with computers?)

Upcoming Schedule – This Week

- Thur, Jan 16th – What is Cloud Computing?
  - **Your Homework 1:** Read Berkeley report
    - *This is not a technical paper and should be a fast read. Just skim it…*
  - **No summary required** for this paper (only!)
- Tue, Jan 21st – What is Cloud Computing?
  - Continuation of Thursday's discussion
  - **Your Homework 2:** Pick 3 papers from the approved reading list that you could present and **email me**

Upcoming Schedule – Next Week

- Thur, Jan 23rd – First paper presentation
  - Presenter: Dr. Shafer *(use an an example)*
  - MapReduce paper *(used for your first project)*
  - **Your Homework 3:** Audience members role
    - Read paper and prepare summary document
- …
- Feb 4th – First student paper presentation
  - **Who volunteers to get a paper out of the way early before the semester gets busy?**

Upcoming Schedule

- The question that you should always be asking yourself is: “What paper(s) should I read **before** the next class?”
  - Check schedule on website

Websites

Main website
(syllabus, schedule, papers, …)

- http://ecs-network.serv.pacific.edu/ecpe-293a

Sakai website
(projects, gradebook)

- http://pacific.rsmart.com/

Website Tour

- Where to find:
  - List of papers to select from
  - Requirements for presenter
  - Requirements for audience members
  - Grading rubrics