

CS 6464:
Advanced Distributed Storage Systems

Spring 2009

Instructor: Hakim Weatherspoon

Who am I?



- Prof. Hakim Weatherspoon
 - (Hakim means Doctor, wise, or prof. in Arabic)
 - Background in Education
 - Undergraduate University of Washington
 - Played Varsity Football
 - » Some teammates collectively make \$100's of millions
 - » I teach!!!
 - Graduate University of California, Berkeley
 - Some class mates collectively make \$100's of millions
 - I teach!!!
 - Background in Operating Systems
 - Peer-to-Peer Storage
 - Antiquity project - Secure wide-area distributed system
 - OceanStore project – Store your data for 1000 years
 - Network overlays
 - Bamboo and Tapestry – Find your data around globe
 - Tiny OS
 - Early adopter in 1999, but ultimately chose P2P direction

Goals for Today

- Be brief!
- Why take this course?
- Goals of distributed storage systems
- Intro to cloud computing
- How does this class operate?

Why take this course?

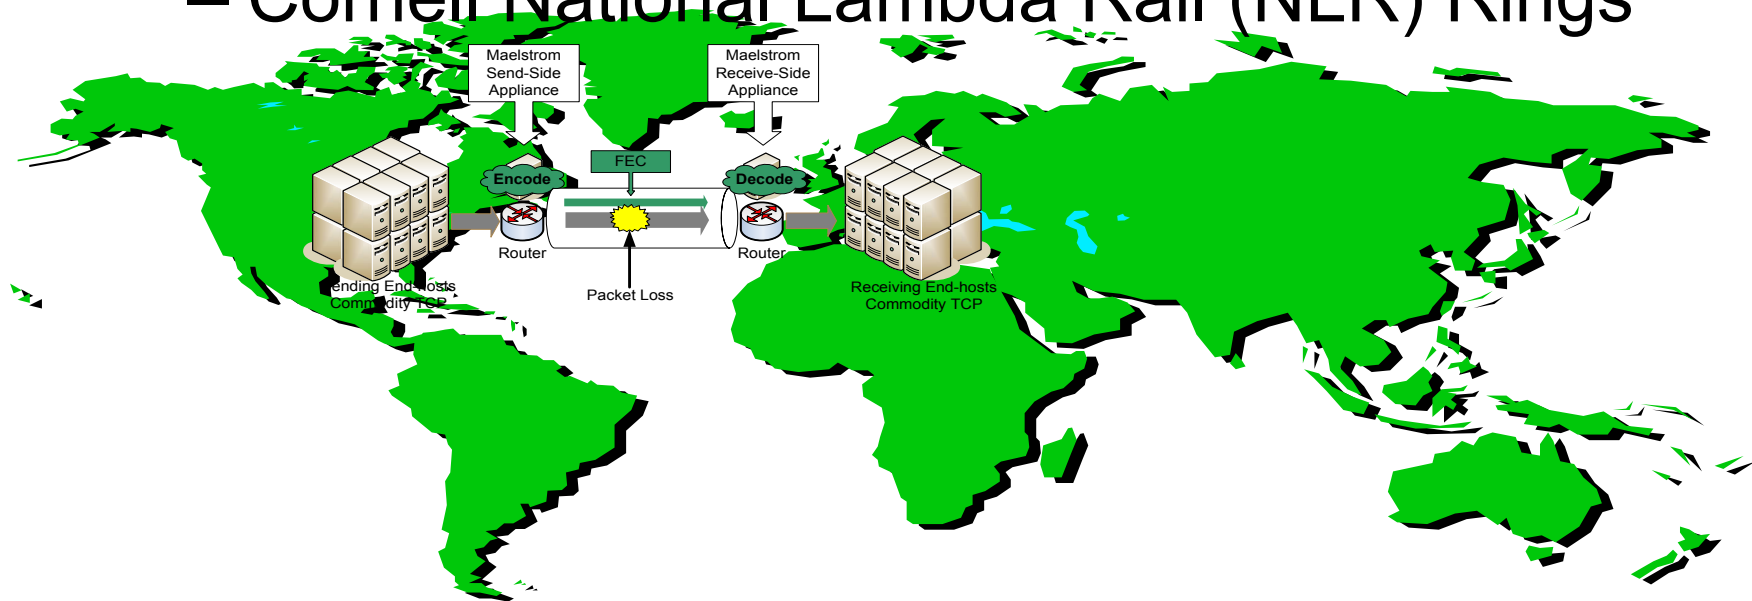


Why take this course?

The Google logo, featuring the word "Google" in its signature multi-colored font (blue, red, yellow, green, red) with a trademark symbol (TM) to the upper right.The YouTube logo, consisting of the word "You" in black and "Tube" in white inside a red rounded rectangle, all enclosed within a thin black L-shaped border.The Yahoo! Mail logo, with "YAHOO!" in red and "MAIL" in black, both in a bold, sans-serif font.The Flickr logo, with the word "flickr" in a blue, lowercase, sans-serif font, where the "r" is pink, set against a light gray rectangular background.

Why take this course?

- We have access to state-of-the-art equipment
 - Cornell National Lambda Rail (NLR) Rings



Why take this course?

- We have access to state-of-the-art infrastructure



Distributed Storage Systems


Distributed Storage System Goals

- Problem: No single server is capable of
 - Holding hundreds of terabytes of data
 - Serving millions of requests per second
- Goals
 - Stable performance – load balance
 - Scalable performance
 - High availability
 - Durability
 - security

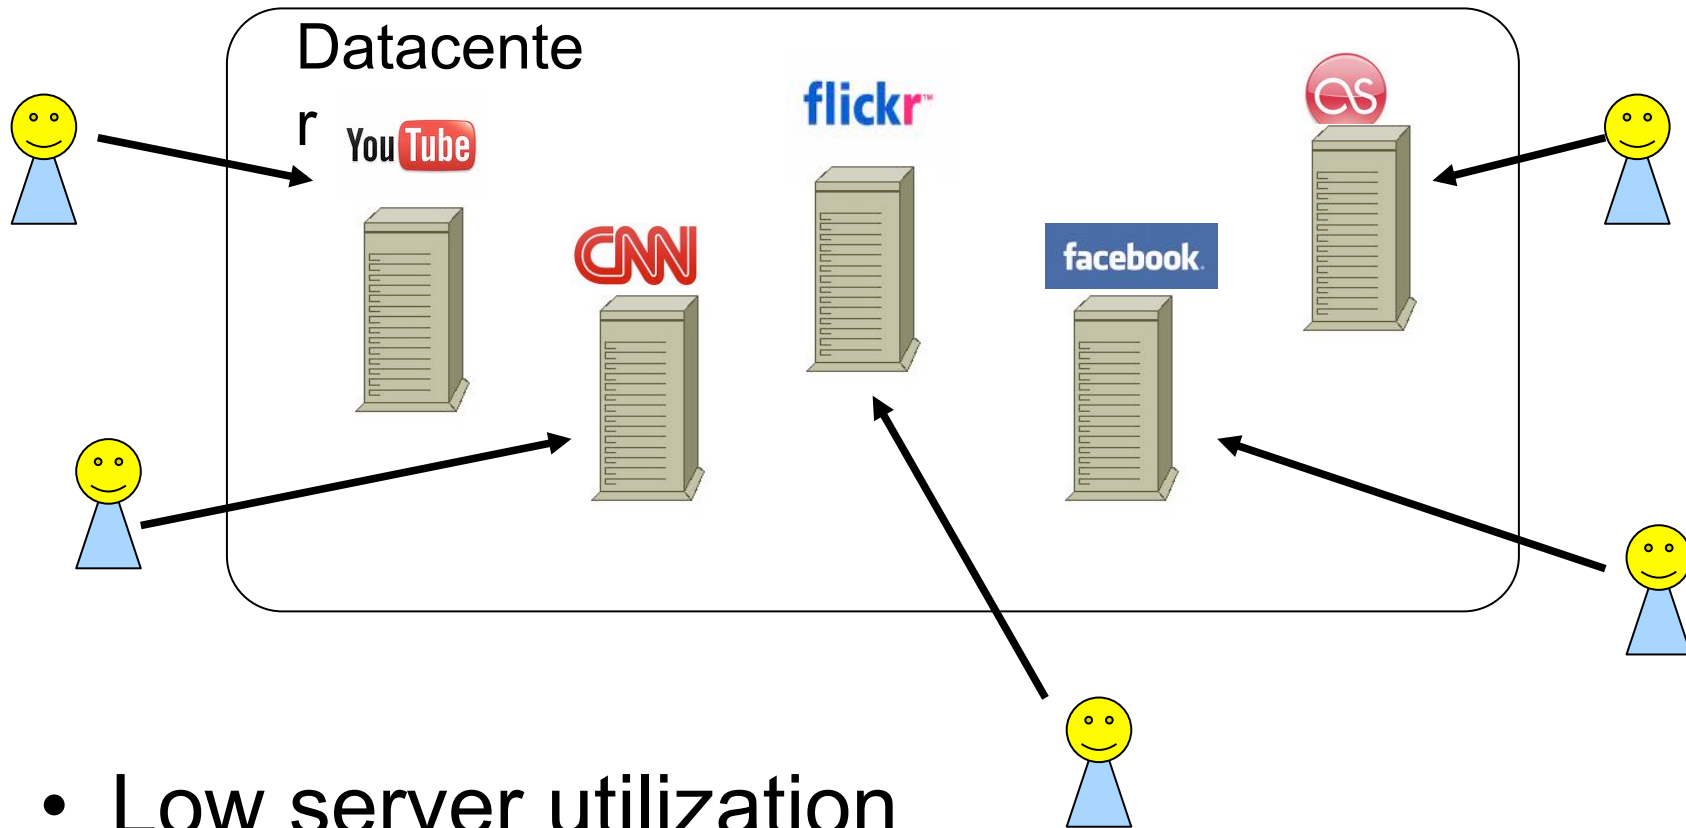
Cloud Computing

Cloud Computing Intro

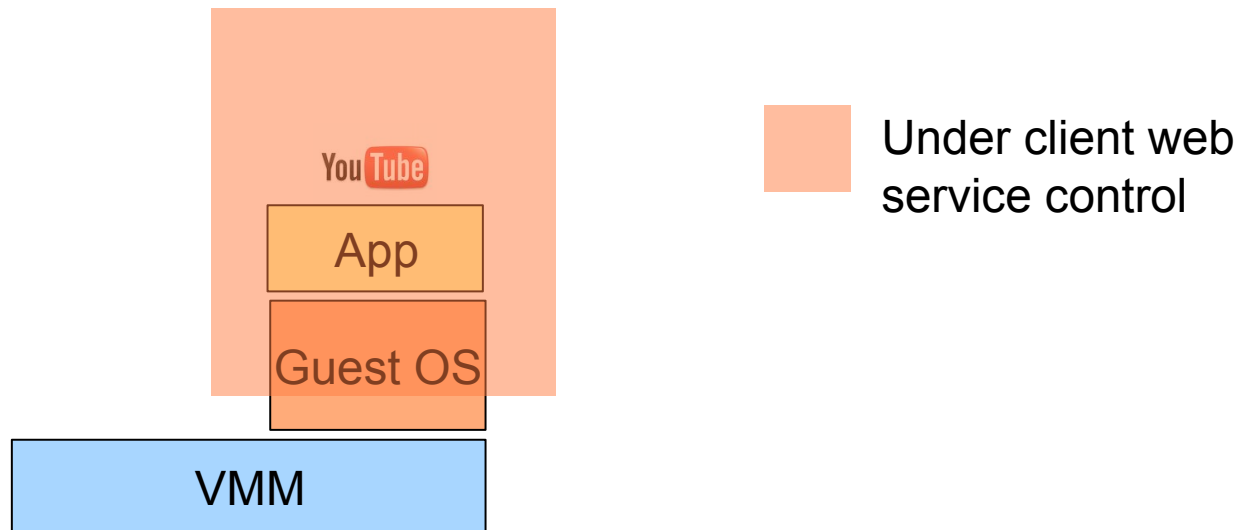
- Web service expensive to deploy
 - Provision for peak load



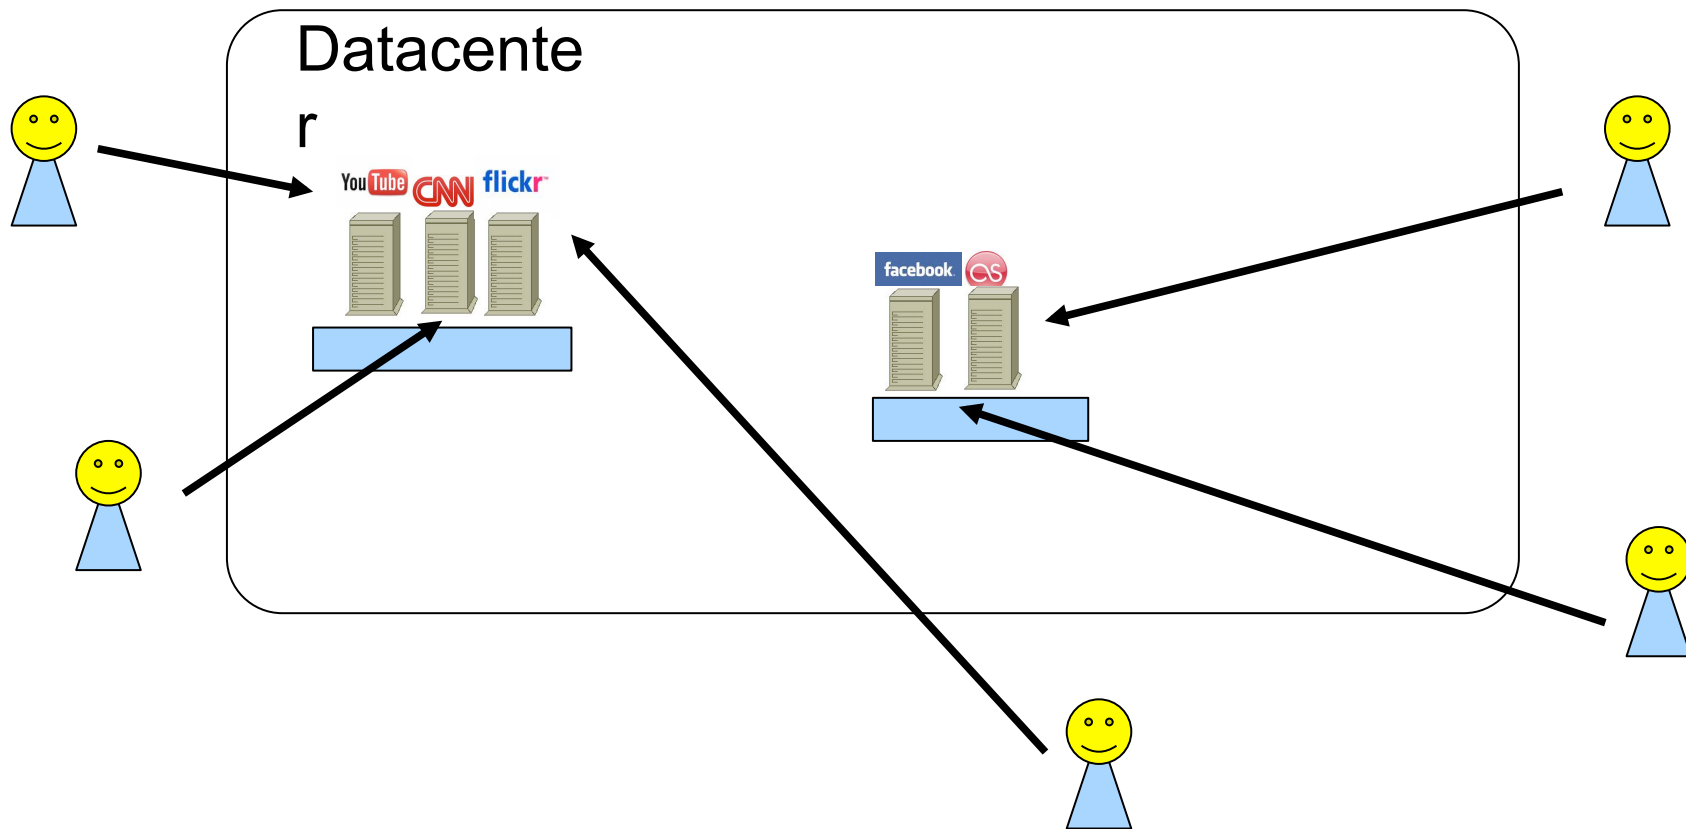
Cloud Provider Leases Resources



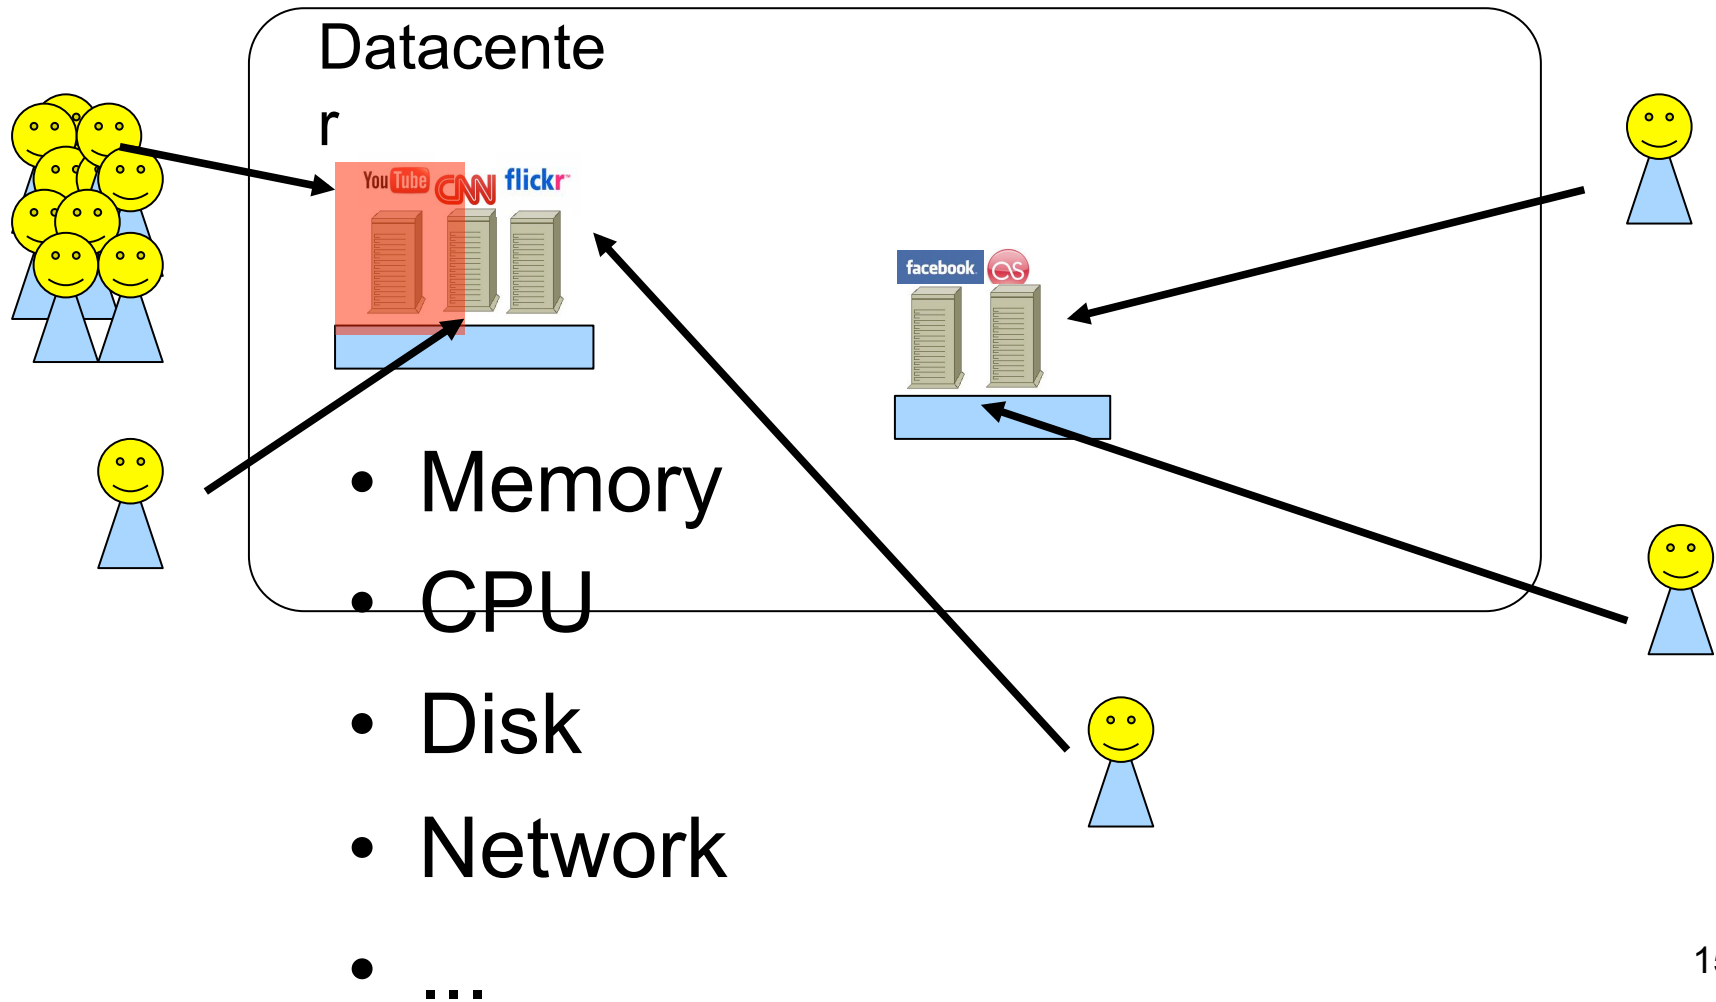
A Service Running in the Cloud



Server Consolidation Through Virtualization



Oversubscription



Observations

- Memory is the local bottleneck resource
 - Swap solves this problem using disk
 - Disk I/O is slow
- Large datacenter environment
 - Managed
 - Guests run standard VMs
 - Very fast network
- Few guests overloaded at one time
 - Memory is not a global bottleneck resource

How class operates

How this class operates

- Instructor: Hakim Weatherspoon
 - hweather@cs.cornell.edu
 - Office Location: 4105C Upson
- TA: Tudor Marian
 - tudorm@cs.cornell.edu
- Lectures:
 - CS 6464: Tu, Th: 10:10 – 11:25 PM, Hollister 362

Course Help

- Course staff, office hours, etc:
 - <http://www.cs.cornell.edu/courses/cs6464/2009sp>

CS 6464: Overview

- Prerequisite:
 - Mastery of CS 4410 material
 - Fundamentals of OS design
 - How parts of the OS are structured
 - What algorithms are commonly used
 - What are the mechanisms and policies used
- Class Structure
 - Papers Readings (whole semester)
 - Paper Presentations (whole semester)
 - Labs (first 1/3)
 - Research Project (second 2/3)

CS 6464: Class Structure

- Paper readings
 - Read paper before each class
 - Write a review and turn in at beginning of class
 - Review has three components
 - Summary, 2 – 3 strengths, and 2 – 3 weaknesses
- Paper presentations
 - Each person will present a paper 1-3 times, depending on class size
 - Have presentation prepared a week ahead of time and show slides or “chalk talk” to professor

CS 6464: Class Structure

- Labs (first 1/3 of semester)
 - 2 – 3 labs
 - Using Amazons EC2/S3 infrastructure
 - Building your own distributed file system
- Research Project
 - Initial project proposal – due beginning of March
 - Midterm survey paper – due beginning of April
 - Final demo/presentation – due beginning of May
 - Final project report – due sbeginning of May

Grading

- CS 6464: Operating Systems
 - Midterm ~ 30%
 - Final ~ 50%
 - Assignments ~ 10%
 - Subjective ~ 10%
 - Regrade policy
 - Submit written request to lead TA. TA will pick a different grader
 - Submit another written request, lead TA will regrade directly.
 - Submit another written request for professor to regrade.
- CS 4411: Systems Programming
 - Six projects ~ 100%
- This is a rough guide

Academic Integrity

- Submitted work should be your own
- Acceptable collaboration:
 - Clarify problem, C syntax doubts, debugging strategy
- Dishonesty has no place in any community
 - May NOT be in possession of someone else's homework/project
 - May NOT copy code from another group
 - May NOT copy, collaborate or share homework/assignments
 - University Academic Integrity rules are the general guidelines
- Penalty can be as severe as an 'F' in CS 6464

Next time

- Read *Cumulus* and write review
 - *Cumulus: Filesystem Backup to the Cloud*, Michael Vrable, Stefan Savage and Geoffrey M. Voelker
 - http://www-cse.ucsd.edu/Dienst/UI/2.0/Describe/ncstrl.ucsd_cse/CS2008-0927
- Lab 0
 - Using Amazon's EC2/S3 infrastructure
- Check website for updated schedule