CSE 704 Data Center Computing Intro

Steve Ko

Administrative Information

- Organizer: Steve Ko
- Ph.D., 2009, UIUC
- Interest: distributed systems, networking, and operating systems
- Office: 210 Bell
- Office hours: Wed 12pm – 3pm
- Email: stevko@buffalo.edu

Seminar Overview

- Data Center Computing
  - Computing activities that utilize data centers
  - An attempt to look at the whole spectrum
  - "What technologies do you use when you access a Web service?"
- Components
  - Front-end, processing, storage, networking, and virtualization

Seminar Credits

- 1 Credit
  - Reading papers
  - Writing reviews
  - Presenting one paper (or two)
  - Participating in discussions
- 3 Credits
  - Additional research project
  - Cannot be used for your master's project
  - Meet me after class

Reading Papers

- 24 papers, 2 papers per week
- Recommended reading: "How to Read a Paper" by S. Keshav (only 2 pages!)
- “A three-pass approach”

Reading Papers

- First-pass
  - Read the title, abstract, and intro
  - Read the titles of all sections and subsections
  - Read the conclusion
  - Goal: five C's
    - Category: which category is this paper in?
    - Context: related papers?
    - Correctness (on the assumptions)
    - Contributions
    - Clarity

Reading Papers

- Second-pass
  - Read with greater care, but still ignore details (e.g., proofs)
  - Jot down key points, ideas, background readings, etc.
- Third-pass
  - Virtually re-create the paper
  - Put down the paper, start from the same assumptions, and re-create the work.

Writing Reviews

- Format
  - What is the research problem?
    - Do you agree that it's a problem? Why or why not?
  - What are the main approaches/ideas? Strengths and weaknesses?
  - Other comments
    - E.g., what you liked, what you didn't understand, possible future work, compare & contrast with other papers, hidden assumptions, etc.
- Don't need to be long

Presentation

- Prepare PowerPoint slides
- Lead a discussion for an hour (presentation + questions/answers/discussions)
- Incorporate others' comments (e.g., what people didn't understand well)
- Schedule a time with me for a practice run
- You can use other people's slides
  - Make sure you acknowledge them

Assignment for Today

- Email me whether or not you're going to stay
- If you decide to stay, do the rest
- Look through the schedule
- Pick two papers you'd like to present
  - 1st choice & 2nd choice
  - FCFS
  - Exception: OpenFlow & NOX (from "Networking – 1" on 10/6) should be presented together with some demo
- Email me your choices

Topics Overview

- What makes these possible?

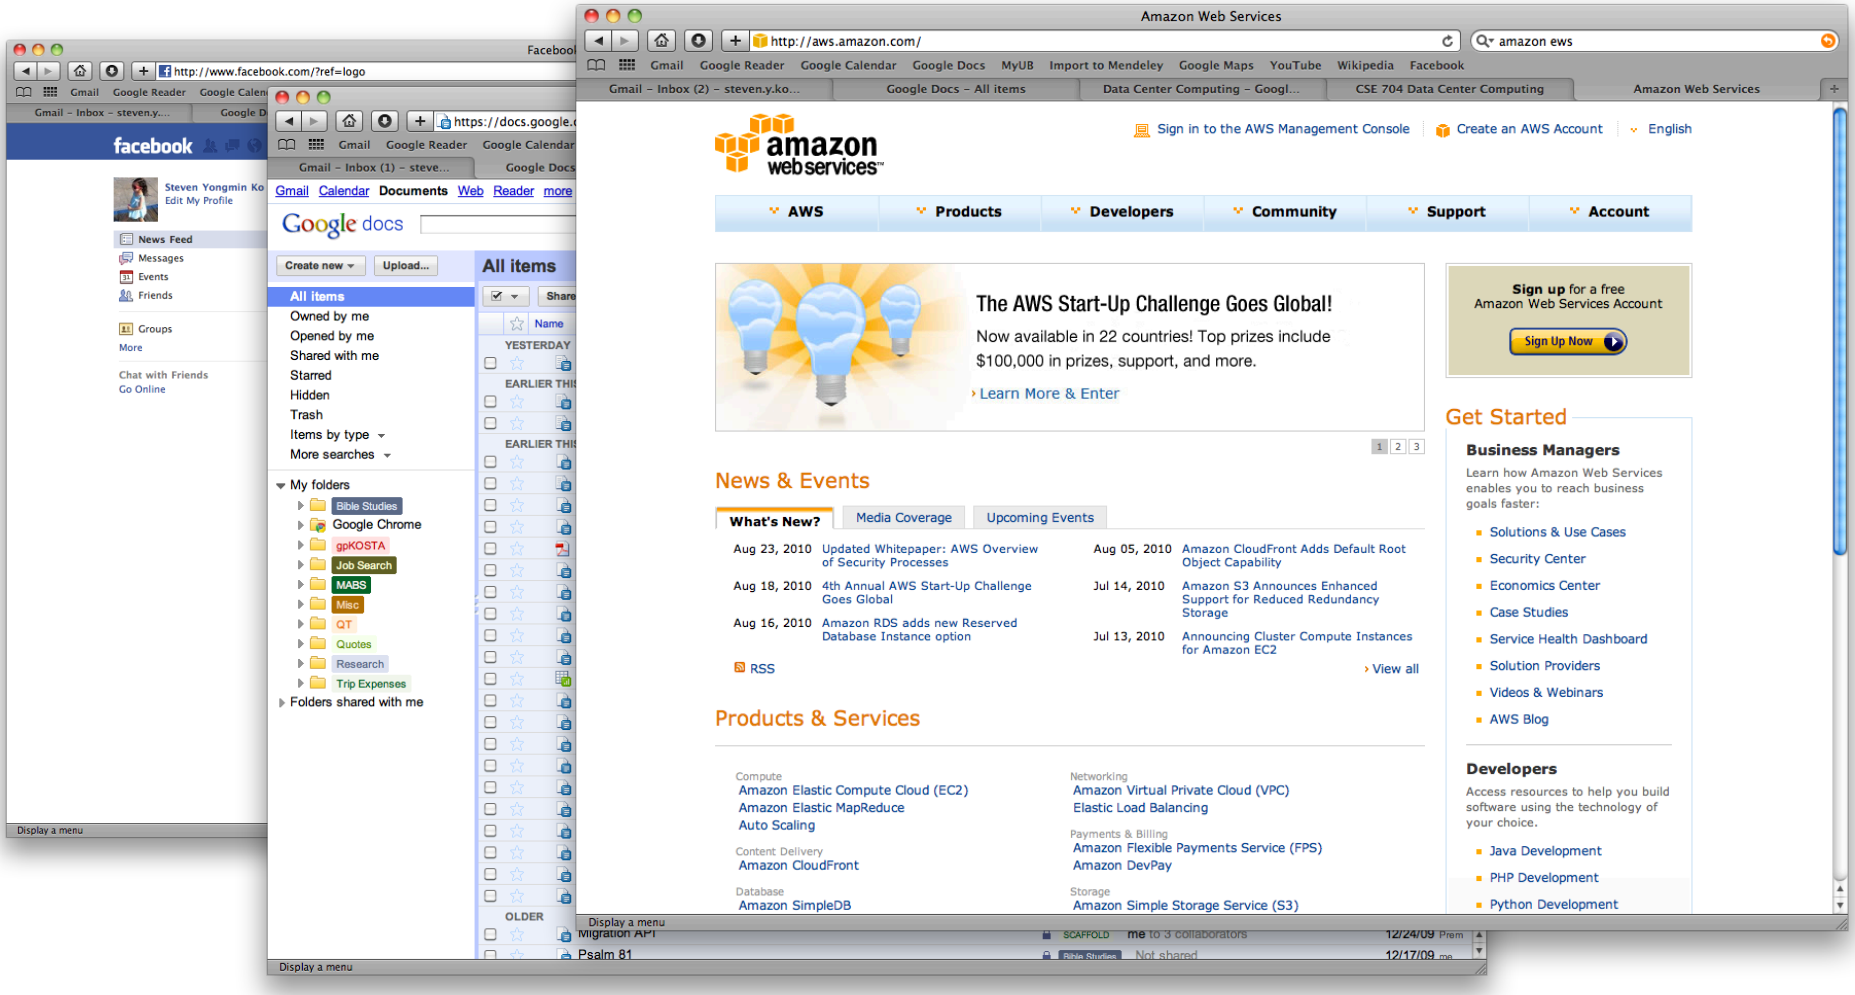

Topics Overview

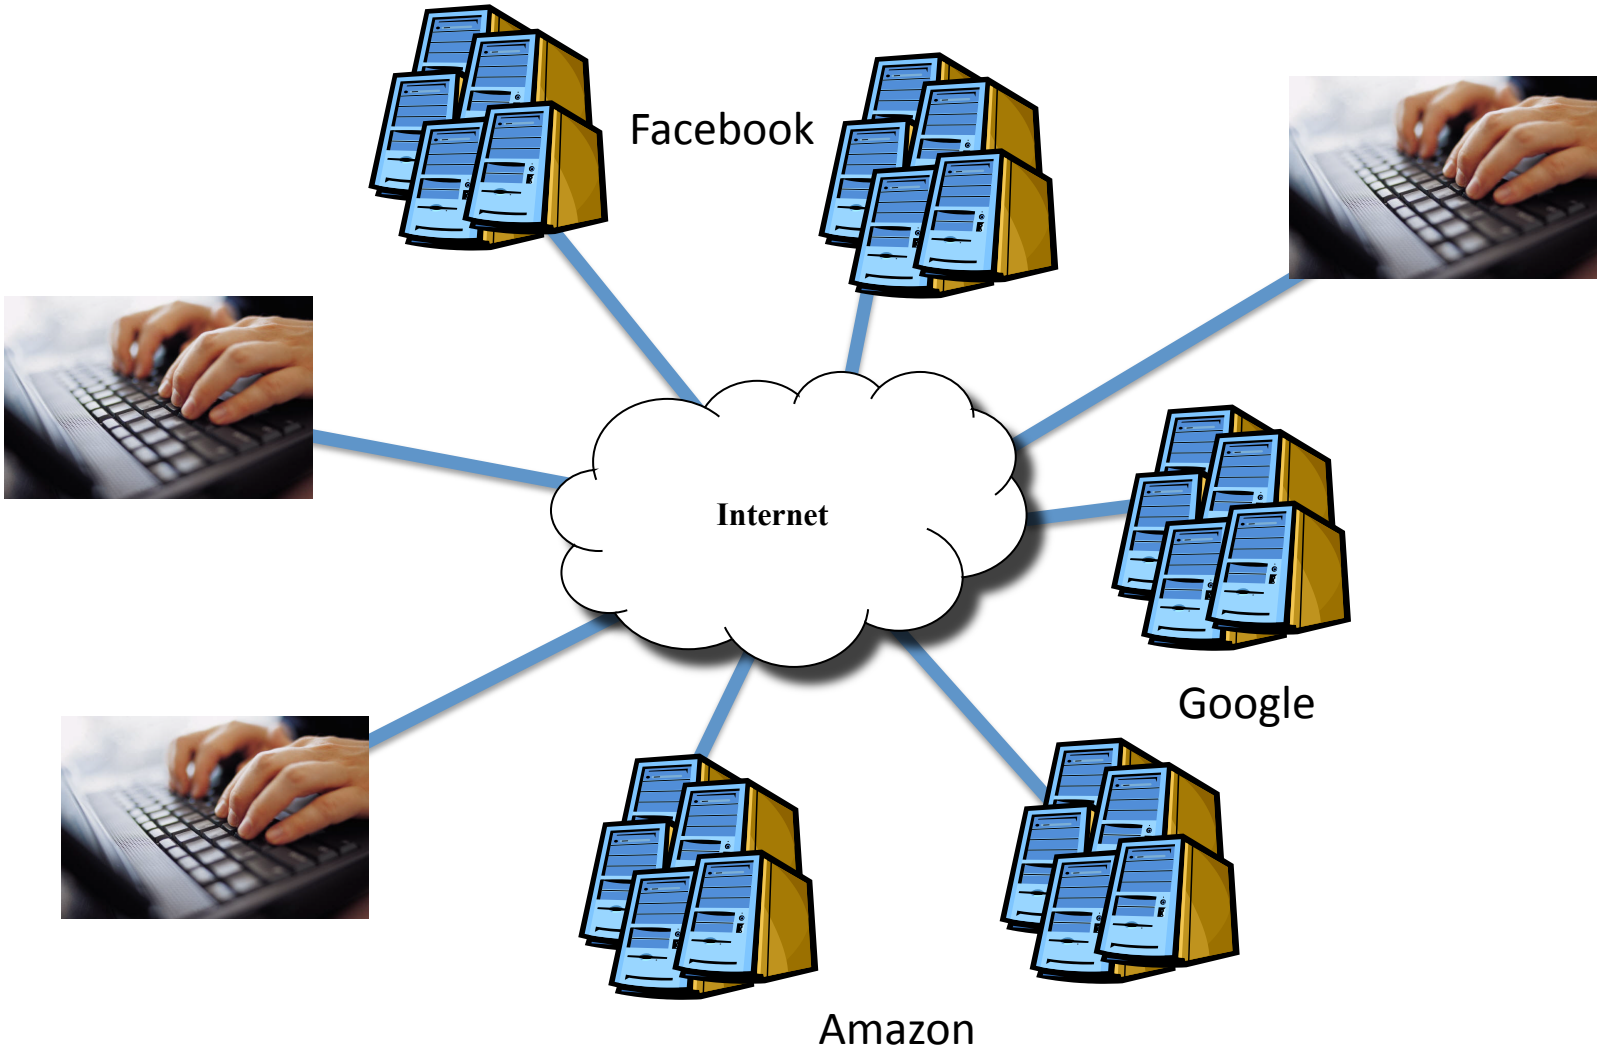

Data Centers

Data Centers

- Hundreds of locations in the US

Inside

- Servers in racks
  - Usually ~40 blades per rack
  - ToR (Top-of-Rack) switch
- Incredible amounts of engineering efforts
  - Power, cooling, etc.

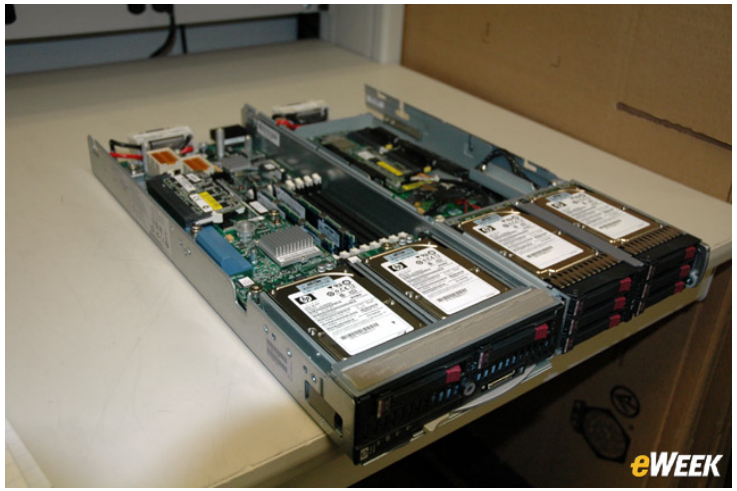

Inside

- Network

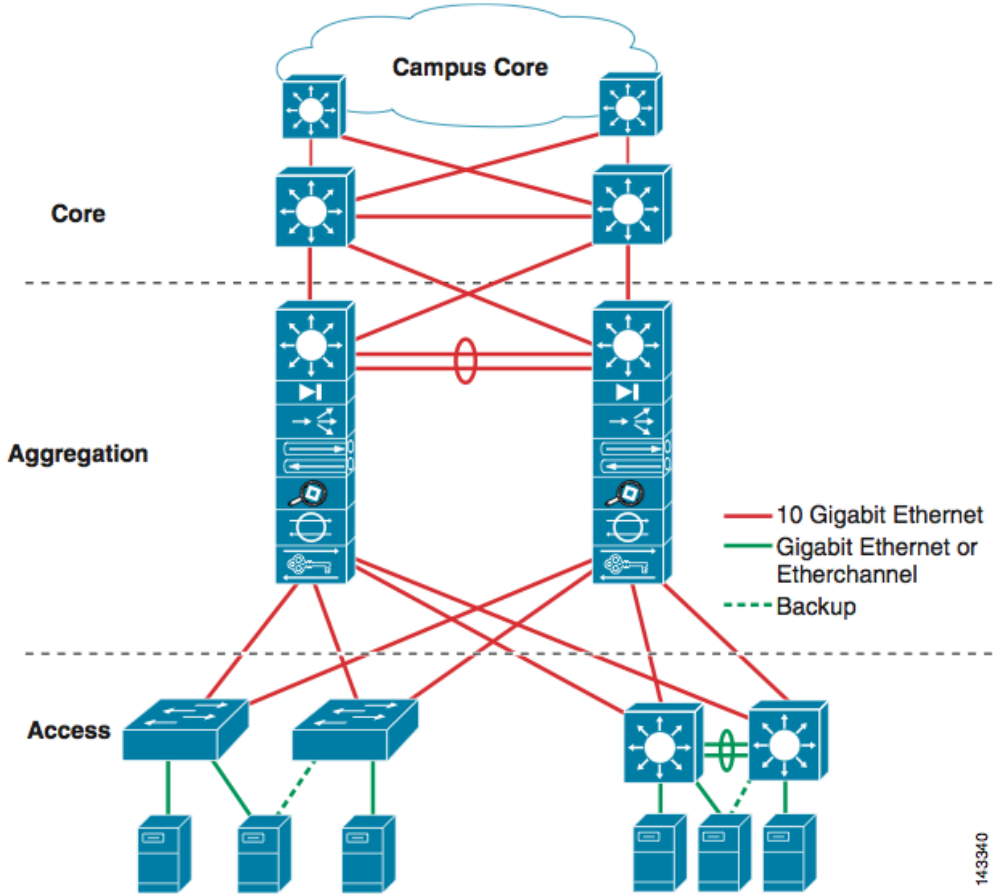

Inside

- 3-tier for Web services

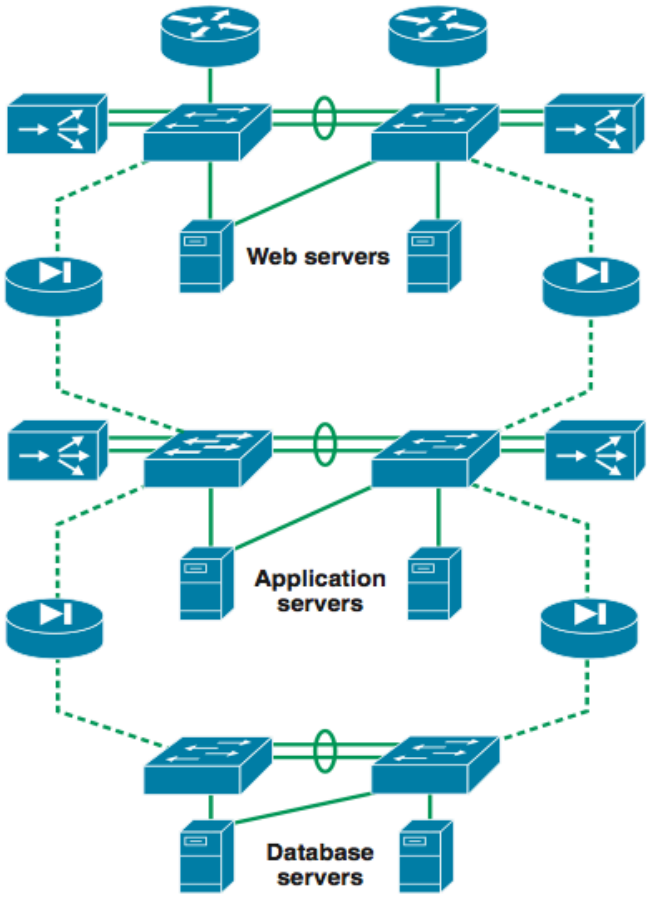

Topics Overview

- What makes these possible?

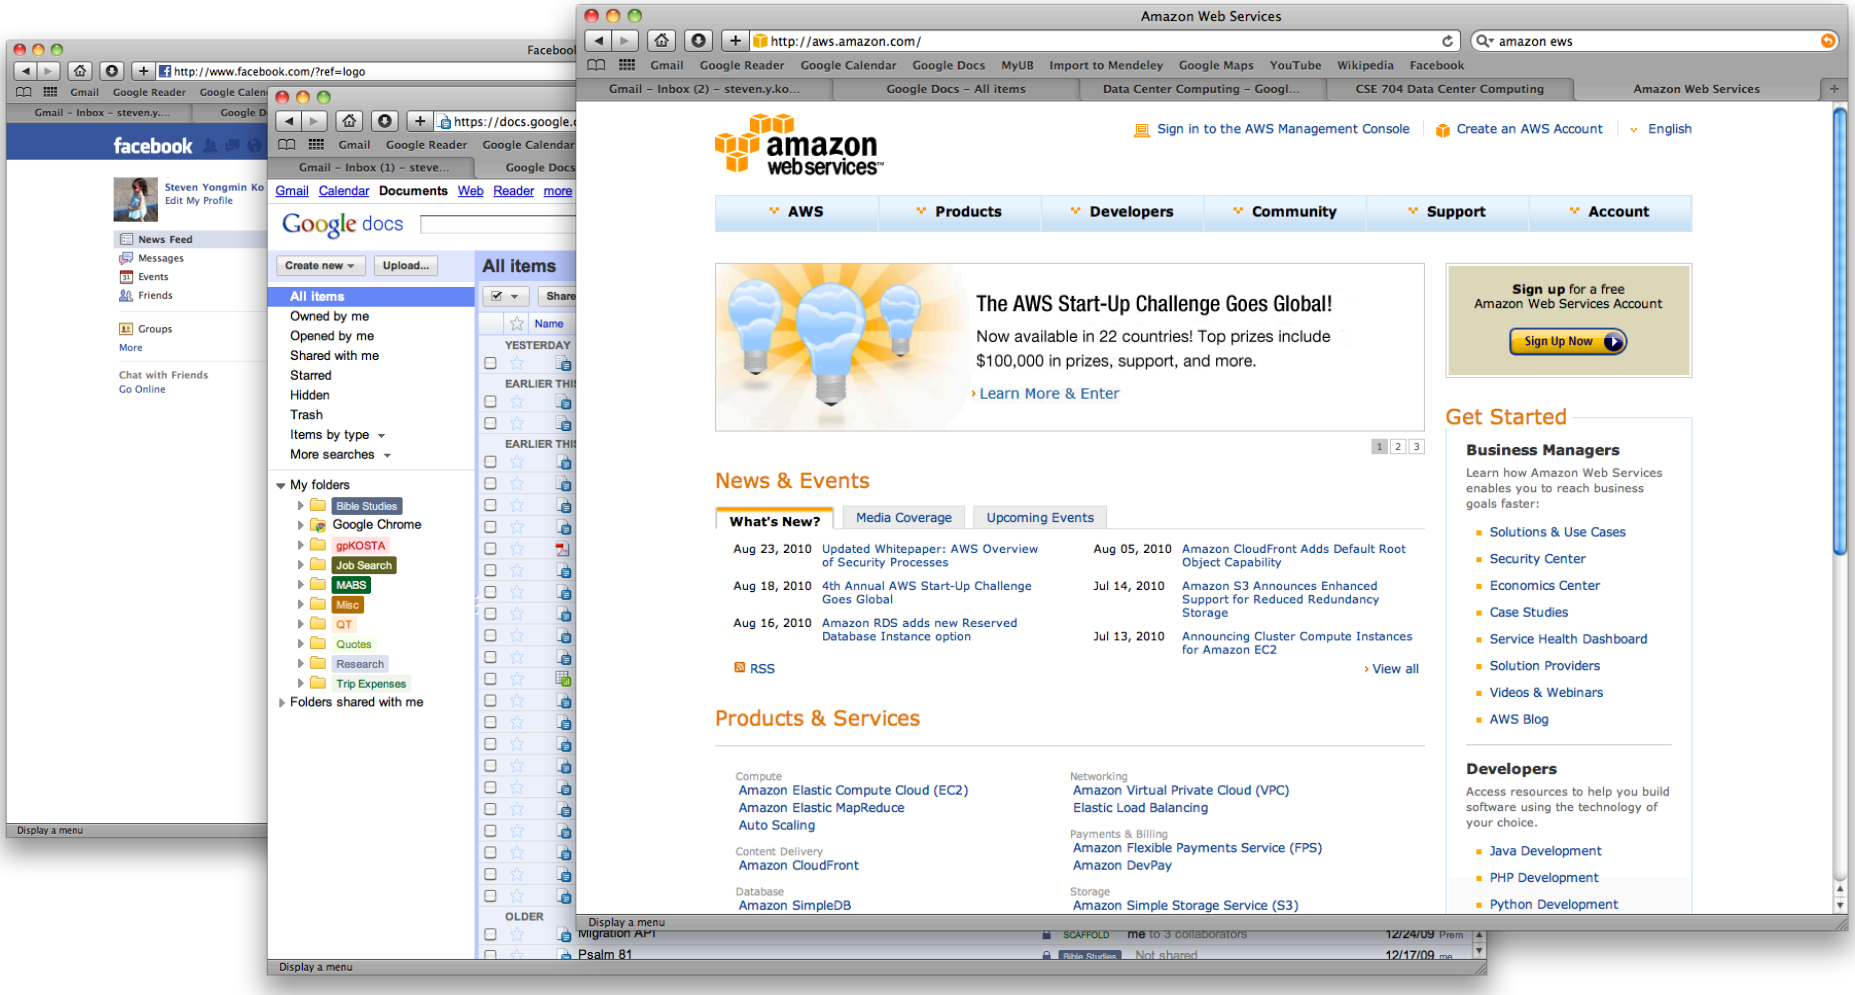

Components

- Front-end Web browsers
- File and storage systems
- Data processing frameworks
- Networking
- Virtualization
- Maybe not an exhaustive list, but a good set…

Front-End Web Browsers

- Why Web browsers?
  - Practically, they are the OSes in the current generation of computing
- They run applications
  - Maps, email clients, etc. (AJAX programs) have hundreds of thousands of LOC
  - All traditional OS problems exist
    - Protection, reliability, privacy, performance, etc.
- Exciting new area of research

A Glimpse into the Issues

- How much do you trust javascripts?

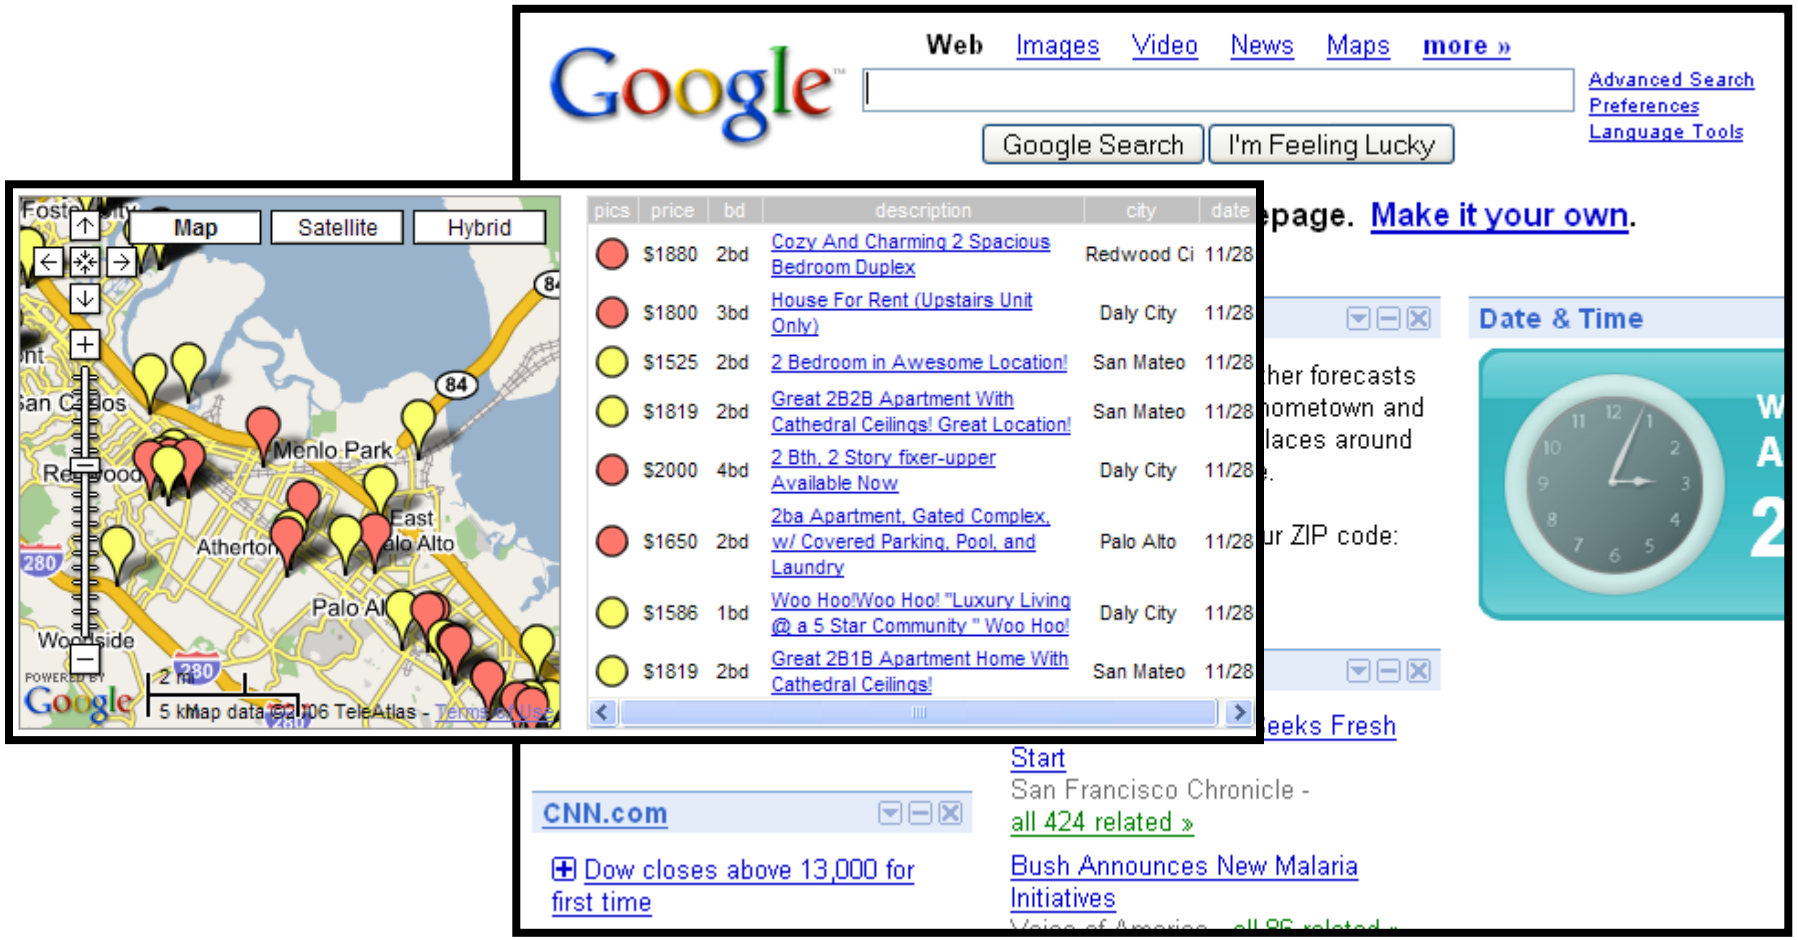

A Glimpse into the Issues

- SOP (Single-Origin Policy)

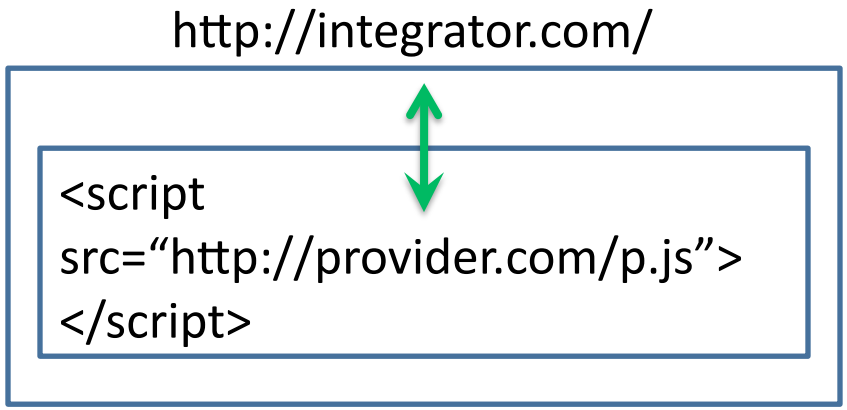

- How do you control the level of trust?

File and Storage Systems

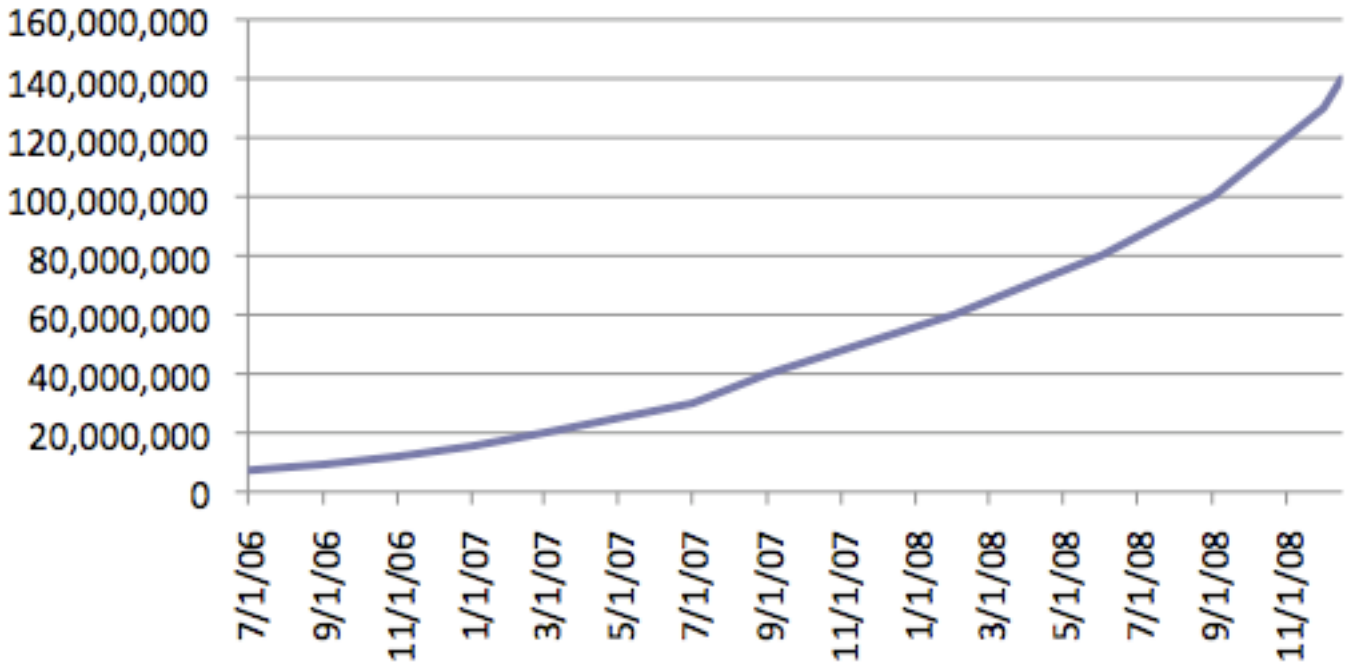

Facebook Statistics

- 13 M users update their statuses at least once each day
- 2.5 M users become fans of Pages each day
- 700 M new photos per month
- 4 M new videos per month
- 15 M pieces of content shared per month
- 2 M new events per month
- 19 M active groups

File and Storage Systems

- How do you store?
- How do you not lose?
- How do you provide good access latency?
- How do you maintain?
- …

Data Processing Frameworks

- Google
  - 20+ billion web pages
    - ~20KB each = 400 TB
  - ~ 4 months just to read the data
  - And growing…
    - 1999 vs. 2009: ~ 100X
- Yahoo!
  - US Library of Congress every day (20TB/day)
  - 2 billion photos
  - 2 billion mail + messenger sent per day
  - And growing…

Data Processing Frameworks

- How do you process this large amounts of data?
  - Page rank, ad click statistics, search query trend, user profiling, etc.
- Again, ~ 4 months to read the data

Virtualization

- VMM: a piece of software that exposes hardware interfaces (ISA, I/O, etc.)

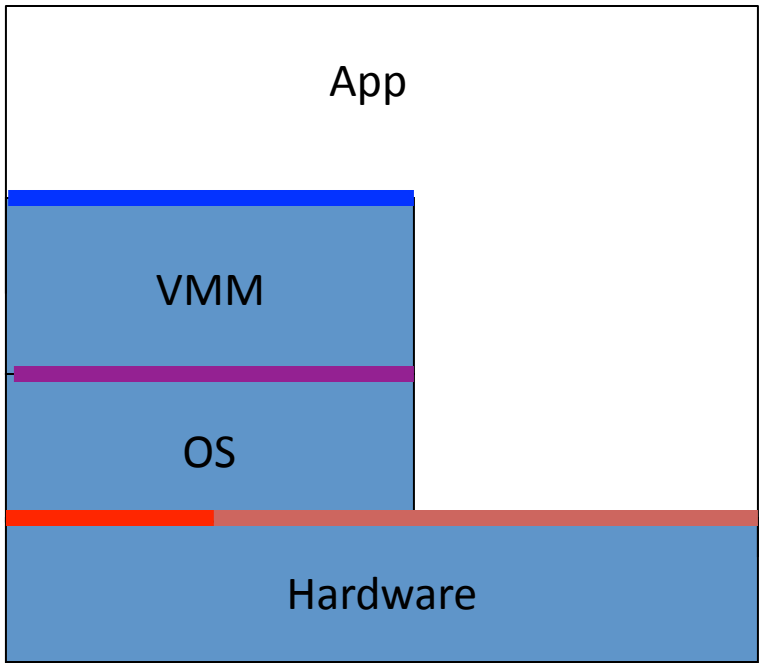

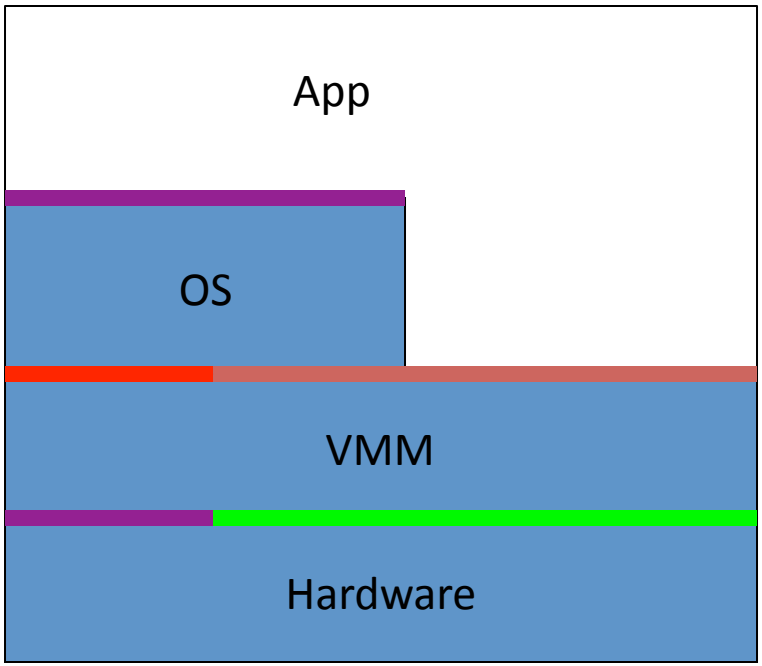

Virtualization

- Many data centers are virtualized
  - Ease of management (start, stop, migrate, etc.)
  - Consolidation (multiplexing one physical machine)
- How did we get here?

Networking

- Obviously, we need a network…
- But, what kind?

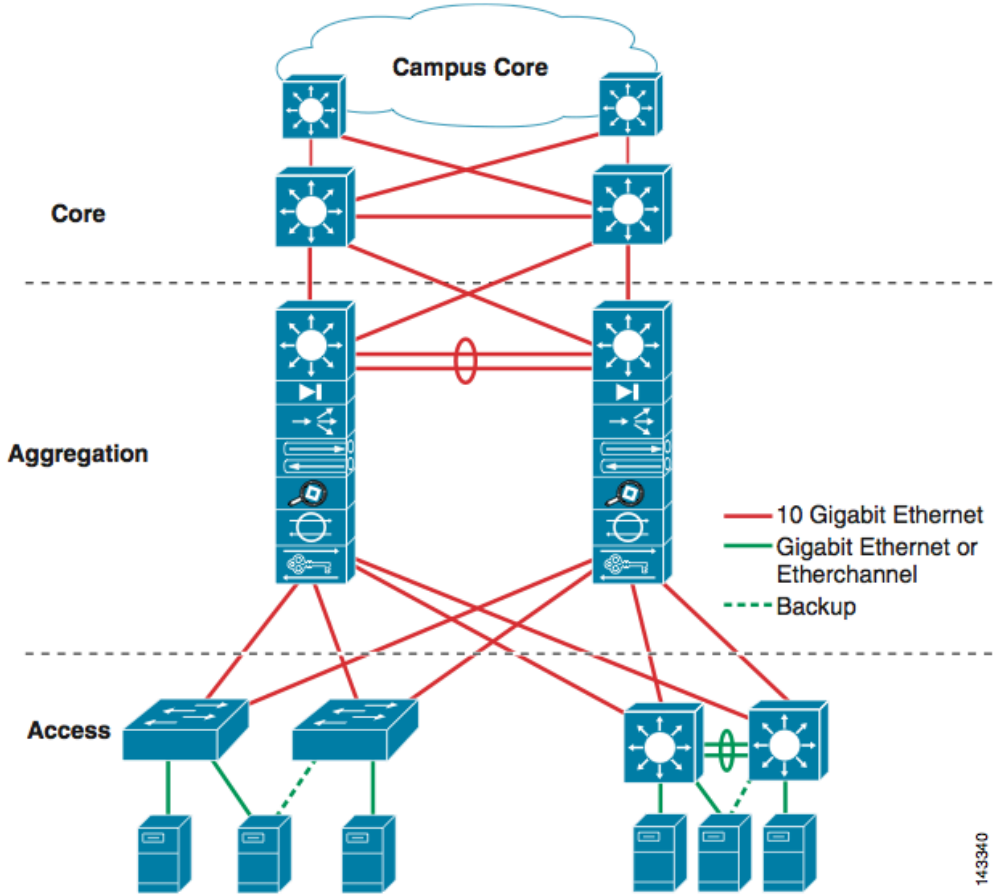

Networking Issues

- Oversubscription
- Management
- TCP performance
- Etc.

Assignment for Today

- Email me whether or not you're going to stay
- If you decide to stay, do the rest
- Look through the schedule
- Pick two papers you'd like to present
  - 1st choice & 2nd choice
  - FCFS
  - Exception: OpenFlow & NOX (from "Networking – 1" on 10/6) should be presented together with some demo
- Email me your choices