

Today's Objectives

- Review
- Recalibrating
- CoDeeN/PlanetLab

Oct 18, 2017

Sprenkle - CSCI325

1

Lucy's Talk

- Research
- Grad school

Oct 18, 2017

Sprenkle - CSCI325

2

Web Server

- Convention: packages named with lowercase

Oct 18, 2017

Sprenkle - CSCI325

3

Project 2: Tiny Bookstore

- What lessons have you learned so far?
- If you started over, how would your process change?

Oct 18, 2017

Sprenkle - CSCI325

4

Friday: Data Center

- Meet here, head over toward Law School

Oct 18, 2017

Sprenkle - CSCI325

5

Recalibrating: Midterm → Exam

- What I want to cover before exam:
 - Email
 - RAID
 - NFS
 - MapReduce (3rd project)
- Earliest: Nov 3, Latest: Nov 10
 - Phil Bennett from Microsoft – Oct 27
 - Priya from Google – network architecture

Oct 18, 2017

Sprenkle - CSCI325

6

CODEEN

Oct 18, 2017

Sprenkle - CSCI325

7

CoDeeN Discussion

- CoDeeN Content Distribution Network
- What is a content distribution network?
 - Goals? Benefits?
- What is PlanetLab?
- Discuss (general) tradeoffs of centralized vs decentralized systems
- What are approaches to security?

Oct 18, 2017

Sprenkle - CSCI325

8

Security

- Prevention
- Detection/Recovery
 - Once you notice it, shut it down and recover

Oct 18, 2017

Sprenkle - CSCI325

9

<https://www.planet-lab.org/>



PLANETLAB

An open platform for developing, deploying, and accessing planetary-scale services

Oct 18, 2017

Sprenkle - CSCI325

10

PlanetLab

- PlanetLab consists of 1353 (Linux) nodes at 717 sites



<https://www.planet-lab.org/>

Oct 18, 2017

Sprenkle - CSCI325

11

PlanetLab Requirements

- Provide global platform for short-term experiments and long-running services
- Convince sites to host PlanetLab machines that run code written by unknown researchers from other organizations
- Support autonomy and decentralized control of machines
 - This isn't entirely true...
- Scale to support many users

Oct 18, 2017

Sprenkle - CSCI325

12

PlanetLab Design

- “Distributed virtualization”
- Each user is allocated a *sliver* on each resource
 - Proportional fair sharing of resources (CPU, memory)
- Aggregation of all of your slivers = your *slice*
 - A slice is basically like your PlanetLab login/account on the remote PlanetLab hosts
 - Your “slice” of global PlanetLab resources
- Node managers at each physical host instantiate virtual machines for each sliver
 - You can’t see what anyone else is doing

Oct 18, 2017

Sprenkle - CSCI325

13

Using PlanetLab

- Each slice gets SSH access to PlanetLab resources
- Slices are initially empty (no slivers)
- If you want access to a specific machine, add it to your slice
- No common file system (!)
- Minimal software installed by default
 - No Java, Python
 - If you want it, you have to install it

Oct 18, 2017

Sprenkle - CSCI325

14

From authors

CODEEN

Oct 18, 2017

Sprenkle - CSCI325

15

Content Distribution Networks

- Replicates Web content broadly
- Redirects clients to “best” copy
 - Load, locality, proximity
- Offloads work from origin servers
- Multiplexes load spikes
 - Reduces overprovisioning
- Ex: Akamai, Mirror Image, Speedera

Aug 12, 2003

CoDeeN Overview - IRIS/PlanetLab

16

What Does It Do?

- An Academic Content Distribution Network
 - Redirects/caches HTTP requests
 - Based on our OSDI 2002 paper on CDN performance
- An Open Proxy Network
 - Probably the largest in existence

Aug 12, 2003

CoDeeN Overview - IRIS/PlanetLab

17

Looking Ahead

- Tiny Bookstore - Tuesday
- Read MapReduce paper by Monday midnight

Oct 18, 2017

Sprenkle - CSCI325

18