COS 448: Innovating Across Technology, Business, & Markets

Spring 2011
MW 1100-1220 in CS105
Prof. JP Singh

TA: Muneeb Ali

Road Map: Recap

Design Principles for Scalable Systems:

- Divide & Conquer
- Asynchrony
- Concurrency, etc.

Trends-1: Cloud Computing

- What technology is available, where is it heading

Scalable Systems

- How to build scalable systems
- Part-1: Scalable Infrastructure
- Part-2: Scalable Software Design

Road Map: Today's Lecture

Scalable Systems

- Example company; grows from 10 users to 100 million users
- Ad serving company spin out
- How both managed growth in a different way
- The design tradeoffs made along the way

A Brief History of Cloud Computing

- A quick recap of technological progress in the last decade

Future of Cloud Computing

- Where are we heading …

Leveraging the Cloud: Example

The Beginning

- Muneeb and JP meet at Cafe Grumpy

Leveraging the Cloud: Example

The Beginning

- Muneeb and JP meet at Cafe Grumpy
- Let's start a blog about local news – The High Line Times
- www.highlinetimes.us
- Where to host?

ROAD RUNNER
HIGH SPEED ONLINE

Meraki Wifi

Old Macbook Pro

High Line Times

The Beginning

- Where to host?

ROAD RUNNER
HIGH SPEED ONLINE

Meraki Wifi

- 54 Mbps (bits)
  (802.11g)

Old Macbook Pro

- CPU: 2.4 Ghz
- Mem: 2 GB
- Disk: 200GB

High Line Times

The Beginning

- Start getting users!
- First problem?

ROAD RUNNER
HIGH SPEED ONLINE

Meraki Wifi

- 54 Mbps (bits) (802.11g)

Old Macbook Pro

- CPU: 2.4 GHz
- Mem: 2 GB
- Disk: 200GB

High Line Times

The Beginning

- Start getting users!
- First problem? **Uplink**

TIME WARNER CABLE Business Class®

Meraki Wifi

- 54 Mbps (bits) (802.11g)

Old Macbook Pro

- CPU: 2.4 GHz
- Mem: 2 GB
- Disk: 200GB

High Line Times

The Beginning

- Start getting **thousands** of users!
- Problems? Opportunities?

TIME WARNER CABLE Business Class®

Meraki Wifi

- 54 Mbps (bits) (802.11g)

Old Macbook Pro

- CPU: 2.4 GHz
- Mem: 2 GB
- Disk: 200GB

High Line Times

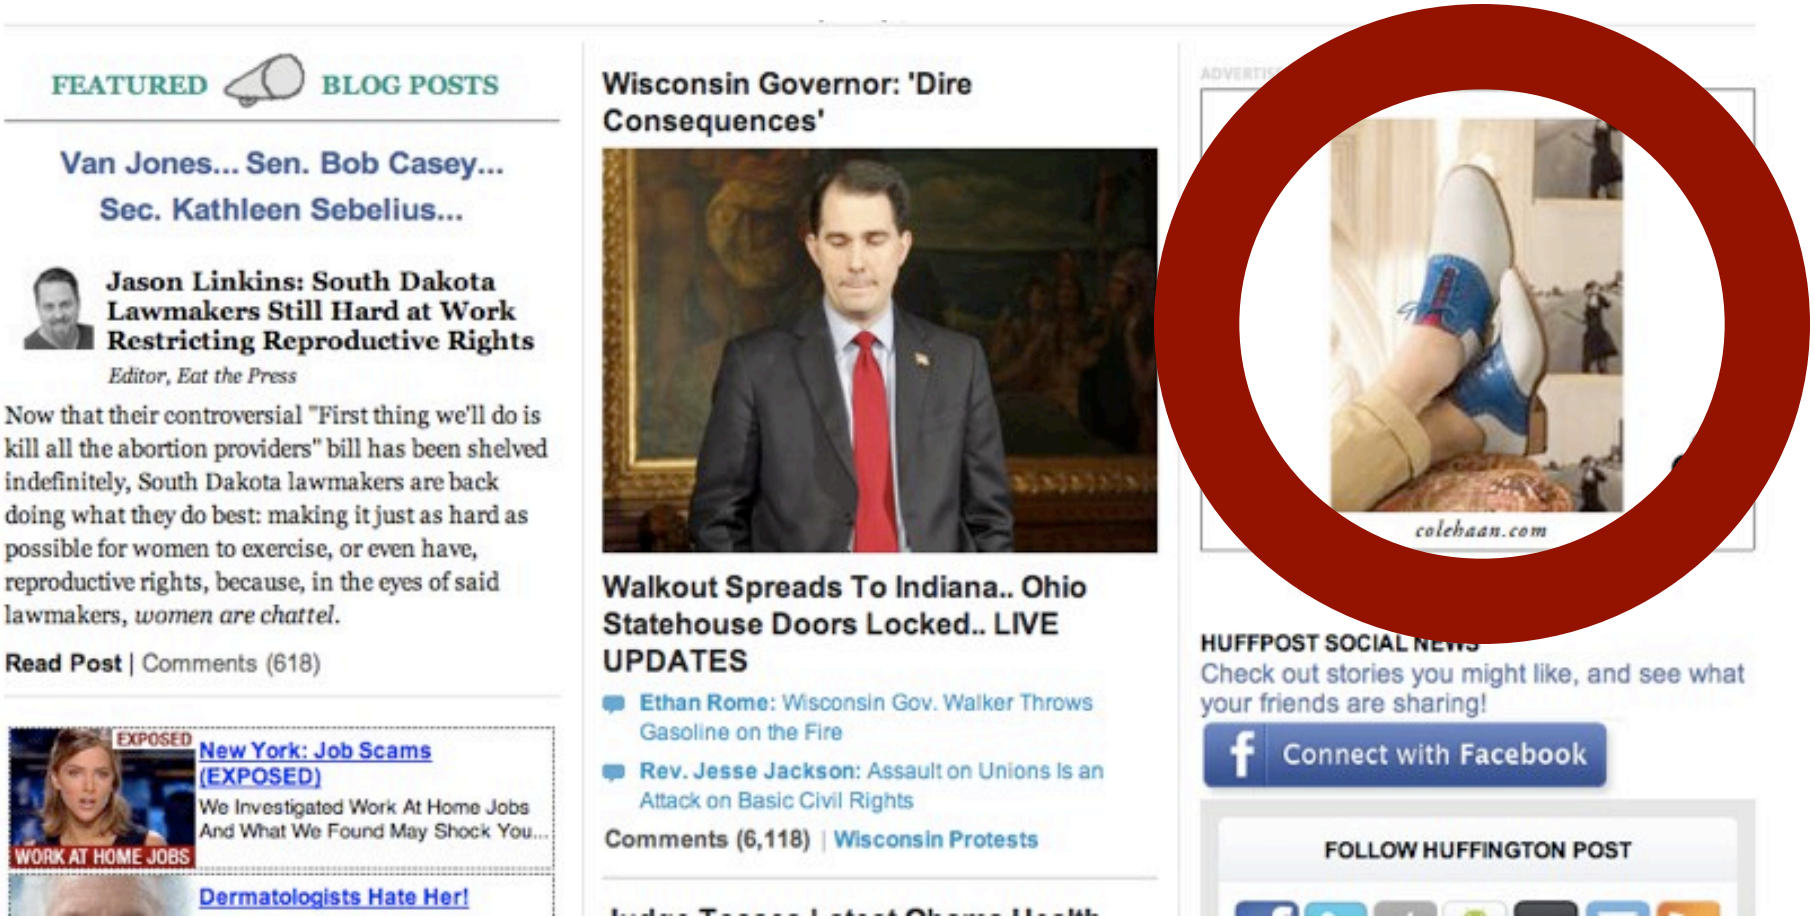

High Line Times

The Growth

- First investment (JP's money)
- $3,000 on server
- Problems?

TIME WARNER CABLE Business Class®

Sun Fire Server

- CPU: Dual-core 2.6 GHz each
- Mem: 8 GB (upto 16 GB)
- Disk: 500 GB
- Network: 1 Gbps Ethernet
- OS: Fedora 13

High Line Times

The Growth: Problems

- Slow Network
- Server makes noise (need proper space + cooling)
- Software bugs

TIME WARNER CABLE
Business Class®

Sun Fire Server

- CPU: Dual-core 2.6 GHz each
- Mem: 8 GB (upto 16 GB)
- Disk: 500 GB
- Network: 1 Gbps Ethernet
- OS: Fedora 13

High Line Times

The Growth: Datacenters

- Rent rackspace
- Apps4rent (NY/NJ)

COLOCATION SILVER PLAN

Class A Data Center
1U Rackspace
100 MBPS Port
1 MBPS Bandwidth included
$ 99 Setup
2 IPs included

$125/Month

High Line Times

The Growth: Software

Muneeb: Your code crashes my Fedora

JP: I don't like Fedora to begin with

Muneeb: Which one do you prefer?

JP: Ubuntu, duh

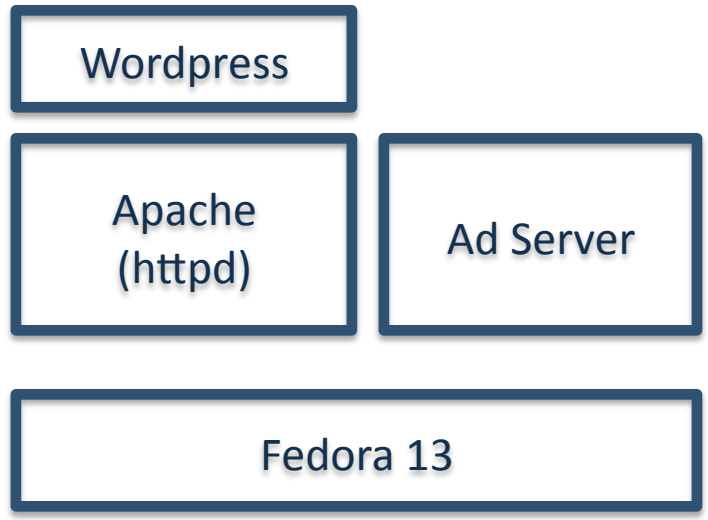

High Line Times

The Growth: Software

Muneeb: Your auctions take too long

JP: I know! They are so awesome

JP: I need a greater share of CPU

Muneeb: You're the boss!

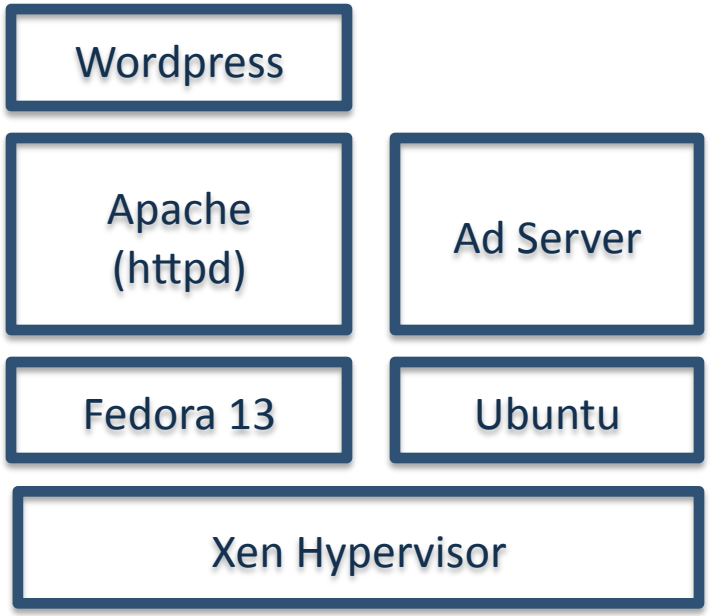

High Line Times

Further Growth: 100k Users

- Problems?

High Line Times

Further Growth: 100k Users

- Problems? More servers, more bandwidth, worried about backups

- CPU: Dual-core 2.6 GHz each
- Mem: 8 GB (upto 16 GB)
- Disk: 500 GB

Cost per server

$3,000 server
$150/month hosting

linode.com

	Linode 512	Linode 768	Linode 1024
RAM	512MB	768MB	1024MB
Storage	16GB	24GB	32GB
Transfer	200GB	300GB	400GB
Price	$19.95	$29.95	$39.95

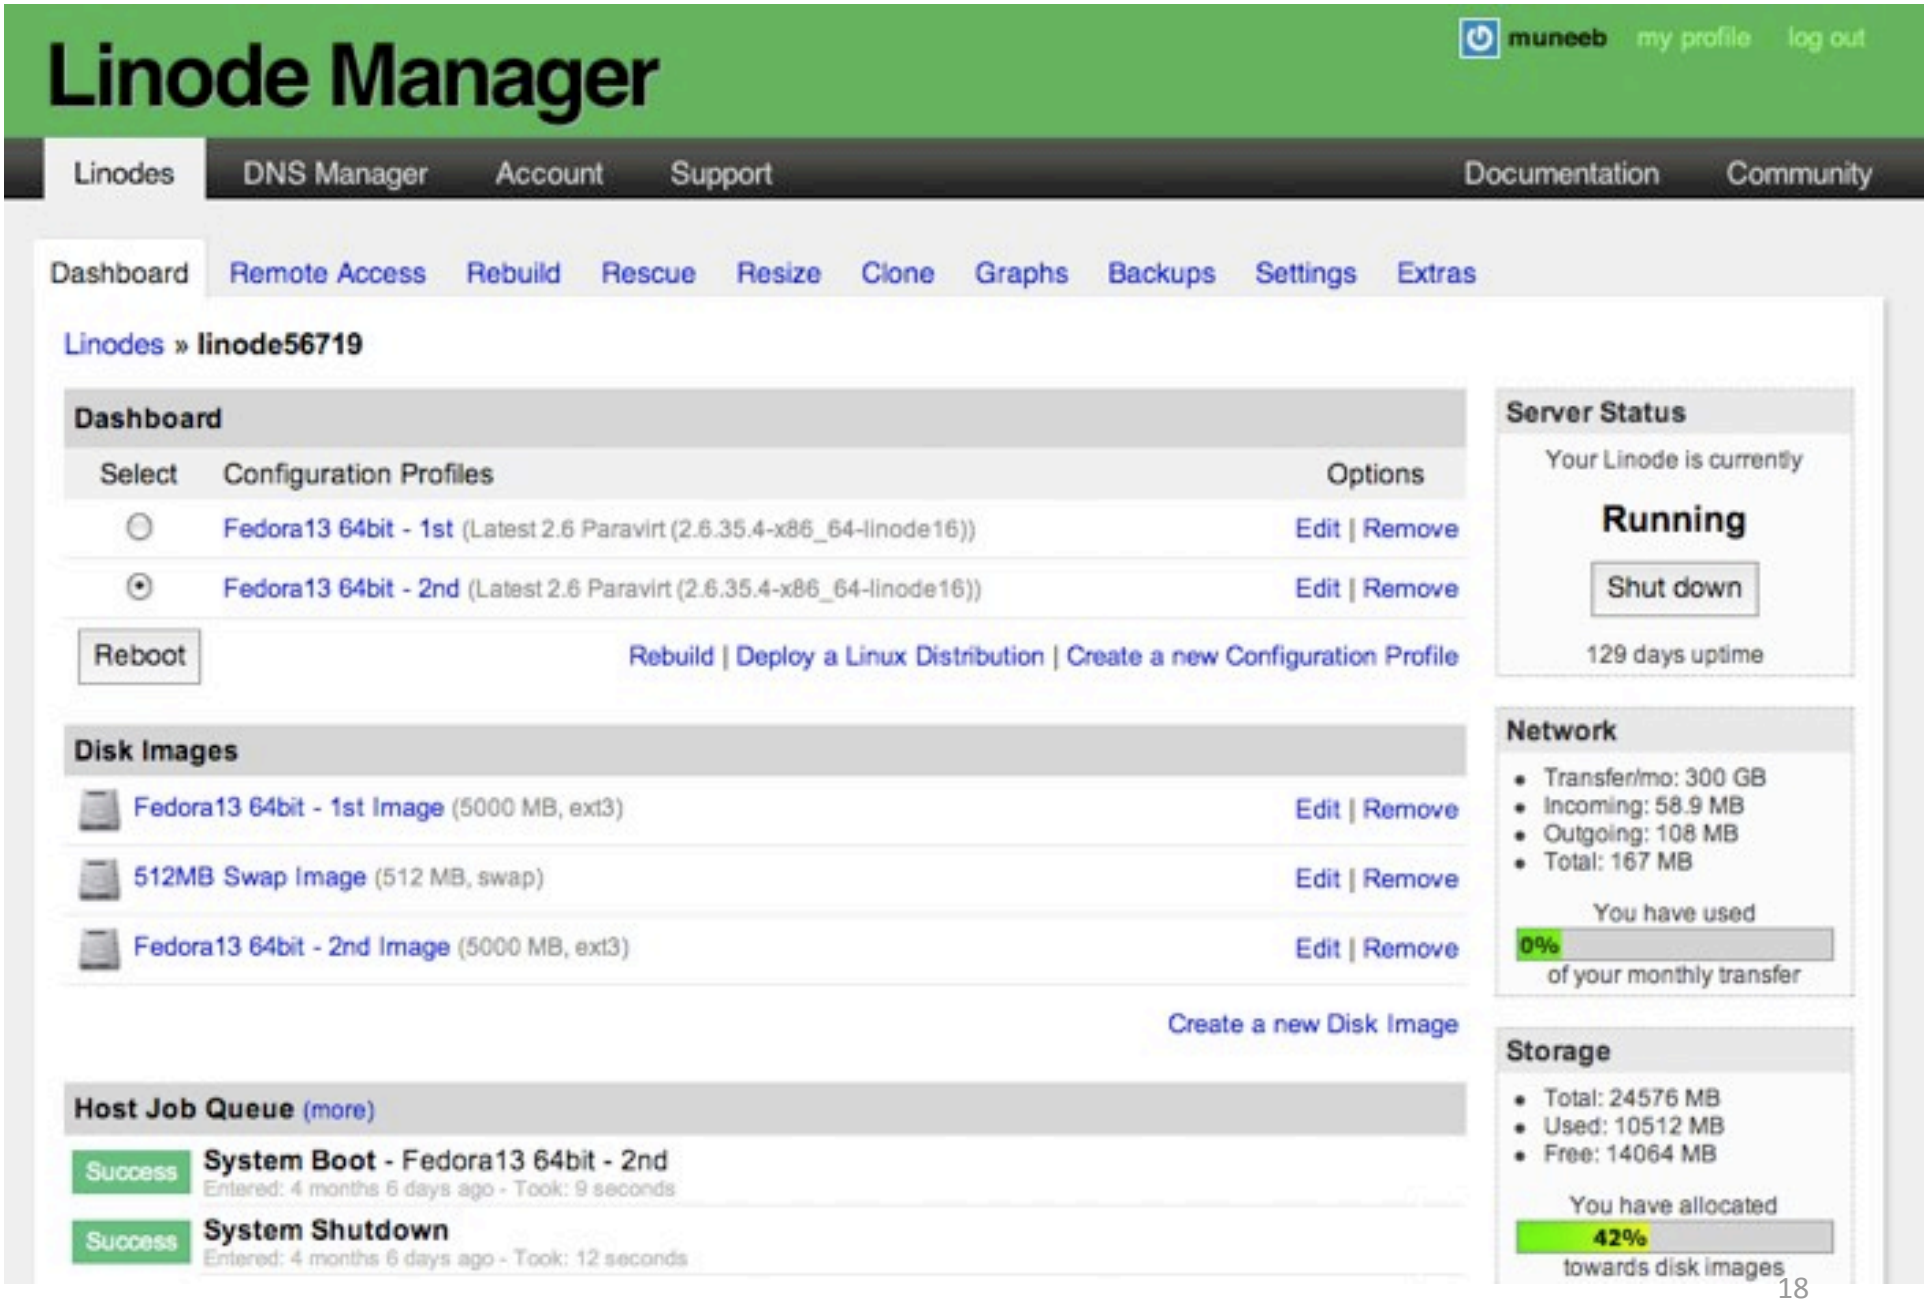
Linode Manager
muneeb
my profile
log out
Linodes
DNS Manager
Account
Support
Documentation
Community
Dashboard
Remote Access
Rebuild
Rescue
Resize
Clone
Graphs
Backups
Settings
Extras
Linodes » linode56719
Dashboard
Select
Configuration Profiles
Options
Fedora13 64bit - 1st (Latest 2.6 Paravirt (2.6.35.4-x86_64-linode16))
Edit | Remove
Fedora13 64bit - 2nd (Latest 2.6 Paravirt (2.6.35.4-x86_64-linode16))
Edit | Remove
Reboot
Rebuild | Deploy a Linux Distribution | Create a new Configuration Profile
Disk Images
Fedora13 64bit - 1st Image (5000 MB, ext3)
Edit | Remove
512MB Swap Image (512 MB, swap)
Edit | Remove
Fedora13 64bit - 2nd Image (5000 MB, ext3)
Edit | Remove
Create a new Disk Image
Host Job Queue (more)
Success
System Boot - Fedora13 64bit - 2nd
Entered: 4 months 6 days ago - Took: 9 seconds
Success
System Shutdown
Entered: 4 months 6 days ago - Took: 12 seconds
Server Status
Your Linode is currently
Running
Shut down
129 days uptime
Network
Transfer/mo: 300 GB
Incoming: 58.9 MB
Outgoing: 108 MB
Total: 167 MB
You have used
0%
of your monthly transfer
Storage
Total: 24576 MB
Used: 10512 MB
Free: 14064 MB
You have allocated
42%
towards disk images

2010	2010	2010	2010	2010	2010	2010	2010	2010	2010	2011	LAST	LAST
MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	JAN	30D	24H

Last 30 Days

CPU

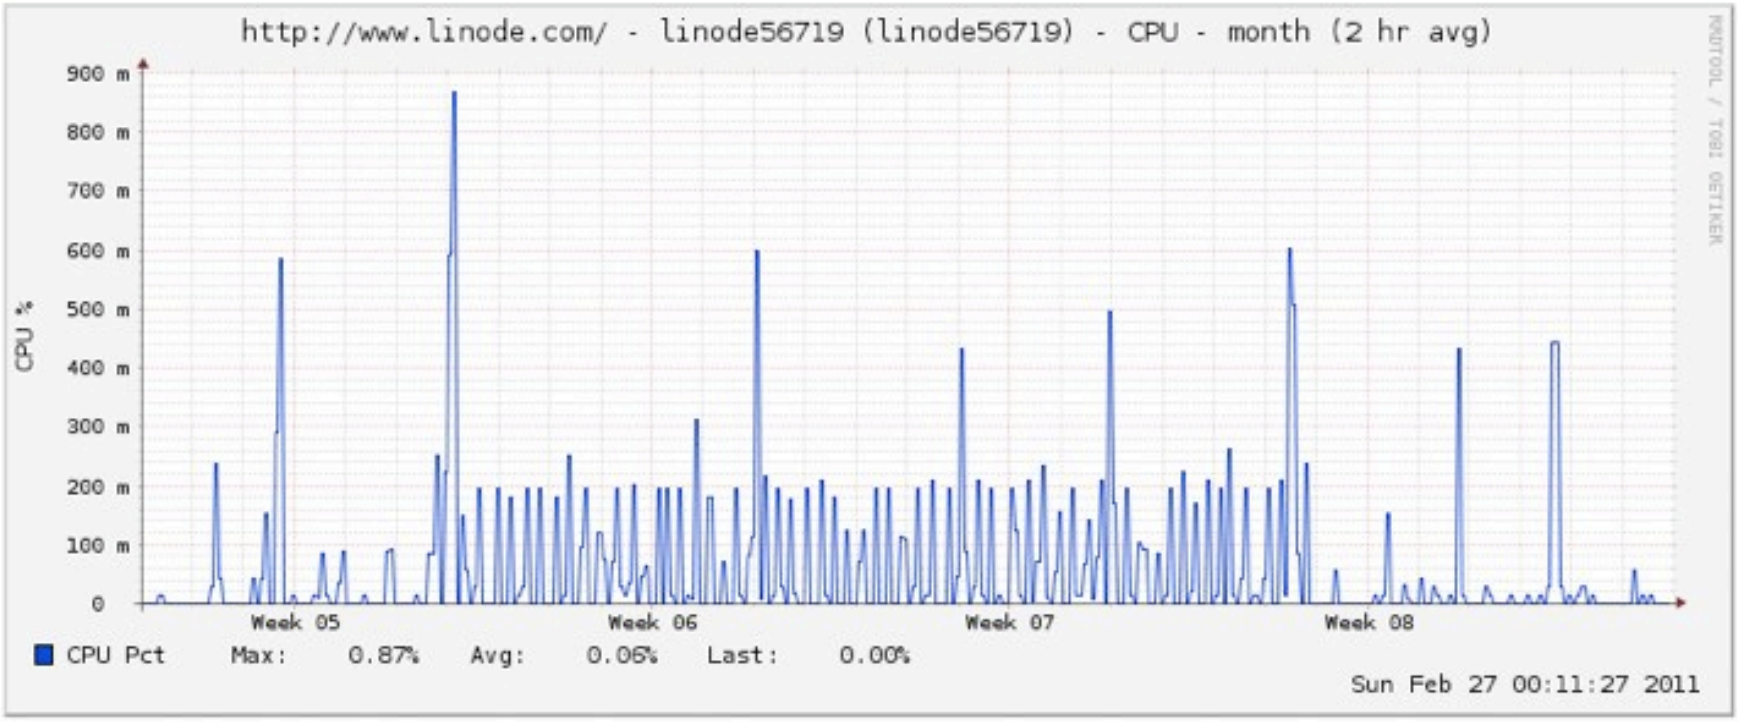

Network

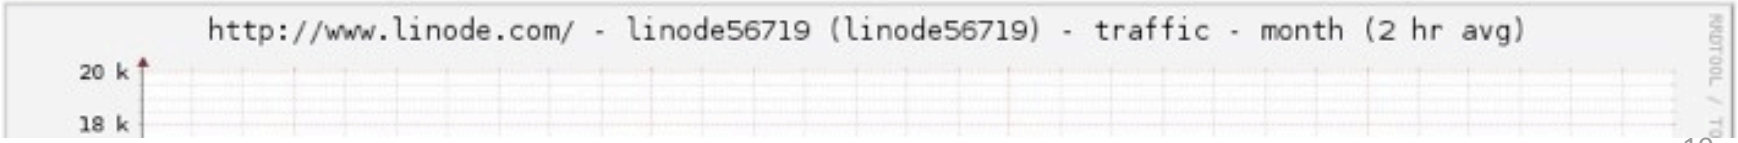

High Line Times

The Spinout

- Company separates online advertising and JP takes the lead on that
- JP decides to use Amazon's infrastructure
- Muneeb keeps the webservices at Linode (and further expands)

Web Services: Architecture

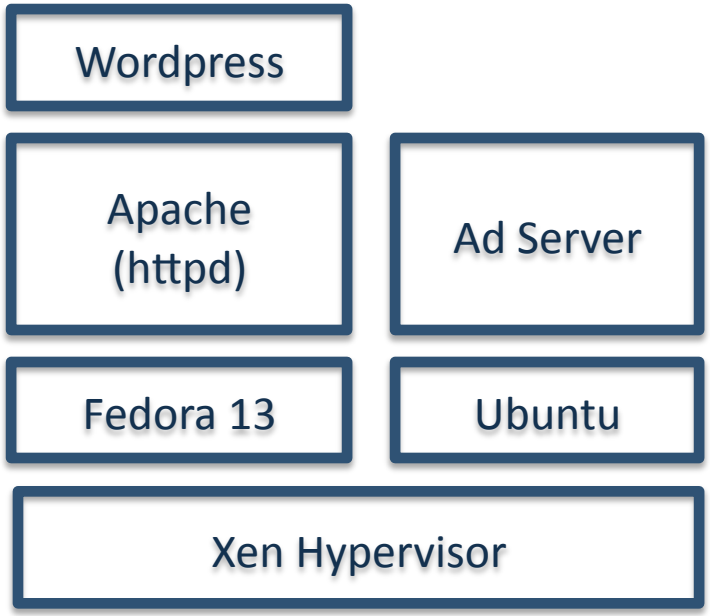

linode.com

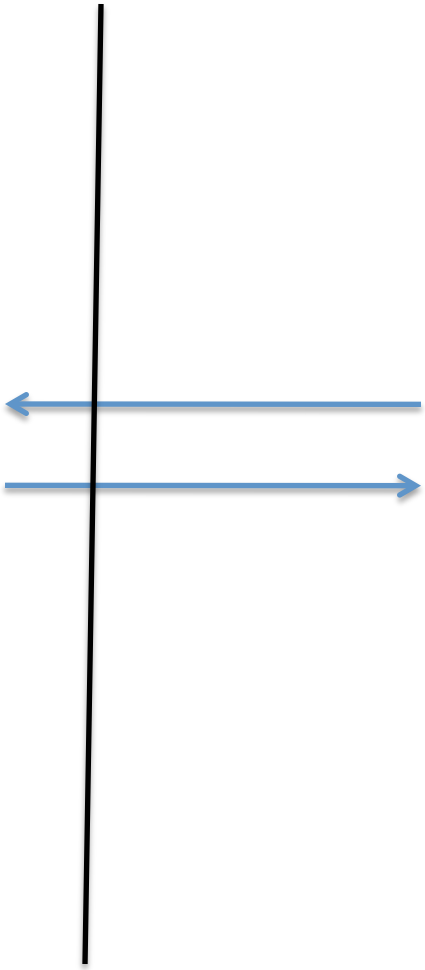

User

Web Services: Architecture

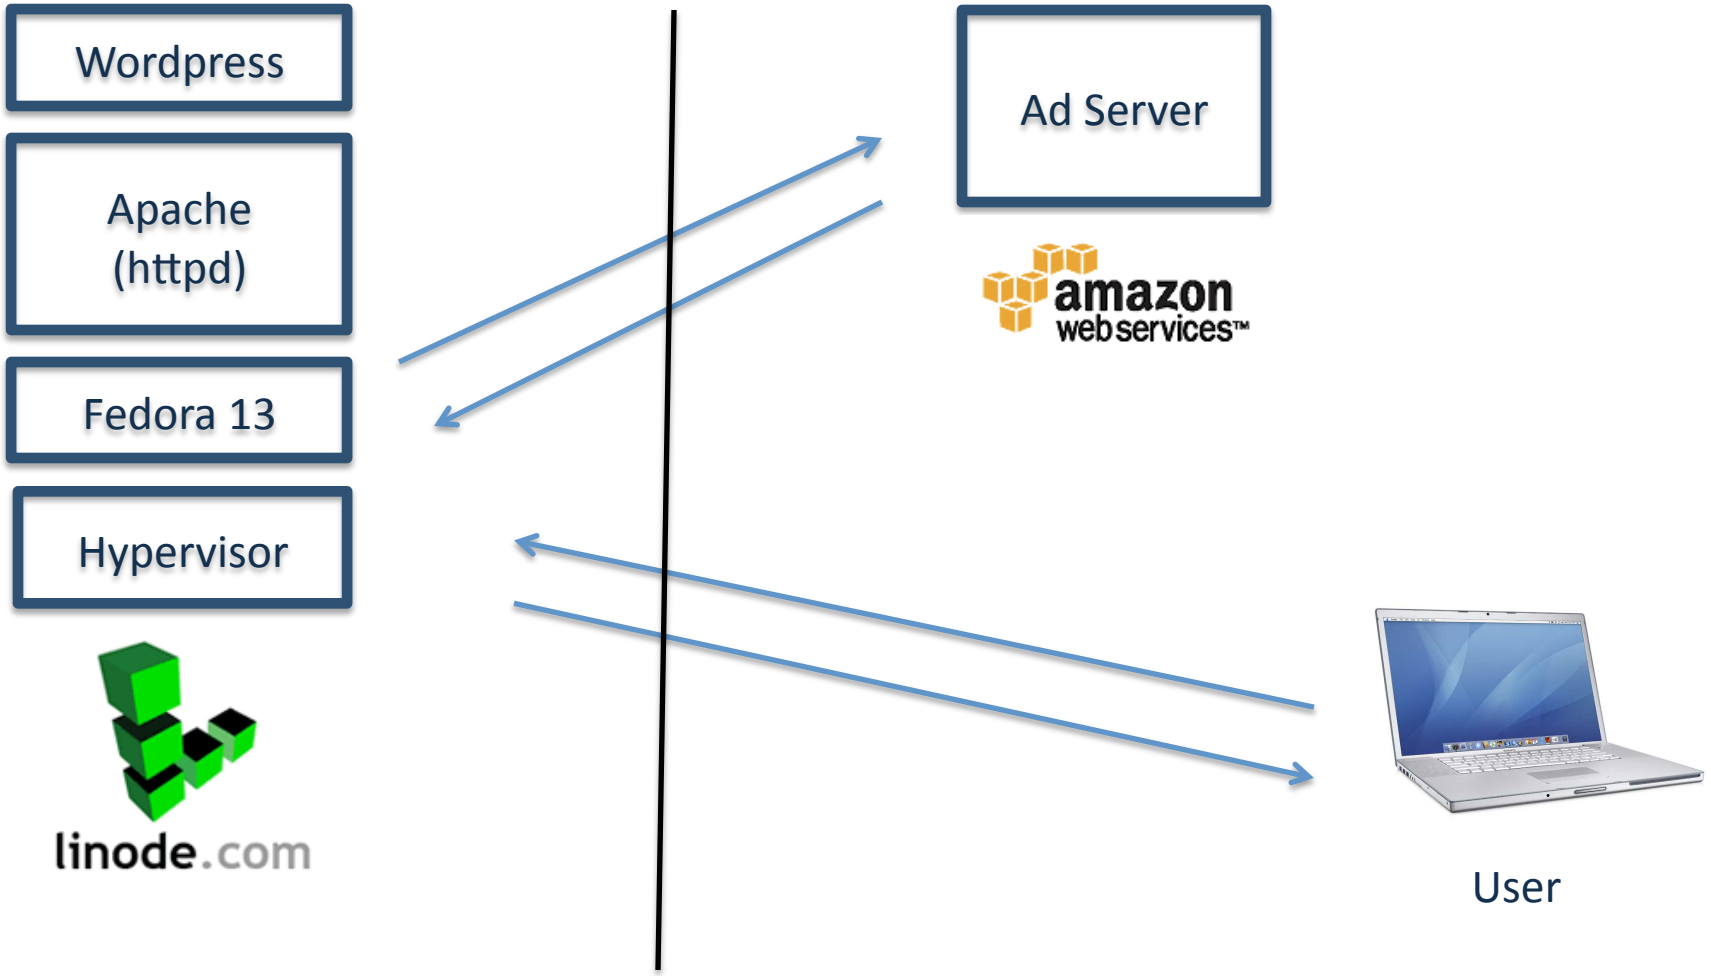

Web Services: Architecture

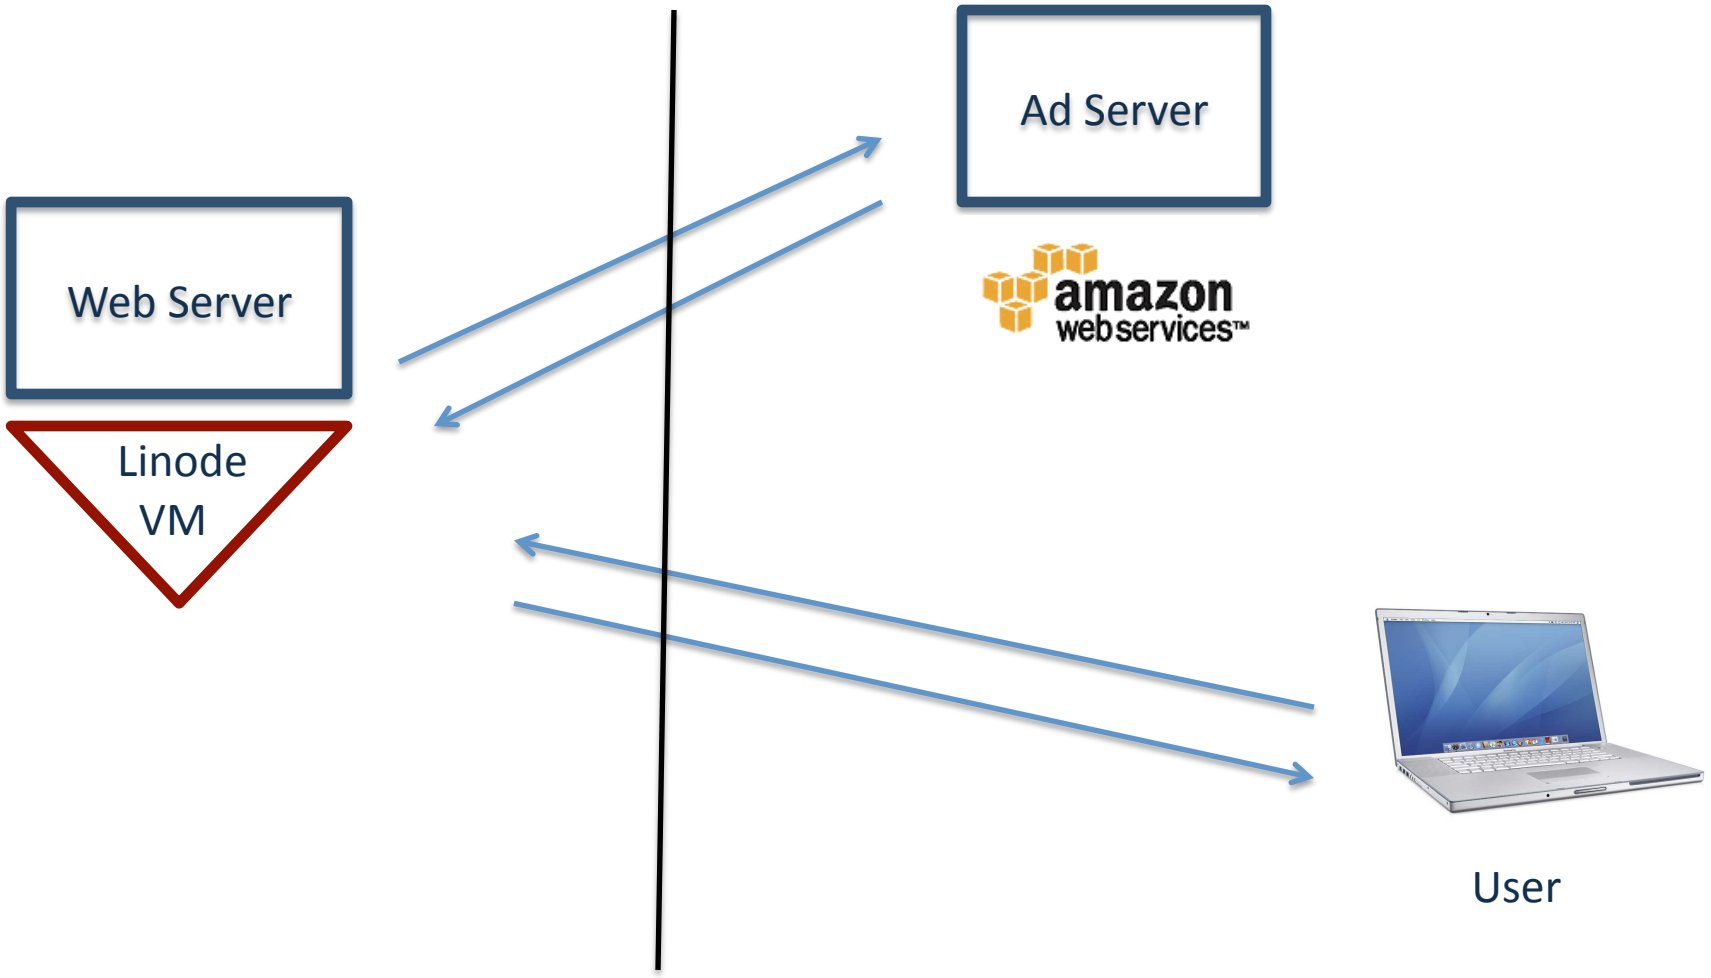

Web Services: Scale

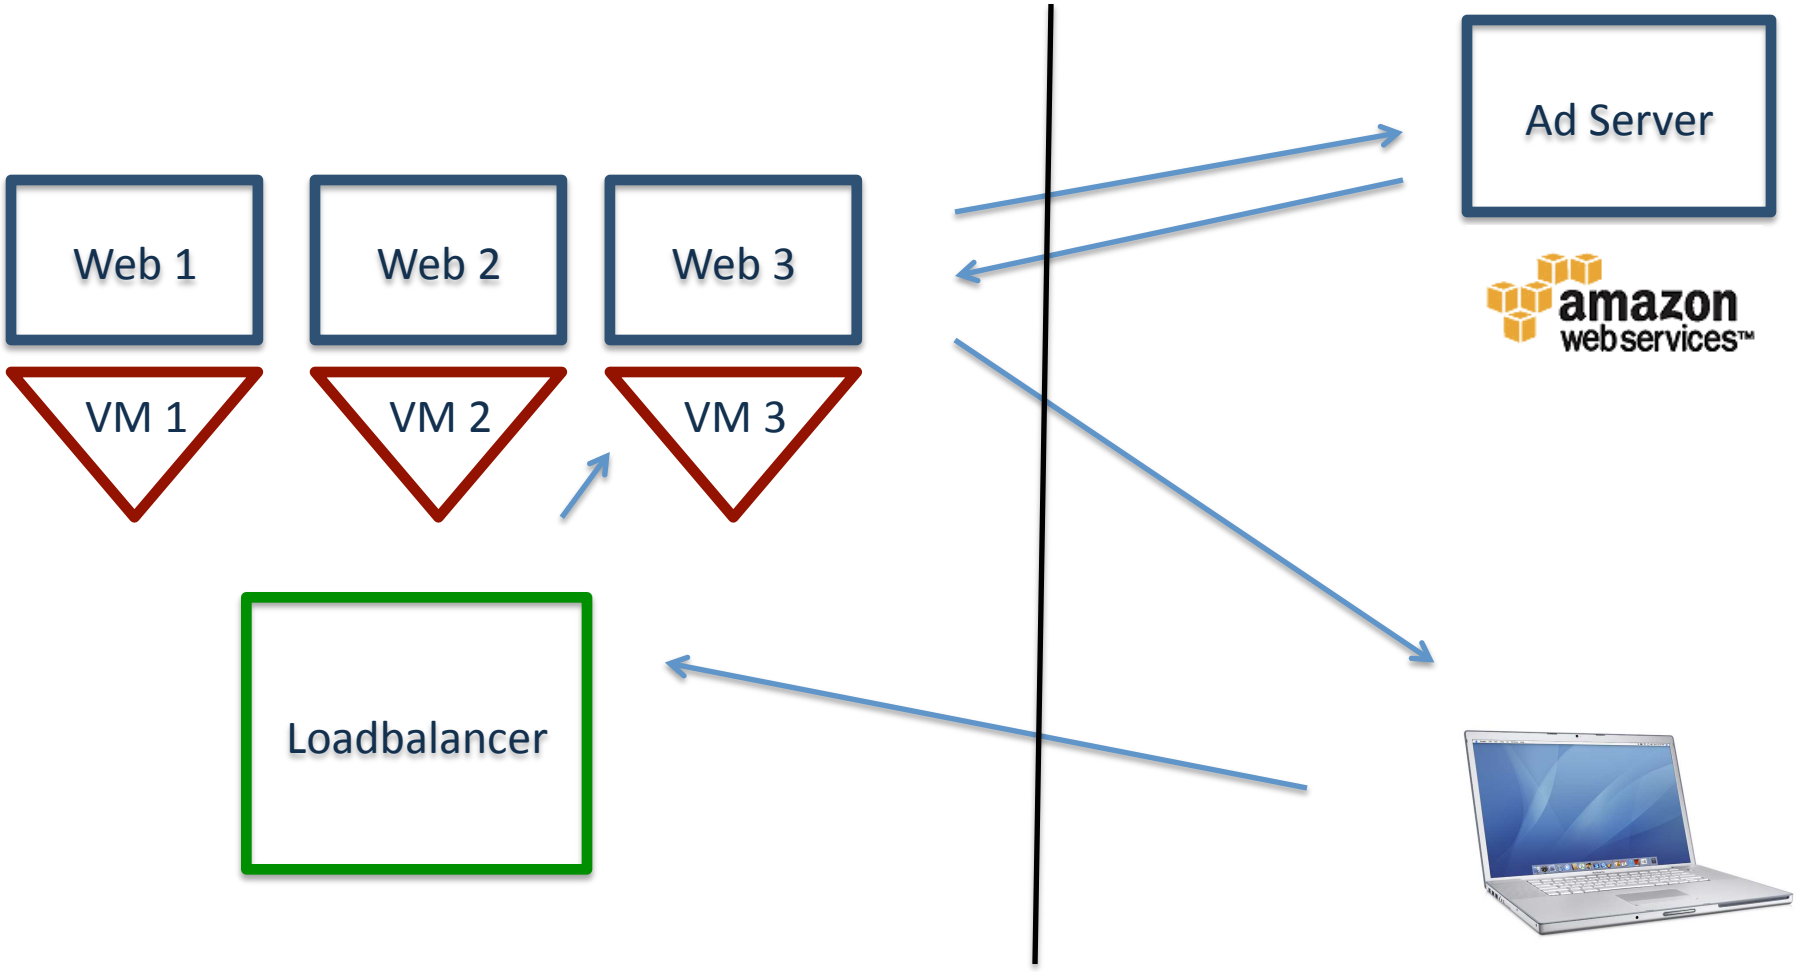

Web Services: Scale

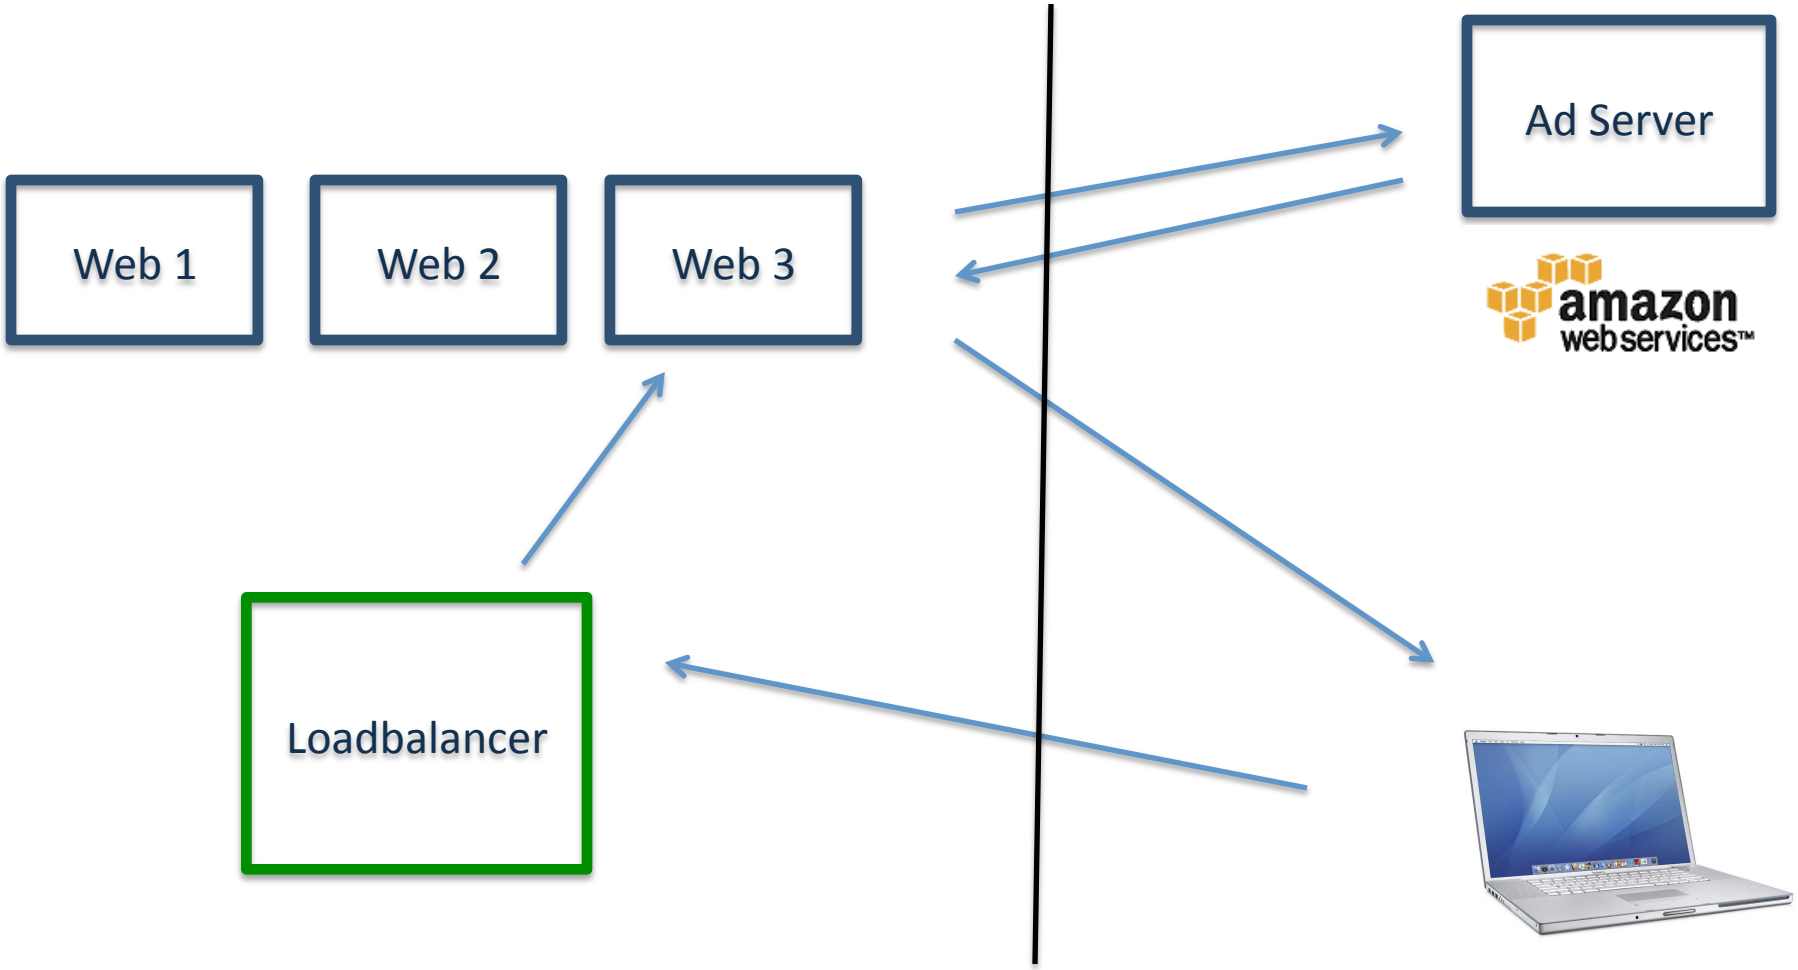

Web Services: Scale

DB 1

DB 2

Mail 1

Web 1

Web 2

Web 3

Loadbalancer

Ad Server

amazon web services™

User

Web Services: Geo-scale

Facility	Hostname	Test Download
London, UK	london1.linode.com	100MB-london.bin
Newark, NJ	newark1.linode.com	100MB-newark.bin
Atlanta, GA	atlanta1.linode.com	100MB-atlanta.bin
Dallas, TX	dallas1.linode.com	100MB-dallas.bin
Fremont, CA	fremont1.linode.com	100MB-fremont.bin

Web Services: Geo-scale

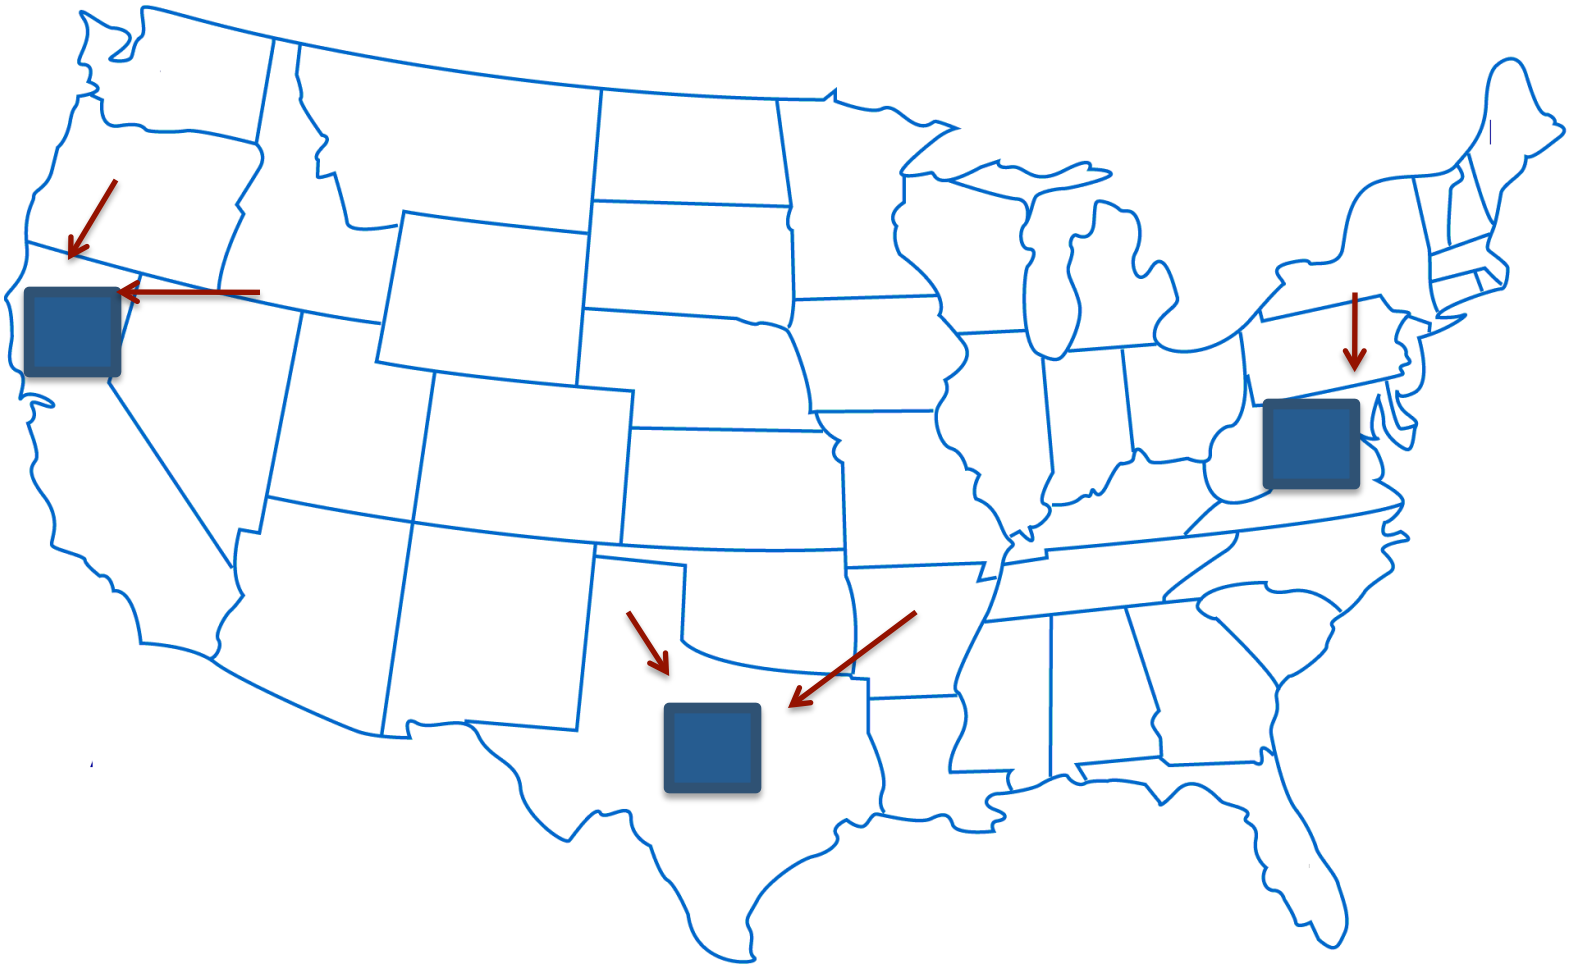

Online Ads: Architecture

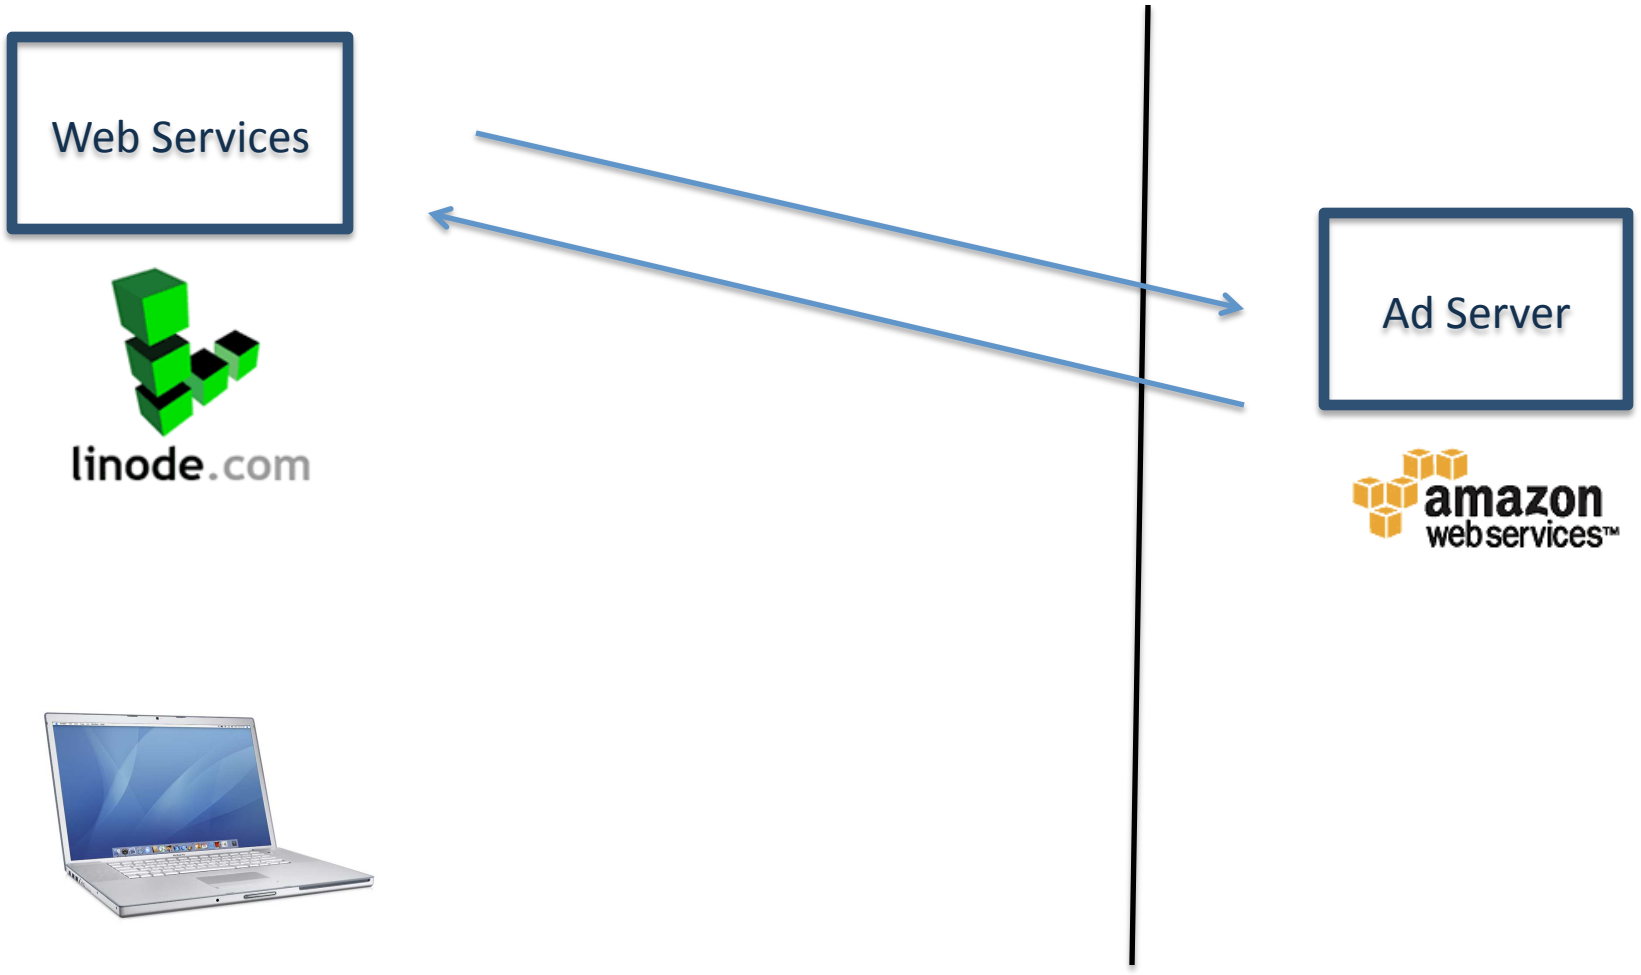

Online Ads: Architecture

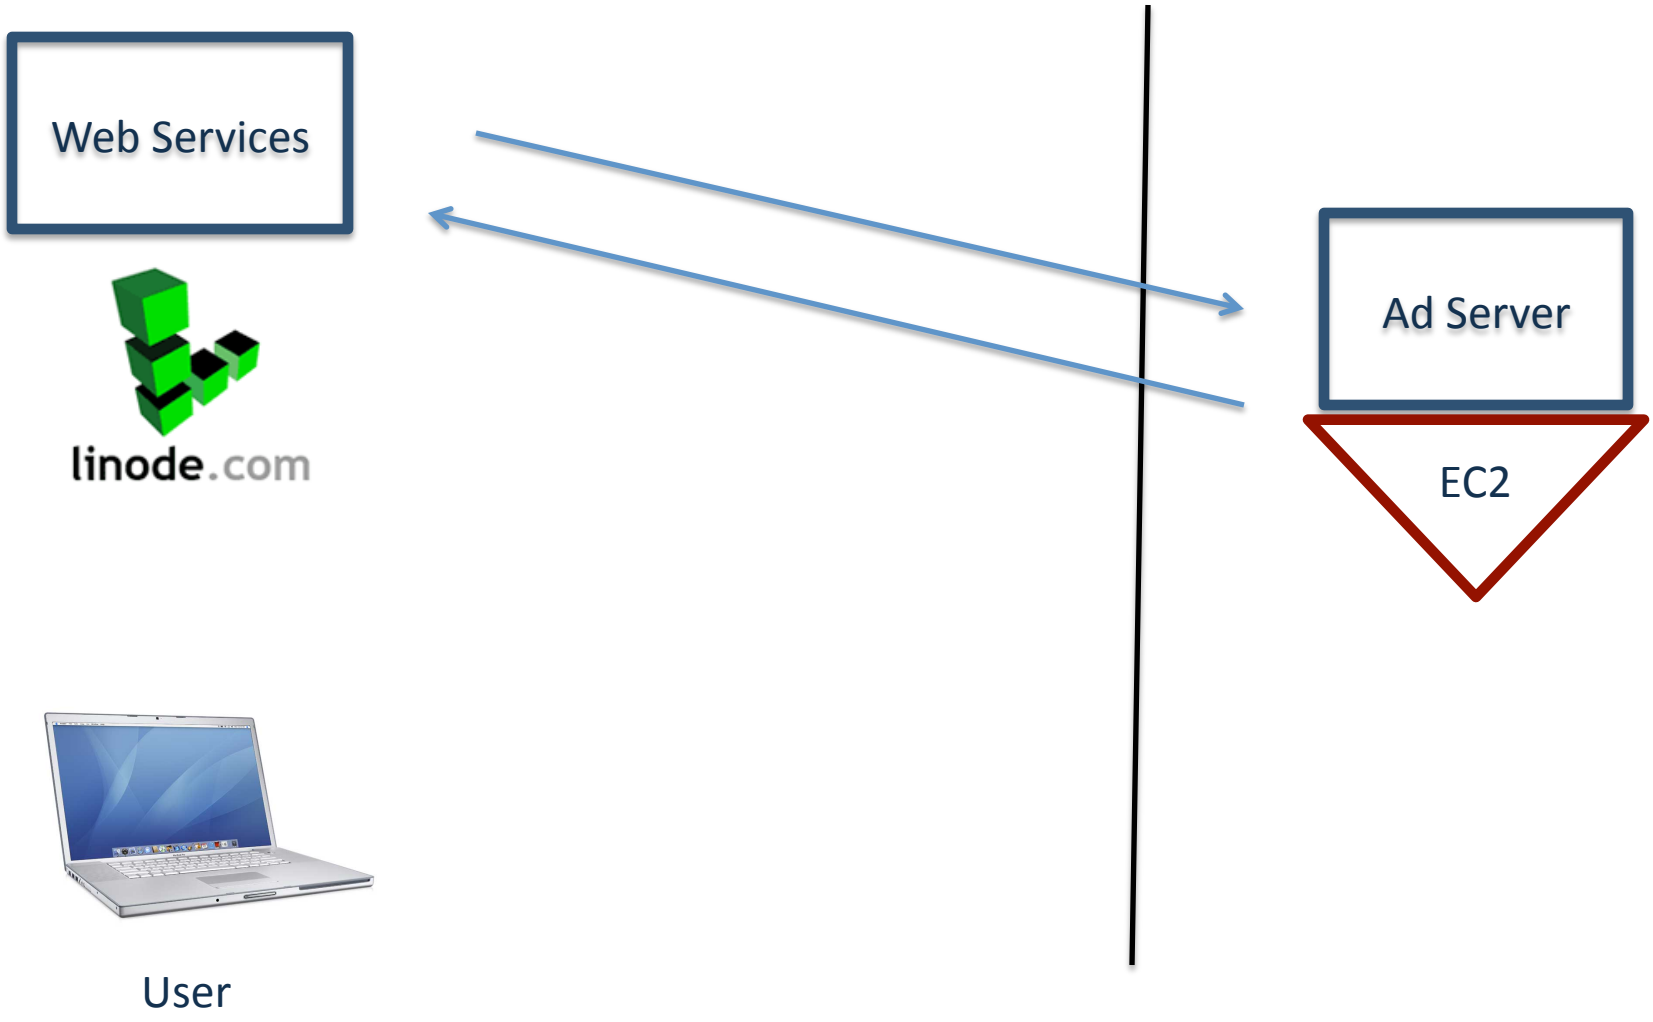

Online Ads: Scale

Web Services

User

Ad 1

Ad 2

EC2

EC2

Online Ads: Scale

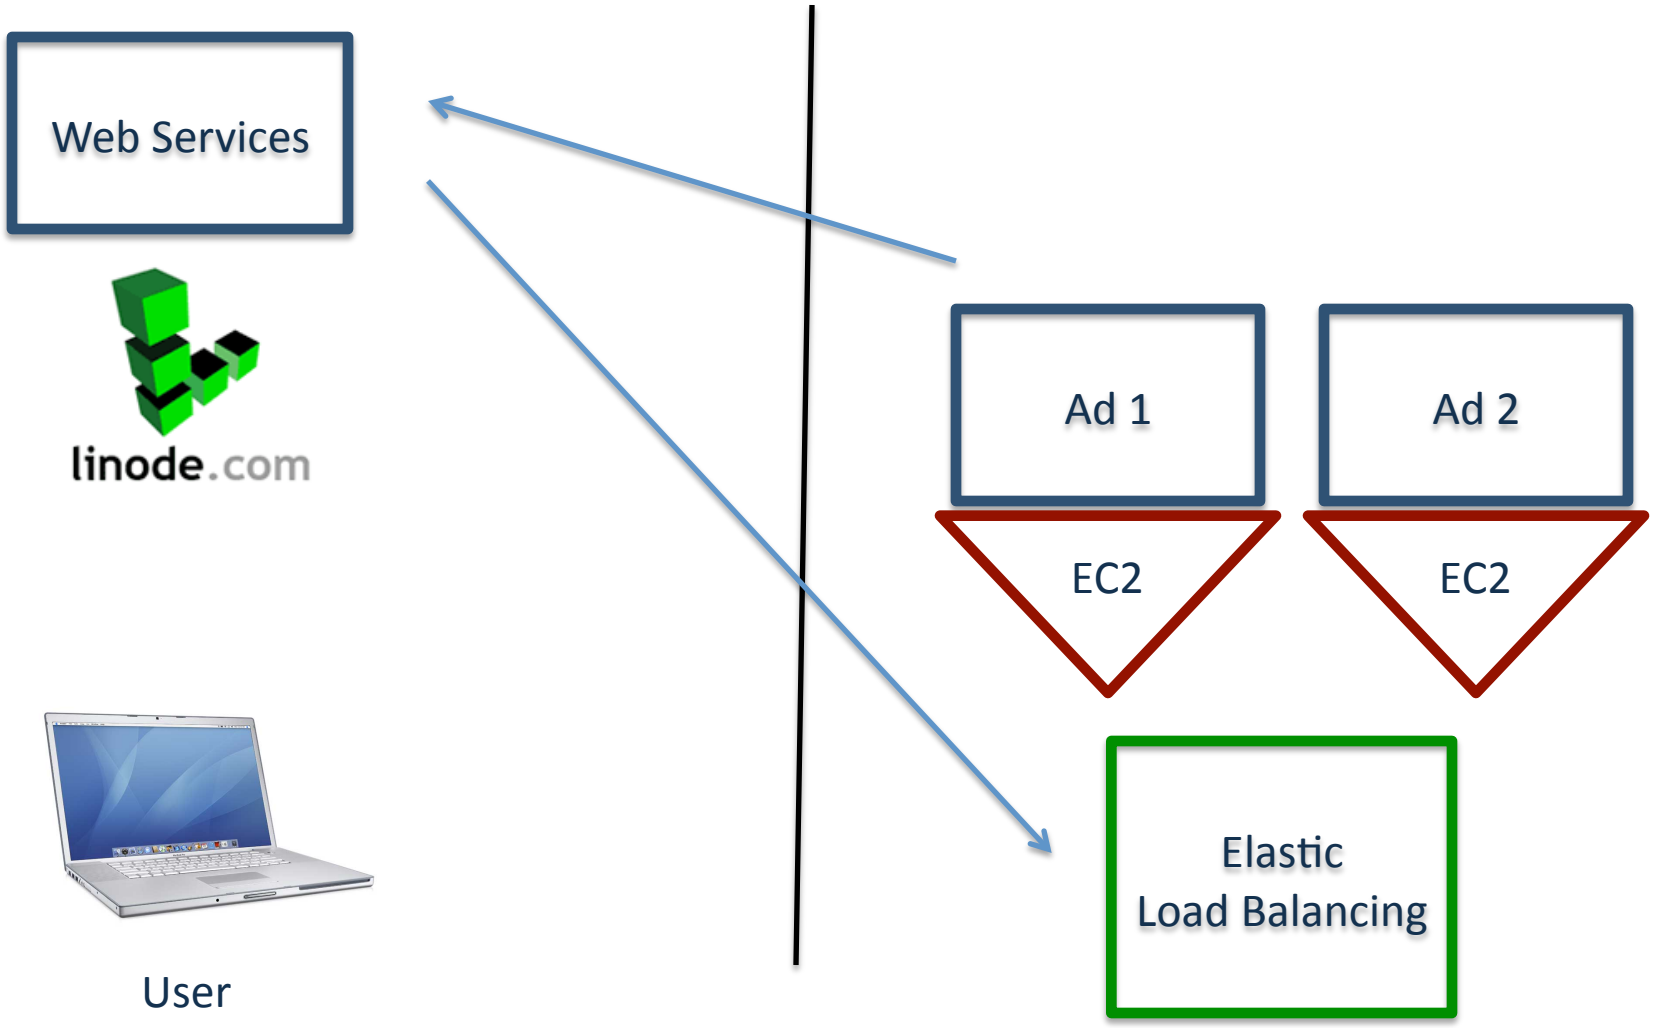

Online Ads: Scale

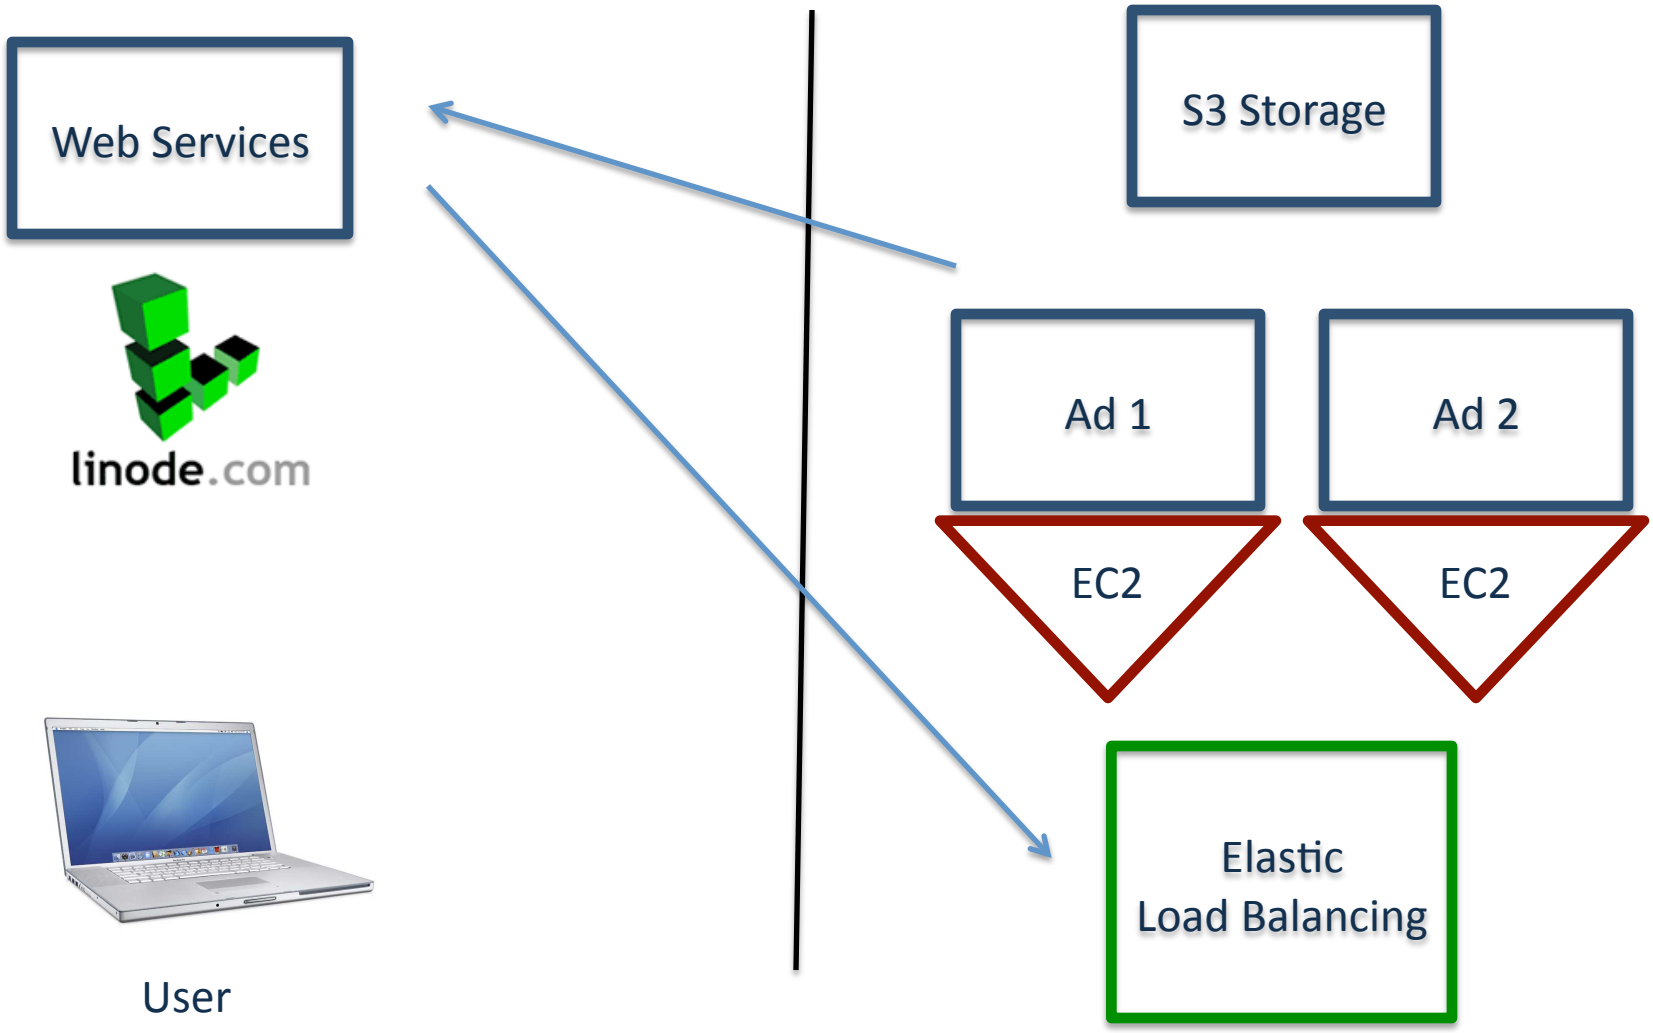

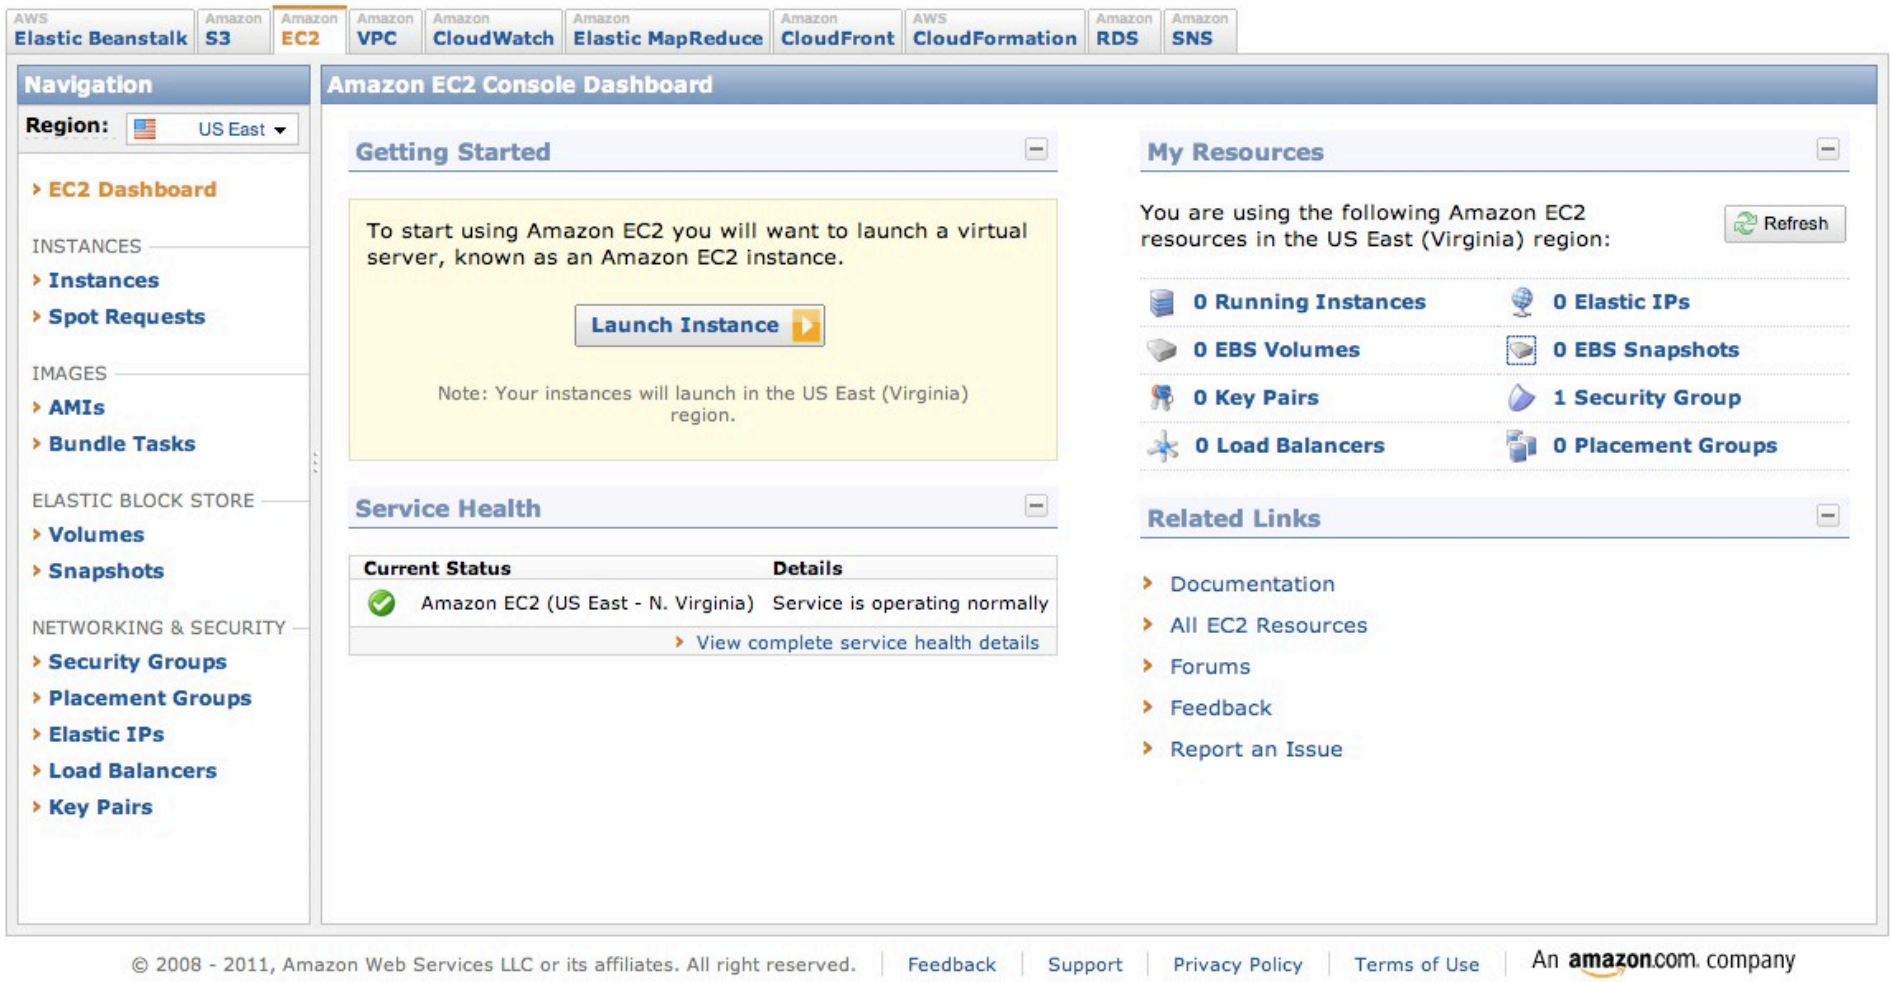
AWS Elastic Beanstalk
Amazon S3
Amazon EC2
Amazon VPC
Amazon CloudWatch
Amazon Elastic MapReduce
Amazon CloudFront
AWS CloudFormation
Amazon RDS
Amazon SNS
Navigation
Region: US East
EC2 Dashboard
INSTANCES
Instances
Spot Requests
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORKING & SECURITY
Security Groups
Placement Groups
Elastic IPs
Load Balancers
Key Pairs
Amazon EC2 Console Dashboard
Getting Started
To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.
Launch Instance
Note: Your instances will launch in the US East (Virginia) region.
Service Health
Current Status
Details
Amazon EC2 (US East - N. Virginia)
Service is operating normally
View complete service health details
My Resources
You are using the following Amazon EC2 resources in the US East (Virginia) region:
Refresh
0 Running Instances
0 Elastic IPs
0 EBS Volumes
0 EBS Snapshots
0 Key Pairs
1 Security Group
0 Load Balancers
0 Placement Groups
Related Links
Documentation
All EC2 Resources
Forums
Feedback
Report an Issue
© 2008 - 2011, Amazon Web Services LLC or its affiliates. All right reserved.
Feedback
Support
Privacy Policy
Terms of Use
An amazon.com company

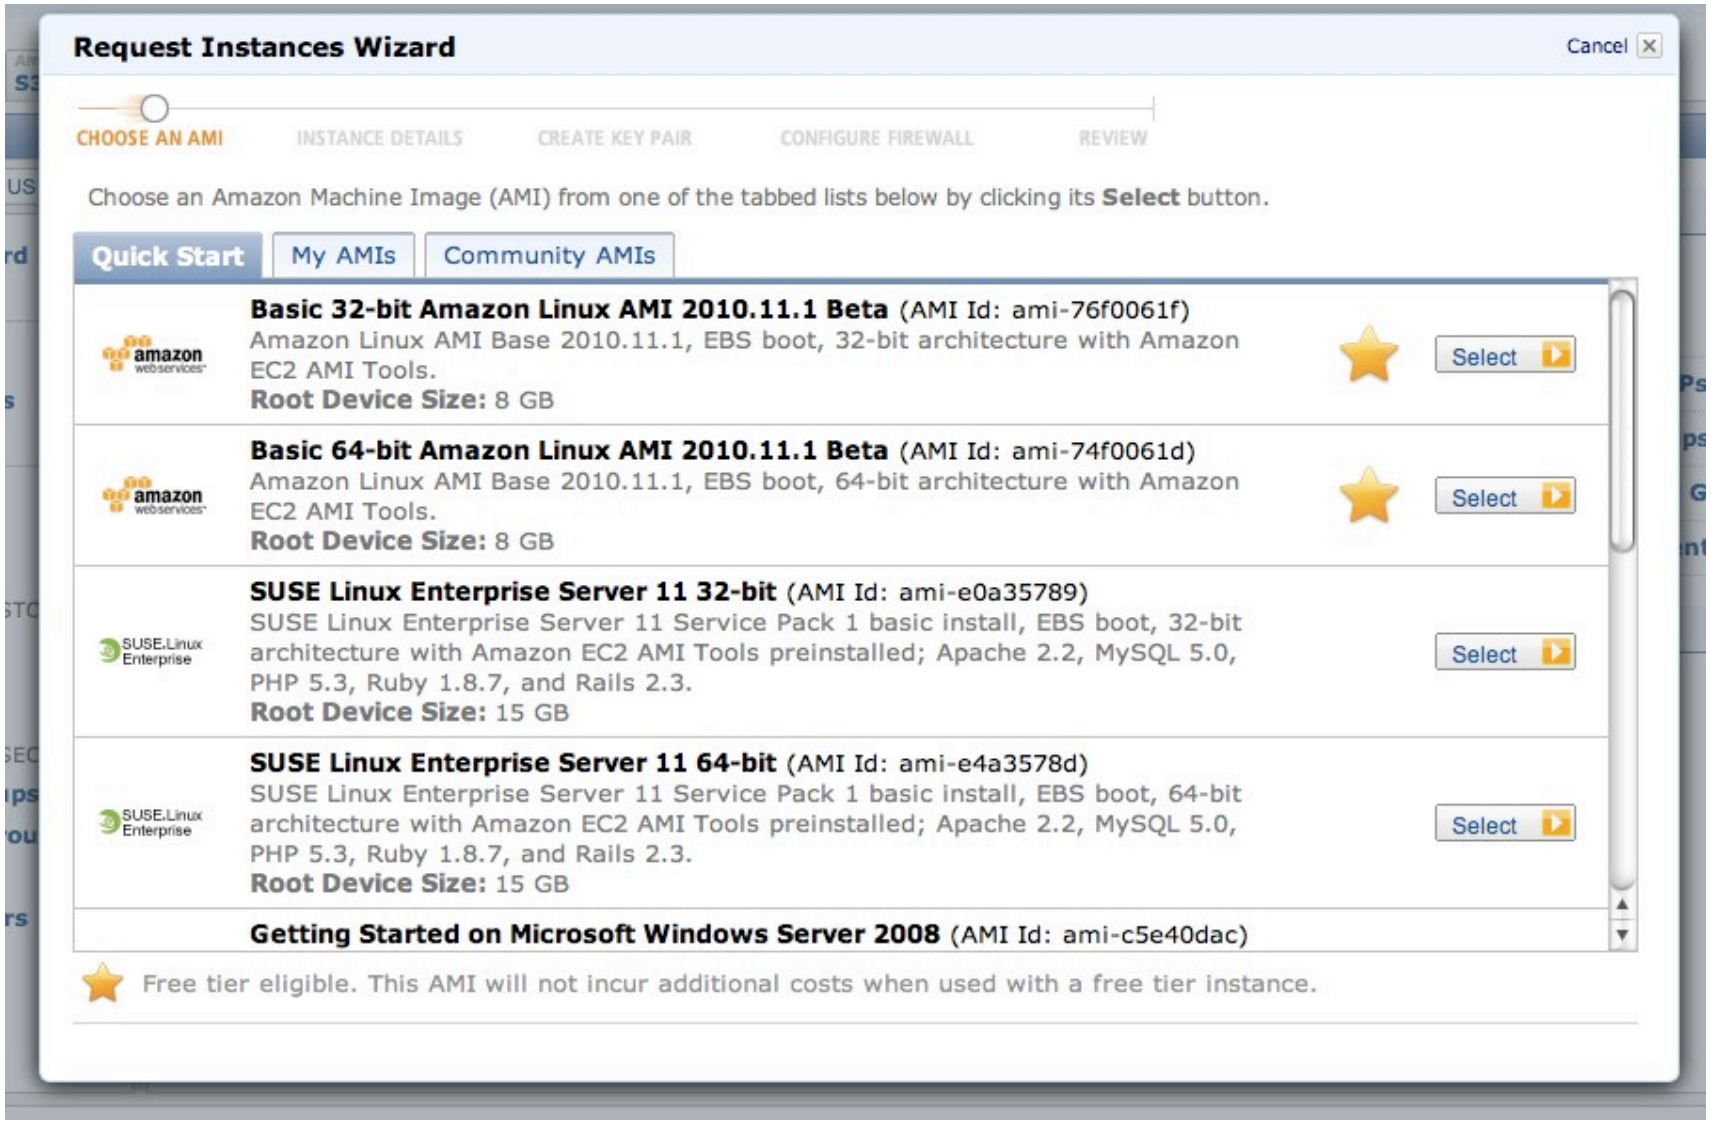
Request Instances Wizard
Cancel
CHOOSE AN AMI
INSTANCE DETAILS
CREATE KEY PAIR
CONFIGURE FIREWALL
REVIEW
Choose an Amazon Machine Image (AMI) from one of the tabbed lists below by clicking its Select button.
Quick Start
My AMIs
Community AMIs
Basic 32-bit Amazon Linux AMI 2010.11.1 Beta (AMI Id: ami-76f0061f)
Amazon Linux AMI Base 2010.11.1, EBS boot, 32-bit architecture with Amazon EC2 AMI Tools.
Root Device Size: 8 GB
Select
Basic 64-bit Amazon Linux AMI 2010.11.1 Beta (AMI Id: ami-74f0061d)
Amazon Linux AMI Base 2010.11.1, EBS boot, 64-bit architecture with Amazon EC2 AMI Tools.
Root Device Size: 8 GB
Select
SUSE Linux Enterprise Server 11 32-bit (AMI Id: ami-e0a35789)
SUSE Linux Enterprise Server 11 Service Pack 1 basic install, EBS boot, 32-bit architecture with Amazon EC2 AMI Tools preinstalled; Apache 2.2, MySQL 5.0, PHP 5.3, Ruby 1.8.7, and Rails 2.3.
Root Device Size: 15 GB
Select
SUSE Linux Enterprise Server 11 64-bit (AMI Id: ami-e4a3578d)
SUSE Linux Enterprise Server 11 Service Pack 1 basic install, EBS boot, 64-bit architecture with Amazon EC2 AMI Tools preinstalled; Apache 2.2, MySQL 5.0, PHP 5.3, Ruby 1.8.7, and Rails 2.3.
Root Device Size: 15 GB
Select
Getting Started on Microsoft Windows Server 2008 (AMI Id: ami-c5e40dac)
Free tier eligible. This AMI will not incur additional costs when used with a free tier instance.

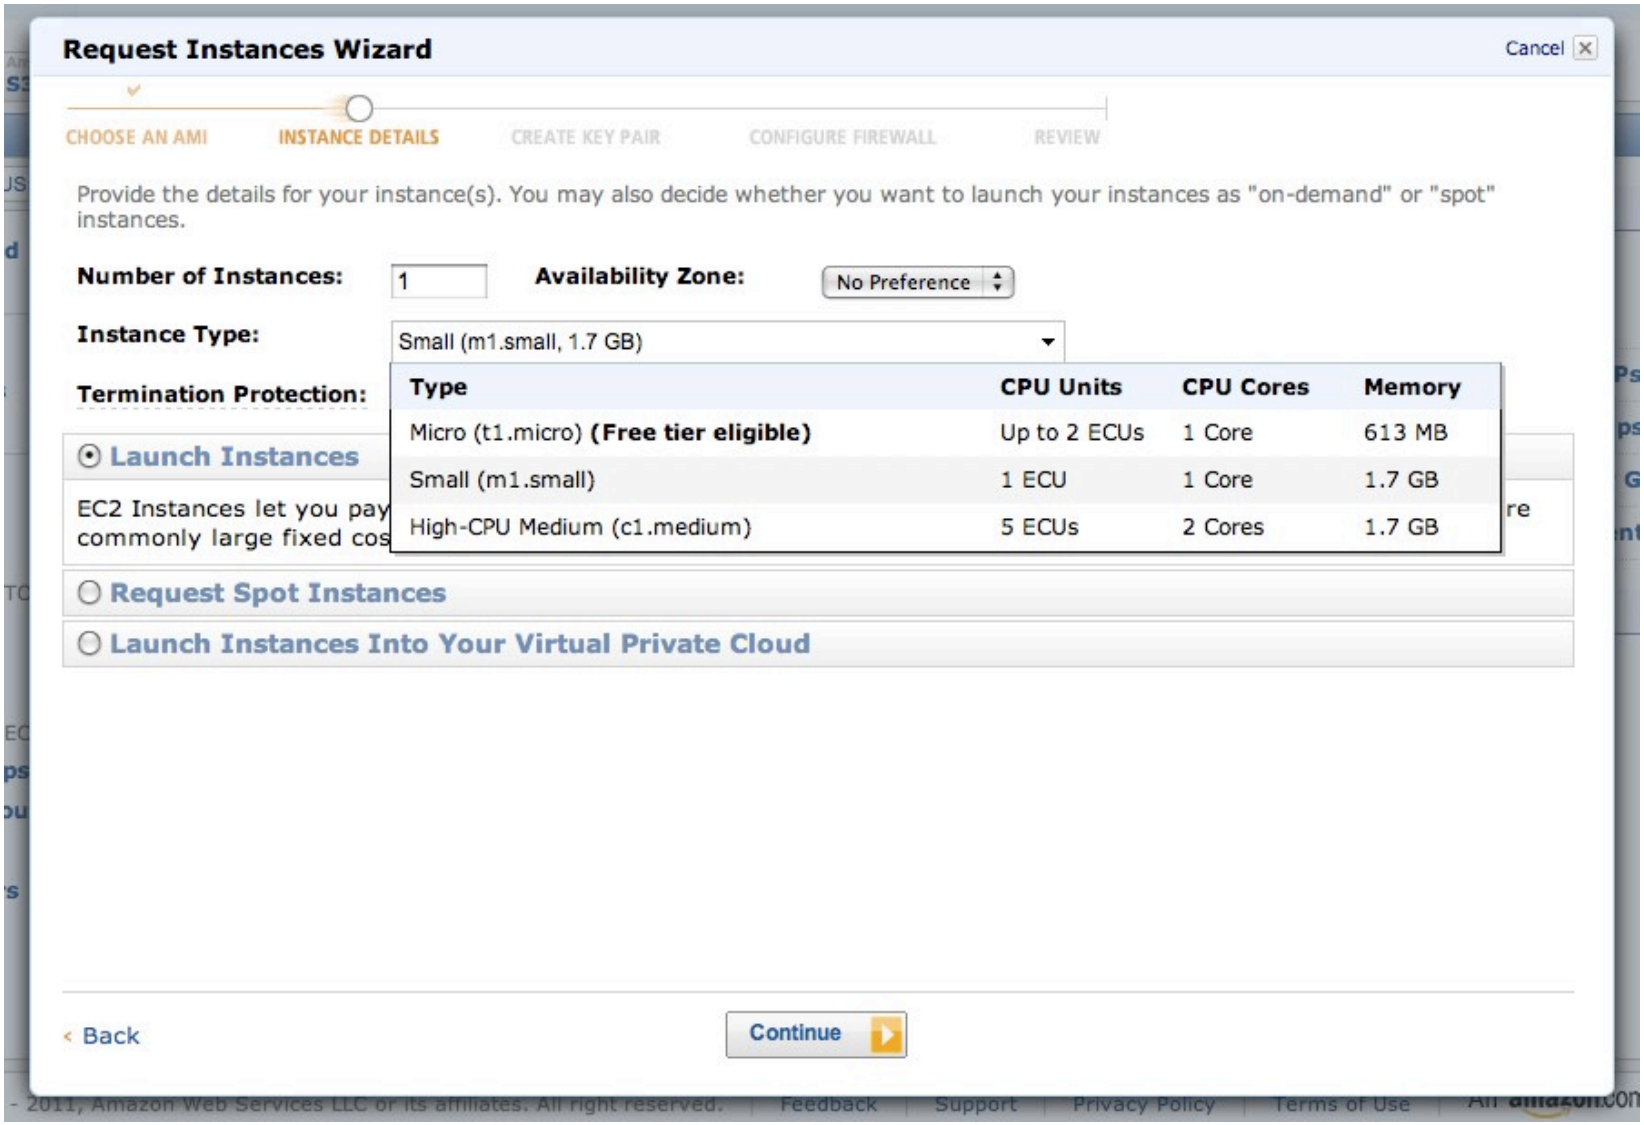
Request Instances Wizard
Cancel
CHOOSE AN AMI
INSTANCE DETAILS
CREATE KEY PAIR
CONFIGURE FIREWALL
REVIEW
Provide the details for your instance(s). You may also decide whether you want to launch your instances as "on-demand" or "spot" instances.
Number of Instances:
1
Availability Zone:
No Preference
Instance Type:
Small (m1.small, 1.7 GB)
Termination Protection:
Type
CPU Units
CPU Cores
Memory
Micro (t1.micro) (Free tier eligible)
Up to 2 ECUs
1 Core
613 MB
Small (m1.small)
1 ECU
1 Core
1.7 GB
High-CPU Medium (c1.medium)
5 ECUs
2 Cores
1.7 GB
Launch Instances
EC2 Instances let you pay
commonly large fixed cos
re
Request Spot Instances
Launch Instances Into Your Virtual Private Cloud
Back
Continue
- 2011, Amazon Web Services LLC or its affiliates. All right reserved.
Feedback
Support
Privacy Policy
Terms of Use

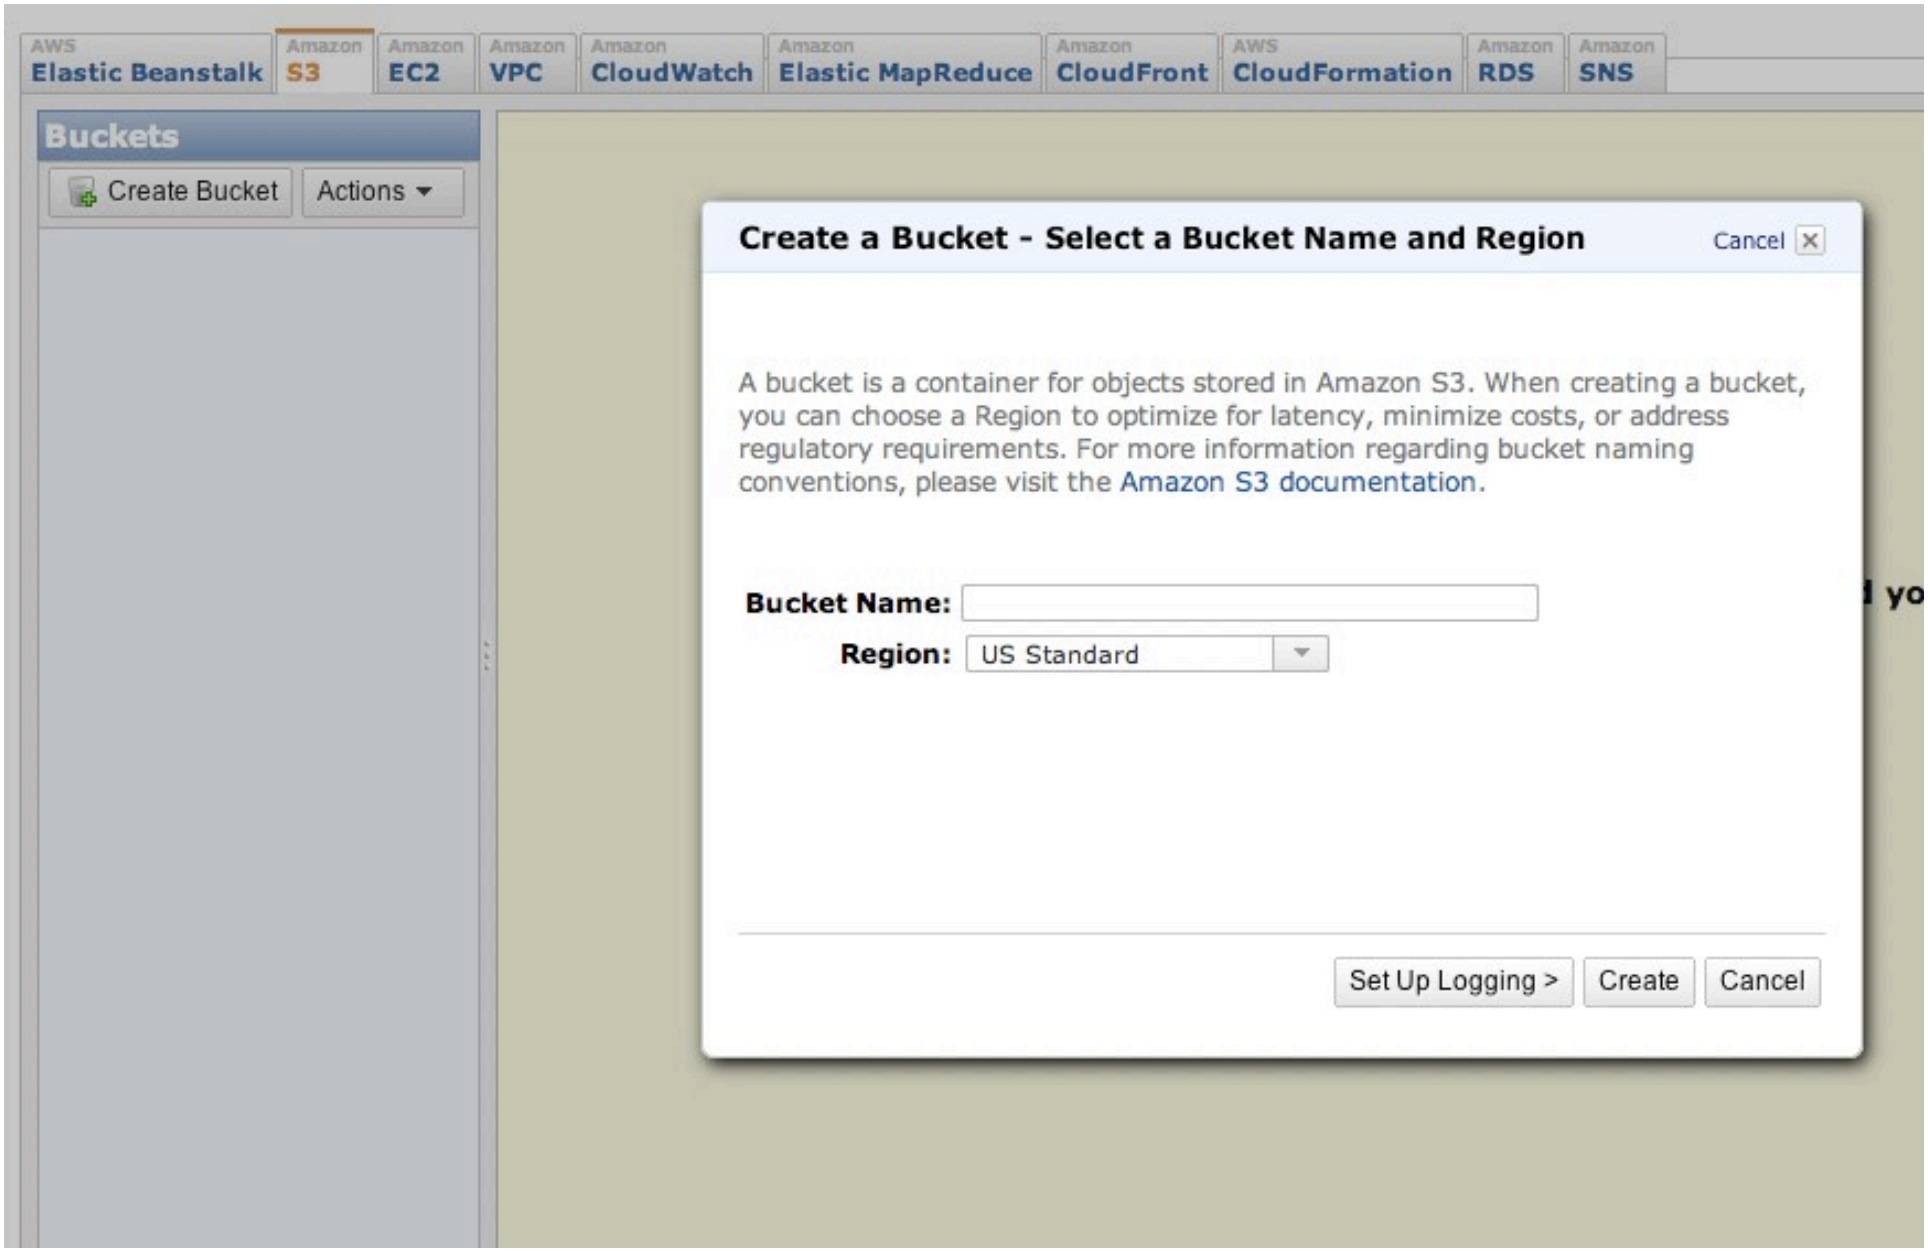
AWS Elastic Beanstalk
Amazon S3
Amazon EC2
Amazon VPC
Amazon CloudWatch
Amazon Elastic MapReduce
Amazon CloudFront
AWS CloudFormation
Amazon RDS
Amazon SNS
Buckets
Create Bucket
Actions
Create a Bucket - Select a Bucket Name and Region
Cancel
A bucket is a container for objects stored in Amazon S3. When creating a bucket, you can choose a Region to optimize for latency, minimize costs, or address regulatory requirements. For more information regarding bucket naming conventions, please visit the Amazon S3 documentation.
Bucket Name:
Region:
US Standard
Set Up Logging >
Create
Cancel

Amazon: SimpleDB

Build Data Set

- Choose a region for "domain"
- Currently Virginia, California, Ireland, and Singapore
- **CreateDomain**, **DeleteDomain**, etc API to manage domains
- **Put**, **Batch Put**, and **Delete** to manage data inside domain

Retrieve Data

- Use **GetAttributes** to retrieve specific item
- Use **Select** to query data that meets specific criteria

Pay only for what you consume

Difference from S3

- Database (structured data) instead of raw data
- E.g., an image goes into S3, but list of users/logins go to SimpleDB

Amazon: Storing Data

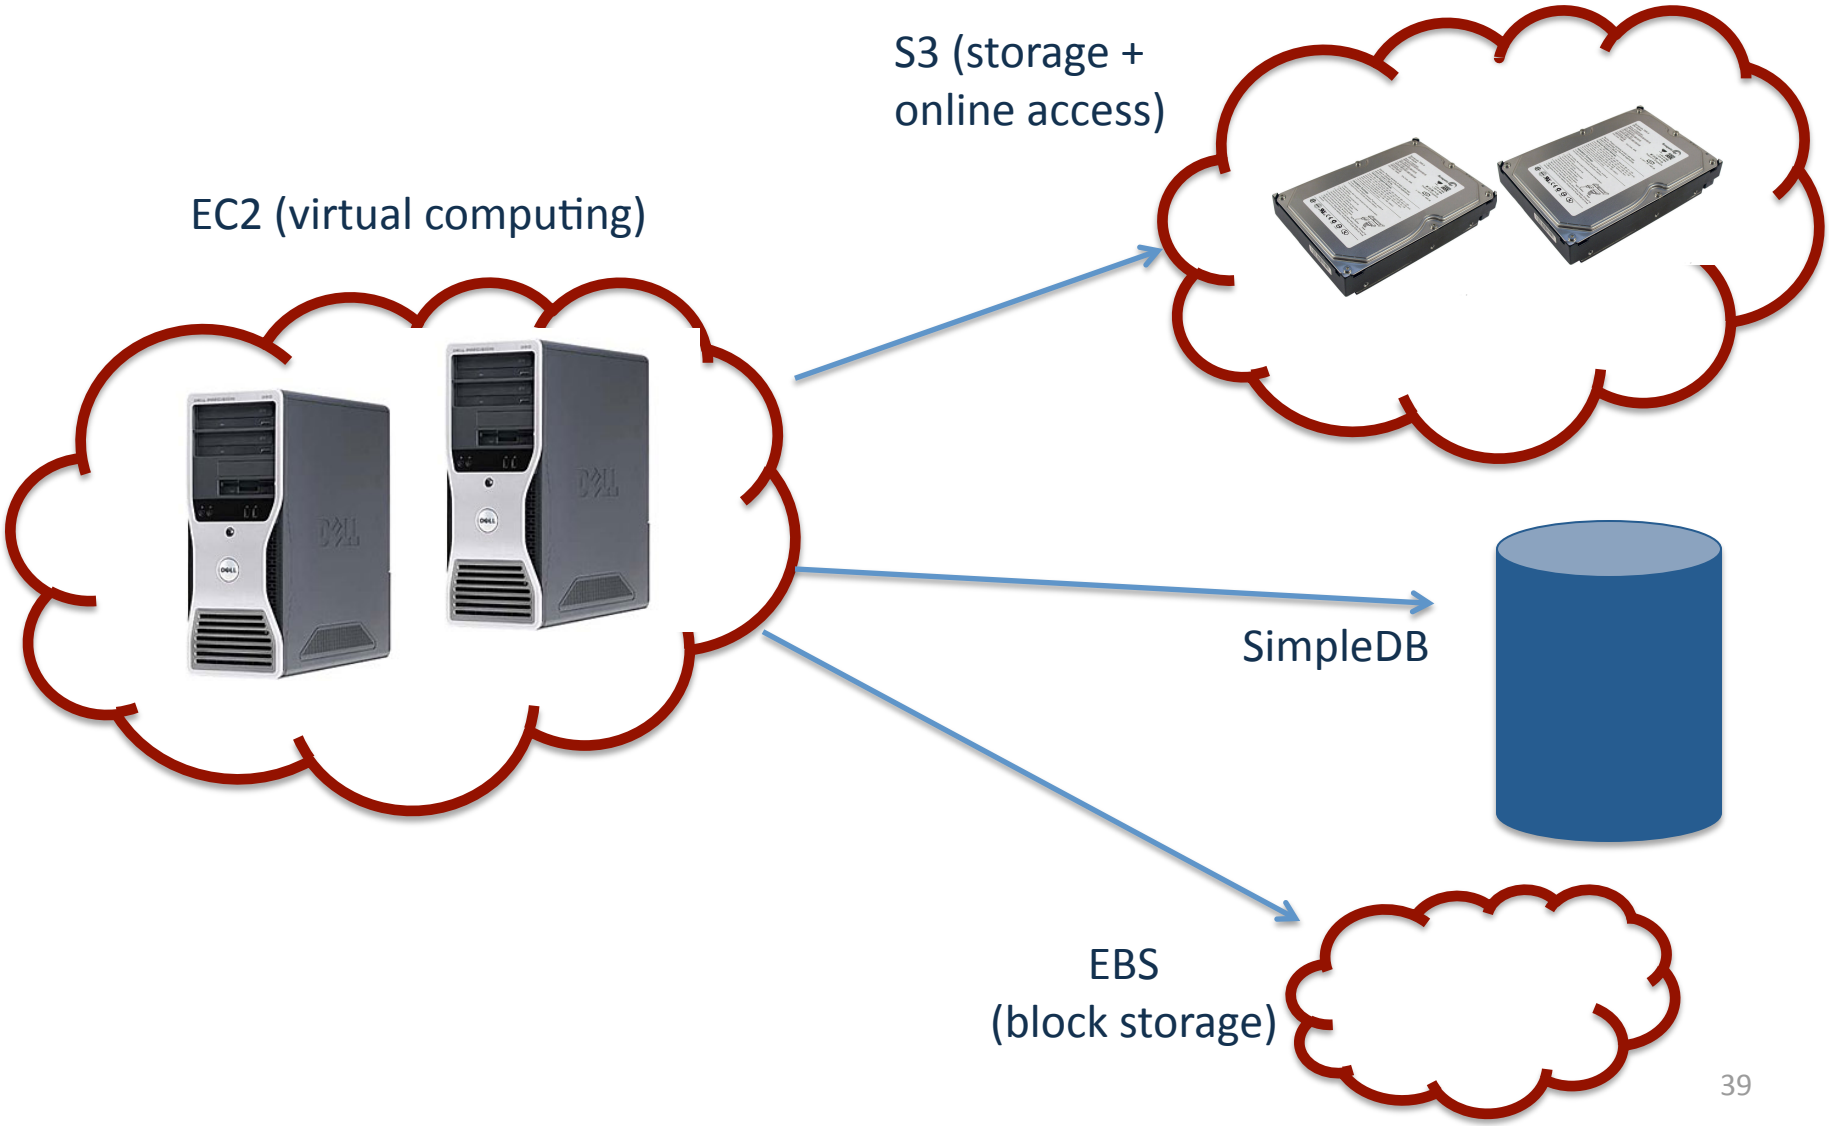

Big Picture

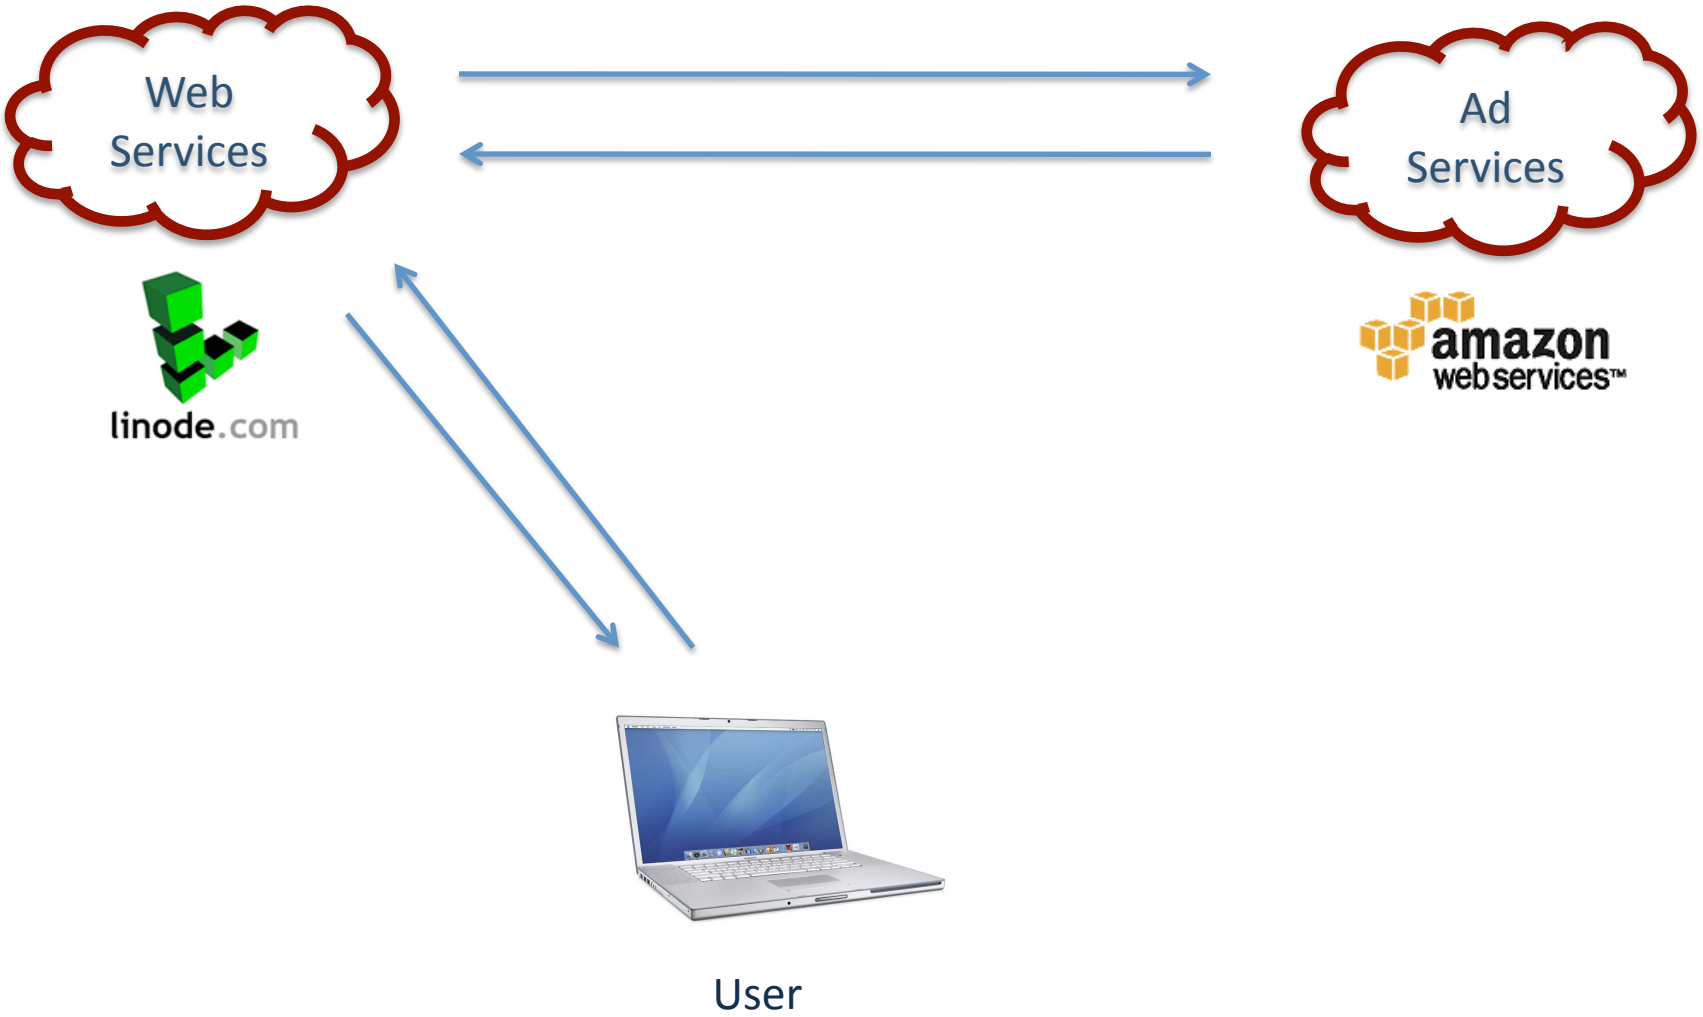
Web Services
Ad Services
linode.com
amazon web services™
User

High Line Times: CDN

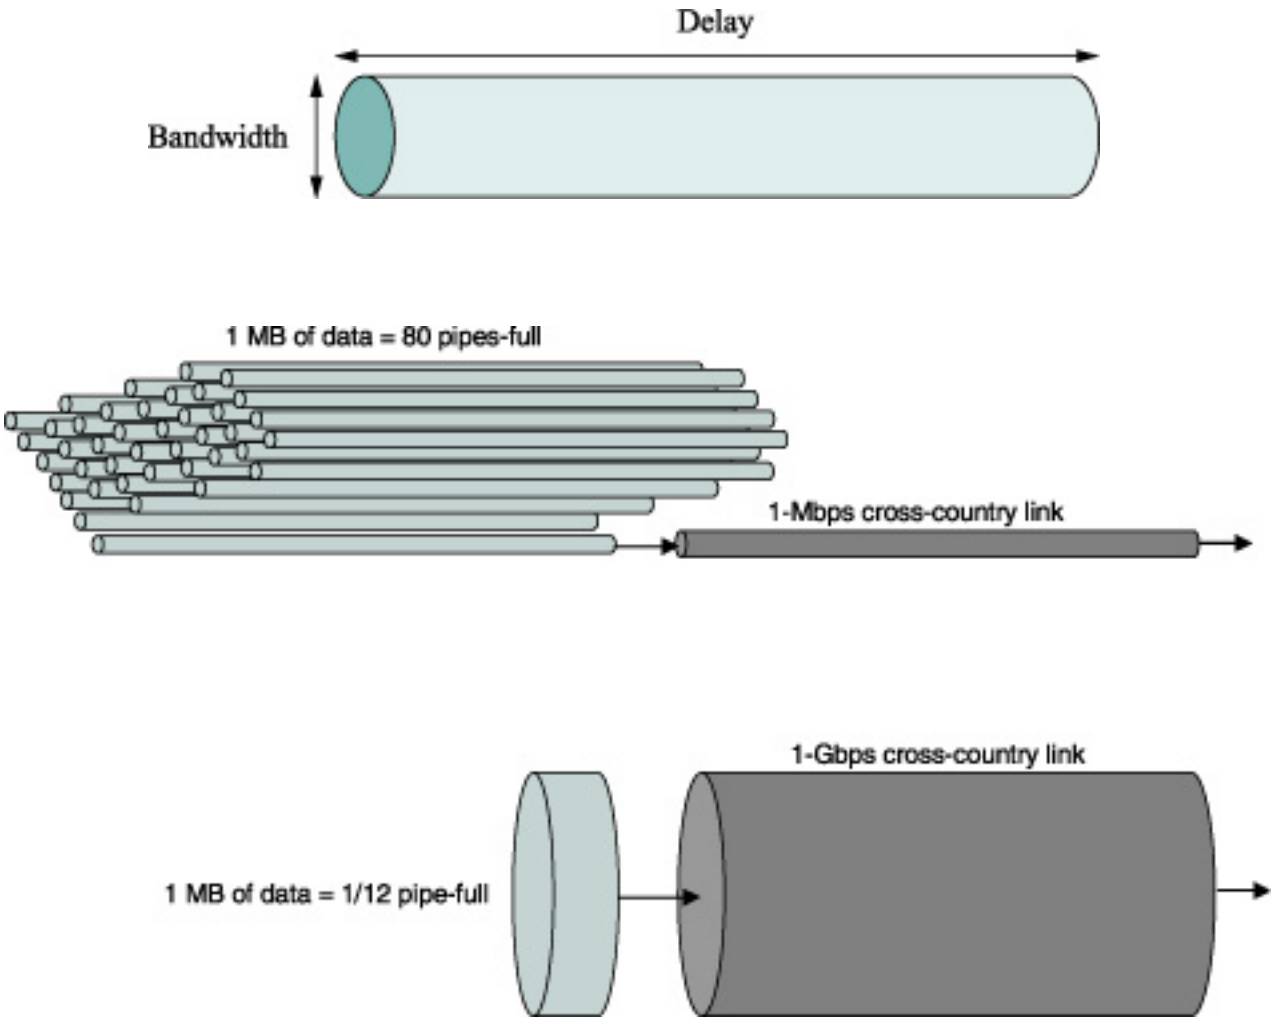

High Line Times: CDN

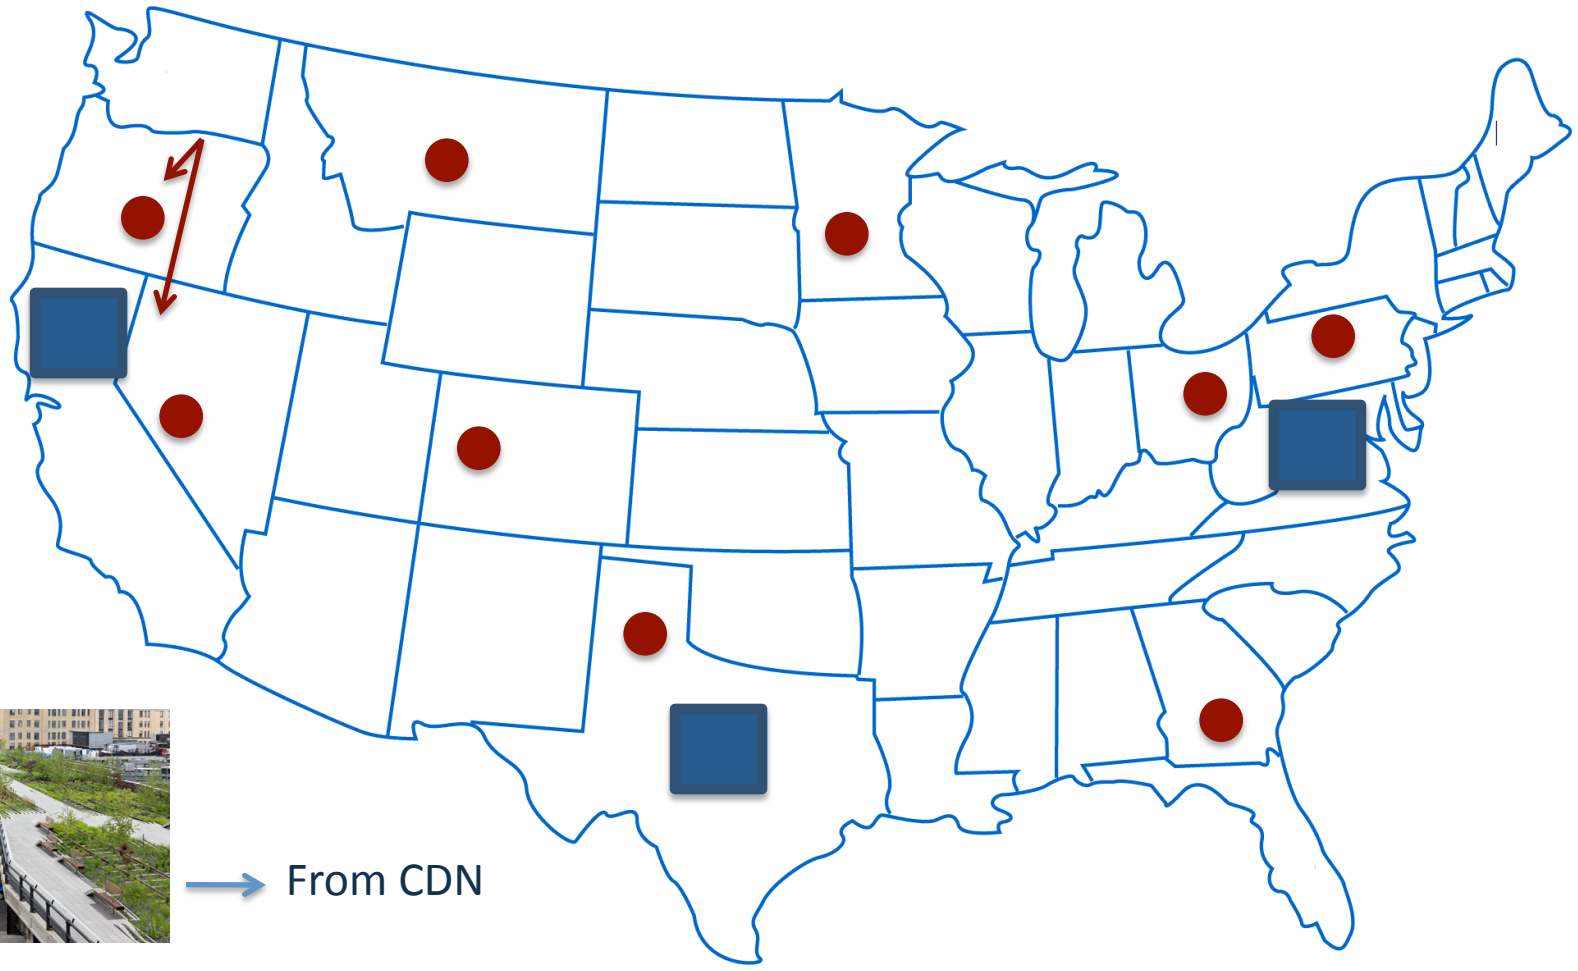

→ From CDN

HTML/TXT → From Datacenter

High Line Times

100 Million Users

- Build own datacenter?

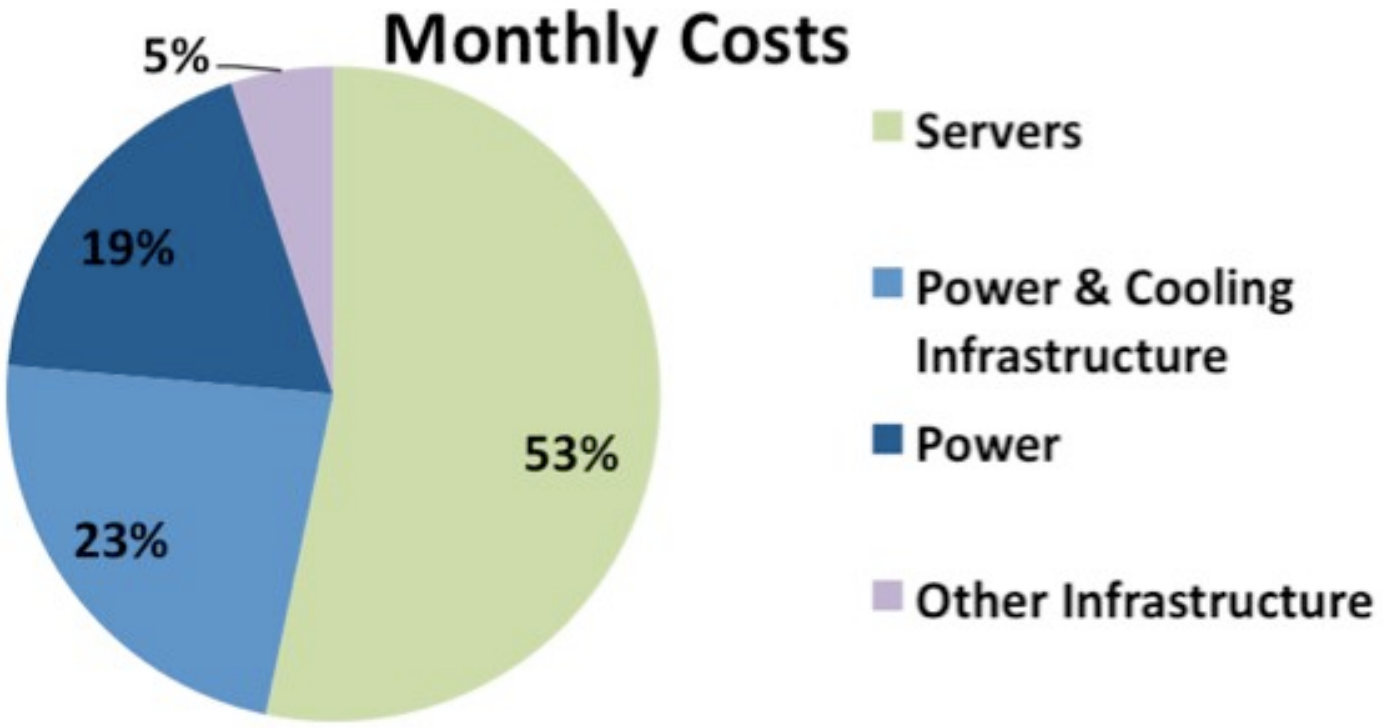

3 Year Server & 15 Year Infrastructure Amortization

With 50k servers, $2 million/month electricity bill

Road Map: Today's Lecture

Scalable Systems

- Example company; grows from 10 users to 100 million users
- Ad serving company spin out
- How both managed growth in a different way
- The design tradeoffs made along the way

A Brief History of Cloud Computing

- A quick recap of technological progress in the last decade (Google)

Future of Cloud Computing

- Where are we heading …

1999 to 2009: Google

1999 to 2009: Google

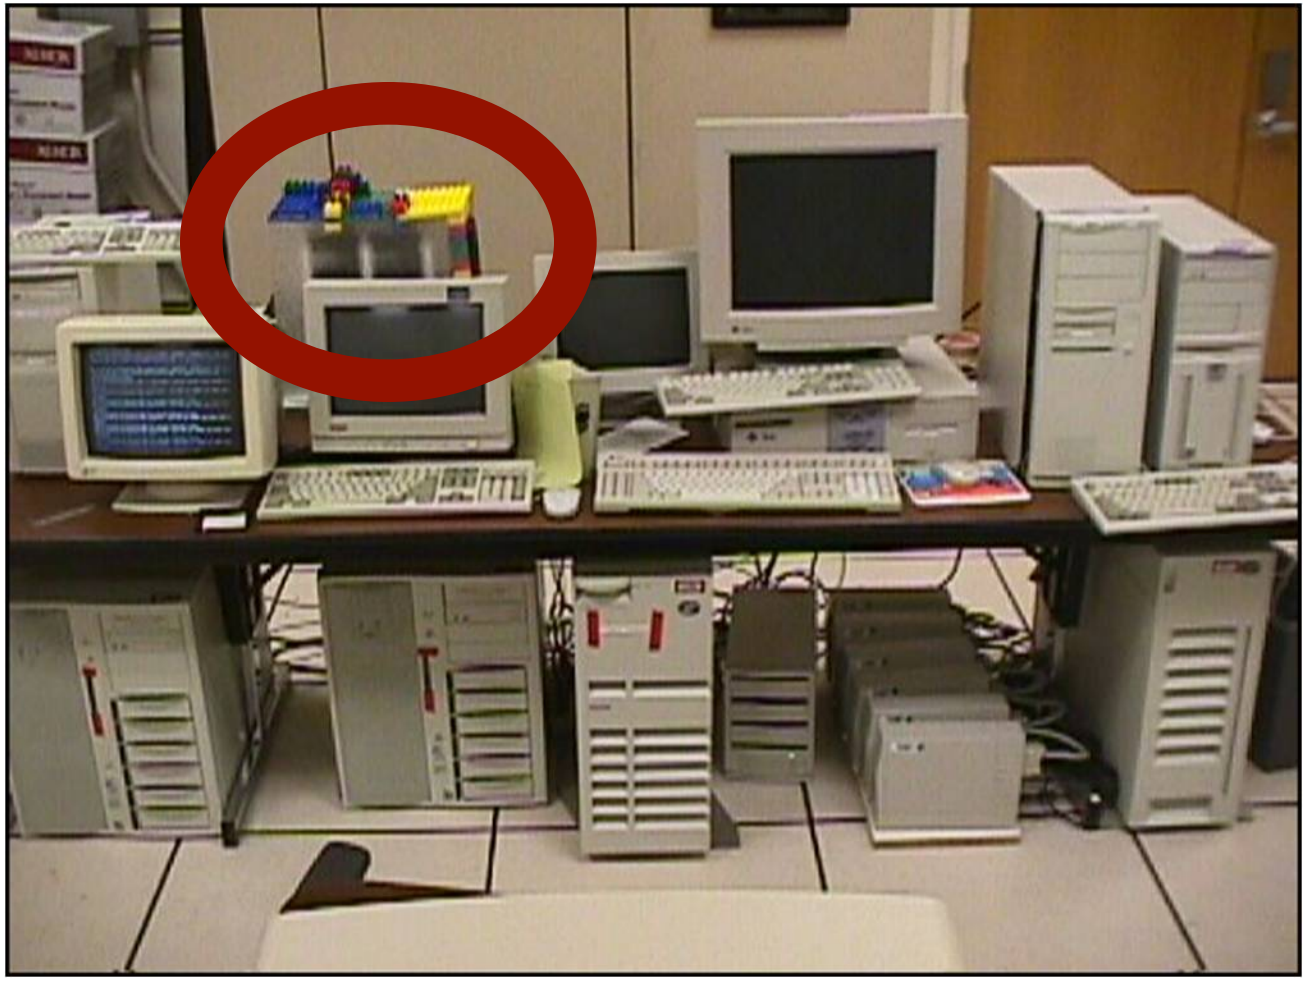

Google servers 1997

(Images courtesy Jeff Dean)

1999 to 2009: Google

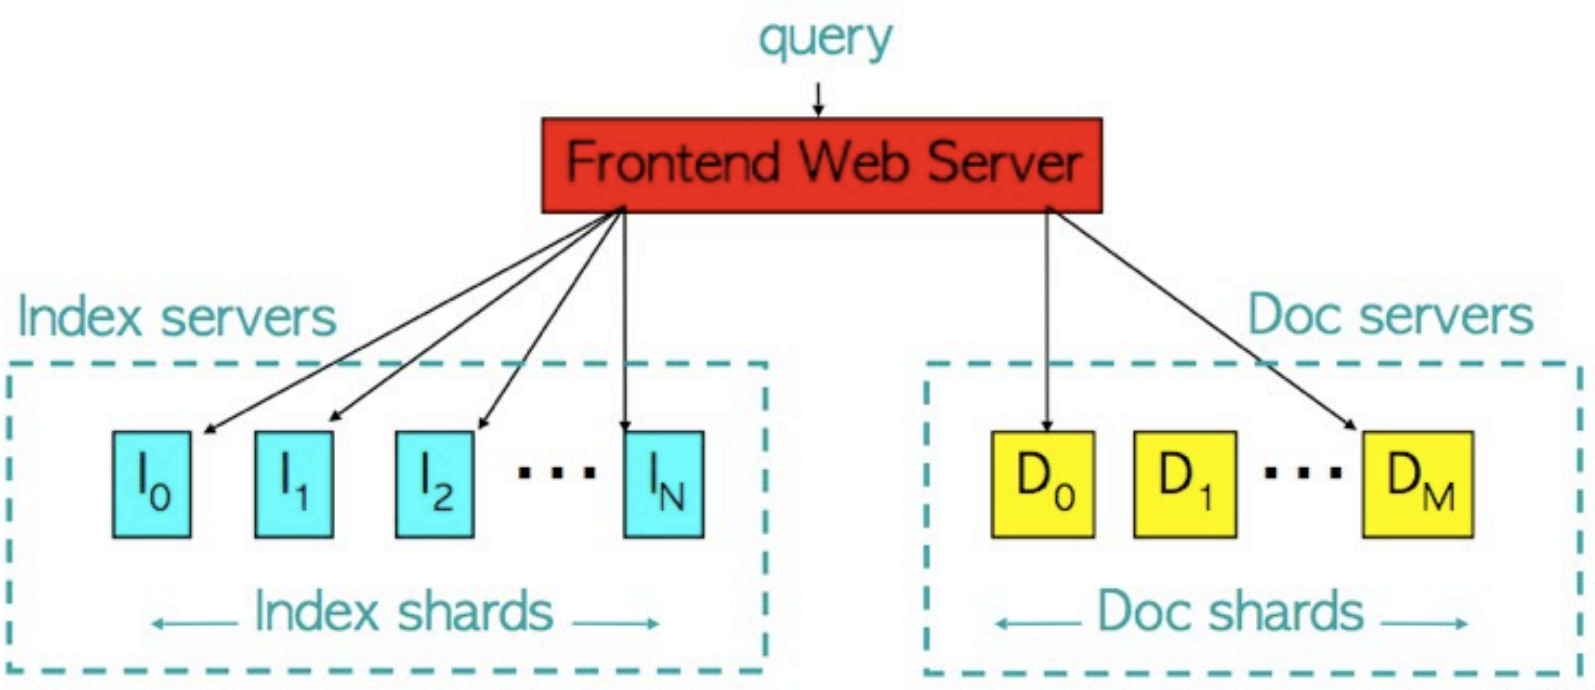

Google architecture 1997

(Images courtesy Jeff Dean)

1999 to 2009: Google

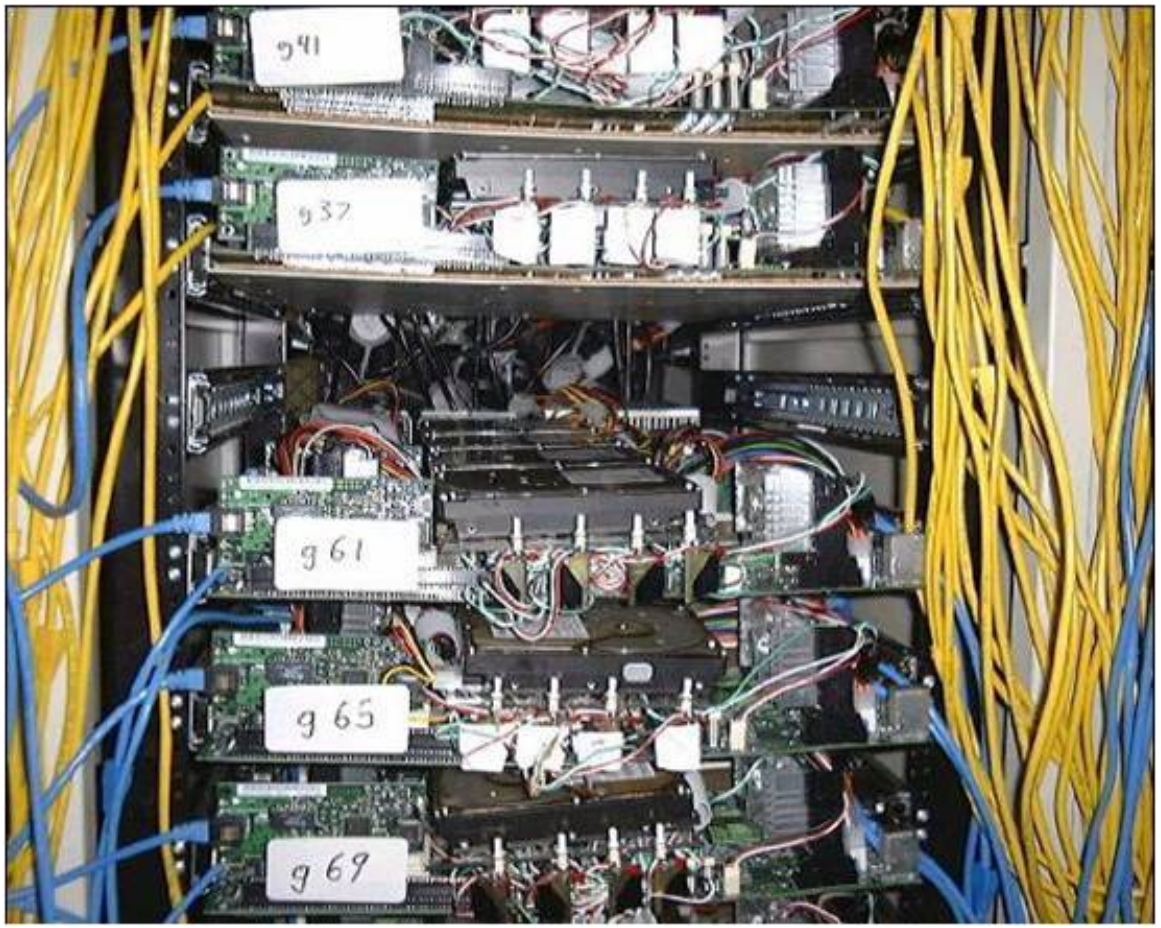

Google servers 1999

(Images courtesy Jeff Dean)

1999 to 2009: Google

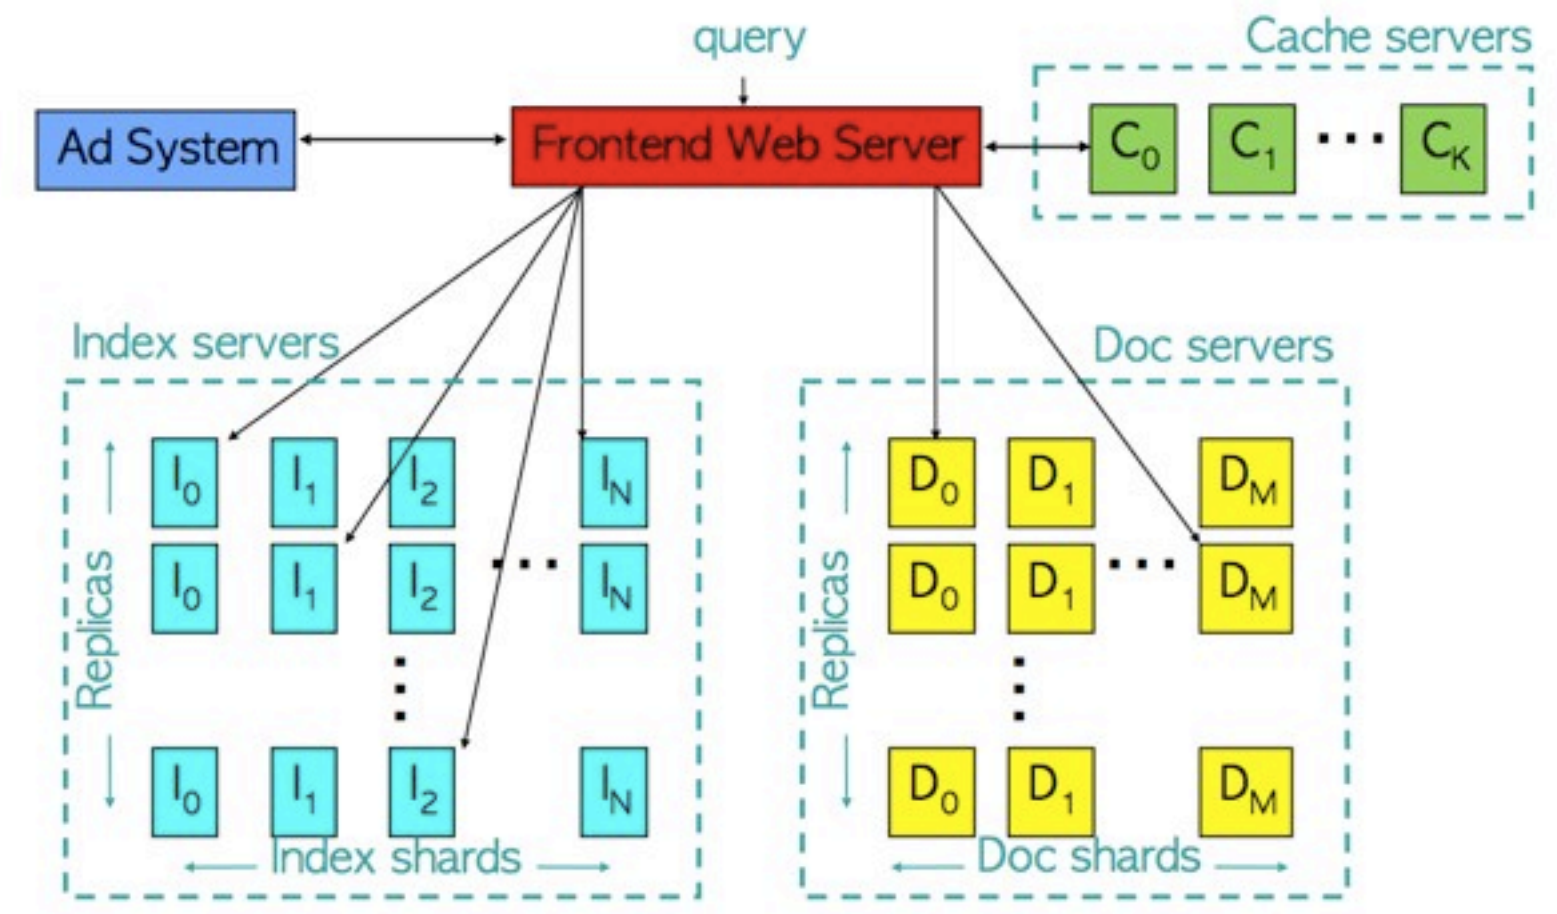

Google architecture 1999

(Images courtesy Jeff Dean)

1999 to 2009: Google

Google servers 2000

(Images courtesy Jeff Dean)

1999 to 2009: Google

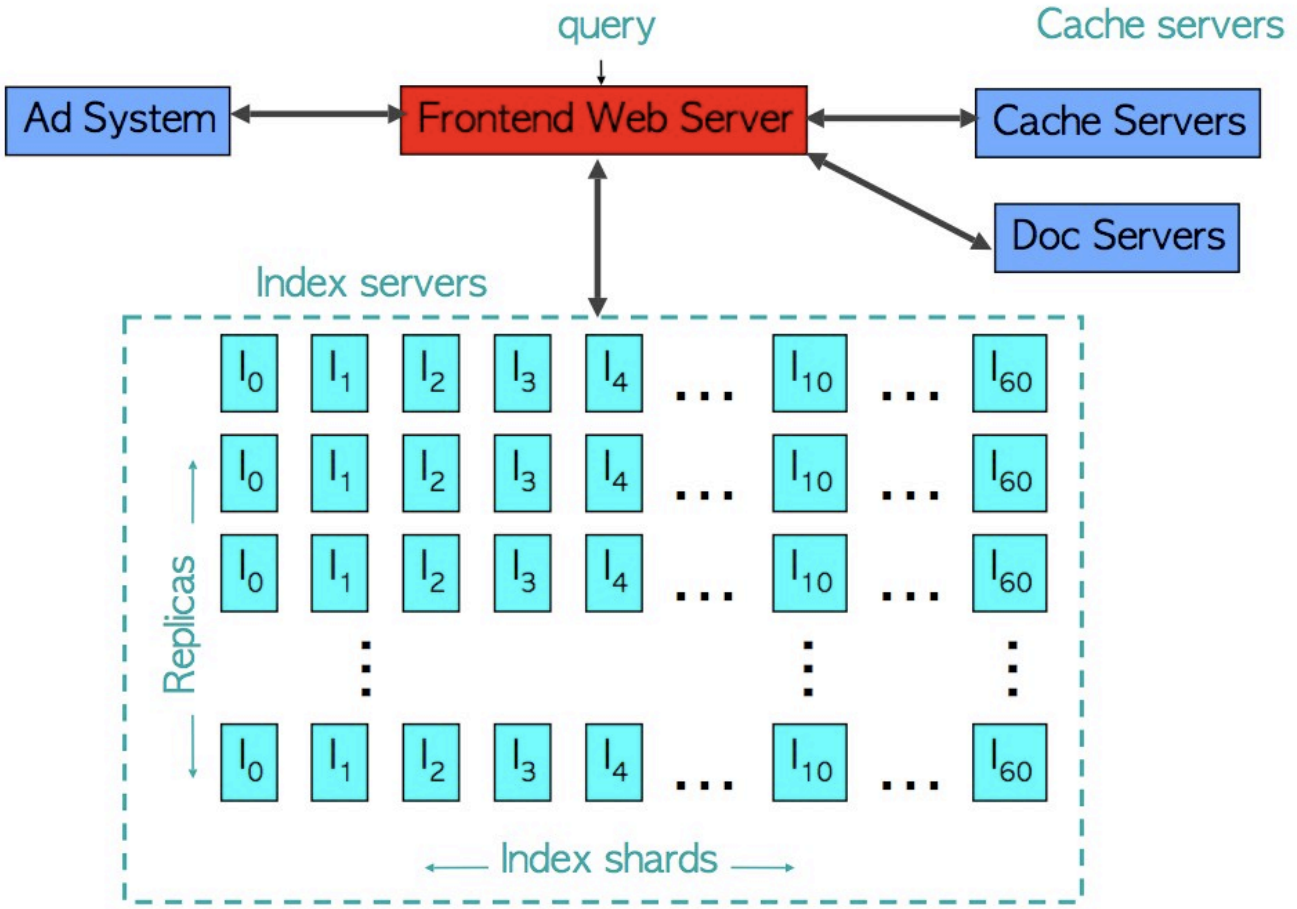

Google architecture 2000

(Images courtesy Jeff Dean)

1999 to 2009: Google

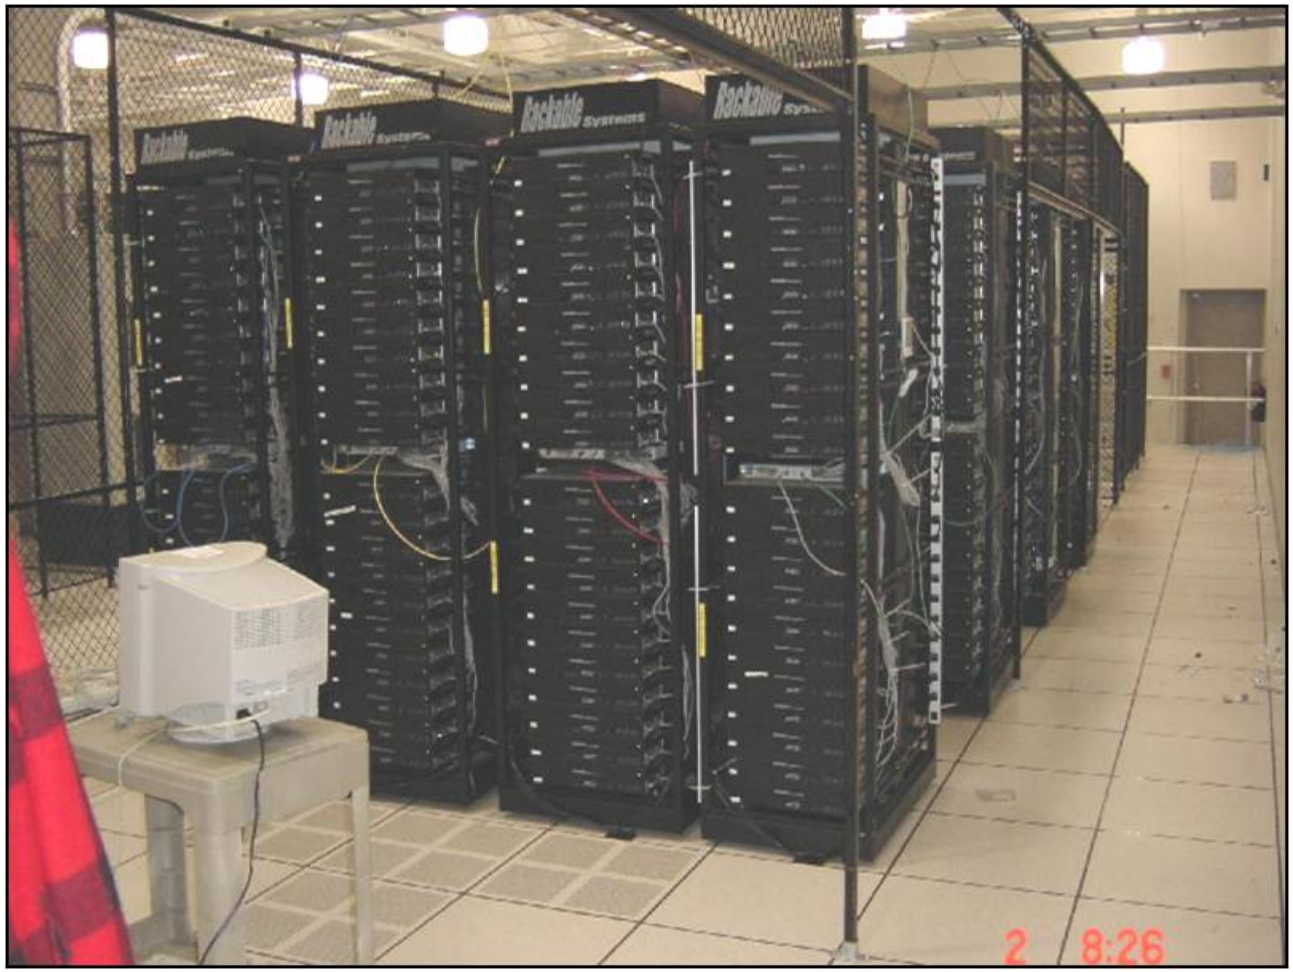

Google servers 2001

(Images courtesy Jeff Dean)

1999 to 2009: Google

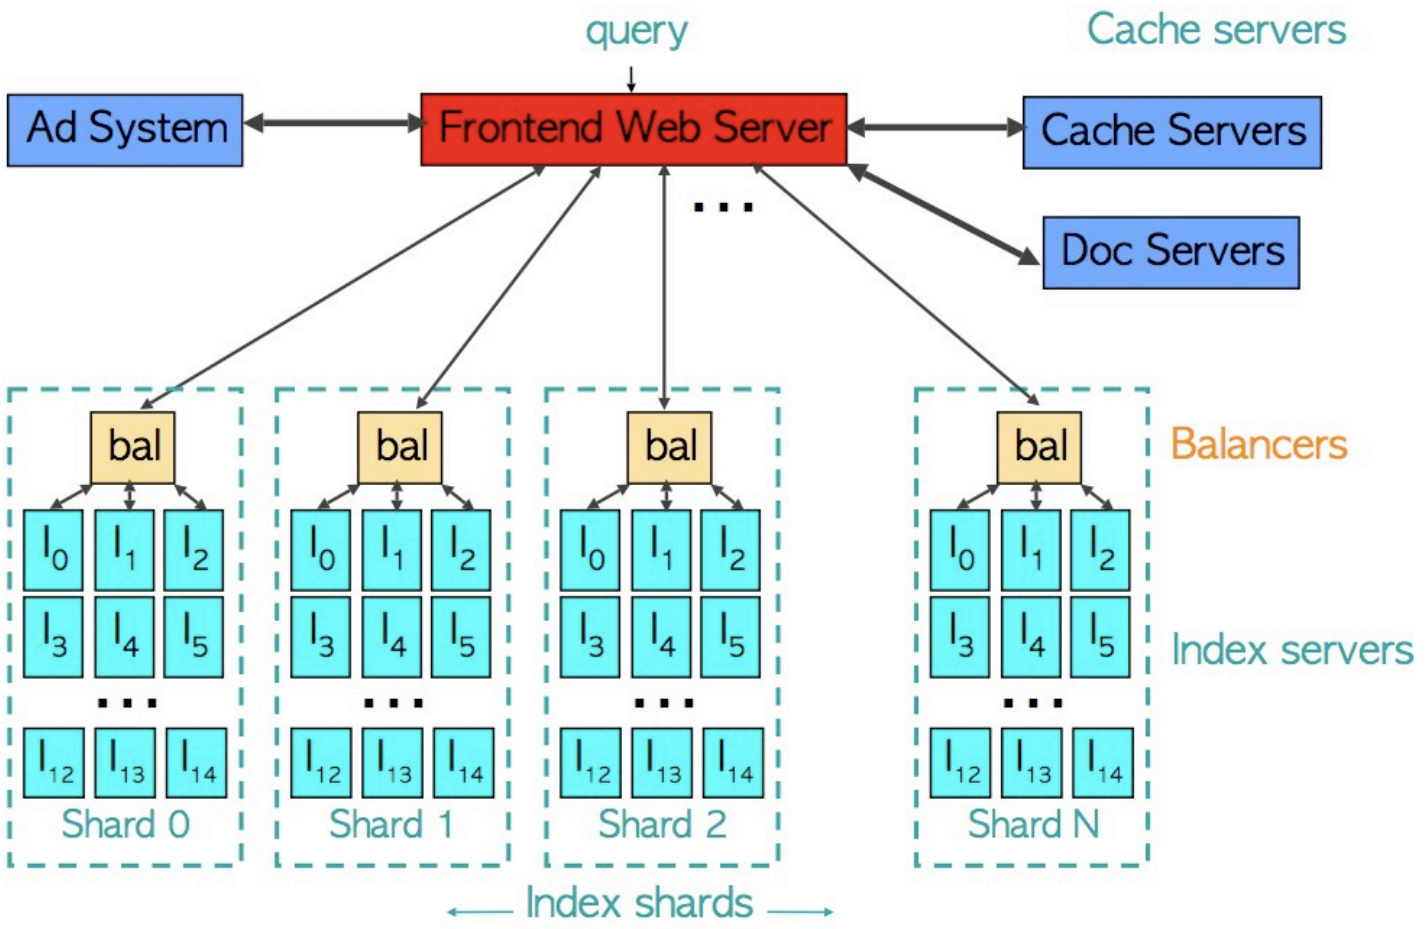

Google architecture early 2000s

(Images courtesy Jeff Dean)

1999 to 2009: Google

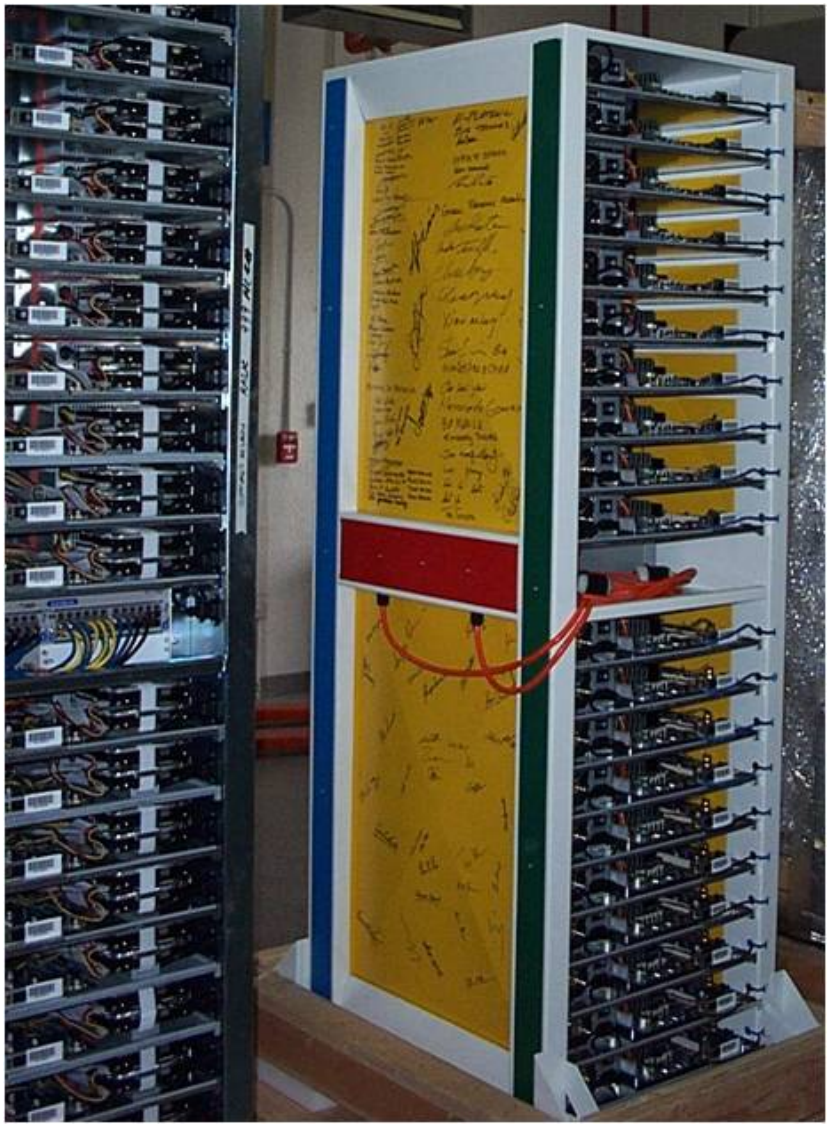

Google servers 2006

- In-house rack design
- PC-class motherboards
- Low-end storage and networking hardware
- Linux
- in-house software

(Images courtesy Jeff Dean)

1999 to 2009: Google

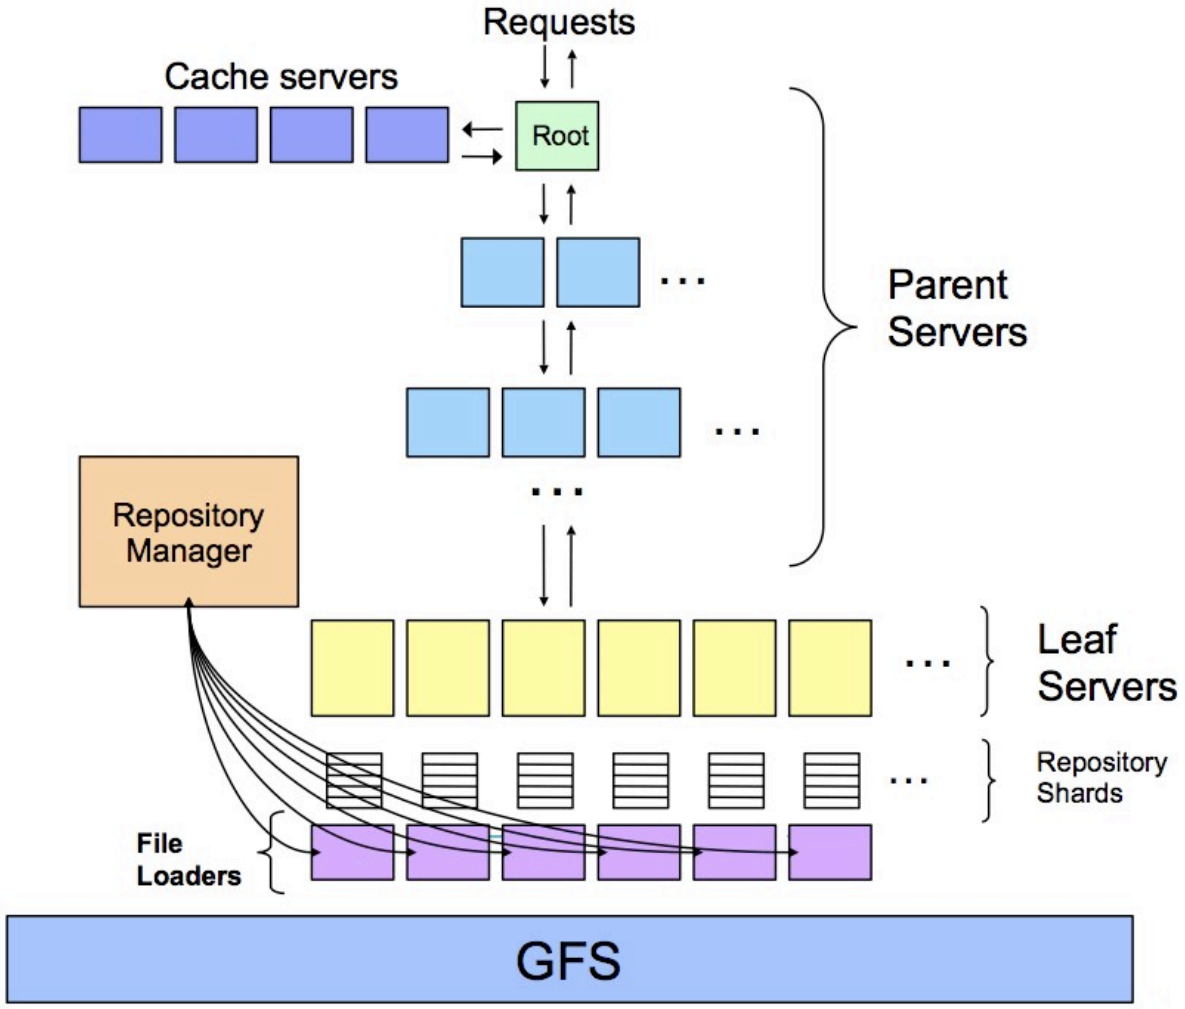

Google architecture mid 2000s

(Images courtesy Jeff Dean)

1999 to 2009: Google

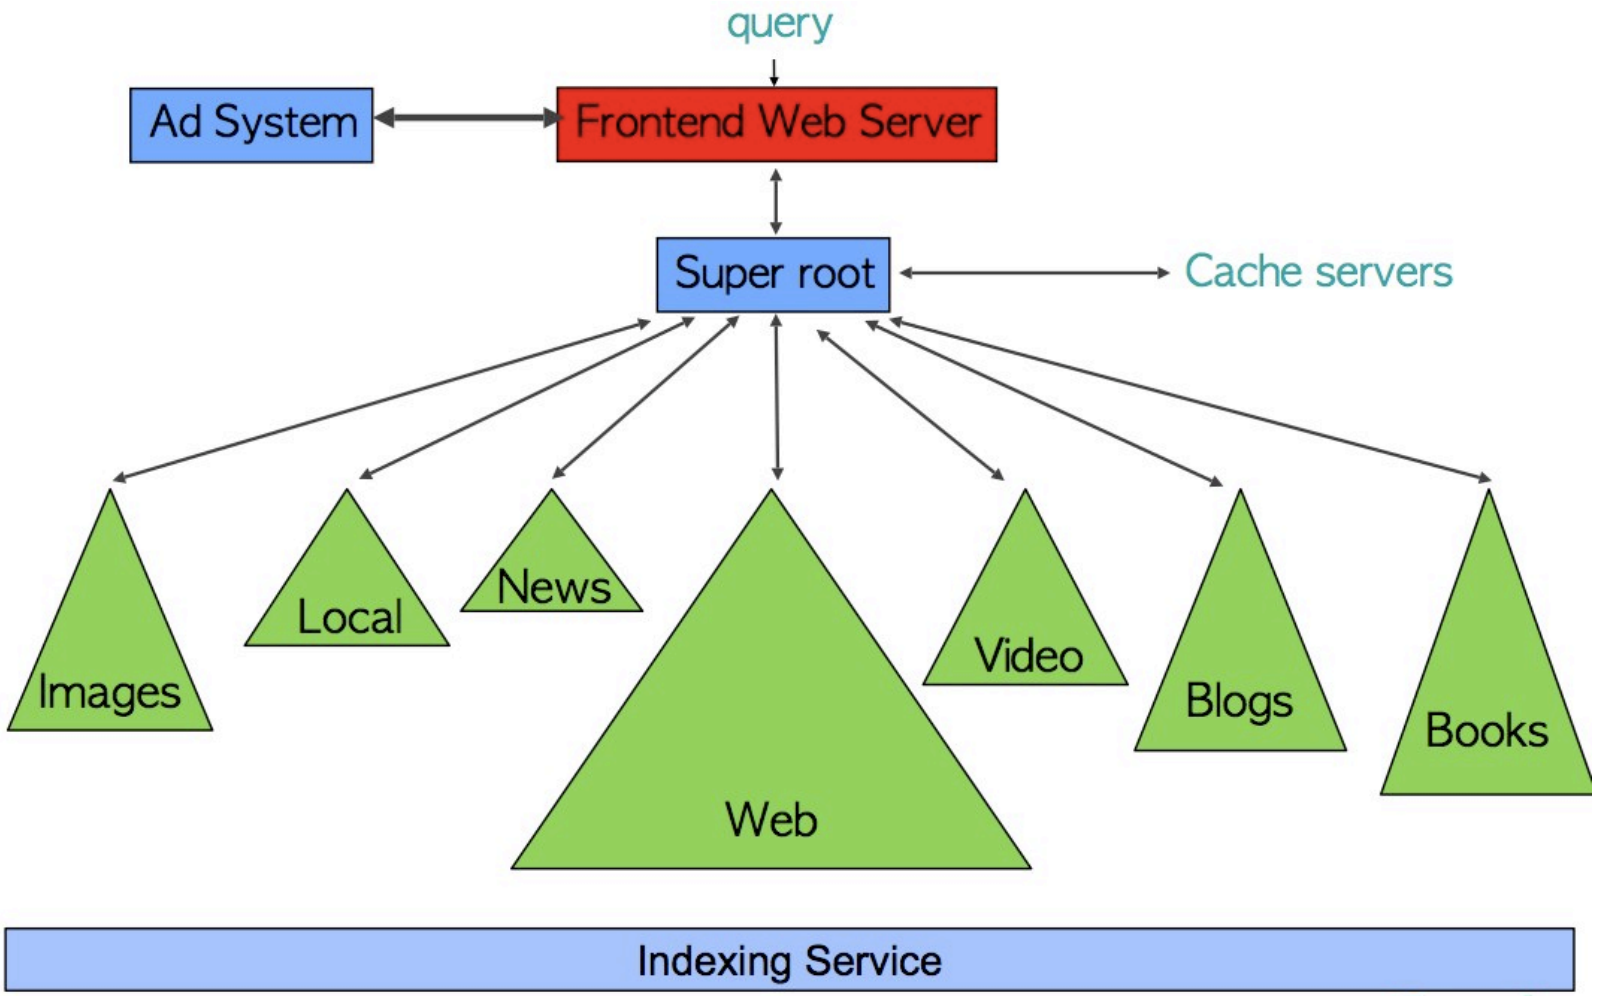

Google architecture late 2000s

(Images courtesy Jeff Dean)

Road Map: Today's Lecture

Scalable Systems

- Example company; grows from 10 users to 100 million users
- Ad serving company spin out
- How both managed growth in a different way
- The design tradeoffs made along the way

A Brief History of Cloud Computing

- A quick recap of technological progress in the last decade

Future of Cloud Computing

- Where are we heading …

Future of Cloud Computing

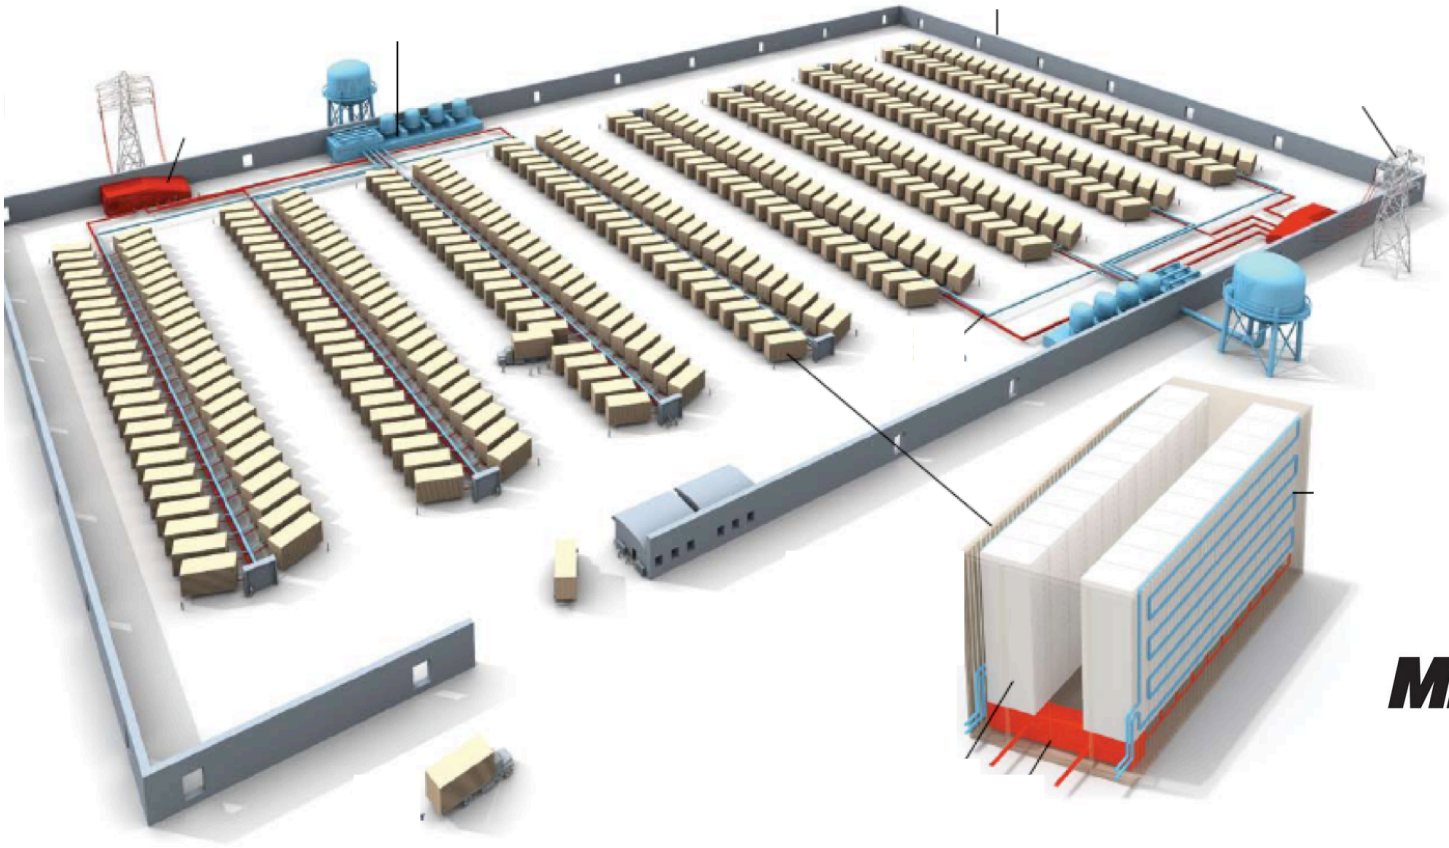

Microsoft®

Future of Cloud Computing

Micro-datacenters

Future of Cloud Computing

Virtualized Network Infrastructure