

**RAD Lab**  
UC Berkeley



# **Above the Clouds: A Berkeley View of Cloud Computing**

Armando Fox, UC Berkeley  
Reliable Adaptive Distributed Systems Lab

© 2009





- What is distributed computing?
- What is warehouse-scale computing?
- What is cloud computing?
- Why should you care?

# What is distributed computing?

- The first demonstration of how to build really large Internet sites out of *clusters* of computers was done by:
  - (a) Stanford
  - (b) Berkeley
  - (c) Yahoo!
  - (d) Google
  - (e) IBM



# Big Computers c. 1996

## Sun E-10000 “supermini”

- Up to 64 processors @250MHz
- Up to 64 GB RAM
- Up to 20 TB Disk
- Used by eBay, among others




## PC

- 200 MHz CPU, 32MB RAM, 4 GB disk



# UC Berkeley Networks Of Workstations (1994-1999)




NOW



NOW-0

1994

Four  
HP-735's






NOW-1

1995

32 Sun SPARC-  
stations







NOW-2

1997

60 Sun SPARC-2



- A Google datacenter built c.2005 would be designed to house approximately \_\_\_\_\_ computers.
  
- (a) 1,000
- (b) 5,000
- (c) 10,000
- (d) 50,000
- (e) 100,000

Challenge: how do you  
program a NOW? (or: what is  
it good for?)



# The Killer App for NOWs

- Prof. Eric Brewer, Armando Fox, Steve Gribble, Paul Gauthier, Yatin Chawathe: *Cluster-Based Scalable Network Servers* in Symposium on Operating Systems Principles, 1997
- *Non-goal*: build best/fastest search engine
  - But led to Inktomi, first *truly scalable* search engine that took advantage of NOW ideas
- Goal: show general techniques for programming NOW's for Internet services



# Access Is the Killer App!

UC Berkeley, 1994-1999

- Project Daedalus: Profs. Katz & Brewer
- Data, services in ~~infrastructure~~ cloud
  - search, email, personal comms, productivity...
- Mobile access *anywhere, anytime*
- Many “firsts”:
  - server architecture with auto-scaling
  - cluster-based Internet service: Inktomi
  - mobile Web: TopGun Wingman on Palm




**Challenge: deploying the service!**



# 2005: *Datacenter* is new “server”

- “*Program*” => Web search, email, map/GIS, ...
- “*Computer*” => 1000’s computers, storage, network
- Warehouse-sized facilities and workloads





# RAD Lab 5-year Mission

*Enable 1 person to develop, deploy, operate next-generation Internet application*

- Key enabling technology: Statistical machine learning
- Highly interdisciplinary faculty & students
  - 7 faculty across CS, from theory to systems
  - 2 postdocs, ~30 PhD students, ~5 undergrads

The multi-colored Google logo.

The Microsoft logo in a bold, black, sans-serif font.

The Sun Microsystems logo, featuring a blue square icon with white lines and the word "Sun" in a blue serif font, with "microsystems" in a smaller blue sans-serif font below it.

The Amazon Web Services logo, featuring a yellow cube icon and the text "amazon web services" in a blue sans-serif font.

The Cloudera logo, featuring a blue bird icon and the word "cloudera" in a blue sans-serif font.

The Cisco logo, featuring a blue and white striped icon and the word "CISCO" in a red sans-serif font.

The eBay logo, featuring the word "eBay" in a stylized, multi-colored font.

The Facebook logo, featuring the word "facebook" in a white sans-serif font inside a blue rectangular box.

The Fujitsu logo, featuring the word "FUJITSU" in a red sans-serif font with a red infinity symbol above the "i".

The Intel logo, featuring the word "intel" in a blue sans-serif font inside a blue oval shape.

The HP logo, featuring the letters "hp" in a blue sans-serif font inside a blue square, with the word "invent" in a smaller blue sans-serif font below it.

The NetApp and SAP logos. NetApp is a small logo with the letters "nc" above it. SAP is a logo with the letters "SAP" in a blue sans-serif font inside a blue trapezoidal shape.

The VMware logo, featuring a blue icon of three overlapping squares and the word "vmware" in a blue sans-serif font.

The Yahoo! Research logo, featuring the word "YAHOO!" in a purple sans-serif font with an exclamation point, and the word "RESEARCH" in a smaller purple sans-serif font below it.





# Utility Computing Arrives

- Amazon Elastic Compute Cloud (EC2)
- “Compute unit” rental: \$0.08-0.80/hr.
  - 1 CU  $\approx$  1.0-1.2 GHz 2007 AMD Opteron/Xeon core

“Instances”	Platform	Cores	Memory	Disk
Small - \$0.10 / hr	32-bit	1	1.7 GB	160 GB
Large - \$0.40 / hr	64-bit	4	7.5 GB	850 GB – 2 spindles
XLarge - \$0.80 / hr	64-bit	8	15.0 GB	1690 GB – 3 spindles

- No up-front cost, no contract, no minimum
- Billing rounded to nearest hour; pay-as-you-go storage also available
- A new paradigm for deploying services?

But...

What *is* cloud computing,  
exactly?



## “It’s nothing (new)”

*“...we’ve redefined Cloud Computing to include everything that we already do... I don’t understand what we would do differently ... other than change the wording of some of our ads.”*

– *Larry Ellison, CEO, Oracle*  
*(Wall Street Journal, Sept. 26, 2008)*



## *Above the Clouds:*

# *A Berkeley View of Cloud Computing*

[abovetheclouds.cs.berkeley.edu](http://abovetheclouds.cs.berkeley.edu)

- 2/09 White paper by RAD Lab PI's and students
- Goal: stimulate discussion on *what's really new*
  - Clarify terminology
  - Comparison with conventional computing
  - Identify challenges & opportunities
- Why can we offer new perspective?
  - Strong engagement with industry
  - Users of cloud computing in our own research and teaching
- Over 60,000 downloads



# *Above The Clouds Impact*

- Research collaborations/hires: Amazon, Google, Microsoft, Twitter, Facebook, Cloudera, Yahoo!...
- Invited presentations/advice
  - Google, Fujitsu, IBM, HP, Microsoft, SAP, Juniper, ...
  - World Economic Forum
  - Nat'l Academy of Engineering
  - OpenCirrus Summit
  - UCB Office of the CIO
  - UC Systemwide Cloud Computing Task Force

*UCB is academic leader in cloud computing in both  
research & education*



# What is it? What's new?

- Old idea: Software as a Service (SaaS)
  - Software hosted in the infrastructure vs. installed on local servers or desktops; dumb (but brawny) terminals
- **New:** pay-as-you-go *utility computing*
  - Illusion of infinite resources on demand
  - Fine-grained billing: release == don't pay
  - Earlier examples: Sun, Intel Computing Services
    - longer commitment, more \$\$\$/hour, no storage
  - *Public (utility) vs. private clouds*



- How much data per month, approximately, is processed through Google's *BigTable* storage system?
  - (a) 1 TB (1,000 GB)
  - (b) 100 TB
  - (c) 1 PB (1,000 TB)
  - (d) 100 PB
  - (e) 1 EB (exabyte = 1,000 PB)




# Why Now (not then)?

- The Web “**Space Race**”: Build-out of extremely large datacenters (10,000’s of **commodity** PCs)
- Driven by growth in demand (more users)
  - Infrastructure software: e.g., Google File System
  - Operational expertise
  - Discovered **economy of scale: 5-7x** cheaper than provisioning a medium-sized (100’s machines) facility
- More pervasive broadband Internet
- Free & open source software




# Cloud Economics 101

- Static provisioning for peak - wasteful, but necessary for SLA



“Statically provisioned”  
data center




“Virtual” data center  
in the cloud

 Unused resources

# Risk of Under Utilization

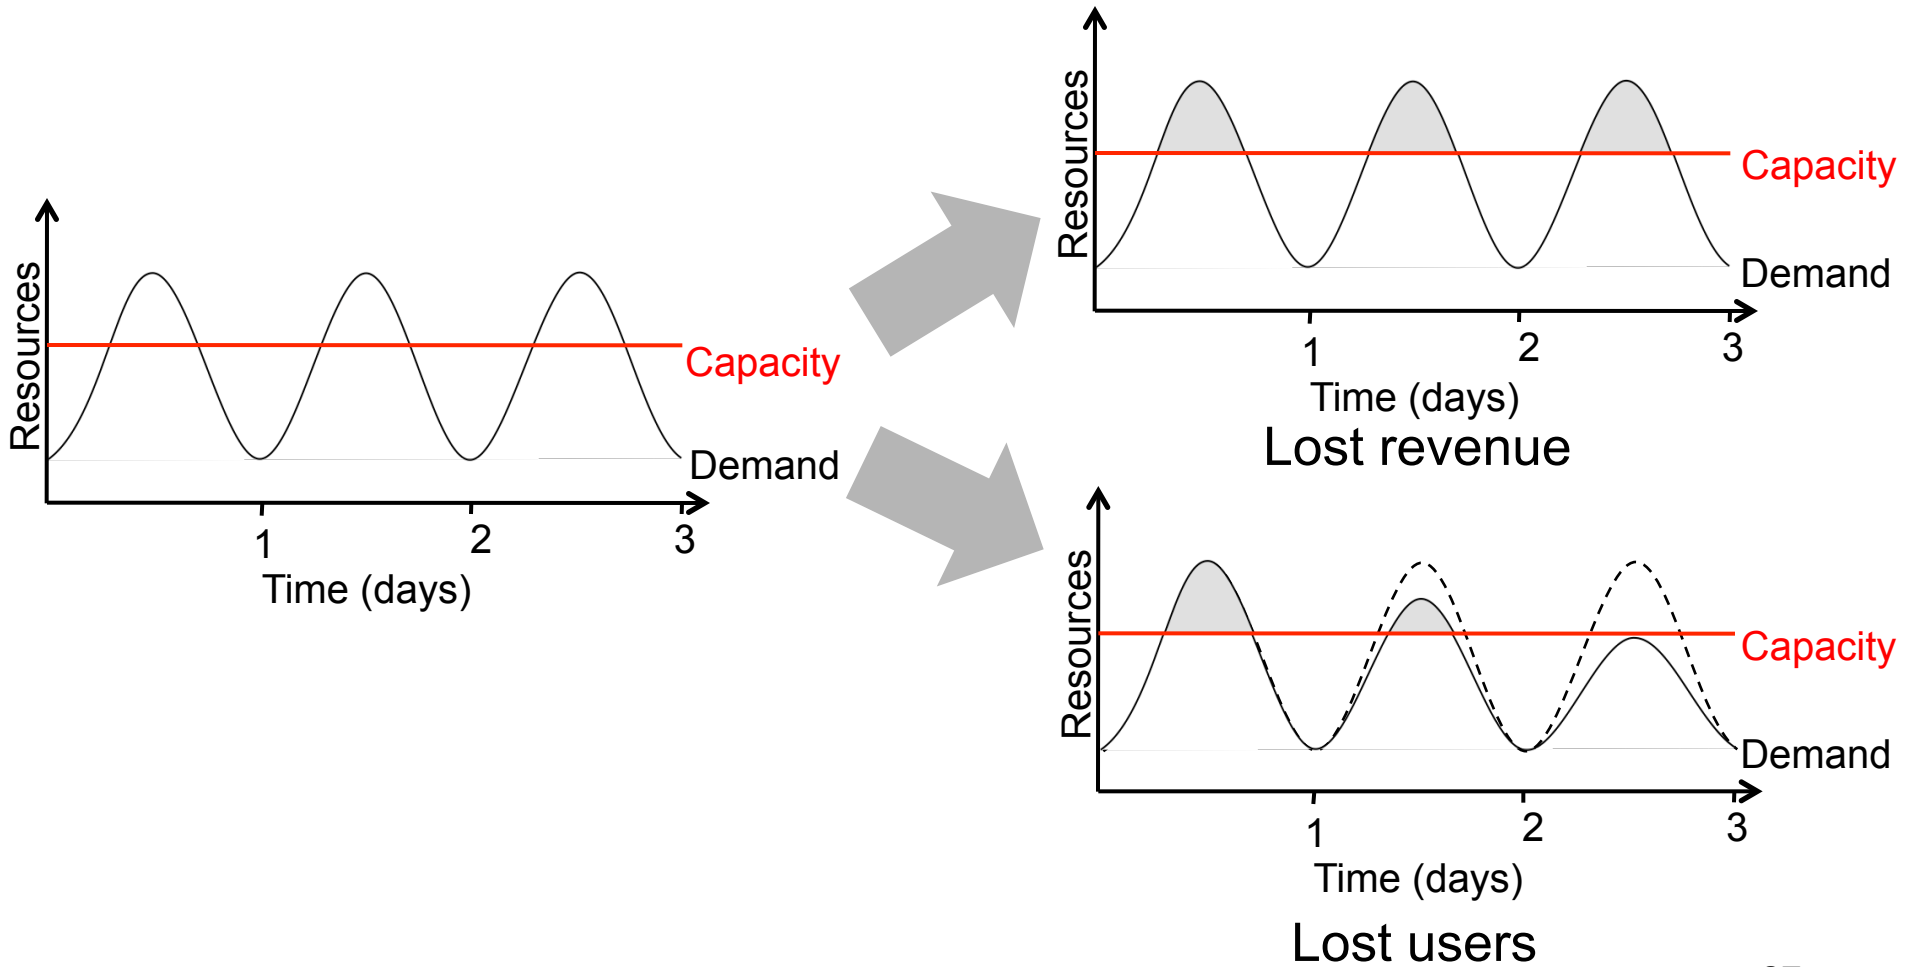
- Underutilization results if “peak” predictions are too optimistic



 Unused resources

Static data center

# Risks of Under Provisioning



What can you do with this?



# Cost Associativity

- 1,000 CPUs for 1 hour same price as 1 CPU for 1,000 hours
- Washington Post converted Hillary Clinton's travel documents to post on WWW
  - Conversion time: **<1 day** after released
  - Cost: less than \$200
- RAD Lab graduate students demonstrate improved MapReduce scheduling—on 1,000 servers

# Risk transfer

- 2001: CNN home page meltdown on 9/11
  - ~10x traffic increase in ~15 minutes
  - result: site had to go offline
- 2008: Animoto
  - traffic doubled every 12 hours for 3 days when released as Facebook plug-in
  - Scaled from 50 to >3500 servers
  - ***...then scaled back down***

# Indexing the Web

To be or  
not to be...

to	A
be	A
or	A
not	A

...or a  
better  
fool ...


or	B
a	B
better	B
fool	B

...better to  
love a  
fool...

better	C
to	C
love	C
a	C
fool	C



Map &  
Combine





# MapReduce in Practice

- Example: spam classification
  - training:  $10^7$  URLs x 64KB data each = 640GB data
  - One heavy-duty server: ~270 hours
  - 100 servers in cloud: ~3 hours (= ~\$255)
- Rapid uptake in other scientific research
  - Large-population genetic risk analysis & simulation (Harvard Medical School)
  - Genome sequencing (UNC Chapel Hill Cancer Ctr)
  - many others... so *what's the downside?*





# Challenges & Opportunities

- Challenges to adoption, growth, & business/policy models
- Both technical and nontechnical
- Most translate to 1 or more *opportunities*
- Complete list in paper; I'll discuss subset



# Challenge: Cloud Programming

- Challenge: exposing parallelism
  - MapReduce relies on “embarrassing parallelism”
- Programmers must (re)write problems to expose this parallelism, if it’s there to be found
- Tools still primitive, though progressing rapidly



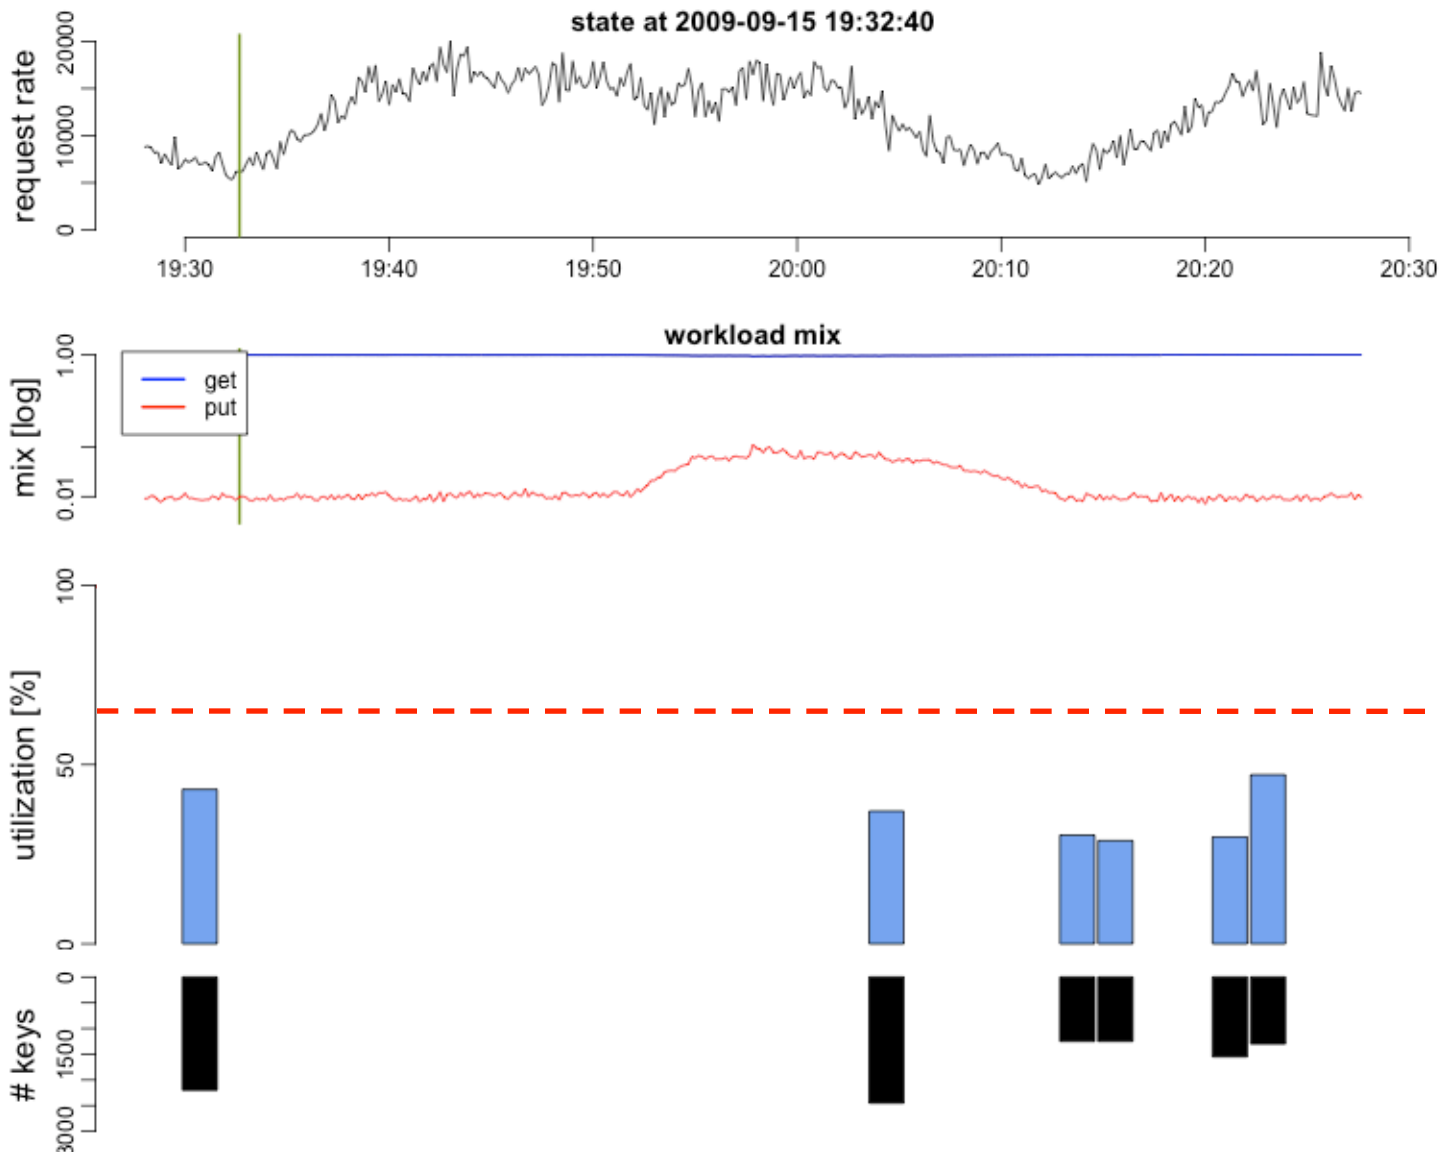
# Challenge: Big Data

Application	Data generated per day
DNA Sequencing (Illumina HiSeq machine)	1 TB
Large Synoptic Survey Telescope	30 TB; 400 Mbps sustained data rate between Chile and NCSA
Large Hadron Collider	60 TB

- Challenge: Long-haul networking is most expensive cloud resource, and improving most slowly
- Copy 8 TB to Amazon over ~20 Mbps network  
=> ~35 days, ~\$800 in transfer fees
- How about shipping 8TB drive to Amazon instead?  
=> 1 day, ~\$150 (shipping + transfer fees)



# Web services in the cloud





# Cloud in Education

1. Berkeley research culture: integrate leading research into teaching at all levels
2. RAD Lab need for “killer apps” to show off infrastructure

Current efforts (student counts approximate):

- Great Ideas in Computer Architecture (reinvented Fall 2010): 190 students
- Software Engineering for SaaS (in its 4<sup>th</sup> iteration): 50+50+50+70 students
- Operating Systems: 70 students
- Intro. Data Science (Spring 2010): 30
- Adv. topics in HCI: 20 students
- Natural language processing: 20 students



# AWS is a great fit for courses...

- New undergraduate teaching opportunities
  - SaaS: make a database fall over—would need 200 servers for ~20 project teams
  - deploy projects publicly, many continue after course
- Better use of resources
  - Heavy usage right before lab deadlines



# Success stories

**PEOPLE DEBATE**

Main Page | Create Account | Create Debate | FAQ | Login

Welcome to PeopleDebate – A new debate forum which empowers users to voice their opinions, to establish credibility, and to highlight the most important ideas by using the up and down arrows.

**PeopleDebate 2008 Presidential Election Debate Results**  
**John McCain** is currently winning the debate with **54%** of the vote!  
 Support your candidate - [Click here register your vote!](#)  
 Then voice your opinions in our [John McCain versus Barack Obama debate!](#)

Select an option, join a debate below, or [click here](#) for more election resources.

**Obama vs. McCain** | Create Debate | Hall of Fame | Read FAQ | Contact Us

**Recently Active Debates** | Most Visited

Some Recent Posts

Who should be president - John McCain or Barack Obama?  
 Answer: **McCain 54%** Visits: 1561 Posts: 112

Should your taxes go to bail out the millionaire or billionaire, that got their selfish selves in this big mess?  
 Answer: **No 100%** Visits: 5 Posts: 1

Does Barack Obama have enough experience?  
 Answer: **Yes 53%** Visits: 2539 Posts: 112

Why do democrats and republicans always have to bash each other? Should they have at least one debate about what needs to be done for America, and how or what they intend to do?  
 Answer: **Yes 100%** Visits: 13 Posts: 1

Question: Should your taxes go to bail out the millionaire or billionaire, that got their selfish selves in this big mess?  
 Argument for "No" By sam\_dunit on 9/30/2008

No! you worked hard to earn enough, so the U.S. revenue people could "take" a large enough part of your check already. We are required to send them tax money. But we shouldn't have to bail out the big loan companies and big careless bankers or any other big business of any kind!!

Question: Why do democrats and republicans always have to bash each other? Should they have at least one debate about what needs to be done for America, and how or what they intend to do?  
 Argument for "Yes" By sam\_dunit on 9/29/2008

I think we have enough bashing, its time to get on with what is going to be done

**LOST AND FOUND 2.0 FOUNDIT™**

HOME | LOST ITEMS | FOUND ITEMS | MY ITEMS | MY NOTIFIER

sign up now  
 find your items  
 find the owner

find your items and protect your items

Cooperate with Berkeley Police and Lost and Found department! Protect your item before it gets lost! Find your item back if it's lost!

User Information

You haven't logged in. Please login or sign up below

Email : fox

Password : \*\*\*\*\*

Log in | Remember me

Sign up | Activate Account | Forgot my password

Lost and Found News

Notifier is working now.

Notifier is working now. If any newly posted found items' title matches notifiers' keywords, an email will be sent to user's email.

Posted at Thu Jun 19 07:22:03 -0700 2008 by Chaohao Wang

Image upload

Users can upload up to 4 pictures for each of their lost items after posting items.

**CommuterPool**

Welcome to **CommuterPool**, a website designed to help you reduce commute pollution and save money too. If you own a car, you can easily find passengers to share the expenses or if you are looking for a ride, chances are somebody share your same route.

**How does it work?**

- Create an account by clicking "Register"
- Enter your starting point and destination
- Browse among user that shares the same route
- Choose your ride!

LOGIN → BROWSE → RIDE

Discover and Post Ride

Find a Ride or Start a New Carpool

Street Address: \_\_\_\_\_  
 Apartment, suite, unit, etc: \_\_\_\_\_

LOGIN

Discover and Post Ride

Street Address: \_\_\_\_\_  
 Apartment, suite, unit, etc: \_\_\_\_\_

RIDE

**hesperian**  
 Publishing for community health and empowerment

Home | About Hesperian | Action for Health | Current Projects | Publications and Resources | Hesperian Store | Get Involved | Contact Us

**Hesperian Celebrates Clinton Global Initiative Commitment and \$2.7 Million Grant from the Bill and Melinda Gates Foundation**

President Bill Clinton recently highlighted Hesperian's Commitment to empower communities in confronting the death and disease that a lack of safe drinking water and sanitation cause at the Clinton Global Initiative gathering in New York. Hesperian is also delighted to announce that it is the recipient of a \$2.7 million grant from the Bill and Melinda Gates Foundation to update and expand one of our most important titles, Where There Is No Doctor. [More...](#)

[Hesperian announces our New Spanish-language Web Site: http://espanol.hesperian.org!](#)

Now available: *Ayudar a los niños sordos*, the Spanish translation of Helping Children Who Are Deaf.

And coming soon: *Un manual de salud para mujeres con discapacidad*. Visit [Books in Progress](#) to learn more.

Buy Books

ESPAÑOL

Search Site

DONATE NOW!

free downloads

click here

What's New

Read our Weblog  
 Hesperian recently implemented something new on our website.




# Summary

- Cloud computing *democratizes access* to “supercomputer-class” capability
  - All you need is a credit card
- Puts students, academia on more level playing field to have high impact in industry
- The next Google, eBay, Amazon, etc. can come from a small team of entrepreneurs even *without* heavy dose of \$\$ up front



# Going back to NOW...

- **2000**: using medium-sized clusters for Internet services  
=> several PhD's



- **2010**: CS169 students do it in 6-8 weeks and deploy on cloud computing
- **2020**: ?

**Thank you!**