RAD Lab
UC Berkeley

Above the Clouds: A Berkeley View of Cloud Computing

Armando Fox, UC Berkeley
Reliable Adaptive Distributed Systems Lab

Image: John Curley http://www.flickr.com/photos/jay_que/1834540/

RAD Lab

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

RAD Lab
UC Berkeley

What is distributed computing?

RAD Lab

- 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

RAD Lab

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

RAD Lab

UC Berkeley *Networks Of Workstations* (1994-1999)

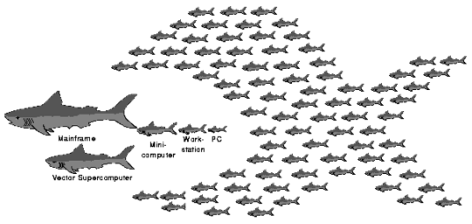

NOW

NOW-0

1994

Four HP-735's

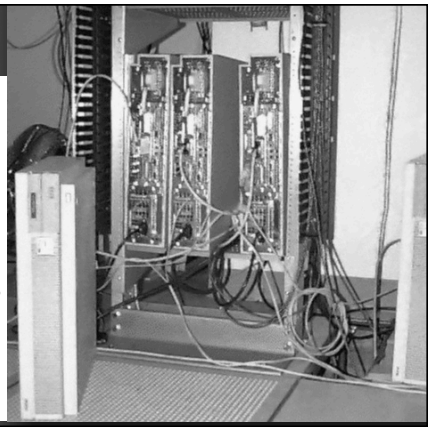

NOW-1

1995

32 Sun SPARC-stations

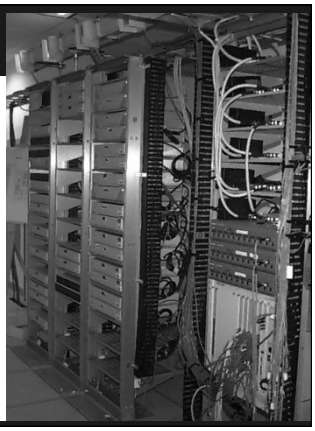

NOW-2

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

10

RAD Lab
UC Berkeley

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

11

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

12

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

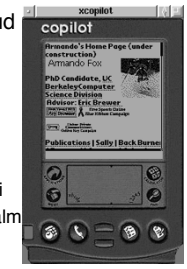

Challenge: deploying the service!

14

2005: *Datacenter* is new "server"

- *"Program"* => Web search, email, map/GIS, …
- *"Computer"* => 1000's computers, storage, network
- Warehouse-sized facilities and workloads

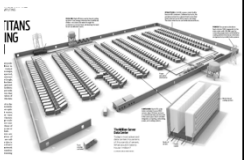

photos: Sun Microsystems, CNET, & datacenterknowledge.com

15

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

Google Microsoft Sun microsystems
amazon web services cloudera cisco eBay facebook FUJITSU
intel hp invent NetApp SAP vmware YAHOO! RESEARCH

16

Utility Computing Arrives

- Amazon Elastic Compute Cloud (EC2)
- "Compute unit" rental: $0.08-0.80/hr.
  - 1 CU ≈ 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?

17

But...

What *is* cloud computing, exactly?

18

"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)*

19

Above the Clouds: A Berkeley View of Cloud Computing

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

20

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

22

- 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)

23

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

24

Cloud Economics 101

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

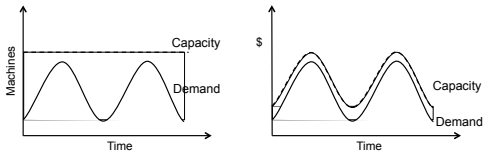

"Statically provisioned" data center

"Virtual" data center in the cloud

Unused resources

25

Risk of Under Utilization

- Underutilization results if "peak" predictions are too optimistic

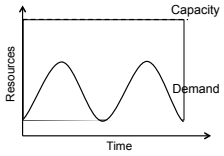

Unused resources

Static data center

26

Risks of Under Provisioning

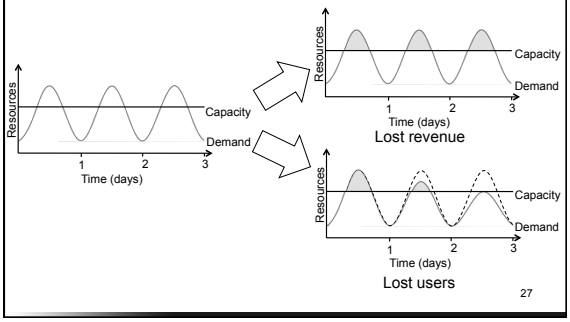

27

What can you do with this?

28

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

29

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***

30

Indexing the Web

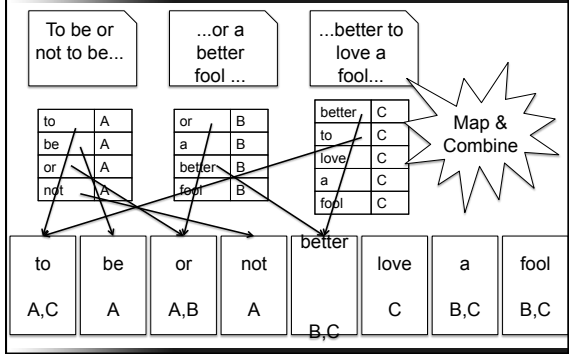

MapReduce in Practice

- 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

33

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)

Source: Ed Lazowska, eScience 2010, Microsoft Cloud Futures Workshop, lazowska.cs.washington.edu/cloud2010.pdf

Web services in the cloud

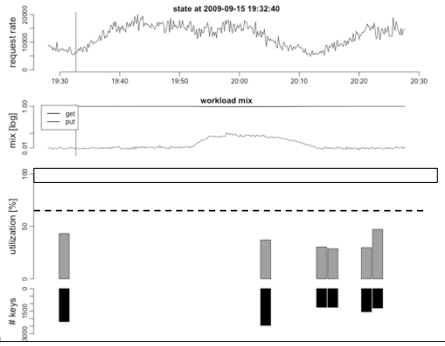

Cloud in Education

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 4th 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

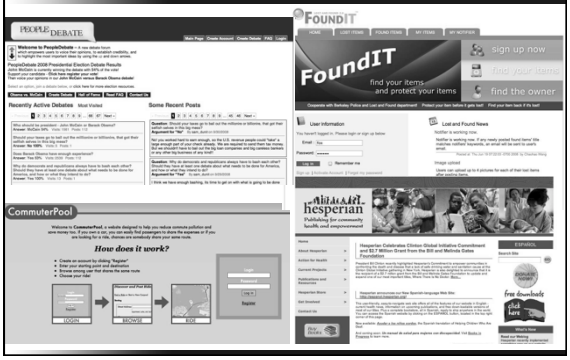

Summary

- 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

40

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: ?

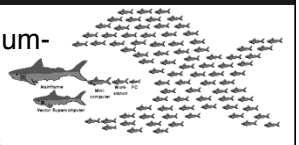

NOW

RAD Lab
UC Berkeley

Thank you!

42