MapReduce for Scalable and Cloud Computing

CS6323

Adapted from NETS212, U. Penn, USA

Overview

- Networked computing
  - The need for scalability; scale of current services
  - Scaling up: From PCs to data centers
  - Problems with 'classical' scaling techniques
- Utility computing and cloud computing
  - What are utility computing and cloud computing?
  - What kinds of clouds exist today?
  - What kinds of applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'
  - Some cloud computing challenges

Cloud Computing Concept

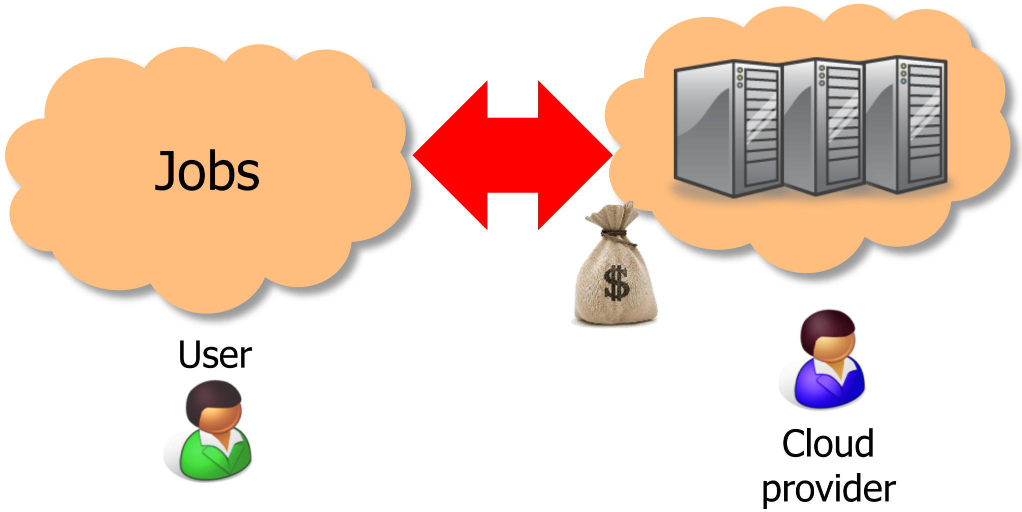

Cloud Computing Concept

- Network structure hidden from user
- No concept of programming paradigm

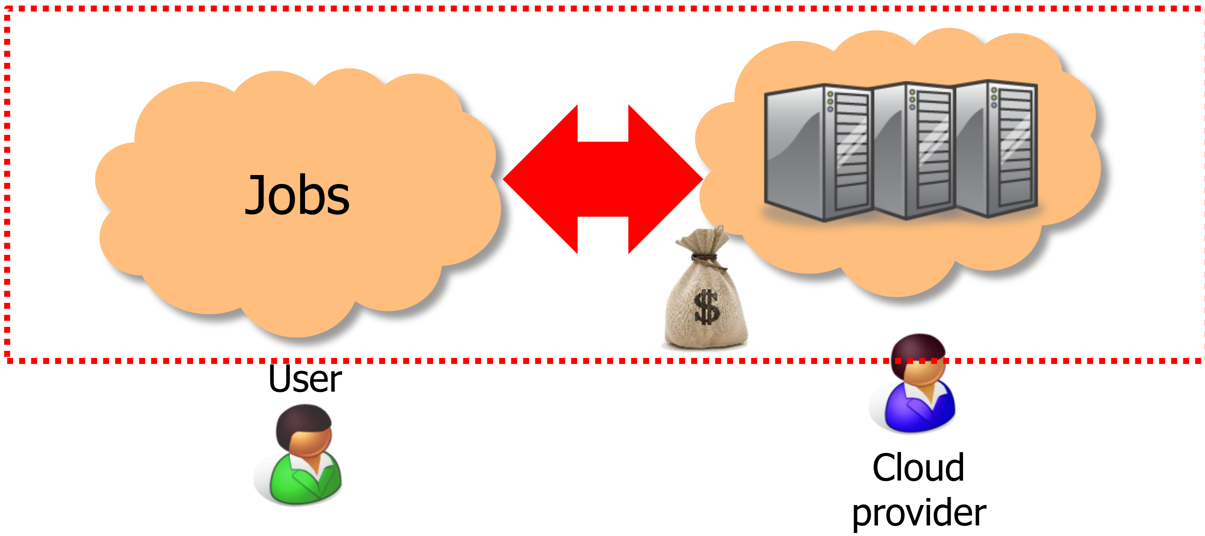

Focus on programming paradigm

How many users and objects?

- Flickr has >6 billion photos
- Facebook has 1.15 billion active users
- Google is serving >1.2 billion queries/day on more than 27 billion items
- >2 billion videos/day watched on YouTube

How much data?

- Modern applications use massive data:
  - Rendering 'Avatar' movie required >1 petabyte of storage
  - eBay has >6.5 petabytes of user data
  - CERN's LHC will produce about 15 petabytes of data per year
  - In 2008, Google processed 20 petabytes per day
  - German Climate computing center dimensioned for 60 petabytes of climate data
  - Google now designing for 1 exabyte of storage
  - NSA Utah Data Center is said to have 5 zettabyte (!)
- How much is a zettabyte?
  - 1,000,000,000,000,000,000,000 bytes
  - A stack of 1TB hard disks that is 25,400 km high

How much computation?

- No single computer can process that much data
  - Need many computers!
- How many computers do modern services need?
  - Facebook is thought to have more than 60,000 servers
  - 1&1 Internet has over 70,000 servers
  - Akamai has 95,000 servers in 71 countries
  - Intel has ~100,000 servers in 97 data centers
  - Microsoft reportedly had at least 200,000 servers in 2008
  - Google is thought to have more than 1 million servers, is planning for 10 million (according to Jeff Dean)

Why should I care?

- Suppose you want to build the next Google
- How do you…
  - … download and store billions of web pages and images?
  - … quickly find the pages that contain a given set of terms?
  - … find the pages that are most relevant to a given search?
  - … answer 1.2 billion queries of this type every day?
- Suppose you want to build the next Facebook
- How do you…
  - … store the profiles of over 500 million users?
  - … avoid losing any of them?
  - … find out which users might want to be friends?
- Stay tuned!

Overview

- Networked Computing
  - The need for scalability; scale of current services
  - Scaling up: From PCs to data centers NEXT
  - Problems with 'classical' scaling techniques
- Utility computing and cloud computing
  - What are utility computing and cloud computing?
  - What kinds of clouds exist today?
  - What kinds of applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'
  - Some cloud computing challenges

Clusters

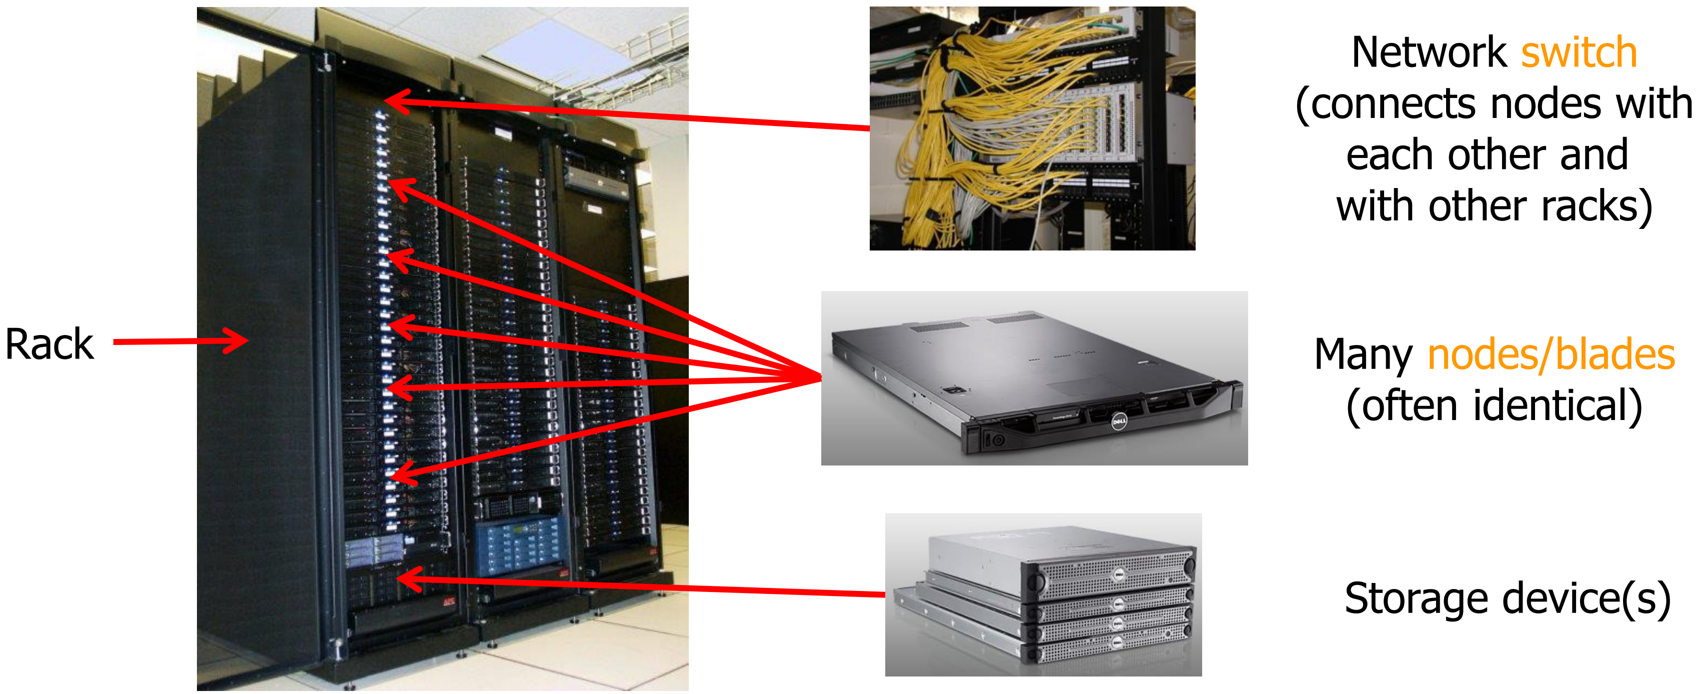

- Characteristics of a cluster:
  - Many similar machines, close interconnection (same room?)
  - Often special, standardized hardware (racks, blades)
  - Usually owned and used by a single organization

Power and cooling

- Clusters need lots of power
  - Example: 140 Watts per server
  - Rack with 32 servers: 4.5kW (needs special power supply!)
  - Most of this power is converted into heat
- Large clusters need massive cooling
  - 4.5kW is about 3 space heaters
  - And that's just one rack!

Scaling up

PC

Server

Cluster

Data center

- What if your cluster is too big (hot, power hungry) to fit into your office building?
  - Build a separate building for the cluster
  - Building can have lots of cooling and power
  - Result: Data center

What does a data center look like?

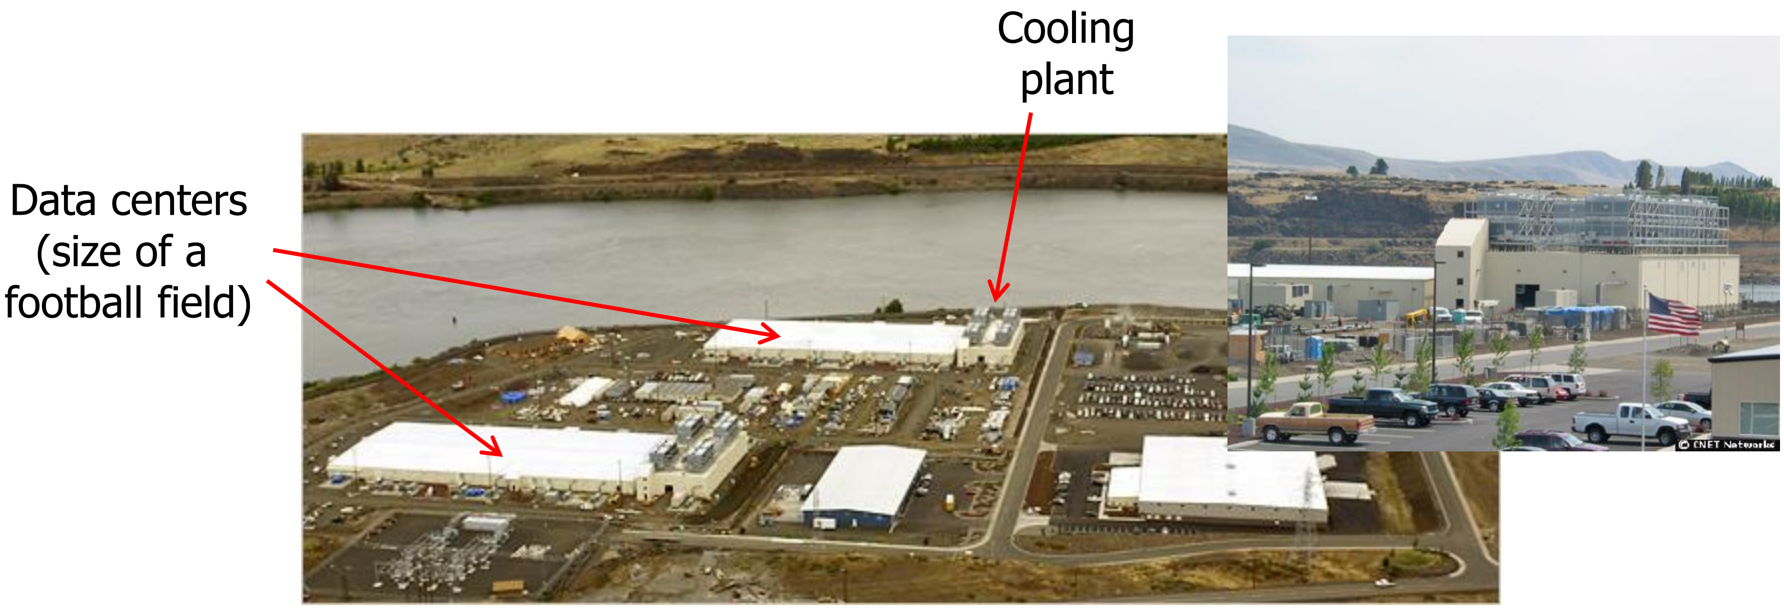

Google data center in The Dalles, Oregon

- A warehouse-sized computer
  - A single data center can easily contain 10,000 racks with 100 cores in each rack (1,000,000 cores total)

What's in a data center?

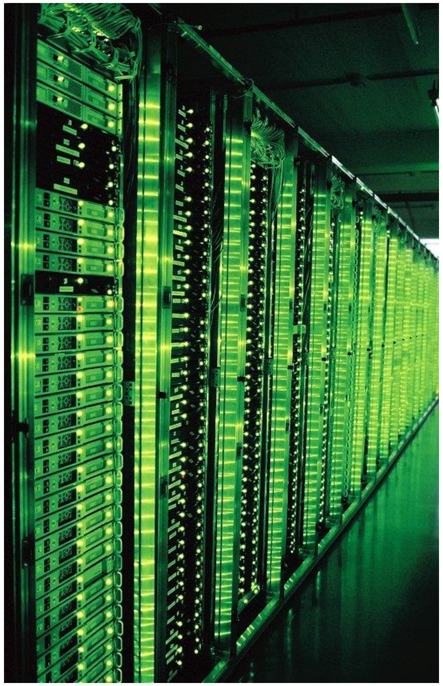

Source: 1&1

- Hundreds or thousands of racks

What's in a data center?


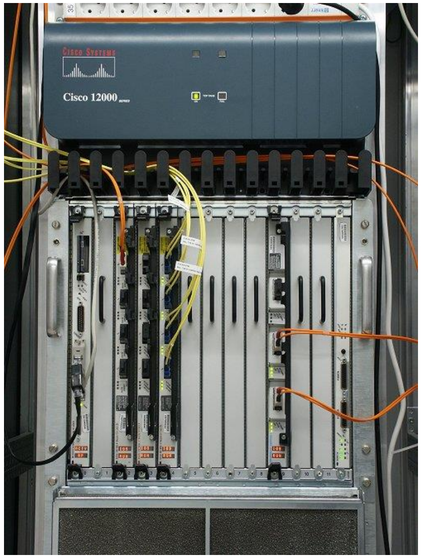

Source: 1&1

- Massive networking

What's in a data center?

Source: 1&1

- Emergency power supplies

What's in a data center?

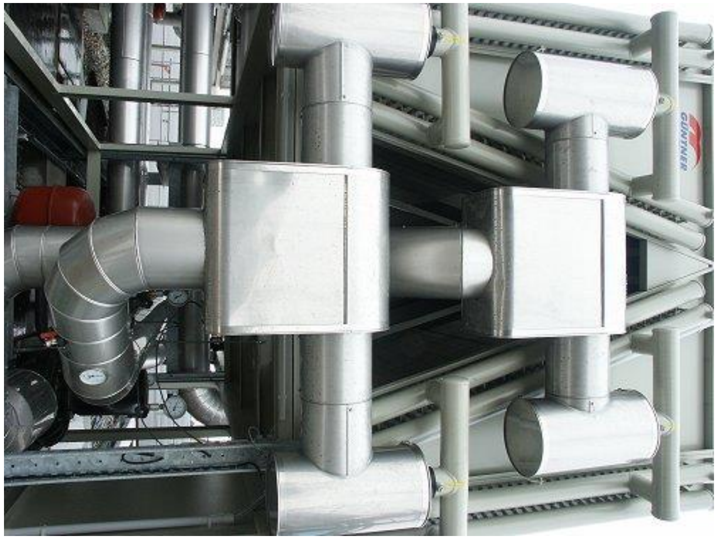

Source: 1&1

- Massive cooling

Energy matters!

Company	Servers	Electricity	Cost
eBay	16K	~$0.6*10^5$ MWh	~$3.7M/yr
Akamai	40K	~$1.7*10^5$ MWh	~$10M/yr
Rackspace	50K	~$2*10^5$ MWh	~$12M/yr
Microsoft	>200K	>$6*10^5$ MWh	>$36M/yr
Google	>500K	>$6.3*10^5$ MWh	>$38M/yr
USA (2006)	10.9M	$610*10^5$ MWh	$4.5B/yr

Source: Qureshi et al., SIGCOMM 2009

- Data centers consume a lot of energy
  - Makes sense to build them near sources of cheap electricity
  - Example: Price per KWh is 3.6ct in Idaho (near hydroelectric power), 10ct in California (long distance transmission), 18ct in Hawaii (must ship fuel)
  - Most of this is converted into heat → Cooling is a big issue!

Scaling up


PC


Server


Cluster


Data center

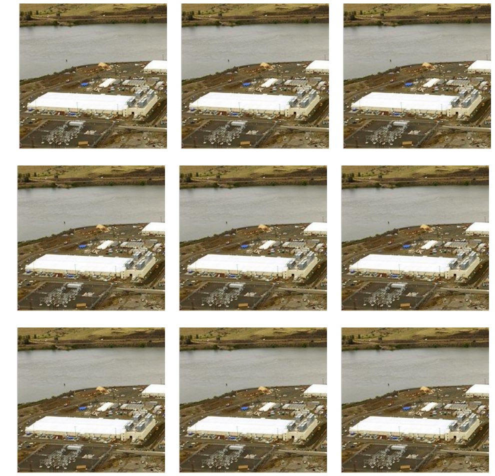
Network of data centers

- What if even a data center is not big enough?
    - Build additional data centers
    - Where? How many?

Global distribution

- Data centers are often globally distributed
  - Example above: Google data center locations (inferred)
- Why?
  - Need to be close to users (physics!)
  - Cheaper resources
  - Protection against failures

Trend: Modular data center

- Need more capacity? Just deploy another container!

Overview

- Networked Computing
  - The need for scalability; scale of current services
  - Scaling up: From PCs to data centers
  - Problems with 'classical' scaling techniques NEXT
- Utility computing and cloud computing
  - What are utility computing and cloud computing?
  - What kinds of clouds exist today?
  - What kinds of applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'
  - Some cloud computing challenges

Problem #1: Difficult to dimension

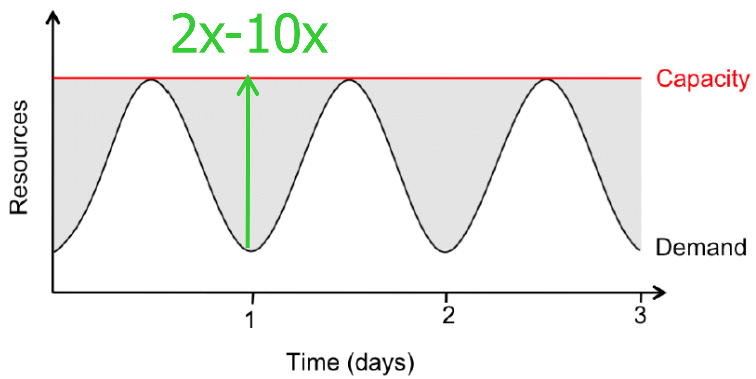

Provisioning for the peak load

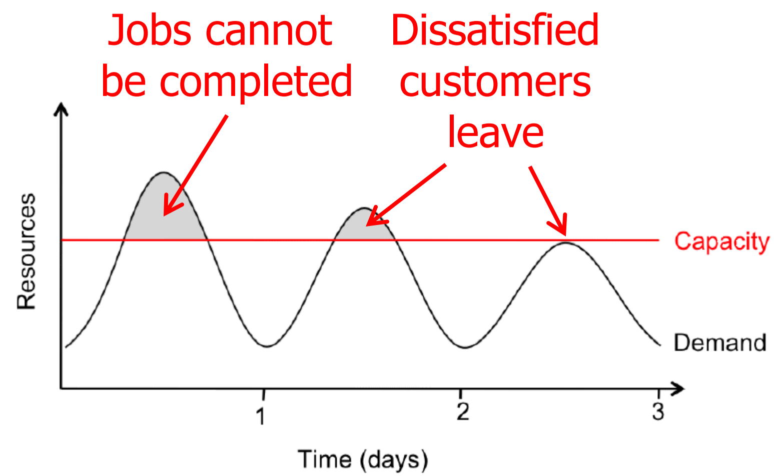

Provisioning below the peak

- Problem: Load can vary considerably
  - Peak load can exceed average load by factor 2x-10x
  - But: Few users deliberately provision for less than the peak
  - Result: Server utilization in existing data centers ~5%-20%!!
  - Dilemma: Waste resources or lose customers!

Problem #2: Expensive

- Need to invest many $$$ in hardware
  - Even a small cluster can easily cost $100,000
  - Microsoft recently invested $499 million in a single data center
- Need expertise
  - Planning and setting up a large cluster is highly nontrivial
  - Cluster may require special software, etc.
- Need maintenance
  - Someone needs to replace faulty hardware, install software upgrades, maintain user accounts, ...

Problem #3: Difficult to scale

- Scaling up is difficult
  - Need to order new machines, install them, integrate with existing cluster - can take weeks
  - Large scaling factors may require major redesign, e.g., new storage system, new interconnect, new building (!)
- Scaling down is difficult
  - What to do with superfluous hardware?
  - Server idle power is about 60% of peak $\rightarrow$ Energy is consumed even when no work is being done
  - Many fixed costs, such as construction

Recap: Computing at scale

- Modern applications require huge amounts of processing and data
  - Measured in petabytes, millions of users, billions of objects
  - Need special hardware, algorithms, tools to work at this scale
- Clusters and data centers can provide the resources we need
  - Main difference: Scale (room-sized vs. building-sized)
  - Special hardware; power and cooling are big concerns
- Clusters and data centers are not perfect
  - Difficult to dimension; expensive; difficult to scale

Plan for today

- Computing at scale
  - The need for scalability; scale of current services
  - Scaling up: From PCs to data centers
  - Problems with 'classical' scaling techniques
- Utility computing and cloud computing
  - What are utility computing and cloud computing? NEXT
  - What kinds of clouds exist today?
  - What kinds of applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'
  - Some cloud computing challenges

The power plant analogy

Waterwheel at the Neuhausen ob Eck Open-Air Museum

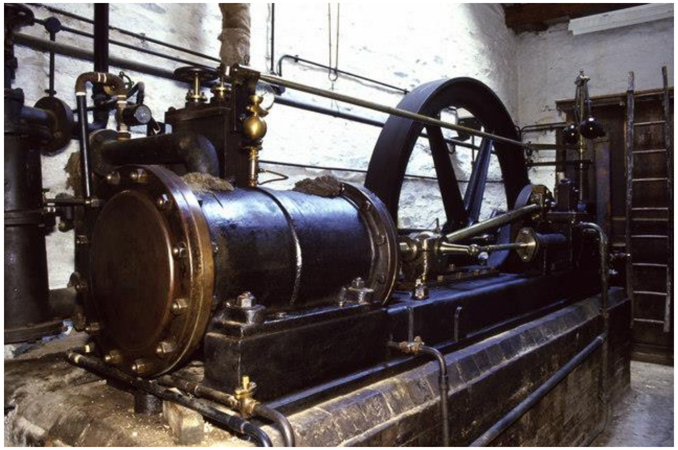

Steam engine at Stott Park Bobbin Mill

- It used to be that everyone had their own power source
  - Challenges are similar to the cluster: Needs large up-front investment, expertise to operate, difficult to scale up/down…

Scaling the power plant

- Then people started to build large, centralized power plants with very large capacity…

Metered usage model

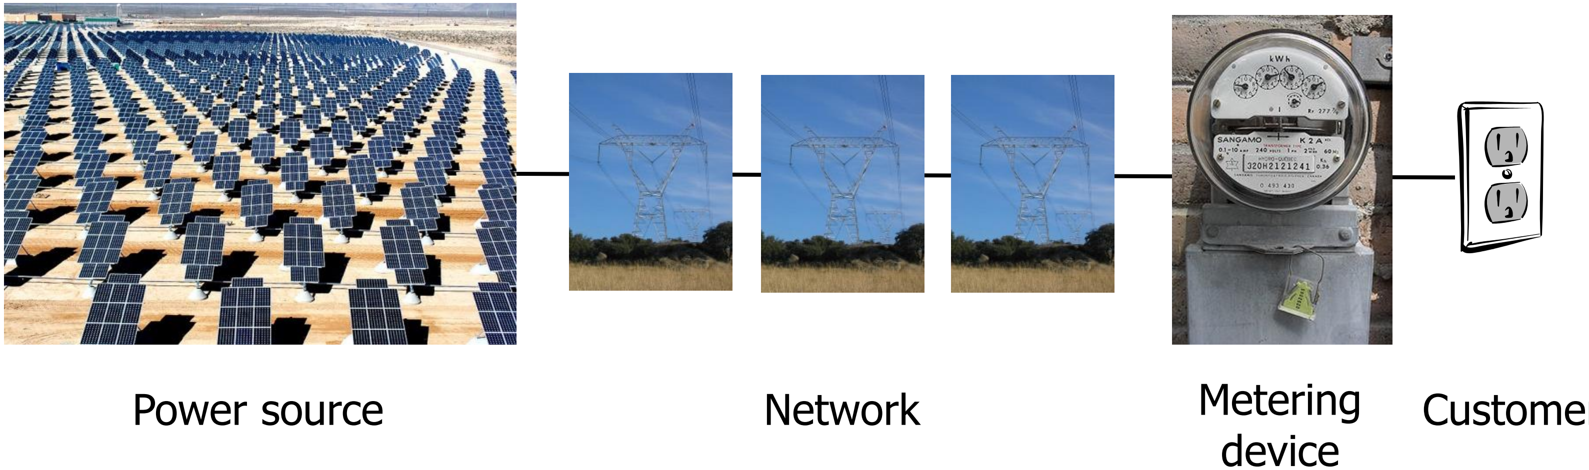

- Power plants are connected to customers by a network
- Usage is metered, and everyone (basically) pays only for what they actually use

Why is this a good thing?

	Electricity	Computing
Economies of scale	Cheaper to run one big power plant than many small ones	Cheaper to run one big data center than many small ones
Statistical multiplexing	High utilization!	High utilization!
No up-front commitment	No investment in generator; pay-as-you-go model	No investment in data center; pay-as-you-go model
Scalability	Thousands of kilowatts available on demand; add more within seconds	Thousands of computers available on demand; add more within seconds

What is cloud computing?

http://www.dilbert.com/fast/2013-06-29/

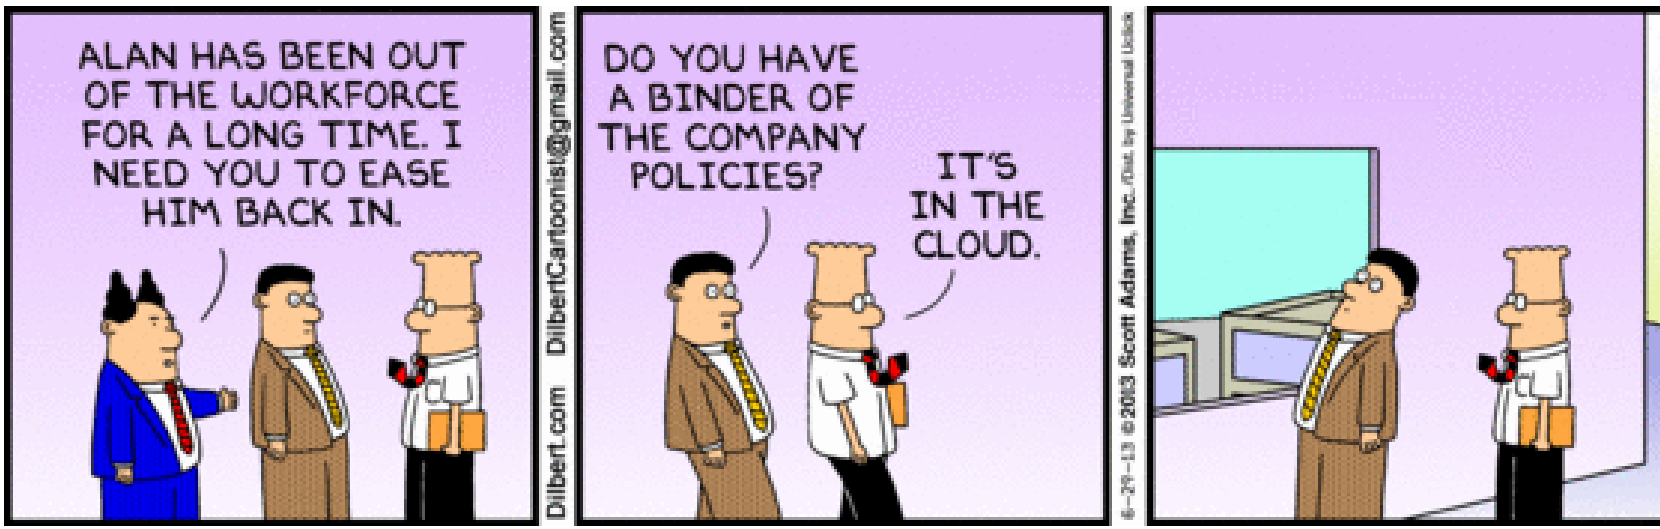

What is cloud computing?

The interesting thing about Cloud Computing is that we've redefined Cloud Computing to include everything that we already do.... I don't understand what we would do differently in the light of Cloud Computing other than change the wording of some of our ads.

Larry Ellison, quoted in the Wall Street Journal, September 26, 2008

A lot of people are jumping on the [cloud] bandwagon, but I have not heard two people say the same thing about it. There are multiple definitions out there of "the cloud".

Andy Isherwood, quoted in ZDnet News, December 11, 2008

So what is it, really?

- According to NIST:

  Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.
- Essential characteristics:
  - On-demand self service
  - Broad network access
  - Resource pooling
  - Rapid elasticity
  - Measured service

Other terms you may have heard

- Utility computing
  - The service being sold by a cloud
  - Focuses on the business model (pay-as-you-go), similar to classical utility companies
- The Web
  - The Internet's information sharing model
  - Some web services run on clouds, but not all
- The Internet
  - A network of networks.
  - Used by the web; connects (most) clouds to their customers

Plan for today

- Computing at scale
  - The need for scalability; scale of current services
  - Scaling up: From PCs to data centers
  - Problems with 'classical' scaling techniques
- Utility computing and cloud computing
  - What are utility computing and cloud computing?
  - What kinds of clouds exist today? NEXT
  - What kinds of applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'
  - Some cloud computing challenges

Everything as a Service

- What kind of service does the cloud provide?
  - Does it offer an entire application, or just resources?
  - If resources, what kind / level of abstraction?
- Three types commonly distinguished:
  - Software as a service (SaaS)
    - Analogy: Restaurant. Prepares&serves entire meal, does the dishes, …
  - Platform as a service (PaaS)
    - Analogy: Take-out food. Prepares meal, but does not serve it.
  - Infrastructure as a service (IaaS)
    - Analogy: Grocery store. Provides raw ingredients.
  - Other xaaS types have been defined, but are less common
    - Desktop, Backend, Communication, Network, Monitoring, …

Software as a Service (SaaS)

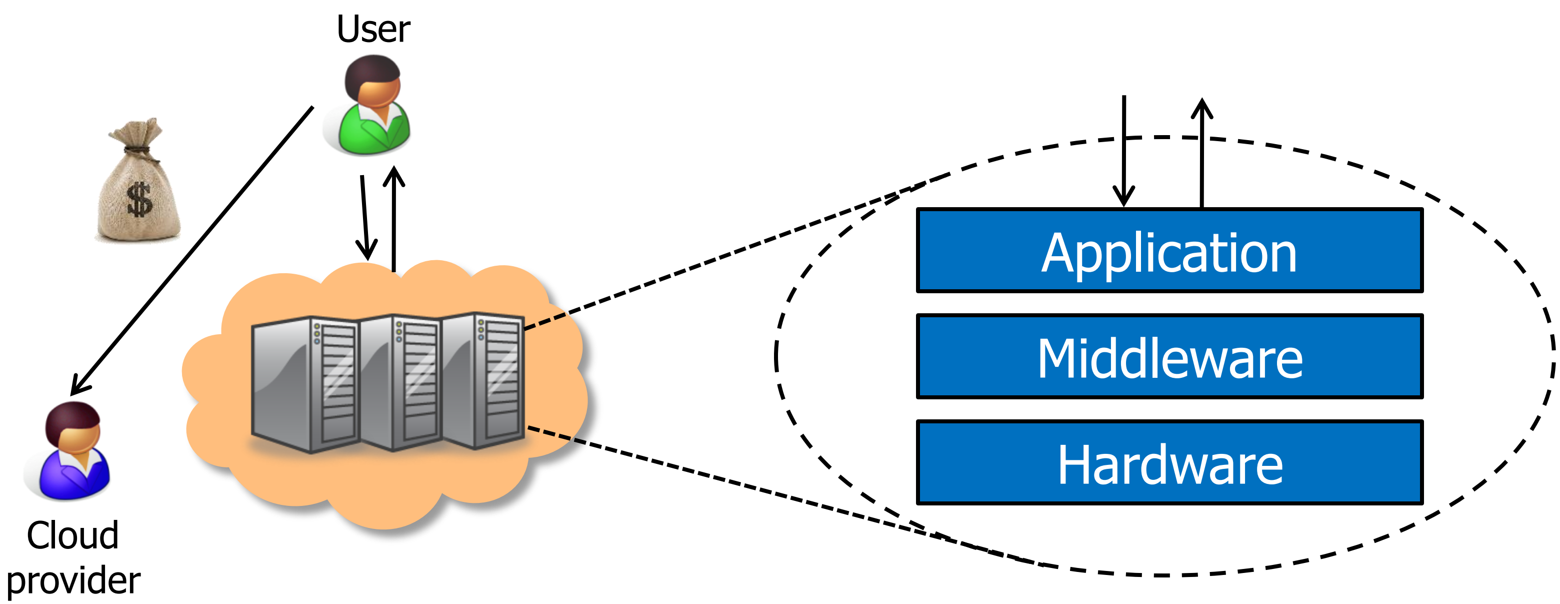

- Cloud provides an entire application
  - Word processor, spreadsheet, CRM software, calendar…
  - Customer pays cloud provider
  - Example: Google Apps, Salesforce.com

Platform as a Service (PaaS)

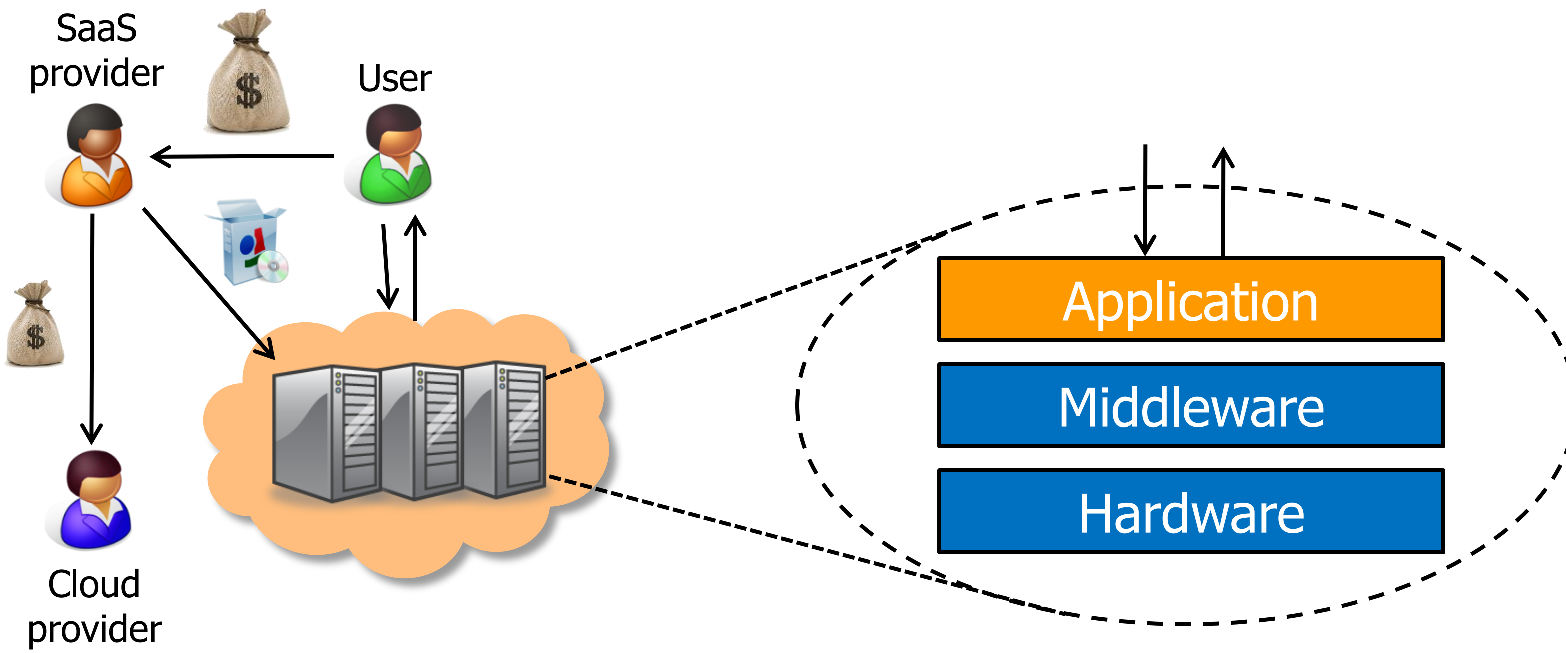

- Cloud provides middleware/infrastructure
  - For example, Microsoft Common Language Runtime (CLR)
  - Customer pays SaaS provider for the service; SaaS provider pays the cloud for the infrastructure
  - Example: Windows Azure, Google App Engine

Infrastructure as a Service (IaaS)

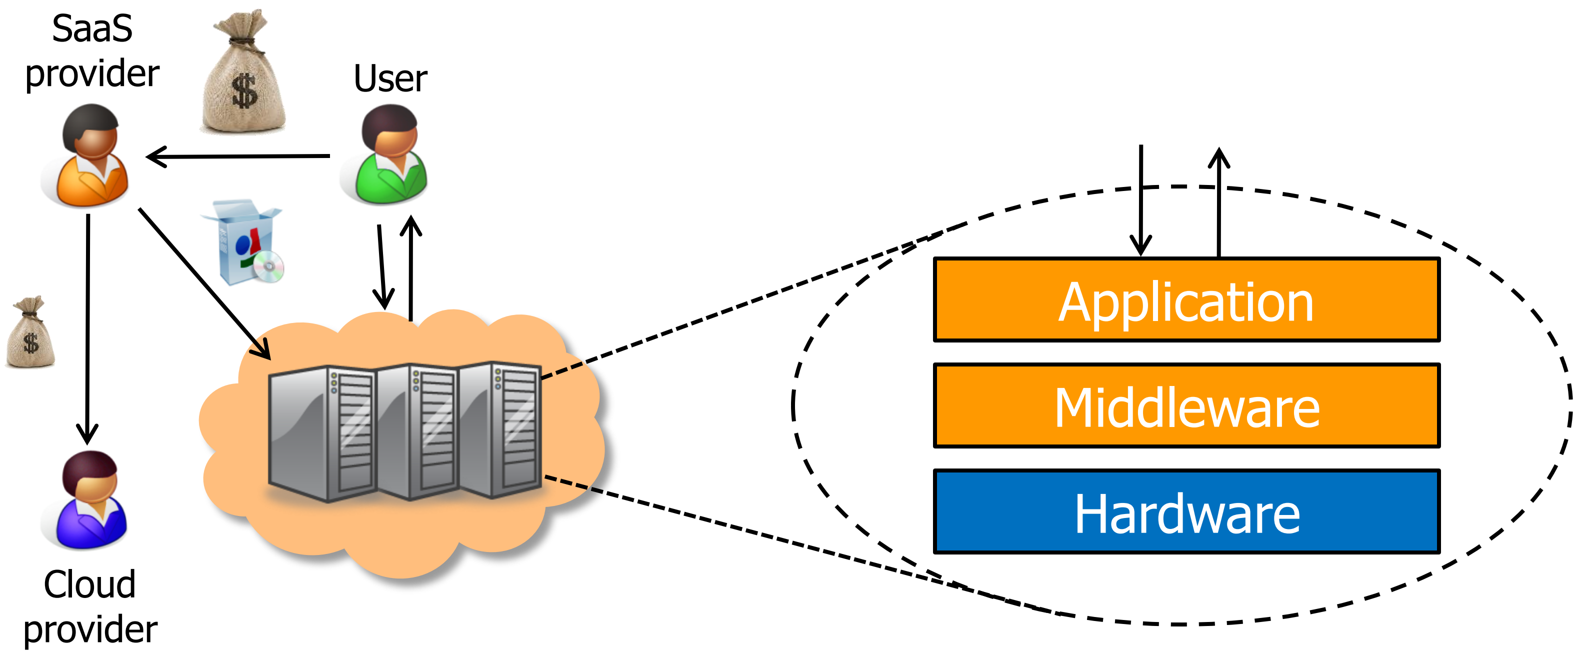

- Cloud provides raw computing resources
  - Virtual machine, blade server, hard disk, ...
  - Customer pays SaaS provider for the service; SaaS provider pays the cloud for the resources
  - Examples: Amazon Web Services, Rackspace Cloud, GoGrid

Private/hybrid/community clouds

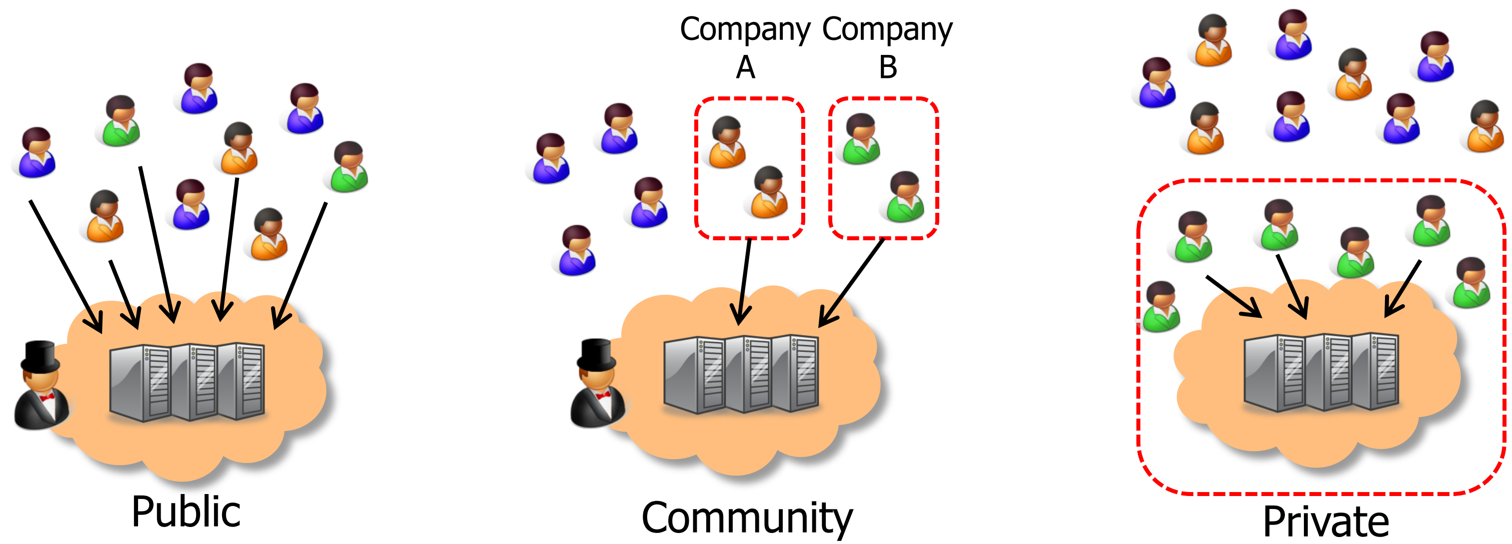

- Who can become a customer of the cloud?
  - Public cloud: Commercial service; open to (almost) anyone. Example: Amazon AWS, Microsoft Azure, Google App Engine
  - Community cloud: Shared by several similar organizations. Example: Google's "Gov Cloud"
  - Private cloud: Shared within a single organization. Example: Internal datacenter of a large company.

New Paradigm for Large Scale Data Processing

- Need to process lots of data (> 1 TB)
- Parallelize across hundreds/thousands of CPUs
- … Want to make this easy
- Approach
  - Divide and conquer
  - Split processing across multiple CPUs
  - Integrate outputs
- Network is “virtual”
  - No fixed structure, but dynamic structure
- Traditional networks (telephone, etc.)
  - Fixed structure

Divide and Conquer

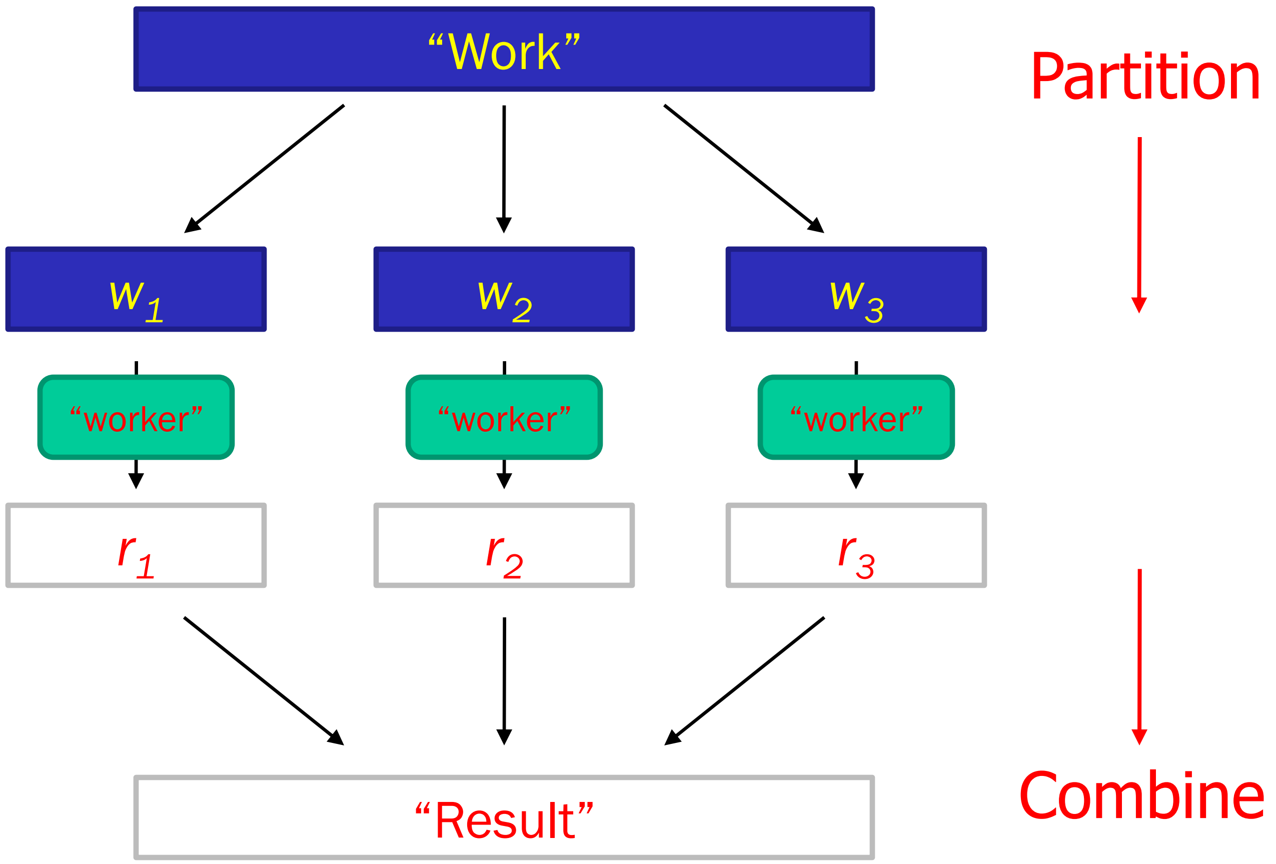

MapReduce Framework

- Divide and conquer
- Makes use of thousands of CPUs

- Network structure is demand-driven
  - Access CPUs/storage based on dynamic jobs
  - Framework (e.g., Hadoop) creates *network on demand*

Parallelization Challenges

- How do we assign work units to workers?
- What if we have more work units than workers?
- What if workers need to share partial results?
- How do we aggregate partial results?
- How do we know all the workers have finished?
- What if workers die?

Examples of cloud applications

- Application hosting
- Backup and Storage
- Content delivery
- E-commerce
- High-performance computing
- Media hosting
- On-demand workforce
- Search engines
- Web hosting

Other examples

- DreamWorks is using the Cerelink cloud to render animation movies
  - Cloud was already used to render parts of *Shrek Forever After* and *How to Train your Dragon*

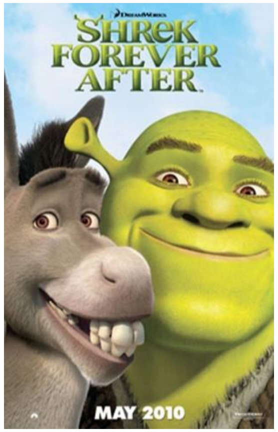

- CERN is working on a "science cloud" to process experimental data
- Virgin atlantic is hosting their new travel portal on Amazon AWS

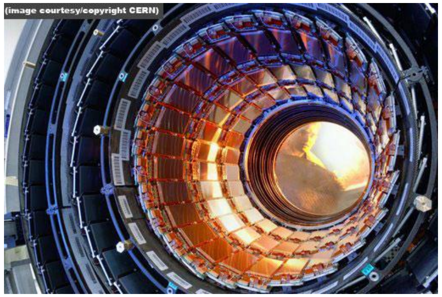

Recap: Cloud applications

- Clouds are good for many things...
  - Applications that involve large amounts of computation, storage, bandwidth
  - Especially when lots of resources are needed quickly (Washington Post example) or load varies rapidly (TicketLeap example)

- ... but not for all things

Stay tuned

Next lecture:
Details of MapReduce