COSC – 121

Computer Systems

Wenchao Zhou

A Case Study: Cloud Computing

Lecture Today

- Scalable 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 applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'

2016 What happens in an INTERNET MINUTE?

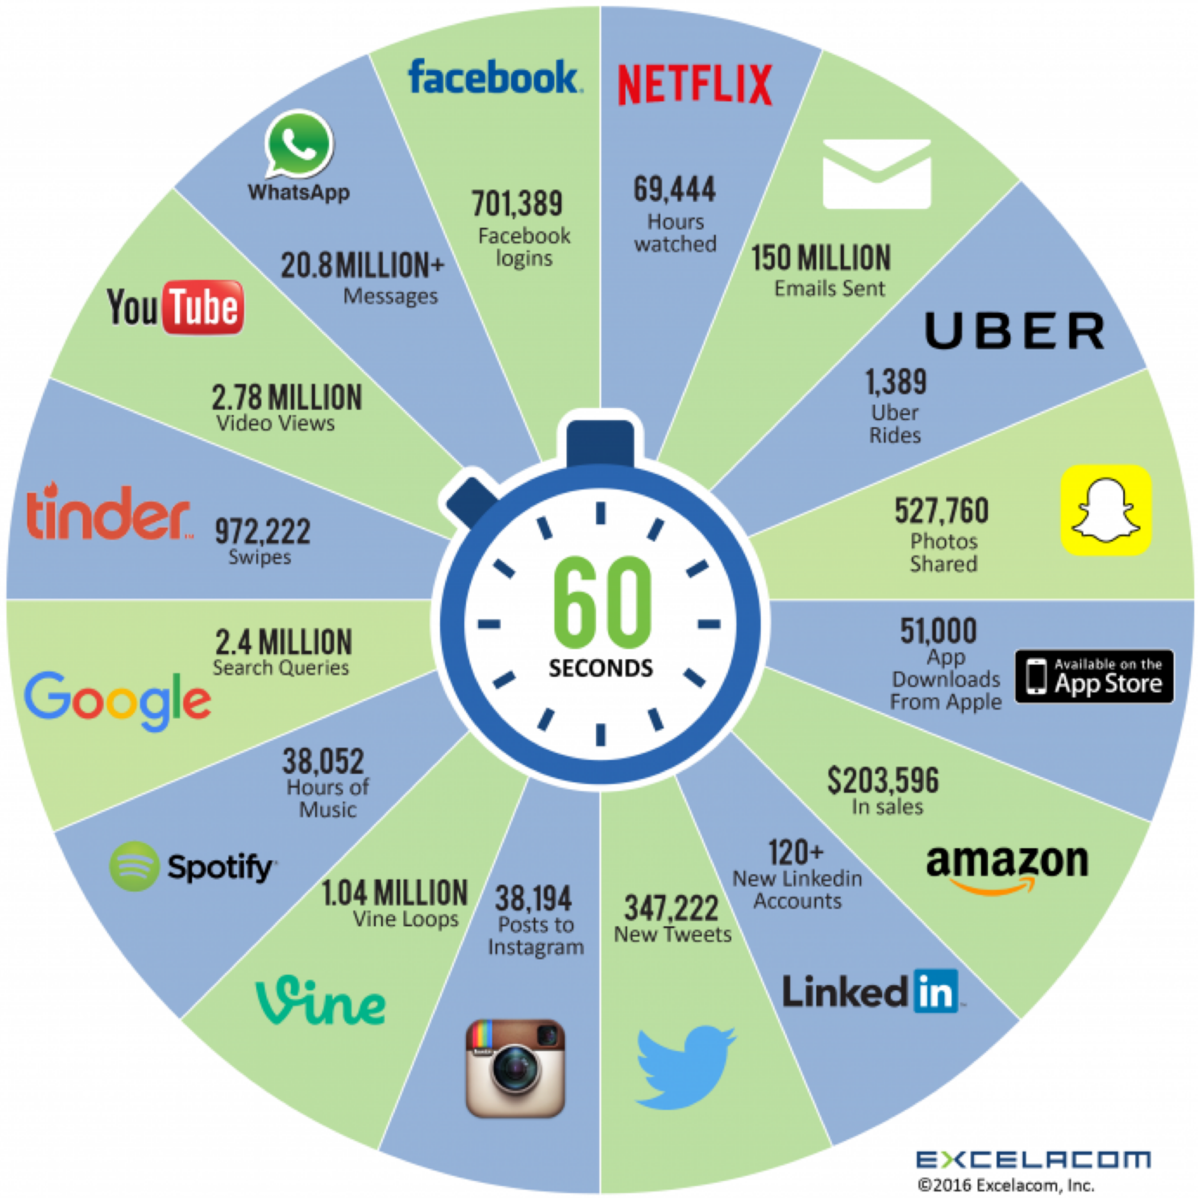

2019 This Is What Happens In An Internet Minute

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
  - Google designs for 1 exabyte of storage (!)
- How much is a petabyte?
  - 1,000,000,000,000,000 bytes
  - A stack of 1TB hard disks that is 25.4 meters 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 10 million servers

Scaling up

PC

Server

Cluster

- What if one computer is not enough?
  - Buy a bigger (server-class) computer
- What if the biggest computer is not enough?
  - Buy many computers

Clusters

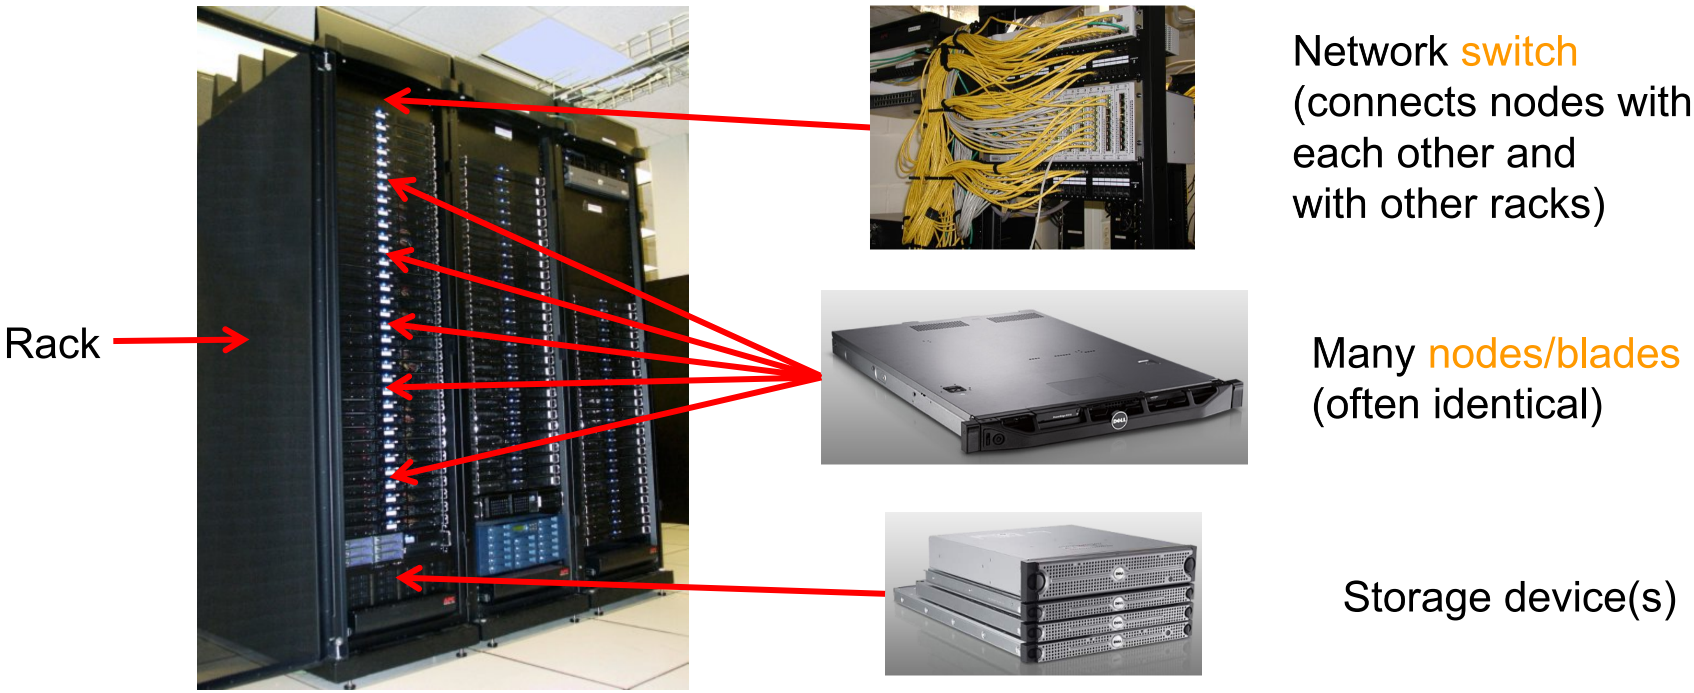

- Characteristics of a cluster:
  - Many similar machines, close interconnection (same room?)
  - Often special, standardized hardware (racks, blades)
  - Usually owned and used by a single organization

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)

Scaling up

PC

Server

Cluster

Data center

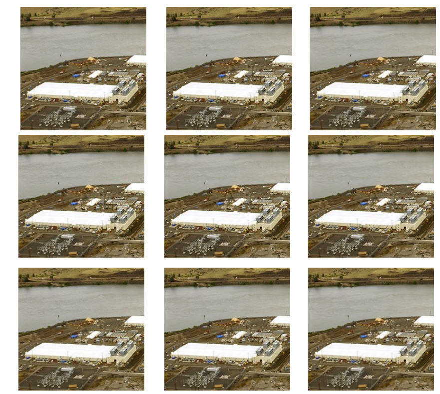

Network of data centers

- 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

Rest of the Lecture

- Scalable 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 applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'

Problem #1: Difficult to dimension

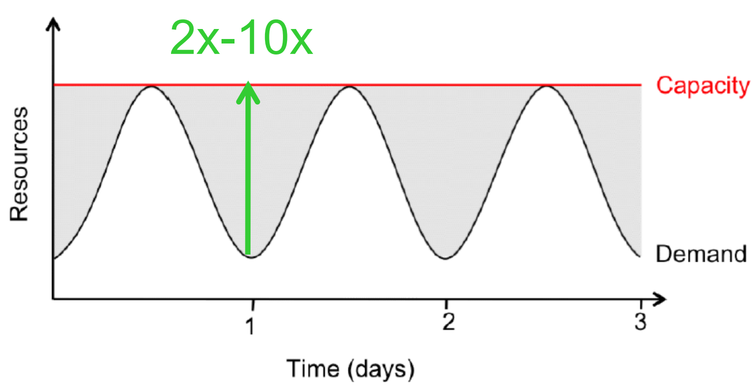

Provisioning for the peak load

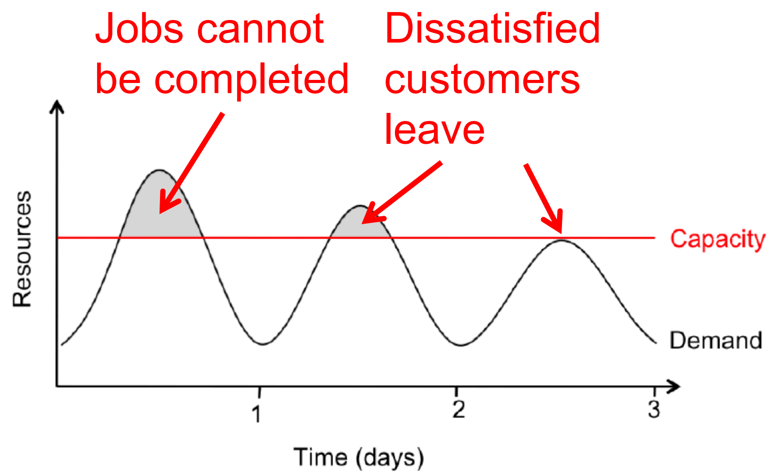

Provisioning below the peak

- Problem: Load can vary considerably
  - Peak load can exceed average load by factor 2x-10x [Why?]
  - 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 → Energy is consumed even when no work is being done
  - Many fixed costs, such as construction

Recap: Scalability

- Modern applications require huge amounts of processing and data
  - Measured in petabytes, millions of users, billions of objects
  - Need special hardware, algorithms, tools at this scale
- Clusters and data centers can provide the resources
  - 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

Rest of the Lecture

- Scalable 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 applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'

The power plant analogy

Waterwheel at the Neuhausen ob Eck Open-Air Museum

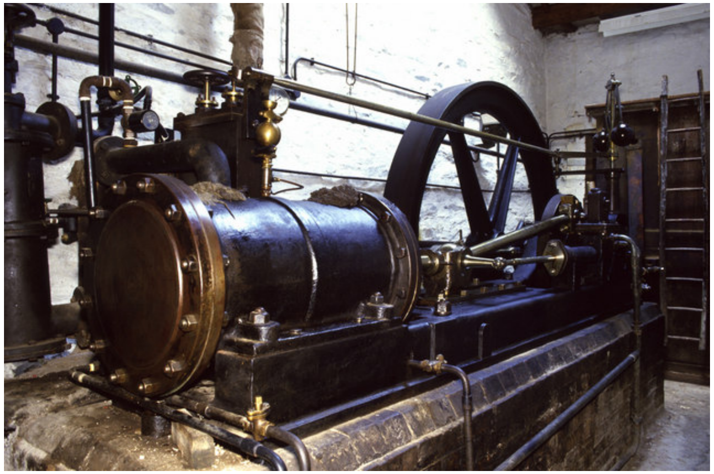

Steam engine at Stott Park Bobbin Mill

- 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

So what is cloud computing, 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

Rest of the Lecture

- Scalable 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 applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'

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

Case study: The Washington Post

- March 19, 2008: Hillary Clinton's official White House schedule released to the public
  - 17,481 pages of non-searchable, low-quality PDF
  - Very interesting to journalists, but would have required hundreds of man-hours to evaluate
  - Peter Harkins, Senior Engineer at The Washington Post: Can we make that data available more quickly, ideally within the same news cycle?
  - Tested various Optical Character Recognition (OCR) programs; estimated required speed
  - Launched 200 EC2 instances; project was completed within nine hours (!) using 1,407 hours of VM time ($144.62)
  - Results available on the web only 26 hours after the release

http://aws.amazon.com/solutions/case-studies/washington-post/

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

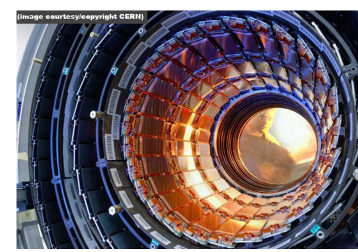

Cloud applications

- Clouds are good for many things...
  - Applications that involve large amounts of computation, storage, bandwidth
  - Especially when resources are needed quickly (Washington Post example) or load varies rapidly (CERN LHD example)
- ... but not for all things
  - Sometimes it is problematic, e.g., because of auditability requirements
  - Example: Processing medical records
    - HIPAA (Health Insurance Portability and Accountability Act) privacy and security rule
  - Example: Processing financial information
    - Sarbanes-Oxley act

Rest of the Lecture

- Scalable 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 applications run on the cloud?
  - Virtualization: How clouds work 'under the hood'

What is virtualization?

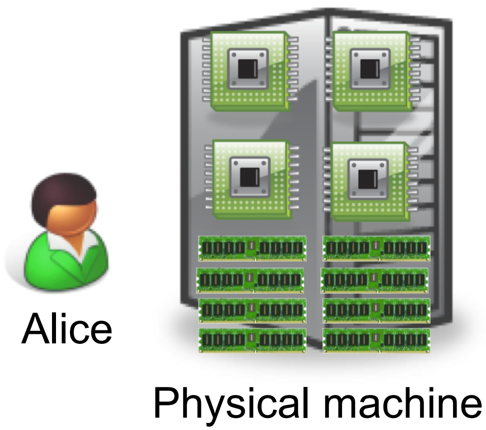

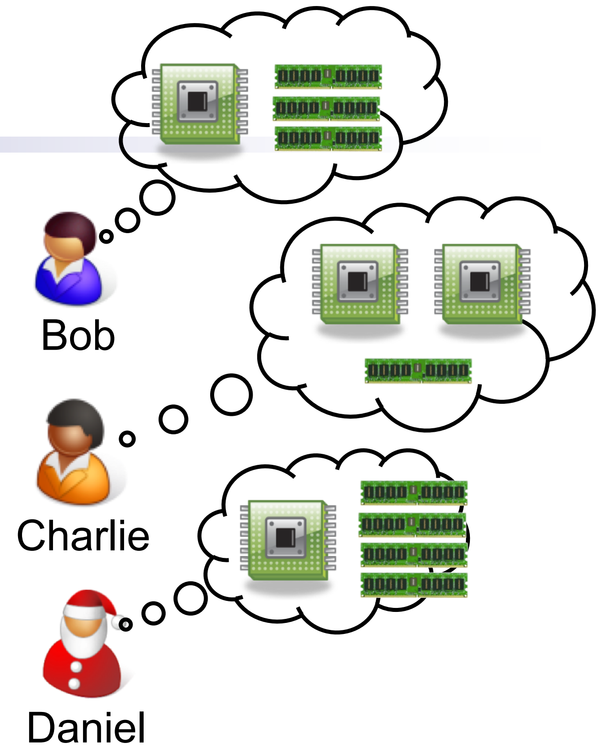

- Suppose Alice has a machine with 4 CPUs and 8 GB of memory, and three customers:
  - Bob wants a machine with 1 CPU and 3GB of memory
  - Charlie wants 2 CPUs and 1GB of memory
  - Daniel wants 1 CPU and 4GB of memory

What is virtualization?

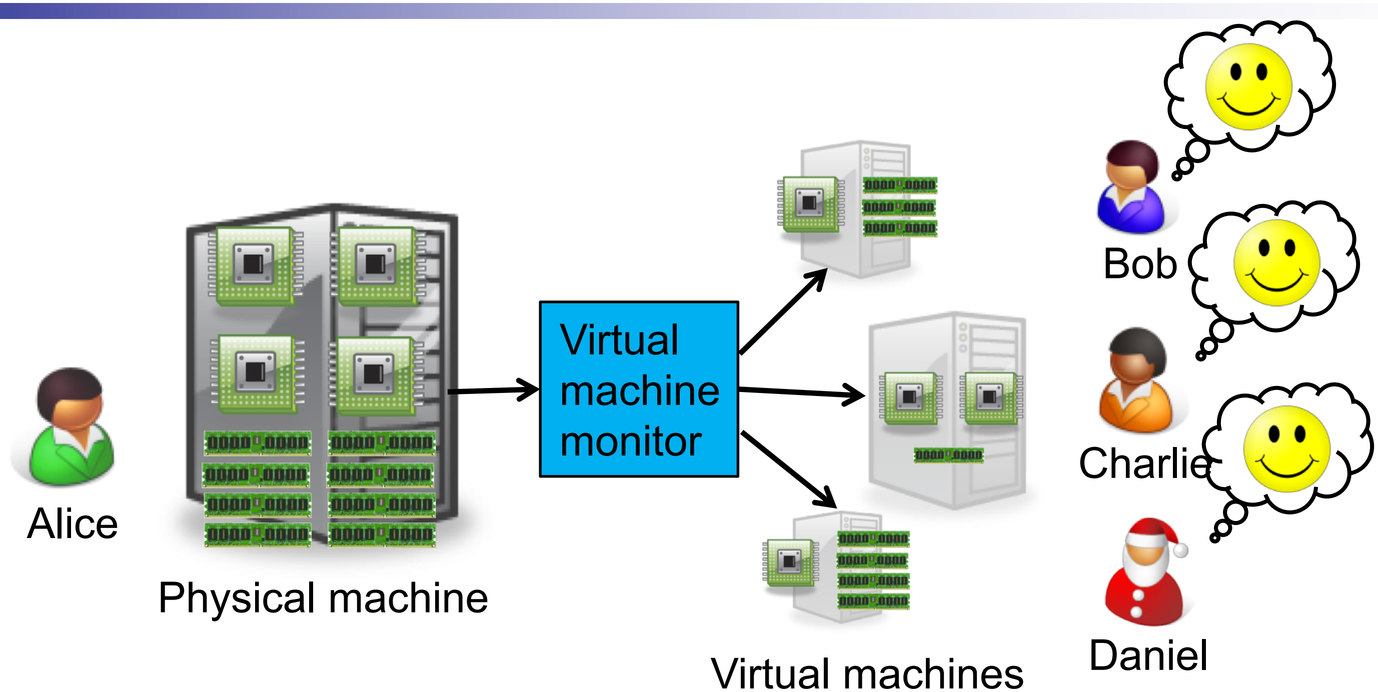

- Alice can sell each customer a virtual machine (VM) with the requested resources
  - From each customer's perspective, it appears as if they had a physical machine all by themselves (isolation)

Benefit: Migration

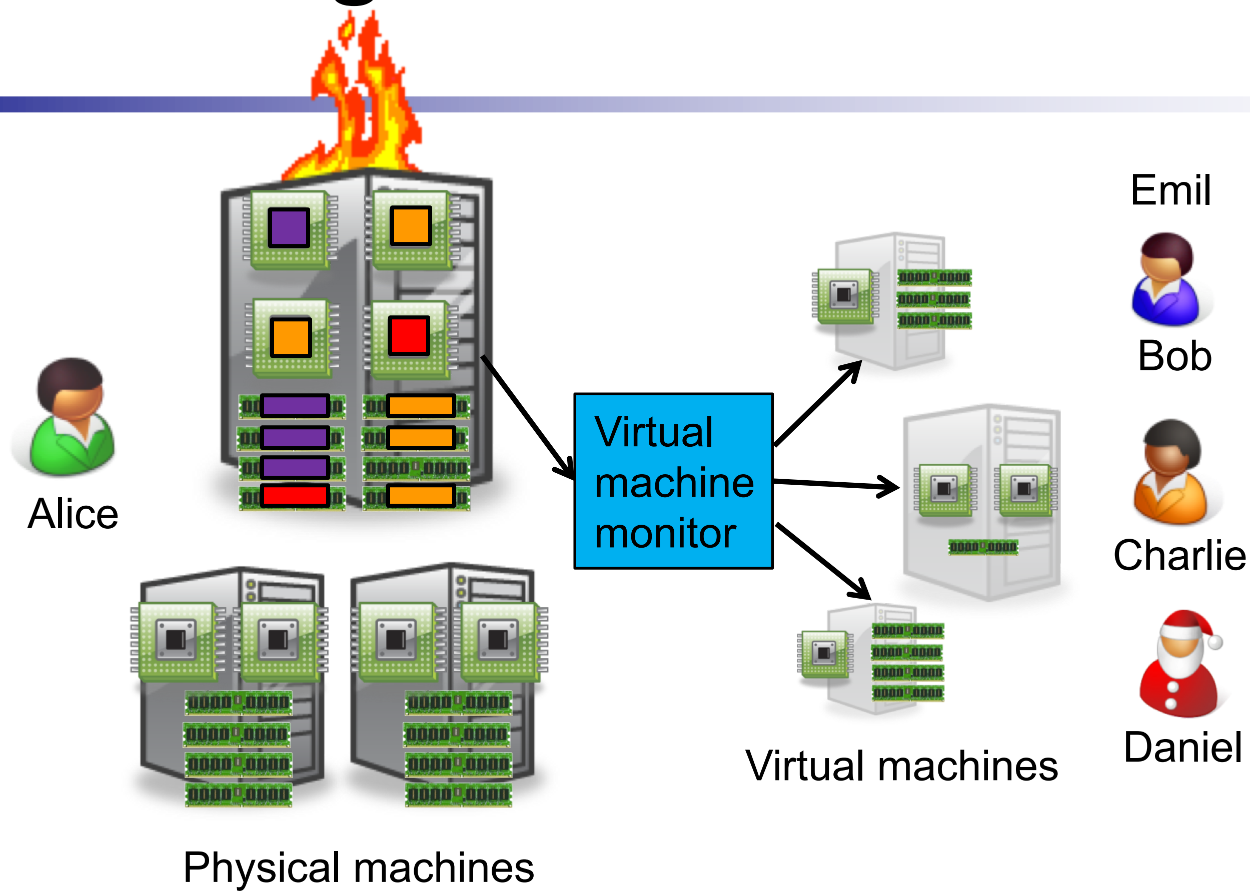

- What if the physical machine fails?
  - Alice can migrate the VMs to different physical machines without any customers noticing

Benefit: Time sharing

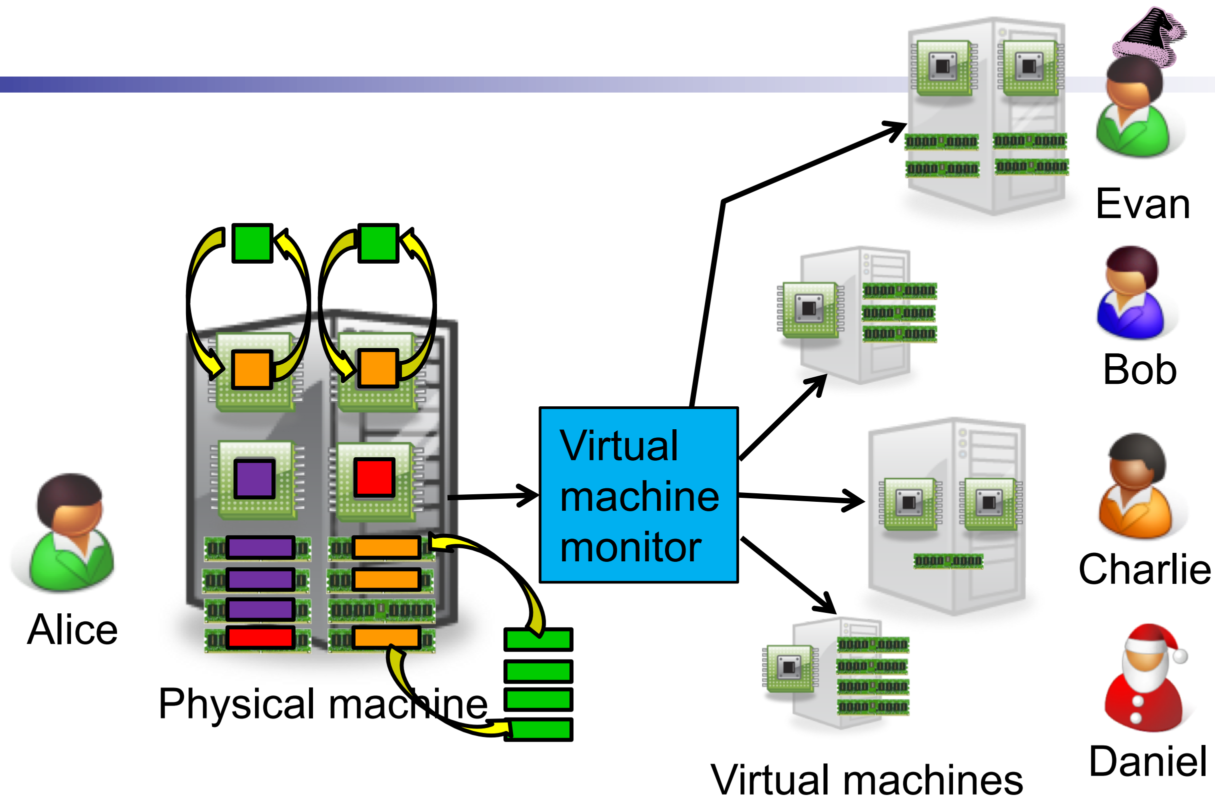

- What if Alice gets another customer?
  - Multiple VMs can time-share the existing resources
  - Result: Alice has more virtual CPUs and virtual memory than physical resources (but not all can be active at the same time)

Recap: Virtualization in the cloud

- Gives cloud provider a lot of flexibility
  - Can produce VMs with different capabilities
  - Can migrate VMs if necessary (e.g., for maintenance)
  - Can increase load by overcommitting resources
- Provides security and isolation
  - Programs in one VM cannot influence programs in another
- Convenient for users
  - Complete control over the virtual 'hardware' (can install own operating system own applications, ...)