Développement logiciel pour le Cloud (TLC)

Introduction

Quentin Dufour

Goals

- Understand functionalities offered by Cloud computing
- Understand which issues are solved (or not) by the cloud
- Understand how cloud computing platforms are organized internally
- Understand how software developers can make use of these offerings

Course and Instructors

- CM course
  - Quentin Dufour quentin.dufour@inria.fr
- Practical Sessions
  - Quentin Dufour quentin.dufour@inria.fr

Introduction to Cloud Computing

Développement logiciel pour le Cloud (TLC)

Quentin Dufour

(credits: Davide Frey, François Taiani, Guillaume Pierre)

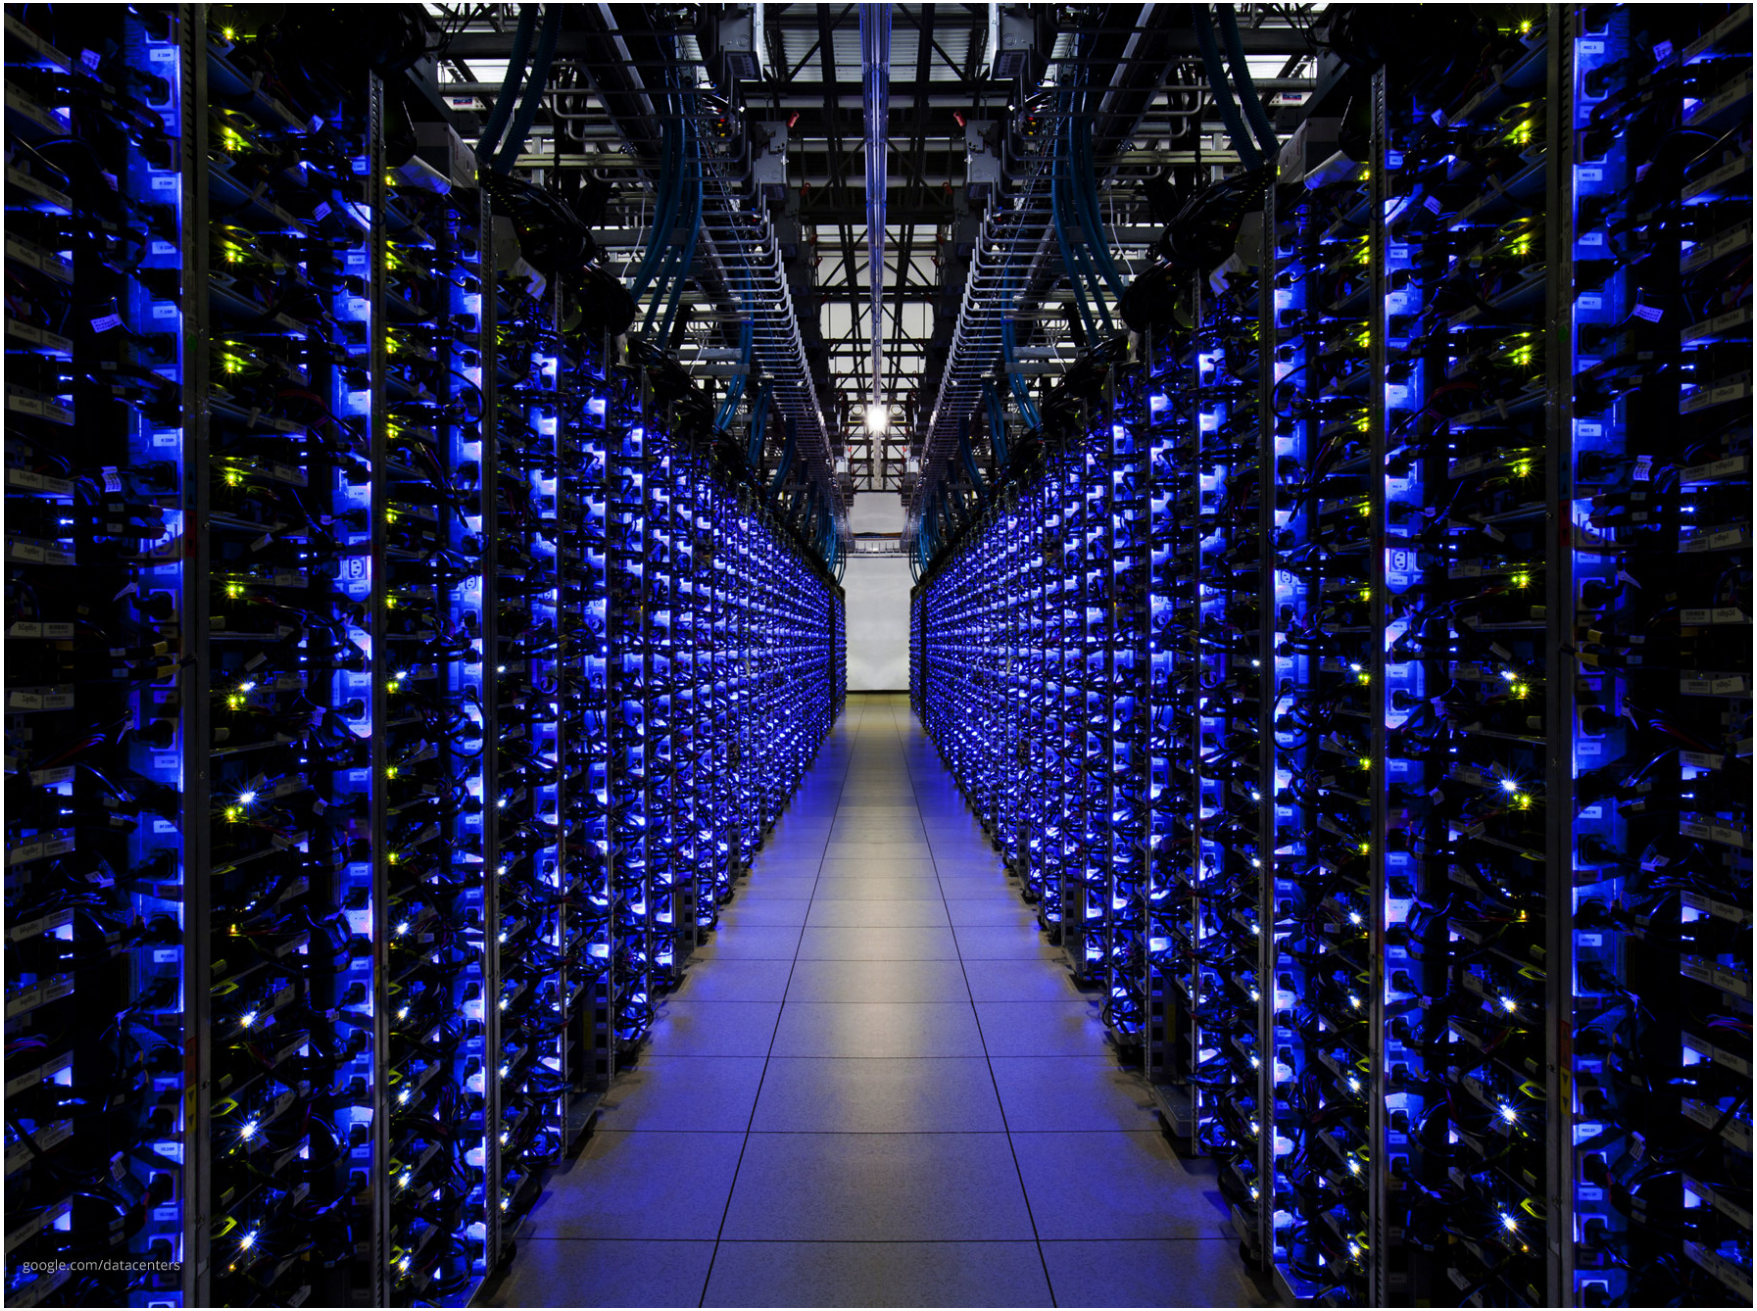
google.com/datacenters

https://lafibre.info/scaleway/dc5/
https://lafibre.info/scaleway/dc3-iliad/

Traditional system architectures

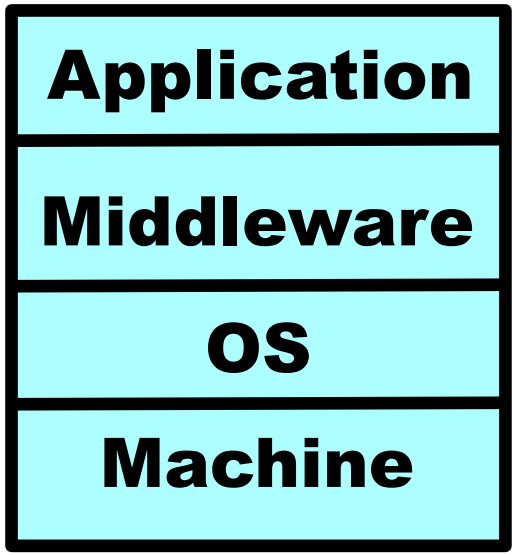

Traditional architecture

Traditional system architectures

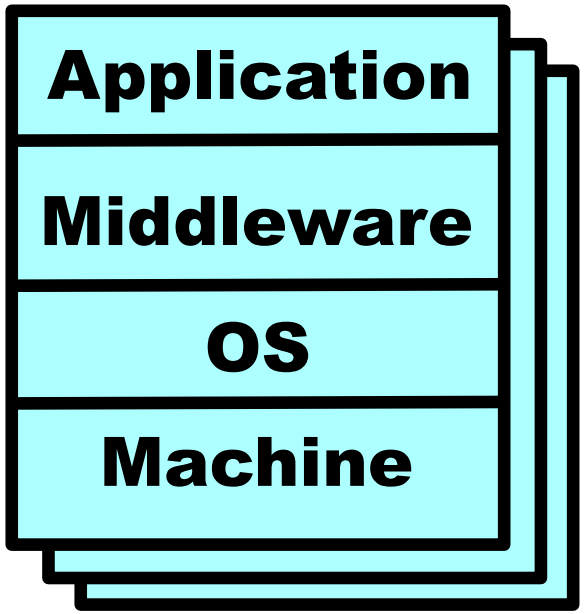

Traditional architecture

The varying capacity problem

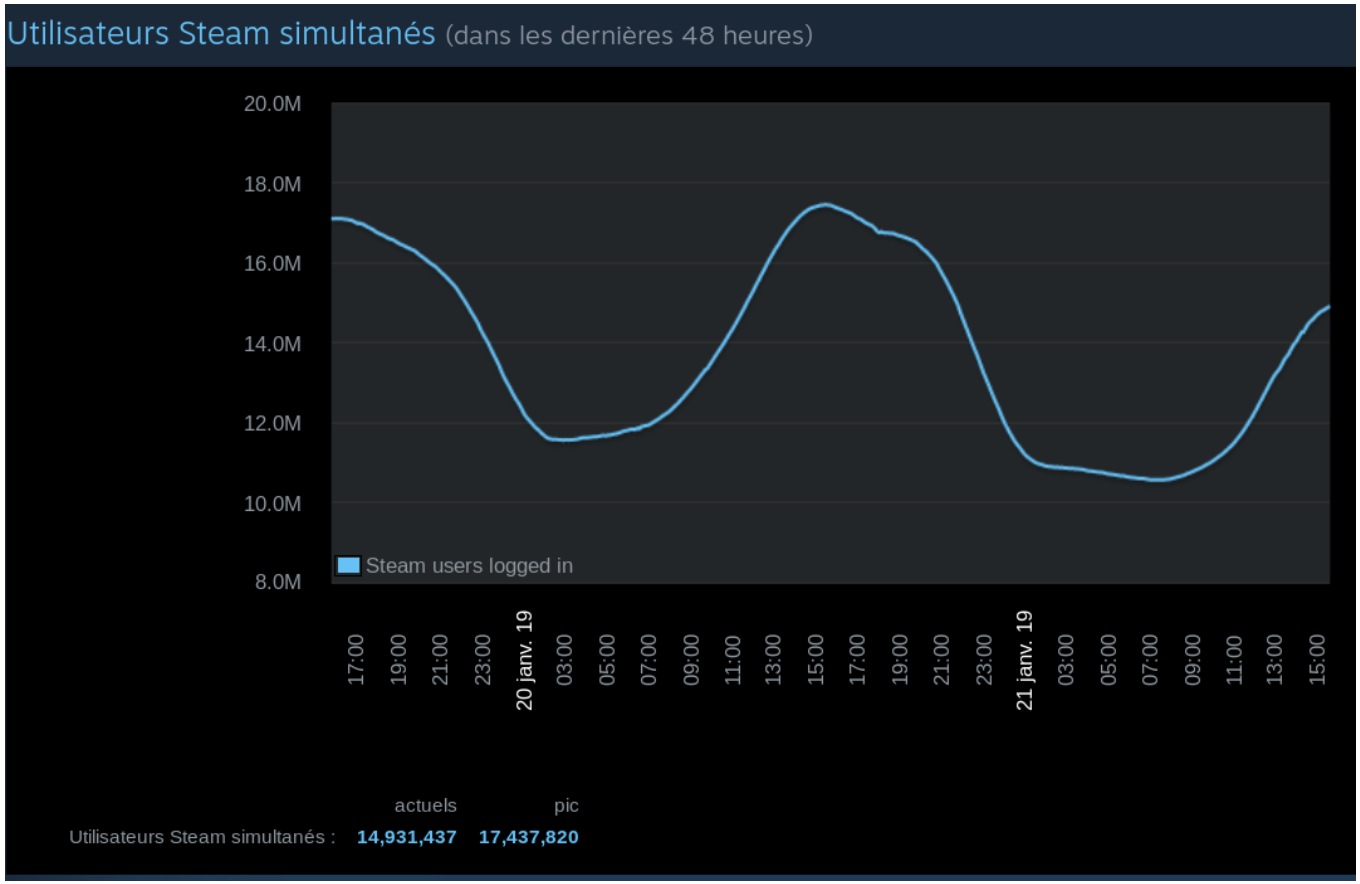

The varying capacity problem

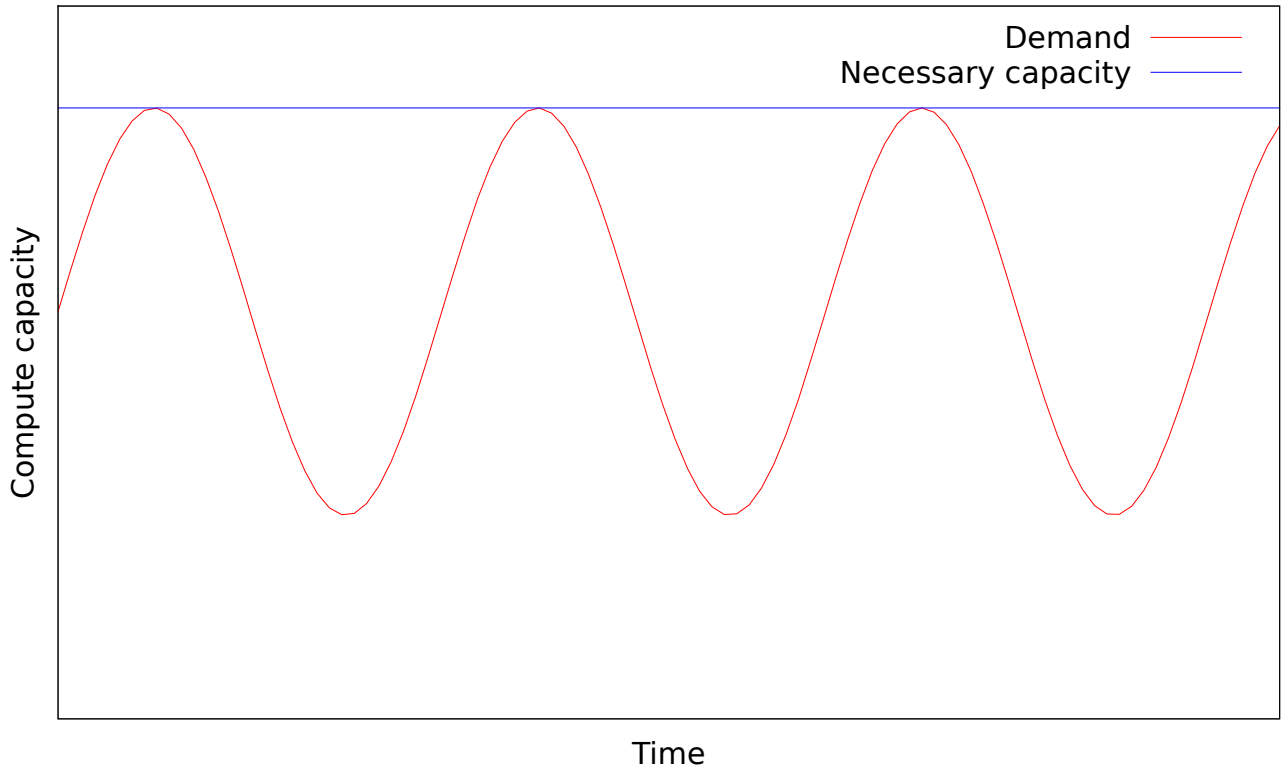

The varying capacity problem

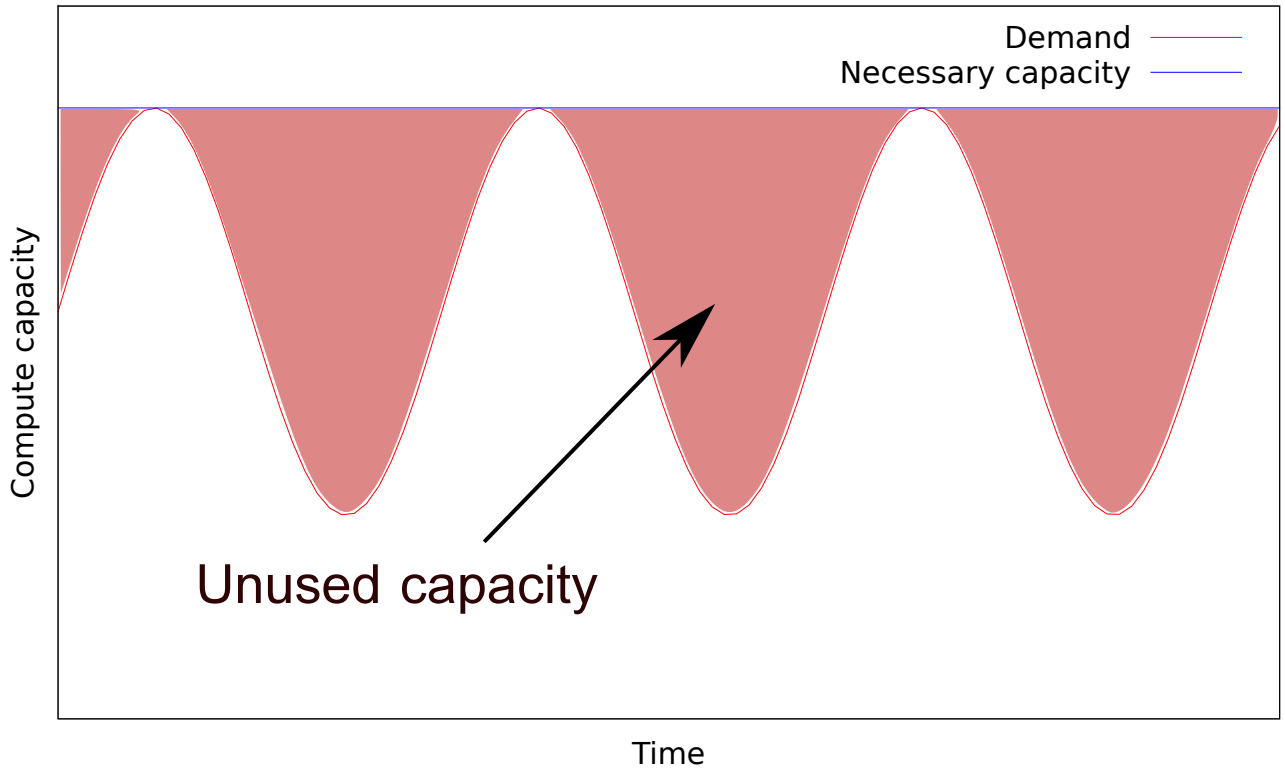

The varying capacity problem

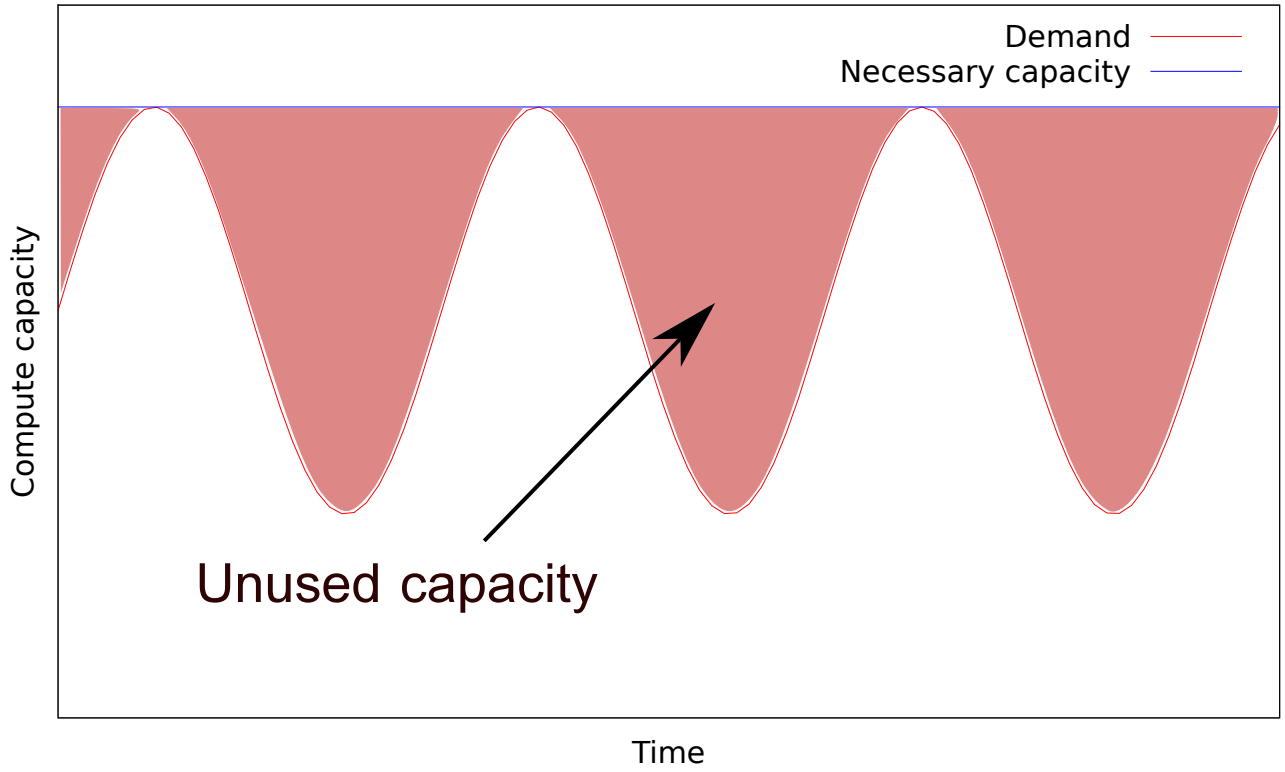

What if demand increases beyond the capacity?

Cloud Computing

- ☹ Difficult to vary capacity!
- ☹ Manual resource management

- ☺ Resources available on demand
- ☺ Resource management is fully automated
- ☺ Pay only for what you use

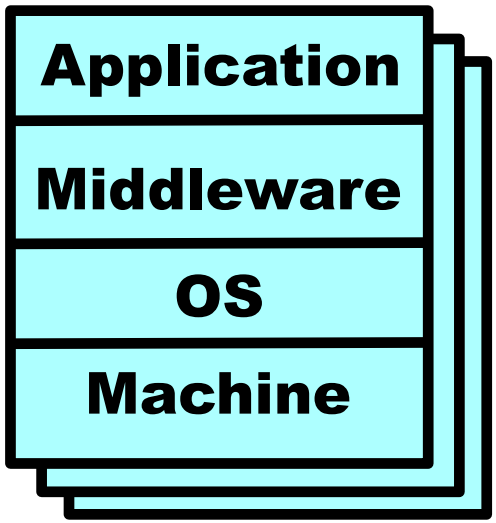

Traditional architecture

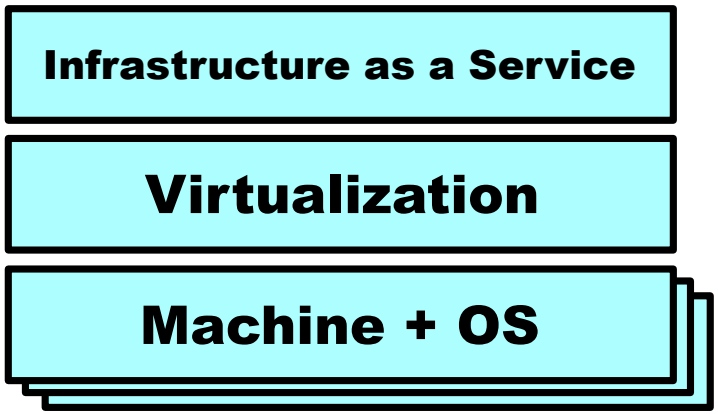

Cloud architecture

Cloud Computing

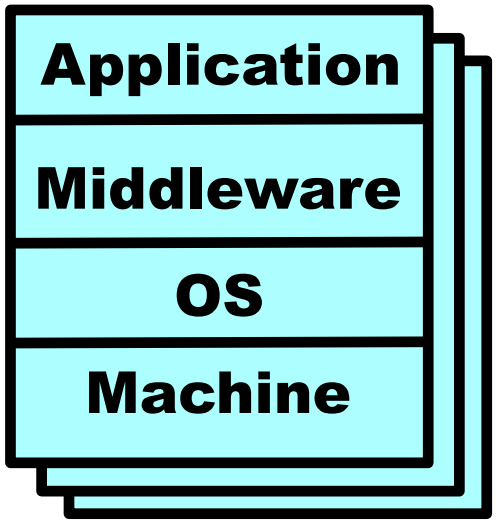

Traditional architecture

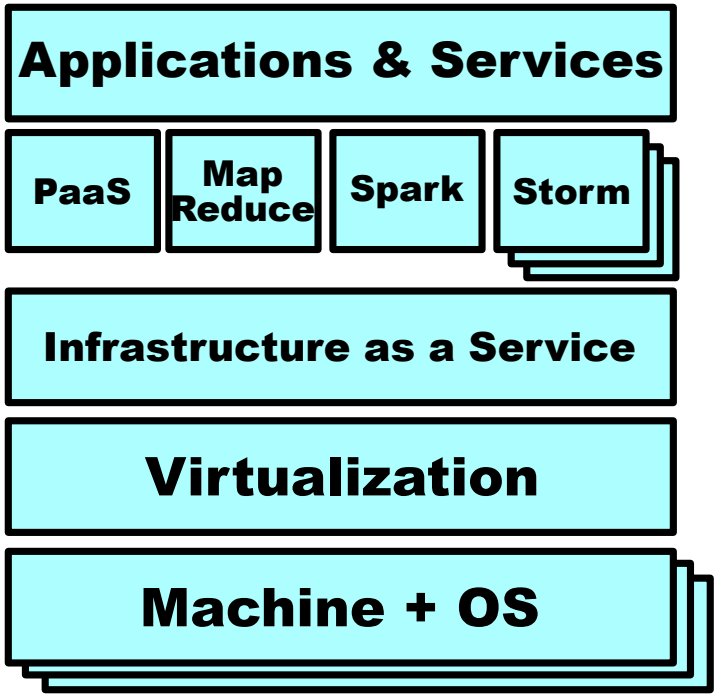

Cloud architecture

Cloud Computing

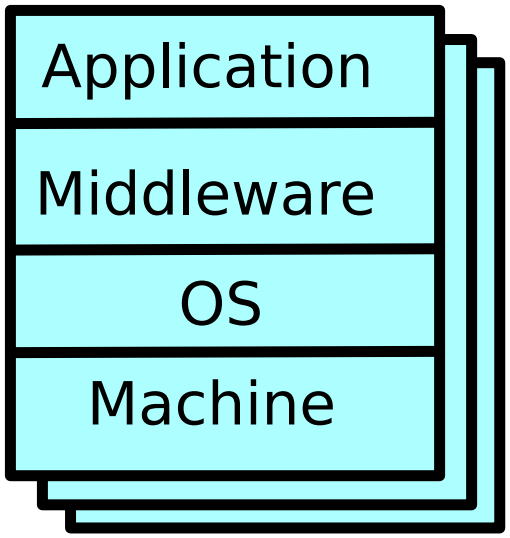

Traditional architecture

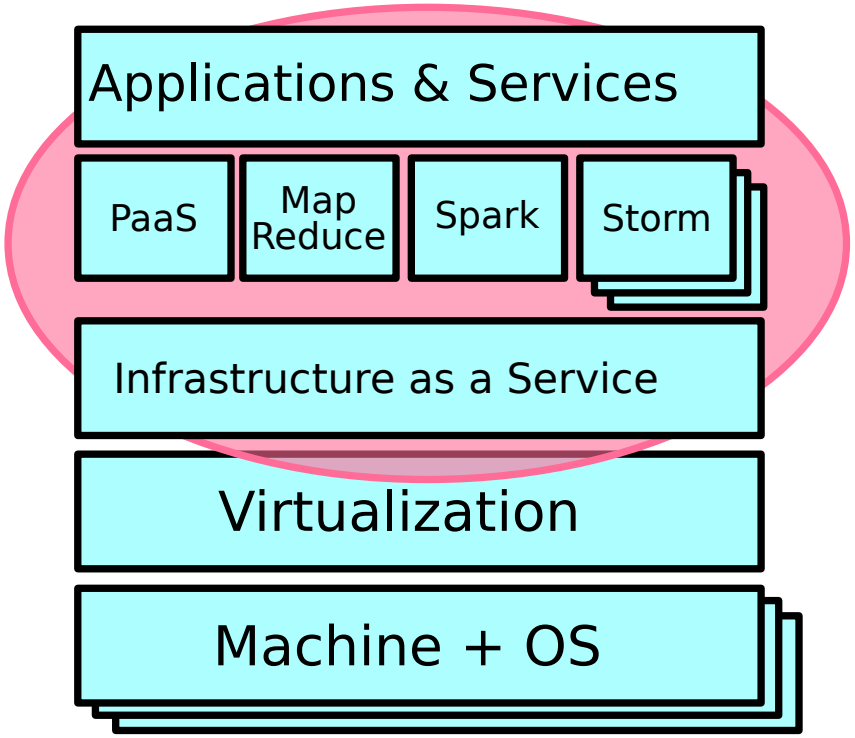

Cloud architecture

What defines Cloud computing exactly?

1. Computing is considered as a service customers use, not as hardware they own
   - This is called utility computing
2. Cloud providers offer a collection of compute/storage/network services via the Internet
   - Customers cannot get physical access to the devices
   - The actual location of devices is (almost) irrelevant
3. The cloud hides the complexity and details of the physical infrastructure from its users
   - Users only see a simple API + a graphical interface
4. Services are available on demand
   - Always available, anywhere, anytime
5. Pay-per-use
   - Pay only for the resources you actually use. You can release resources any time and stop paying immediately.

Who uses the cloud?

Everybody USES THE CLOUD. . .

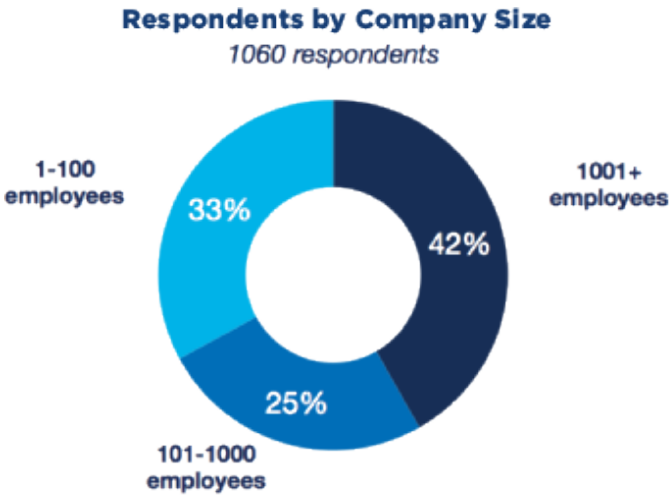

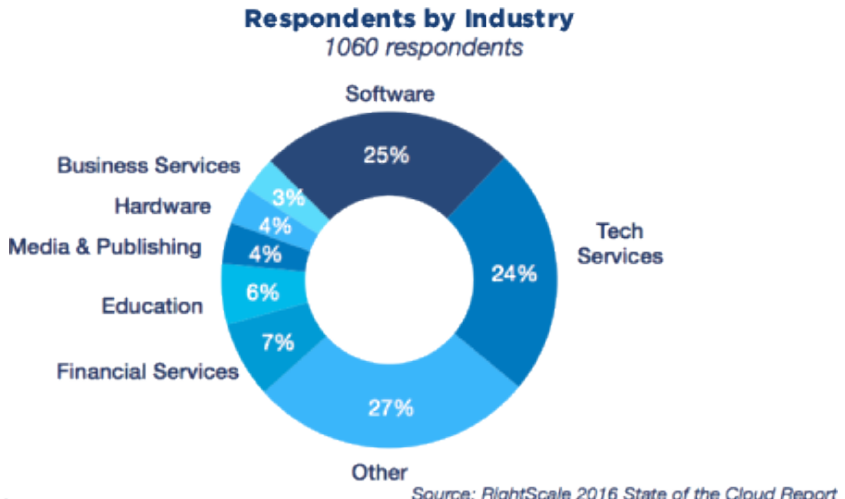

Virtualization

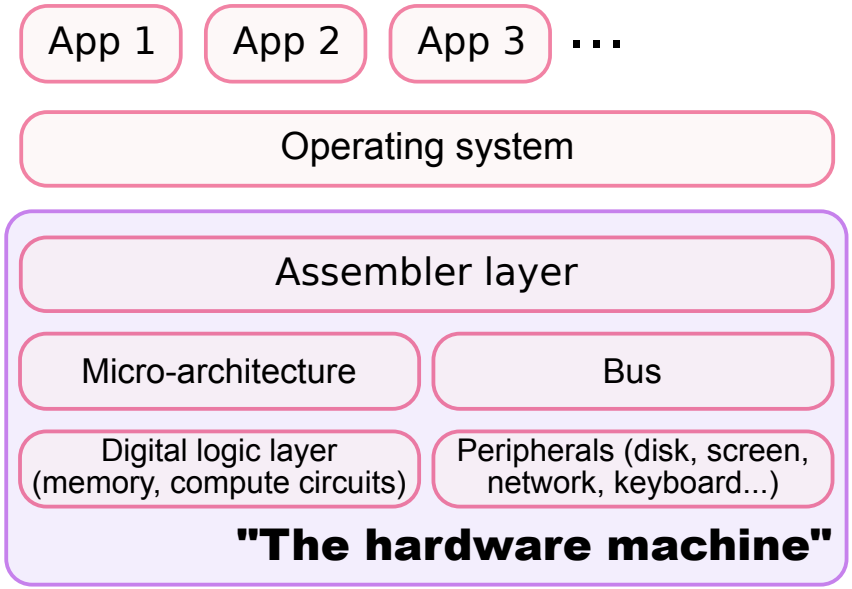

Traditional machine architecture:

- Applications
- Operating system
- Hardware

Virtualization

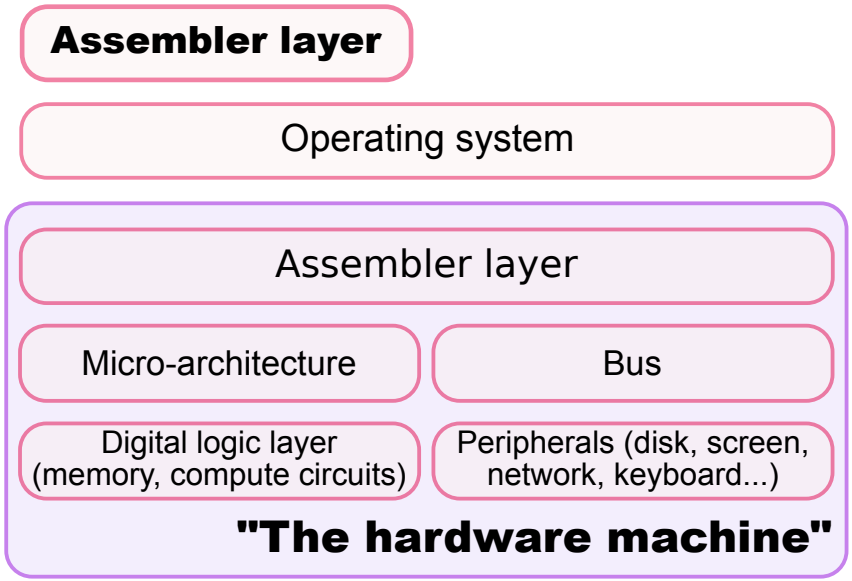

Let's create a "special application" which behaves exactly the same as the assembler layer...

Virtualization

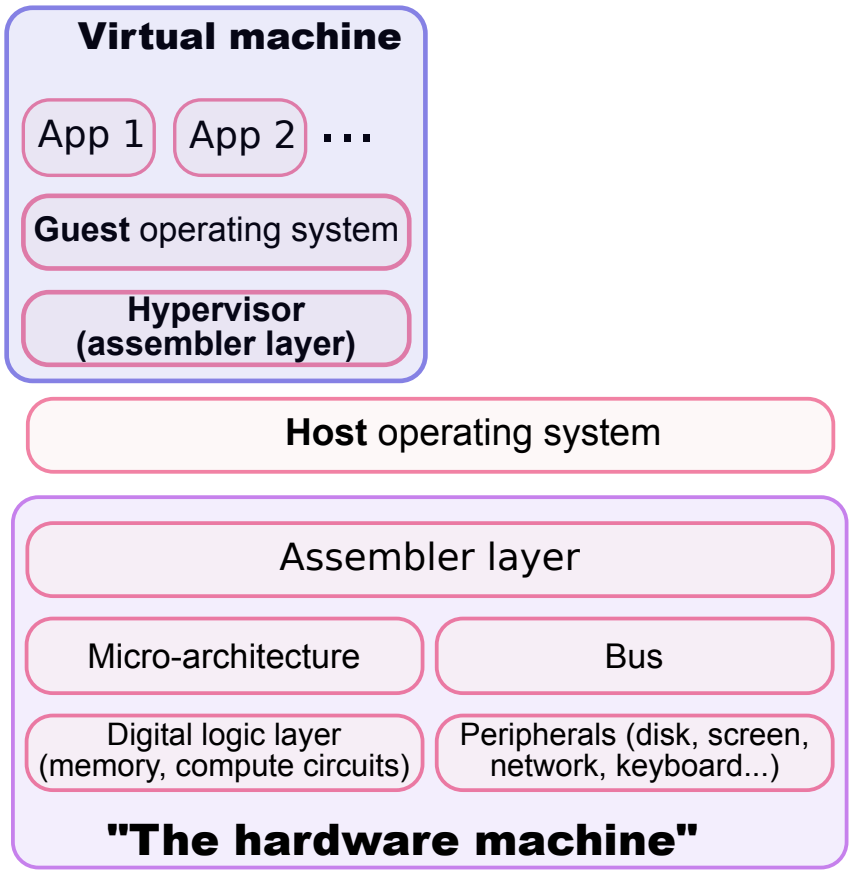

We can execute any operating system on top of it...

... and any application over the guest operating system

$\Rightarrow$ We have a virtual machine

Virtualization

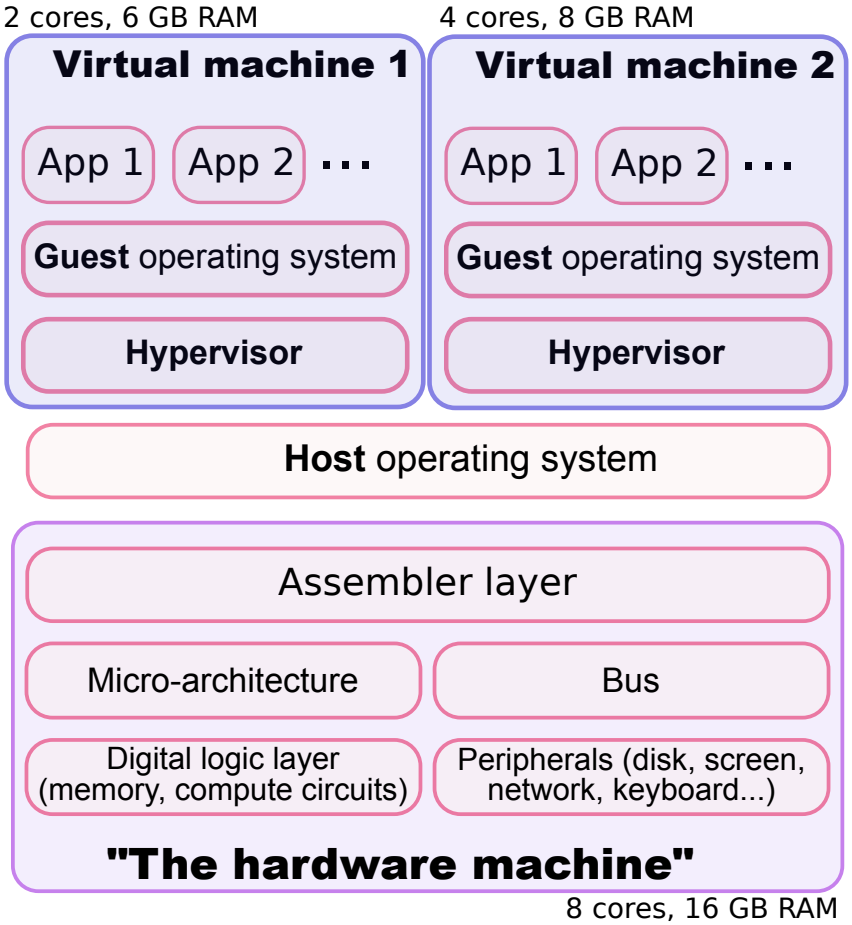

We can run multiple virtual machines on the same physical machine:

- Each virtual machine runs in full isolation from the other VMs
- Each virtual machine owns a subset of the hardware resources of the physical machine

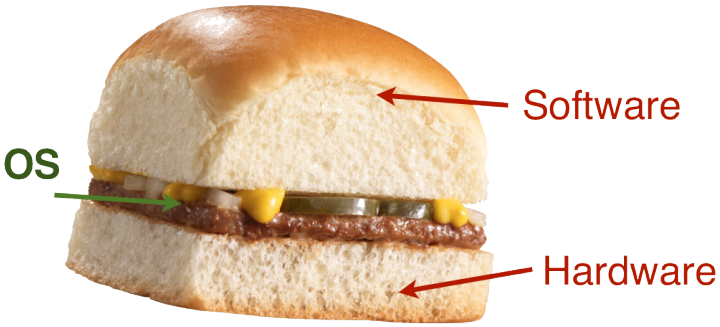
OS
Software
Hardware

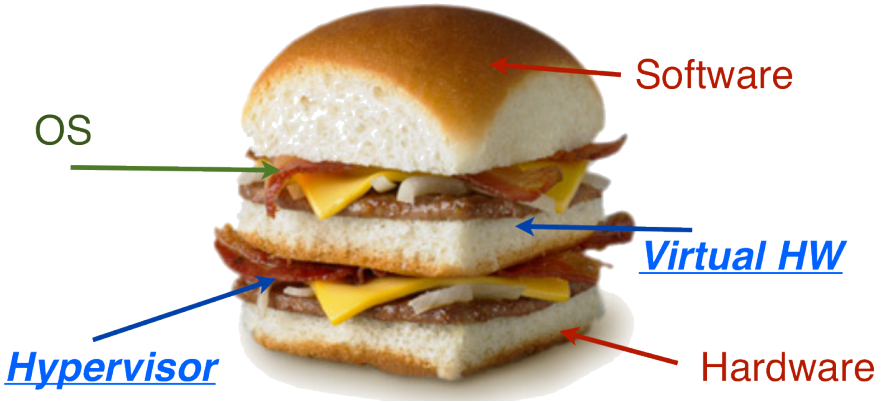
OS
Software
Virtual HW
Hypervisor
Hardware

Why is virtualization interesting for cloud providers?

Isolation: I can create multiple VMs on the same machine and give each VM to a different user (they will not see nor interfere with each other)

Customization: Each user can customize their VMs according to their own requirements.

Consolidation: Few applications can really exploit a large server machine to its maximum capacity. With virtualization I can split this capacity in smaller units and thereby increase my resource utilization.

Management: Virtualization simplifies resource management: I can measure how many resources each user is using, migrate VMs from one host to another, etc.

Virtualization technologies

Virtualization technologies are now totally mainstream:

- Commercial: VMware, Microsoft App-V, ...
- Open-Source: KVM, VirtualBox, Xen, ...

Paravirtualization vs. full virtualization:

- Paravirtualization works on any hardware platform but it requires special support in the guest OS.
- Full virtualization does not require special support in the guest OS.
  - Originally done through binary translation of system calls
  - Today it exploits special features of modern CPUs,

Virtualization technologies

Virtualization technologies are now totally mainstream:

- Commercial: VMware, Microsoft App-V, ...
- Open-Source: KVM, VirtualBox, Xen, ...

Paravirtualization vs. full virtualization:

- Paravirtualization works on any hardware platform but it requires special support in the guest OS.
- Full virtualization does not require special support in the guest OS.
  - Originally done through binary translation of system calls
  - Today it exploits special features of modern CPUs,

Virtual machines are often considered too heavyweight ☹

- Startup time $\sim$ 2–5 minutes
- Each guest OS needs lots of memory
- Each OS needs to execute lots of background stuff
- Impossible to run 100+ VMs on a single machine...

Containers (zones, jails, etc.)

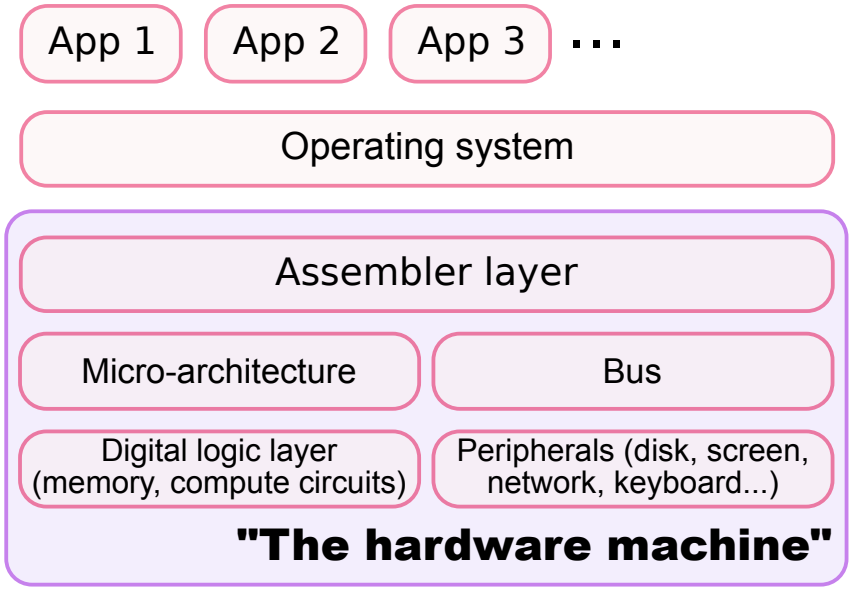

Traditional machine architecture:

- Applications
- Operating system
- Hardware

Containers (zones, jails, etc.)

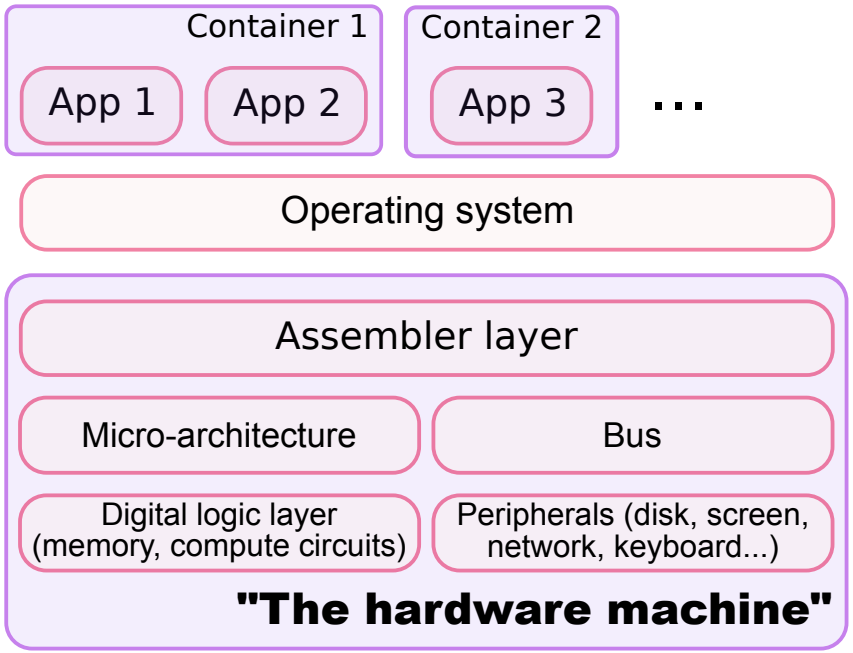

Let's create groups of processes which belong together:

- Process groups are totally isolated from each other
- Each process group has its own hardware resource limits (CPU, RAM, ...)
- Each process group has its own network access policy

$\Rightarrow$ We have containers

Container technologies

- Containers are an operating system feature
  - No need for special CPU support
  - Fully supported in Linux, Windows, BSD, Solaris...
  - Built with composable primitives on Linux (namespaces, cgroups, etc.) so add an extra software layer to simplify management like Docker, LXC, rkt
- Containers are less customizable than VMs
  - Container owners cannot choose their OS
  - But was that really necessary in the first place? Not always.
- Containers are much more lightweight than VMs
  - No need to run lots of (mostly identical) operating systems next to each other
  - Containers often start in less than 1 second
  - Can easily run hundreds of containers on a mid-sized machine

Containers provide indispensable flexibility for micro-service architectures

Public clouds vs. private clouds

Public clouds:

- Resources are owned by an external company (Amazon, Microsoft, Google, . . .)
- Pay per use (credit card)

Private clouds:

- Resources are owned by your own company
- Internal resource accounting

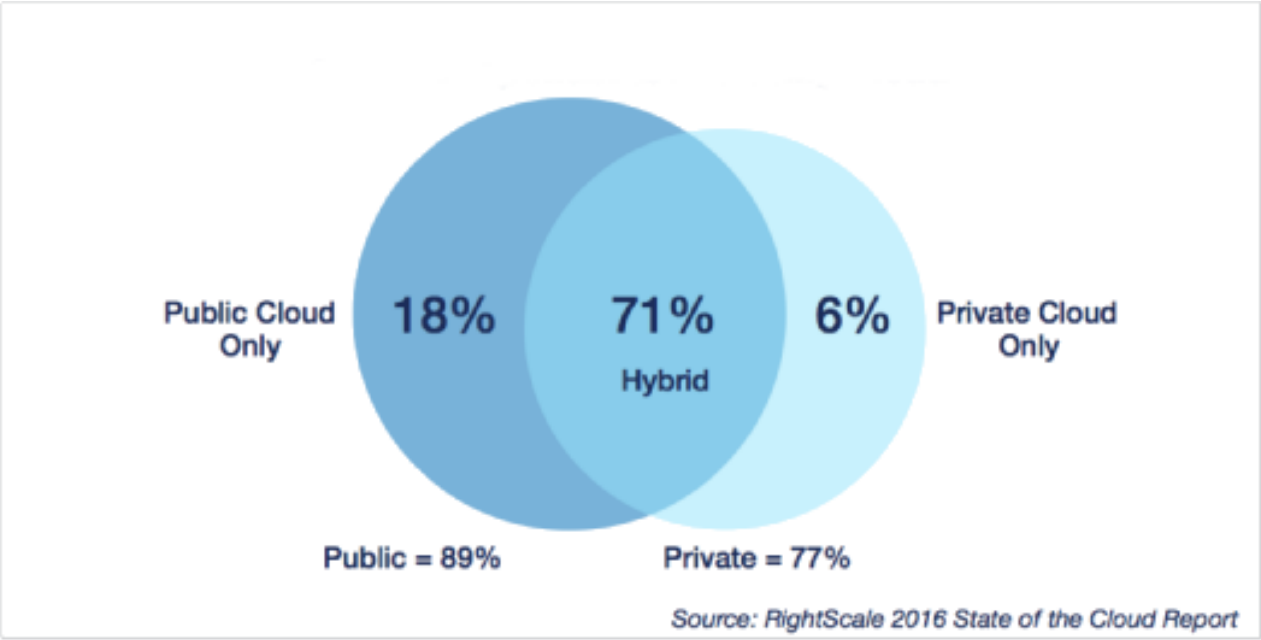

Enterprises are increasingly relying on the cloud

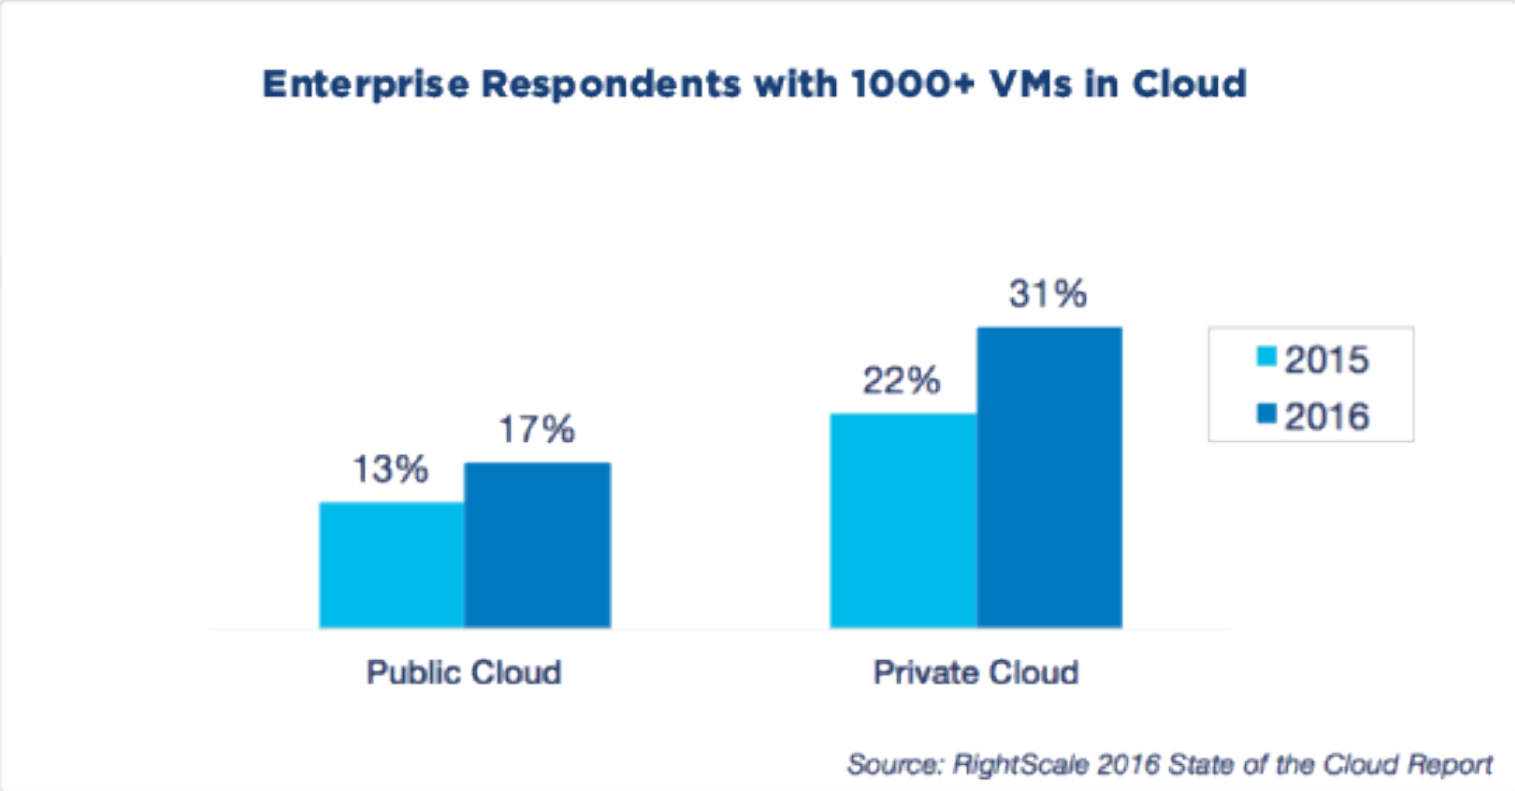

Cloud service layers

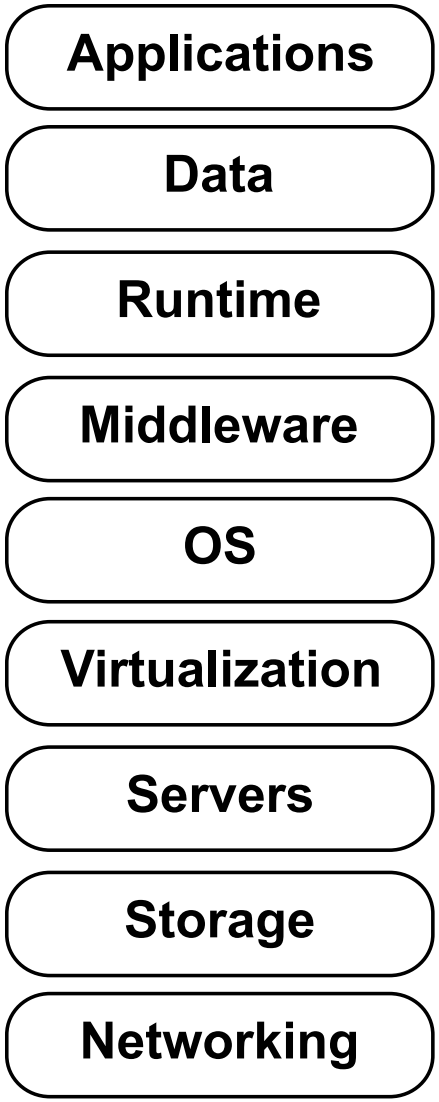

Cloud service layers

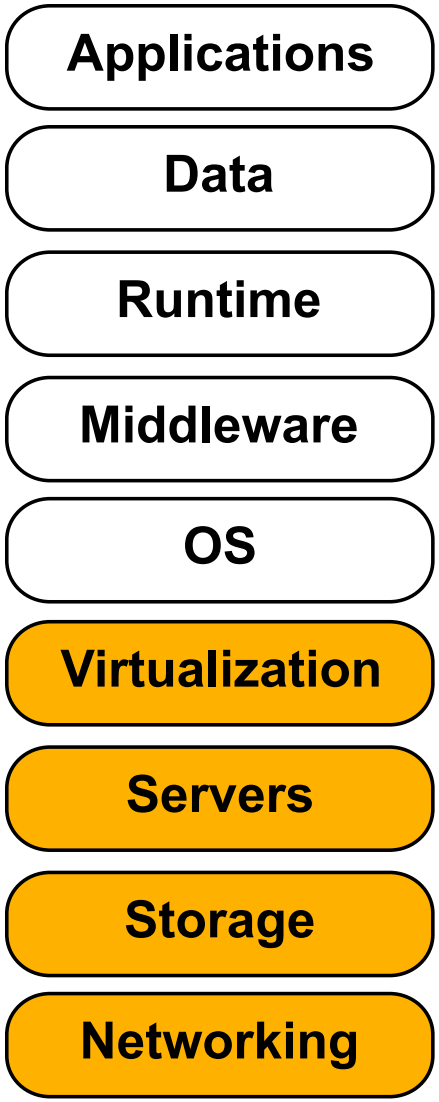

IaaS

Cloud service layers

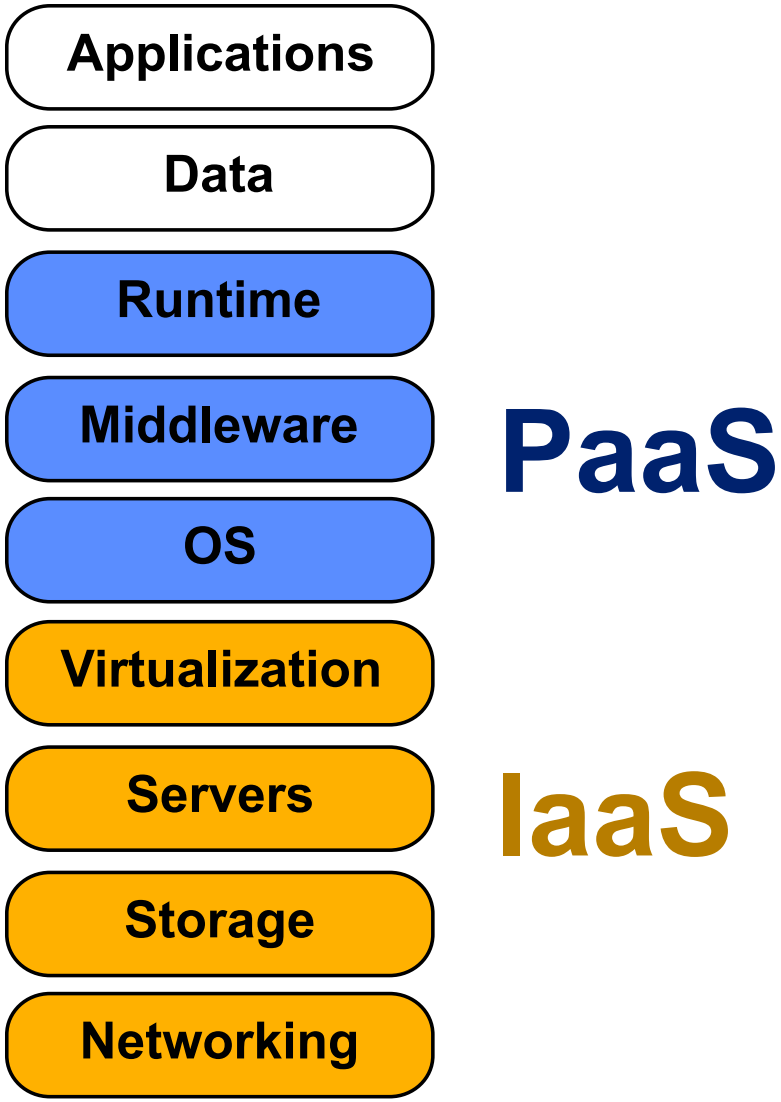

Cloud service layers

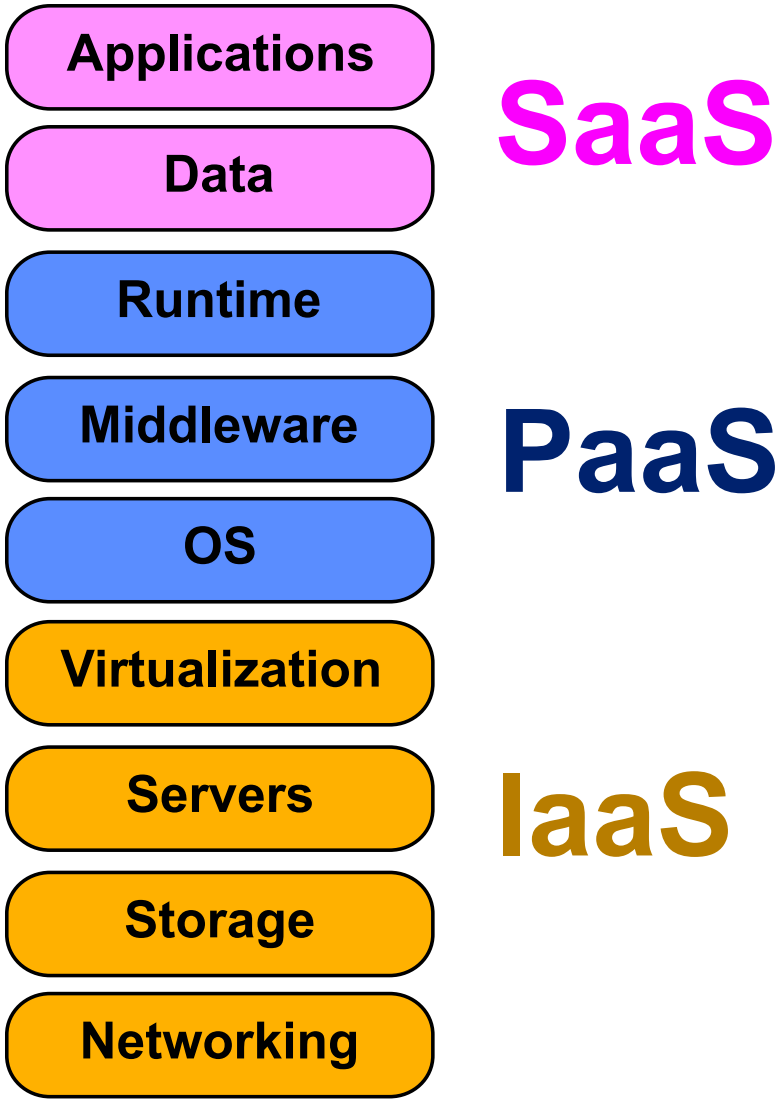

Cloud service layers

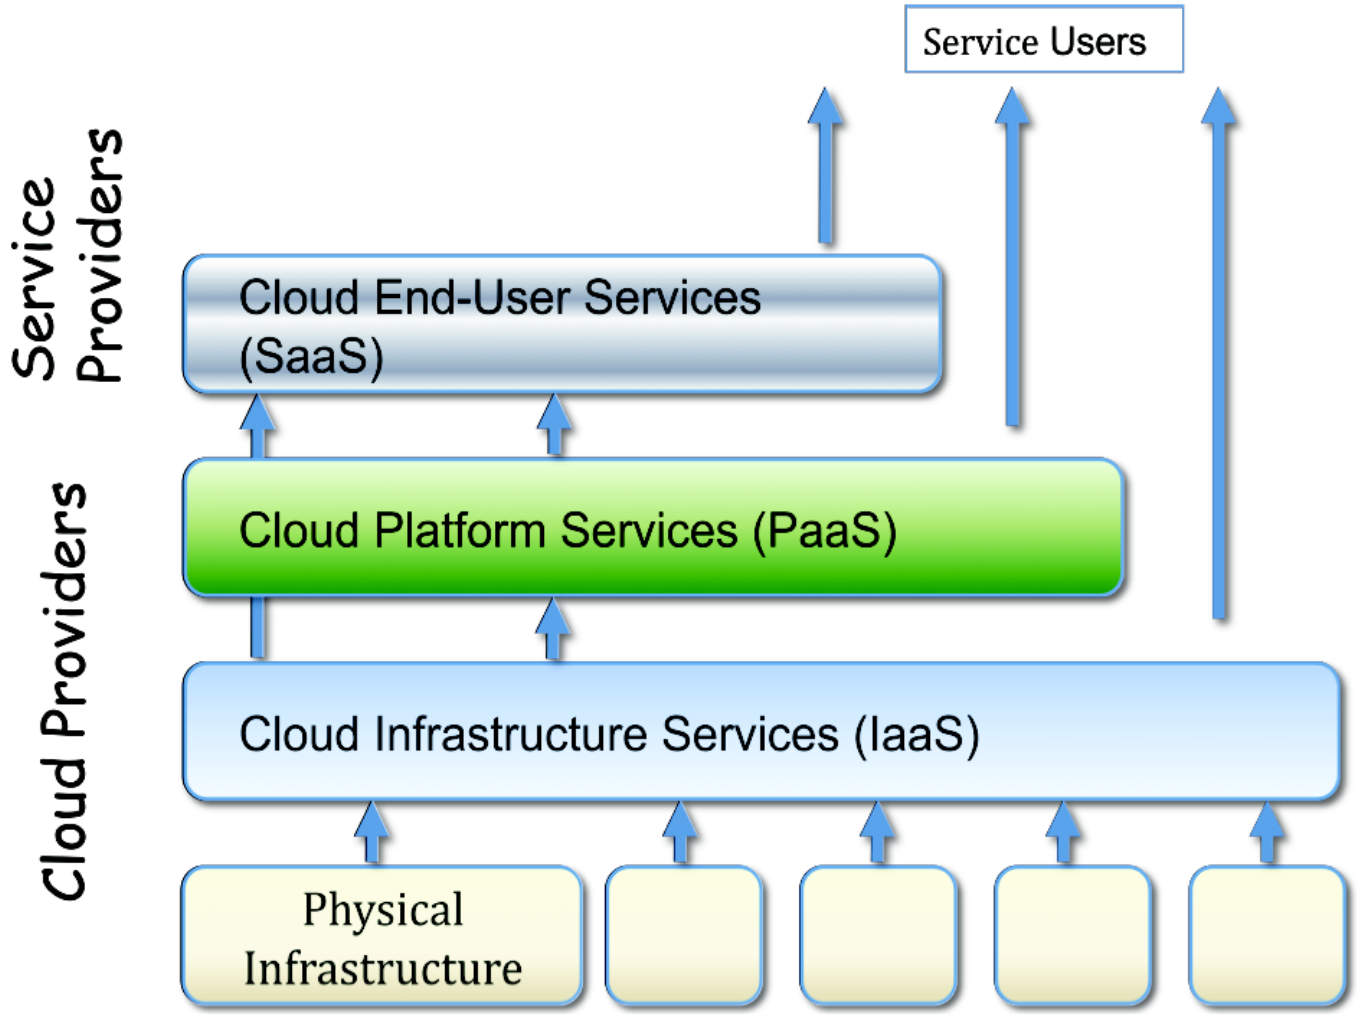

The main Cloud service layers

- Infrastructure-as-a-Service (IaaS) offers basic computing services
  - Computing: "Create a new machine for me"
  - Data storage: "Store/retrieve this data for me"
  - Communication
- Platform-as-a-Service (PaaS) offers high-level services for developers
  - Runtime environments: "Here is my Web application, run it for me"
  - Big data frameworks: "parallelize this program for me"
  - Databases: "I need a SQL database"
  - Development tools: "Give me a git repository"
- Software-as-a-Service (SaaS) offers high-level services for end users
  - Mail: Gmail
  - Office applications: Google docs
  - Enterprise applications: human resource applications, finance...
  - Payment services: Paypal
  - Data management: Dropbox, iCloud
  - Music on demand: iTunes, Spotify

(Unmodified) enterprise applications

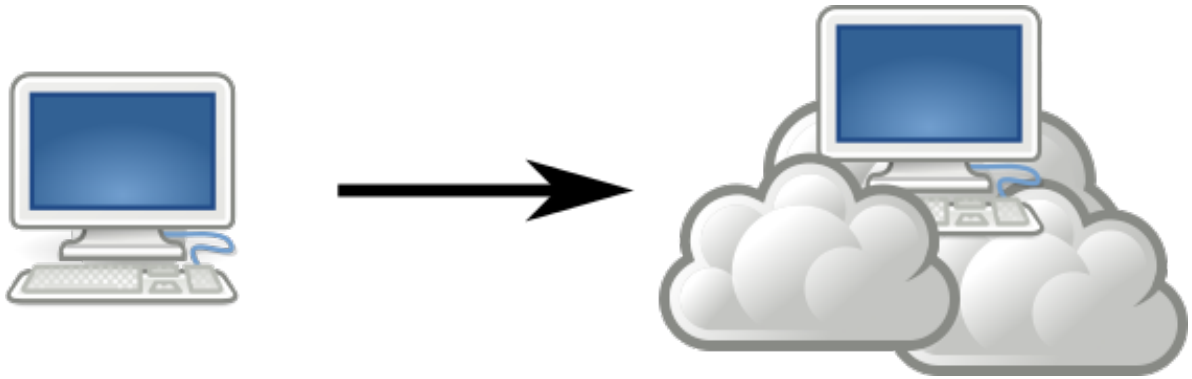

Not extremely exciting. . . ☹

(Unmodified) enterprise applications (bis)

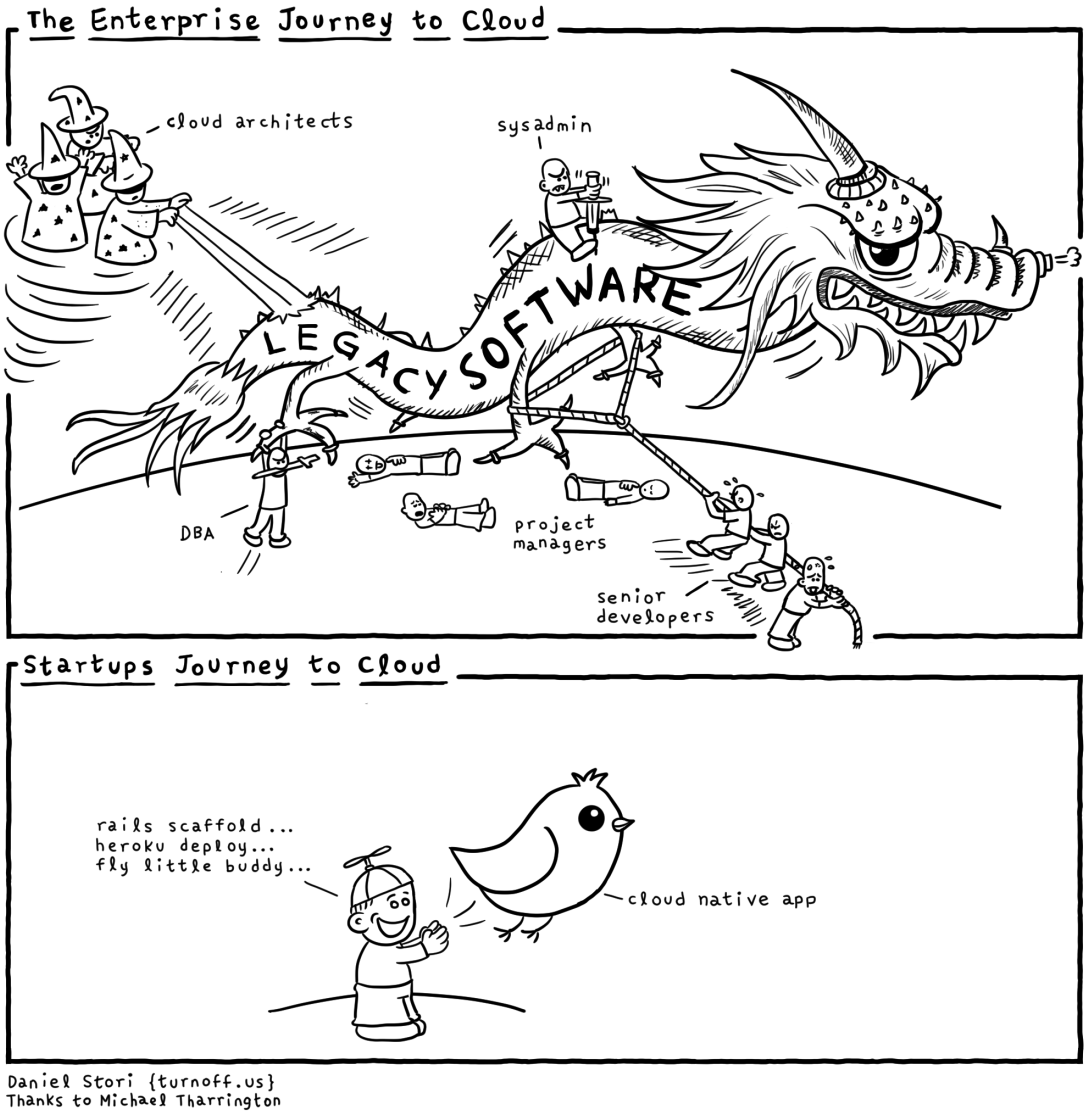

Data storage services

Dropbox

iCloud

Amazon S3

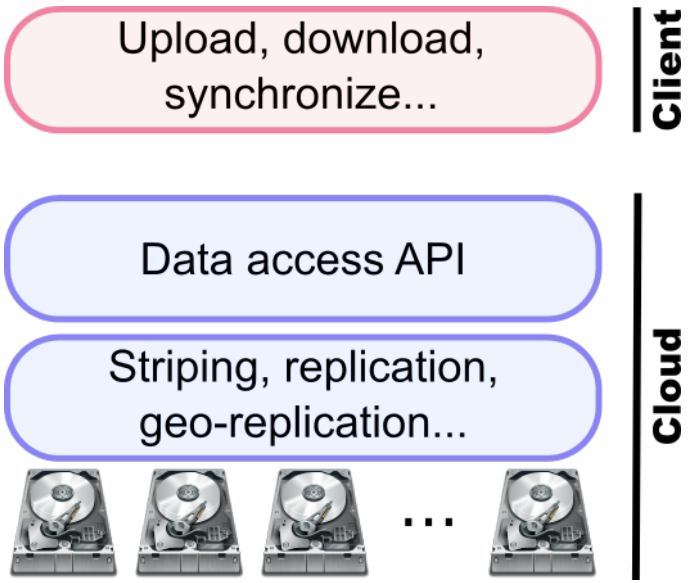

Web/mobile applications

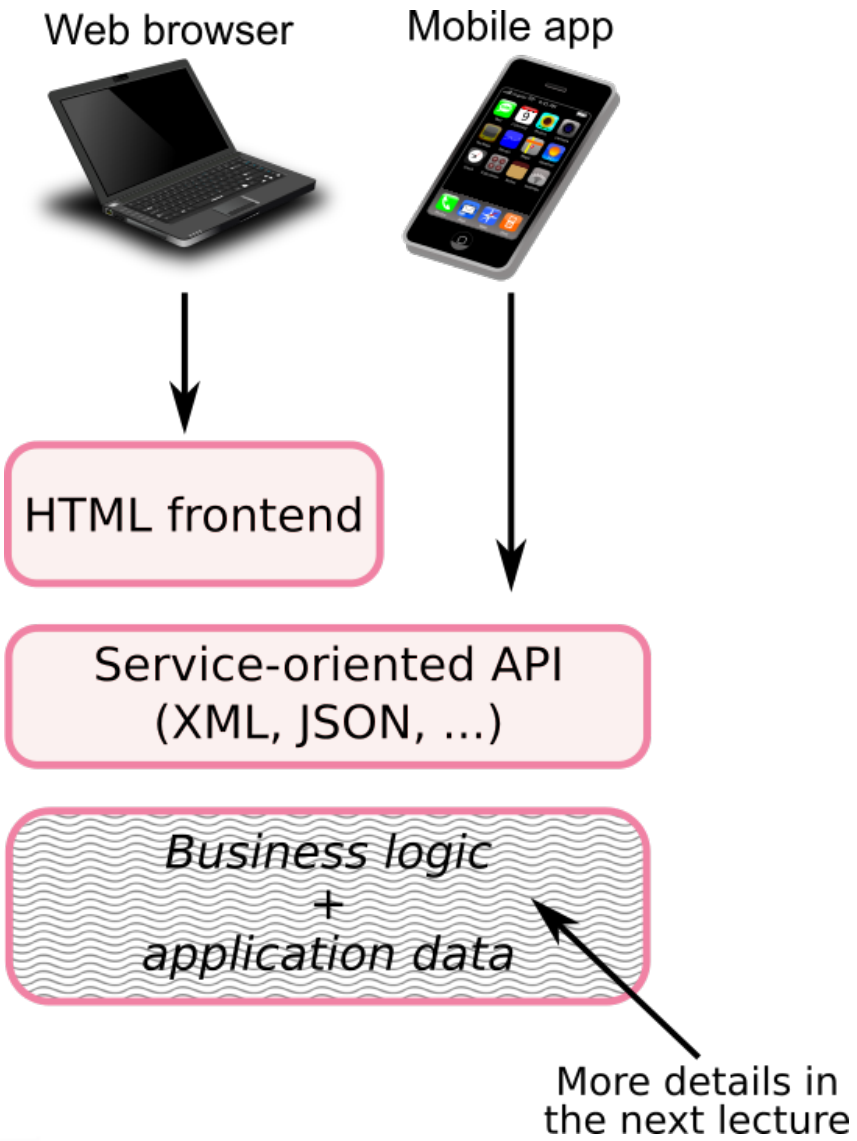

Data analytics

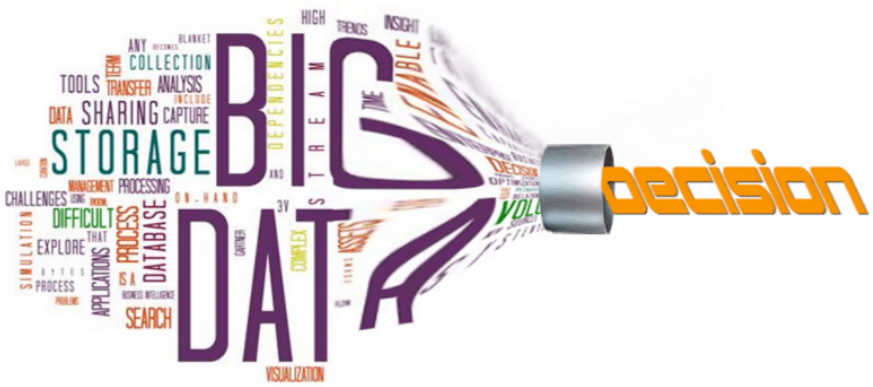

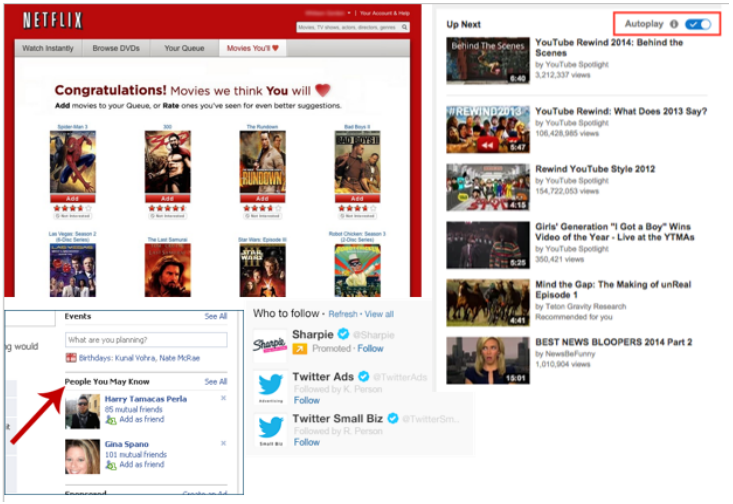

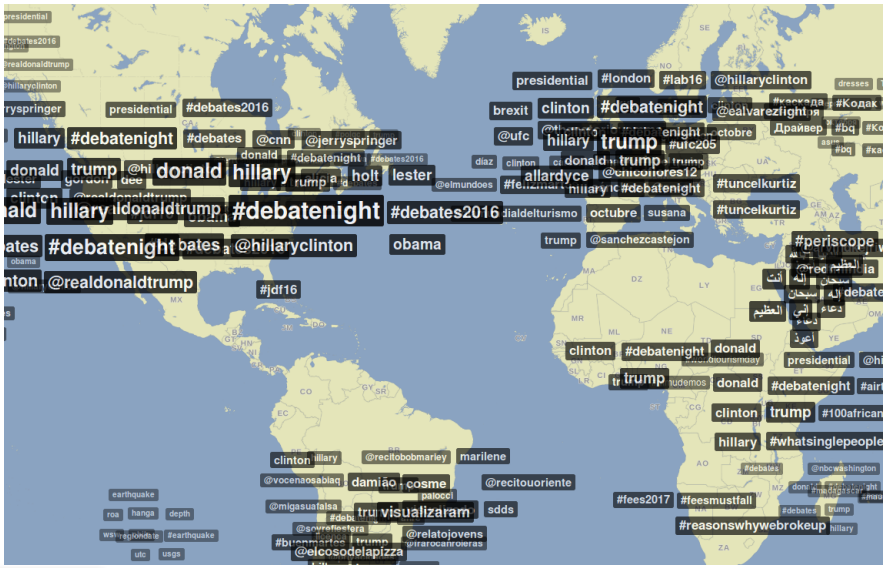

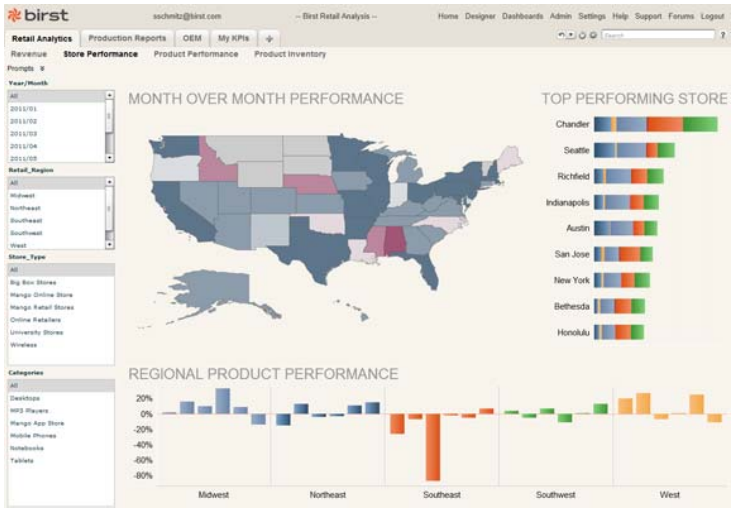

Cloud benefits

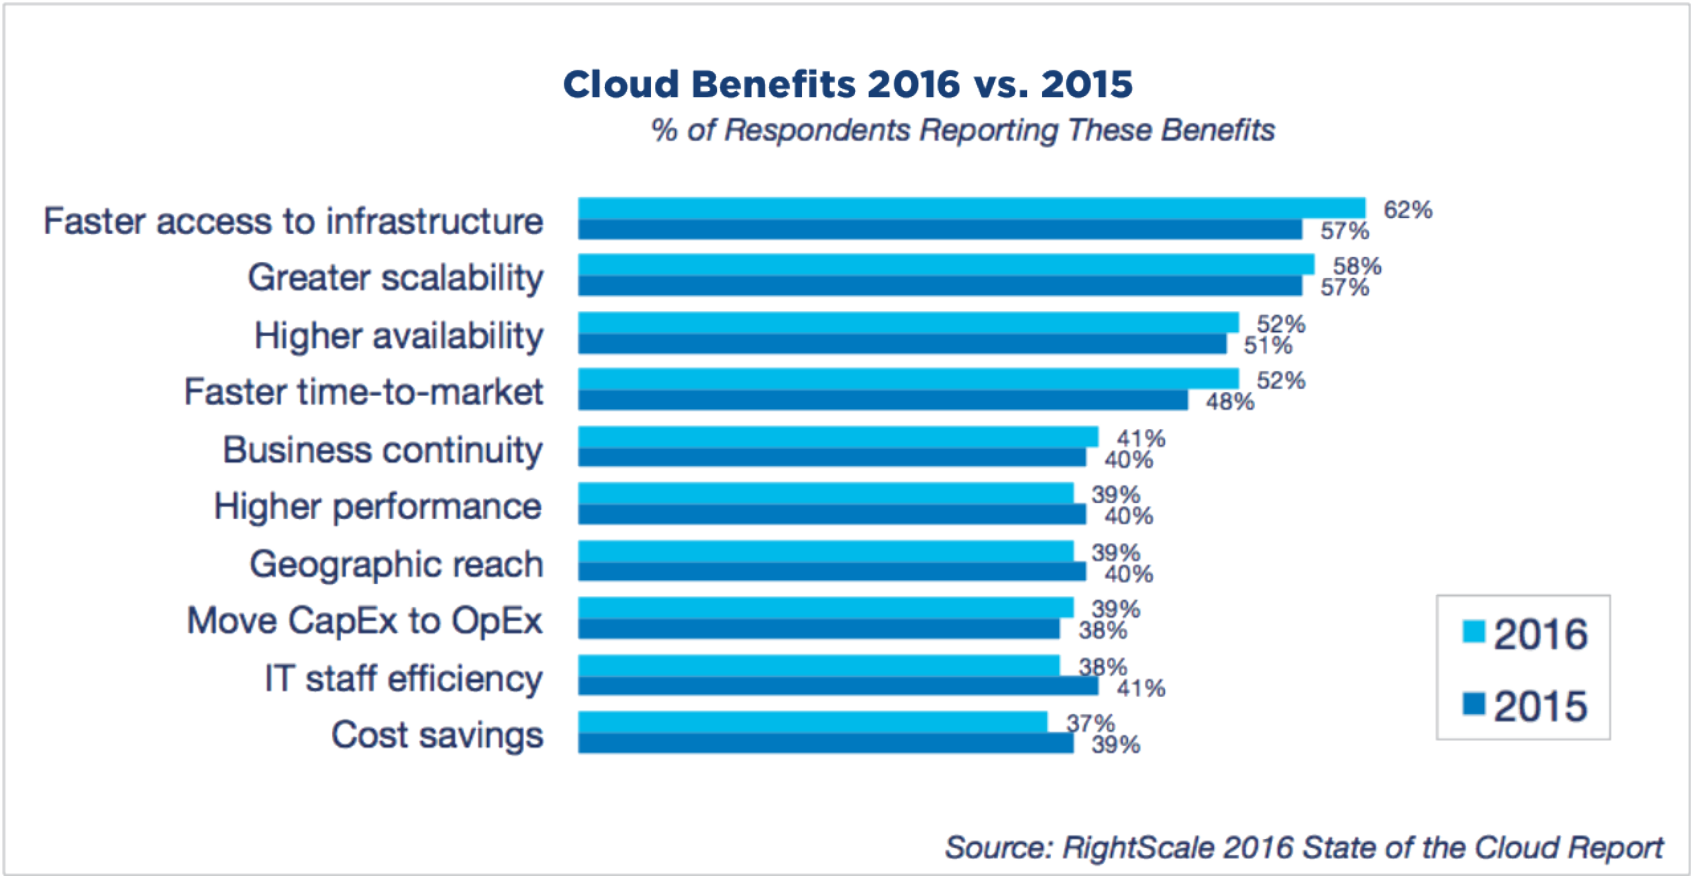

Cloud challenges

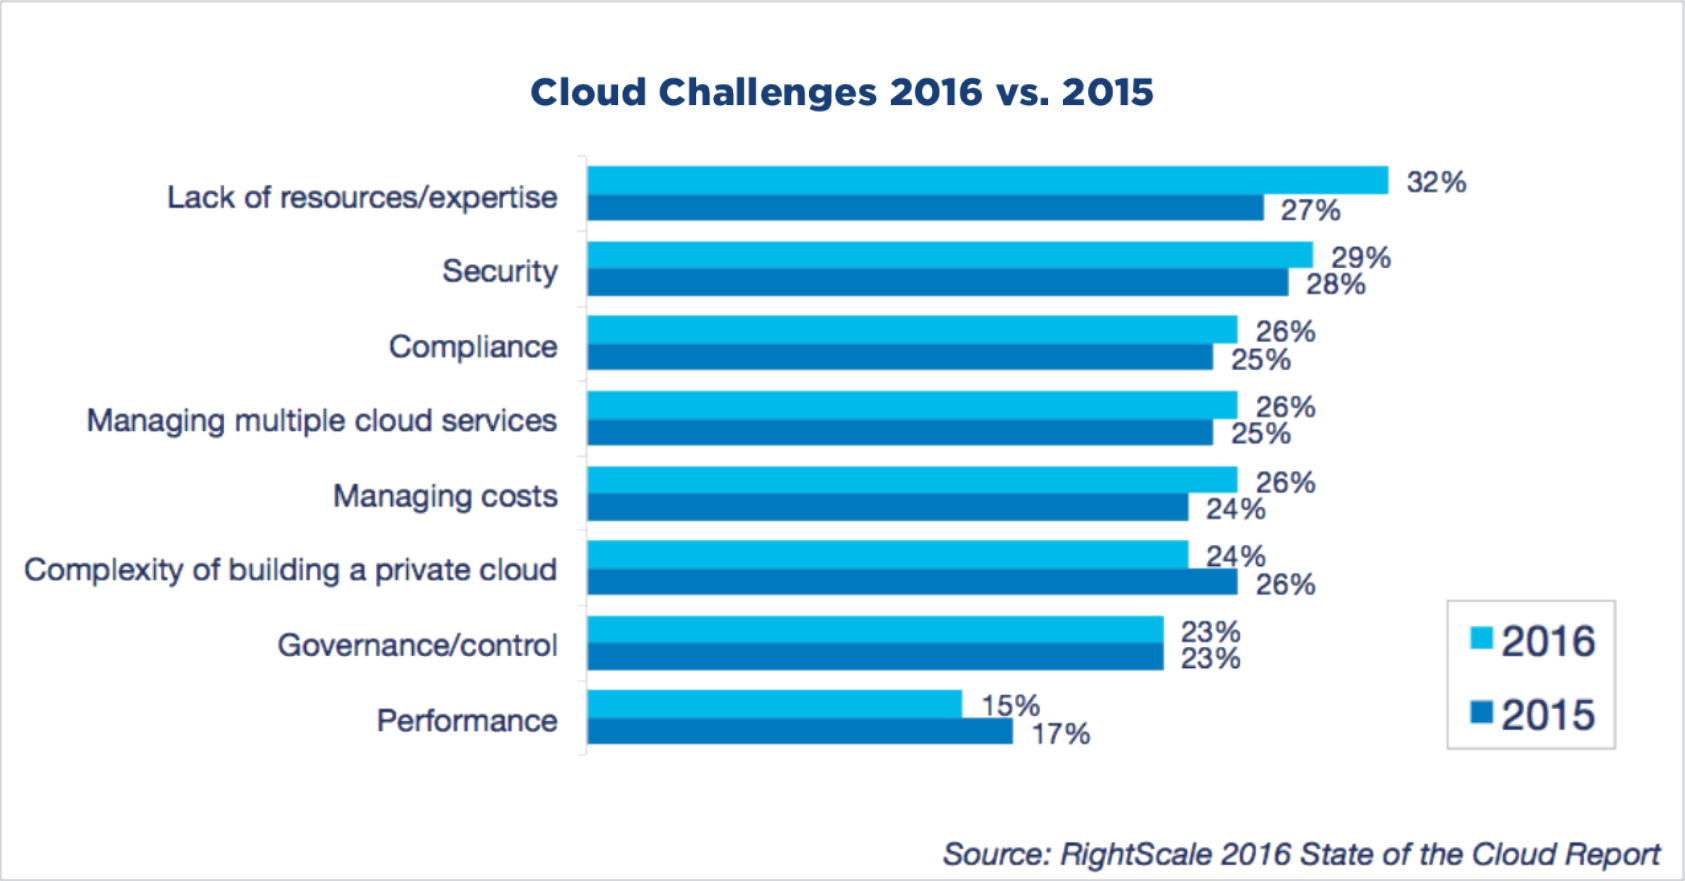

Cloud challenges